# Supporting Science at Scale: CILogon Hosted COmanage for ACCESS

Jim Basney
Scott Koranda
Benn Oshrin

# What is COmanage?

- "COmanage" is the *project* name
- Two *products* within the project: Registry and Match
  - Registry v4.4.0
  - Registry PE v5.0.0 (early adopter release)
  - Match v1.2.1
    - An identity de-duplication tool

# What is COmanage Registry?

- An entity Registry (People, Groups, Services, Servers, etc)
- A lifecycle management tool
  - Enrollment / Onboarding
    - Management of identifiers, roles, groups, attributes, etc
  - Expiration / Offboarding
- Can be deployed to support a number of use cases
  - Enterprise Registry
  - Virtual Organization (VO) Registry
  - Guest / Affiliate Management

# Virtual Organization Reference Architecture

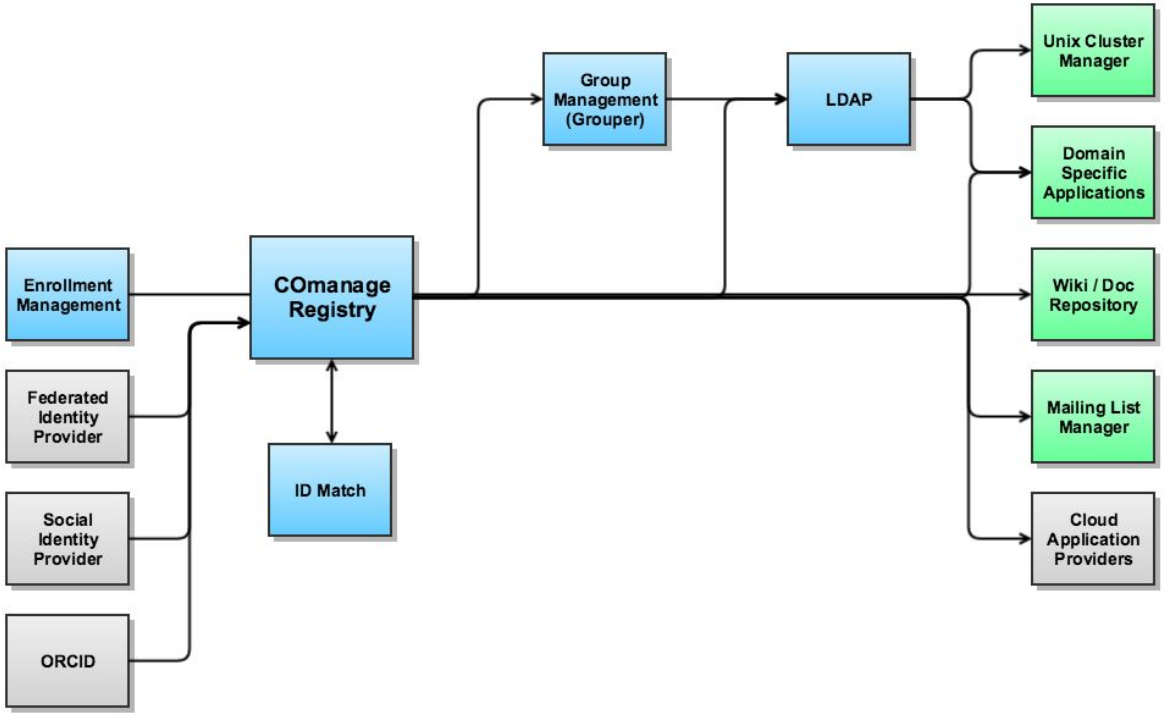

# Extending Registry Capabilities Through Plugins

- Out of box capabilities are typically designed around the 80/20 rule
- Support for additional business logic is available through Plugins
  - Plugins can be shipped but disabled, available from the community, or locally written
  - Multiple types of plugins supported
    - Authenticator, Cluster, Dashboard Widget, Data Filter, Enrollment Flow
    - Identifier Assignment, Identifier Validation, Invitation Confirmer, Job
    - LDAP Schema, Normalization, Organization Source
    - Organizational Identity Source, Provisioner, Vetter

