Open Science Grid

Moving the National and Global Ecosystem forward: Dynamic Resources (Clouds, HPC), Network Monitoring, Security, etc.

August 19, 2014

Michael Ernst
OSG Applications

Open Science Grid

Evolution of OSG

OSG's focus is expected to shift towards an ecosystem that is designed to provide support for an even richer variety of science domains

- OSG's value is in building community, leveraging efforts to produce synthesis, sharing resources to maximize science return on investment
- OSG's services, though very successful, are largely building on components that were invented and implemented in the early days of the Grid
  - Compute Element (CE), Storage Element (SE), central information services, X509 certificates are in our way and need to be replaced
- OSG will become a much less static system. Resource providers and consumers will be continuously flowing in and out
- Building on the concept of the application taking the environment to the resources at execution time (Virtual Machines, OASIS/CVMFS + parrot)
- We have tools for programmatic/light(er) weight provisioning for computational resources (e.g. BOSCO, OSG Connect, AutoPyFactory for VM lifecycle management in private and commercial cloud environments)
- We have tools to deliver data from a distance (Data Store Federations over Xrootd and http)
- Divergence of HTC into CPU (co-processor dominated) and heavy I/O (data filtering)
- Capabilities to serve data intensive disciplines
  - Large communities increasing use of high performance/high throughput capabilities

“Mental Model” of OSG DHTC Facility

Why is Resource Provisioning Important?

- Open Science Grid supports already a variety of science domains but will only grow with flexible and nimble mechanisms researchers can easily adopt
  - Large (e.g. LHC) and small stakeholders need opportunistic computing and expect resource growth for programs that may last decades
- OSG will evolve in influence by being indispensable to many smaller sciences
  - Key Issue: Time needed to register resources and adopt applications
    - Adoption time should be less than 10% of usage time, ideally applications need to be up in one day
    - OSG Connect is one way to address the issue
    - Demand driven, policy based provisioning of Cloud resources is another

