

Globus Authentication in Practice

Derek Simmel <dsimmel@psc.edu>

SAC-PA2 Workshop
June 14, 2018

Globus Authentication in Practice

Overview

- Introduction: Pittsburgh Supercomputing Center, XSEDE, Globus
- Globus Authentication Ecosystem
- Authentication with Globus Toolkit Services at PSC
- Federated Authentication with XSEDE for PSC HPC systems
- Globus ID and Globus Auth
- Observations and Challenges

Pittsburgh Supercomputing Center

Joint effort of **Carnegie Mellon University** and the **University of Pittsburgh**

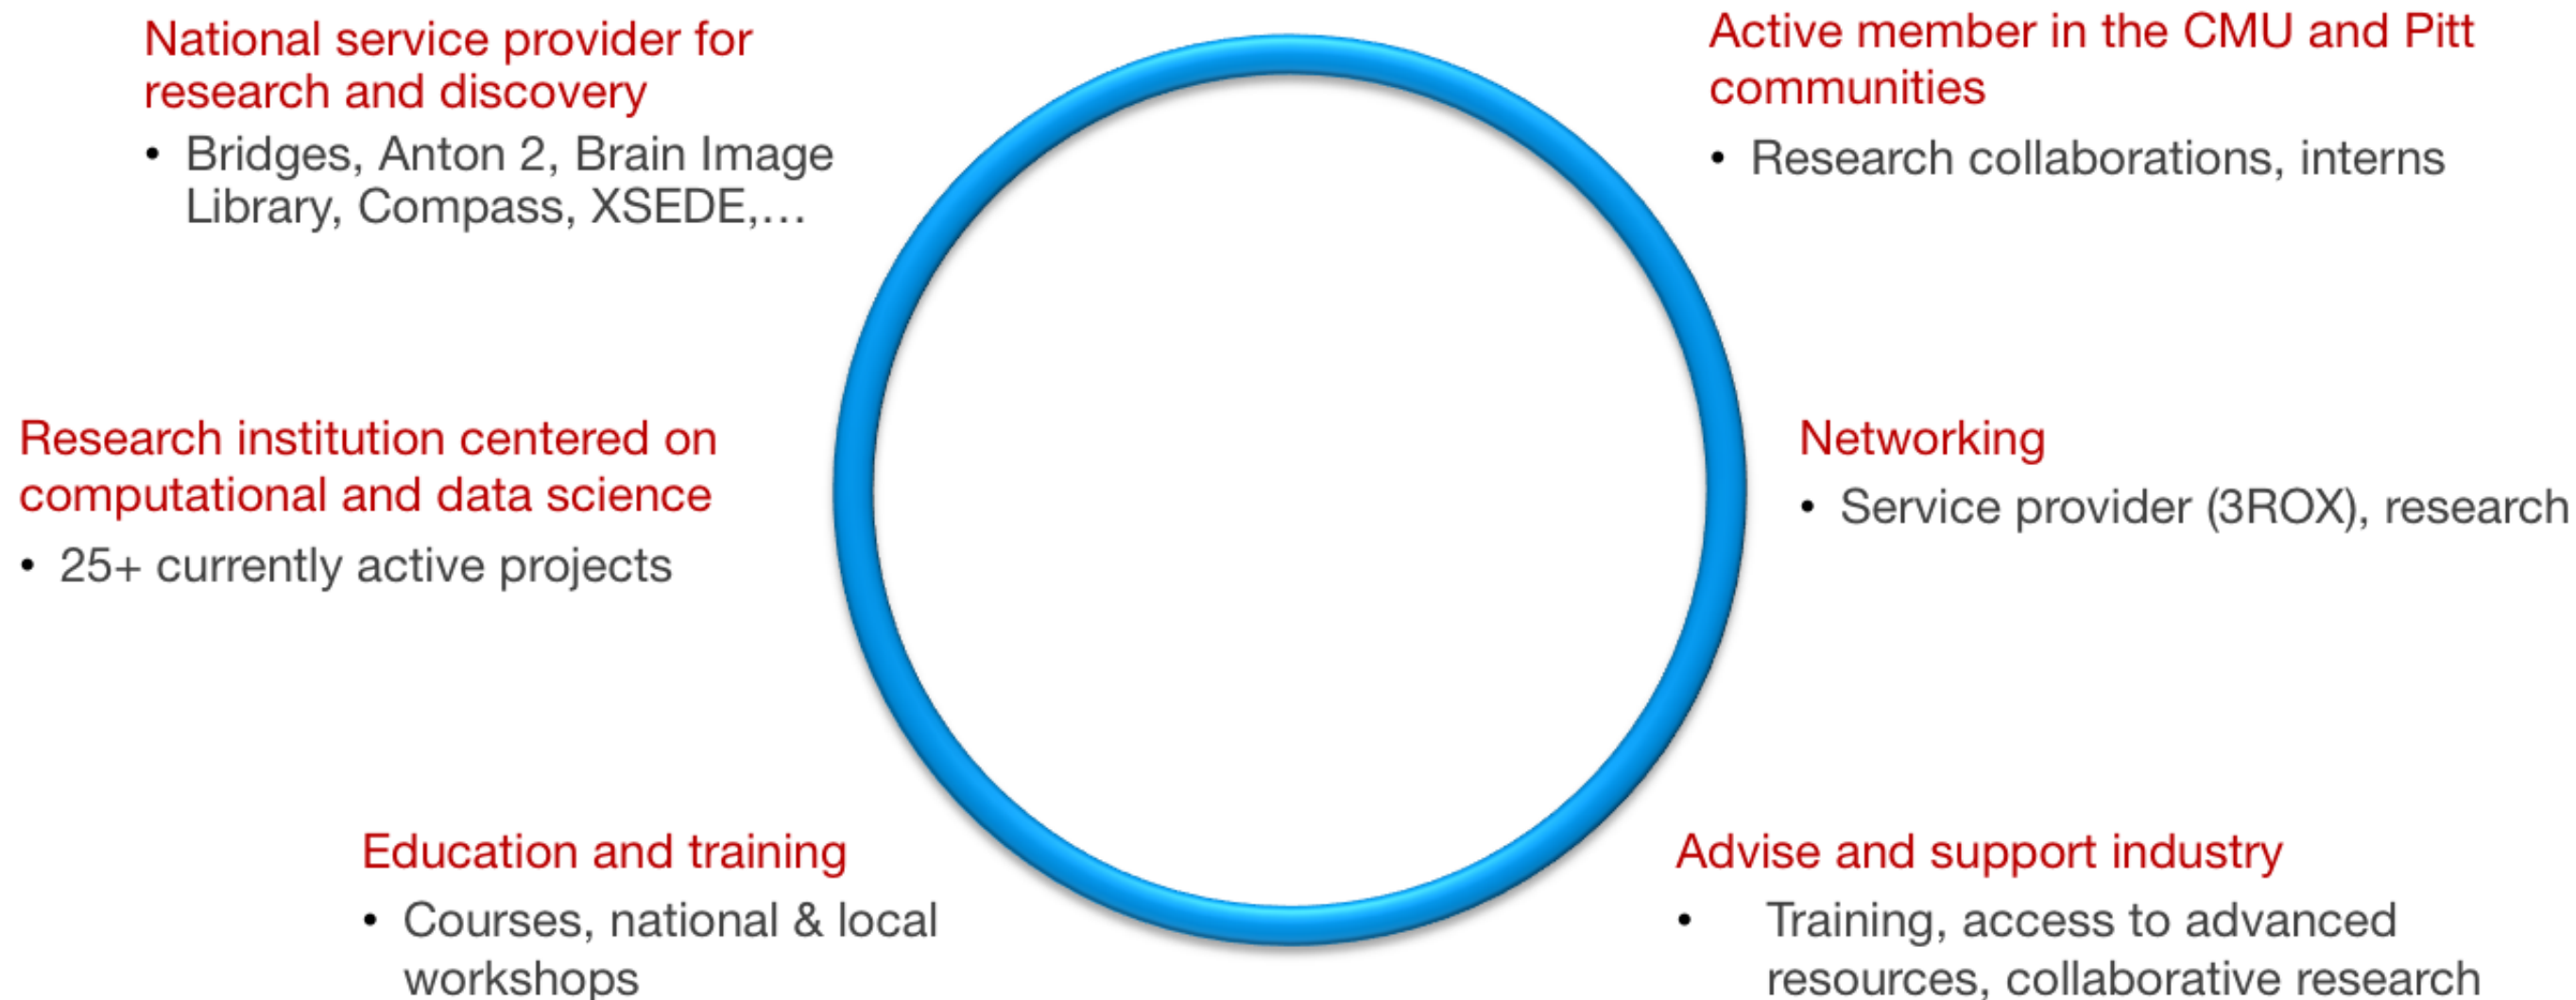
32 years of national leadership in:

- High-performance computing and data analytics (HPC, HPDA)
- 19 high-performance computers, including 9 that were/are “serial #1”
- Research groups: Artificial Intelligence & Big Data, Biomedical Applications, Public Health Applications, User Support for Scientific Applications, Networking, Security
- Software architecture, implementation, and optimization
- Networking and network optimization
- Enabling ground-breaking science, computer science, and engineering
- Leading research in AI, biology, public health, neuroscience, filesystems, networking, HPC software engineering, chemistry, materials science, engineering, physics, statistics, ...



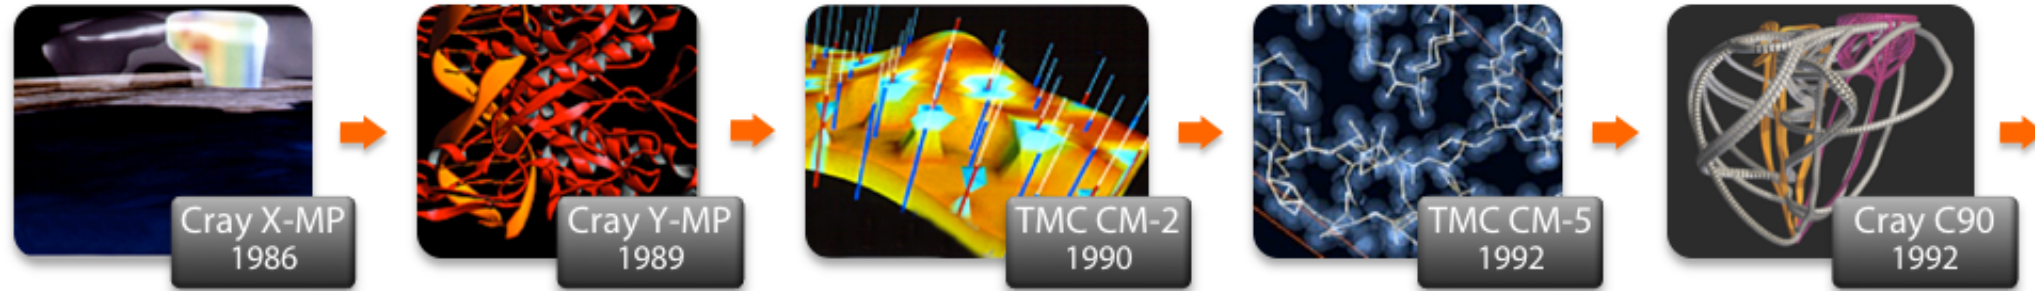
Supported by: NSF, NIH, DOE, DoD, the Commonwealth of Pennsylvania, D. E. Shaw Research, National Energy Technology Laboratory, Bill and Melinda Gates Foundation, UNICEF, USAID, Grable Foundation, Buhl Foundation, and industry



PSC's Multiple Roles



PSC's Systems Enable New Science




PSC Bridges



- Available at no cost for U.S. research
- *Pittsburgh Research Computing Initiative: easy application process for CMU & Pitt PIs*
- 28,628 CPU cores
- 64 NVIDIA P100 GPUs + 32 K80 GPUs
- 10 PB persistent storage + 7.3 PB node-local storage
- 274 TB memory, up to 12 TB per node
- 44M core-hours, 442k GPU-hours, and 343k TB-hours allocated quarterly
- Serving 1,260 projects and 5,814 users

- Funded by NSF award #ACI-1445606 (\$17.2M), Bridges emphasizes usability, flexibility, and interactivity
- Available at no charge for open research and course support and by arrangement to industry
- Popular programming languages and applications: Python, R, MATLAB, Java, Hadoop, Spark, ...
- 846 compute nodes containing 128GB (800), 3TB (42), and 12TB (4) of RAM each; 96 GPUs (nVidia P-100, K-80)
- Dedicated nodes for persistent databases, web servers, and distributed services
- The first deployment of the Intel Omni-Path Architecture fabric

PSC Bridges



Bridges Virtual Tour:
<https://psc.edu/bvt>

The eXtreme Science & Engineering Discovery Environment

XSEDE

Extreme Science and Engineering
Discovery Environment



- National Science Foundation-funded cyberinfrastructure (NSF OAC 1548562)
- Connects major U.S. High Performance Computing (HPC) centers
- Centralizes resource allocation to users of U.S. NSF-funded HPC resources
 - PSC, NCSA, SDSC, TACC + 15 partners
- Serves as a federated identity provider (IdP) across participating service providers

