Service-Oriented Science: Scaling eScience Impact

Or, "Science 2.0"

Ian Foster

Computation Institute

Argonne National Lab & University of Chicago

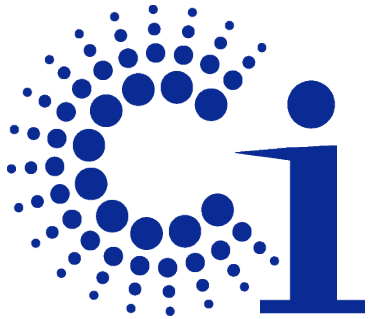

"Web 2.0"

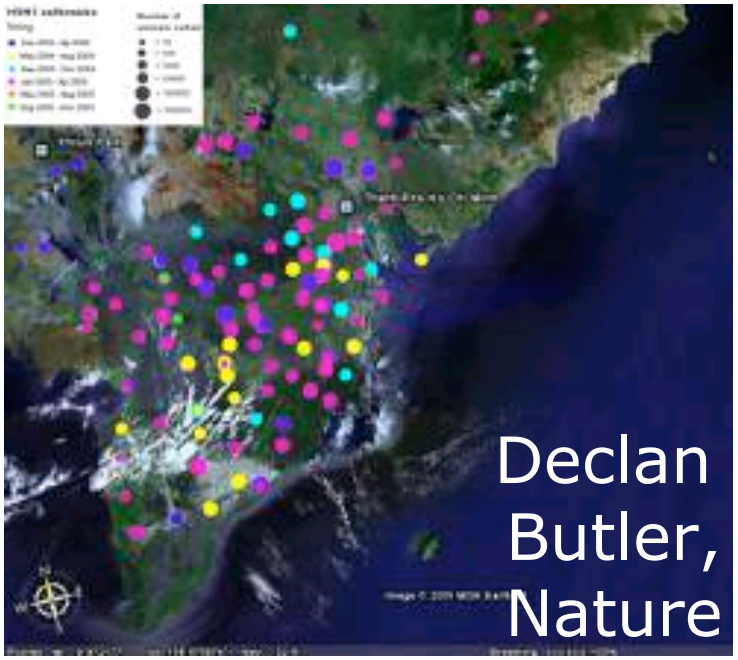

- Software as services
  - Data- & computation-rich network services
- Services as platforms
  - Easy composition of services to create new capabilities ("mashups")—that themselves may be made accessible as new services
- Enabled by massive infrastructure buildout
  - Google projected to spend $1.5B on computers, networks, and real estate in 2006
  - Many others are spending substantially
- Paid for by advertising

the globus alliance
www.globus.org

Science 2.0: E.g., Virtual Observatories

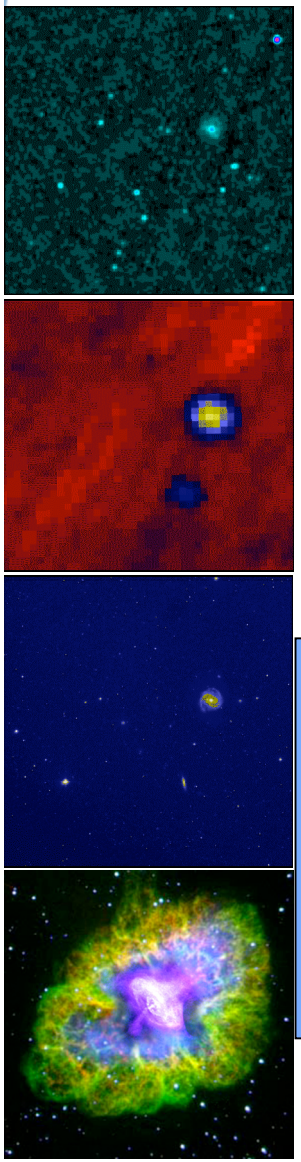

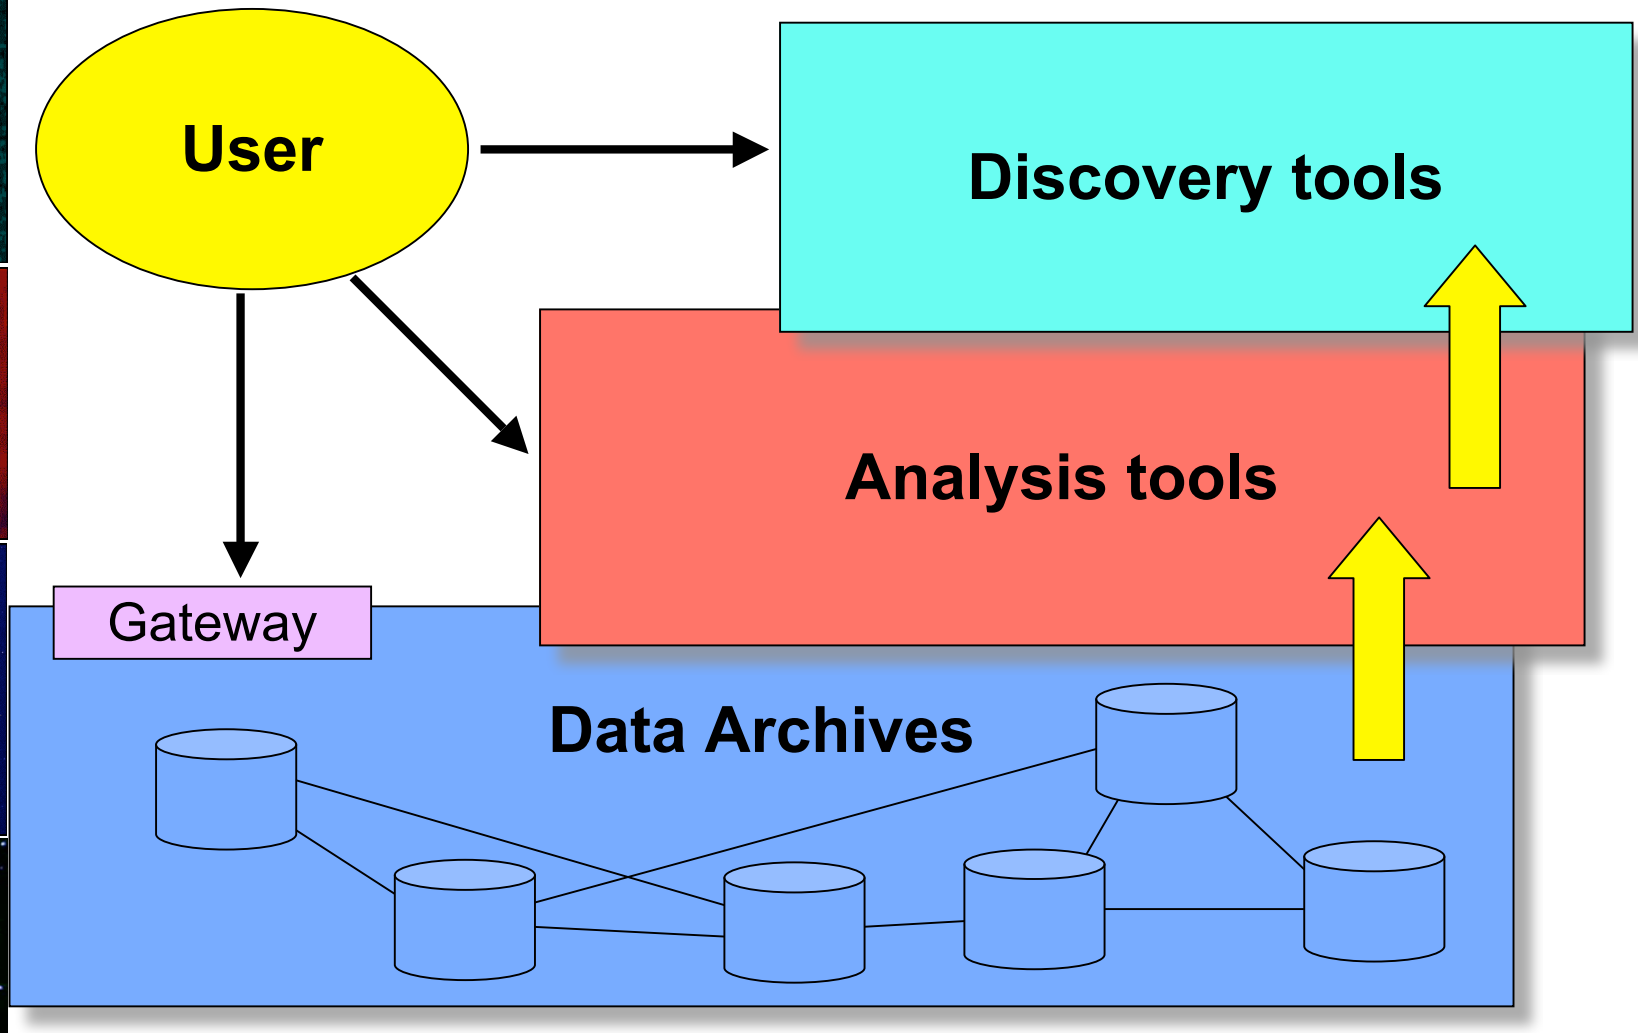

Figure: S. G. Djorgovski

the globus alliance
www.globus.org

Science 2.0

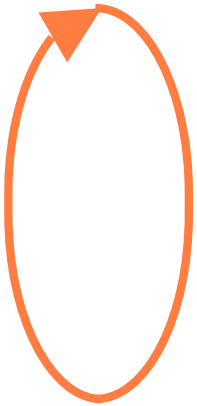

People **create** services (data or functions) …
which I **discover** …
& **compose** to create a new function **…**
and then **publish** as a new service.

!!

→ *I find "someone else" to **host** services, so I don't have to become an expert in operating services & computers!*

→ *I hope that this "someone else" can **manage** security, reliability, scalability, …*

TeraGrid™
EMPOWERING DISCOVERY

"Service-Oriented Science", *Science*, 2005

Creating Services

- Take an "application"
  - An arbitrary executables
  - A procedure in Java or another language
  - A Web service
- And provide a "Web Service" interface
  - Address authentication & authorization
  - Pass input data (XML, files, …?)
  - Invoke the application
  - Permit monitoring & control
  - Return output data (XML, files, …?)

