PELICAN PLATFORM

Sharing Data in Pelican

Authentication, User, and Group Management

Brian Aydemir

Morgridge Institute for Research

HTC 26

Getting or putting objects at an origin

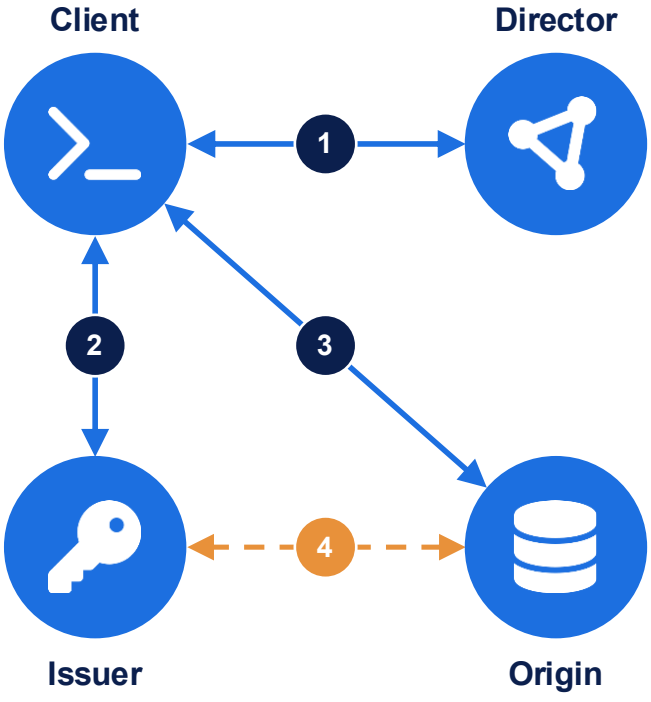

1 **Discover**

The client asks the Director for the object, and is sent to the Origin and an issuer (if the operation requires an access token).

2 **Authorize**

If a token is needed, the client gets a signed one from that issuer.

3 **Transfer**

The client downloads or uploads the object directly with the Origin, presenting its token when required.

4 **Verify**

The Origin checks any token against the issuer's keys, then serves or stores the object.

Three ways people share data via OSDF today

Public

Open data, no token needed.

/ncar-gdex
/routeviews

Proprietary

Internal datasets, centrally-managed token service.

/icecube/wipac
/igwn/virgo

Personal

Data staged for jobs by a user, AP manages tokens.

/chtc/staging/baydemir

Who implements the policy? Only a sysadmin

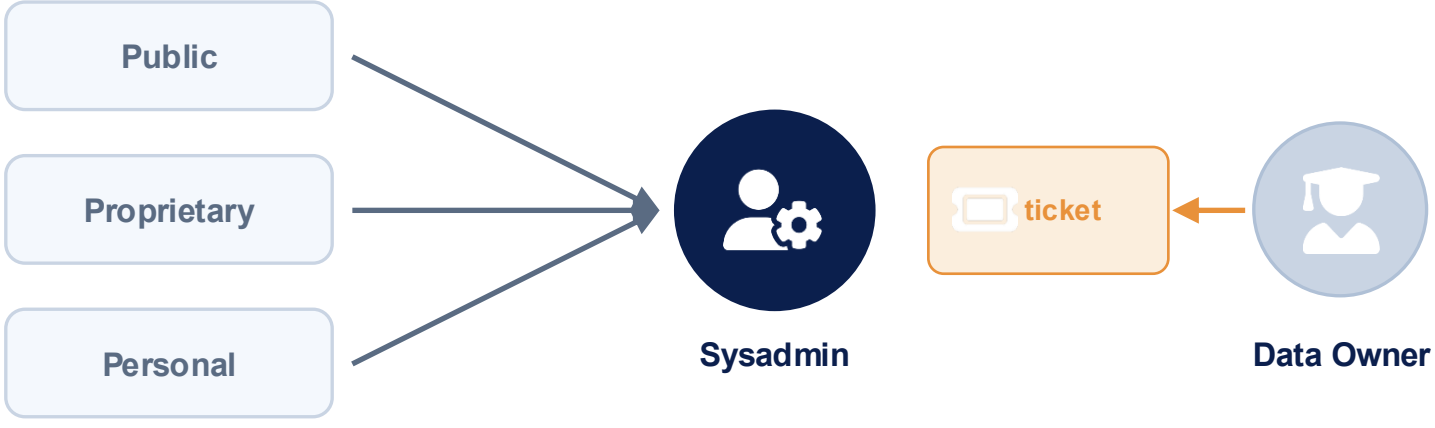

The sysadmin is the bottleneck for the data owner who determines the policy.

