Globus Nexus: An identity, profile, and group management platform for science gateways and other collaborative science applications

Rachana Ananthakrishnan, Josh Bryan, Kyle Chard, Ian Foster, Tom Howe, Mattias Lidman, Steven Tuecke
Computation Institute
Argonne National Laboratory & University of Chicago
Chicago, IL 60637, USA

Abstract—**Globus Nexus is a flexible and powerful Platform-as-a-Service to which developers can outsource identity, group, and profile management needs. By providing these frequently important but always challenging capabilities as a service, accessible over the network, Globus Nexus streamlines web application development and makes it easy for individuals, teams, and institutions to create collaborative web applications such as science gateways for the science community. We introduce the capabilities of this platform and review representative applications.**

Keywords—identity, group, authentication, authorization, profile, science gateways, platform

I. Introduction

Developers of science gateways [1] frequently need to assign identities to their users, manage user profiles, and organize users into groups for authorization and other purposes. They may also want to allow users to authenticate to the gateway by presenting campus credentials. But providing high-quality implementations of such capabilities can be extremely challenging, due to the complexity of the associated security protocols.

Globus Nexus allows developers of science gateways to outsource identity, profile, and group management functions to a third party platform, Globus Nexus, which the University of Chicago operates for the research community. This platform addresses four major obstacles to the creation and operation of high-quality collaborative applications:

1. Identity provisioning: Create and manage identities for gateway users.
2. Identity hub: Link different user identities, so that for example a user can authenticate to a gateway with a campus (InCommon) credential.
3. Group hub: User-managed group creation and management functions. Groups can then be used in authorization decisions.
4. Profile management: User-managed profile attributes and visibility for those attributes. Profile attributes can be used in authorization decisions, for example to determine who is allowed to join a group.

Globus Nexus provides developers with powerful and flexible management interfaces and REST APIs. Web interfaces can easily be "skinned" to match the interface of an institution or project. Users of Globus Nexus capabilities encounter intuitive web interfaces with a common look and feel across different services. Both developers and users benefit from high-quality implementations, support for a wide range of security protocols, and a highly reliable platform based on replicated state and services distributed over multiple commercial cloud data centers.

Globus Nexus is part of the Globus Online family of services (Figure 1) [2]. Its capabilities are used by Globus Online file transfer, synchronization, and sharing. Its identity management capabilities permit Globus Online users to connect, via Globus Connect, to storage systems at locations such as NERSC, the University of Michigan, XSEDE, and Cornell University. Its group management capabilities permit users to manage access to their data on such storage systems.

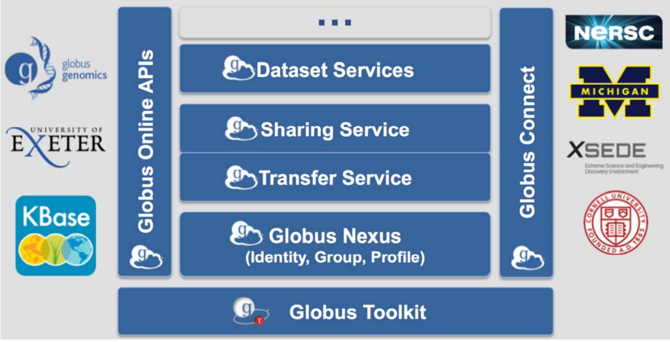

Figure 1: Globus Online services provide users with research data management functions (right) but also implement REST APIs that permit their use as a platform (left).

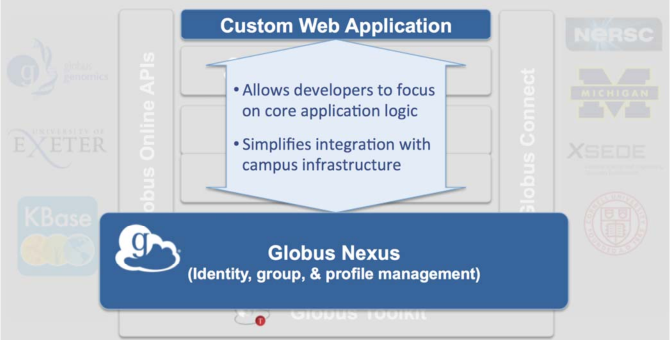
Custom Web Application
• Allows developers to focus on core application logic
• Simplifies integration with campus infrastructure
Globus Nexus
(Identity, group, & profile management)

Figure 2: The Globus Nexus platform facilitates the development of collaborative science applications.