Bringing in New Types of Resources

Open Science Grid

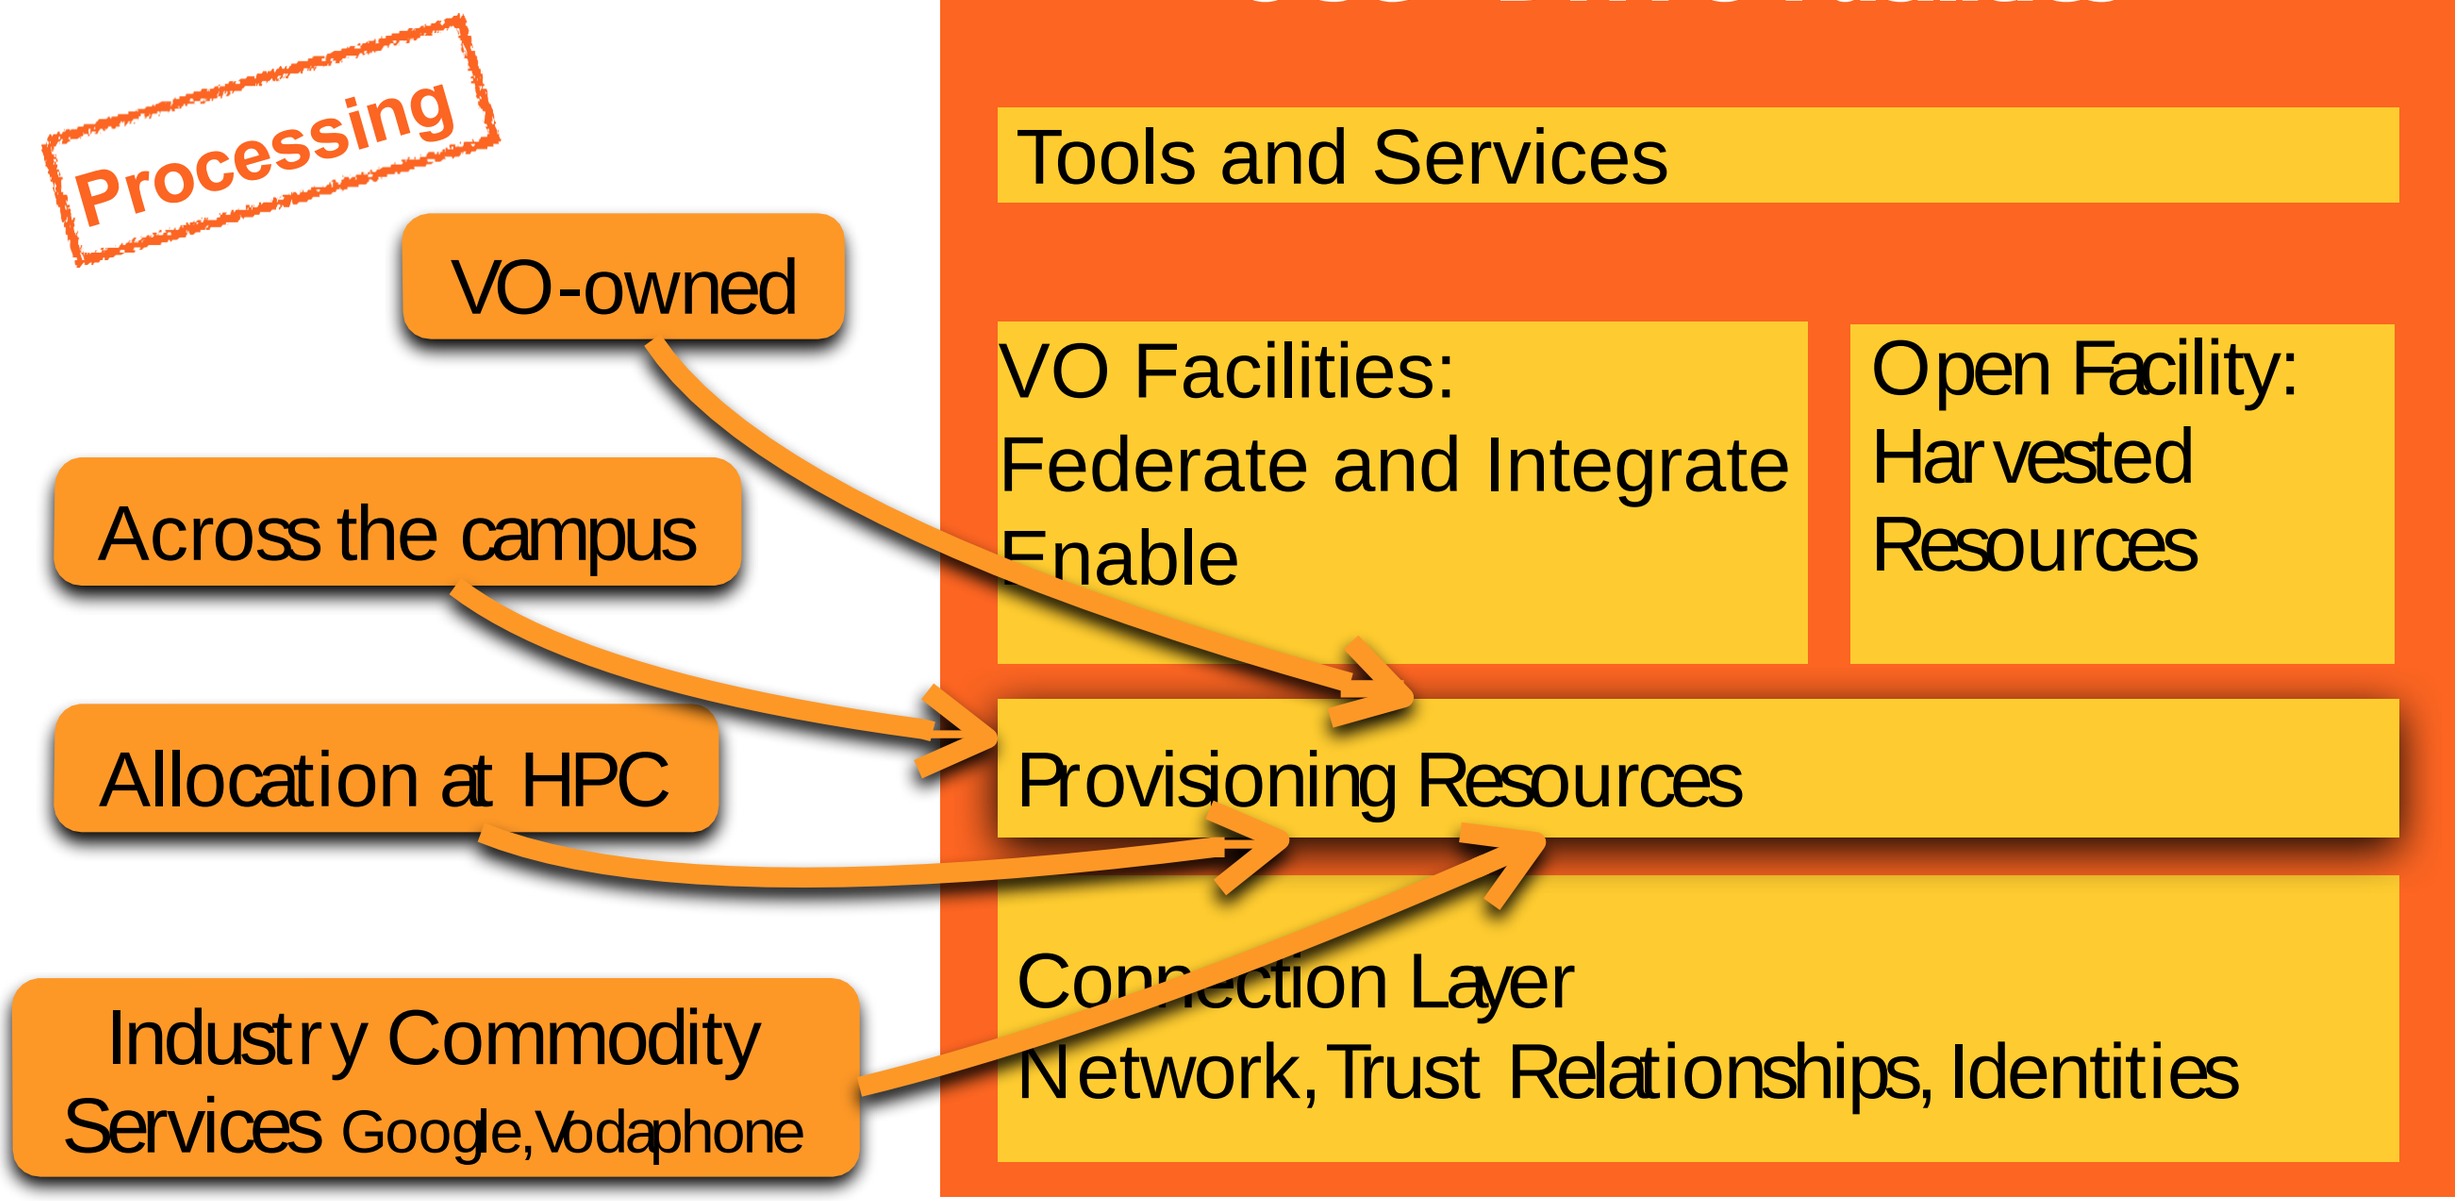

- Focus on *Resource Provisioning*
  - ★ Statically federated resources need to be integrated with dynamically allocated resources causing new challenges for resource planning, acquisition, provisioning

A Provisioning Model for the OSG

Multi-year experience has shown that the Job Management Overlay, pilot based approach is far superior to submitting payload jobs to sites

- Pilot job infrastructure based on “pilot factory”
- Pilot factories are part of the Resource Provisioning System
- Resources provisioned becoming more diverse and policies more complex (e.g. EC2-like IaaS provisioning)

Resource Provisioning is in transition to a new era

- Observing increased use of EC2-compatible “Infrastructure as a Service (IaaS)”
  - Instead of a batch slot the resource acquired is a virtual machine
  - Increasing user interest in using sites via common interfaces like SSH instead of grid interfaces
  - Stakeholders are asking for increasingly complicated resource acquisition policies

Resource Provisioning to Campus Users

Open Science Grid

- Can OSG assist VOs provision resources with different costs?
  - Clouds are a “business model” that Resource Providers use to provide resources
  - Other examples are CPU power, network locality, allocation
- OSG could act as an intermediary, between VOs and resources
- OSG does not yet have a flexible “provisioning” capability
  - Discussing a model that can be extended to “enabling acquisition” of commercial cycles on behalf of an OSG member
  - Maybe even be extended to the provisioning of opportunistic resources

Using Cloud Resources effectively: A Policy-based Cloud Scheduler, developed at BNL