The future: Delegating authority

Authority should flow through people, with each person acting independently of the one before.

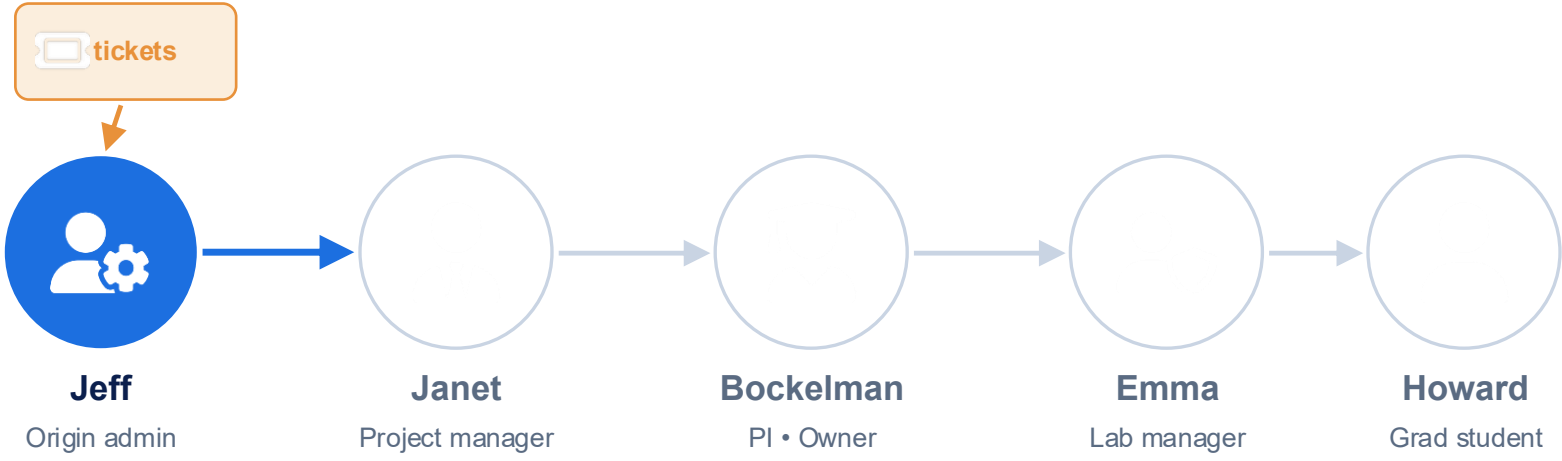

What if the origin had a notion of someone other than “sysadmin”?

Built-in issuer: Every Pelican origin can mint its own access tokens.

**This feature is under active development.*

Collection: Slice of a namespace

A collection binds a namespace prefix to an owner and a policy.

Bockelman Lab
/wi-ren/bockelman-lab

Janet creates the collection and binds it to a prefix, but its owner sets the policy.

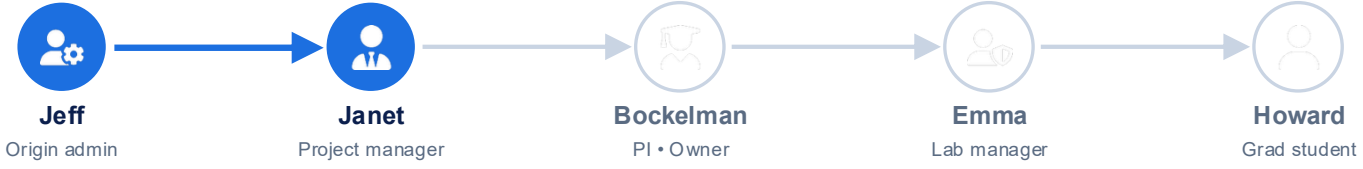

**This feature is under active development.*

Owner: A collection belongs to a person

The collection's owner has (almost) full control over the collection.

Person, or user: The origin records a unique internal ID, a human-readable display name, and one or more OIDC-based identities.

