

# The Taming of the Grid: Lessons Learned in the National Fusion Collaboratory

Kate Keahey





# Overview

- Goals and Vision
- Challenges and Solutions
- Deployment War Stories
- Team
- Summary

## Goals

- “ enabling more efficient use of experimental facilities  
more effective integration of experiment, theory and  
modelling”
- Fusion Experiments
  - Pulses every 15-20 minutes
  - Time-critical execution
- We want:
  - More people running more simulation/analysis codes  
in that critical time window
  - Reduce the time/cost to maintain the software
  - Make the best possible use of facilities
    - > Share them
    - > Use them efficiently
- Better collaborative visualization (outside of scope)

# Overview of the Project

- Funded by DOE as part of the SciDAC initiative
  - 3 year project
  - Currently in its first year
- First phase
  - SC demonstration of a prototype
- Second phase
  - More realistic scenario at Fusion conferences
  - First shot at research issues
- Planning an initial release for November timeframe
- Work so far:
  - Honing existing infrastructure
  - Initial work on design and development of new capabilities

# Vision

- Vision of the Grid as a set of “network services”
  - Characteristics of the software (problems)
    - > Software is hard to port and maintain (large, complex)
    - > Needs updating frequently and consistently (physics changes)
    - > Maintenance of portability is expensive
    - > “Software Grid” as important as “Hardware Grid”
    - > Reliable between pulse execution for certain codes
    - > Prioritization, pre-emption, etc.
  - Solution:
    - > provide the application (along with hardware and maintenance) as a remotely available “service to community”
    - > Provide the infrastructure enabling this mode of operation

## What prevents us?

- **Issues of control and trust**
  - How do I enter into contract with resource owner?
  - How do I ensure that this contract is observed?
  - Will my code get priority when it is needed?
  - How do I deal with a dynamic set of users?
- **Issues of reliability and performance**
  - Time-critical applications
    - > How do we handle reservations and ensure performance?
  - Shared environment is more susceptible to failure
  - No control over resources
  - But a lot of redundancy