Creating Services: E.g., GT4 Web Services Container

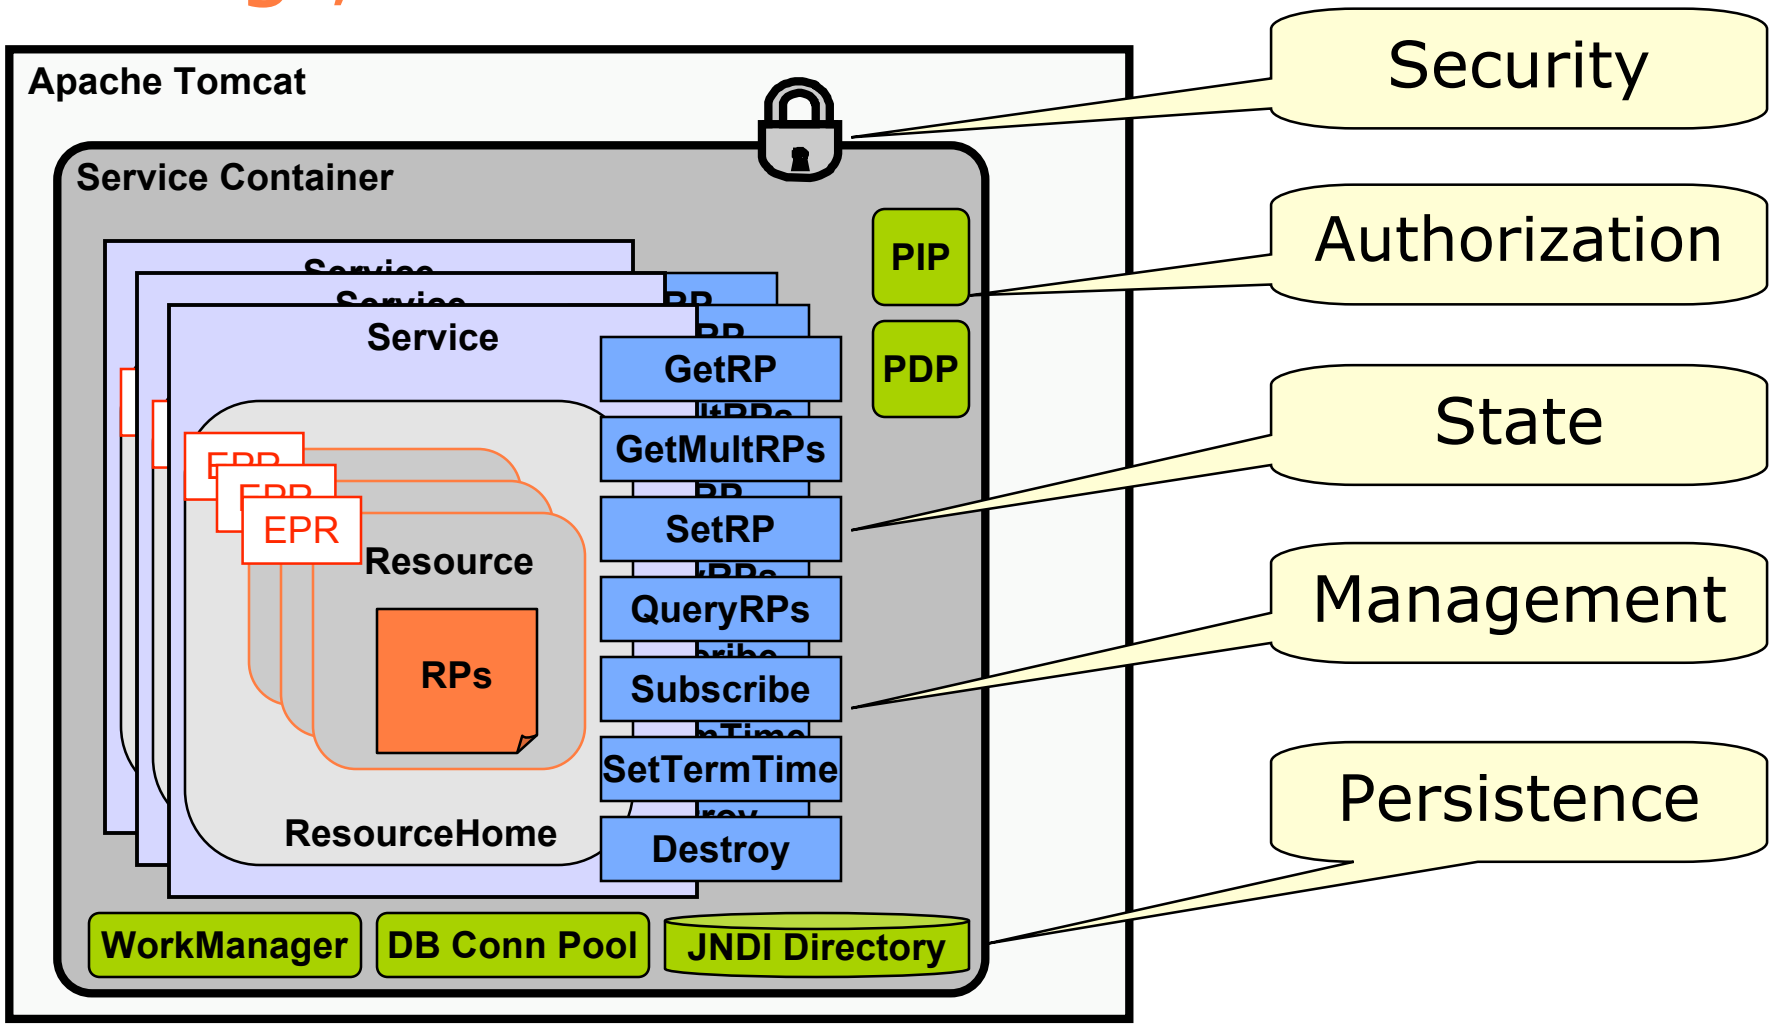

Creating Services: E.g., Introduce Authoring Tool

- Define service
- Create skeleton
- Discover types
- Add operations
- Configure security
- Modify service

  → ***targets GT4***

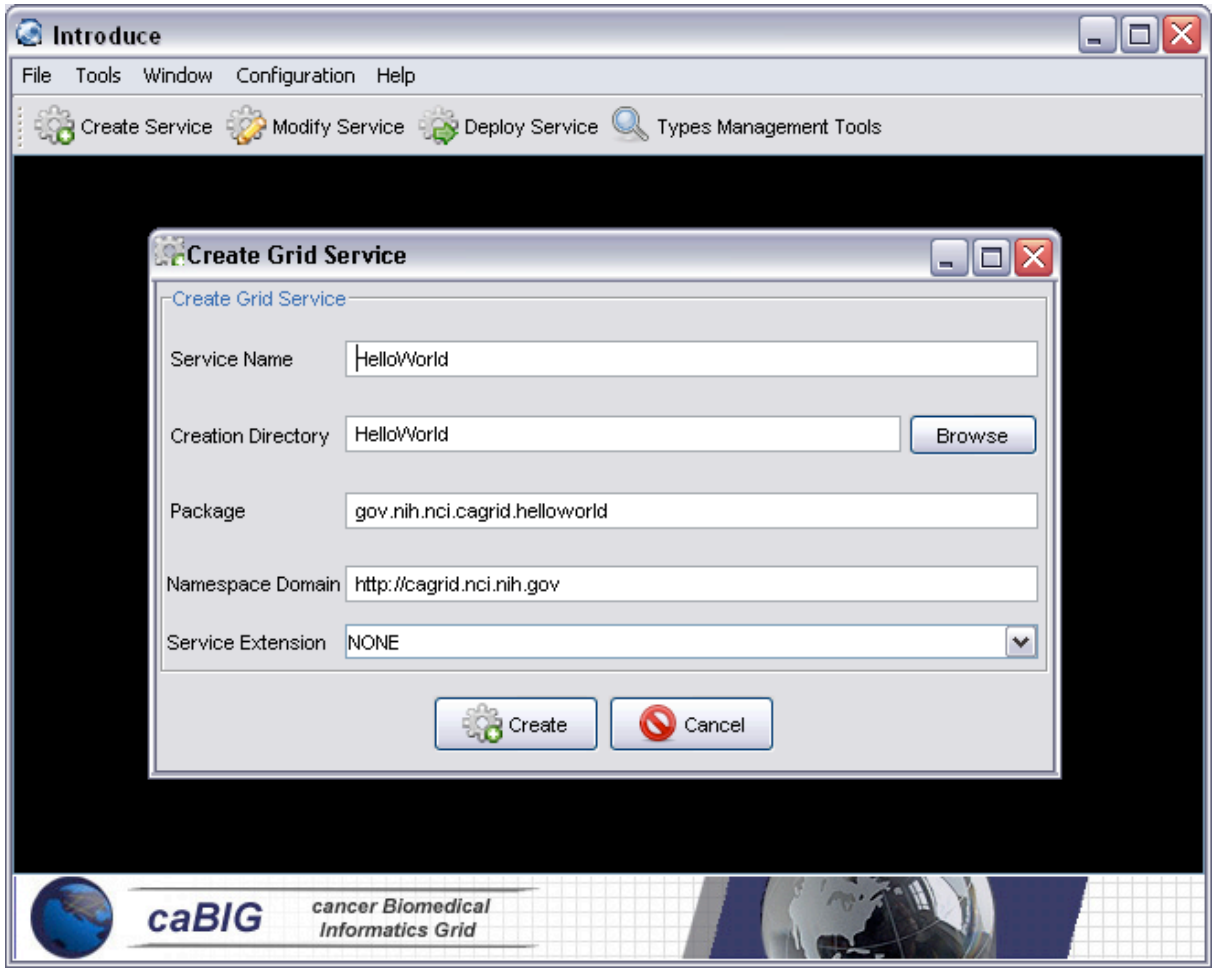

See also: SOAPLab, OPAL, pyGlobus, Gannon, etc.

Introduce: Hastings, Saltz, et al., Ohio State University

Composing Web Services: E.g., BPEL Workflow System

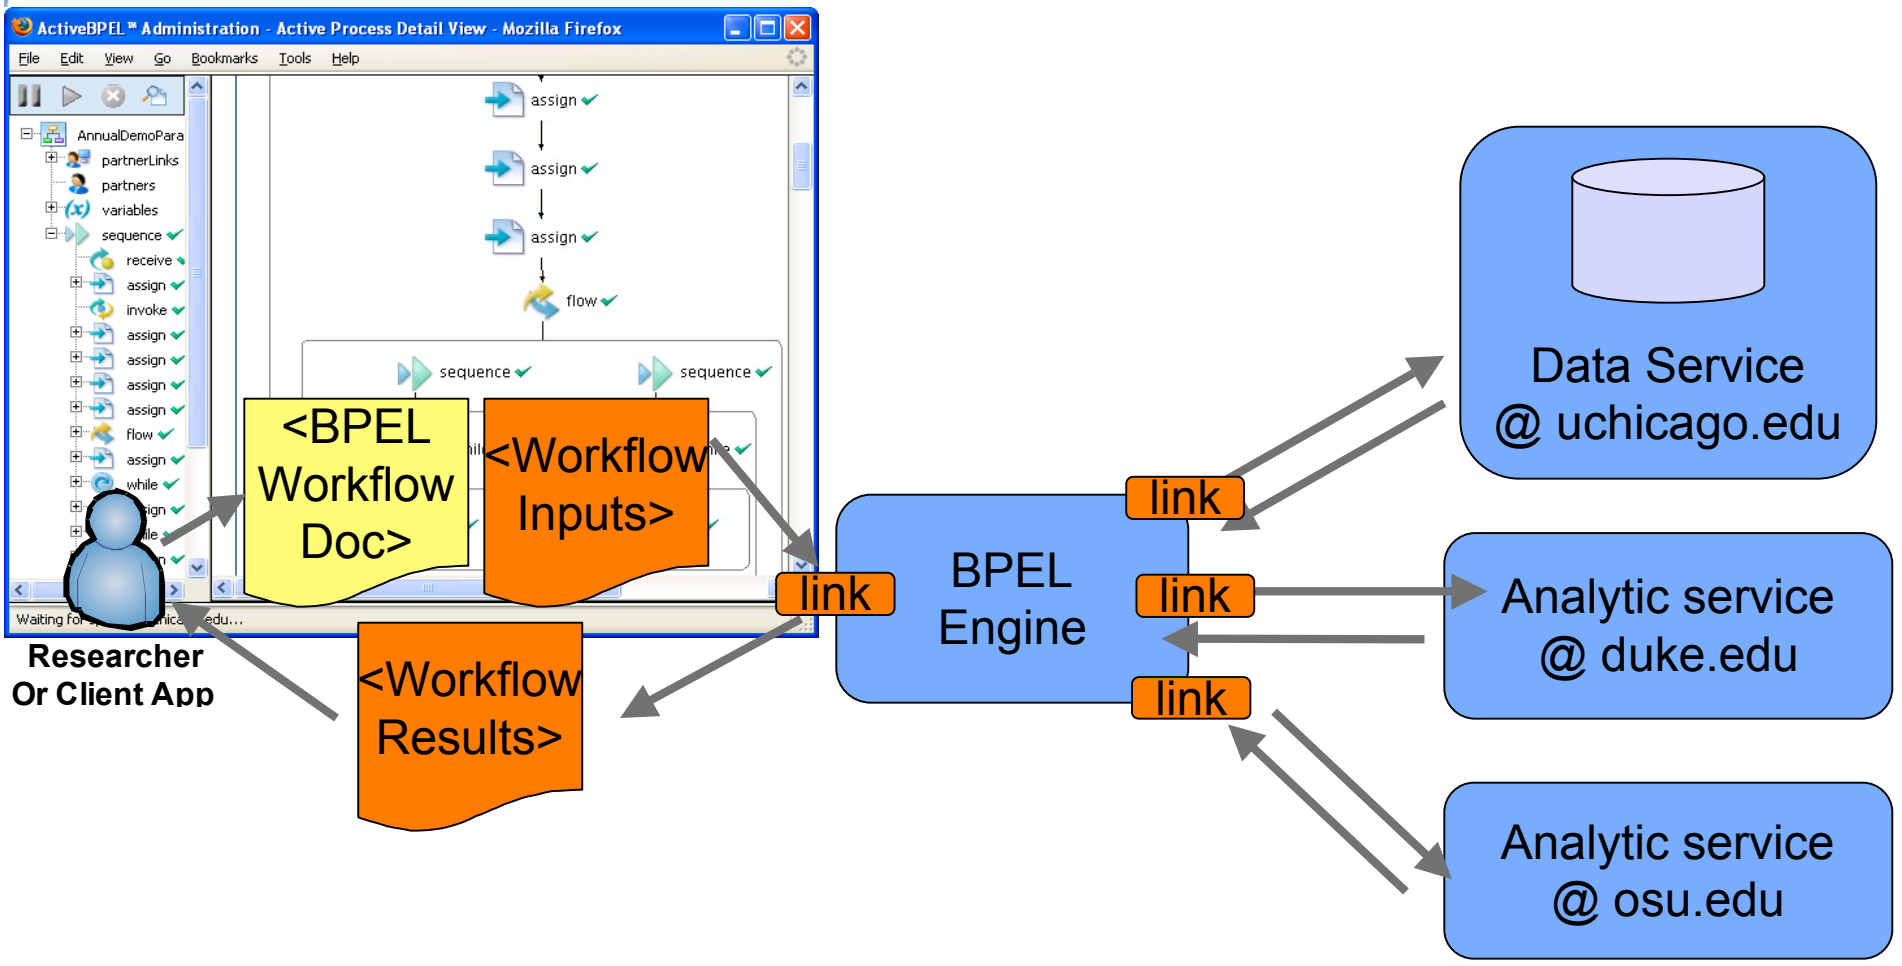

caBIG™ cancer Biomedical Informatics Grid™ an initiative of the National Cancer Institute

caBiG: https://cabig.nci.nih.gov/; BPEL work: Ravi Madduri et al.

