WS on Federated Identity for Openstack - Gn4 Jr4 and OSO
Rome, October 17-18, 2016
https://eventr.geant.org/events/2527

Start: *Monday, October 17 at 1 PM*
Close: *Tuesday, October 18 at 4:30 PM*

GARR Offices will be open starting at 9:30 am for Early Birds;
A light welcome lunch will be served at 1 PM upon arrival.

Time slot	Title - Speaker
Day 1 - Monday, October 17	
13:00 - 14:00	*Light welcome lunch*
14:00 - 14:20	**Opening** - GARR
14:20 - 14:50	**The INDIGO Authentication and Authorization Infrastructure** *Andrea Ceccanti / INFN CNAF*
14:50 - 15:20	**KeyRock and Wilma : Identity management based on Openstack in FIWARE architecture**. *Joaquin Salvachua & Alvaro Alonso - Universidad Politécnica de Madrid*
15:20 - 15:50	**The EGI AAI Checkin service** *Mario David - EGI / LIP*
15:50 - 16:00	*Coffee break*
16:00 - 16:30	**3 approaches for federation authorization at CSC** *Kalle Happonen - CSC*
16:30 - 17:00	**eduGAIN integration at CERN** *Jose Castro Leon - CERN*
17:00 - 17:20	**GARR Cloud** **Authentication and Authorization model** *M.Reale, D.Vaghetti - GARR*
17:20 - 17:40	**How user management today works in SWITCHengines.** *Saverio Proto - SWITCH*
17:40 - 18:00	**Openstack WGs how to contribute in the Community** *Saverio Proto - Stig Telfer*

Day 2 - Tuesday, October 18	
9:00 - 11.00	**Pain point session**: *every participant will tell what is not happy about Openstack user management and federation*.
11:00 - 11:15	*Coffee break*
11:15 - 12:45	**Questions and Answers session:** *Ask in this slot questions to others : how do you handle this topic ? Talk to others to share solutions and find your answer*
12:45 - 14	*Light Lunch*
14:00 - 14:30	**Proposal**: write a user story about the NREN use case for the Product WG using the template: https://github.com/openstack/openstack-user-stories/blob/master/user-story-template.rst Etherpad https://etherpad.openstack.org/p/scientific-wg-rome-federated-identity-user-story
14:30 - 15:00	Additional presentation 2
15:30 - 16:00	Wrap up
16:30 - 17:00	Closing

Proposed perspectives for discussions:

1) How comfortable do we feel with the available tools
2) What tools / patches did we develop internally
3) What would be nice to have provided by openstack
4) What are the most annoying issues
5) What Openstack federation means for you? ID federation through keystone, common marketplace for images/templates, or something more? Federated cloud,

Items for a common discussion on Federated Identity and Openstack:

1) Configuration of non-Keystone based AuthN
    a) Direct IDP plugins towards SAML, OIDC, OAuth2
    b) Through IDP/SP proxies

    c) Any other option required by our user communities ?
2) DevOps solutions for managing both sides
3) Generation and revocation of user accounts
    a) Ephemeral users (see also *shadow users* now in Keystone)
4) Mapping user attributes to domains, groups, roles
5) Using openstack domains
6) Attribute Authorities and Group Management
7) Providing admin delegation to cloud managers
8) Security : how to handle it (tools, approach..)
9) Federation of Openstack from different organizations
    a) Who is aiming at doing it ?
    b) Which deployment model

Useful resources:

Slides from this workshop

- https://gbox.garr.it/garrbox/index.php/s/ixZq9HclDtREYc4
-

Links to past edition of the workshop

- Agenda https://wiki.geant.org/display/gn41sa7/Agenda
- Etherpad https://sandstorm.cloud.switch.ch/shared/Hkz58gpVi61In33NPbjwlre5K_Pc-zzN9tvY1O-ZrBd

Scientific Working group

- https://etherpad.openstack.org/p/scientific-wg

Ops tooling repos

- https://wiki.openstack.org/wiki/Osops
- *https://github.com/openstack/osops-tools-contrib*

- Cern keystone
- https://gitlab.cern.ch/cloud-infrastructure/keystone
-
- Cern cloud infrastructure
- https://gitlab.cern.ch/groups/cloud-infrastructure

Ansible Rackspace Keystone to Keystone federation
https://developer.rackspace.com/blog/keystone-to-keystone-federation-with-openstack-ansible/

Questions / Pain Point session

Kalle: Does anyone have a good solution to manage machine/API/automation credentials for federated or non-federated users?

