



Computation Institute

# Grid, cloud, and beyond

What we have learned about computing on demand

Ian Foster

Computation Institute

Mathematics and Computer Science Division

Department of Computer Science

Argonne National Laboratory & The University of Chicago

**Thirty Years of Parallel Computing at Argonne: A Symposium**

**May 14-15, 2013**



[www.ci.anl.gov](http://www.ci.anl.gov)  
[www.ci.uchicago.edu](http://www.ci.uchicago.edu)

# The original grid vision



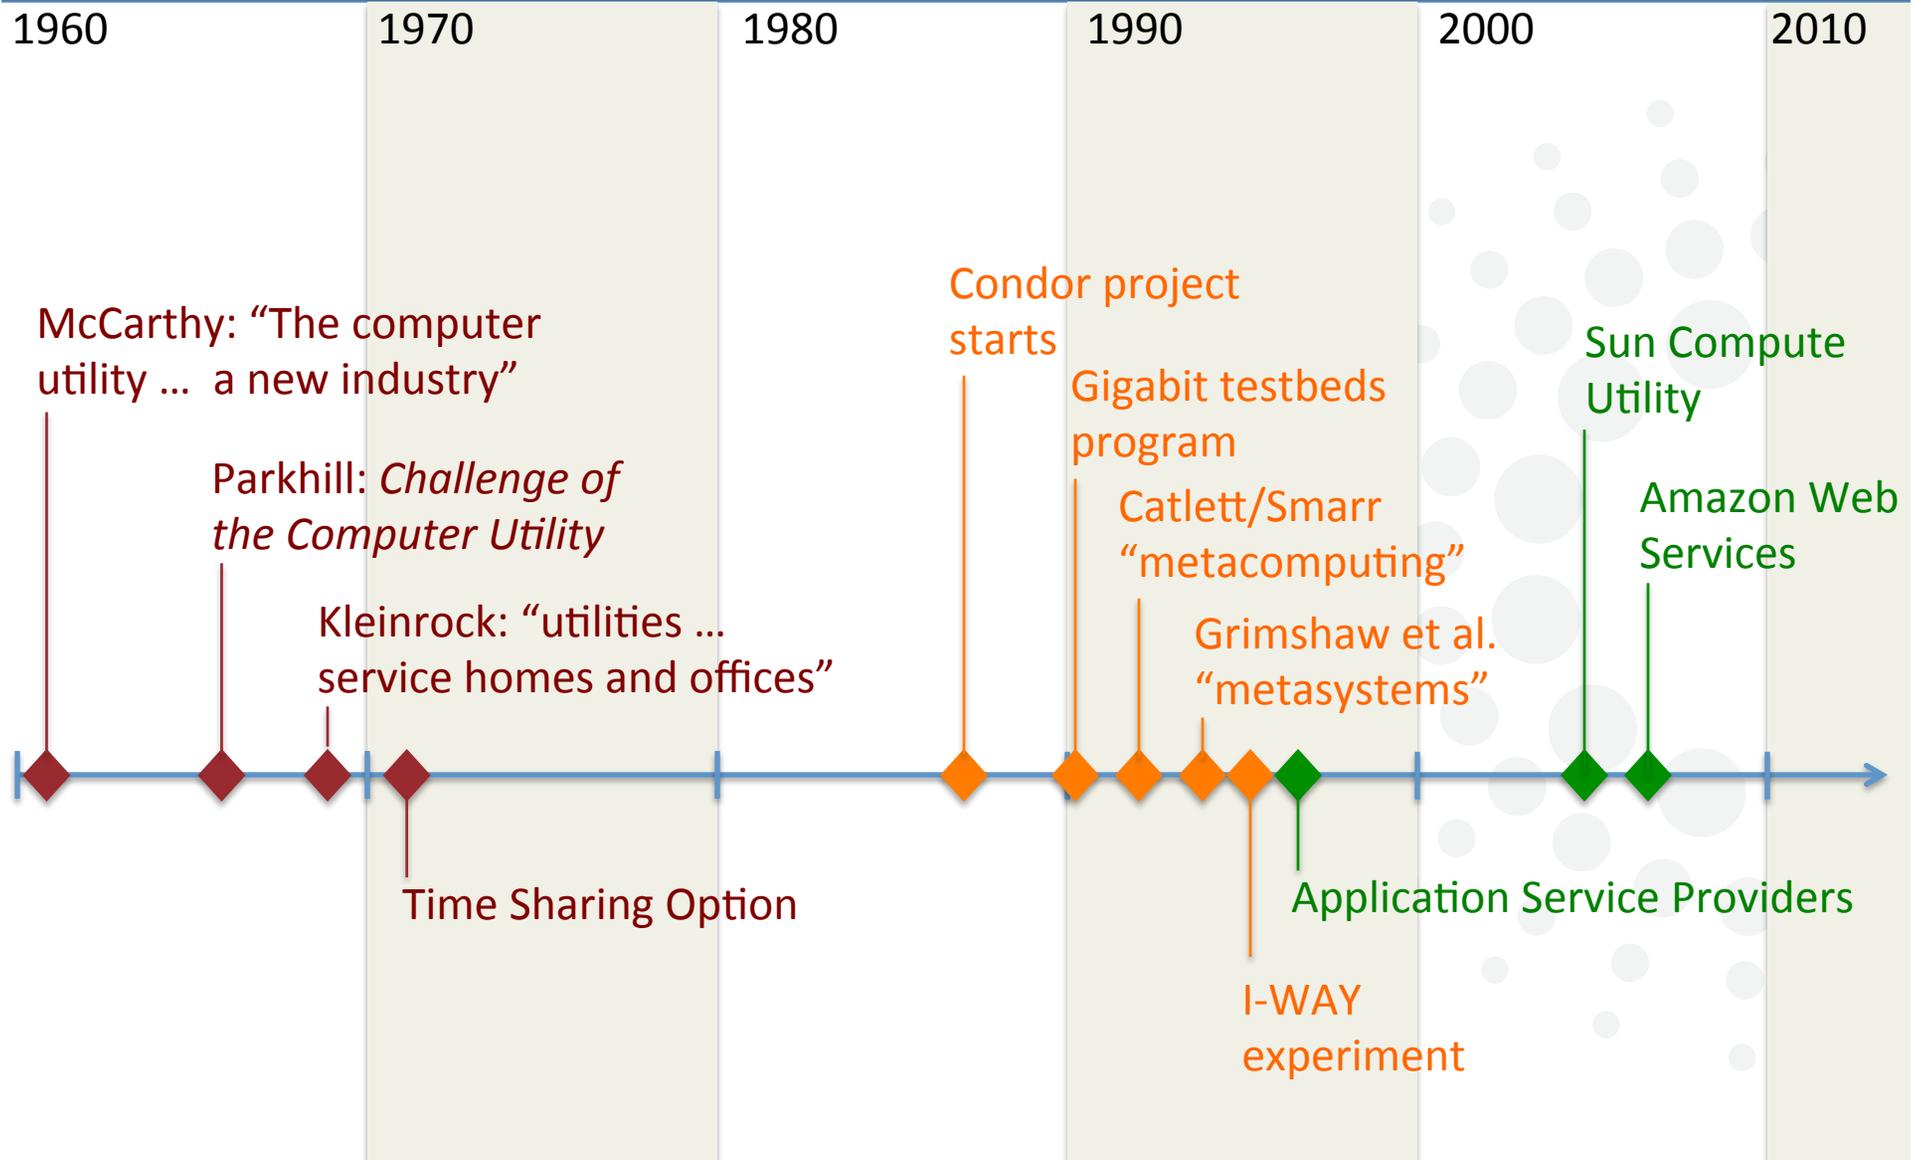
Accelerate discovery and innovation by providing on-demand access to computing