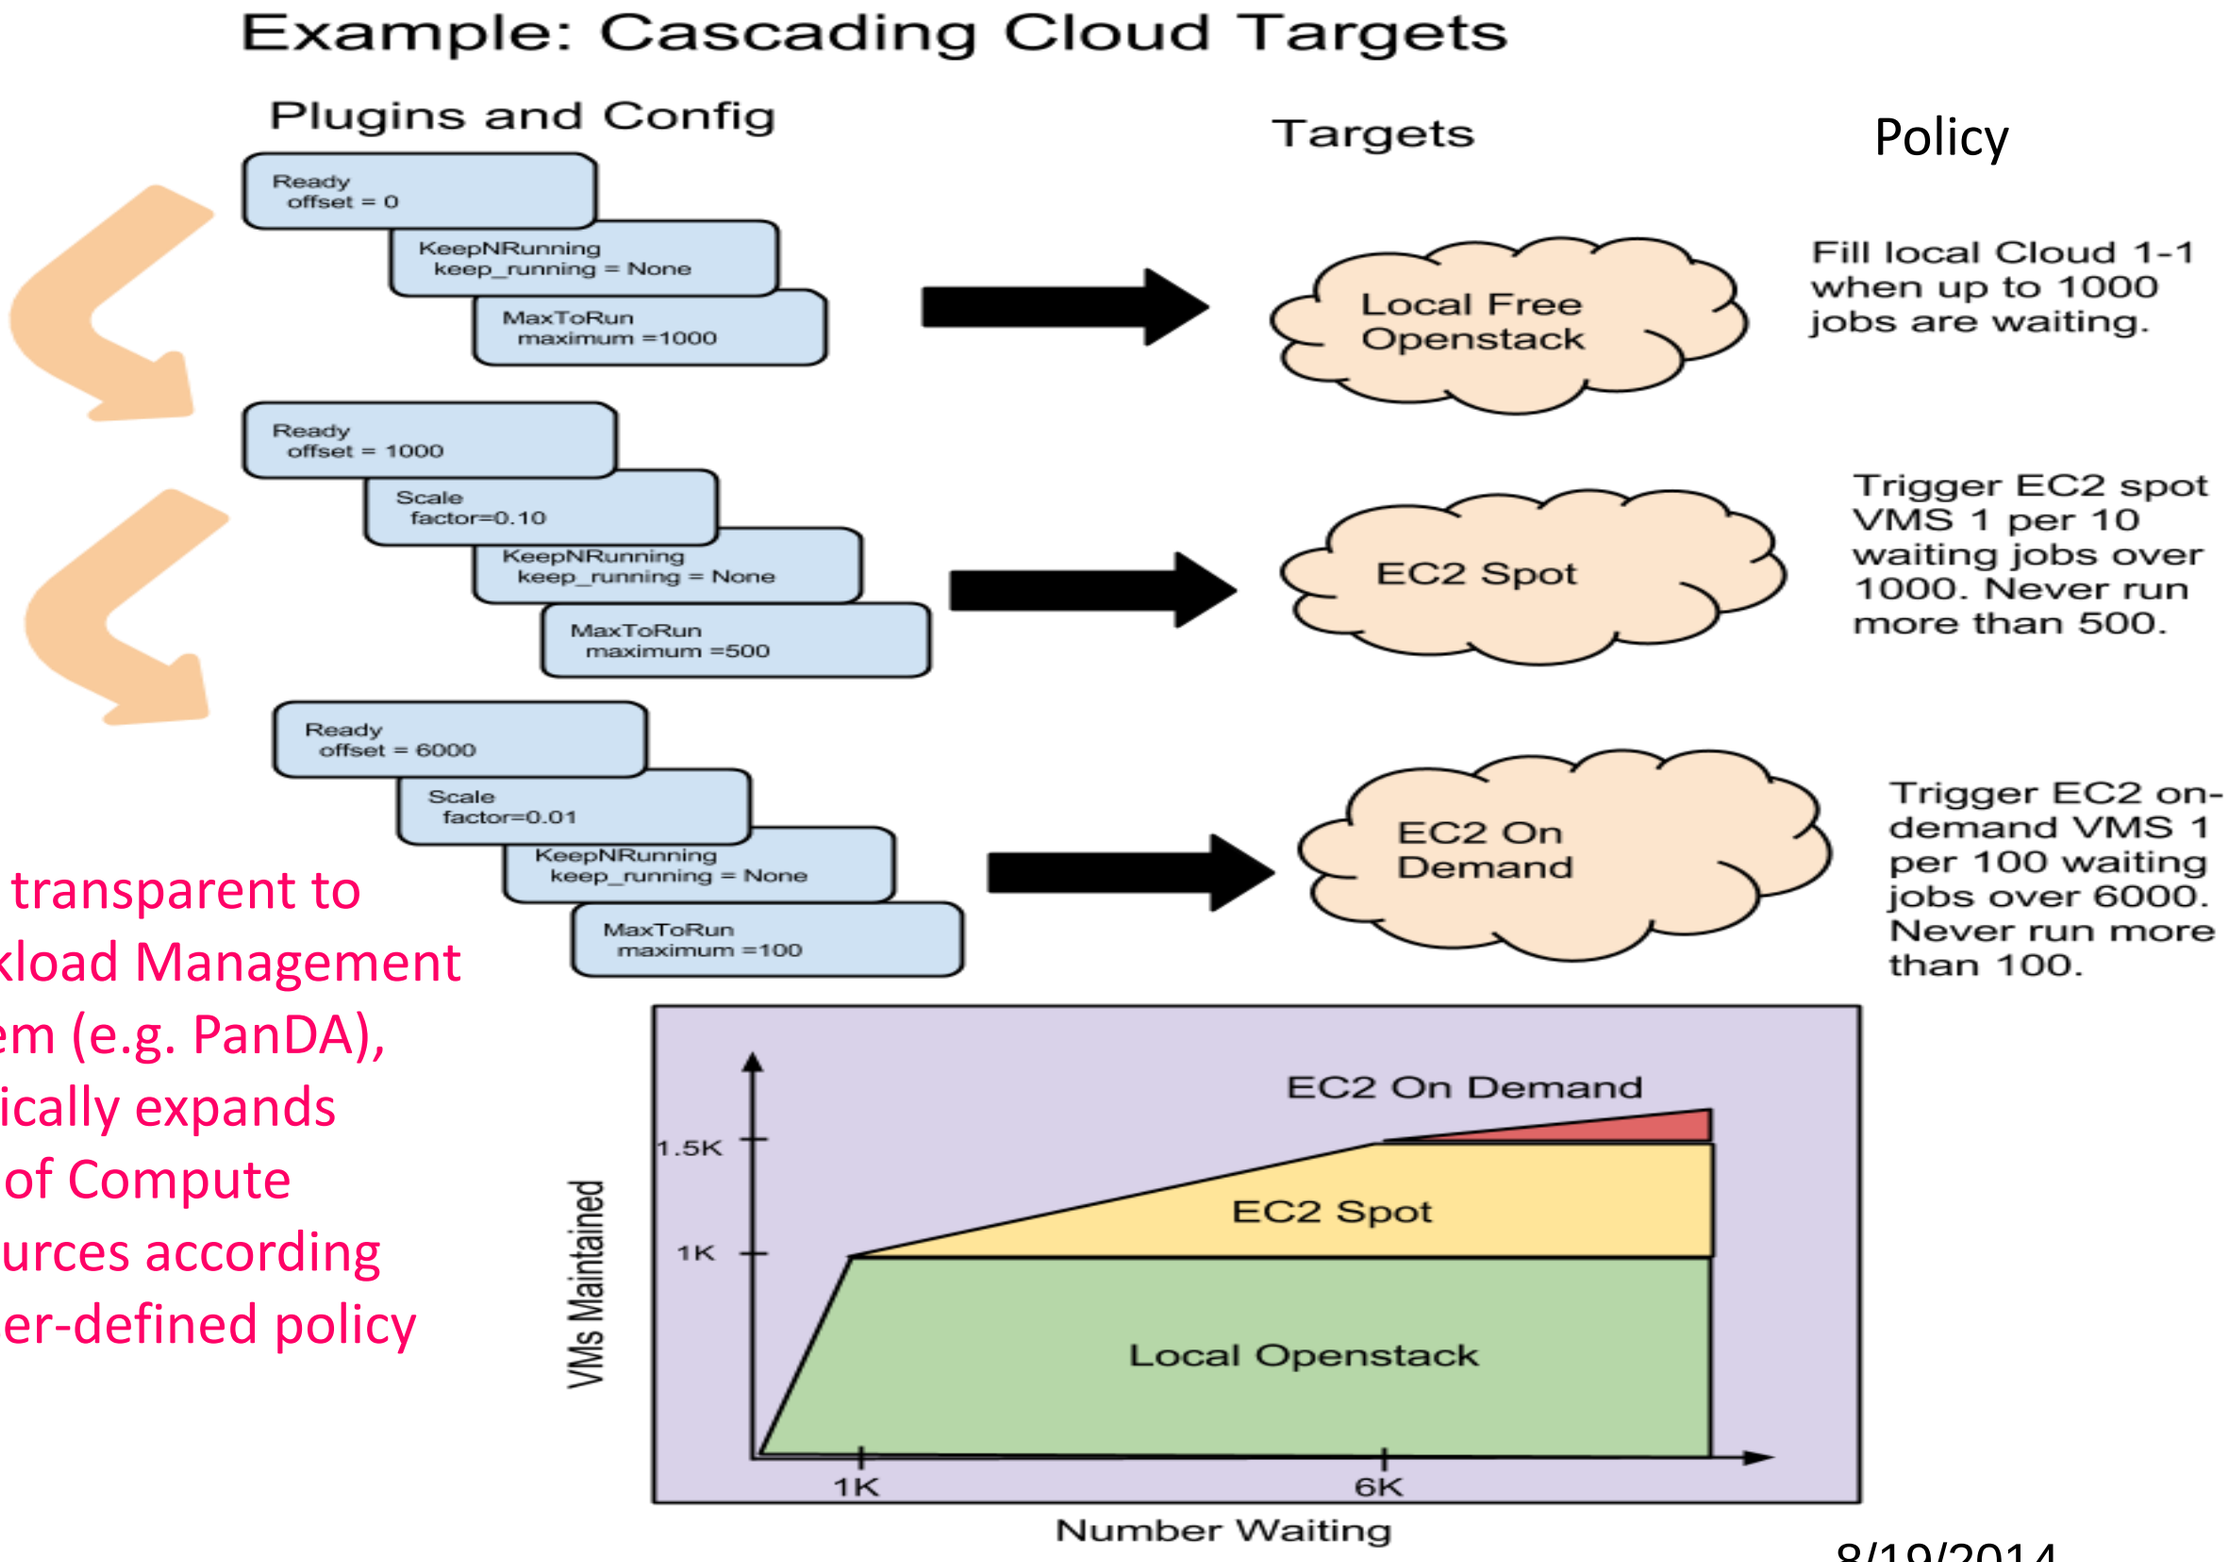

Fully transparent to Workload Management System (e.g. PanDA), Elastically expands Pool of Compute Resources according to user-defined policy