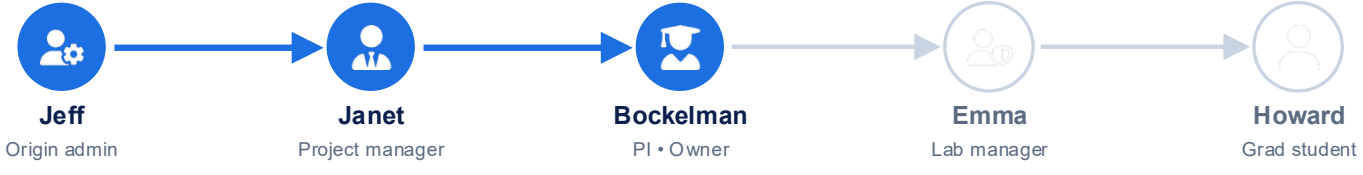

**This feature is under active development.*

Invite links: Self-service onboarding

A one-time link that mediates a hand-off, without an admin.

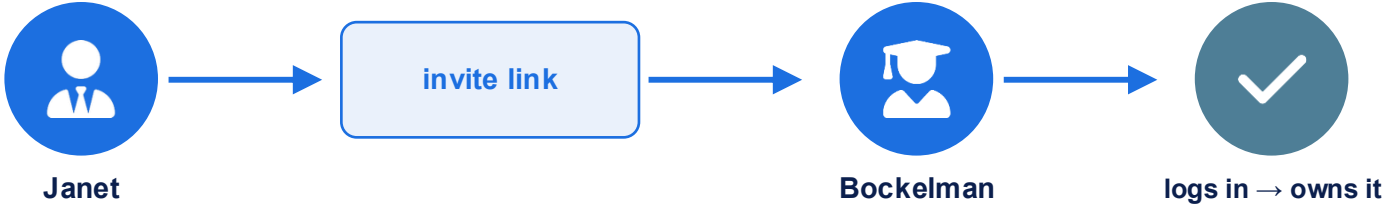

Redeeming a link: Log in, (origin automatically creates a user), sign Acceptable Use Policy, accept or reject the access granted by the link.

**This feature is under active development.*

Groups: Granting access

A group is a set of users, with an owner and an admin (another user or group!).

Collection: bockelman-lab

Access control lists (policy)

GROUP	ROLE
bockelman-lab-readers	read
bockelman-lab-writers	write
bockelman-lab-admins	admin

Group: bockelman-lab-admins

Emma
invited to admins group

Emma can invite people
to the readers and writers groups.

But she cannot

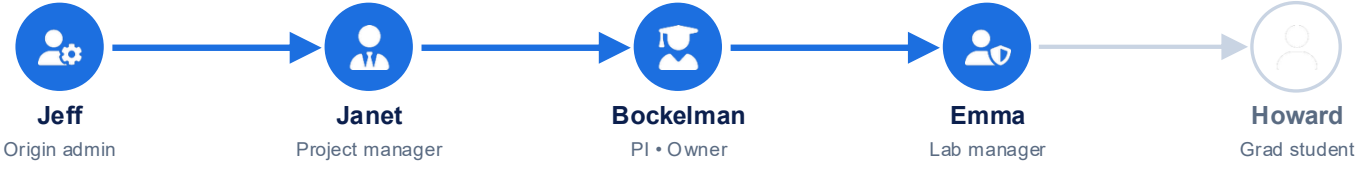

**This feature is under active development.*

What membership buys: read / write

Joining a group grants you access associated with the group's role scoped to the collection's objects.

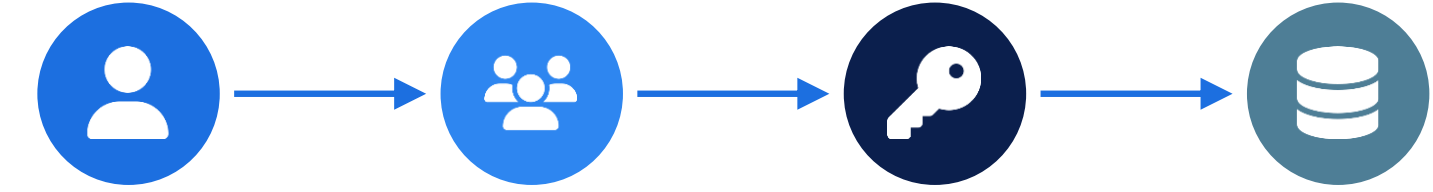

readers → download. **writers** → download / upload.