# IAM for Research Collaborations

CILogon: 10+ year sustained effort to enable secure logon to scientific cyberinfrastructure (CI)

for seamless identity and access management (IAM)

using federated identities (SAML, OIDC, OAuth, JWT, X.509, LDAP, SSH, etc.) so researchers log on with their existing credentials from their home organization

supporting 20,000+ active users from 500+ organizations around the world

with onboarding/offboarding/attributes/groups/roles managed consistently across multiple applications

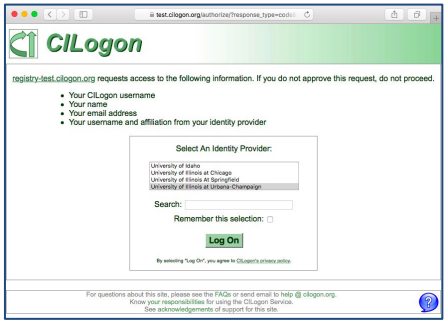


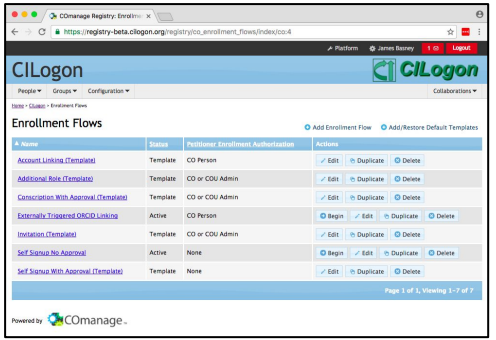

supporting access to science applications on HPC clusters, in Jupyter notebooks, using Globus, via REST APIs, and many other interfaces

using existing identity providers from the researcher's home organization (SAML/ADFS) or external sources (Google, GitHub, Microsoft, ORCID)