Open Science Grid

Elastic Cluster

Ready to serve
- Peak Demands
- users w/o dedicated resources

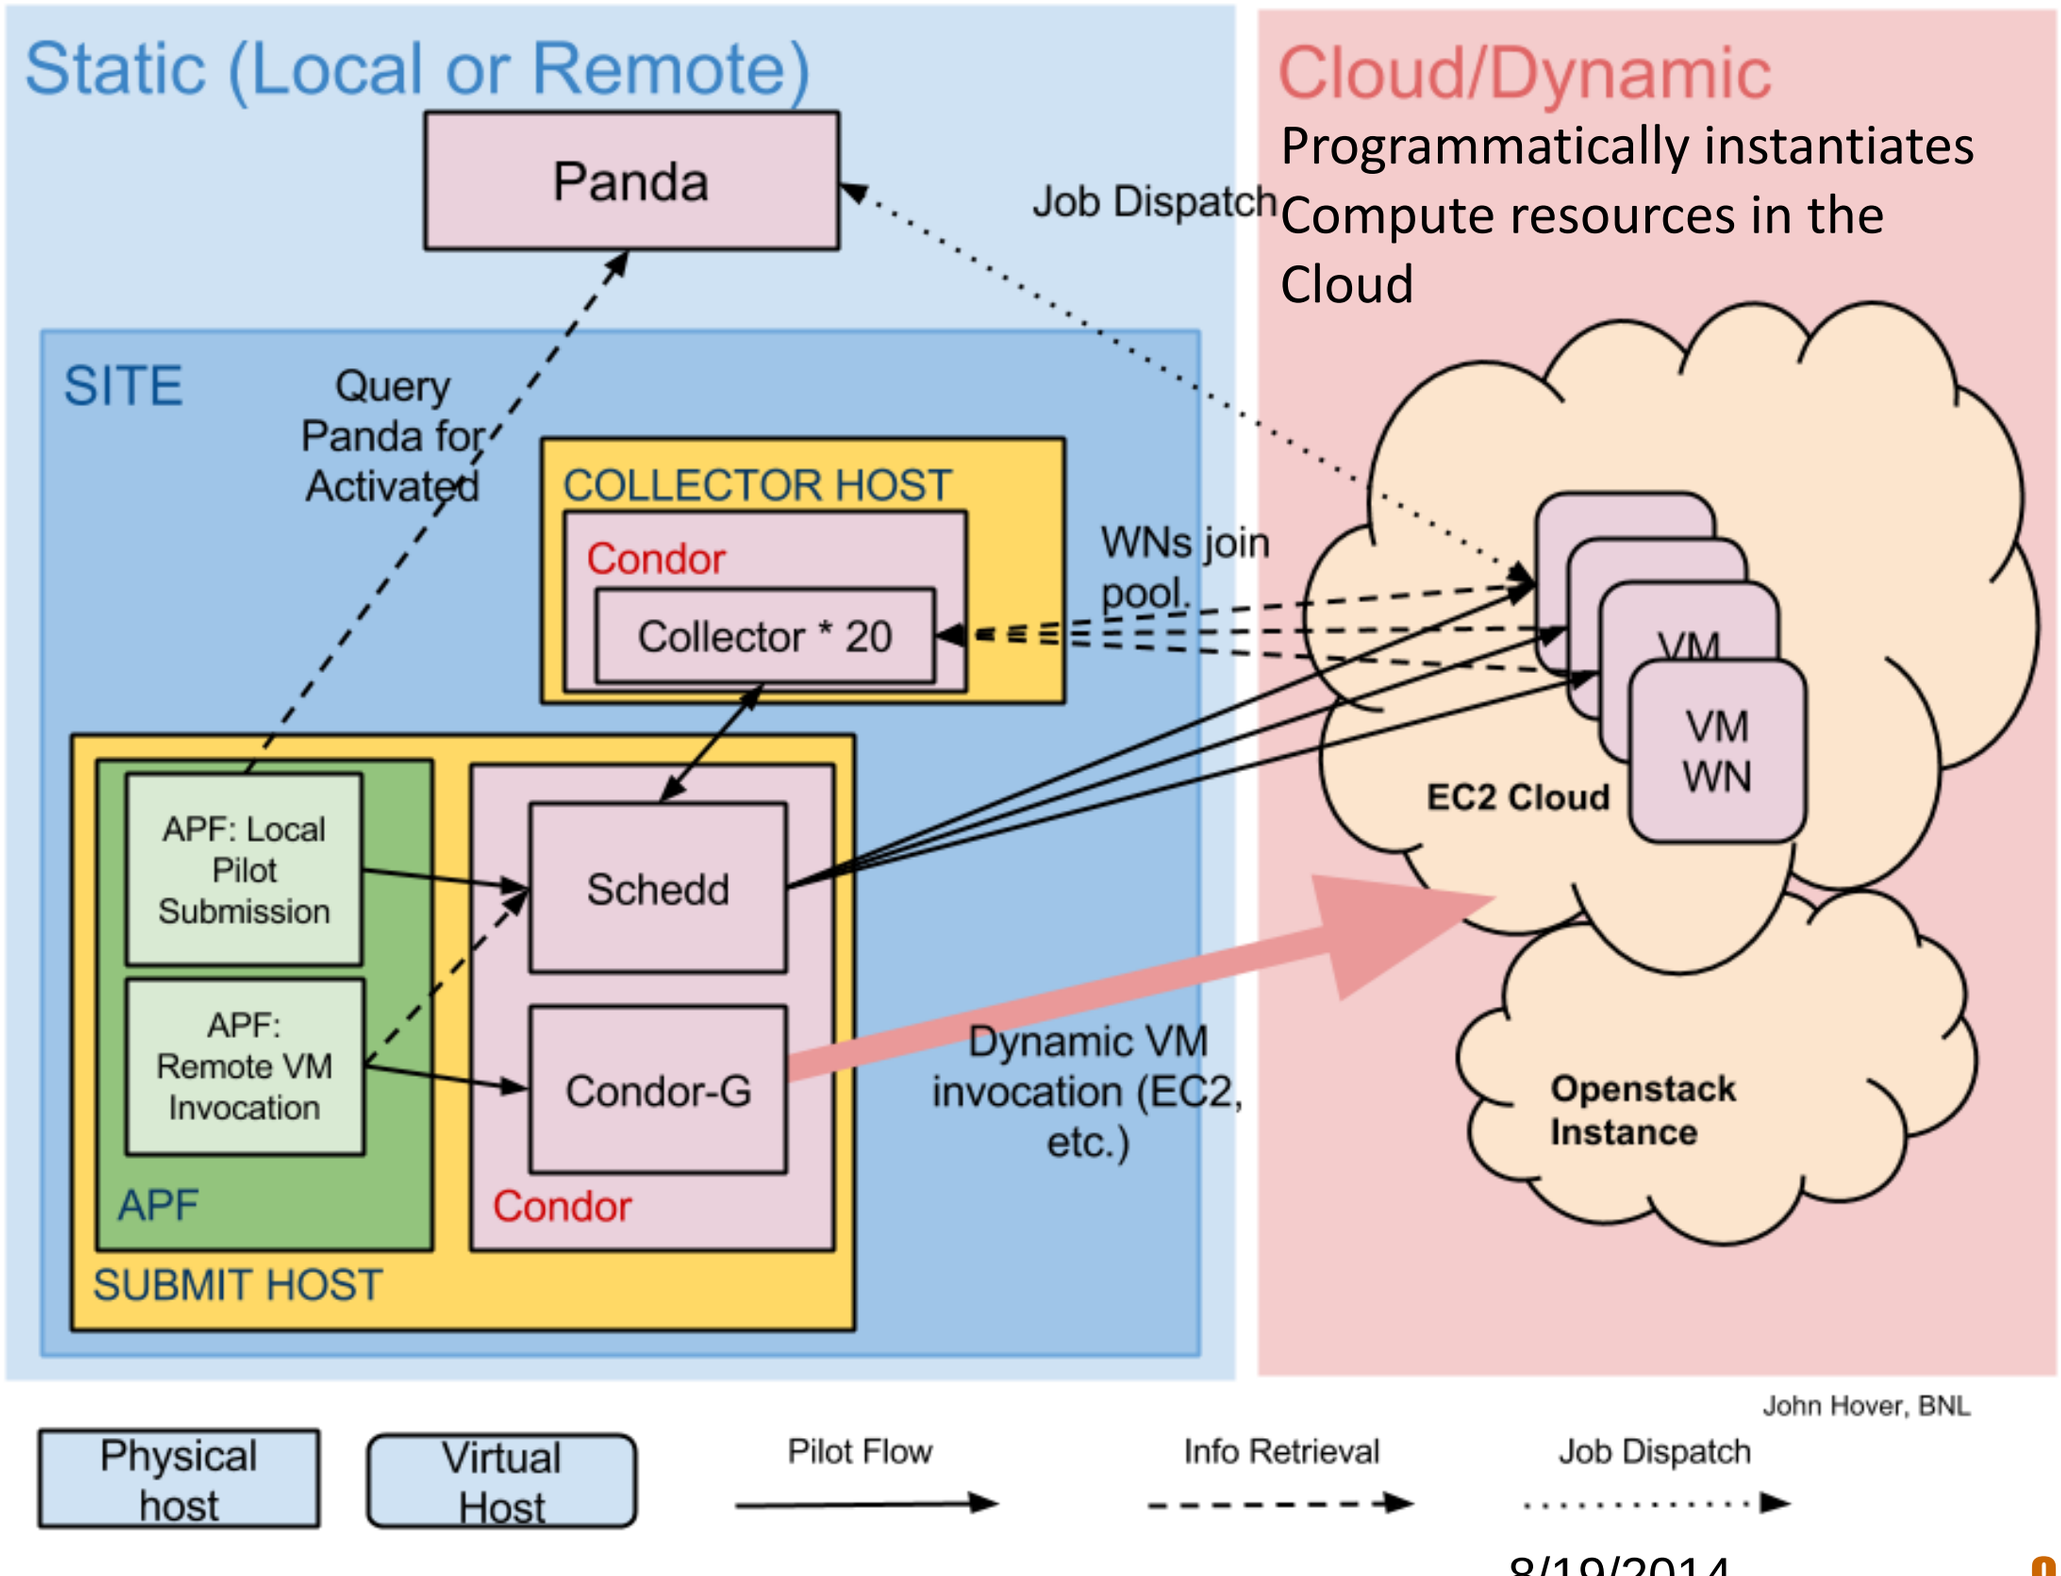

Open Science Grid

OSG's possible Role in Cloud Resource Provisioning

- For Academic and commercial/for fee cloud resources OSG could assume the role of a Reseller or Broker
  - Act as an allocation clearing house
  - Added Value for Users
    - OSG provides auditing and billing services
    - OSG acquires capacities on behalf of users, thereby taking care of the infrastructure users need
    - OSG becomes a reference for best practices, orchestration of virtualized resources, create VM images users can trust, provide monitoring and auditing capabilities
    - OSG provides End-User Support

- Use Cloud Technologies to virtualize OSG resources, e.g. with OpenStack
  - We won't have to maintain two different infrastructures any more
  - E.g. could be used to knit together LHC Tier-2 resources

Open Science Grid

Bringing in new Resources: HPC Usage by ATLAS and CMS

- **HPC resources are very valuable to HEP experimental computing, and HEP generally is a big user**
  - The US national HPC allocation for HEP is comparable to global LHC computing
- In the past, HEP experiments have used HPCs little – this is no longer true
  - HEP now competitor for allocated and opportunistic cycles
- HEP is compute-limited in their science – HPCs can enrich the science
- Increasing concurrency in HEP codes makes HPCs a natural target