- “the average computing environment remains inadequate for [many] computationally sophisticated purposes”
- “if mechanisms are in place to allow reliable, transparent, and instantaneous access to high-end resources, then ... it is as if those resources are devoted to them”

[*The Grid*, Chapter 2, 1998]



# The on-demand idea certainly isn't new



“A simple back of the envelope calculation shows that [McCarthy’s] idea can never work.”

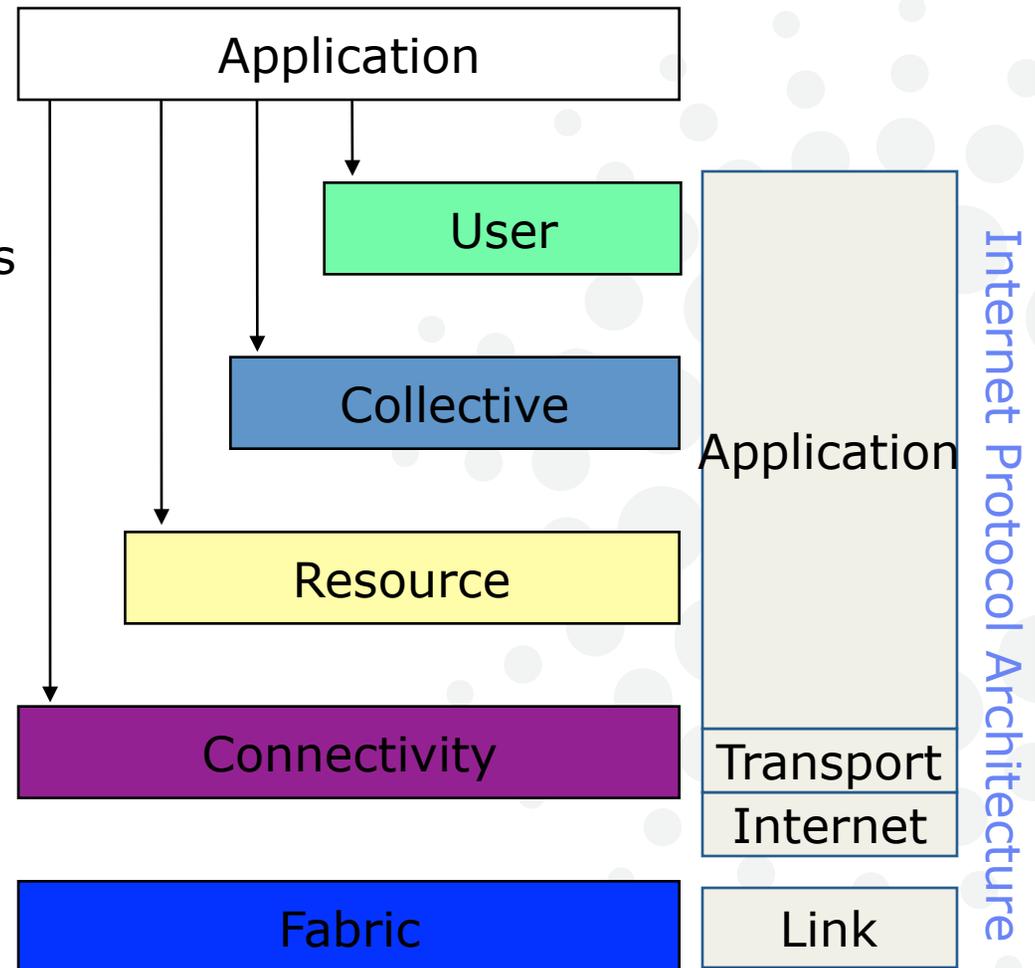




# Focus on abstractions and mechanisms



- “Specialized services”: user- or appln-specific distributed services
- “Managing multiple resources”: ubiquitous infrastructure services
- “Sharing single resources”: negotiating access, controlling use
- “Talking to things”: communication (Internet protocols) & security
- “Controlling things locally”: Access to, & control of, resources



# “The Anatomy of the Grid,” 2001



The ... problem that underlies the Grid concept is coordinated resource sharing and problem solving in **dynamic, multi-institutional virtual organizations**. The sharing that we are concerned with is not primarily file exchange but rather direct access to computers, software, data, and other resources, as is required by a range of collaborative problem-solving and resource-brokering strategies emerging in industry, science, and engineering. This sharing is, necessarily, highly controlled, with resource providers and consumers defining clearly and carefully just what is shared, who is allowed to share, and the conditions under which sharing occurs. **A set of individuals and/or institutions defined by such sharing rules form what we call a virtual organization (VO).**

