Open Issues on Distributed Authorization for Cloud Infrastructure

Mihály Héder, MTA SZTAKI	mihaly.heder@sztaki.mta.hu
Szabolcs Tenczer MTA SZTAKI	burgosz@sztaki.hu
Michal Procházka, CESNET	michalp@ics.muni.cz
Maarten Kremers, SURFnet	maarten.kremers@surfnet.nl
Andrea Biancini, GARR	andrea.biancini@garr.it
Slávek Licehammer, CESNET	slavek@ics.muni.cz

Keywords: SAML, Attribute Authorities, OpenStack, OpenNebula

Problem statement

The efficiency of cloud infrastructure is partially based on the economics of scale. The more resources are concentrated the cheaper a unit will become. However, the efficiency of authentication and authorization does not scale the same way. Experience shows that subsidiarity is more sustainable in the long term than centralization. In order to harvest the full potential of cloud systems they need to be facilitated with a distributed authorization system in which communities with a given resource quota in the cloud can manage their own members and provision resources for them.

In this paper we describe how distributed authorization can be achieved in SAML federations using SAML IdPs and SAML Attribute Authorities. We show how we enabled OpenStack and OpenNebula in a SAML federation with multiple Virtual Organization and Group management services.

The main problems that arise in such project are the naming, scoping and filtering of authoritative attributes; the need of a persistent identifier that is a consequence of the architecture; provisioning and de-provisioning federated users and resources in clouds; technical problems of the non-parallelity of AA queries.

Showcases

We have deployed an OpenStack instance - https://openstack.hbit.sztaki.hu that is now part of eduGAIN as a service provider. This instance relies on three different attribute authorities (AA) for authorization information: a HEXAA[1], a PERUN[2], and a Grouper[3] deployment. Besides, it is open for any users from any IdPs from eduGAIN, therefore authentication is separated from authorization. These AAs implement Virtual Organization/Group

[1] http://www.hexaa.eu
[2] http://perun.cesnet.cz
[3] http://www.internet2.edu/products-services/trust-identity-middleware/grouper/

management. The goal of this deployment was to show how the membership information can be provisioned from multiple sources in order to grant access to OpenStack resources.

OpenStack was made to be compatible with SAML federations with the recent Keystone WebSSO Specification[4] which was complemented by regsite[5] (Heder et. al. 2015). Regsite uses the information gathered on SAML to register users and projects/groups in OpenStack.

The attributes that are carrying the authorization information are *eduPersonEntitlement* and *isMemberOf*. The Shibboleth SP makes sure that any of the attribute authorities can issue a certain set of entitlements only by attribute policy configuration and that they cannot release a value that is reserved for another AA. Therefore, an implicit *scoping* of attribute information is achieved - while this approach works, it is clearly a local solution right now. The scoping of authoritative information, e.g. eduPersonEntitlement needs to be standardized. In this topic, substantial effort has been made by Niels van Dijk[6].

Another area in which this OpenStack deployment needs to be improved is the de-provisioning of users. While user and project de-registration can be done by regsite, standardization of notification of regsite by the attribute authority about the fact that de-provisioning has to be done is necessary.

For OpenNebula, initial SAML integration[7] was created by MTA SZTAKI from v.3.6 to v.4.10.0 (The current version is v.4.14.2) Unfortunately, the integration code is no longer supported. The latest version of the integration supports Shibboleth SP, therefore it is possible to use any SAML IdP-s and any number of AA-s. Based on an *eduPersonEntitlement* attribute, the user is associated with a group within OpenNebula. Main method of authorization in OpenNebula cloud is group based. A user can only belong to one primary group at a time, and optionally to several secondary groups. However, in our experience the best practice is only to rely on primary groups and we have developed a custom group chooser screen to implement that.

OpenNebula with YAVOM[8] (HEXAA's predecessor) has a 4 year long successful operational history at MTA SZTAKI. However, the integration requires patching of the Sunstone web frontend's code, and maintaining the patch requires significant effort. This is in contrast with OpenStack, as thanks to OpenStack's highly modular design, the integration code in regsite is fully self-contained. The OpenNebula patch is not currently upgraded to be compatible with newer versions of OpenNebula. OpenNebula's status is assessed and requirements for a future SAML integration are laid out in a white paper (Parak 2015).

[4] https://specs.openstack.org/openstack/keystone-specs/specs/kilo/websso-portal.html

[5] https://github.com/burgosz/openstack-horizon-shibboleth

[6] https://github.com/surfnet-niels/edupersonEntitlement-and-IsMemberOf-scoped-semantics

[7] https://github.com/burgosz/opennebula-sunstone-shib

[8] YAVOM stands for Yet Another Virtual Organization Manager. It was made by MTA SZTAKI and presented on TNC 2013: https://tnc2013.terena.org/getfile/95 HEXAA is a completely rewritten successor of YAVOM, with some design principles kept.