978-1-4799-0898-1/13/$31.00 ©2013

As a platform, Globus Nexus provides (see Figure 2) a set of interfaces that application developers can use to invoke capabilities that would be difficult for them to implement and operate themselves.

II. Globus Nexus Capabilities

A. *Identity Provisioning*

Globus Nexus can act as an identity provider for a project, providing convenient Web interfaces for identity creation and providing email validation. The DOE Systems Biology Knowledge Base (kBase.us) is an example of a project that uses Globus Nexus for identity provisioning: see Figure 3. We currently manage ~700 identities for kBase.

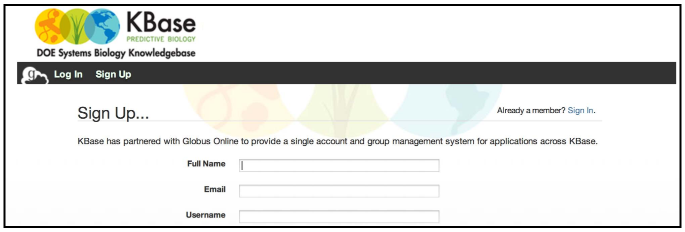

Sign Up...

Already a member? Sign In.

KBase has partnered with Globus Online to provide a single account and group management system for applications across KBase.

Full Name

Email

Username

Figure 3: kBase sign up page, showing part of sign-up dialog.

B. *Identity Hub*

Having created a Globus identity, it is straightforward to link identities from other federated identity providers: see Figure 4. For example, InCommon (via the SAML protocol) [3], Google (via OpenID), XSEDE (via OAuth MyProxy [4]), an IGTF-certified X.509 certificate authority, or SSH.

Having linked an identity, the user can then use that identity to authenticate to Nexus as the Globus identity: see Figure 5. Thus, for example, it is straightforward for a user to authenticate to Globus Nexus with their InCommon campus identity: a frequently requested feature for science gateways.

Globus Nexus can act as a federated identity provider to other services, via the OAuth protocol. Various groups have leveraged this capability to enable authentication to XSEDE, Jira, Zendesk, and Globus data management services.

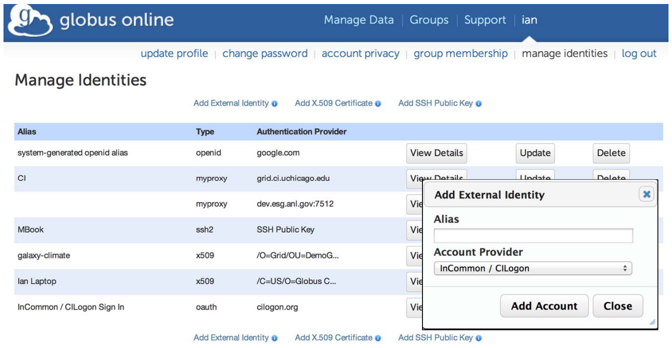

Manage Identities

Alias	Type	Authentication Provider
system-generated openid alias	openid	google.com
CI	myproxy	grid.ci.uchicago.edu
	myproxy	dev.esg.anl.gov:7512
MBook	ssh2	SSH Public Key
galaxy-climate	x509	/O=Grid/OU=DemoG...
Ian Laptop	x509	/C=US/O=Globus C...
InCommon / CILogon Sign In	oauth	cilogon.org

Add External Identity

Alias

Account Provider: InCommon / CILogon

Figure 4: Globus Nexus identity hub, showing user engaged in linking an InCommon identity.

Nexus can also cache, on the user's behalf, delegated credentials obtained from a third-party service. Thus, for example, a gateway that must access an XSEDE service repeatedly on a user's behalf need not interact with the user repeatedly after a first authentication.

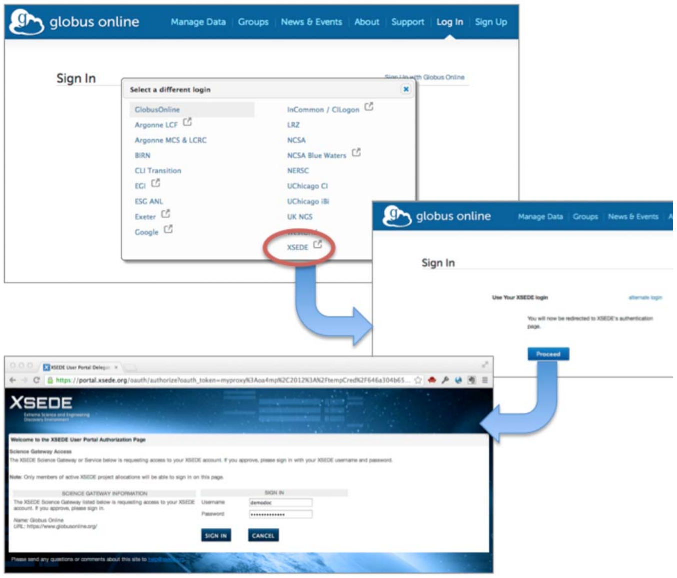

You will now be redirected to XSEDE's authentication page.