- The convergence with HEP's computing comes from the HPC side as well
  - HPC has a growing number of data intensive use cases: future architectures have to take this into account
  - This was e.g. a message to NERSC in their last HEP requirements review
    - ATLAS Event Service expects to be an early user of NERSC's next machine
- HPCs have holes HEP can fill, even when their occupancy is high: this is the opportunistic usage ATLAS is targeting, collaborating particularly with ORNL
  - Just as there's room for sand in a full jar of rocks, there's room for HEP
- ATLAS and CMS porting appropriate applications (event generators, detector simulation)
  - ATLAS is extending PanDA to HPCs, and putting HPCs into ATLAS production

HPC Usage by ATLAS

ATLAS Event Generators and Full Simulation now running at ANL, NERSC and ORNL

- In the process of adding ATLAS Analysis Framework (Athena)
- Allocations at ANL (60M hours), ORNL/Titan (10M hours) and NERSC (2M hours)

➢ Potential to become solution adopted by other science domains

➢ OSG's role could be to help researchers with an allocation to manage and provision HPC resources to their workflows

- Alpgen v2.14 (fully ATLAS-ified)
- Pythia v8.18 & v6.428
- Geant4 v10
- Sherpa v2.1

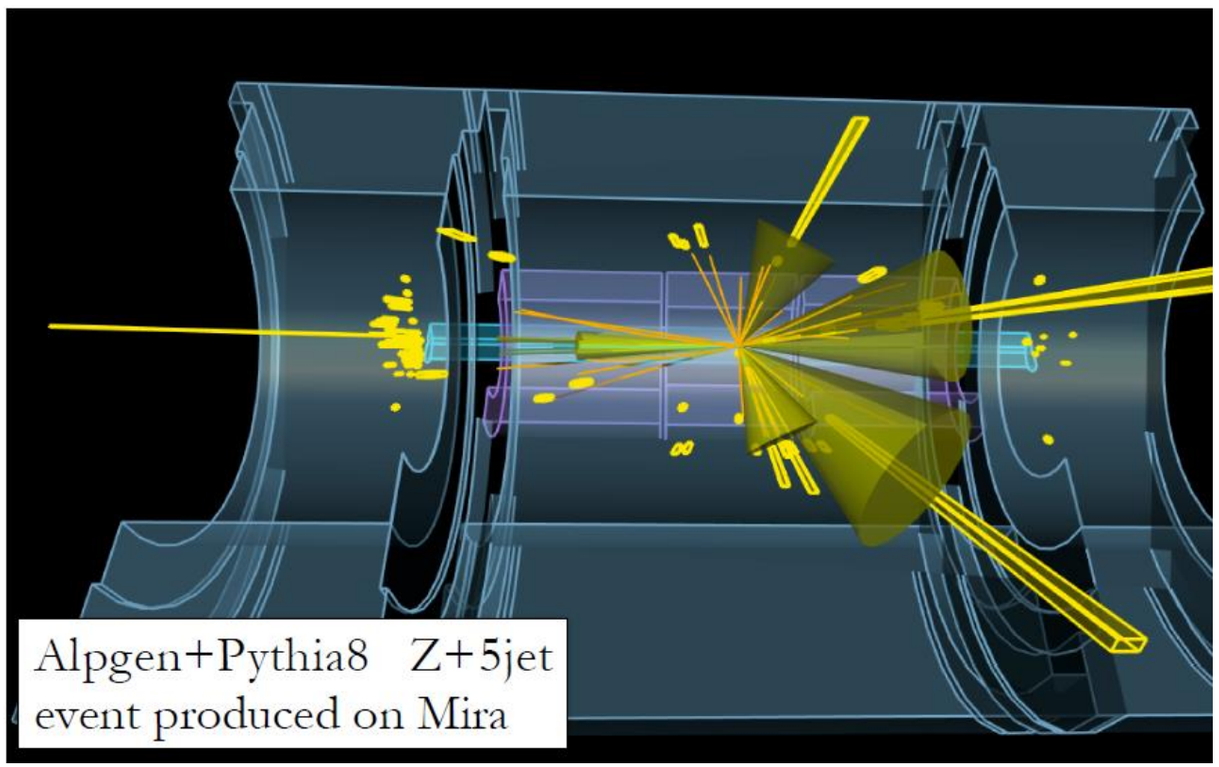

Alpgen+Pythia8 Z+5jet event produced on Mira

- **45,000 $Z \rightarrow \tau\tau$+5jet** Alpgen+Pythia events
- Would have required **12,250 24hr Grid jobs**
- **Saved 50k CPU-hours** compared to the Grid due to work duplication in grid workflow