# Examples (from AotG, 2001)

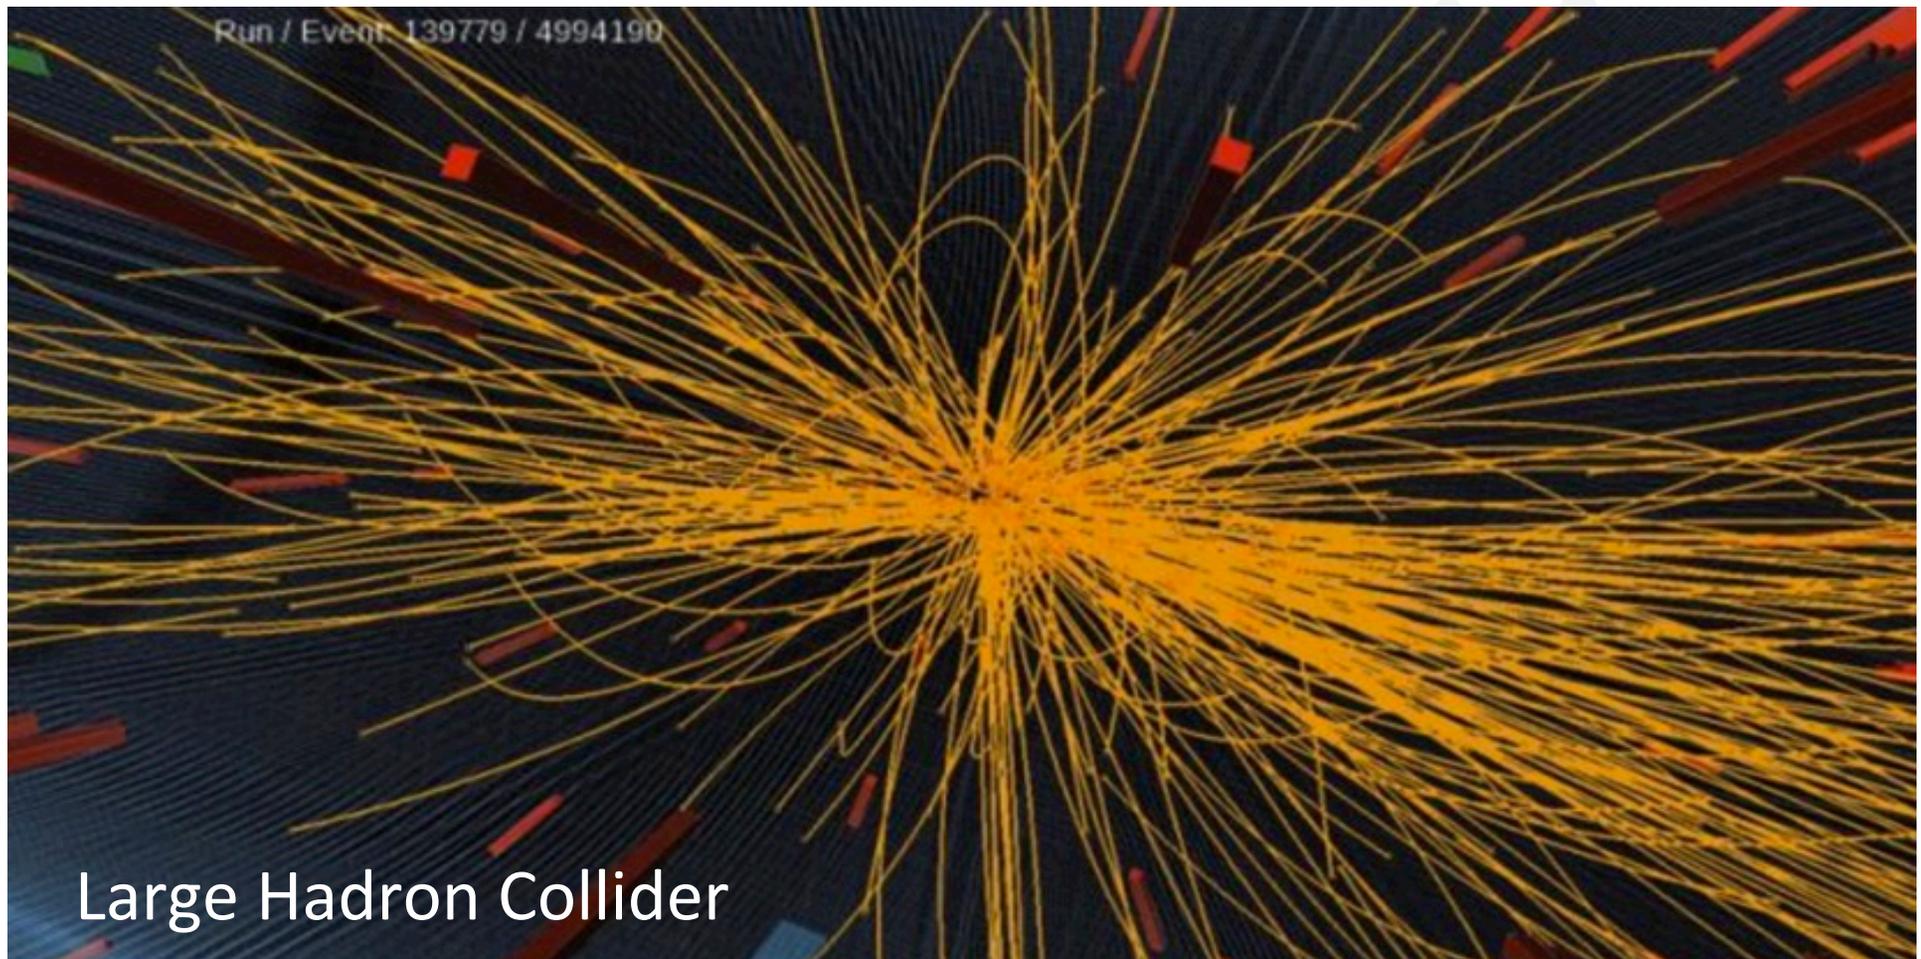


- “The application service providers, storage service providers, cycle providers, and consultants engaged by a car manufacturer to perform scenario evaluation during planning for a new factory”
- “Members of an industrial consortium bidding on a new aircraft”
- “A crisis management team and the databases and simulation systems that they use to plan a response to an emergency situation”
- “Members of a large, international, multiyear high-energy physics collaboration”

# Grid technology accelerates discovery



Higgs discovery “only possible because of the **extraordinary achievements** of ... **grid computing**” —Rolf Heuer, CERN DG



# 3,761

GridFTP servers  
reporting usage

# 169 million GRAM jobs submitted

*by the Open Science Grid in 2012*

# 300,000 jobs/day reported

# 382 million operations

# 29 petabytes transferred

*during February 2013*



Open Science Grid



# 100 MyProxy servers

# 2,000,000 requests per week

# 1,000 GSI-OpenSSH servers

# 1,000,000 login requests per week

# Things we got right



- Close partnerships with application groups with substantial problems
- Focus on resource models and low-level mechanisms vs. all-encompassing frameworks
- Definition of data movement & security protocol conventions to encourage interoperability
  - E.g., GridFTP, Grid Security Infrastructure, Storage Resource Manager (Open Grid Forum, IETF, etc.)
- Virtual organizations as an organizational principle for collaborative work