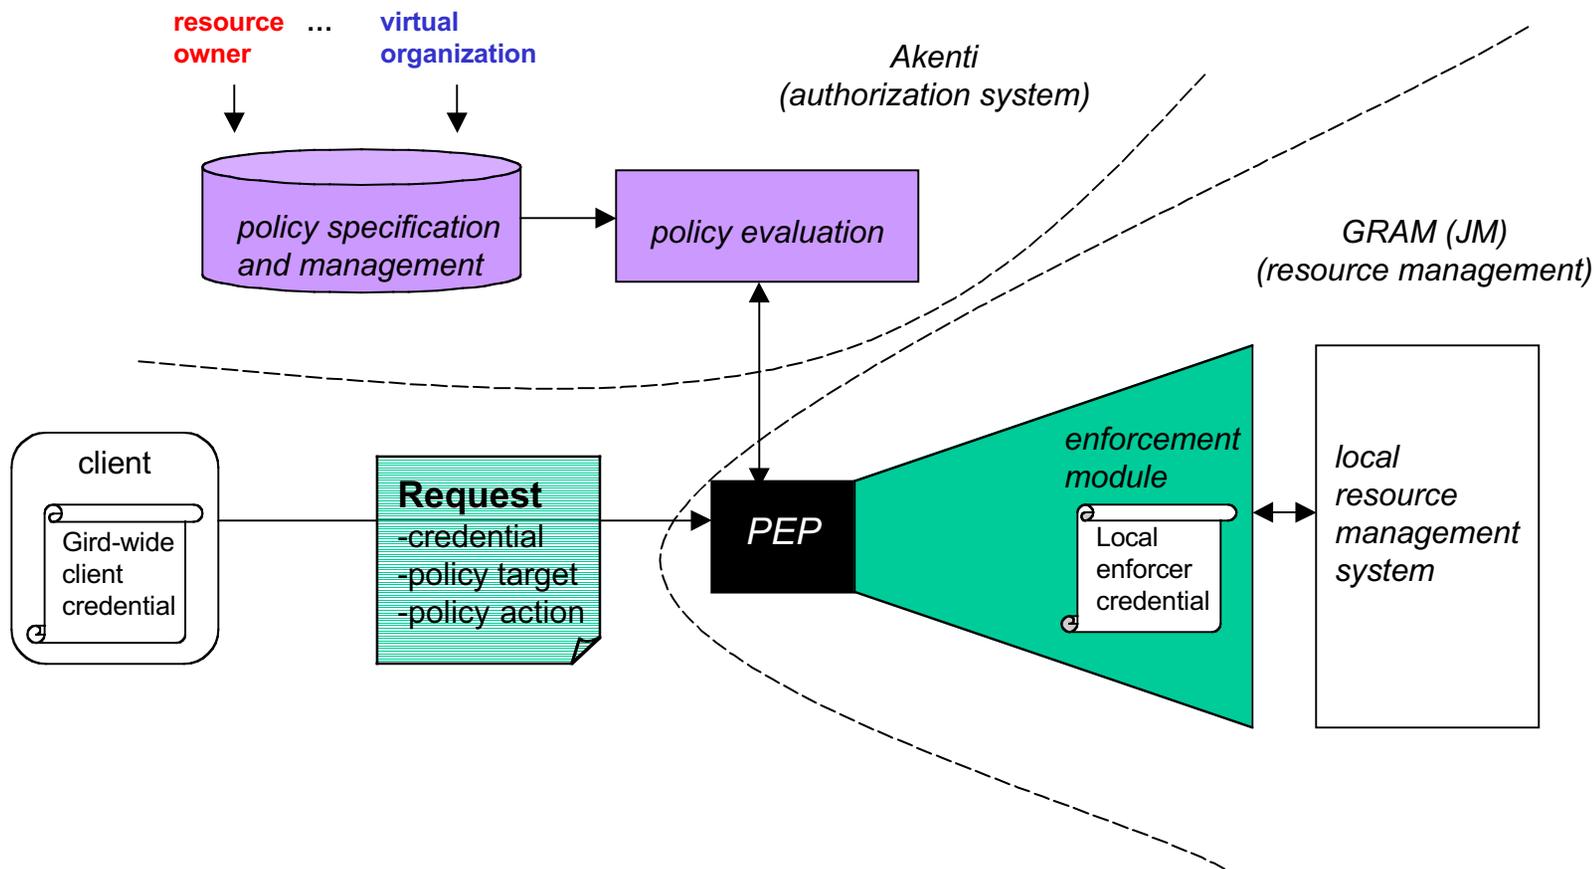
## Other Challenges

- Service Monitoring
- Resource Monitoring
- Good understanding of quality of service
  - Application-level
  - Composition of different QoS
- Accounting
- Abstractions
  - How do “network services” relate to OGSA Grid Services?
- Implementational and deployment issues
  - firewalls

## Issues of Trust: Use policies

- Requirements
  - Policies coming from different sources:
    - > A center should be able to dedicate a percentage of its resources to a community
    - > Community may want to grant different rights to different groups of users
  - A group within a VO may be given management rights for certain groups of jobs
  - Managers should be able to use their higher privileges (if any) to manage jobs
  - Shared/dynamic accounts dealing with dynamic user community problem

# Issues of Trust (cntd.)



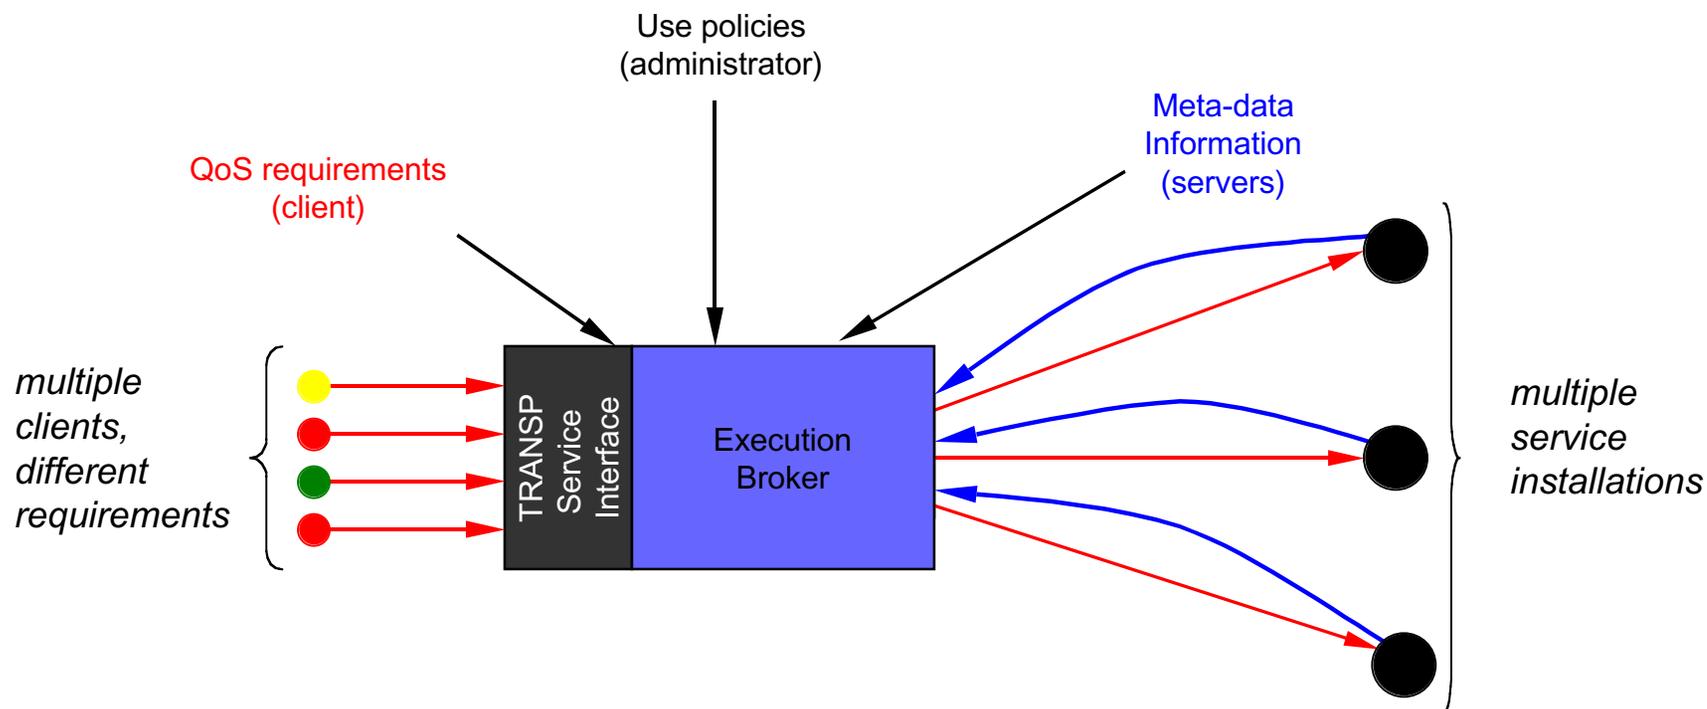
## Issues of Trust (cntd.)