https://portal.xsede.org/oauth/authorize?oauth_token=myproxy%3Aoa4mp%2C2012%3A%2FtempCred%2F646a304b65...

Welcome to the XSEDE User Portal Authorization Page

Science Gateway Access

The XSEDE Science Gateway or Service below is requesting access to your XSEDE account. If you approve, please sign in with your XSEDE username and password.

Note: Only members of active XSEDE project allocations will be able to sign in on this page.

SCIENCE GATEWAY INFORMATION

The XSEDE Science Gateway listed below is requesting access to your XSEDE account. If you approve, please sign in.

Name: Globus Online
URL: https://www.globusonline.org/

SIGN IN

Username: demodoc
Password: ••••••••••••••

Please send any questions or comments about this site to [illegible]

Figure 5: Using a linked identity for authentication to Globus Nexus. Here, an XSEDE identity is used. The OAuth protocol is used to delegate authentication to XSEDE.

We use an example from the BIRN project [5] to illustrate the power of the Globus Nexus identity hub: see Figure 6. BIRN uses Globus Nexus for identity provisioning. In this example, Dr. Smith has created a Globus identity (via a BIRN-tailored interface) to which she has linked her campus identity and XSEDE identity. Dr. Smith can then:

- Authenticate to BIRN with her campus identity
- Query a BIRN catalog (using their BIRN identity)
- Request data transfer from BIRN to campus (BIRN and campus identities)
- Request transfer from BIRN to XSEDE (BIRN and XSEDE identities)
- Repeat these tasks without repeated authentication, thanks to the use of cached credentials.

C. *Group Hub*

Having created a set of identities, it is natural to want to group them for authorization and related purposes. Globus Nexus provides powerful group management functions, with a particular emphasis on putting users in control of group creation, membership, and properties. Any authorized user can use intuitive Web and REST interfaces to create a group, define its properties (e.g., admission policies, visibility), and invite other users to join. Groups can be used in authorization

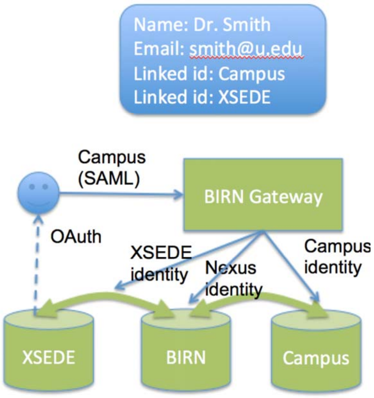

Figure 6: Globus Nexus identity hub example, showing use of multiple identities for different purposes.

decisions. As with other Globus Nexus services, interfaces can be skinned to meet the needs of specific communities.

Figure 7 shows an user view of the Globus Nexus group management interface, here skinned for kBase. The kBase project automatically enrolls every user who signs up for a kBase identity in the **kbase_user** group. Subgroups defined within that group are used to organize kBase users who participate in specific kBase project functions.

Figure 8 shows a different Globus Nexus group management interface, here indicating all groups that include the user "ian." The subscreen shows the policies that apply to that group, which govern visibility (the group is only visible to members) and membership (users must be invited to join).

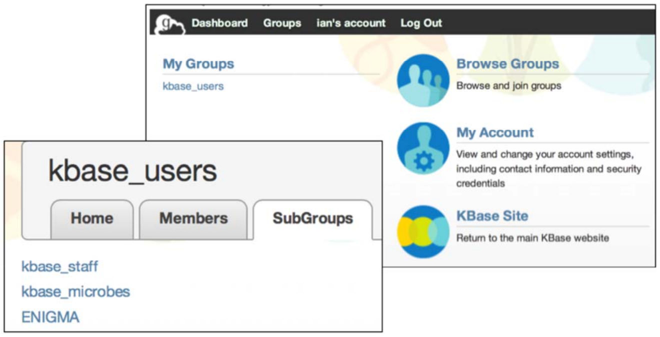

Browse and join groups

View and change your account settings, including contact information and security credentials

Return to the main KBase website

Figure 7: kBase group management interfaces.