Jose: you can get alternate credentials. Revoke them at any time, these kind of solutions
Davide: what about leveraging OAuth2 - plugin on KeyRock ? can you use OAuth2 plugin with the vanilla keystone ? Yes. (not yet pushed upstream)

Andrea: you provision when you get the OAuth2 token, at the time, or you can provision before. Getting user information somewhere else .
Provisioning interface gets user information - and a notification mechanism - “notify me when a new group is created” … or “do something when user management changes”

Jose; I have access to my projects , but I want alternative credentials to act on my behalf - I am already federated, I have access to the project, but I need cron-jo access

Kalle: if you want auto scaling, you don’t want to put your own credentials to that VM
Andrea: we plan to have for this in INDIGO the ability to create access tokens, a scoped tokens, to some CLIs - at another level, it is not inside keystone.

Jose: OAuth2 plugin in the keystone itself. This is what we are talking about. This is another option w.r.t. INDIGO approach
Jose: If CERN IDP provides something to support OAUth2, but if they do not provide anything I need alternative credentials.

Saverio: anyone doing billing for federation ?
Kalle: we integrate all our federation into our local projects, project lifecycle into CSC way. We use our central billing, everything is also in our own central DB. It is easy in this way. Probably not the common way of doing it…

Jose: at CERN we have representative of external users, a contact person, he gets in touch for security, and also sending the bills to the external users.

Saverio: at SWITCH we have some keywords in the tenant name [mailing lists tags] for example, then we have own software querying API continuously about resource usage from every projects, VMs per minute, object storage, we can then send bills.
MarioR: do you also send proactive consume alarms to customers ?
Saverio: no. so far we never had an issue. If our users did not really understand, commission, space…...we can help. 500 keystone users registers in our keystone, we never had real issues, + 120 students last month - it is responsibility of the people having the course, to warn us when we can delete the VMs and the storage…

Andrea: this is a good use case : change in state of membership of a group, you have some automated procedures to react to this change

Jose: when you have and IDP, and something changes on the IDP side, you are not informed, so how can you trigger any action ?
Saverio: since we started the billing, things improved.
Jose: about the hard coding of the tenant name stuff, you can use extra attributes to use if you want ? Saverio: we thought this would have been not a stable feature...

We have an accounting group, users have several attributes. You do not need to hardcode it in the tenant name.

Saverio: it helps against errors…. It is not elegant, but it will survive every openstack upgrade - not completely nice
Jose: attributes are exposed in the API calls, so if you get the tenants you get them
Save: it is not in the name - It is in the tenant description
Jose: in future people will possibly be able to change tenant name….for any reason….the ID won't change - you map VMs to IDs and not to tenant name

Stig: Do you have a solution different amount for different CPU architectures ?
Mario: flavours inside glance ?

Saverio: We have some flavors can be scheduled only on some CPU architecture...price differentiates depending on flavors

Jose: we manage billing on flavors. .

MarioD: Any benchmarking to the CPUs beforehand ? Saverio: no. in our use case, it's more for the disks, some machines have disks on CEPH backend, some on SSD on the hypervisor...even if a machine has 30 GBs, for example, the flavor tells us the storage backend type…. SSD is more expensive than CEPH…

Kalle: based on performances

If we introduced the nested projects, when I query a VM, can I retrieve the description of the project highest in hierarchy ? I need it for my billing system.
Jose: not immediate. All this nested behavior is on the keystone side.

Save: how many layers of nesting can you have ? **Jose**: as many as you want. This thing is ..nested quotas instead of nested projects - in my case like ATLAS get 30 0000 cores….they will partition them for the child …...each child will repartition their stuff-accounting works on the top level one, the allocated quota for each of the children - you charge the top level one anyhow….you don't really need ot.

Saverio: a tenant has all info to do billing….if it is a nested project, the tenant could be created by the user itself...so I have less visibility on whom to bill…

Jose: force properties inheriting you can do it - export all properties to child project - you can ask this to the keystone guys….add all info on the children

MarioD: you always go in the federated case via Web….? what alternatives do you have ?
Stefan: ECP does not work in a proxy env…… **Davide**: it is tricky .
ECP is not standardize to work in a proxy env - how do you tell the PX where you want to go to ? **Davide:** you can give ICP just one hop. AFAIK: you can give the ECP client 1 hop. It works, I tested.
Stefan: it takes you to the proxy, but then you cannot go any further.
ECP cannot do proxy env, DI4R in Cracow, I asked….RADIUS allows this as RFC, yes...with a bit of magic in freeRADIUS it works….but moonshot clients currently does not support …. User@..... in a proxy fashion…
Scott Cantor….he is behind the Shib consortium…..principal pushers at ITF for SAML ECp standardization….if they cannot make ECP support proxies you're in troubles (SoL) .