# Things we got wrong



- Unrealistic expectations that supercomputer centers could become “cloud providers”
- Excessive reliance on Web Services
- Didn't account for European vs. U.S. competition (and Microsoft vs. IBM etc.)
- Overly focused on big science

# What has changed? { Causation or just correlation? Discuss ...



- Thousands of people learned about the joys of large-scale distributed systems
- Virtual organization concepts and technologies
- Now routine to move 100s of terabytes (e.g., GridFTP moves >1 petabyte per day)
- High throughput computing is mainstream (e.g., Condor and Globus run millions of jobs per day)
- Large Hadron Collider found the Higgs
- Earth System Grid supports >25,000 users
- Commercial cloud computing has exploded

# Looking forward



- Exploding data volumes and earlier successes mean that many more people face challenges of big data, big compute, big collaboration
- Networks are several orders of magnitude faster than when Grid started
- Commercial cloud providers provide a substrate on which powerful new capabilities can be built with new economies of scale

# Complexity is large and growing



Time



Run experiment

Collect data

Move data

Check data

Annotate data

Share data

Find similar data

Link to literature

Analyze data

Publish data



# Process automation for science



Time

Run experiment

Collect data

Move data

Check data

Annotate data

Share data

Find similar data

Link to literature

Analyze data

Publish data



Research IT  
as a service

# MG-RAST

metagenomics analysis server



LOGIN

[REGISTER](#)

[PASSWORD](#)

[FORGOT?](#)

login



[Browse Metagenomes](#)

search for metagenomes



[Register](#)



[Contact](#)



[Help](#)



[Upload](#)



[News](#)

## About

MG-RAST (the Metagenomics RAST) server is an automated analysis platform for metagenomes providing quantitative insights into microbial populations based on sequence data.

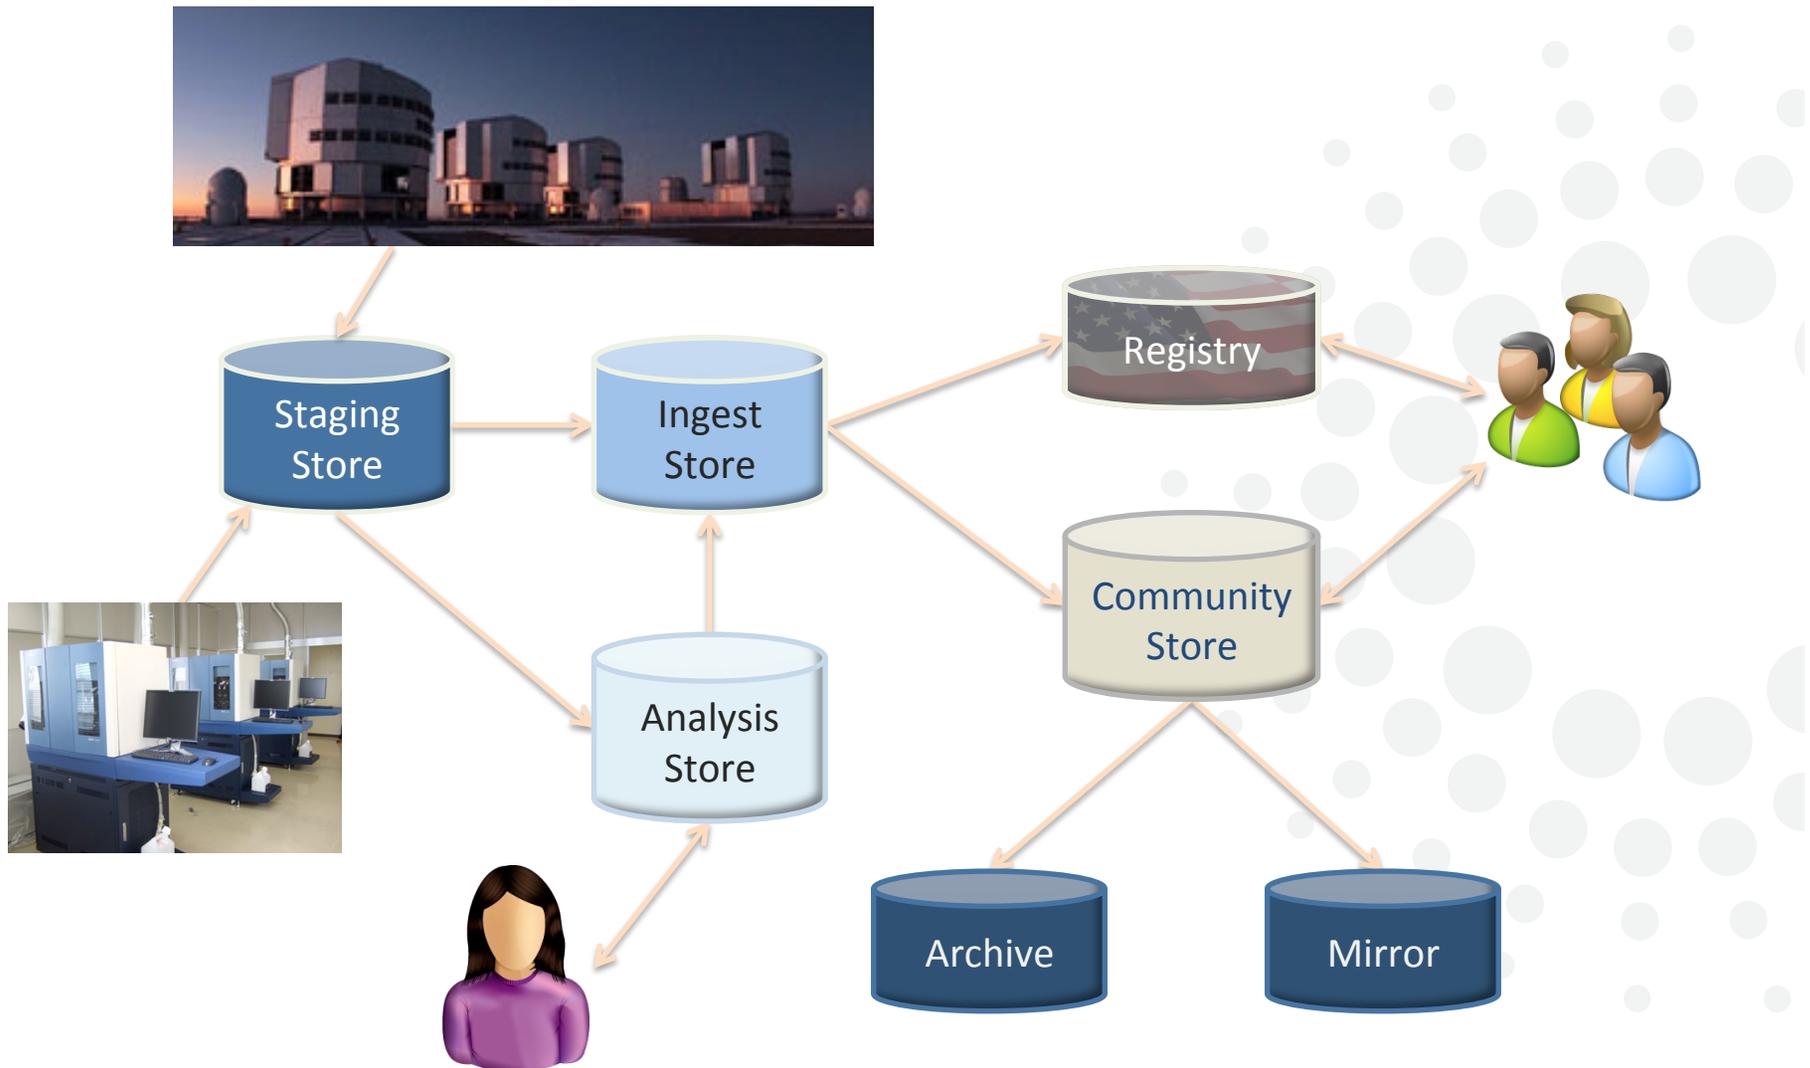
# of metagenomes	78,948
# base pairs	26.25 Tbp
# of sequences	240.38 billion
# of public metagenomes	12,542

