

# Scientific OpenStack & Convergent Federation

Federate all the (OpenStack) things

*Stig Telfer & John Garbutt,  
OpenStack Scientific Working Group*

# Machine-to-Machine interfaces matter

1. Users will want to use the CLI, Ansible, or other automation for anything beyond trivial usage examples
2. *Does this favour OIDC over SAML as authentication technology?*
3. Users may want deferred access to resources
4. *Is there a role for OAuth?*

# Mapping People between Organisations

1. ID Proxies (Indigo IAM, Keycloak) seem like the way to avoid problems when a federation's membership changes.
2. What is best practice for managing the mapping of remote users to local roles?
  - a. Keystone federated user mapping helps with the mechanics
3. What is the best practice for tracking the activities of federated users?
  - a. *Does CADF provide all the linkage we need or is more needed?*

# Mapping Resources Between Clouds

1. How can we make different clouds look more similar?
  - a. Flavors?
  - b. Images?
2. How much does this matter for users?

# Linking Resources Between Clouds

1. If an organisation is distributed, it will probably have resources allocated in many federated clouds
2. *How can those resources be interlinked conveniently and securely?*
3. Is there a discussion around the concept of a “Fog” cloud in this forum?

# It Might be Open and Free but Somebody is Paying

1. Resource utilisation must be accounted for, presumably at project level
2. How can a project know it's resource usage across multiple federated clouds?
3. How can a project PI control costs (financial or otherwise?)
4. How can users leave an organisation in a federated cloud?

# New OpenStack Capabilities Might Help

- Unification of users from different sources in keystone (in Newton)
- Many demos of cross-cloud workloads
- Application credentials may help with machine-machine communication
- Improving Policy/Role management
- os-traits working to standardise qualitative resource requests in flavors
- Tricircle is looking at connecting different network regions together

# Sharing Best Practice

1. How do we learn from past experiences?
2. *How do we avoid reinventing?*
3. How do we share knowledge on best practice?
4. *How do we keep learning?*

# OpenStack Scientific Working Group / SIG

1. A way for people with research computing use cases to collaborate
2. Advocate the research computing use case within OpenStack
3. Advocate OpenStack within the scientific computing community
4. [https://wiki.openstack.org/wiki/Scientific\\_working\\_group](https://wiki.openstack.org/wiki/Scientific_working_group)