The Globus-Based LIGO Data Grid

LIGO Gravitational Wave Observatory

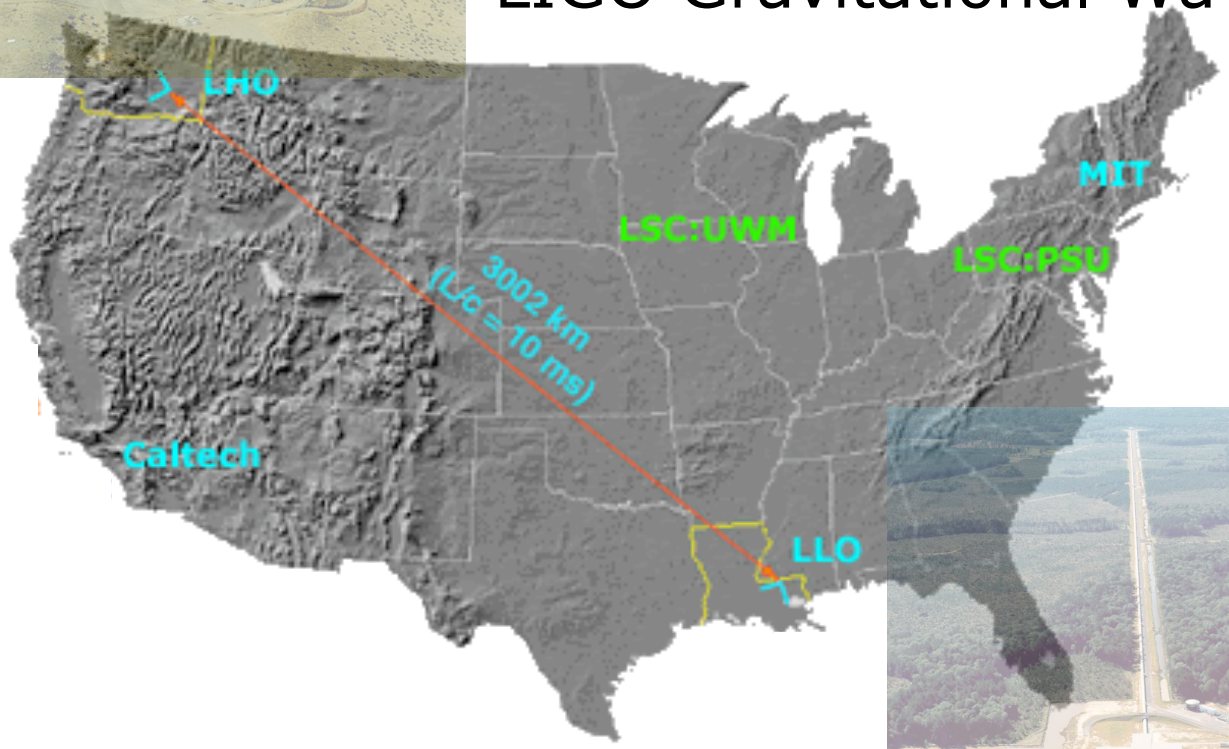

Replicating >1 Terabyte/day to 8 sites

>40 million replicas so far

MTBF = 1 month

www.globus.org/solutions

Data Replication Service

- Pull "missing" files to a storage system

Data Movement

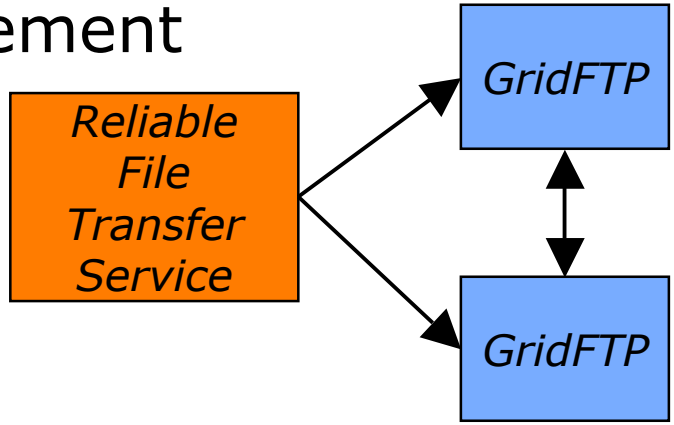

"Design and Implementation of a Data Replication Service Based on the Lightweight Data Replicator System," Chervenak et al., 2005

Data Replication Service

- Pull "missing" files to a storage system

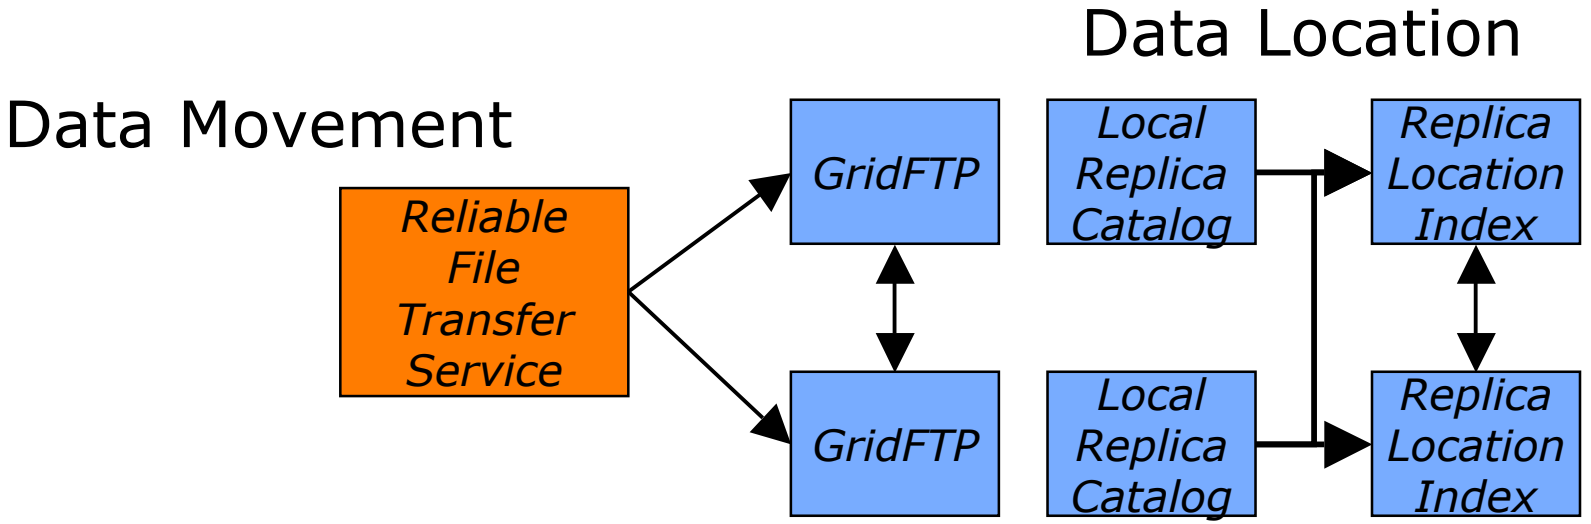

"Design and Implementation of a Data Replication Service Based on the Lightweight Data Replicator System," Chervenak et al., 2005

Data Replication Service

- Pull "missing" files to a storage system

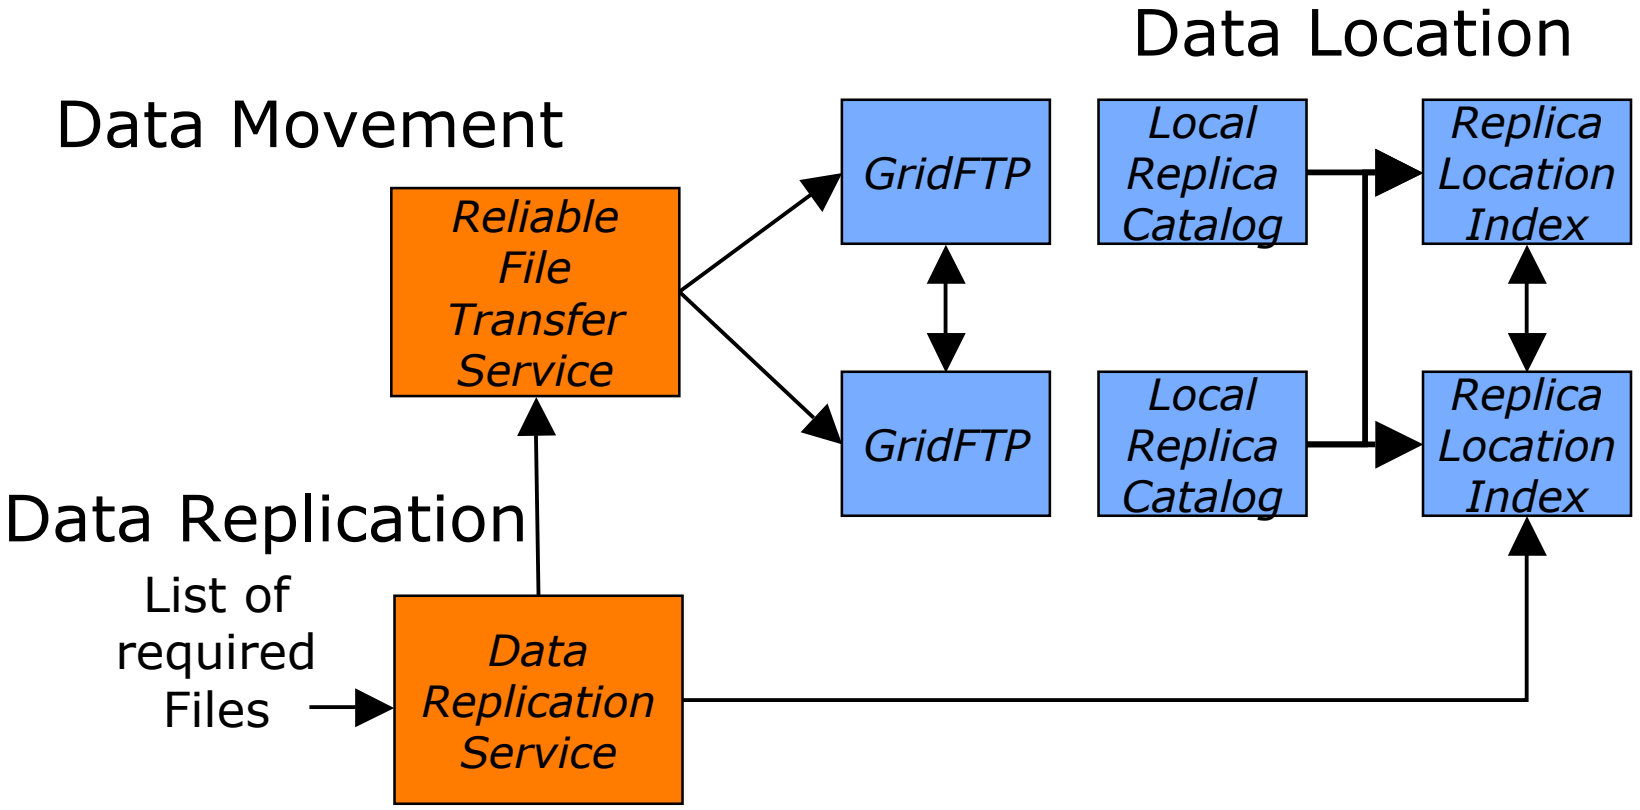

"Design and Implementation of a Data Replication Service Based on the Lightweight Data Replicator System," Chervenak et al., 2005

the globus alliance
www.globus.org

Science 2.0

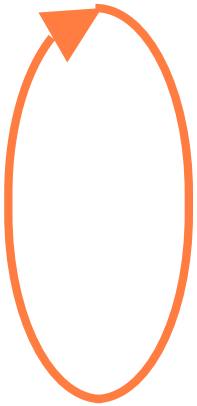

People **create** services (data or functions) …
which I **discover** …
& **compose** to create a new function **…**
and then **publish** as a new service.

!!

→ *I find "someone else" to **host** services, so I don't have to become an expert in operating services & computers!*

→ *I hope that this "someone else" can **manage** security, reliability, scalability, …*

TeraGrid™
EMPOWERING DISCOVERY

"Service-Oriented Science", *Science*, 2005

The Importance of "Hosting" and "Management"

Tell me about this star

Tell me about these 20K stars

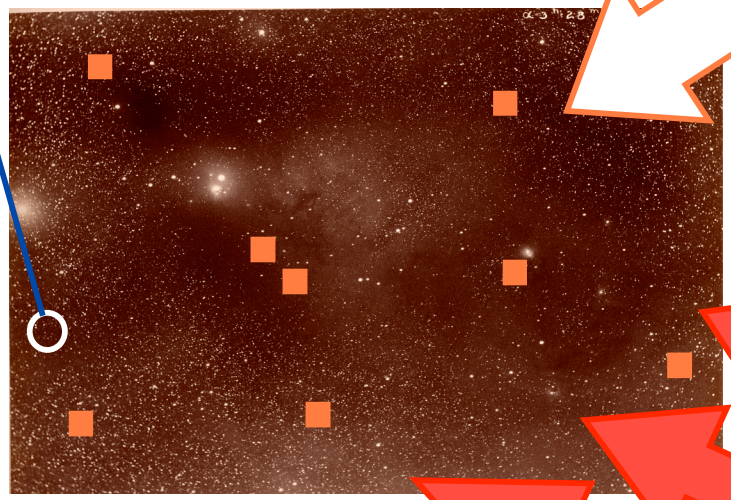

Support 1000s of users

E.g., Sloan Digital Sky Survey, ~10 TB; others much bigger

Integration & Decomposition: A Two-Dimensional Problem

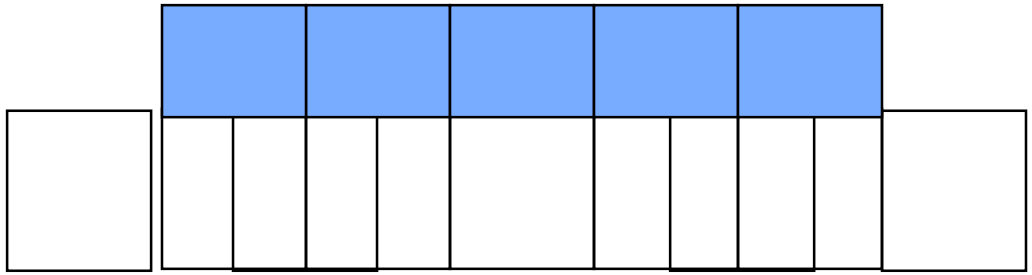

Function

Resource

- **Decompose** across network

  Clients **integrate** dynamically
  - Select & compose services
  - Select "best of breed" providers
  - Publish result as new services