The server primarily provides upload, quality control, automated annotation and analysis for prokaryotic metagenomic shotgun samples. MG-RAST was launched in 2007 and has over 8000 registered users and 78,948 data sets. The current server version is 3.3.3.3. We suggest users take a look at [MG-RAST for the impatient](#).

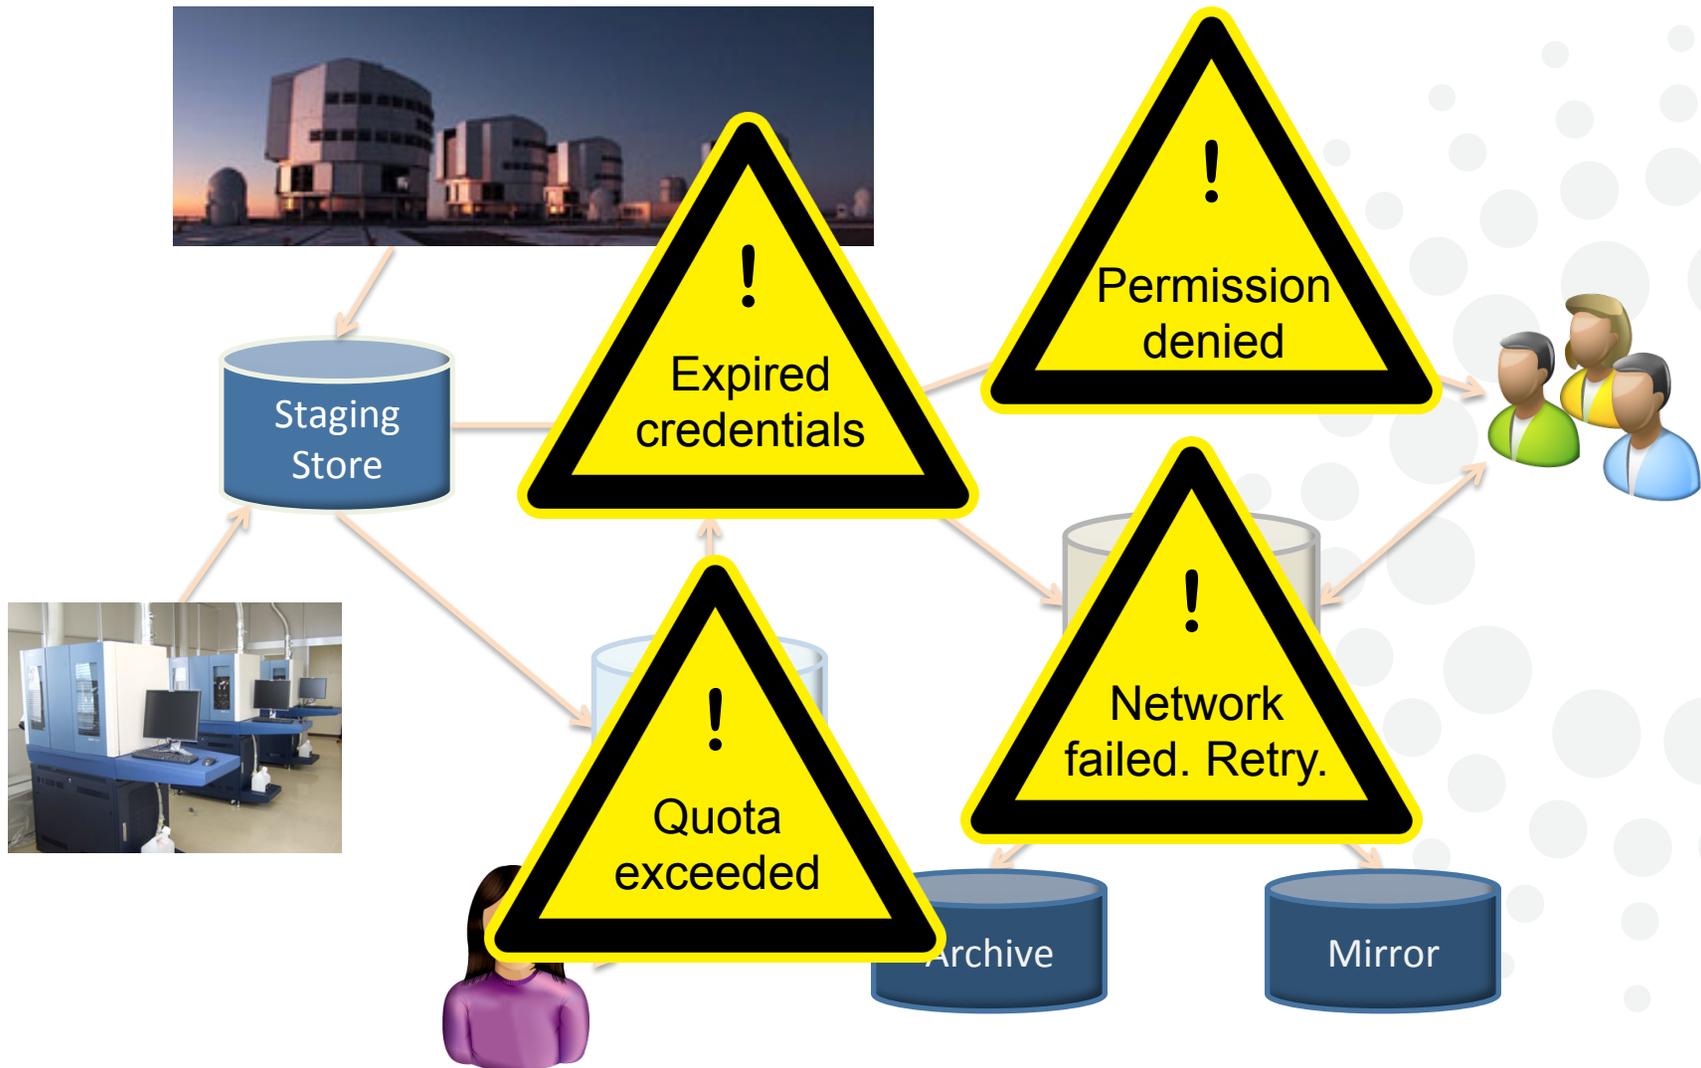
[Updates](#)