III. Related Work

Many commercial service providers offer identity management and authentication services. Social network identities in particular, are commonly used for authentication by other services, for example Facebook Connect or Google Accounts. In research domains, CILogin is an example of an authentication framework that enables users to authenticate using a campus identity. Amazon Identity and Access Management (IAM) [6] enables user management across Amazon services and resources. It has been recently extended to include federation of public identities such as Facebook and Google. Like Nexus, these capabilities allow service developers to integrate their different identities into their services. The Atlassian Crowd [7] service provides identity management capabilities for web applications. It enables user identities to be sourced from several directories (e.g., LDAP) and exposes different authentication interfaces that can be embedded in external applications (e.g., OpenID). Both IAM and Crowd are commercial applications that require subscriptions; they are also designed to support commercial identity providers.

Group management and authorization services are also available. For example, Google Groups provides user defined groups that can be used for authorization to Google services. Previous work from Grid computing such as the Virtual Organization Management Service (VOMS) [8] provides group based authorization capabilities using short-lived proxy credentials to Grid resources. Atlassian Crowd also provides user defined groups that can be incorporated in external applications. Grouper [9] is perhaps most similar to Globus Nexus in its group management capabilities. Nexus is distinguished by its focus on user-driven group management.

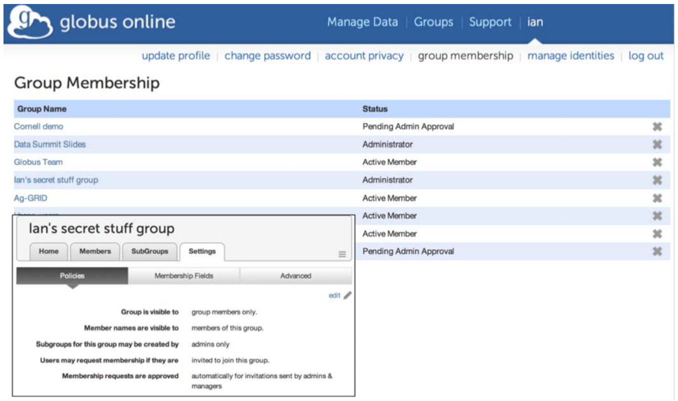

Group Membership

Group Name	Status
Cornell demo	Pending Admin Approval
Data Summit Slides	Administrator
Globus Team	Active Member
Ian's secret stuff group	Administrator
Ag-GRID	Active Member
	Active Member
	Active Member
	Pending Admin Approval

Group is visible to: group members only.
Member names are visible to: members of this group.
Subgroups for this group may be created by: admins only
Users may request membership if they are: invited to join this group.
Membership requests are approved: automatically for invitations sent by admins & managers

Figure 8: A user view of the groups visible to user "ian."

Acknowledgments

We thank the Globus Online team for their work implementing and operating Globus Nexus services. We also thank the XSEDE architecture team for their contributions to our understanding of requirements. This work was supported in part by the NIH through NIGMS grant U24 GM104203 Bio-Informatics Research Network Coordinating Center (BIRN-CC) and by the DOE through grant DE-AC02-06CH11357.

References

1. Wilkins-Diehr, N., *Science Gateways – Common Community Interfaces to Grid Resources.* Concurrency and Computation: Practice and Experience, 2007. **19**(6): p. 743–749.
2. Foster, I., *Globus Online: Accelerating and democratizing science through cloud-based services.* IEEE Internet Computing, 2011(May/June): p. 70-73.
3. Barnett, W., et al., *A Roadmap for Using NSF Cyberinfrastructure with InCommon*, 2011.
4. *CILogon Service*. September 18, 2013]; Available from: www.cilogon.org/service.
5. Helmer, K.G., et al., *Enabling collaborative research using the Biomedical Informatics Research Network (BIRN).* Journal of the American Medical Informatics Association, 2011.
6. *Amazon Identity and Access Management (IAM)*. September 18, 2013]; Available from: http://aws.amazon.com/iam/.
7. *Atlassian Crowd*. September 18, 2013]; Available from: https://www.atlassian.com/software/crowd/overview.
8. Alfieri, R., et al., *From gridmap-file to voms: managing authorization in a grid environment.* Future Generation Computer Systems, 2005. **21**(4): p. 549-558.
9. *Grouper groups management toolkit*. September 18, 2013]; Available from: www.internet2.edu/grouper/.