IDP IN THE CLOUD

Federated identity management as a service at GARR

Fabio Farina[1], Andrea Biancini[1], Maria Laura Mantovani[1,2],
Marco Malavolti[1], Pasquale Mandato[1], Cristiano Valli[1], Luca Prete[1], Sabrina Tomassini[1]
[1]*Consortium GARR*
[2]*Università degli Studi di Modena e Reggio Emilia*

{fabio.farina, andrea.biancini, marialaura.mantovani, marco.malavolti, pasquale.mandato, cristiano.valli,luca.prete,sabrina.tomassini}@garr.it

Abstract

Identity federations are a benefit for the NRENs' communities and core e-Infrastructure services. Common protocols and attributes offer to the users more services and a simpler management of the credentials, while allowing service providers (SP) to reach a larger group of potential users. However, smaller institutions may struggle when setting up the tools needed to join a federation. The lack of resources, expertise and manpower can discourage the setup of an identity provider (IdP).

The "IdP in the cloud" service has been developed to cope with these issues. The service relies on Infrastructure-as-a Service paradigm and Development-Operation methodology (DevOps) to automate and simplify the creation and the maintenance of an IdP appliance, letting the local account managers free to focus on the users accreditation and policies.

This paper describes the needs that brought GARR and the IDEM federation to define the new service, introduces the principles, the implementation and the early operation. The article reports feedback from the early adopters and outlines the planned roadmap.

1. The IDEM Federation

Consortium GARR, the Italian NREN, created and supports IDEM [1]: an identity federation among the national universities and research institutions. IDEM fosters the effort of its community to define and support a common framework that allows users to access on-line resources through the unique identity their organization provides them. The federation adopts SAML2 [2] as assertion exchange protocol with Shibboleth [3] and simpleSAMLphp [4] as favorite implementations. Thanks to the standard compliancy, IDEM is in line with other NRENs activities, is member of eduGAIN [5] and participates to REFEDS [6].

Many Italian institutions participate to IDEM, counting 48 IdPs and 88 SPs. However, some institutions lack the manpower and the knowledge needed to adopt SAML-based technologies. It is essential to lower the barrier for joining to identity federations, also providing ready to use solutions, as their effectiveness grows with the adoption rate. Therefore, it is important for the whole GARR community that the services accessible through IDEM can be made available also to partners with limited resources, increasing further the federation adoption.

1.1 Federating can be though

Identity federations are not trivial to deploy and configure. This holds in particular for small organizations that cannot, or do not want, allocate time, money, and resources to understand SAML and AAI details. Instead, they would like to focus on account management getting only the pros of the federated services. Joining a federation implies also formal steps, like signed agreements and acceptance of policies that could slow down further the process. The overall complexity hinders the benefits of taking part to a federation, discouraging also new service providers in a vicious circle that keeps the participation below the desired levels [7].

Part of the Italian biomedical research is funded by the Italian Ministry of Health. The Ministry and GARR entered into a multi-year framework agreement for joint projects to provide network services through the GARR-X backbone and advanced federated applications (cloud storage, large files sharing, HD Video-Conferencing). The involved research community is affiliated to 55 research hospitals and is generally made of small groups with 200 researchers on average in each hospital. Their constrained IT resources impose a review of the whole IDEM adoption process, simplifying both the organizational and the technological aspects. Moreover, the identity management within a large organization can suffer from a difficult attribution of responsibility. It is therefore essential to grant an easy access to the federation services for these non-ICT focused users, providing a flexible tool to benefit from sharing applications and research results with transparent rules.

1.2 Make it easy: IdP in the cloud

Wanting to tackle the problem as a customer service optimization problem, the process of joining the IDEM Federation has been formalized using a workflow. The tasks' interdependencies have been analyzed and minimized to favor pipelining and parallel execution. Subsequently, the tasks have been characterized according to the relevant stakeholders: users providing information about the organization, agents for the repetitive tasks, and operators addressing pre- and post-conditions to satisfy the policies required to join the IDEM production federation.
GARR has developed the "IdP in the cloud" service to implement the workflow shown in Figure 1. "IdP in the cloud" focuses on the efficient creation and configuration of new IdP virtual appliances in high availability, with identity management tools and monitoring. The activities in the dashed box will be discussed in the next section.