[Announcing DRISSE our new tool to describe sequencing error \[June 19, 2012\]](#)

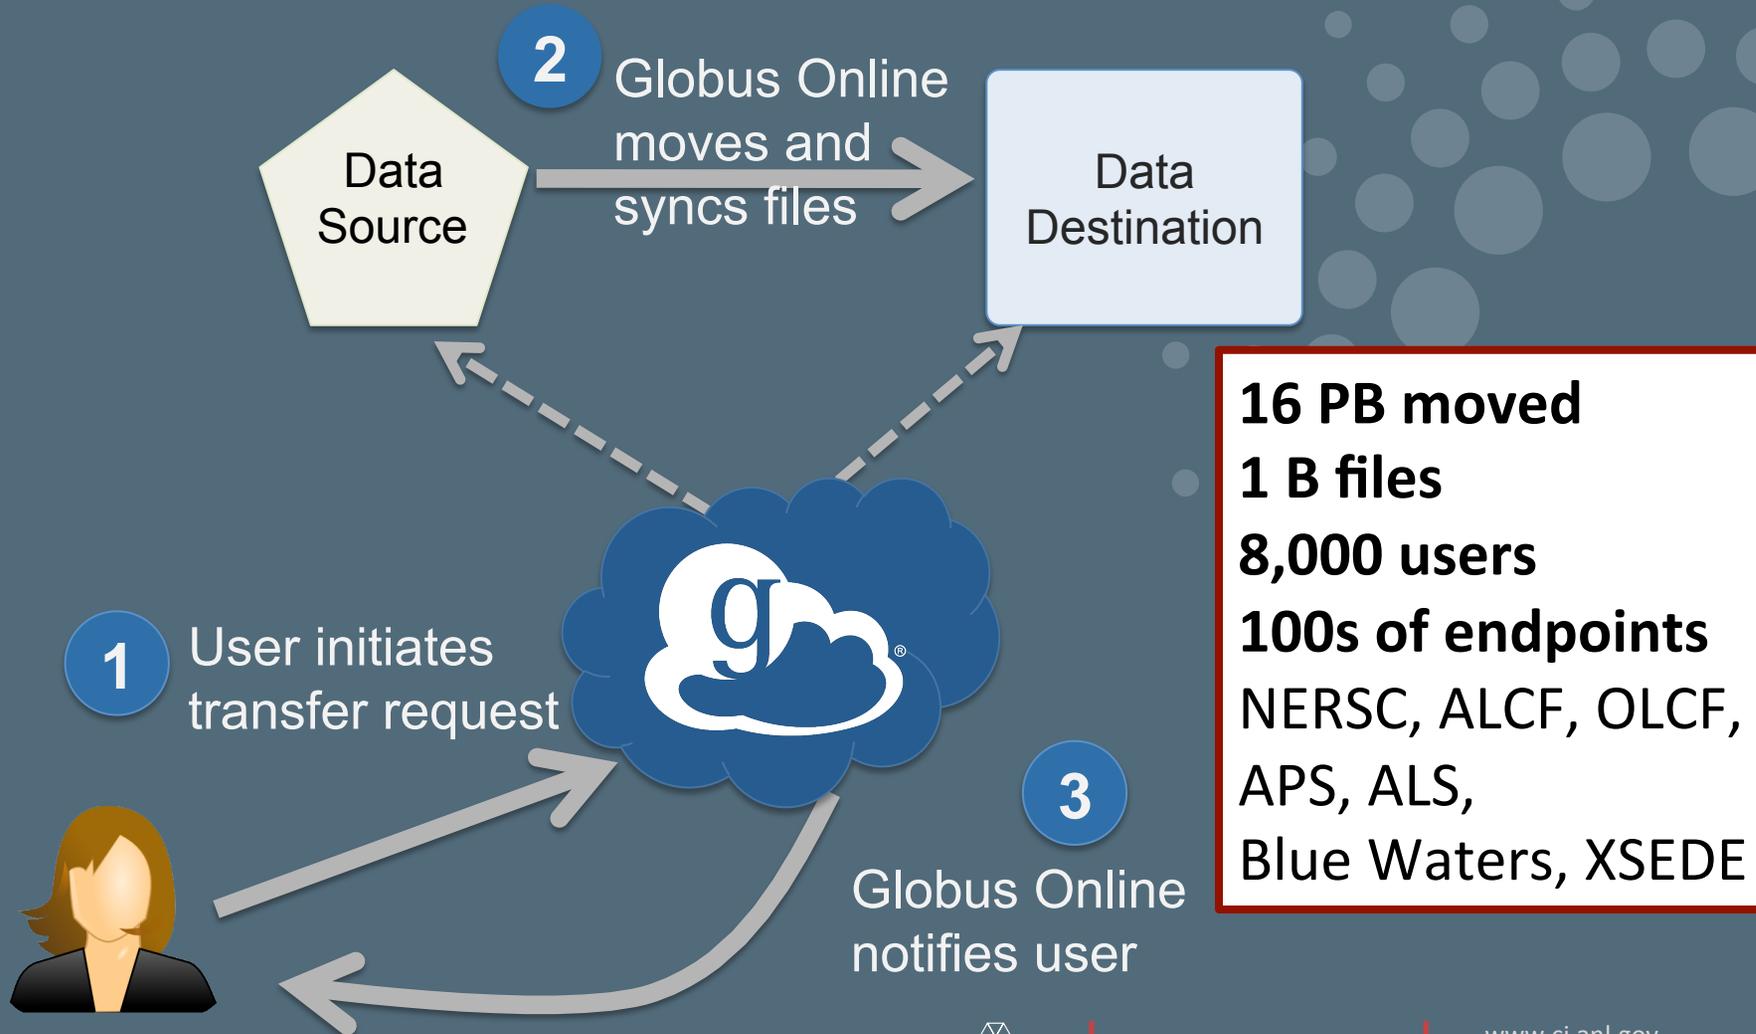
# Managing data should be easy ...