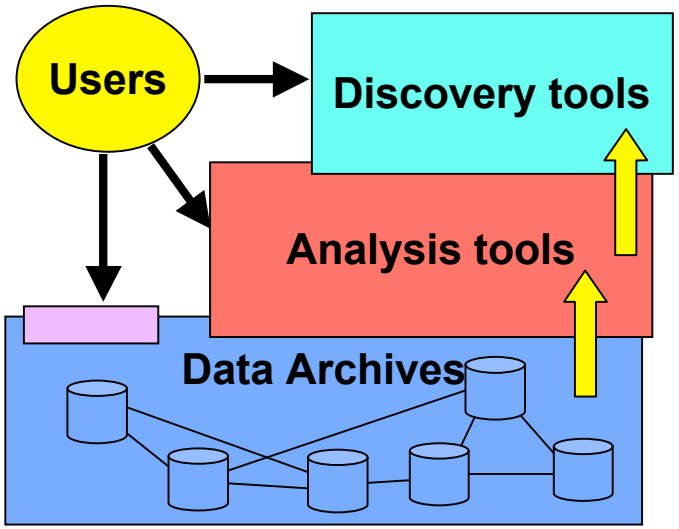

- Decouple **resource** & **service** providers

Fig: S. G. Djorgovski

the globus alliance
www.globus.org

Service Hosting & Management

- Negotiate service level agreements
- Delegate and deploy capabilities/services
- Provision to deliver defined capability
- Configure environment
- Host layered functions

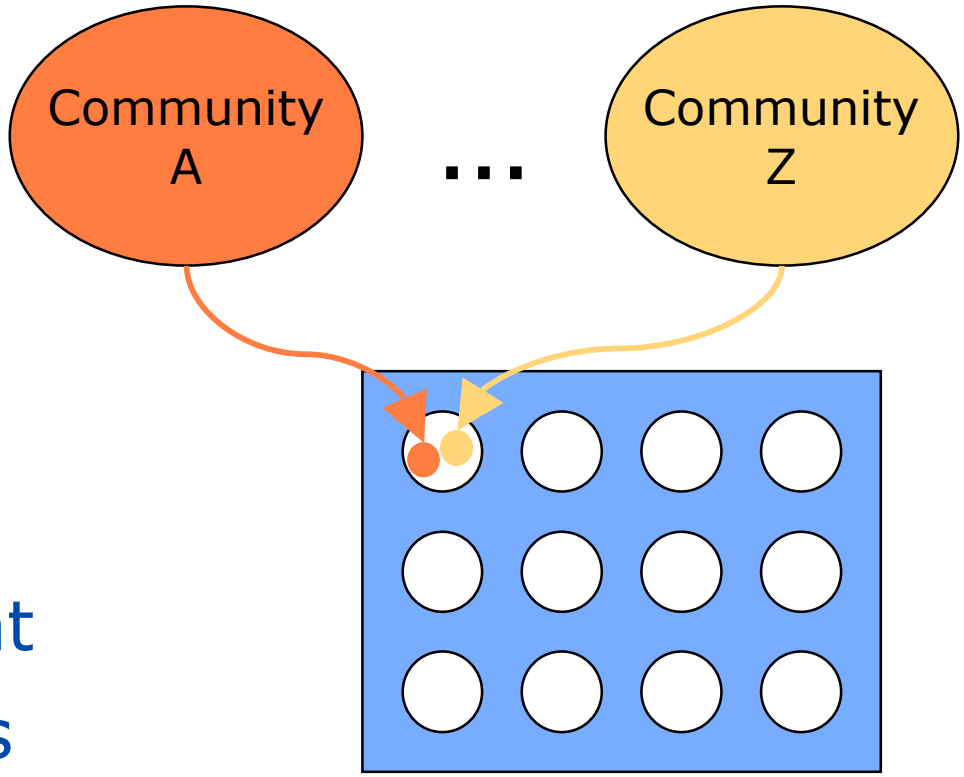

WSRF (or WS-Transfer/WS-Man, etc.), Globus GRAM, Virtual Workspaces

Virtualizing Existing Services into a VO

- Establish service agreement with service
  - E.g., WS-Agreement
- Delegate use to VO user

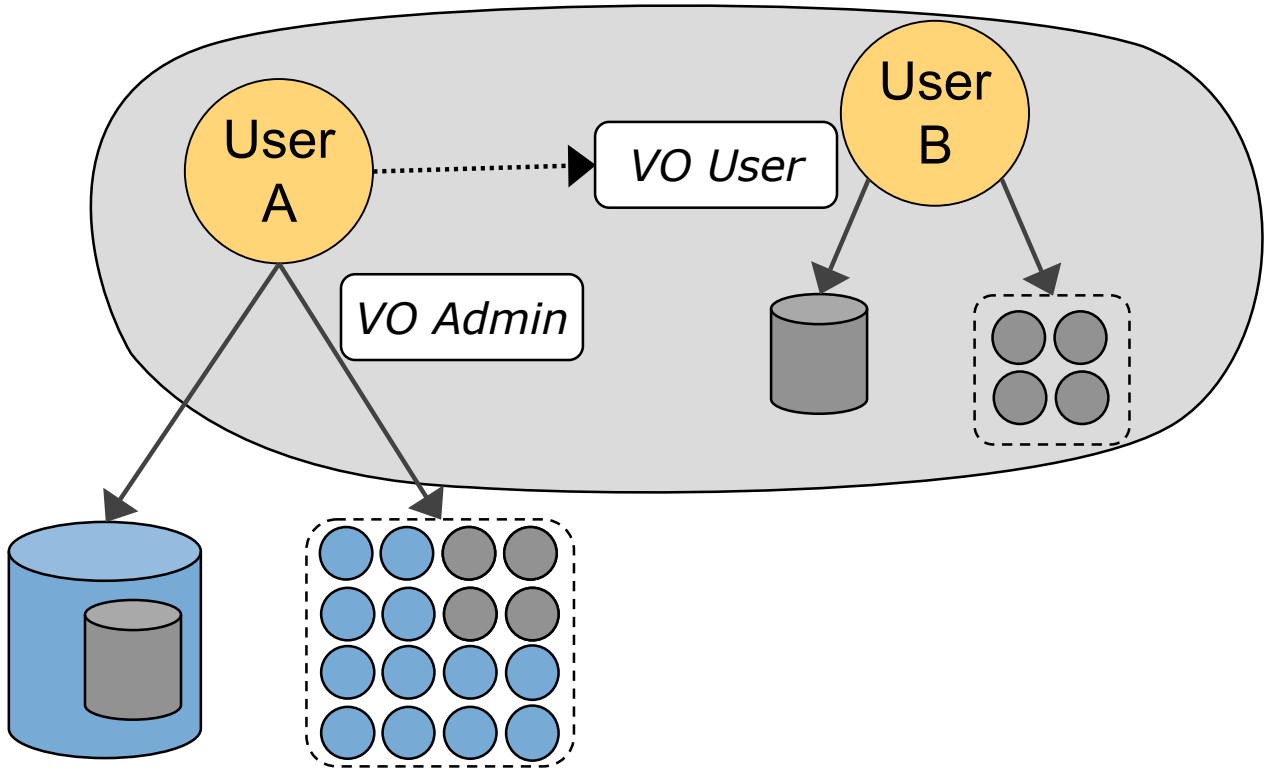

the globus alliance
www.globus.org

Example: Cybershake

Calculate hazard curves by generating synthetic seismograms from estimated rupture forecast

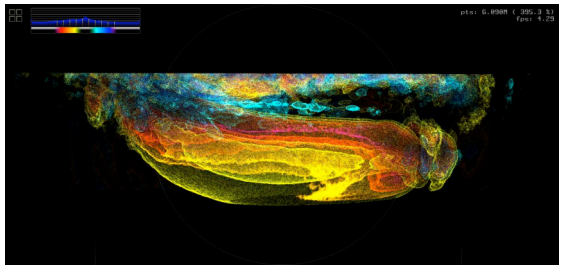

Strain Green Tensor

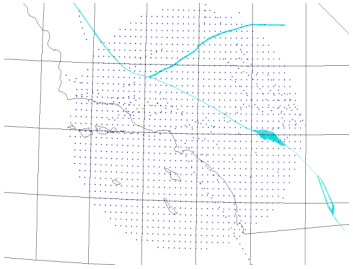

Rupture Forecast

Hazard Map

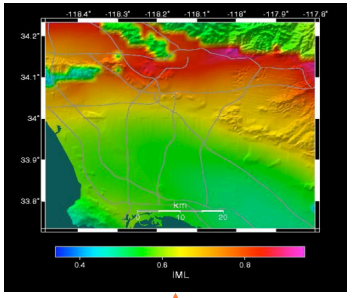

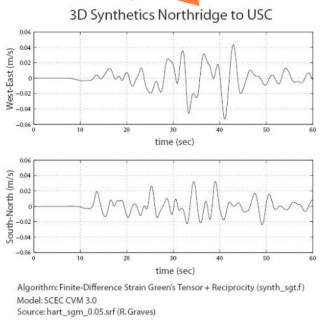

Synthetic Seismogram

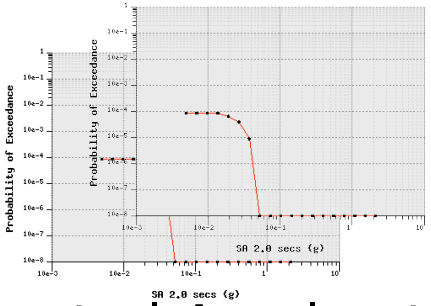

Spectral Acceleration

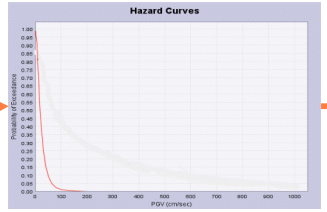

Hazard Curve

Tom Jordan et al., Southern California Earthquake Center

Enlisting TeraGrid Resources

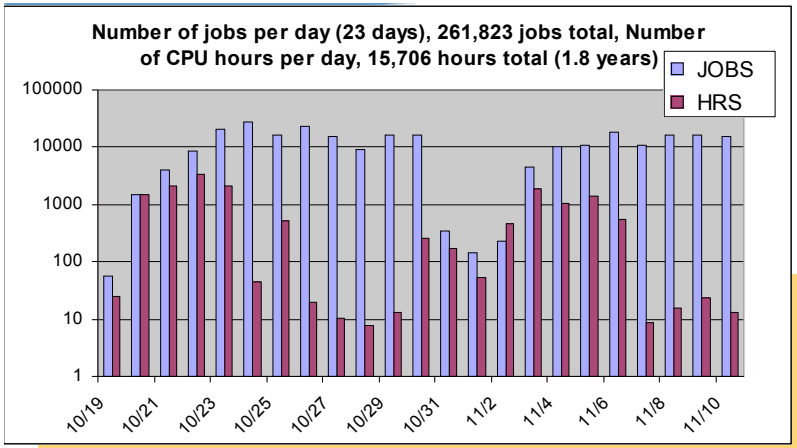

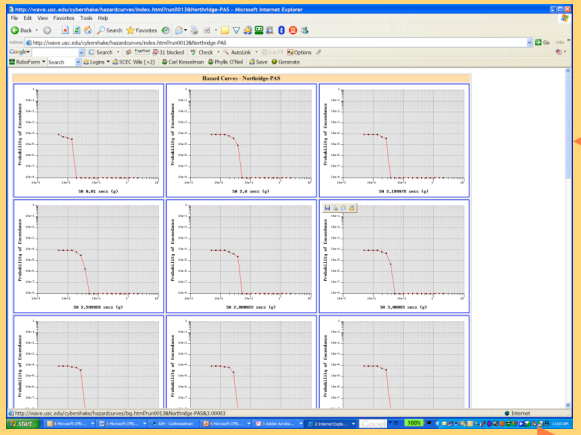

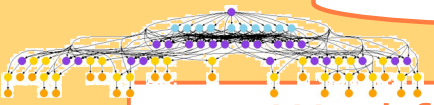

Provenance Catalog

Data Catalog

Workflow Scheduler/Engine

VO Service Catalog

SCEC Storage

TeraGrid Storage

VO Scheduler

TeraGrid Compute

20 TB, 1.8 CPU-year

Ewa Deelman, Carl Kesselman, et al., USC Information Sciences Institute

Dynamic Provisioning: Astro Portal Stacking Service

- Purpose
  - On-demand "stacks" of random locations within ~10TB dataset
- Challenge
  - Rapid access to 10-10K "random" files
  - Time-varying load
- Solution
  - Dynamic acquisition of compute, storage

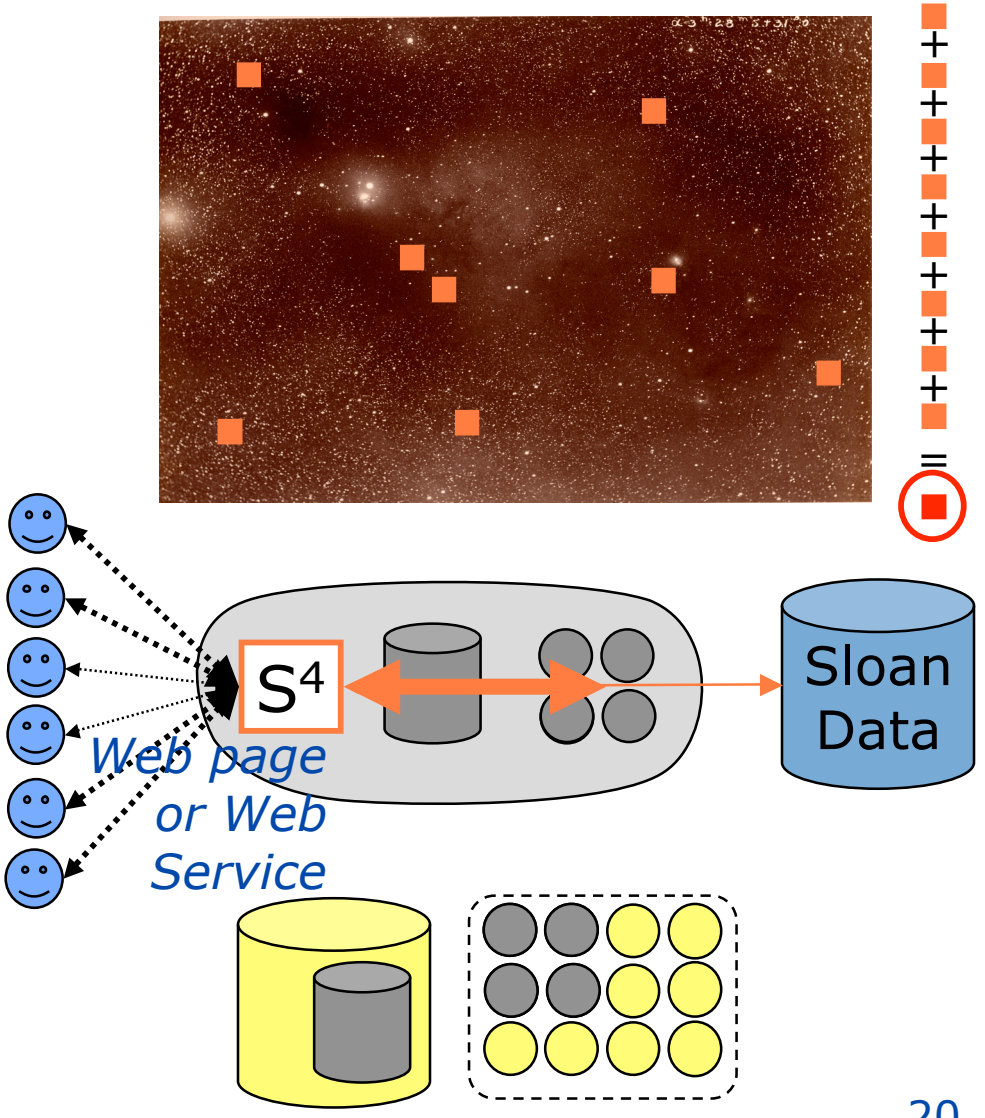

the globus alliance
www.globus.org

Astro Portal Stacking Performance (LAN GPFS)