Globus is a grid infrastructure and middleware project, which has evolved to operate a “Platform as a Service” with subscription-based services

- Started in 1997 at the University of Chicago with Argonne National Lab, with support from DoE, NSF and NIH
- <https://www.globus.org> web-based GridFTP data transfer scheduling service with fault tolerance
- Core functionalities are provided in the Globus Toolkit (GTK)
 - *Grid Security Infrastructure* (GSI): Authentication with X.509 PKI
 - *GridFTP*: FTP-style data transfer with GSI + throughput optimization
 - *GSI-OpenSSH*: OpenSSH with GSI + PSC HPN enhancements
 - *MyProxy CA* and Credential Management service
 - Globus support for GTK retired Dec. 31 2017; Security updates only thru Dec. 31 2018
 - *Grid Community Toolkit* open source available at <https://github.com/gridcf/gct>
 - All current and future development moved to Globus Connect Server (GCS) platform with Globus Auth
- Globus ID: <https://globusid.org> A Globus-managed identity and authorization “consents” management service
- Globus Auth: OpenID Connect + OAuth2 authentication and authorization service

Globus Authentication Ecosystem

- XSEDE Service Providers (PSC, NCSA, SDSC, TACC,...) currently run services based on the Globus Toolkit
 - GSI based authentication using X.509 PKI
 - Enables Single Sign-On for Users across all XSEDE resources
 - XSEDE operates online MyProxy CA services for user authentication
 - GSI-OpenSSH and Globus GridFTP services with host certificates from approved Certificate Authorities (e.g., InCommon IGTF Server CA)
 - Users may use X.509 credentials from other XSEDE-approved CAs
 - CILogon as a Basic LoA authentication service with InCommon and EduGain member Identity Providers
- XSEDE and CILogon OAuth+MyProxy authentication chain to avoid disclosing XSEDE usernames and passwords to Globus services

GSI Authentication Setup

- **/etc/grid-security/certificates/**
 - contains the certificates and associated policy files for all CAs that you choose to trust
 - For XSEDE, we get these from the IGTF distribution of CA certificates for CAs that have been accredited by IGTF (TAGPMA, EUGridPMA, APGridPMA)
- **/etc/grid-security/grid-mapfile**
 - contains pairwise mappings of approved (user) certificate subjects (a.k.a. Distinguished Names, or DNs) to the local username on the system
 - Many-to-One mappings ONLY: many certificate subjects may map to a single username
 - a certificate can only be valid for one specific end entity (user, system, device or service)
 - a user can have several certificates
 - the same certificate subject must NOT be mapped to multiple different users

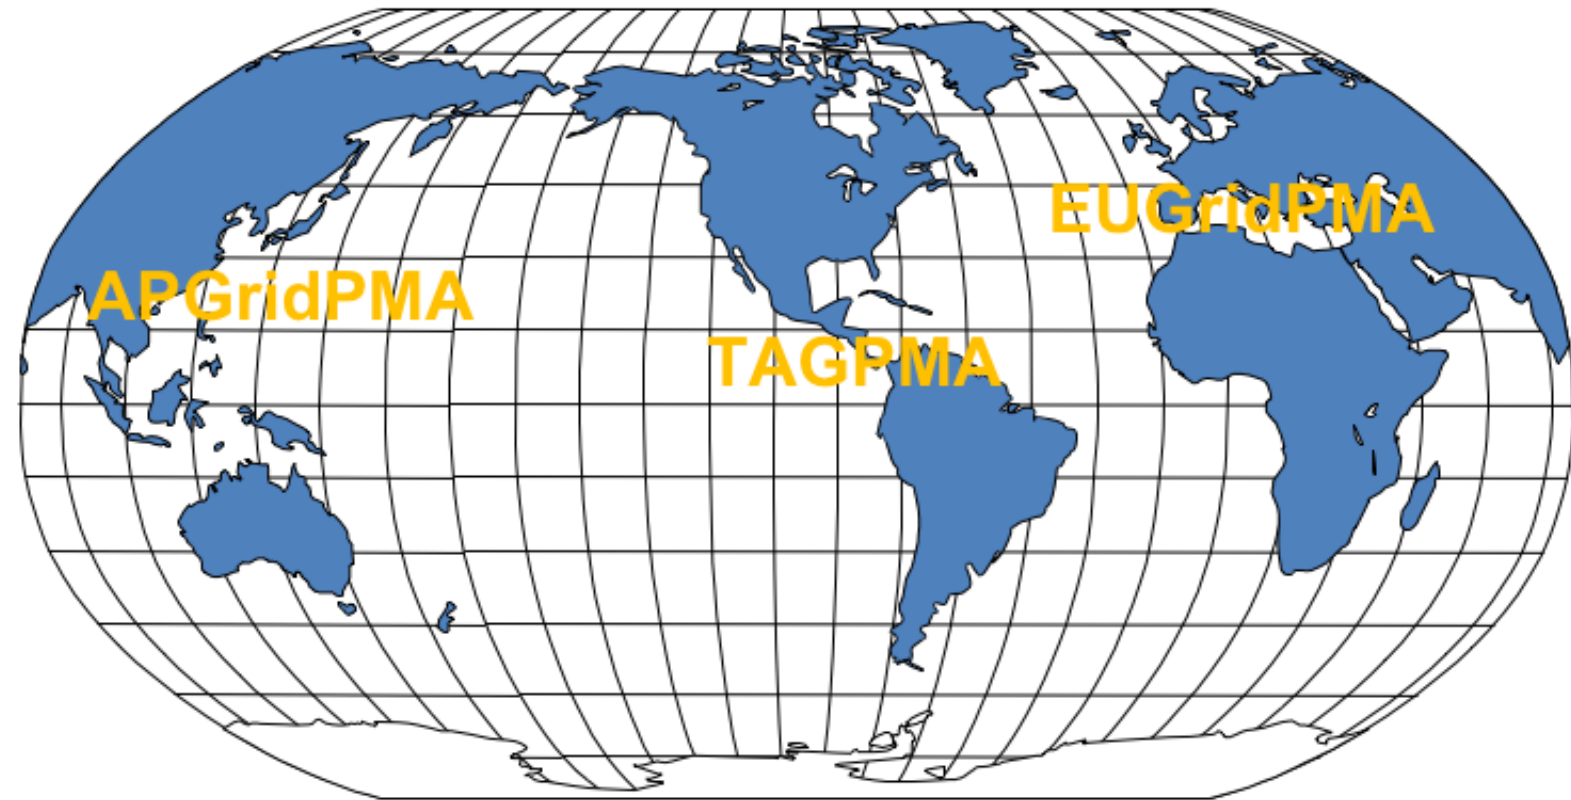
Authentication with Globus Toolkit Services at PSC

- PSC production systems support a variety of methods for authentication
 - PSC local (Kerberos) username and password
 - SSH Keys
 - GSI-OpenSSH with user credentials from XSEDE-approved CAs
 - Federated (XSEDE) username and password with XSEDE DUO MFA
- For Globus file transfers to/from restricted filesystem space, a separate GridFTP service is provided, with authorized credentials limited to those from Identity Providers that locally require multifactor authentication
 - e.g., CILogon with CMU local authentication plus DUO MFA

Interoperable Global Trust Federation (IGTF)

IGTF <https://www.igtf.net>

- Defines and publishes standards for authentication assurance
- Accredits authentication providers with specific profiles
- Distributes trust anchors for accredited authentication providers
- Comprised by members in three regional policy management authorities
 - TAGPMA (Americas)
 - EUGridPMA (EU & Middle East)
 - APGridPMA (Asia/Pacific)

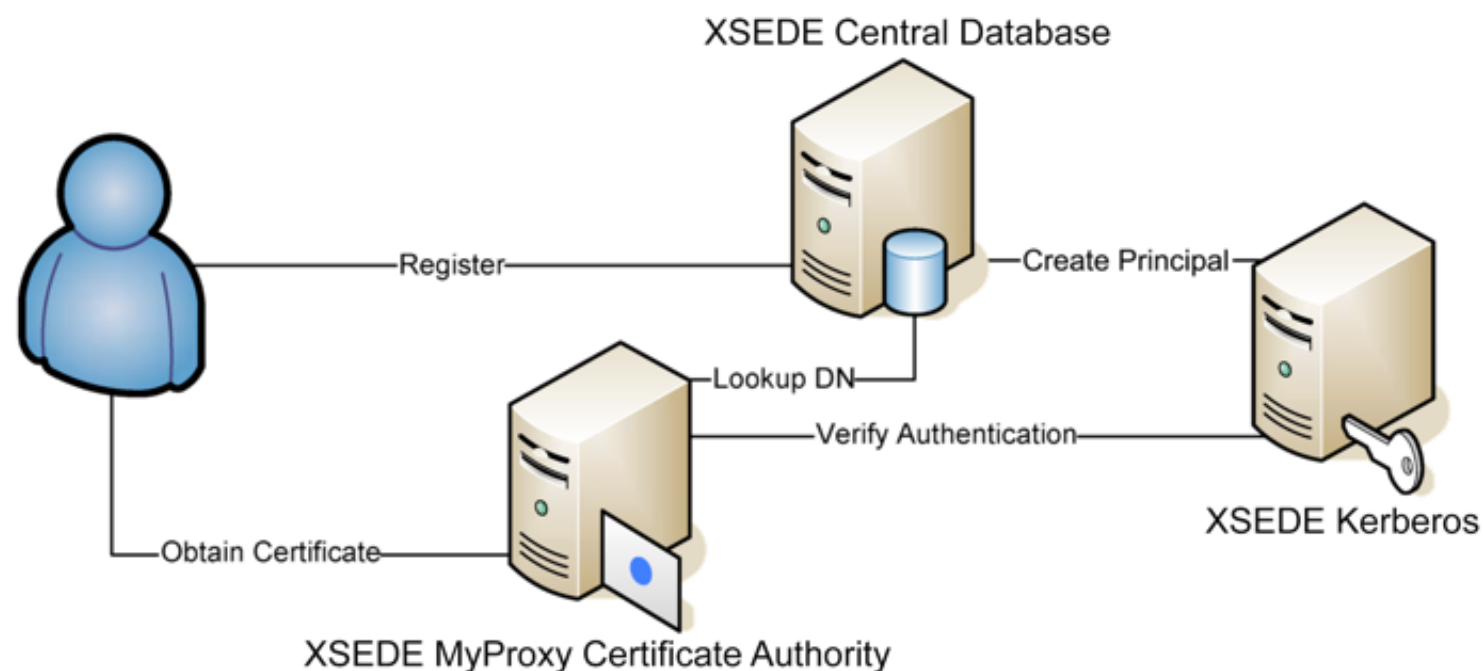


SSH and GSI-OpenSSH

- Secure Shell (SSH) has been the de facto interactive login and data transfer service on Linux systems for over 20 years.
- GSI-OpenSSH is a customized edition of the popular OpenSSH implementation of SSH, developed by the National Center for Supercomputing Applications (NCSA).
- GSI-OpenSSH adds Globus Grid Security Infrastructure (GSI) authentication utilizing X.509 credentials and user mapping:
 - Users' X.509 certificate subjects (a.k.a. Distinguished Names, DNs) are mapped to their local username on the login host via a grid-mapfile
 - GSI-OpenSSH adds a configuration option, **PermitPAMUserChange**, which enables the current username value to be substituted during Linux-PAM processing