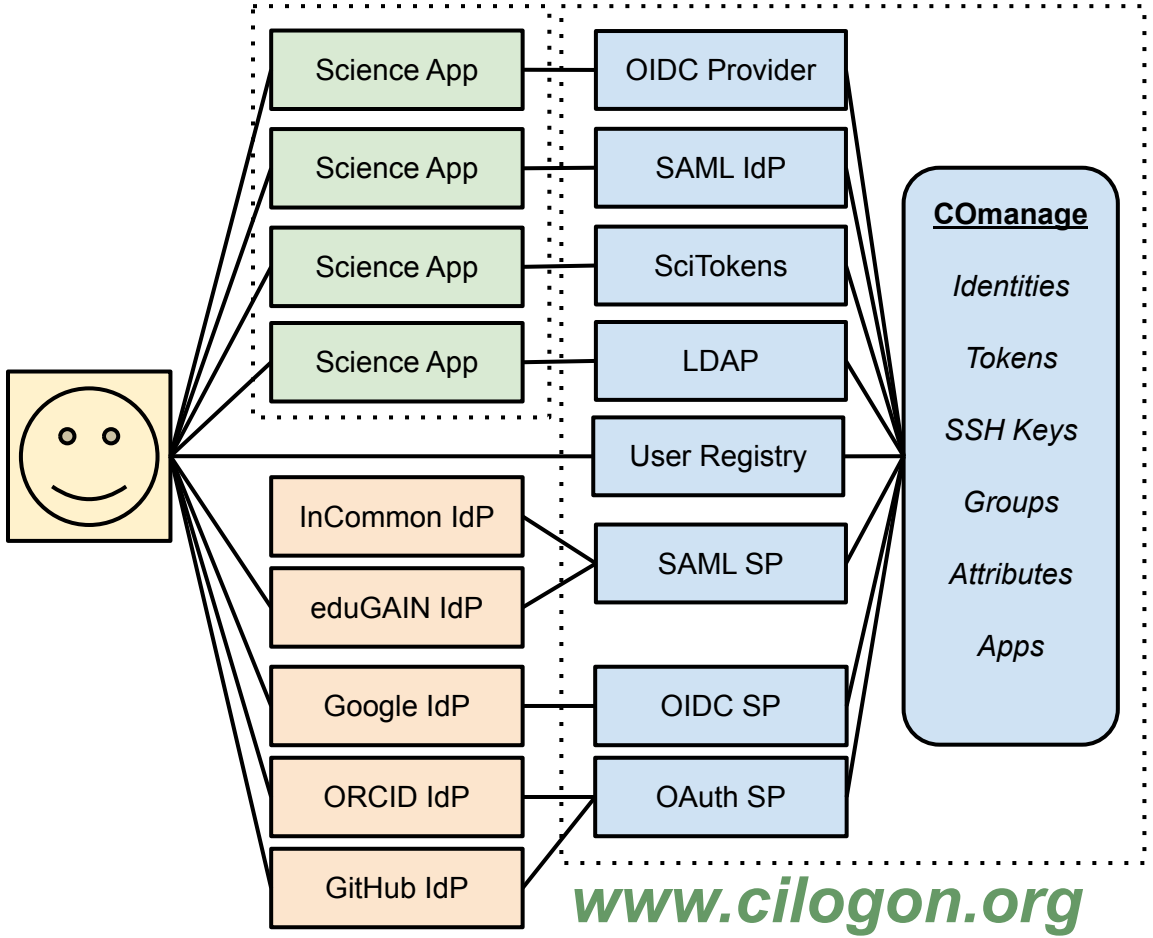

# What is ACCESS?

The **Advanced Cyberinfrastructure Coordination Ecosystem: Services and Support** (ACCESS) is a virtual collaboration funded by the National Science Foundation that facilitates free, customized access to advanced digital resources, consulting, training, and mentorship. ACCESS helps the nation's most creative minds discover breakthroughs and solutions for some of the world's greatest scientific challenges.

ACCESS's virtual cyberinfrastructure allows scientists to interactively share computing resources, data, and expertise. ACCESS resources may be broadly categorized as follows: High-Performance Computing, High Throughput Computing, Visualization, Storage, and Data Services. Many resources provide overlapping functionality across categories.

The NSF’s ACCESS program builds upon the successes of the 11-year XSEDE project, while also expanding the ecosystem with capabilities for new modes of research and further democratizing participation.

# NSF ACCESS Program Overview

# 2022 Transition from XSEDE

- May 1: ACCESS Project Start
- June 24: ACCESS IAM WG Kick-off Meeting
- July 1: identity.access-ci.org created
- August 1: idp.access-ci.org operational
- August 19: registry.access-ci.org operational
- September 1: xsede.org redirects to access-ci.org

# 2022 Transition from XSEDE

- 100,000 registered users / 5,000 active users
- 30+ years of usage history
- Kerberos passwords
- Duo phone registrations

# Standardize on ORCID iDs?

# idp.access-ci.org

- Clone of idp.xsede.org
- s/xsede.org/access-ci.org/

- Your ACCESS ID is the same as your XSEDE Portal account. Please do not create a new ACCESS ID.
- You do not need to change your password or your Duo registration for ACCESS.
- Select the “ACCESS CI (XSEDE)” identity provider to log on using your XSEDE account.