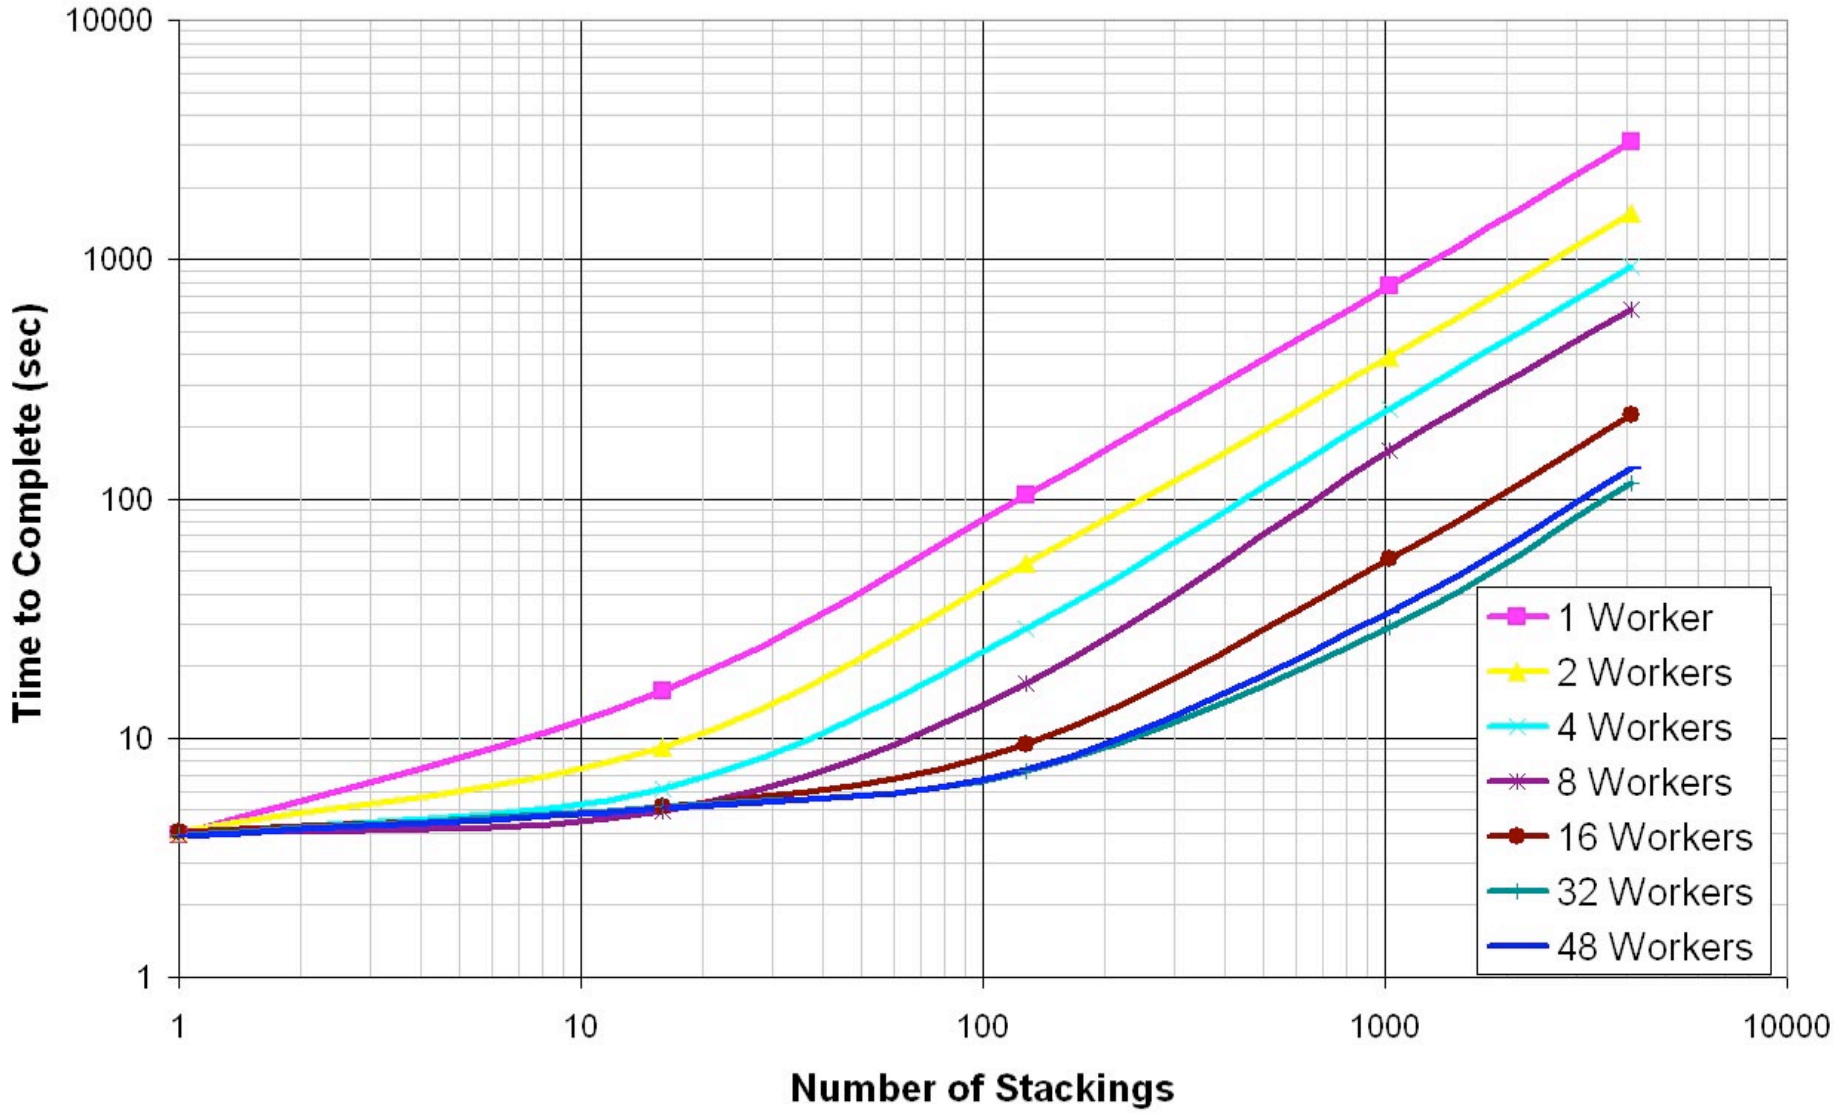

Virtual Workspaces (Kate Keahey et al.)

- GT4 service for the creation, monitoring, & management of **virtual workspaces**
- High-level workspace description
- Web Services interfaces for monitoring & managing
- Multiple implementations
  - Dynamic accounts
  - Xen virtual machines
  - (VMware virtual machines)
- Virtual clusters as a higher-level construct

Composing Resources … Composing Services

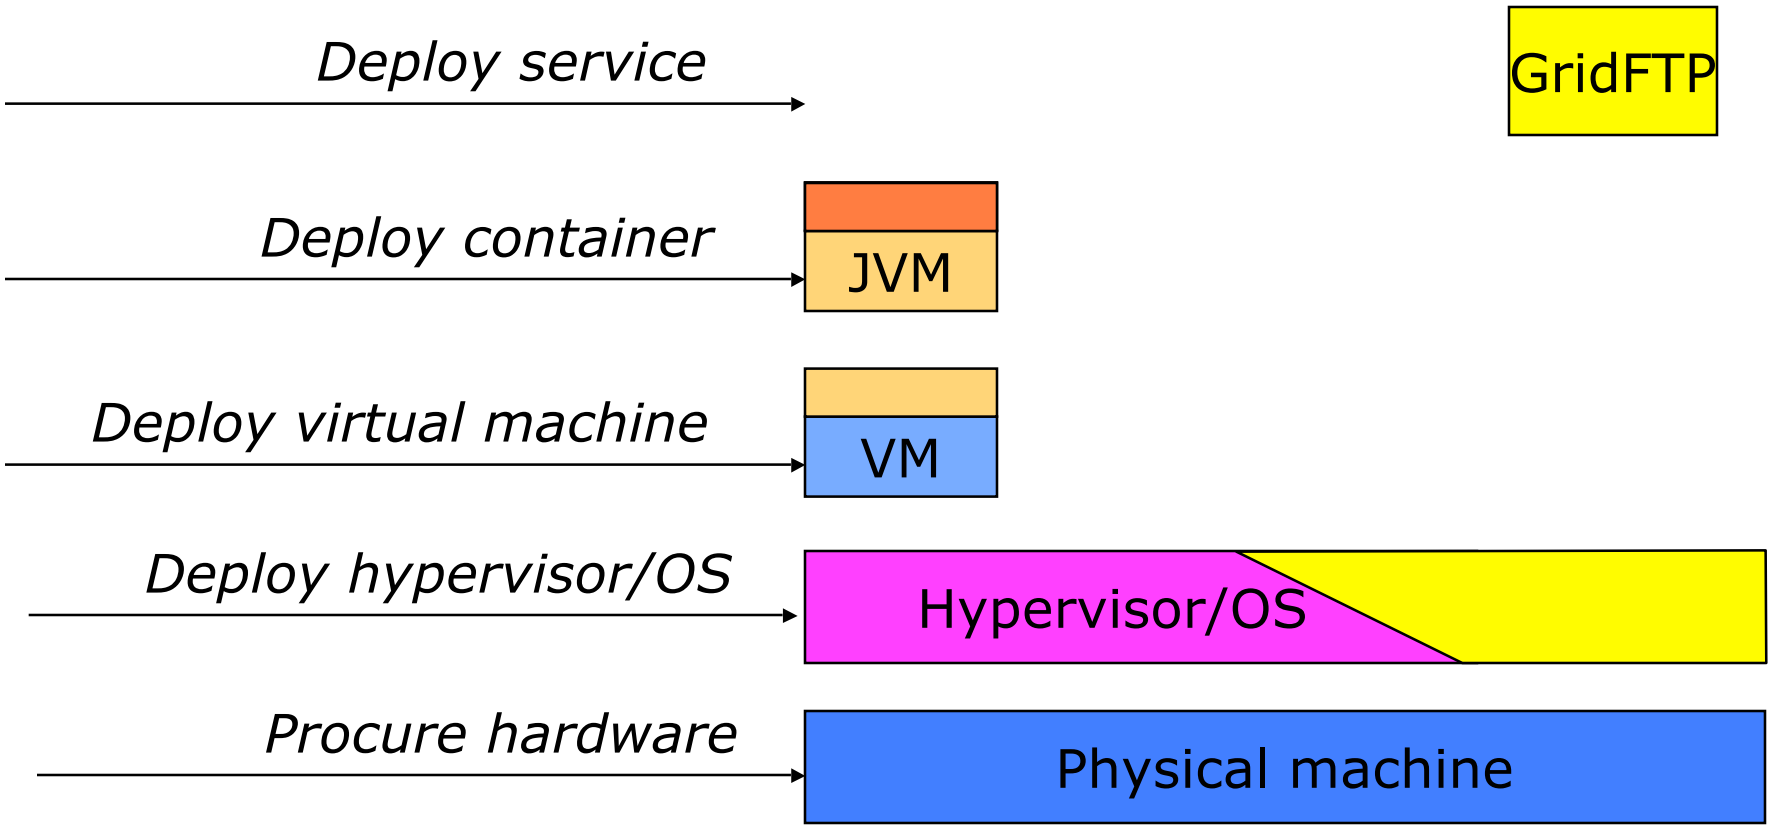

State exposed & access uniformly at all levels
Provisioning, management, and monitoring at all levels