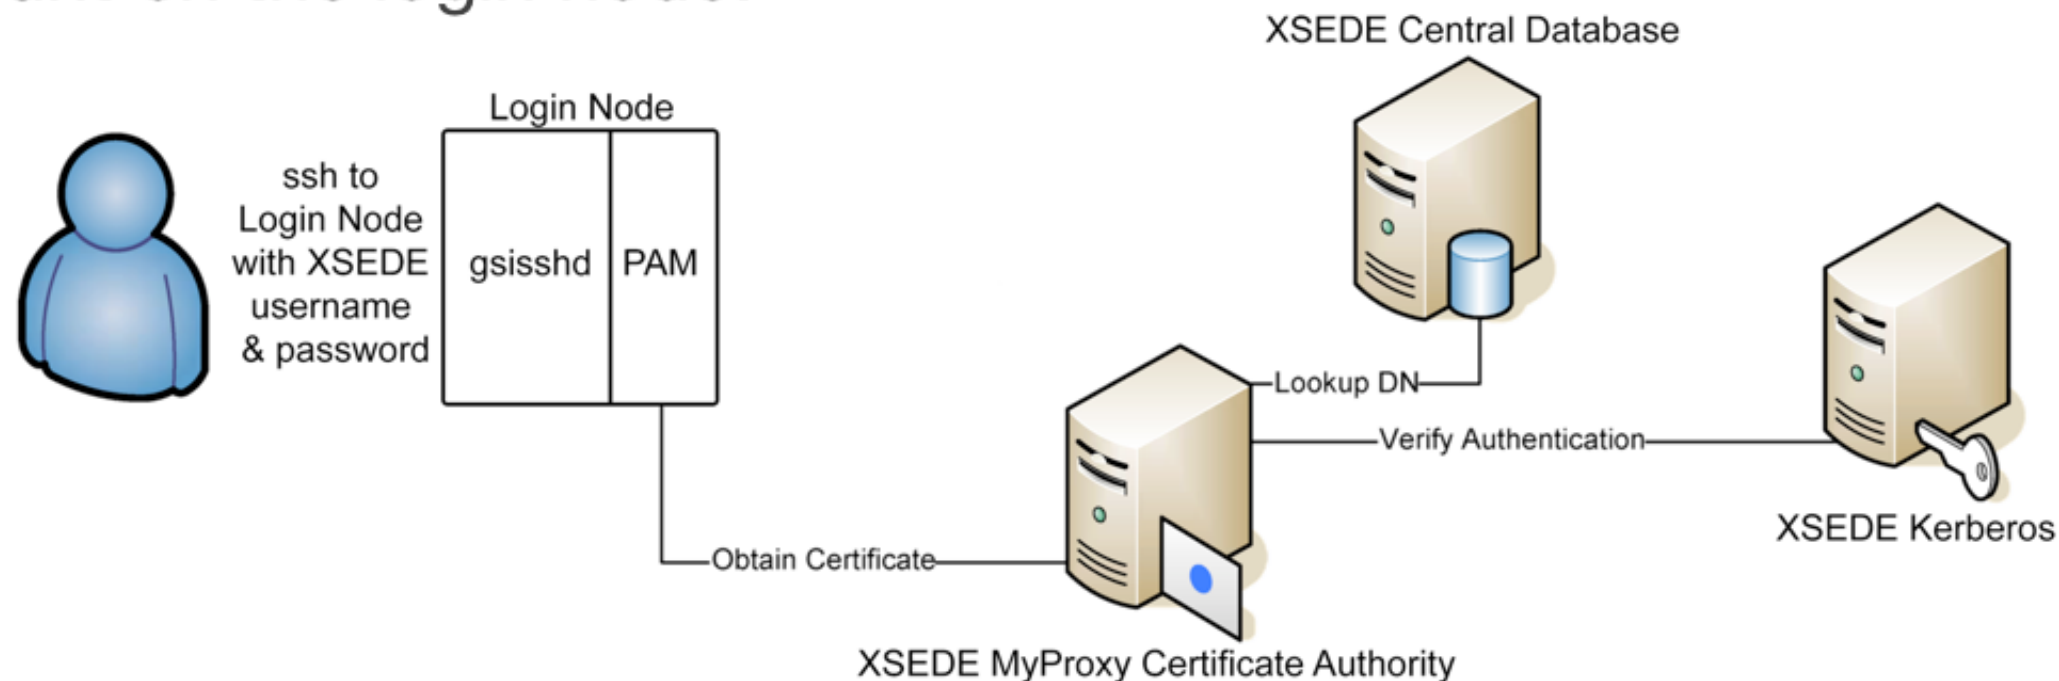
XSEDE Single Sign-On (SSO)

- To facilitate Single Sign-On across XSEDE resources, XSEDE Service Providers deploy Globus GSI-authenticated services, including GSI-OpenSSH and Globus GridFTP.
- To simplify issuance of X.509 user certificates, XSEDE operates a MyProxy Certificate Authority (CA) Service:



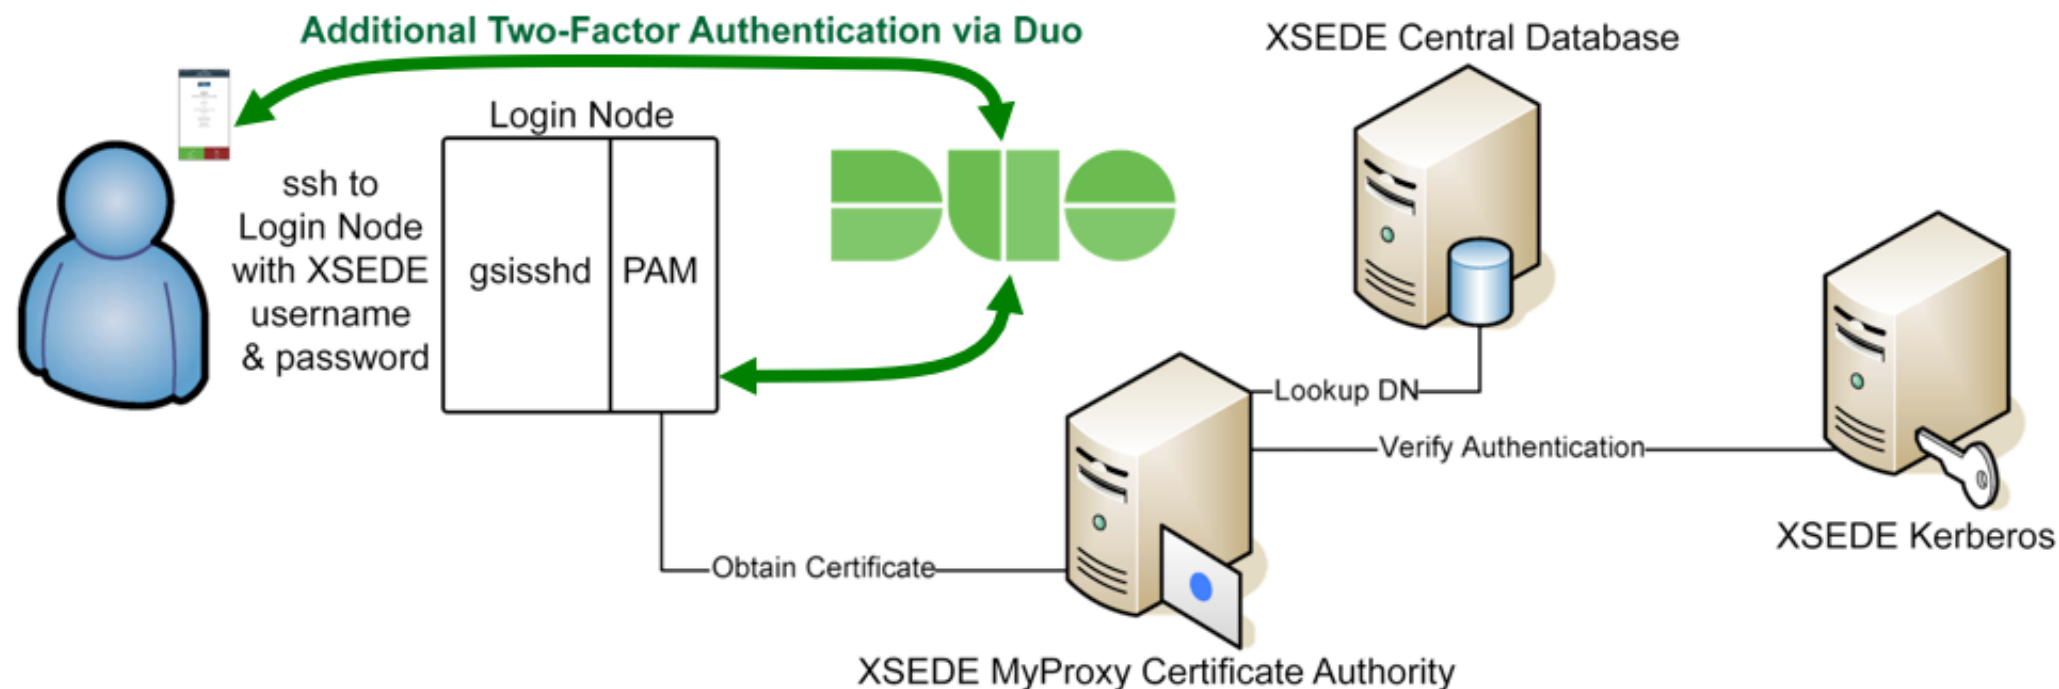
XSEDE-SSO with MyProxy-Enabled GSISSH

- To further simplify XSEDE SSO on login nodes, GSI-OpenSSH is deployed with Linux-PAM for authentication.
- Via PAM, a user certificate is automatically retrieved from the XSEDE MyProxy CA for the user, which is then used to map the user to their local account on the login node.