# Three Ways to Access ACCESS

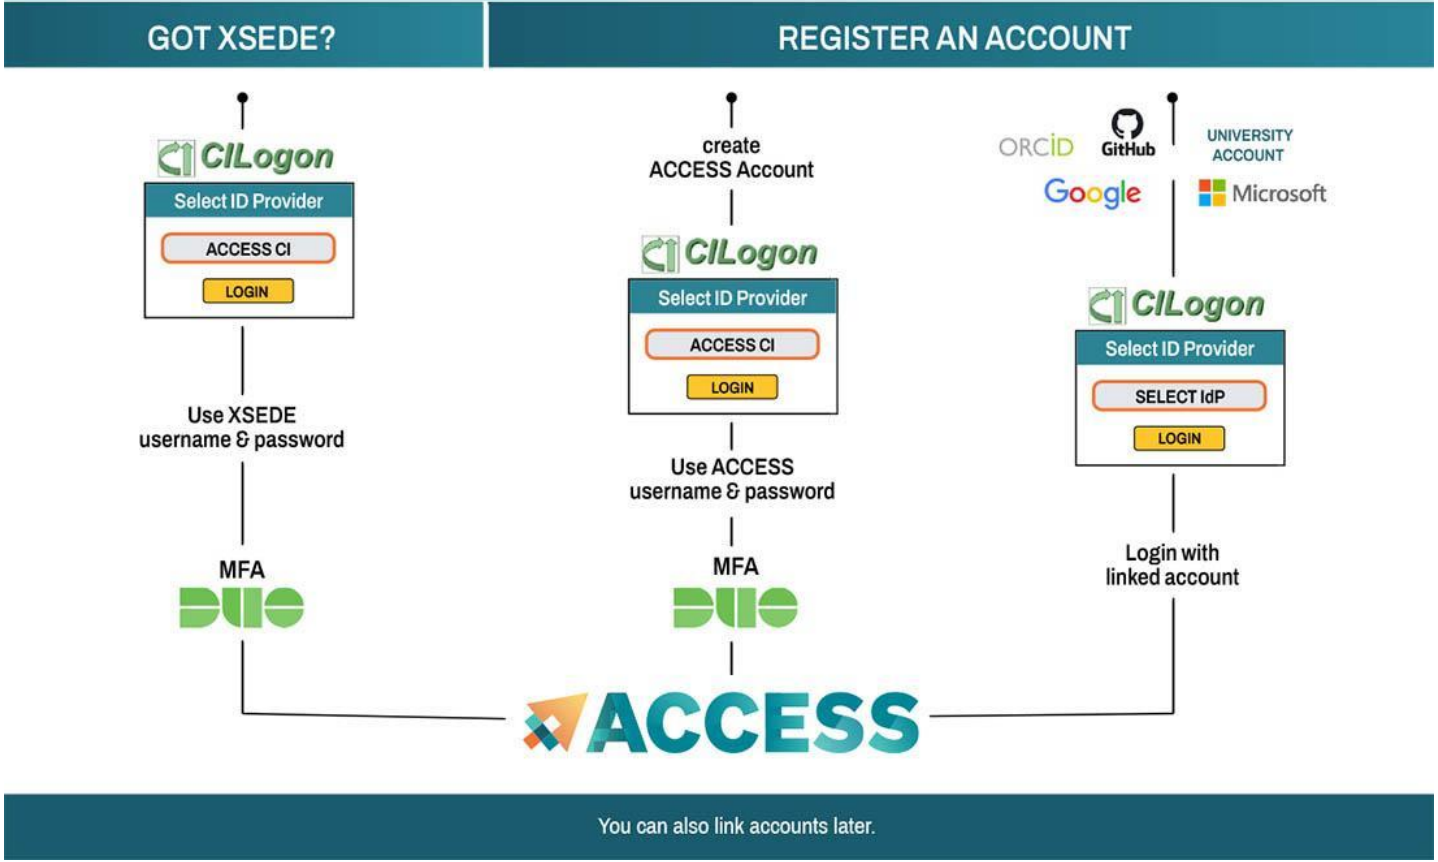

# Transition from XSEDE

xsede.org / portal.xsede.org

↓

access-ci.org / allocations.access-ci.org / support.access-ci.org /
operations.access-ci.org / metrics.access-ci.org / registry.access-ci.org

# Technical Requirements, Challenges, and Solutions

# Requirement: XSEDE Central Database

- XSEDE Central DB ⇒ ACCESS Central DB
  - Custom HTTP API
- Continues as the system of record (SOR)
  - Registry enrollment flows for onboarding new users
  - Some limited workflows allow central DB profile to appear first
  - Duplicate merging and expiration can happen at central DB
- ~ 100K user profiles (initially)

# Solution: AccessDbsyncJob Plugin

- Query central DB HTTP endpoint for profiles
- Synchronize with CO Person records
  - Name (Primary)
  - EmailAddress (Official, Verified)
  - Identifier (AccessID)
  - CO Person Role (Affiliate)
  - Org Identity
    - ACCESS ID as login Identifier
- Used for initial population of 100K CO Person records
- Pull with push handled via Provisioner plugin
- Runs nightly via cron job

```
public $cmPluginType = "job"
```