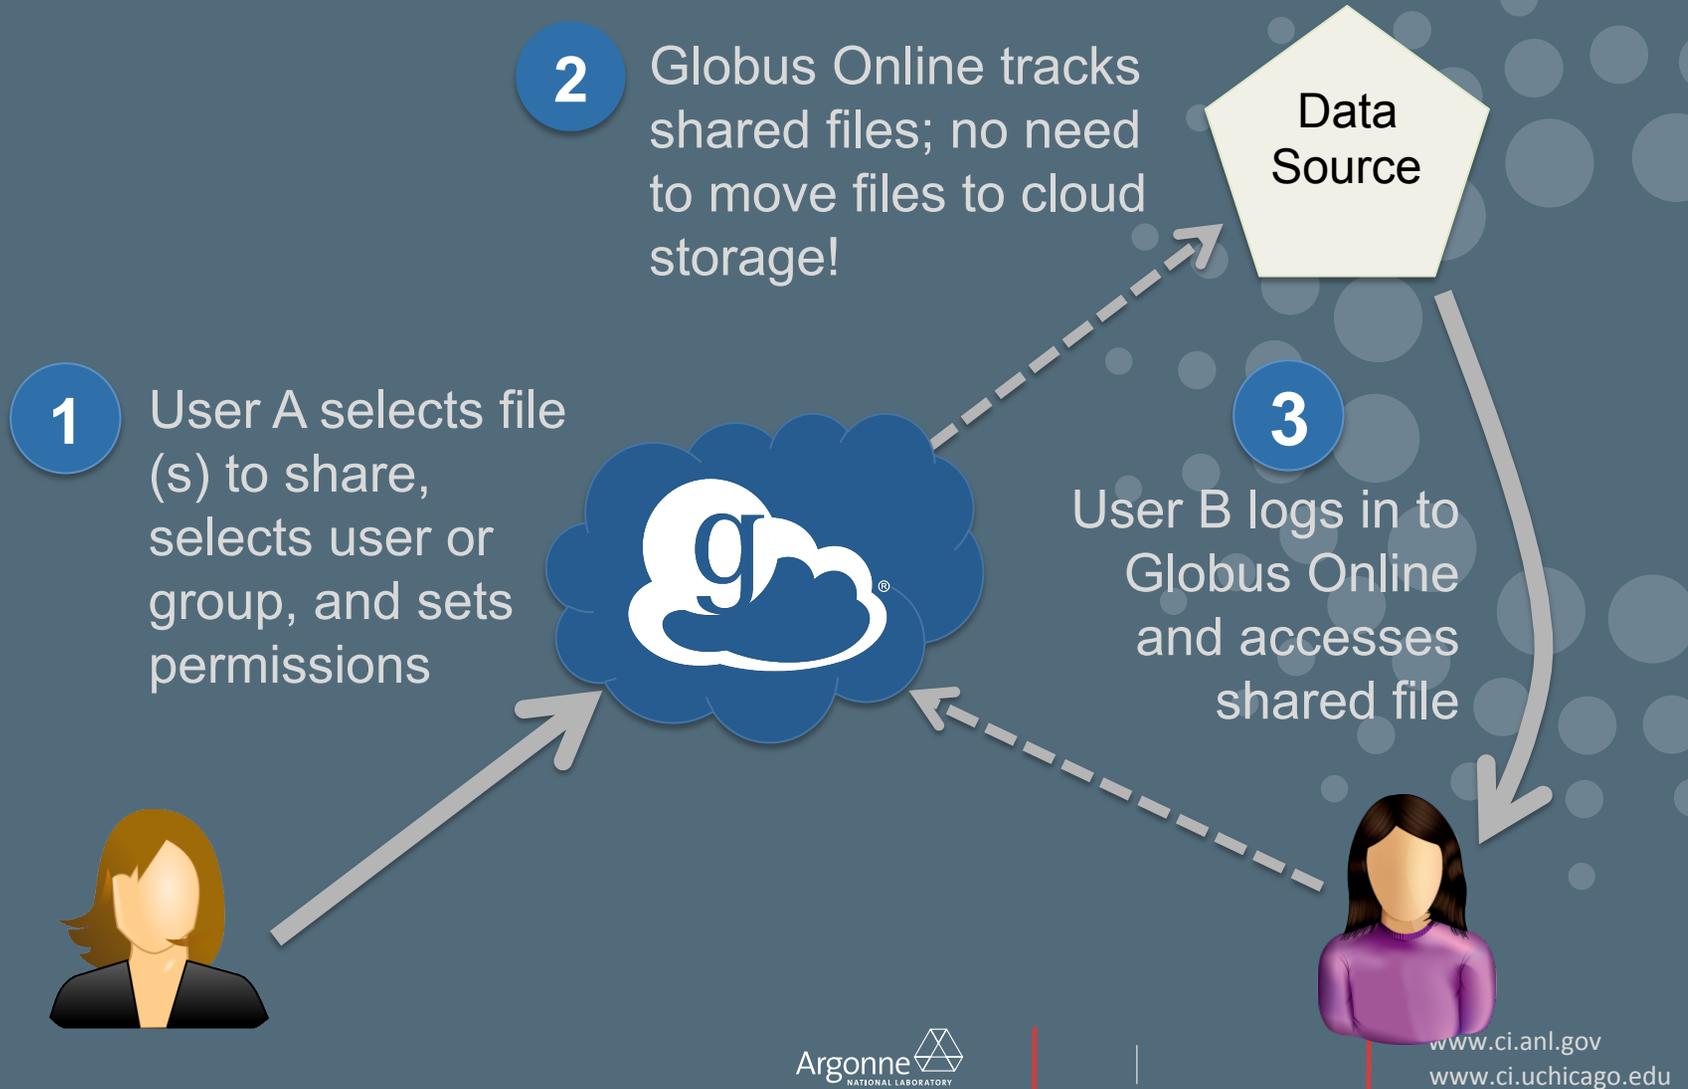
# ... but it's often hard and frustrating!