XSEDE SSO Login with Duo Two-Factor AuthN

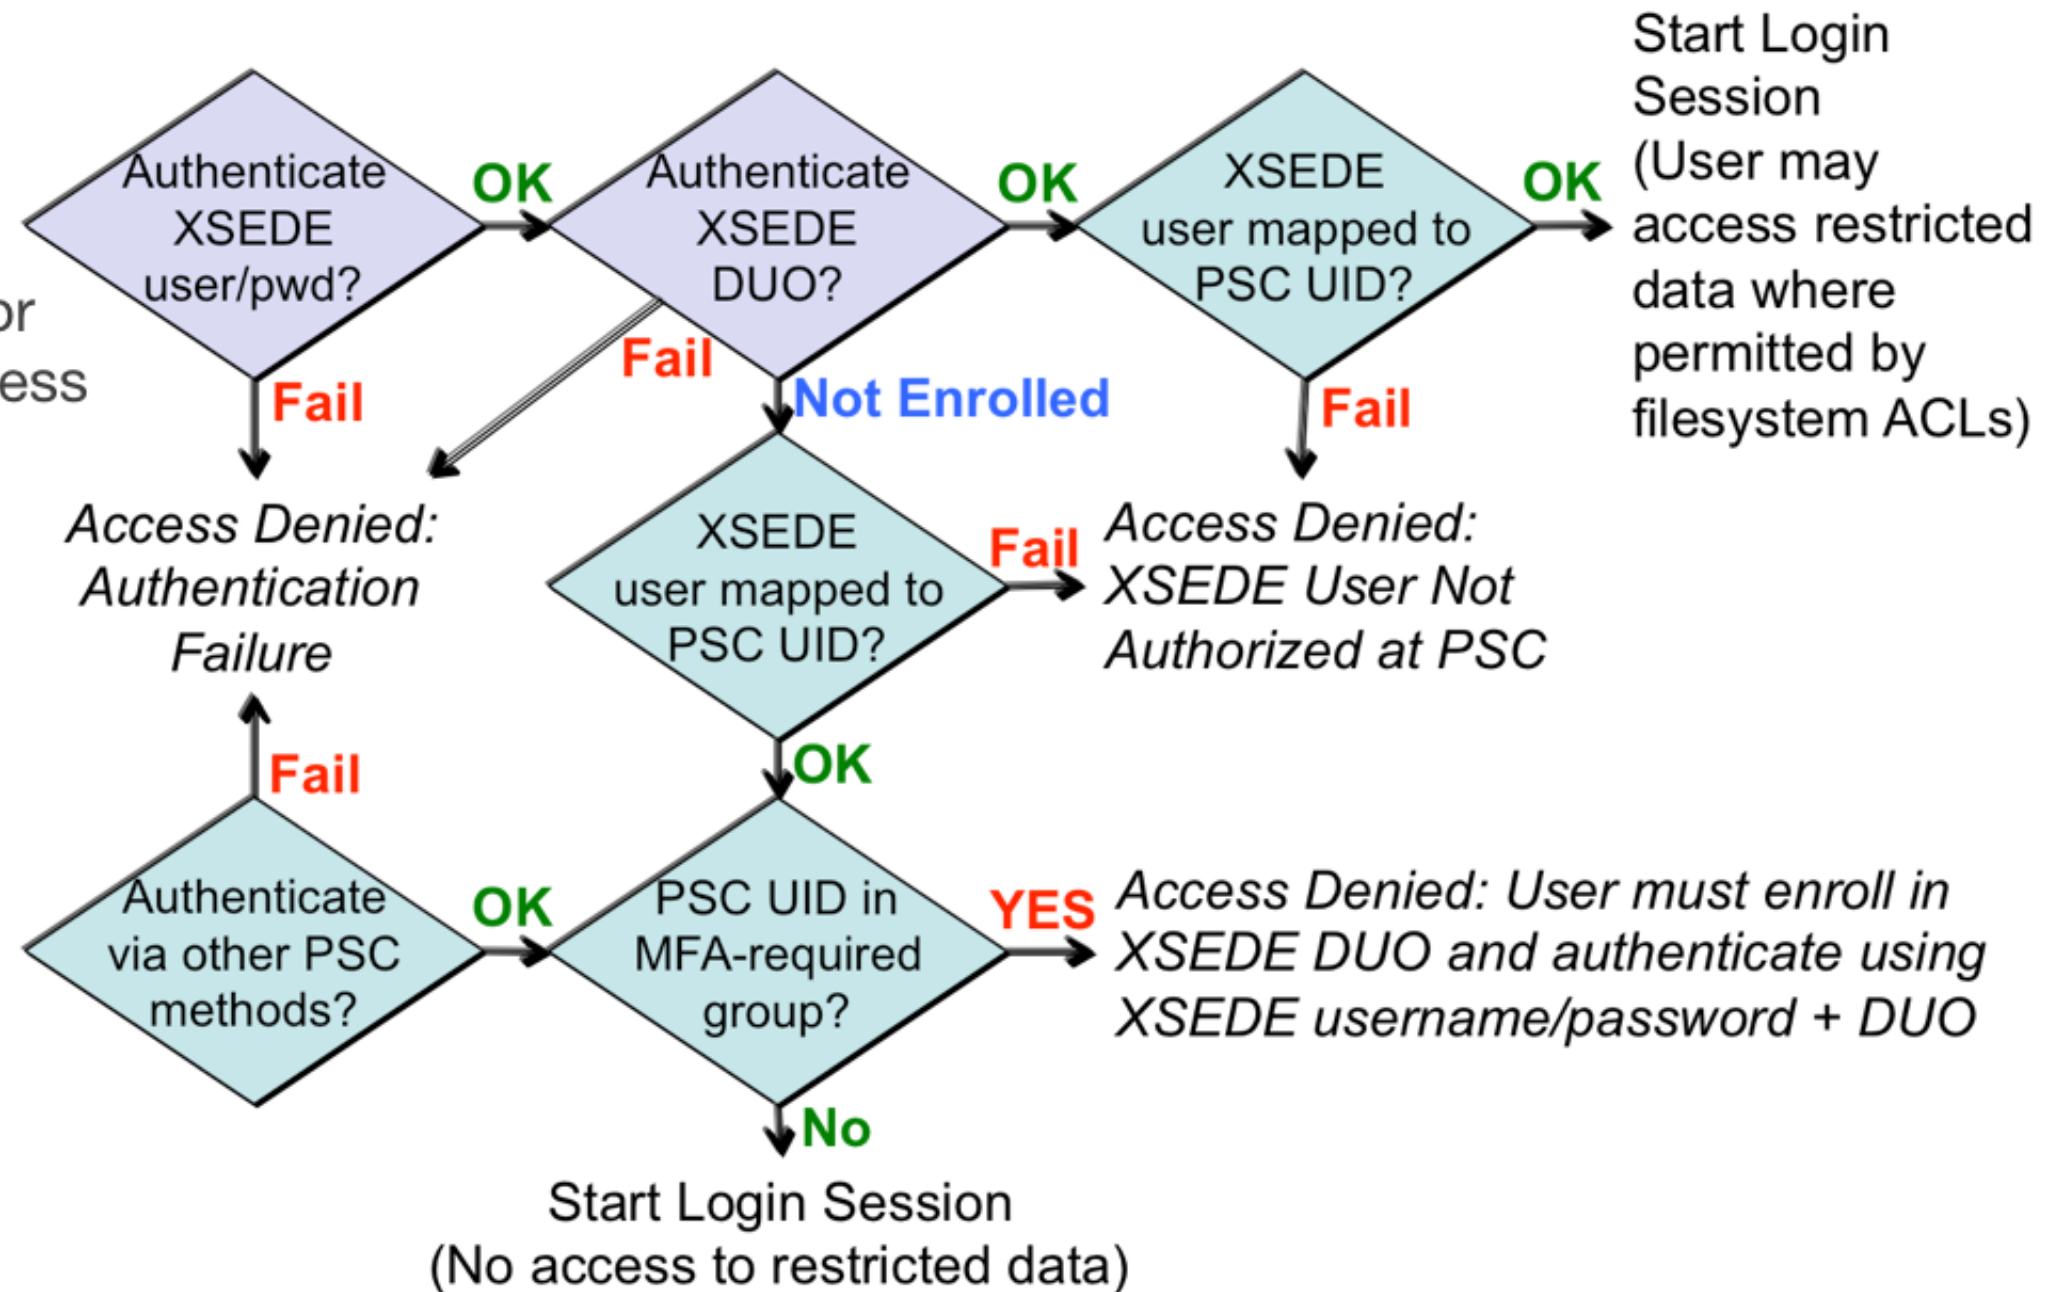
- For stronger user authentication, XSEDE subscribes to the Duo two-factor authentication(TFA) service.
 - XSEDE users enroll in XSEDE Duo TFA via the XSEDE User Portal
- Duo provides a PAM module (pam_duo) to incorporate additional authentication into PAM-compatible services.



PSC Bridges Federated User Authentication with MFA

Some, but not all Bridges users are required to authenticate using multifactor authentication (MFA) for access to restricted datasets

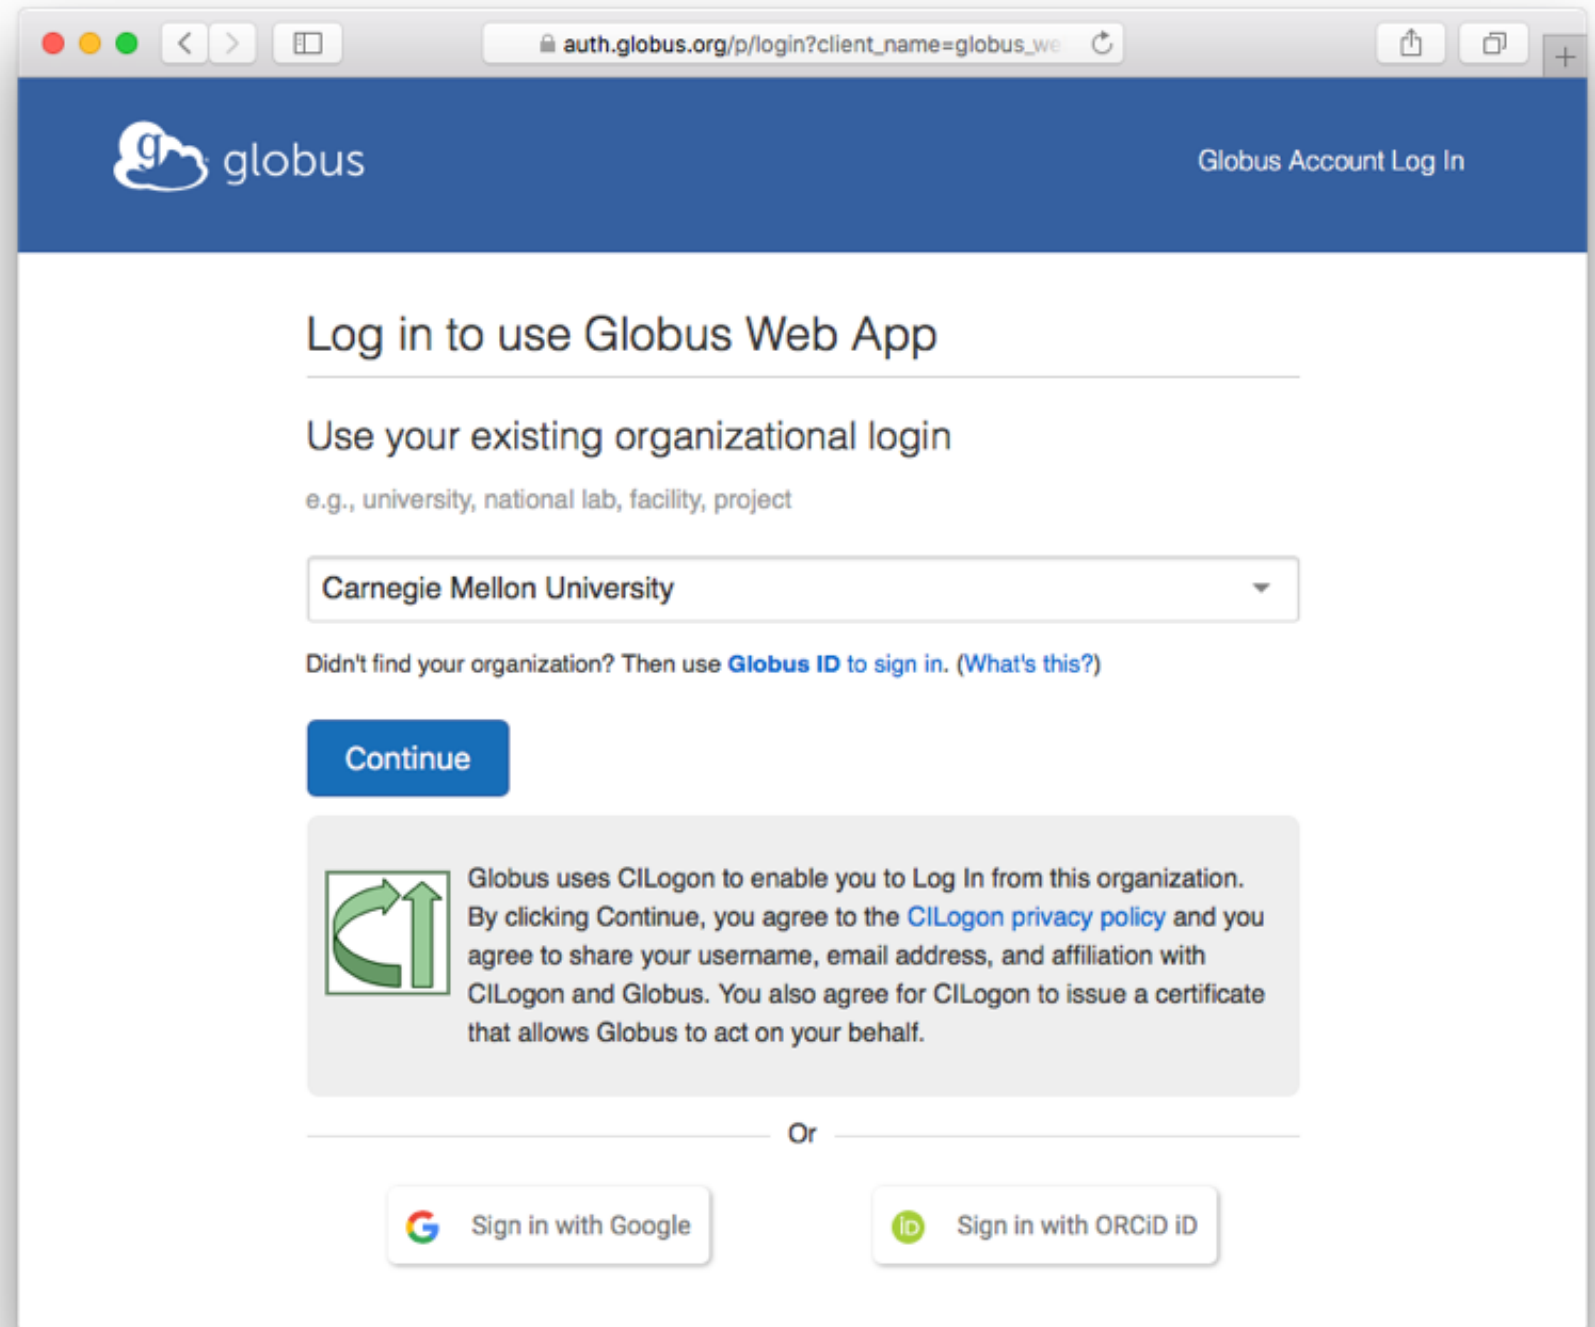
PSC implemented modifications to DUO's SSH integration to facilitate flexible enforcement of multifactor authentication while preserving Single Sign-On for users with XSEDE federated IDs