- Policy language
  - Based on RSL
  - Additions
    - > Policy tags, ownership, actions, etc.
- Experimenting with different enforcement strategies
  - Gateway
  - Sandboxing
  - Services
- Joint work with Von Welch (ANL), Bo Liu
- Work based on GT2
- Collaborating with Mary Thompson (LBNL)

## Issues of Reliable Performance

- Scenario:
  - A GA scientist needs to run TRANSP (at PPPL) between experimental pulses in less than 10 mins
  - TRANSP inputs can be experimentally configured beforehand to determine how its execution time relates to them
    - > Loss of complexity (“physics”) to gain time
  - The scientist reserves the PPPL cluster for the time of the experiment
  - Multiple executions of TRANSP, initiated by different clients and requiring different QoS guarantees can co-exist on the cluster at any time, but when a reservation is claimed, the corresponding TRANSP execution claims full CPU power

# Issues of Reliable Performance (cntd)



- Status: an OGSA-based prototype
  - Uses DSRT and other GARA-inspired solutions to implement pre-emption, reservations, etc.
- Joint work with Kal Motawi

## Deployment (Firewall Problems)

- The single most serious problem: firewalls
  - Globus requires
    - > Opening specific ports for the services (GRAM, MDS)
    - > Opening a range of non-deterministic ports for both client and server
    - > Those requirements are necessitated by the design
  - Site policies and configurations
    - > Blocking outgoing ports
    - > Opening a port only for traffic from a specific IP
    - > Authenticating through the firewall using SecureID card
    - > NAT (private network)
    - > “opening a firewall is an extremely unrealistic request”
- An extremely serious problem: makes us unable to use the Fusion Grid

## Firewalls (Proposed Solutions)

- Inherently difficult problem
- Administrative Solutions
  - Explain why it is OK to open certain ports
    - > Document explaining Globus security (Von Welch)
  - Agree on acceptable firewall practices to use with Globus
    - > Document outlining those practices (Von Welch)
  - Talk to potential “influential bodies”
    - > ESCC: August meeting, Lew Randerson, Von Welch
    - > DOE Science Grid: firewall practices under discussion
- Technical Solutions
  - OGSA work: Von Welch, Frank Siebenlist
  - Example: route interactions through one port
- Do you have similar problems? Use cases?

## Firewalls (Resources)

- New: updated firewall web page
  - <http://www.globus.org/security/v2.0/firewalls.html>
- Portsmouth, UK
  - <http://esc.dl.ac.uk/Papers/firewalls/globus-firewall-experiences.pdf>
- DOE SG Firewall Policy Draft (Von Welch)
- DOE SG firewall testbed
- Globus Security Primer for Site Administrators (Von Welch)

## The NFC Team

- Fusion
  - David Schissel, PI, General Atomics (applications)
  - Doug McCune, PPPL (applications)
  - Martin Greenwald, MIT (MDSplus)
- Secure Grid Infrastructure
  - Mary Thompson, LBNL (Akenti)
  - Kate Keahey, ANL, (Globus, network services)
- Visualization
  - ANL
  - University of Utah
  - Princeton University
- More information at [www.fusiongrid.org](http://www.fusiongrid.org)

## Summary

- Existing infrastructure
  - A lot in relatively little time
  - Caveat: firewalls
- Building infrastructure
  - Network services
    - > A view of a “software grid”
    - > Goal: to provide execution reliable in terms of an application-level QoS
    - > To accomplish this goal we need:
      - Authorization and use policies
      - Resource management strategies