# Requirement: Federated Identity Support

- New for ACCESS (not available for XSEDE or TeraGrid)
- InCommon (eduGAIN)
- Social identities acceptable
  - ORCiD
  - Google
  - GitHub
  - Microsoft
- Account/Identity linking

# Solution: CILogon OIDC SSO and Registry EnvSource Enrollments

- SSO across access-ci.org sites
- Registry EnvSource consumes login identity to create Org ID
  - Pipeline creates CO Person record from Org ID
  - ACCESS ID automatically generated
- Registry linking flow maps linked identities to ACCESS ID
- Remember selected IdP
  - “ACCESS CI (XSEDE)” is the default
- Verify registration
- Enforce denylist

ACCESS

## Consent to Attribute Release

**ACCESS Registry** requests access to the following information. If you do not approve this request, do not proceed.

- Your CILogon user identifier
- Your name
- Your email address
- Your username and affiliation from your identity provider

**Select an Identity Provider**

ACCESS CI (XSEDE)

**Remember this selection**

LOG ON

By selecting "Log On", you agree to the **privacy policy**.

# Requirement: ACCESS Specific Requirements on Federated Identity Onboarding

- No duplicate enrollments
- Collect ACCESS Organization from curated list
- Query user and collect and provision ACCESS ID password
- Automatically link ACCESS IdP account

# Solution: Access01Enroller Plugin

- Detect ACCESS IdP at start step and short circuit enrollment
- Detect known EmailAddress and short circuit enrollment
- Display form to collect ACCESS Organization (type ahead search)
  - Will move to Organization Source (Registry 4.4.0)
- Prompt to collect ACCESS ID password
  - User may decline
  - May set password later also

```
public $cmPluginType = "enroller";
```

# Requirement: Anonymous Registration

- Federated identity is not required to onboard
- No duplicate enrollments

# Solution: Registry Anonymous Enrollment Flow

- Form to collect name and email
  - Access02Enroller plugin detects known EmailAddress and stops flow
- ACCESS ID automatically generated
  - Reminder email sent to user
- Collect ACCESS ID password
  - Kerberos KdcProvisioner plugin ( `public $cmPluginType = "provisioner";` )
  - Kerberos KrbAuthenticator plugin ( `public $cmPluginType = "authenticator";` )
  - Modeled on “out of the box” PasswordAuthenticator plugin

ALLOCATIONS RESOURCES EVENTS & TRAININGS SUPPORT NEWS ABOUT Find info for you Login


If you had an XSEDE account, please enter your XSEDE username and password for ACCESS login.

**ACCESS ID**

korandas

**ACCESS Password**

••••••••••••••••••

**LOGIN**

**Register for an ACCESS ID**

**Forgot your password?**

**Need Help?**

# Solution: AccessdbProvisioner Plugin

- New CO Person records pushed immediately to ACCESS Central DB
- Any CO Person changes pushed to ACCESS Central DB
- No deletion

```
public $cmPluginType = "provisioner";
```

# Challenge: Email Verification Failures

- Email containing URL often identified as SPAM
- Enterprise mail services spoil the nonce (Azure!)
- Users switch browsers/devices to click URL
- Single most cause of Helpdesk tickets

# Solution: EmailVerificationEnroller Plugin

- Display form to collect a code
- Send code to registered email address
- Much less SPAM triggering
- Users do not leave browser/session
- Designed by ACCESS staff
- Developed by SCG

```
public $cmPluginType = "enroller";
```