https://github.com/pscedu/duo_unix_psc

GlobusID

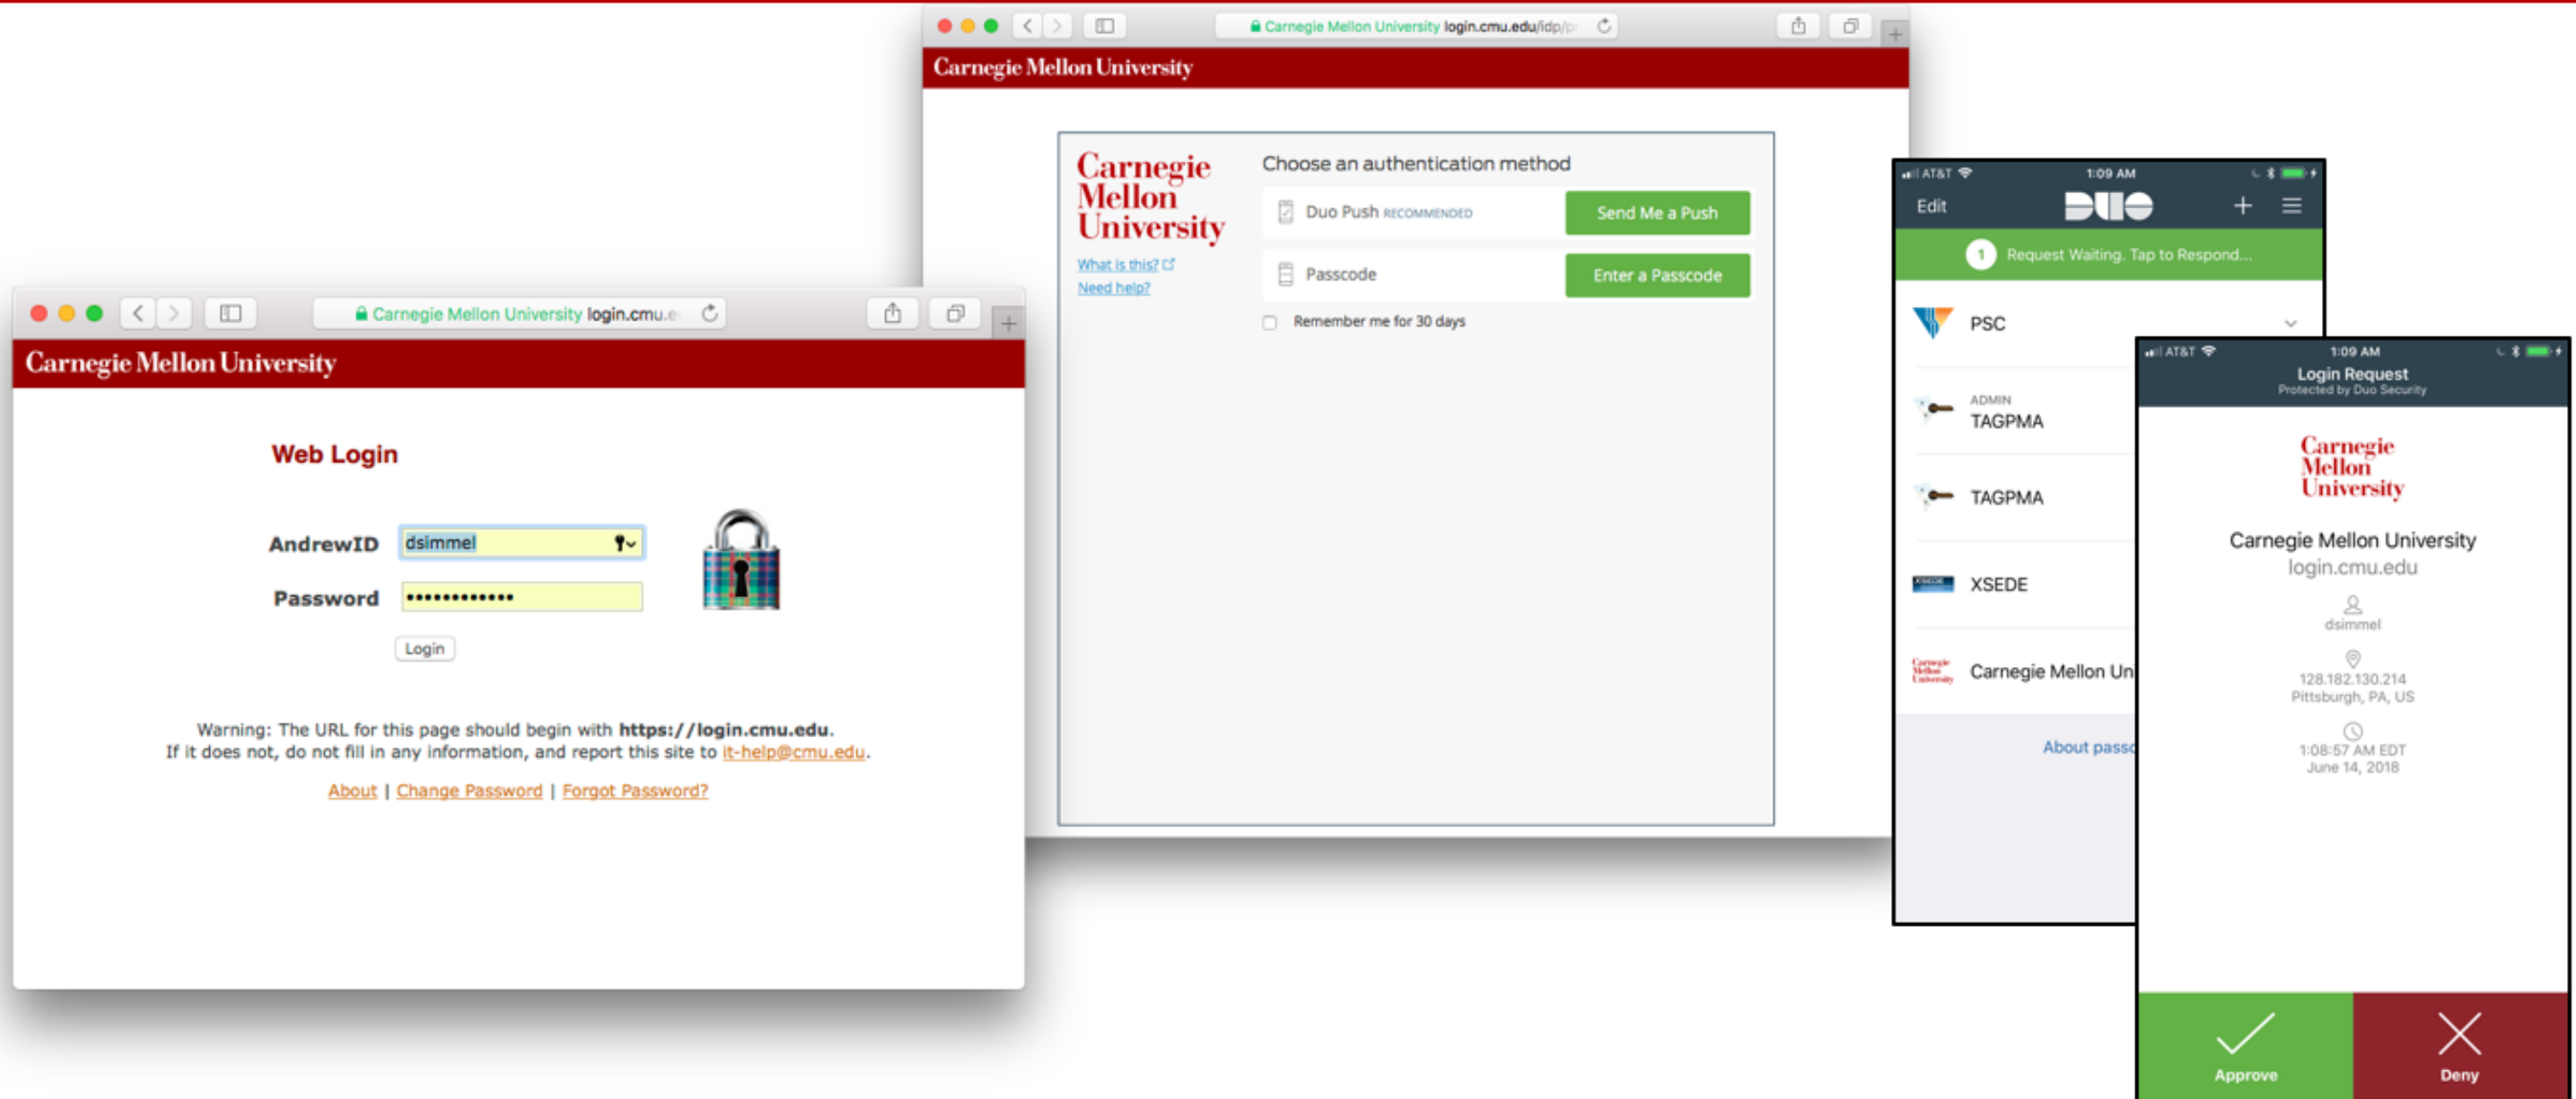
- Globus ID is an identity and authorization management service
- It allows users to employ a number of different identities, and to manage “consents” to expose attributes and credentials to remote services



The screenshot shows a web browser window with the URL `auth.globus.org/p/login?client_name=globus_we`. The page has a blue header with the Globus logo and a "Globus Account Log In" link. The main content area is white and contains the following elements:

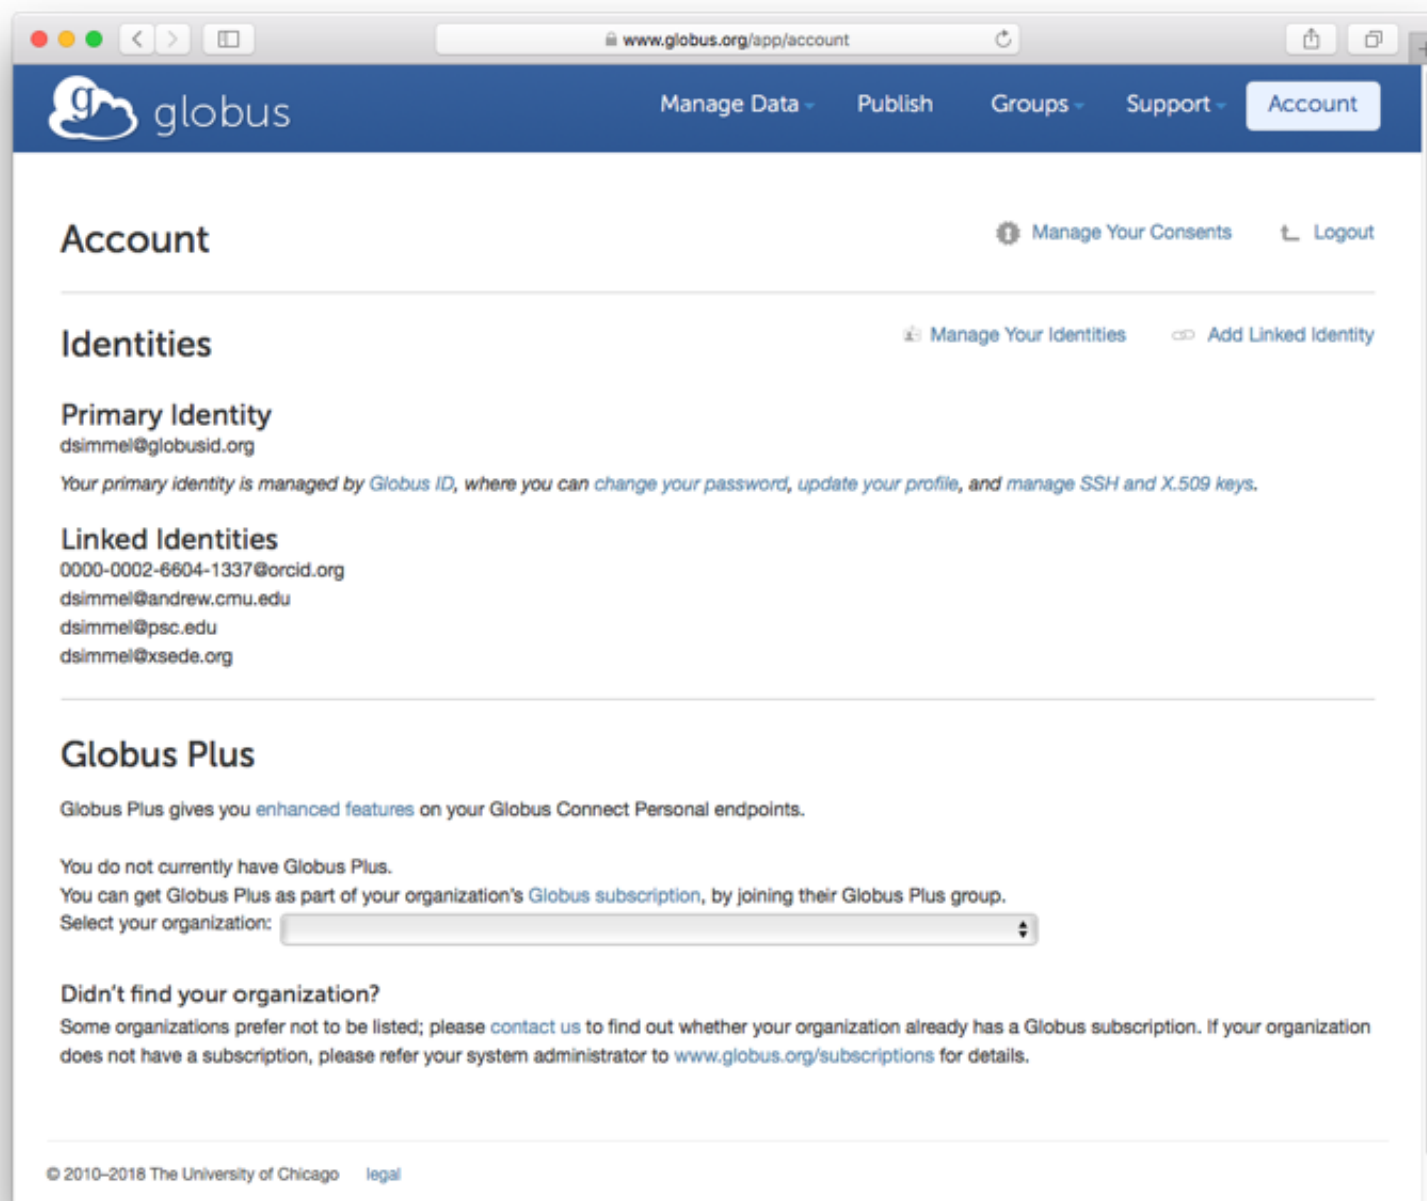
- A heading: "Log in to use Globus Web App"
- A sub-heading: "Use your existing organizational login"
- A note: "e.g., university, national lab, facility, project"
- A dropdown menu showing "Carnegie Mellon University"
- A link: "Didn't find your organization? Then use [Globus ID](#) to sign in. (What's this?)"
- A blue "Continue" button
- A consent box with a green circular arrow icon and the text: "Globus uses CILogon to enable you to Log In from this organization. By clicking Continue, you agree to the [CILogon privacy policy](#) and you agree to share your username, email address, and affiliation with CILogon and Globus. You also agree for CILogon to issue a certificate that allows Globus to act on your behalf."
- A separator line with the word "Or" in the center
- Two buttons at the bottom: "Sign in with Google" (with the Google logo) and "Sign in with ORCID iD" (with the ORCID logo)