Composing Resources … Composing Services

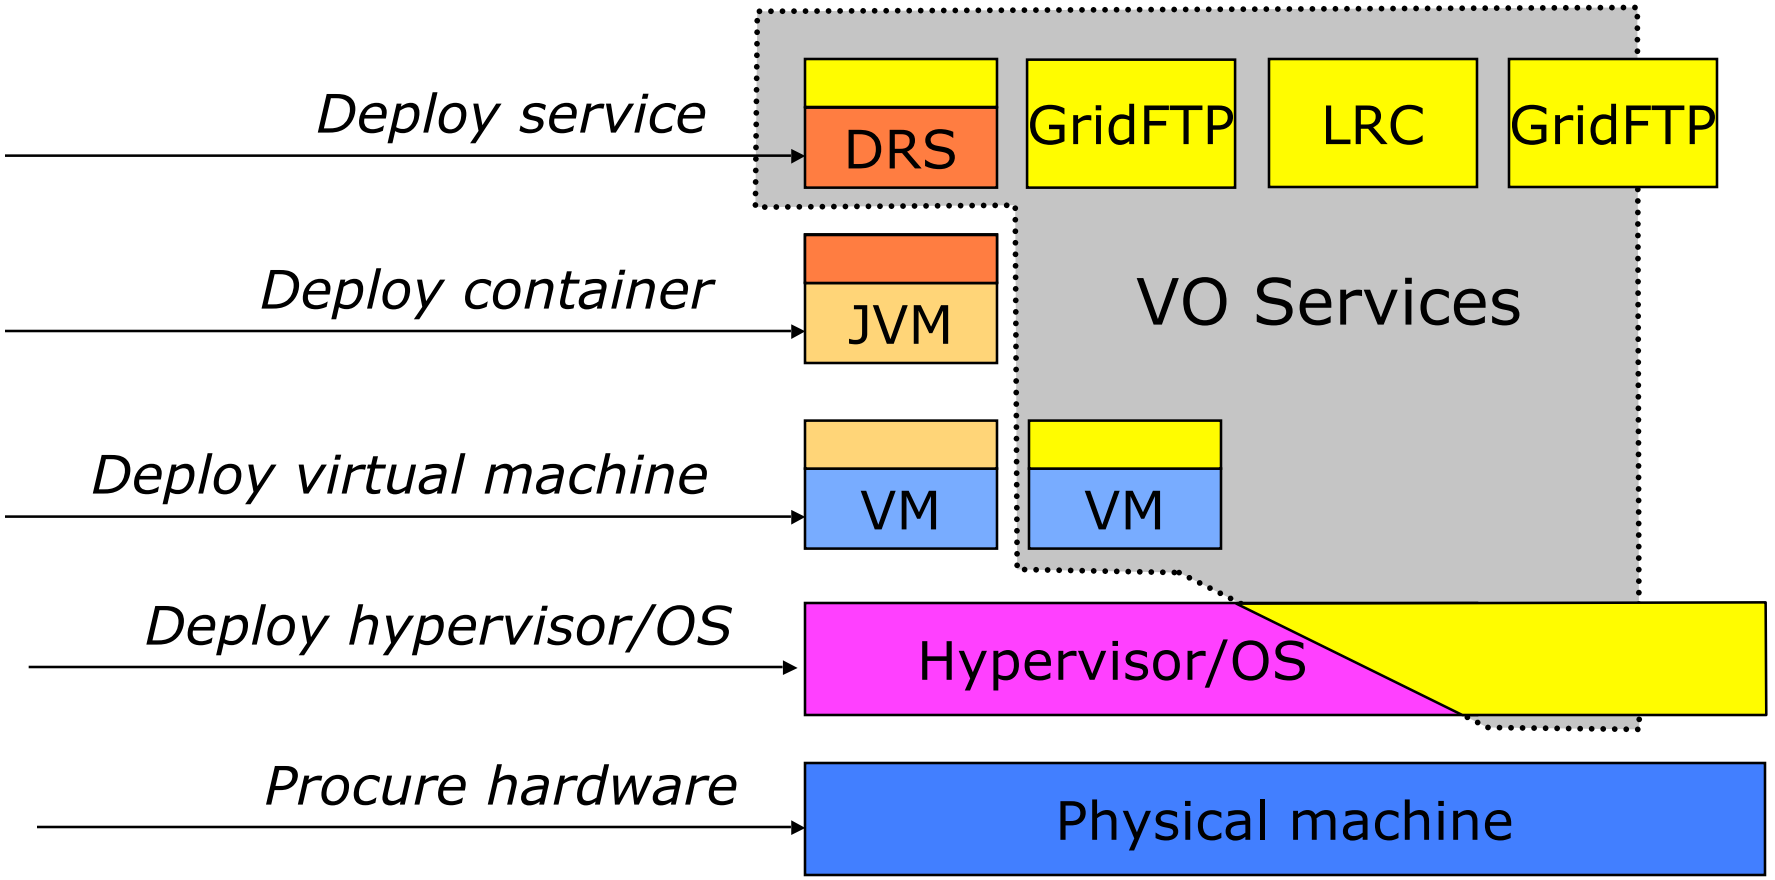

State exposed & access uniformly at all levels
Provisioning, management, and monitoring at all levels

Separation of **Concerns** & **Roles**

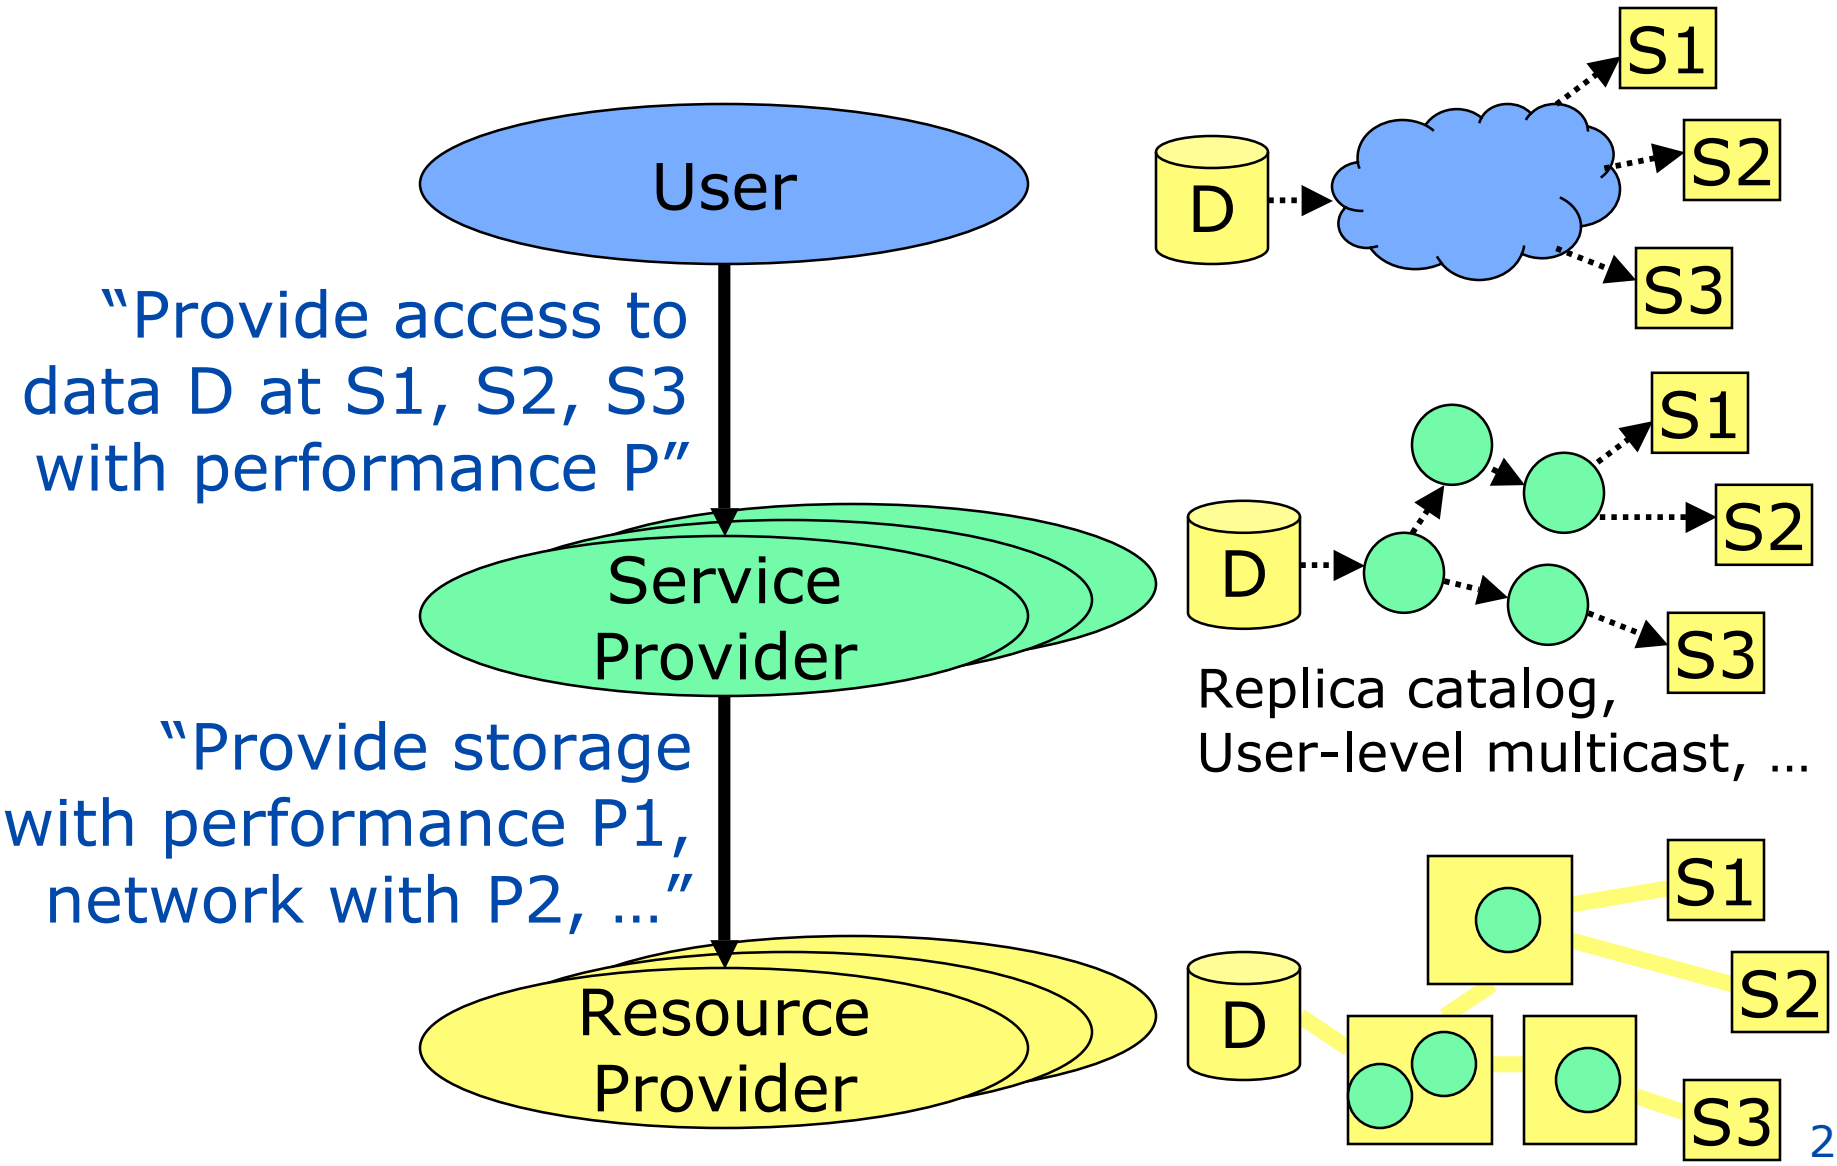

Evolution of Grid Security & Policy

1) Grid security infrastructure
   - Public key authentication & delegation
   - Access control lists ("gridmap" files)

     → *Limited set of policies can be expressed*

2) Utilities to simplify operational use, e.g.
   - MyProxy: online credential repository
   - VOMS, ACL/gridmap management

     → *Broader set of policies, but still ad-hoc*

3) General, standards-based framework for authorization & attribute management

Core Security Mechanisms

- Attribute Assertions
  - **C** asserts that **S** has attribute **A** with value **V**
- Authentication and digital signature
  - Allows signer to assert attributes
- Delegation
  - **C** asserts that **S** can perform **O** on behalf of **C**
- Attribute mapping
  - {A1, A2… An}vo1 ⇨ {A'1, A'2… A'm}vo2
- Policy
  - Entity with attributes **A** asserted by **C** may perform operation **O** on resource **R**

Security Services for VO Policy

- Attribute Authority (**ATA**)
  - Issue signed attribute assertions (incl. identity, delegation & mapping)
- Authorization Authority (**AZA**)
  - Decisions based on assertions & policy