OSG and HPC

Geant4

- Running v10, builtin parallel capability
- Run single Geant4 instance per node, tell Geant how many threads to use locally.
- Have tested an idealized CMS model on Vesta (Mira test machine)
- 16-cores/node each with 4 hardware threads (64 hardware threads)
- Performance tested with a few different thread counts per node:

CMS

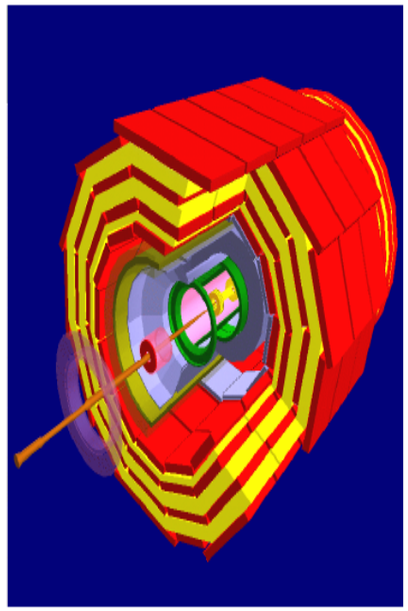

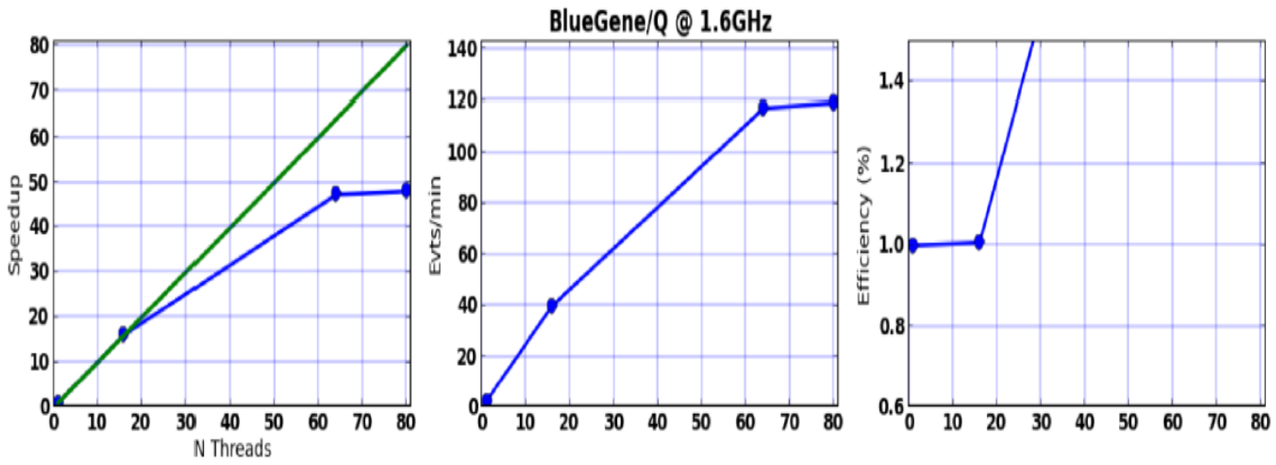

Andrea Dotti (SLAC)

- Extending the OSG eco system to the HPC community
  - OSG is an XD service provider,
  - how about XD providing resources to OSG communities?
  - interface to XSEDE installations
  - Others could probably Piggy-back on Atlas Connect XSEDE initiative
    - glideinWMS and BOSCO access to DOE facilities
- HPC as a data producer
  - e.g. cosmology simulation
  - e.g. HEP simulation

Open Science Grid

How About Data?

- Large VOs (e.g. LHC) offer users sophisticated data infrastructures
  - Dynamic data placements and high throughput robust data transfers
  - Data discovery and metadata tools
- OSG will be supporting such services to the communities
- Need to focus on Connectivity and Network Throughput
  - Including Networks and Federated Identity Management
  - Enabling remote data access, federated storage systems and allowing global access to locally managed data
- OSG has started with something that's based on local data and that's simple
  - OSG Connect already has a "stash" service
  - Anticipate adding data archiving service, e.g. at BNL and Fermilab
- These locally provided data services will become imminently useful through connectivity and federation (and possibly squid-based cashing)
  - A simple general data archival storage facility could have a big impact
  - Sustainable, dependable, accessible archive, offer open access to the data
  - Data preservation activities, long-term guardianship of data
    - DASPOS, NSF funded project and OSG Satellite on Data Preservation

Networking: Research Agenda

(as defined by Academia and Industry)

Future networks will integrate communication, computing, and storage resources in order to support a wide range of discovery techniques and environments.

- These will include domain-specific science gateways and portals, cloud-based workflows, high-performance and high-throughput computing models, and new data service capabilities.

Research and innovation in many domains are necessary to support this evolution.

- Major themes: virtualization, programmability, application integration, and complexity.
- Core question: can global research networks evolve into adaptive, self-organizing, programmable systems that quickly respond to requests of science applications?

Software Defined Networking is a very promising research area.

- Closed, inflexible, proprietary hardware/software systems are re-imagined as open, programmable hardware/software components.
- Easier software evolution, plus potential cost savings through cheaper hardware.
- Software-defined networks have the potential to enable great innovation inside the network, and to benefit science by facilitating virtualization, programmability, integration.