Login via CILogon with CMU IdP and DUO MFA



- CILogon redirects to [CMU] for user authentication, returns an OAuth token

GlobusID Account & Identities



Account

Manage Your Consents Logout

Identities

Manage Your Identities Add Linked Identity

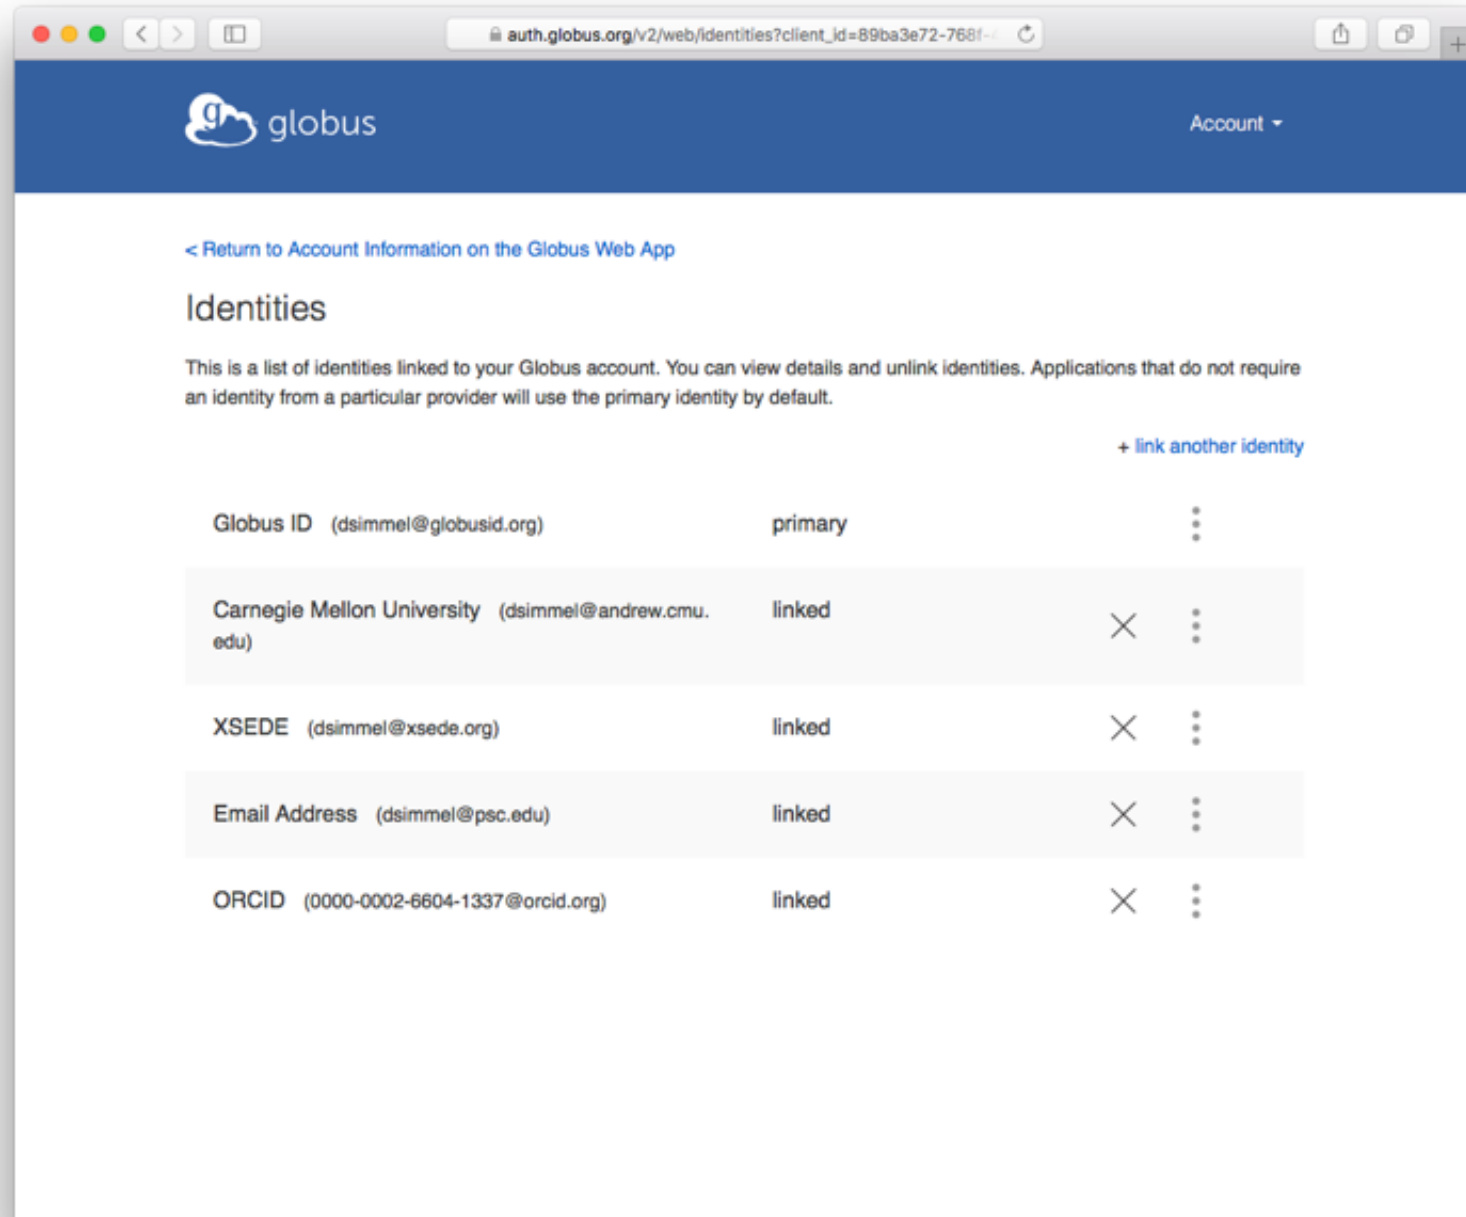
Primary Identity
dsimmel@globusid.org
Your primary identity is managed by Globus ID, where you can change your password, update your profile, and manage SSH and X.509 keys.

Linked Identities
0000-0002-6604-1337@orcid.org
dsimmel@andrew.cmu.edu
dsimmel@psc.edu
dsimmel@xsede.org

Globus Plus
Globus Plus gives you enhanced features on your Globus Connect Personal endpoints.
You do not currently have Globus Plus.
You can get Globus Plus as part of your organization's Globus subscription, by joining their Globus Plus group.
Select your organization:

Didn't find your organization?
Some organizations prefer not to be listed; please [contact us](#) to find out whether your organization already has a Globus subscription. If your organization does not have a subscription, please refer your system administrator to www.globus.org/subscriptions for details.

© 2010–2018 The University of Chicago [legal](#)



Identities

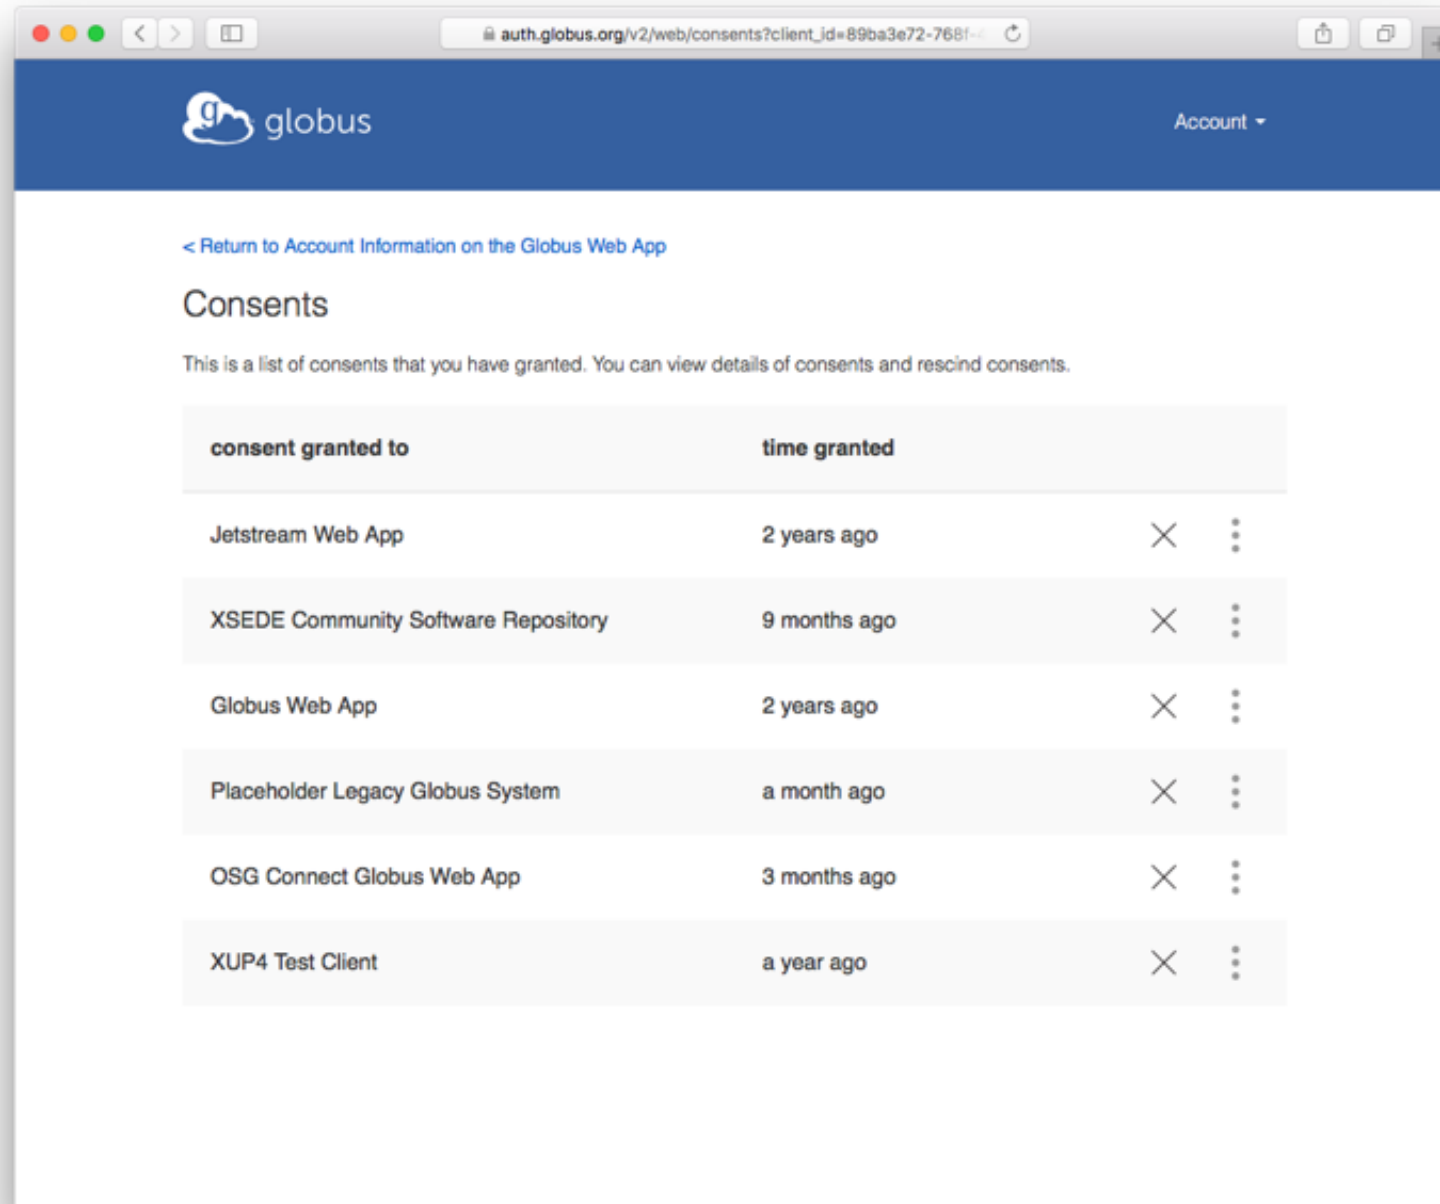
< Return to Account Information on the Globus Web App

This is a list of identities linked to your Globus account. You can view details and unlink identities. Applications that do not require an identity from a particular provider will use the primary identity by default.