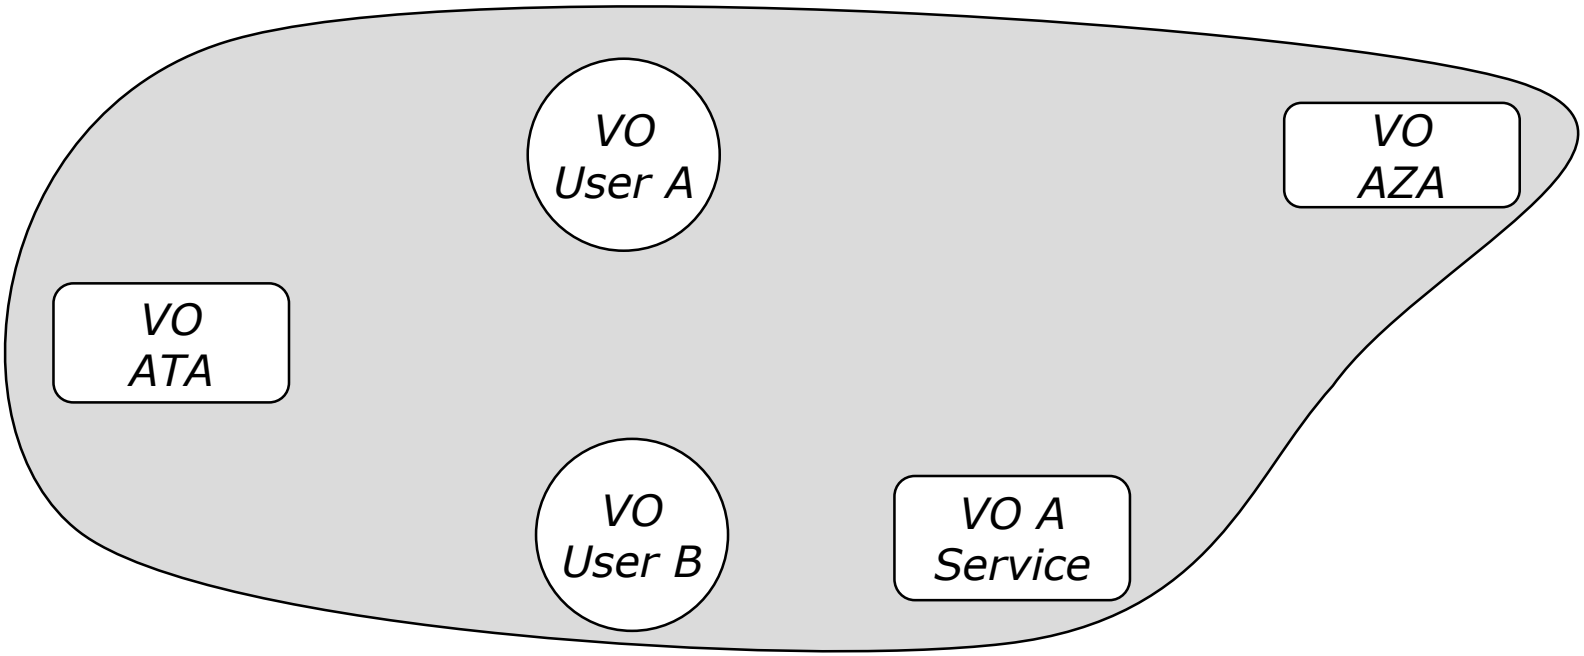

Security Services for VO Policy

- Attribute Authority (**ATA**)
  - Issue signed attribute assertions (incl. identity, delegation & mapping)
- Authorization Authority (**AZA**)
  - Decisions based on assertions & policy

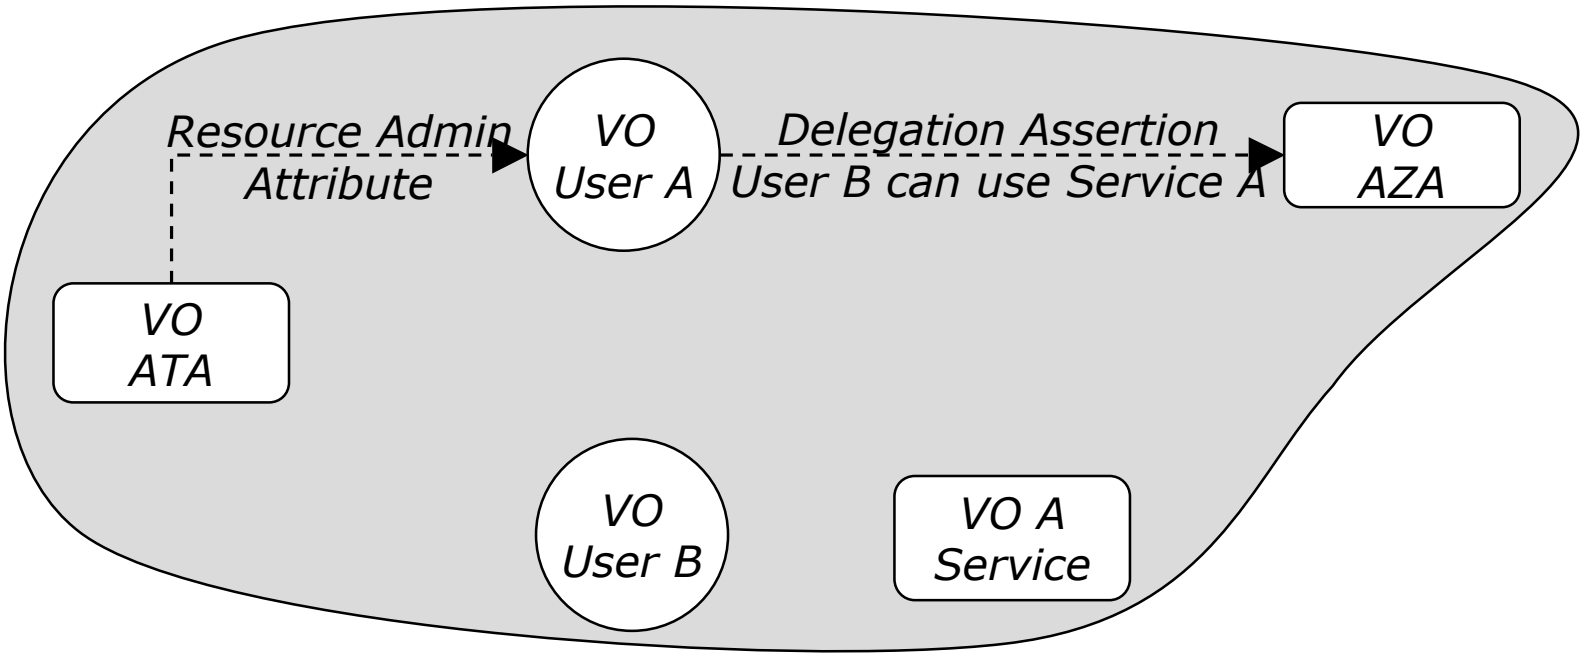

Security Services for VO Policy

- Attribute Authority (**ATA**)
  - Issue signed attribute assertions (incl. identity, delegation & mapping)
- Authorization Authority (**AZA**)
  - Decisions based on assertions & policy

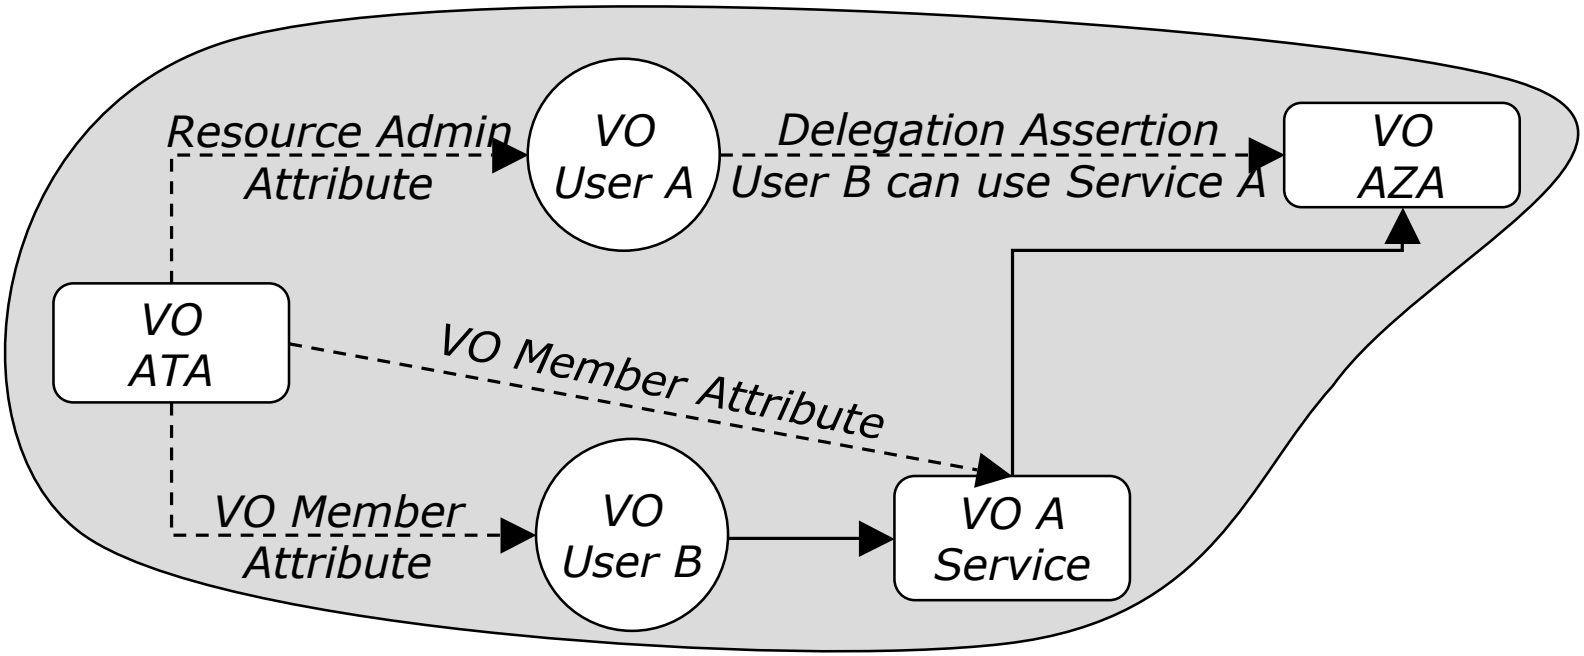

Security Services for VO Policy

- Attribute Authority (**ATA**)
  - Issue signed attribute assertions (incl. identity, delegation & mapping)
- Authorization Authority (**AZA**)
  - Decisions based on assertions & policy