Andrea: Works for the PX, likely won't work for the IDP….if your IDP is behind a proxy like surfnet (hub and spoke fed)..... You won't be able….
Own tests and EGI

MarioD: we have OpenIDCOnnect…. So as Andrea knows…. **Andrea:** that works….

Andrea: You need password credentials at the proxy for that to work.. Our IDP proxy in the middle, gives the user a password, you go there to use CLI, you get there and proceed with the flow to enable CLI applications…
If you want a real flow, you need ECP working. We see people working CLI applic working, not ECP working. If you can manage to provide something that works it is OK

MarioD: we have OpenID Connect and IDP proxy, how does it work when you have a guy in edugain coming through SAML to the IDP ox, is it oK just to have the IDP proxy, or do you need more hops ?

Andrea: you do there the account linking if a users comes there, then register via SAML, then he gets a token for using CLI application

You get through the Web interface, then you create the user, then you get the token, then you have CLI available

After you create a user, you go to some client to the ODIP px, to get the access token…
Andrea: this is the old way. You get a password, either a user password or a scoped token yor you use the scoped token directly
The idea is that you stop there, you do not not go behind that. This normally is enough to support our use case.
Stefan: you have a PX password than an application password on top
Andrea; this is an davanagem, to support CLI applic, you want t away to revoke a token if people are not related to your IDP credentials.

Stefan: but also the other way round, If I a user using your resources….. Data belong to the pharmaceutical… you are going to allow that - the token at your PX has not been revoked…..

If you have this requirement then make sure to link the token the lifetime to the validity of a membership on an external IDP…….
Stefan: if I own an account on the PX, I do not go back to the IDP…

Andrea: Controlling the lifetime of a membership on an IDP proxy, it is another topic.
You can have some kind of mechanism by which IDP PX check the account linking to the external IDP, using a SAML query, or some other provision i/f.

Stefan: these are the kind of questions……..

MarioR: Have ECP supporting more than one hop would be the only cleaner way ?
Stefan: I would say so.

Andrea: have people investing in these standard supporting interface. Dealing with membership and linking membership to a cost organization, it is provisioning and AuthZ, not really AuthN. It is before AuthZ. If you want to do it across domains, It is a provisioning action, linking AuthZ or cleanup procedures……
Commercial give you solutions (ping identity for example), have integrated solutions and this provision stuff….you manage all your resources under big umbrella ensuring AuthZ is consistent across all your infrastructure…..to make it across organization you need standards for this. It could be meaningful to use indigo IAM, you can leverage INDIGO IAM to use the SCIM API.

Stefan: but yes, it would be very good if ECP will support an extra hop !

MarioD: whatever is above the IaaS level, then you need to use the APIs….unless you only have IaaS, anything above that, you need APIs

Davide: it is a limit on some protocols, on radius and OIDC you have a lot of tools, like scoped usernames (webfinger) - if you have OIDC provider PX in front of you, you can still reach the end, and in RADIUS you have domain...and PX know what to do...where to send this AuthN request….with SAML you obtain the scope from the IDP…..that is why it is difficult….in SAML you get the scope once you reach the IDP…..

Lalla; Outside the federation you can configure tenants for different scopes….microsoft and
Davide: Google are using in a very OIDC/or RADIUS way : you have to associate a domain to an IDP (username@yourdomain) they ignore the domain part and get you to the IDP using it.

Stefan: for the use experience, instead have to scroll 10000s IDP lists...it' is easier. Likely IDPs to be yours….

Davide: In the IDP MD you usually have the scope, if that domain matches the the mail domain of the HO, then this is the perfect match

In OAuth token can be opaque….depending on how you integrate it, you can have introspectable token (JWT) the relying party can look into it….or not…..the token introspection point gives AuthZ info linked to the token….the User info point is the second point (ID attributes for the user, ID, when authenticated, claims ,,)

MarioD: if you use the mod_auth there are 2 queries - one to the introspective endpoint and one to the user info endpoint
1 query only to the instorspecitve endpoint if you use the CLI - do not get the user info endpoint -

Andrea: that is why we provide more information on the token introspection endpoint - if you go to this test client info - if you go there, you see information is also at the token introspection endpoint - to have easier integration with OS Keystone…
Jose: you only have to modify your plugin…….**Andrea**: we want to avoid upstream changes...we needed to showcase this in a demo ! it is a temporary solution until we fett full OIDC connect support in the the keystone upstream - we minimized the required changes upstream.

==========

Product working group User story submission

The etherpad with our user story
https://etherpad.openstack.org/p/scientific-wg-rome-federated-identity-user-story

Action point for Saverio:
https://www.openstack.org/summit/barcelona-2016/summit-schedule/events/16855/product-working-group-working-session