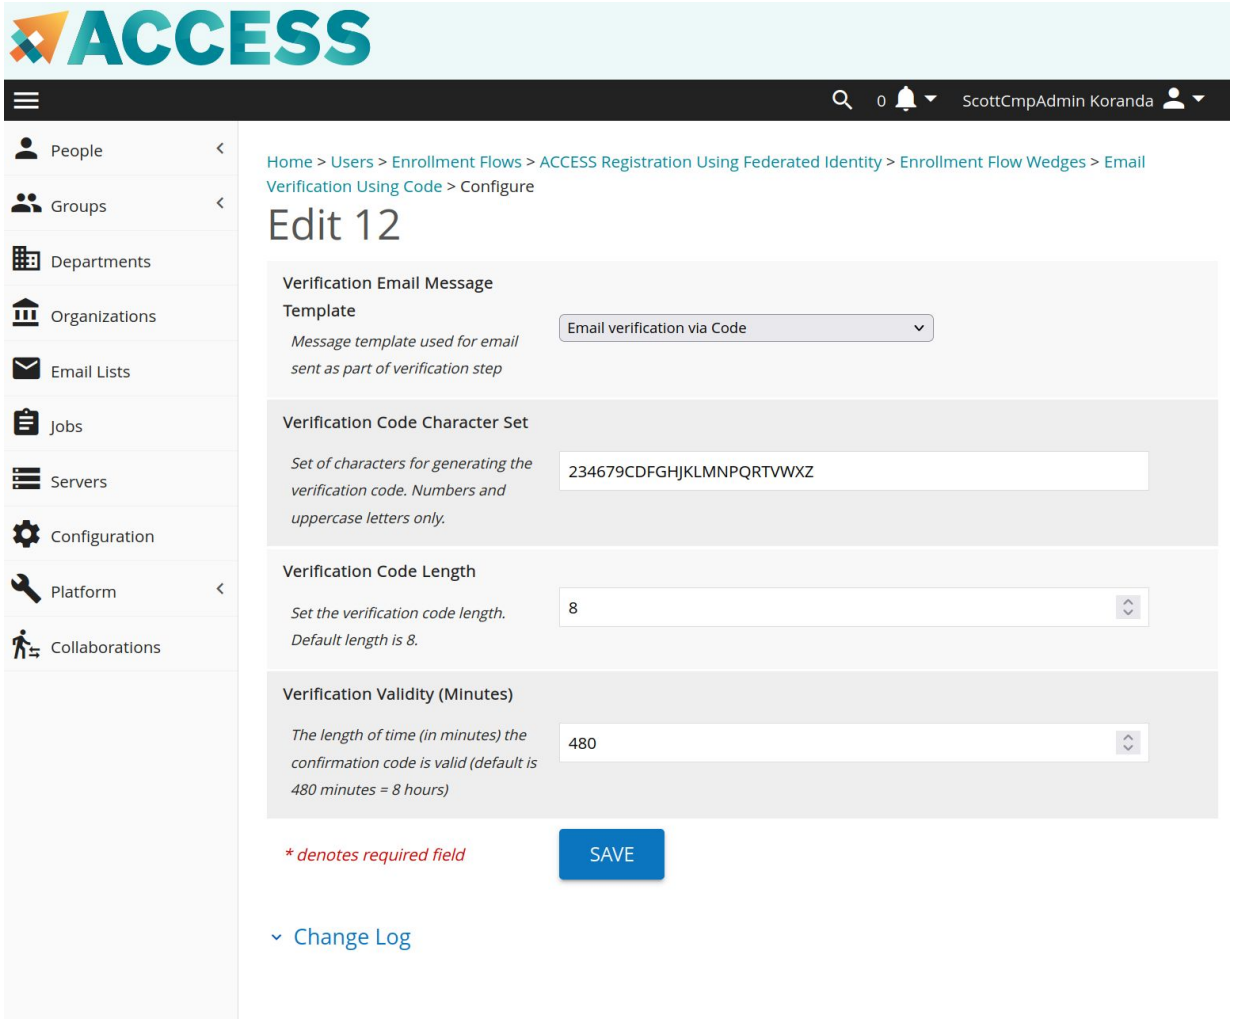
ACCESS
0
ScottCmpAdmin Koranda
People
Groups
Departments
Organizations
Email Lists
Jobs
Servers
Configuration
Platform
Collaborations
Home > Users > Enrollment Flows > ACCESS Registration Using Federated Identity > Enrollment Flow Wedges > Email Verification Using Code > Configure
Edit 12
Verification Email Message Template
Message template used for email sent as part of verification step
Email verification via Code
Verification Code Character Set
Set of characters for generating the verification code. Numbers and uppercase letters only.
234679CDFGHJKLMNPQRTVWXZ
Verification Code Length
Set the verification code length. Default length is 8.
8
Verification Validity (Minutes)
The length of time (in minutes) the confirmation code is valid (default is 480 minutes = 8 hours)
480
* denotes required field
SAVE
Change Log