[+ link another identity](#)

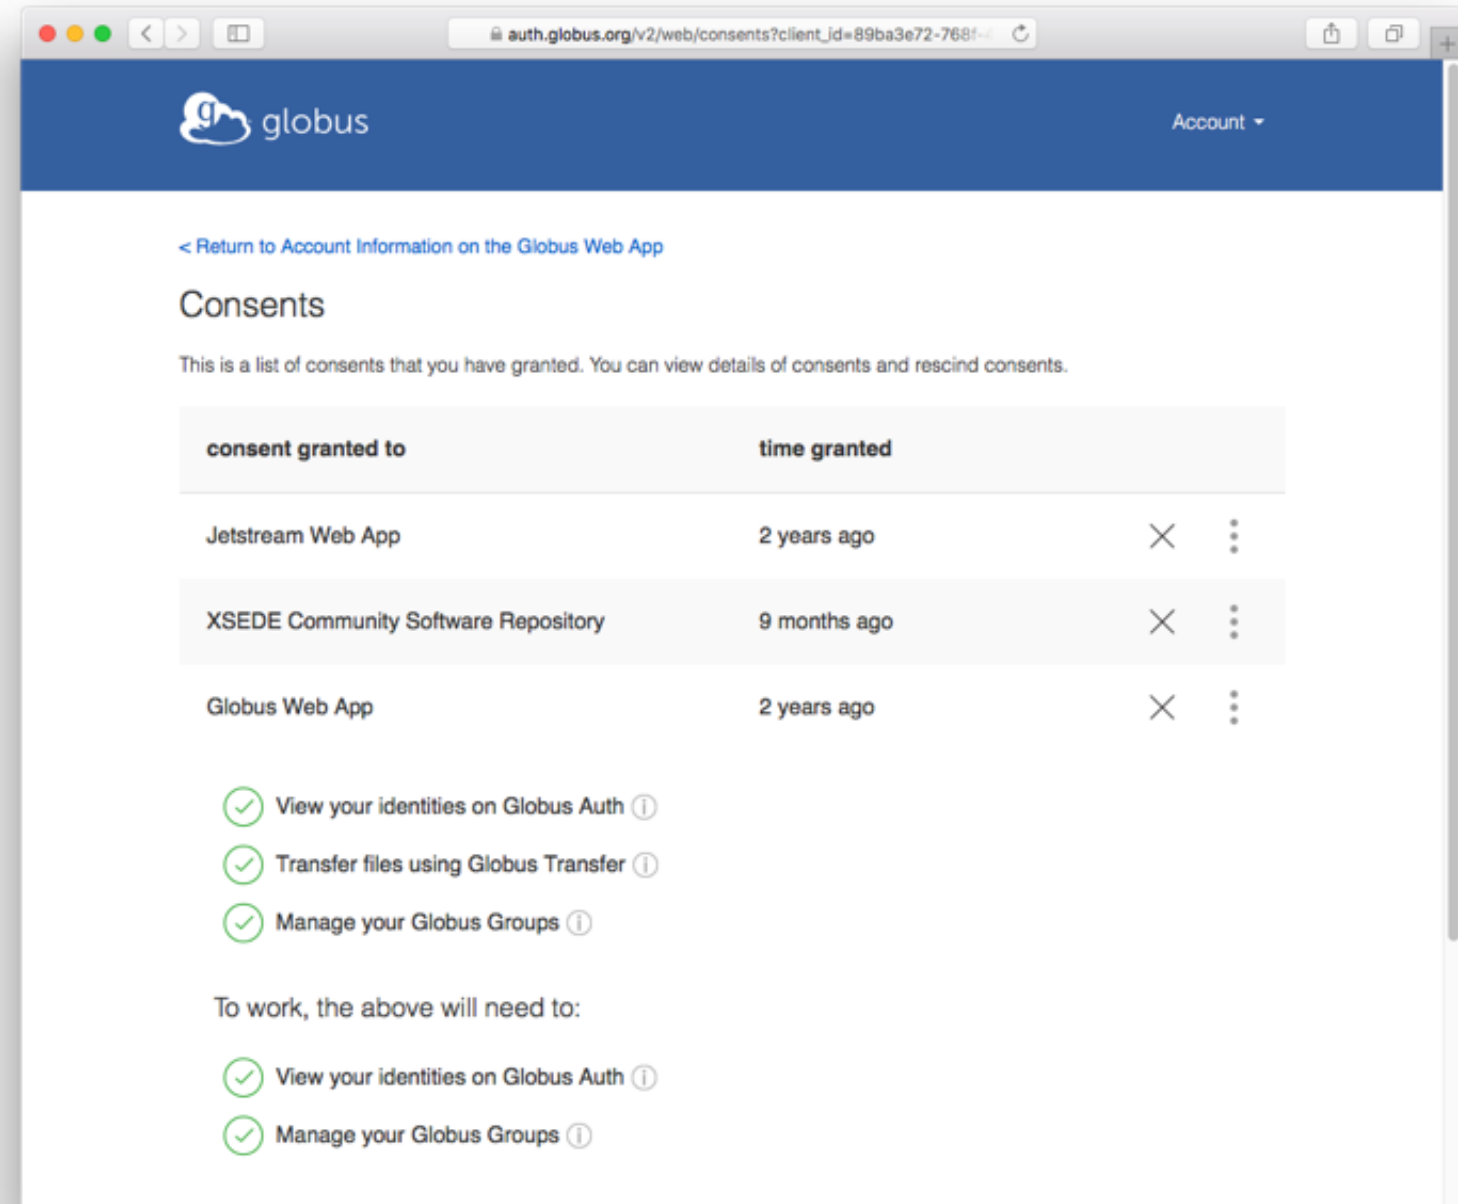
Globus ID (dsimmel@globusid.org)	primary		
Carnegie Mellon University (dsimmel@andrew.cmu.edu)	linked	×	
XSEDE (dsimmel@xsede.org)	linked	×	
Email Address (dsimmel@psc.edu)	linked	×	
ORCID (0000-0002-6604-1337@orcid.org)	linked	×	

GlobusID Authorization “Consents”



The screenshot shows the 'Consents' page in the Globus Auth interface. The page title is 'Consents' and it includes a link to '< Return to Account Information on the Globus Web App'. Below the title, a message states: 'This is a list of consents that you have granted. You can view details of consents and rescind consents.' A table lists the following consents:

consent granted to	time granted		
Jetstream Web App	2 years ago	✕	⋮
XSEDE Community Software Repository	9 months ago	✕	⋮
Globus Web App	2 years ago	✕	⋮
Placeholder Legacy Globus System	a month ago	✕	⋮
OSG Connect Globus Web App	3 months ago	✕	⋮
XUP4 Test Client	a year ago	✕	⋮



The screenshot shows the 'Consents' page in the Globus Auth interface. The page title is 'Consents' and it includes a link to '< Return to Account Information on the Globus Web App'. Below the title, a message states: 'This is a list of consents that you have granted. You can view details of consents and rescind consents.' A table lists the following consents:

consent granted to	time granted		
Jetstream Web App	2 years ago	✕	⋮
XSEDE Community Software Repository	9 months ago	✕	⋮
Globus Web App	2 years ago	✕	⋮

Below the table, there are three green checkmarks indicating granted permissions:

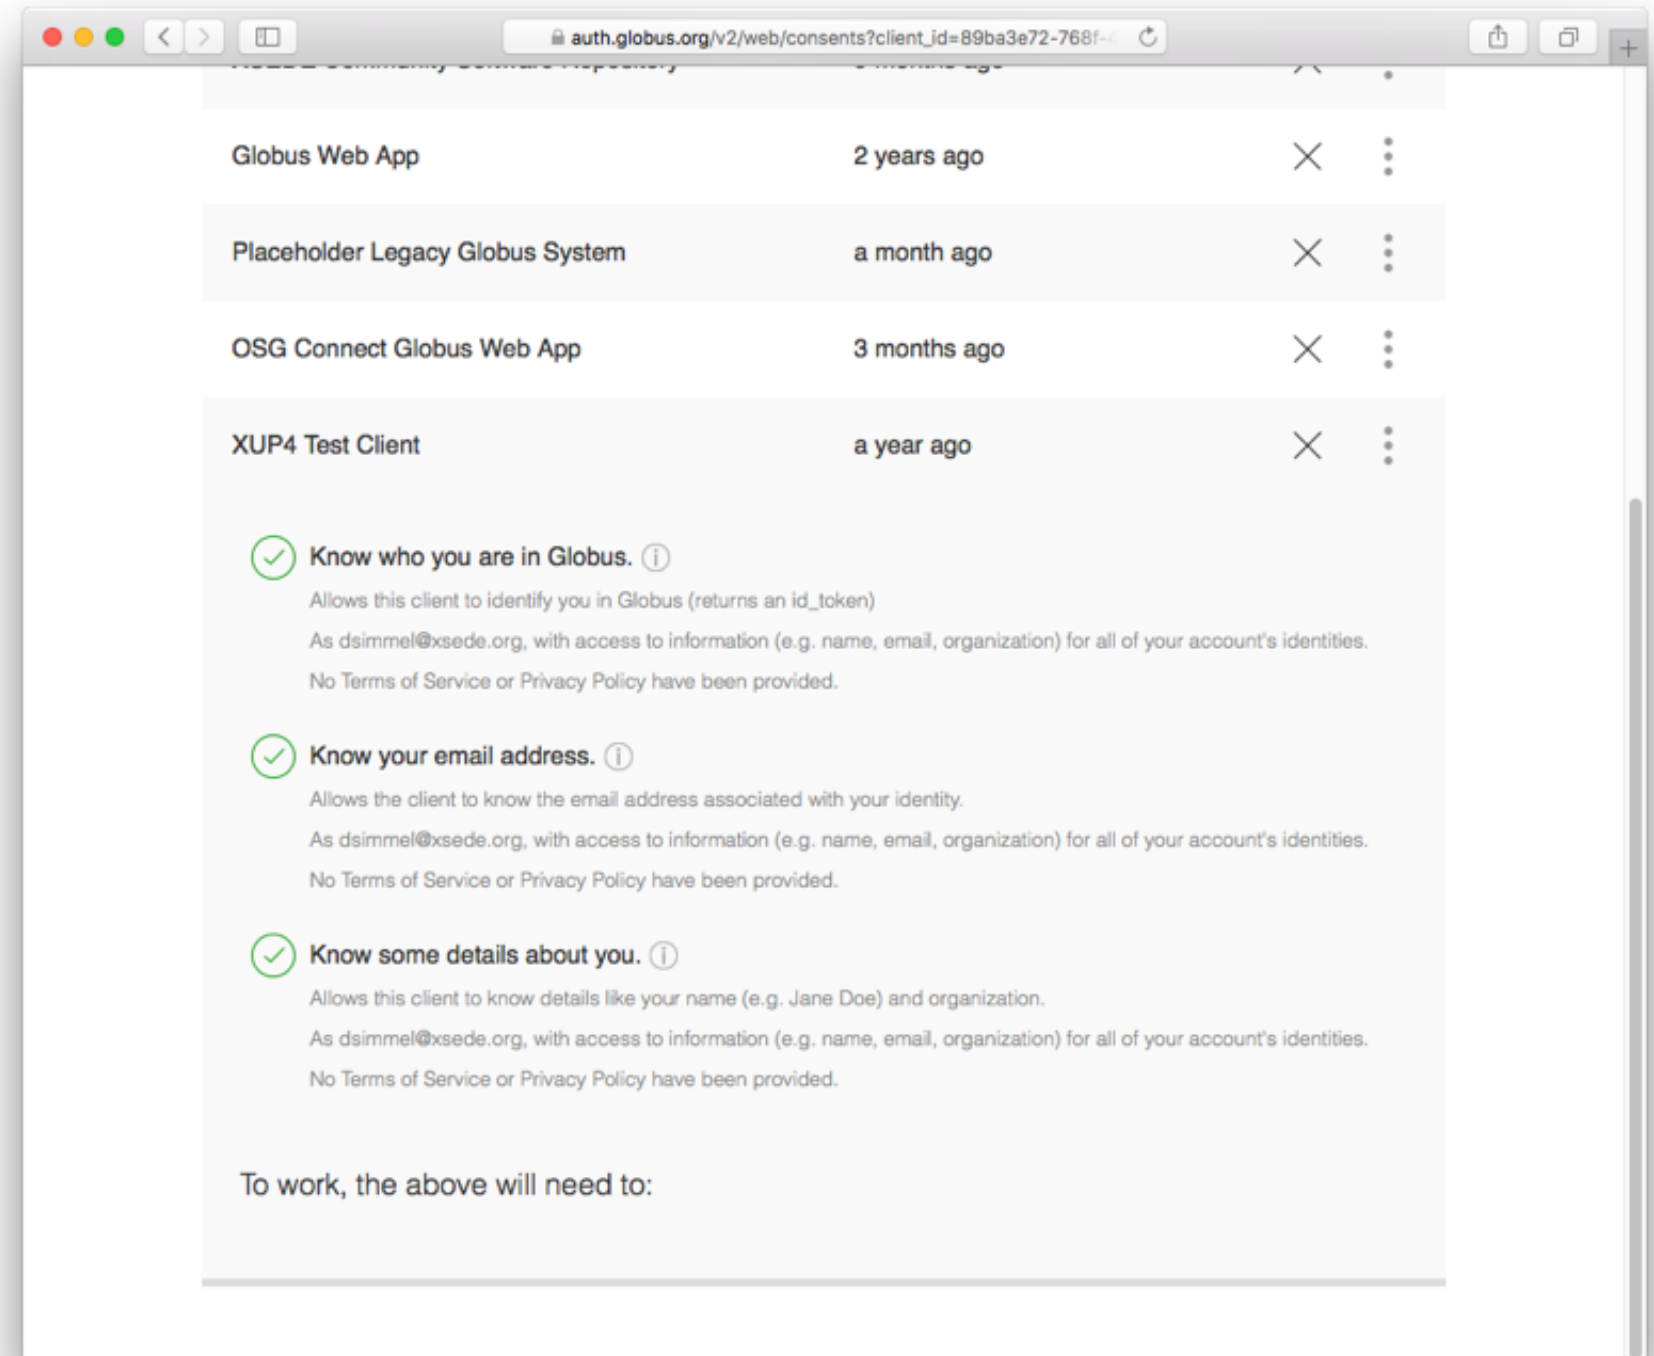
- ✓ View your identities on Globus Auth ⓘ
- ✓ Transfer files using Globus Transfer ⓘ
- ✓ Manage your Globus Groups ⓘ