Network Monitoring Data Repository: Vision and Motivation

Open Science Grid

For Service Provisioning we should treat all the network layers in a holistic and integrated fashion

- Network Services vs. the “Network as a Resource”
  - The next generation of Advanced Networked Applications will require more “flexible control”, “scheduling”, and “deterministic performance” across all the resources in their ecosystem
  - This will require integration and co-scheduling across Network, Middleware, and Application level resources (compute, storage, domain specific instruments)

The Network needs to be available to application workflows as a first class resource in the OSG ecosystem

- While OSG is not expected to be directly involved in Network Research, OSG is an ideal place to provide information about the network as it is used by applications
  - Based on widely deployed perfSONAR probes developed by ESnet and I2, coordinated through the OSG Network Area (and throughout WLCG)
    - OSG aiming at having a role in Governance of perfSONAR development project
  - Inform about usage patterns and network characteristics to optimize application performance, to guide job brokerage in a distributed computing environment, and to optimize the network configuration and to provide trending information
  - Detect and Debug end-to-end transfer performance problems, use trending information for network capacity planning

Open Science Grid

OSG's Security Mission and Impact

Mission

- Protecting OSG users and resources from security breaches
- Preventing loss of effort and resources due to security problems
- Making OSG resources easily accessible to users without compromising their security
- Being a security hub
  - Disseminate security knowledge, best practices, and education

Impact to DOE and NSF Stakeholders

- Advances in identity and access control management
- Transition of Certificate service from DOEGrids CA to OSG CA
- Performed a thorough analysis of OSG infrastructure access control needs and options
  - Can we use a different, more user-friendly technology instead of Certificates?

Security: Challenges Ahead

- Goal is to shield users from certificates and offer more user friendly access control technologies.
- Certificate-free jobs was a big success
  - Has enabled OSG users to gain access to resources they could not use before
  - Have started working on accessing storage w/o Certificate
- Get more VOs to submit jobs without certificates
- Campus grid and federated identities
  - Utilize campus identities seamlessly in OSG.
- Further reducing the number of host certificates, aiming at getting rid of them entirely.

Next Generation Identity Management

OSG moving forward by collaborating with ASCR funded “eXtreme Scale Identity Management for Scientific Collaborations (XSIM) Project

Key Relationships

OSG Satellite

http://opensciencegrid.org/

Share common interests in better understand VOs in order to serve them. Key stakeholder of work.

Center for Trusted Scientific Cyberinfrastructure

http://trustedci.org/

NSF-funded project to help science CI project with Cybersecurity. Will be guided by XSIM’s work.

eXtreme Scale Identity Management (XSIM) for Scientific Collaborations

PI: Von Welch, Sr. Analysts: Craig Jackson, Bob Cowles – Indiana University / CACR

Objectives

- Understand the core elements of the trust relationship between scientific collaborations, resource providers and users.
- Understand how those trust relationships are (or desirably would be) expressed in the IdM systems.
- Validate the model and advance the state of practice through software and applied research.

- **Ultimately: Enable the next generation of trustworthy extreme scale scientific collaborations by understanding and formalizing a model of identity management (IdM) that includes the collaboratory**.

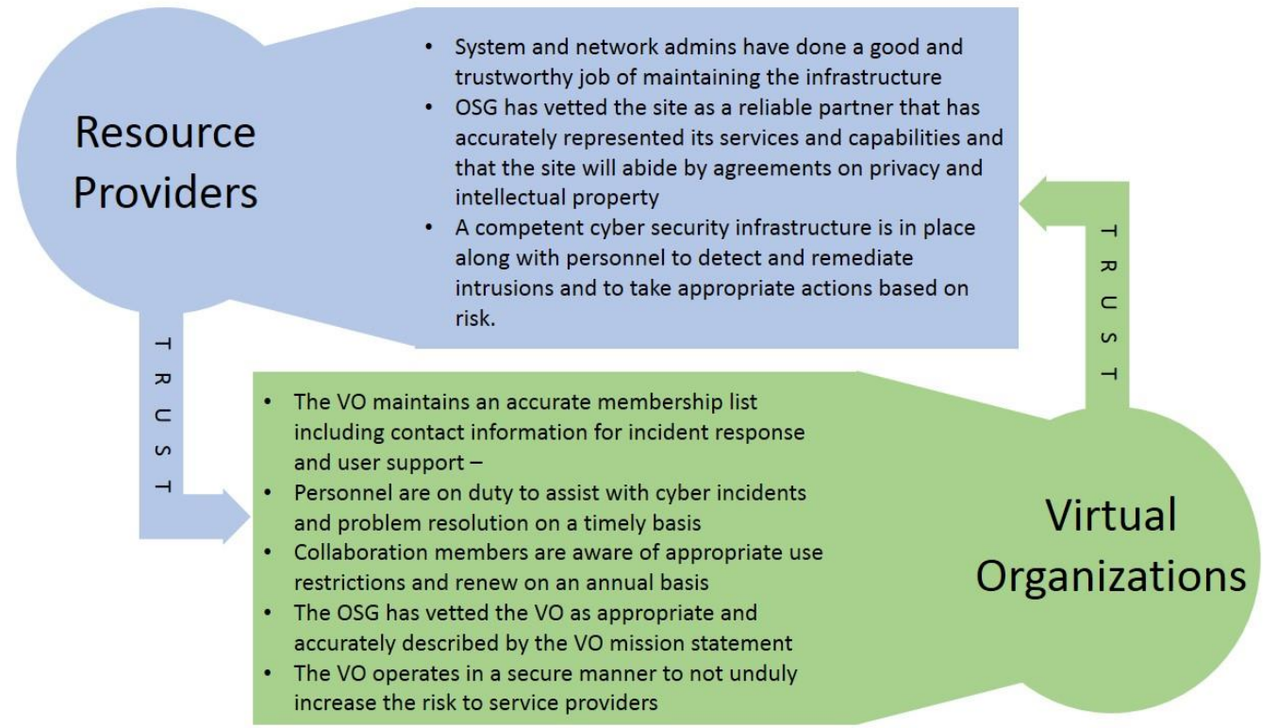

Approach:

- Analyze implementations – study literature of the different collaboratory IdM approaches and interview members of the community.
- Discern the trust model each implementation strived for.
- Enumerate the different relationships between collaborations and their resource providers and the evolution of each (lessons learned)
- Propose a model for an evolutionary step in IdM that describes trust relationships between collaborations, resource providers and users.
- Model must be understandable and useful to non-IdM experts and is accepted by resource providers.
- Refine and extend model based on feedback and experience.

Schedule:

- Project start: September, 2012

- Y1: Publication and presentation of document describing the results of the interviews and the IdM model. (Targeting CHEP and eScience.)

- Y2: Develop software implementing the model and revise the model based on feedback and experience from initial field tests

- Y3: Further development of the model user trust relationships; documentation and packaging of the software.

Open Science Grid

Conclusions

- OSG's Emphasis needs to be on
  - Fast transient resource contribution and provisioning
  - Authentication Management
  - Providing Data and Dynamic Storage
  - Making Network Monitoring Data an Integral Part of the OSG ecosystem
- "OSG Connect" unites most of the above and is the strategic direction of OSG