ACCESS

0

# Enter Code to Verify Your Email Address

An email was sent to scott.koranda@ligo.org containing an alphanumeric code. If you do not receive the email in your Inbox, check your Spam/Junk folder.

DO NOT CLOSE YOUR BROWSER OR NAVIGATE AWAY FROM THIS PAGE. If you have problems, please Open a Help Ticket.

Enter Code:

SUBMIT

# Challenge: LDAP Server Cost

- CILogon OAuth2 server needs to resolve OIDC claim values
- Begin by leveraging Registry LDAP Provisioner
- 130K records
- 24 x 7 x 365 HA
- CILogon operates HA OpenLDAP for first two years
  - Doable, but at what cost?
  - Expertise on ports 389/636 is harder and harder to find

# Solution: AWS DynamoProvisioner Plugin

- AWS DynamoDB serverless, NoSQL, fully managed database
- Single-digit millisecond performance at any scale
- PHP (Registry) and Java (OAuth2 server) client libraries
- October 2024 cost for ACCESS ~ $8 US

```
public $cmPluginType = "provisioner";
```

# Challenge: Relying Party Integration

- Easy and fast integration of OIDC RPs
  - Computing service providers
  - Project infrastructure (e.g. allocations)
- Leverage the CILogon OAuth2 server (OA4MP) API
- 52 RPs as of November, 2024

# Solution: Oa4mpClient Plugin

- Self-service OIDC RP creation/management
- Delegated users (CO Group)
- Standard OIDC configuration parameters
  - Scopes, callback URI, refresh token, …
- Custom mapping of Registry models to claims
  - If it’s part of a CO Person or Org ID record you can map it

```
public $cmPluginType = "other";
```