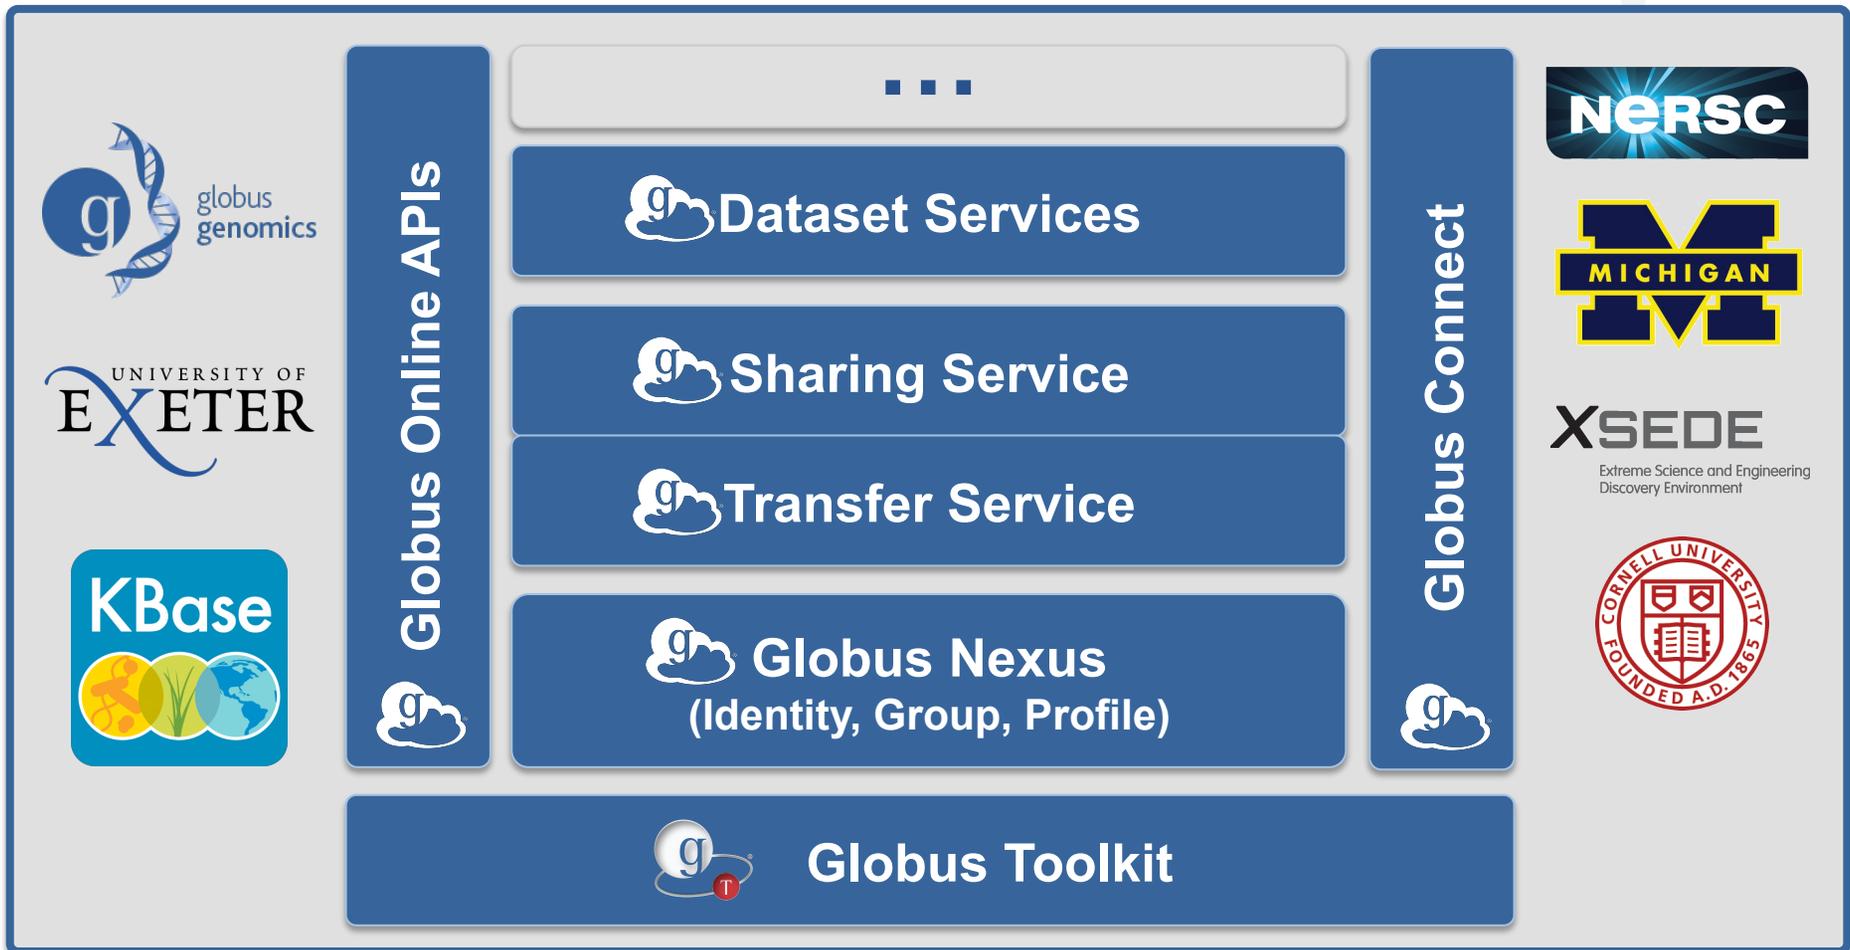
# We started with reliable, secure, high-performance file transfer ...



# ... and then made it simple to share big data off existing storage systems



# And are next expanding to a platform





Accelerate discovery and innovation worldwide by providing **research IT as a service**

Leverage software-as-a-service to

- provide millions of researchers with unprecedented access to powerful tools;
- enable a massive shortening of cycle times in time-consuming research processes; and
- reduce research IT costs dramatically via economies of scale



Computation Institute

# Thanks to ...



U.S. DEPARTMENT OF  
**ENERGY**



THE UNIVERSITY OF  
**CHICAGO**



[www.ci.anl.gov](http://www.ci.anl.gov)  
[www.ci.uchicago.edu](http://www.ci.uchicago.edu)



Computation Institute

# Thank you! Questions?

[foster@anl.gov](mailto:foster@anl.gov)

[foster@uchicago.edu](mailto:foster@uchicago.edu)