Open Issues

While many key features have been implemented in our OpenStack and OpenNebula deployments, some aspects of the integration need improvements in the future. In general, the authorization of non-web cloud interfaces, e.g. command-line interface or VNC needs to be improved. These might be solved by Moonshot in the future.

While the federated access of cloud interfaces with distributed authorization was the necessary first step, a much deeper integration is imaginable. It is envisioned that we can access the virtual resources themselves by relying on SAML federations. This means e.g. the access of VMs SSH interface in a similar SAML-enabled method that is employed to access the cloud management interfaces. Again, this is an area where Moonshot might provide a solution in the future.

Our experience with distributed authorization shows that a standardized naming and scoping scheme for group names and entitlements is a must-have. This will not only facilitate the future deployment of multi-AA setups, but is also essential for the SAML integration of web applications. If authoritative information comes in a standard fashion from the federation then less effort will be needed in the application in order to map this information to local resources.

Similarly, at least a partial standardisation or harmonization of de-provisioning is necessary. Whether this can be built upon SCIM or another existing standard or has to be created from the ground up is a question of the future. De-provisioning is usually overlooked, however in distributed environments where direct connection to the user’s primary account is missing, proper process of de-provisioning is crucial. When user has lost permissions to access the cloud infrastructure then all his running VM-s must be handled appropriately, for example shutted down. This can be done only if the cloud infrastructure is somehow notified about change of the permissions. Therefore we have to deal with different types of transportation of the authorization information.

Acknowledgements

This work has received funding from the European Union’s Horizon 2020 research and innovation program under Grant Agreement No. 691567 (GN4-1).

References

Mihály Héder, Szabolcs Tenczer, Andrea Biancini (2015). Collaboration Between SAML federations and OpenStack clouds. Submitted to International *Journal of Cooperative Information Systems*. arXiv: http://arxiv.org/abs/1510.04017

Boris Parak (2015). *Using OpenNebula with SAML-based Authentication and Authorization*, author manuscript, CESNET. Online (November 2015): https://wiki.metacentrum.cz/w/images/f/f9/UsingOpenNebulawithSAML-basedAuthenticationandAuthorization.pdf

Vitae

Mihály Héder, Ph.D. graduated as a Software Engineer at Technical University Budapest in 2009. Since 2004 he has been working for the *Computer and Automation Research Institute of the Hungarian Academy of Sciences*. His experiences include java-opensaml based development, Shibboleth IdP module development, simpleSAMLphp-based and OIOSAML-based development and the SAML-integration of several applications. In HEXAA project he worked on requirements engineering, software design and integration tasks. He is currently working for GN4-JRA3-T1. He defended his PhD thesis on the Philosophy of Artificial Intelligence in 2014 at TU Budapest.

Michal Procházka, **Ph.D.** works at Masaryk University and CESNET (Czech NREN) mainly focusing on IT security and identity and access management area. Issue of federated identity and the concept of identity federations is one of his major scope within the identity management area. On the same topic he defended his doctoral thesis at Masaryk University in 2015. Since 2011 he has been leading project Perun -- identity and access management system. He is involved in AARC, GEANT GN4p1 and MAGIC projects focusing on AAI topic, he is also co-lead of AAI task in ELIXIR EXCELERATE project.

Maarten Kremers joined SURFnet in 2007 and is currently technical product manager Trust and Identity. He is involved as project member and project manager in the innovation and development of collaboration and identity management infrastructure SURFconext, with focus on group management and authorization. Furthermore he is currently leading the GN4 research task on authorization and attribute management in the federated identity area as well as leadings its predecessor in GN3plus. He holds a MSc degree in Information Management from Tilburg University.

Andrea Biancini has developed his career working for companies in the Finance and Information Technology fields and I dealt with: IT project management, planning and governance on themes of management and control (budget/costs), project portfolio management. He is particularly oriented in favoring the adoption of structured and effective working methodologies. He is also naturally oriented at developing an organic and systemic vision upon processes, resource management and governance activities.
In the last 4 years, Andrea worked for the Italian NREN on different European projects in the field of AAI. Andrea participated to Géant projects starting from GN3+ and has also took part to the AARC project acting for the first 8 months of the project as task leader of the JRA1.4 activity.

Slávek Licehammer graduated from Faculty of Informatics Masaryk University in 2012 with master's degree. Currently he has been studying PhD programme at the same faculty. He has been working since 2010 at CESNET as a researcher in Network identity division and also he has been working since 2012 at Masaryk University as a developer responsible for technical support of AAI activities. Moreover he has been participating in European projects GN4, EGI-ENGAGE and INDIGO, working on AAI related tasks in all of them. He is a technical manager of Perun project.