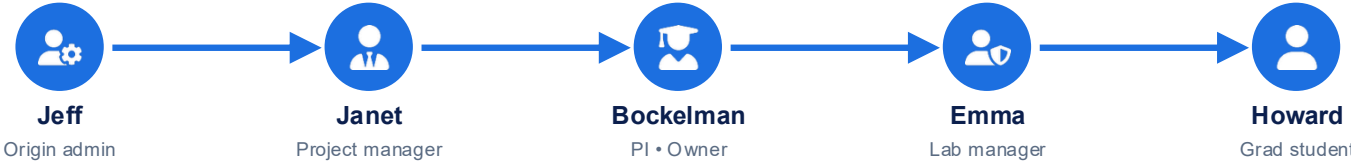

**This feature is under active development.*

Sharing: Hand off part of your access

Create a sub-prefix of a collection, with its own ACLs, but always restricted to your own access.

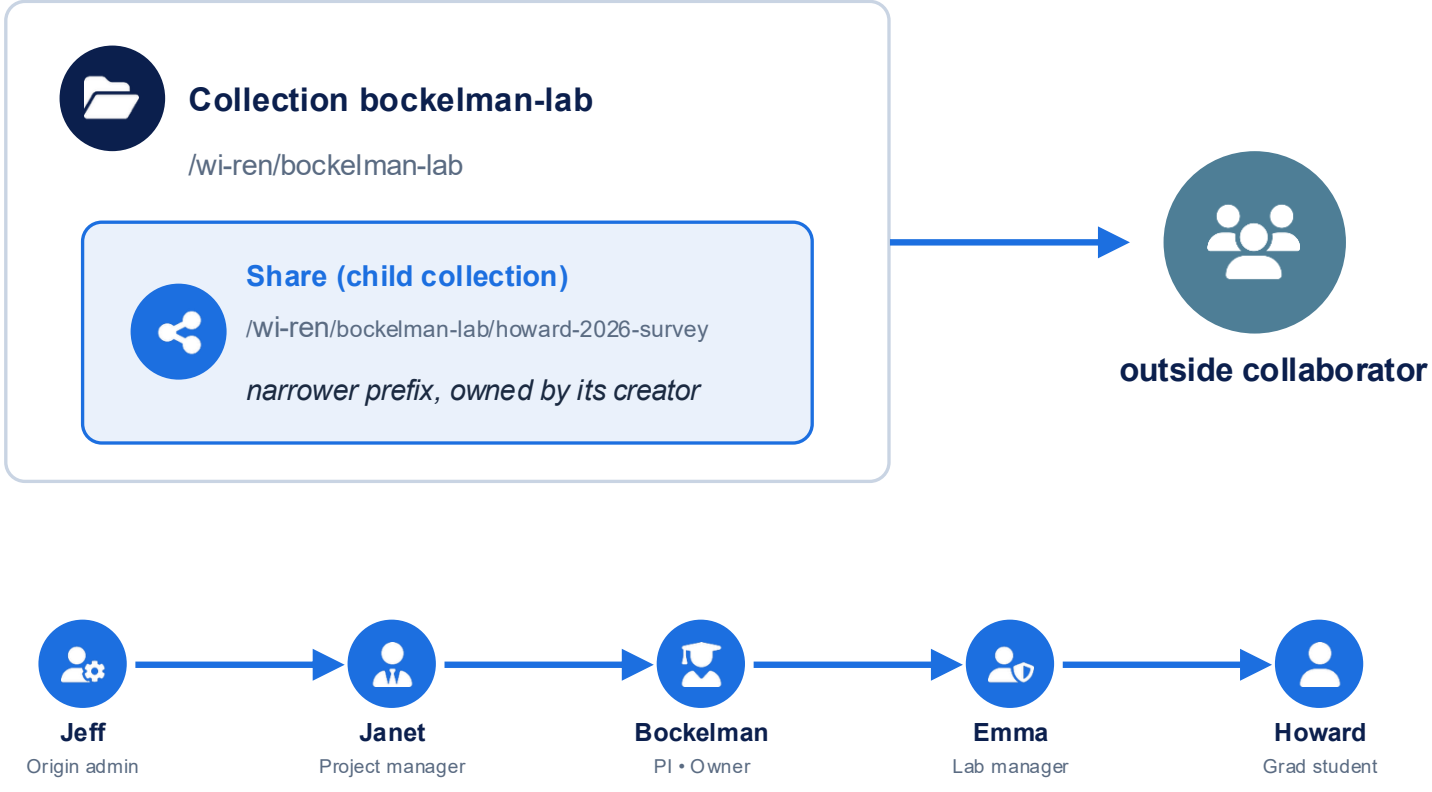

**This feature is under active development.*

Acknowledgements

National Science Foundation

This material is based upon work supported by the National Science Foundation under Cooperative Agreements OAC-2209645. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.