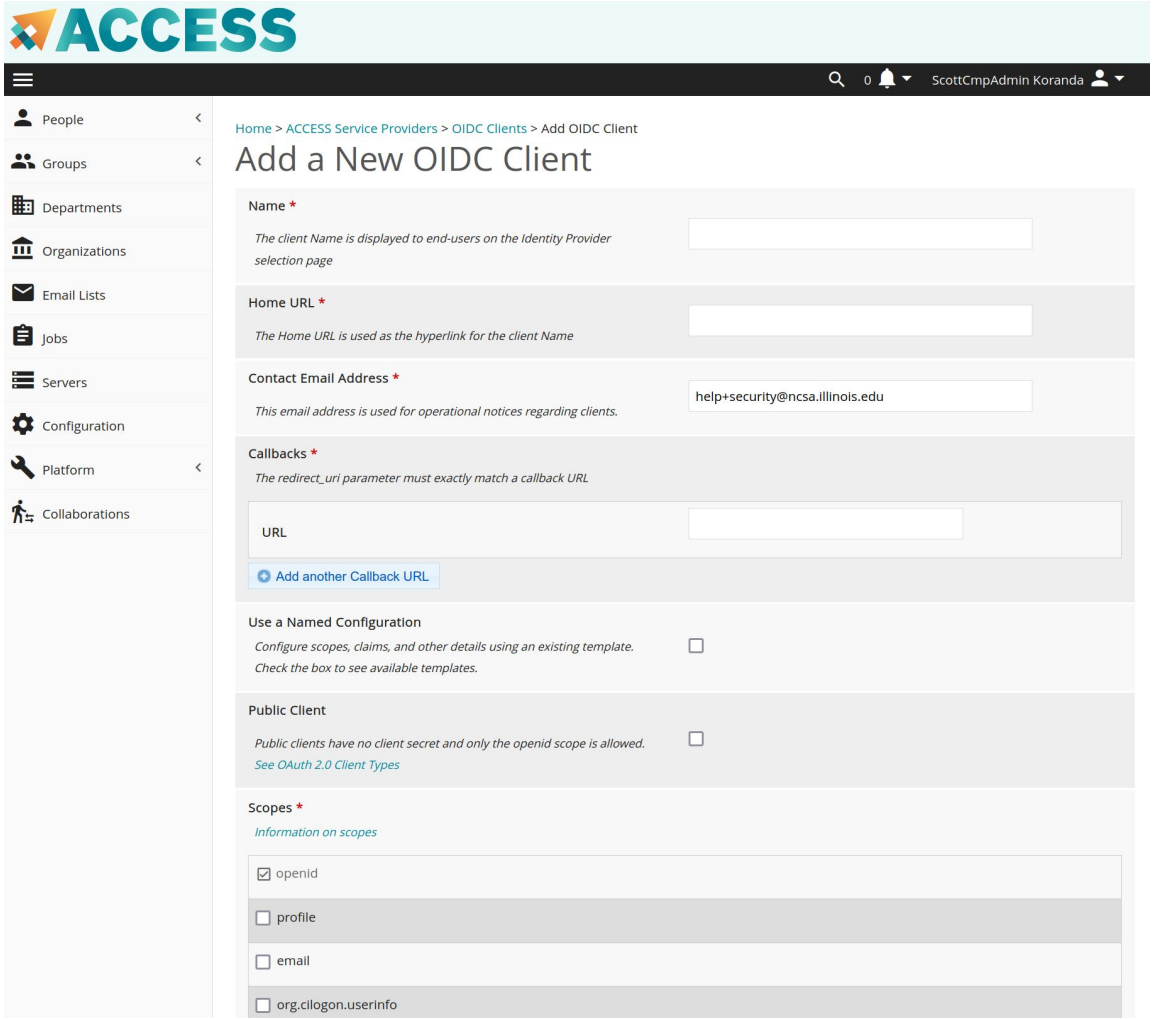
ACCESS
0
ScottCmpAdmin Koranda
People
Groups
Departments
Organizations
Email Lists
Jobs
Servers
Configuration
Platform
Collaborations
Home > ACCESS Service Providers > OIDC Clients > Add OIDC Client
Add a New OIDC Client
Name *
The client Name is displayed to end-users on the Identity Provider selection page
Home URL *
The Home URL is used as the hyperlink for the client Name
Contact Email Address *
help+security@ncsa.illinois.edu
This email address is used for operational notices regarding clients.
Callbacks *
The redirect_uri parameter must exactly match a callback URL
URL
Add another Callback URL
Use a Named Configuration
Configure scopes, claims, and other details using an existing template. Check the box to see available templates.
Public Client
Public clients have no client secret and only the openid scope is allowed.
See OAuth 2.0 Client Types
Scopes *
Information on scopes
openid
profile
email
org.cilogon.userinfo

# So Many Custom Registry Plugins!?!

- Total of 14 custom plugins in production
- Allowed CILogon team to meet initial requirements in 6 weeks (!)
- Allows CILogon team to continue to meet new requirements quickly
- As comparison 12 CILogon subscribers with no custom plugins

# What’s next?

- More SSO integrations (OIDC & SAML)
- Groups/Roles
- Delegated COmanage administration
- Identity unlinking
- SciTokens

# Thank You to Our Collaborators

- Shayna Atkinson (SCG)
- Robin Blair (CILogon)
- Christopher Clausen (NCSA)
- Terry Fleury (ACCESS/CILogon)
- Jeff Gaynor (CILogon)
- Ioannis Igoumenos (SCG)
- Arlen Johnson (SCG)
- Rob Light (ACCESS/PSC)
- Laura Paglione (SCG)
- Derek Simmel (ACCESS/PSC)
- Scott Sakai (ACCESS/SDSC)
- Michael Shapiro (NCSA)
- Nathan Tolbert (ACCESS/NCSA)
- Tom Zeller (SCG)
- Yan Zhan (CILogon)

Thanks!

Contact:
benno@sphericalcowgroup.com
jbasney@illinois.edu
skoranda@illinois.edu