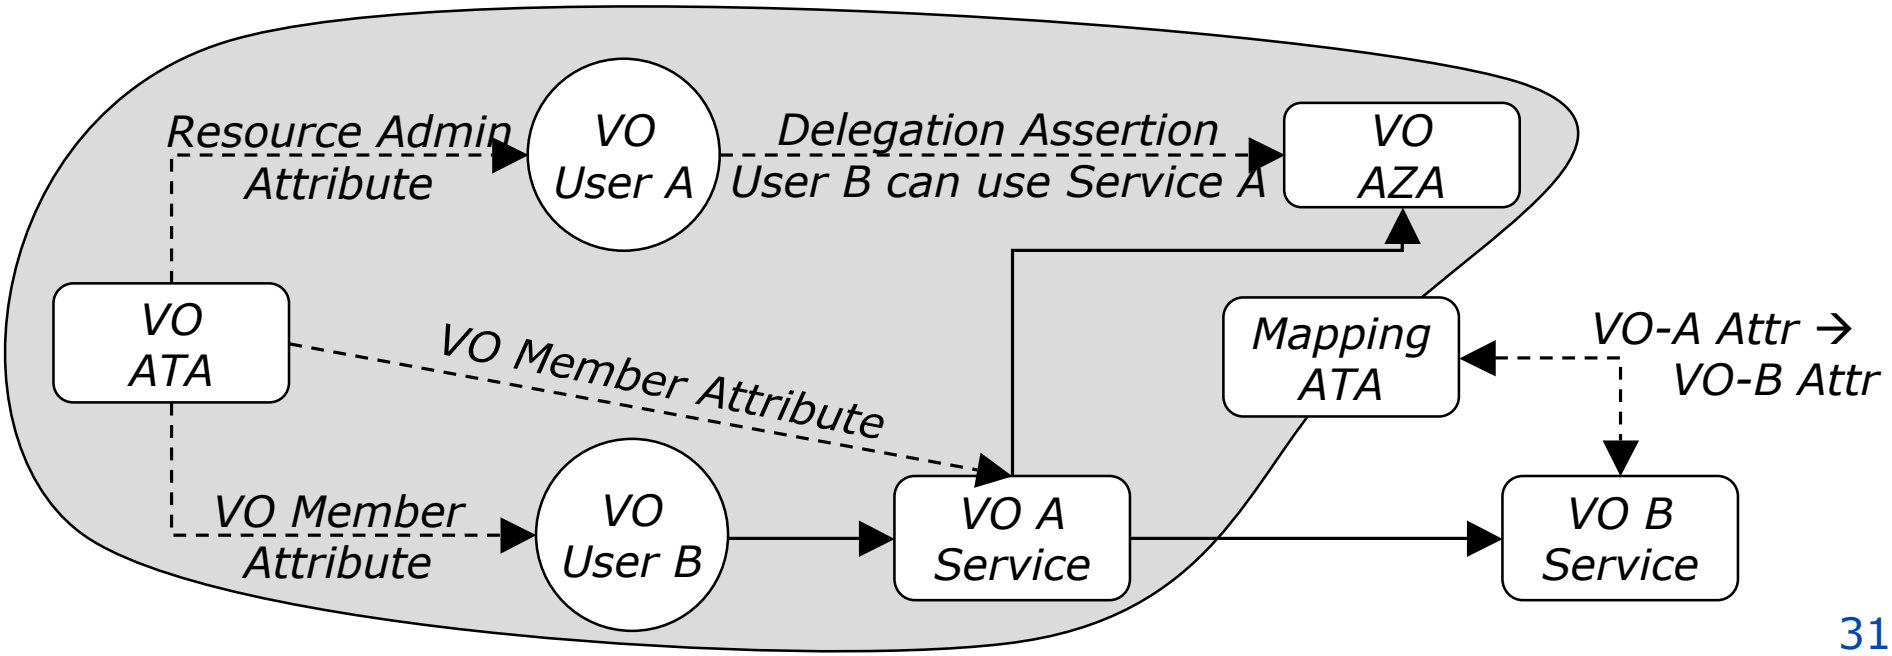

Closing the Loop: GT4 Security Toolkit

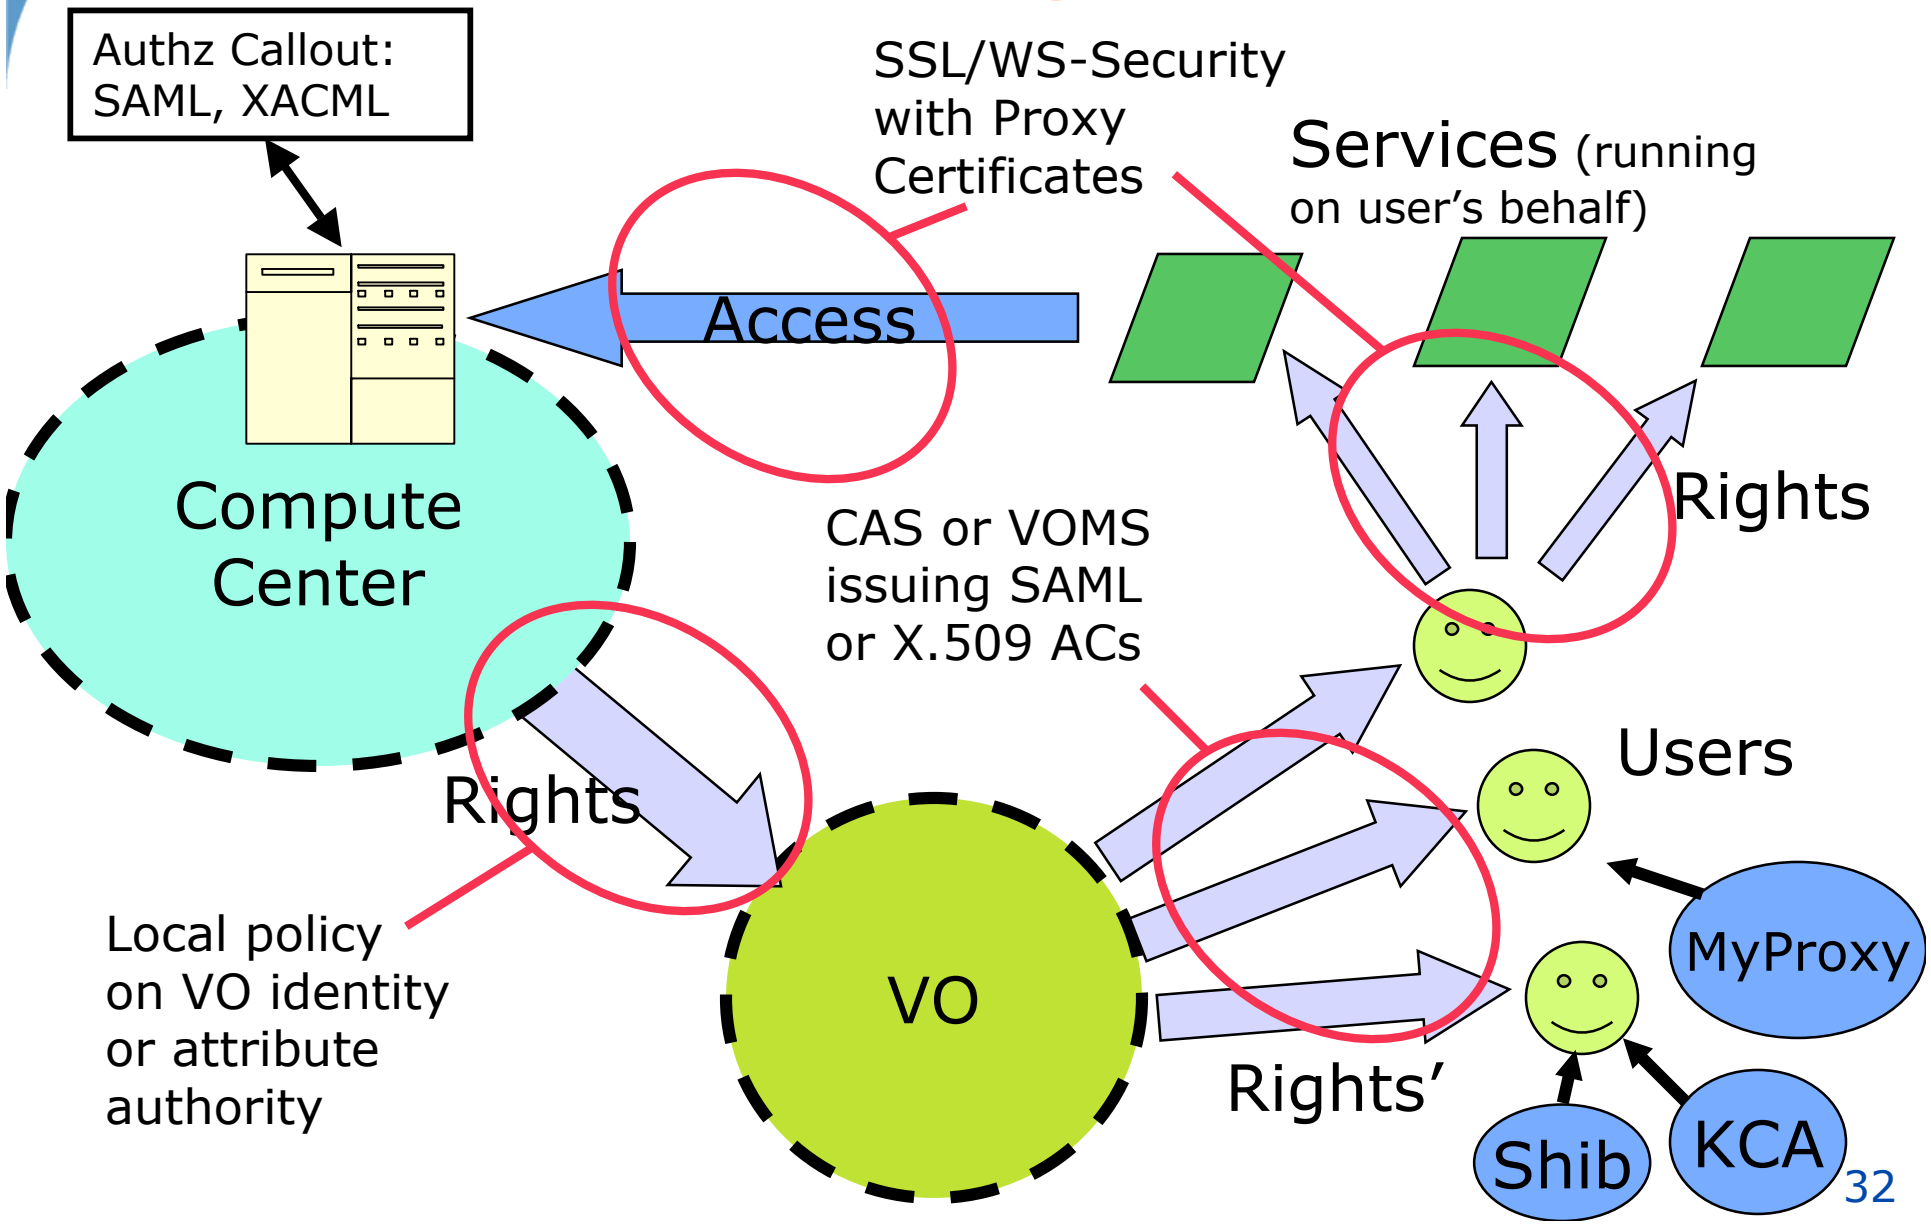

Security Needn't Be Hard: Earth System Grid

- Purpose
  - Access to large data
- Policies
  - Per-collection control
  - Different user classes
- Implementation (GT)
  - Portal-based User Registration Service
  - PKI, SAML assertions
- Experience
  - >2000 users
  - >100 TB downloaded

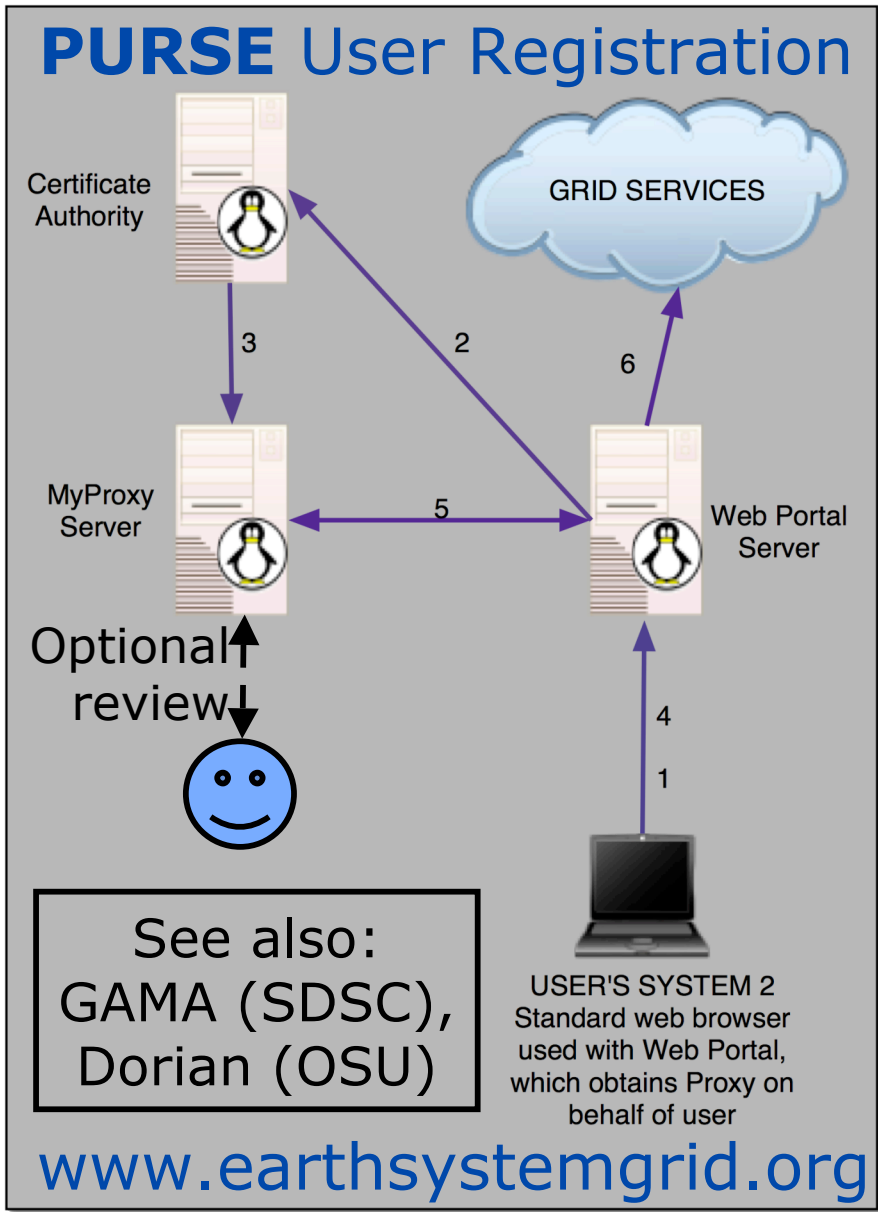

the globus alliance
www.globus.org

Summary: Building Science 2.0 Applications

1) Integrate services from external sources

- Virtualize "services" from providers

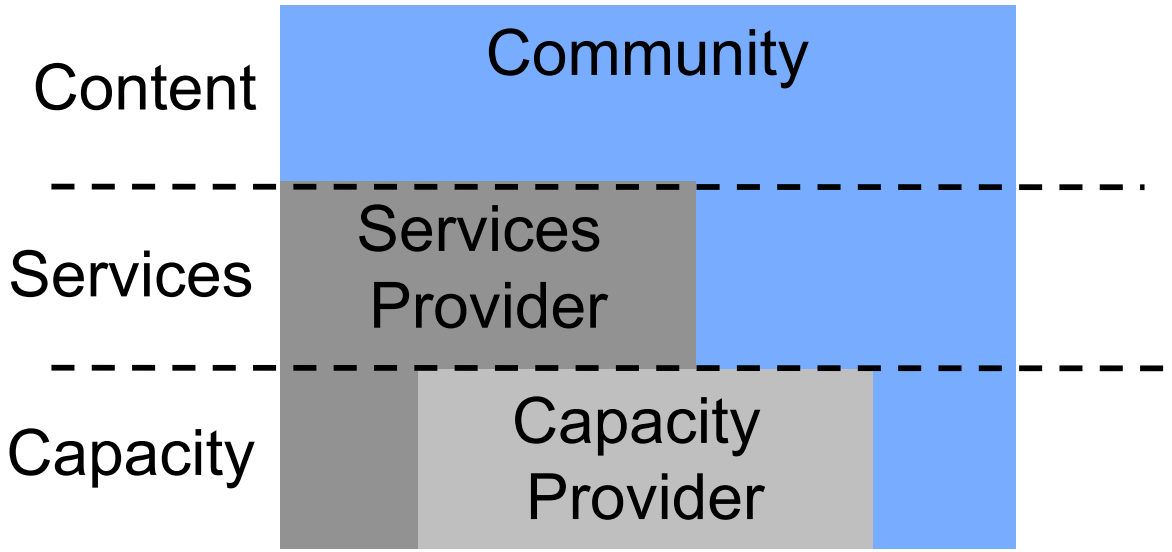

2) Coordinate & compose

- Create new services from existing ones

"Service-Oriented Science", *Science*, 2005

Summary: Science 2.0 Challenges

- A need for new **technologies, skills, & roles**
  - Creating, publishing, hosting, discovering, composing, archiving, explaining … services
- A need for substantial **software development**
  - "30-80% of modern astronomy projects is software"—S. G. Djorgovski, Caltech
- A need for more & different **infrastructure**
  - Computers & networks to host services

the globus alliance
www.globus.org

For More Information

- Globus Alliance
  - www.globus.org
- Dev.Globus
  - dev.globus.org
- Open Science Grid
  - www.opensciencegrid.org
- TeraGrid
  - www.teragrid.org
- Background
  - www.mcs.anl.gov/~foster