To work, the above will need to:

- ✓ View your identities on Globus Auth ⓘ
- ✓ Manage your Globus Groups ⓘ

GlobusID Authorization “Consents” info...

- Users are expected to manage their “consents” for each service that they want to authenticate to
- How do we advise users?
- How will this scale?



Globus Auth Framework

- Globus Auth is an authentication and authorization system built upon OpenID Connect and OAuth protocols
- Globus Connect Services (v.5 and later) will rely on Globus Auth, replacing the GSI PKI model for authentication

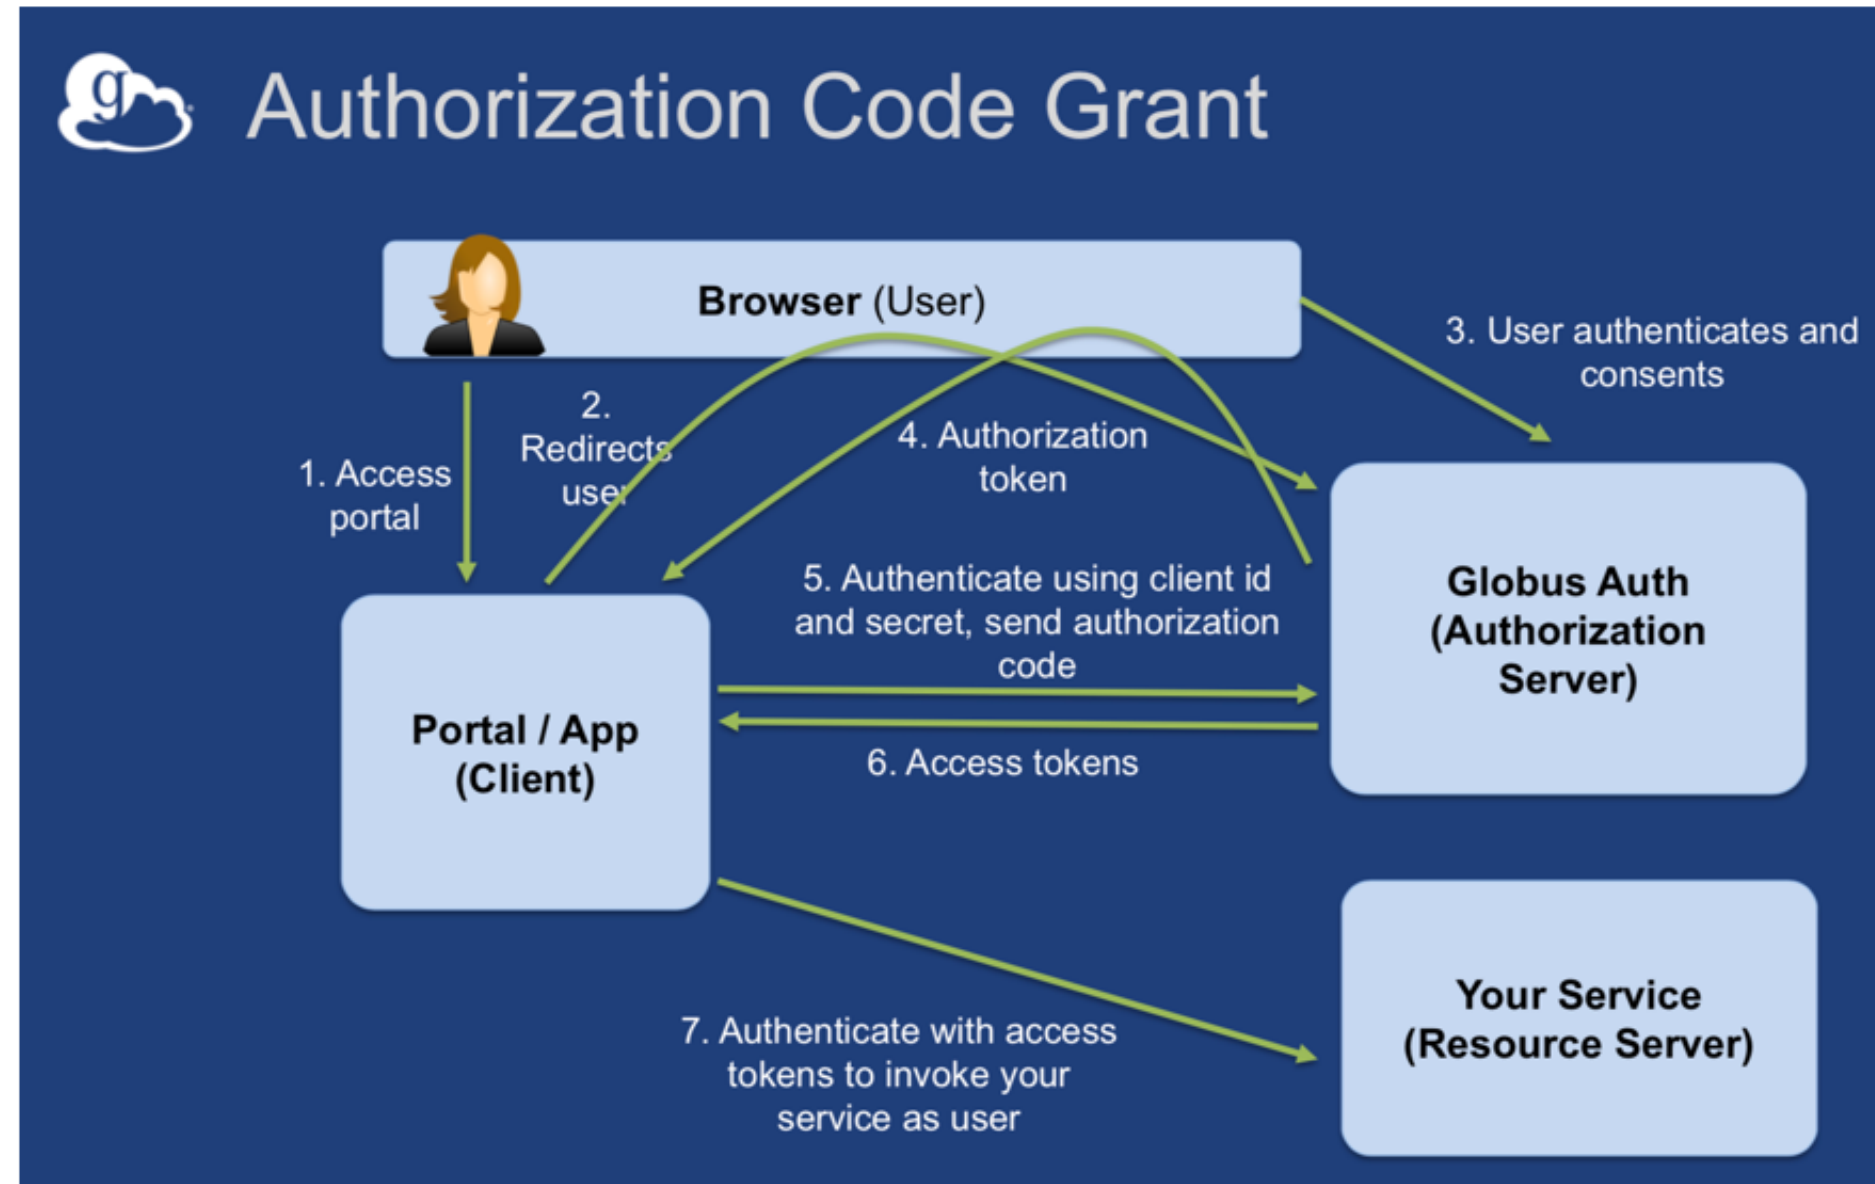


Image credit: Steve Tuecke, "Building the Services Ecosystem" GlobusWorld 2018

https://www.globusworld.org/files/2018/tut08_Auth_Services_Ecosystem.pdf

Observations and Challenges

- Globus Auth-based services remain under development and do not yet support all the functionality currently required and in place using the Globus Toolkit-based services and GSI
- Globus Toolkit has been retired from further development by Globus
 - Many large projects are supporting the “Grid Community Toolkit”
 - WLCG, XSEDE, Open Science Grid
- A PAM-based Globus Auth authentication method is in the works...
 - ...but we haven’t seen it yet
 - Clients (e.g. SSH) will need to be wrapped to manage identity selection
 - How will Service Providers enforce which identities to allow?

Globus Authentication in Practice

- Questions?

References (1)

- Globus
 - <https://www.globus.org/>
- GlobusWorld 2018 presentation slides (April)
 - <https://www.globusworld.org/conf/program>
- Globus Auth API
 - <https://docs.globus.org/api/auth>
- Interoperable Global Trust Federation (IGTF)
 - <https://www.igtf.net>
- InCommon Certificate Service
 - <https://www.incommon.org/certificates/>
- CILogon
 - <http://www.cilogon.org>

References (2)

- PEARC17 paper available in ACM Digital Library at:
 - <https://doi.org/10.1145/3093338.3093392>
- PEARC17 paper and pam_duo patches available at:
 - https://github.com/pscedu/duo_unix_psc
- Instructions for applying the patches to the duo_unix source
 - https://github.com/pscedu/duo_unix_psc/wiki
- Linux Pluggable Authentication Modules (Linux-PAM)
 - <http://www.linux-pam.org>
 - Note that the Offline documentation is more current than the online HTML editions; download the document tar.gz archive that matches your release.
- *pamtester* utility
 - <http://pamtester.sourceforge.net>
 - Available for installation using yum in EPEL repo

References (3)

- OpenSSH
 - <http://www.openssh.com>
- GSI-Enabled OpenSSH (GSI-OpenSSH)
 - <http://grid.ncsa.illinois.edu/ssh/>
- PSC High-Performance Networking patches for OpenSSH (HPN-SSH)
 - <https://github.com/rapier1/openssh-portable>
- MyProxy Credential Management Service
 - <http://grid.ncsa.illinois.edu/myproxy/>
- XSEDE MyProxy-Enabled GSISSH (part of the xsede-user-tfa-ssh package)
 - <https://software.xsede.org/development/xsede-user-tfa-ssh/>
- Duo Unix – Two-Factor Authentication for SSH with PAM support
 - <https://duo.com/docs/duounix>

PEARC₁₈



Practice and Experience
in Advanced Research
Computing

July 22-26, 2018
Pittsburgh PA

Wyndham Grand Pittsburgh Downtown