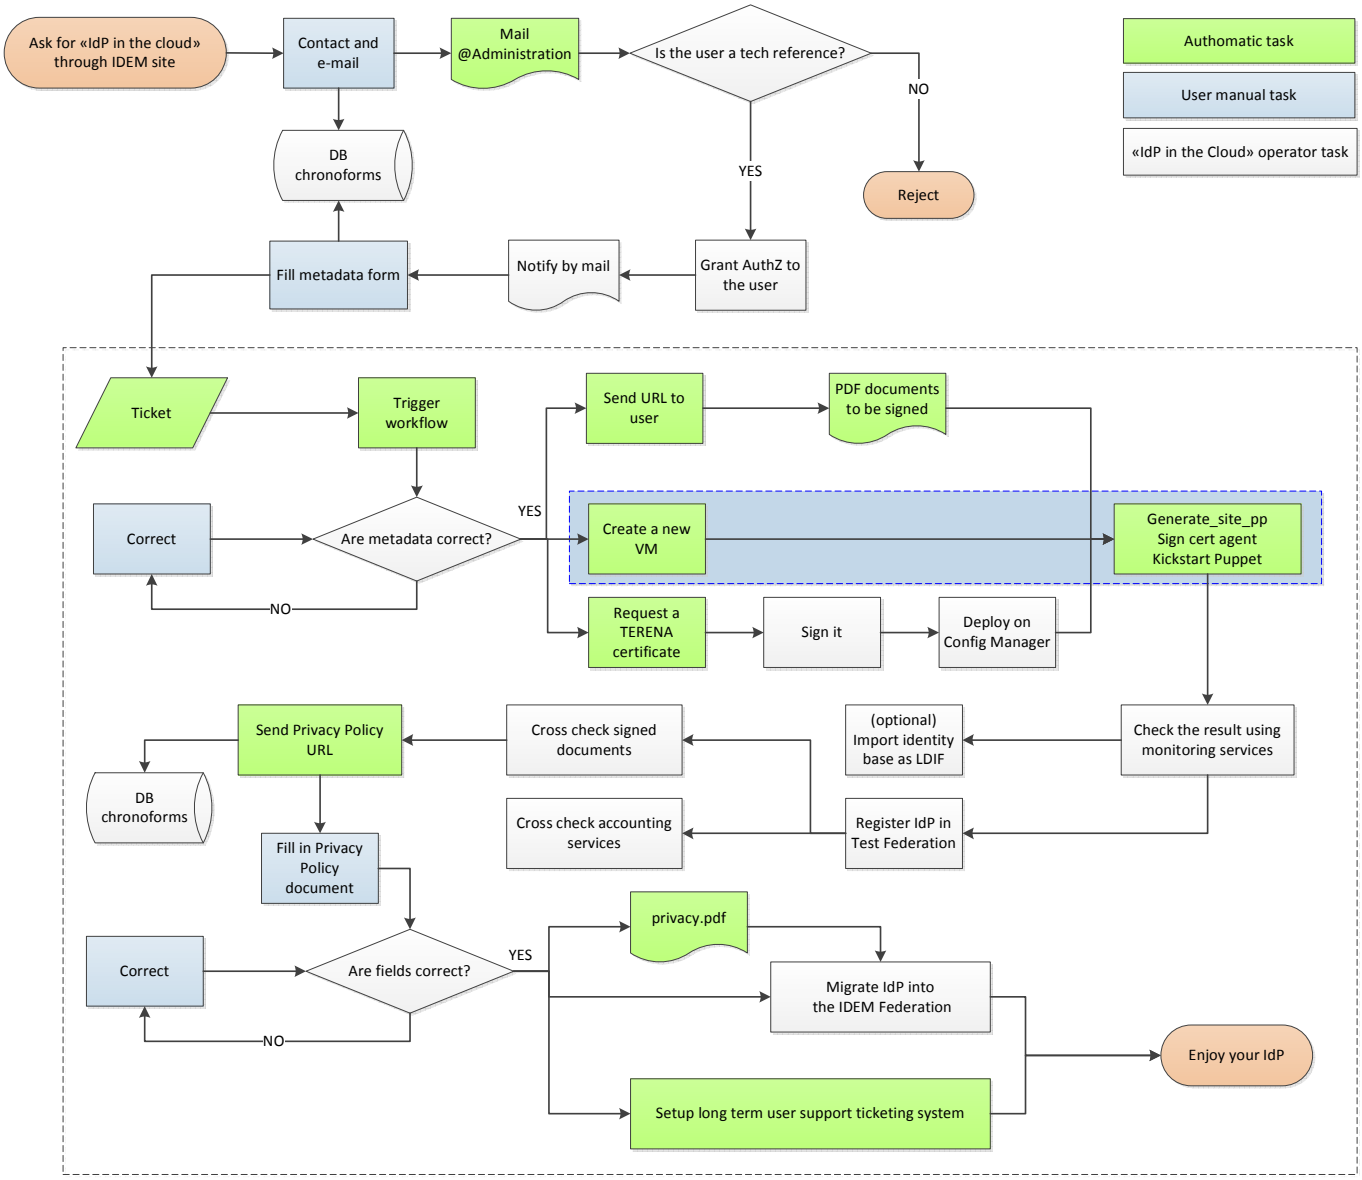

Figure 1 – IdP in the cloud workflow for creating a new Identity Provider virtual machine

The service minimizes the manpower requirements and increases scalability, with few operators being able to administer hundreds of IdPs. The user's tasks are reduced by 80% and the organizations are led to join the IDEM federation with a workflow completion time reduction of 88%. The user duties focus on initial provisioning of information about the Organization and on the daily identity management once the IdP is delivered.

2. The technology: IaaS and DevOps

"IdP in the cloud" leverages IaaS cloud and DevOps agile methodology [8] to provision new IdPs in a few minutes with a PaaS strategy.
GARR harmonized the following tools in a distributed infrastructure:

- OpenStack [9], used to create the VMs that host the IdP and the related networking properties like firewall rules, and IP address management
- GlusterFS file system [10], to ensure resilient geographical replication of both the VM instances and live migration. GlusterFS is used also to persist Organization data
- On top of them, Puppet [11] configuration manager installs and configures the software dependencies to deploy a Shibboleth IdP, an LDAP registry and web interfaces for monitoring and identity management. If the Organization wants to use an external preexisting identity base, Puppet either imports it or connects it to the IdP through a secure VPN channel according the user preference. Auxiliary tools like phpLDAPadmin, uApprove and custom login pages are deployed.
- Monitoring and alarming rely on Nagios, Collectd and Splunk. Both the user and the service operators are constantly informed on the status and the resource consumption of the IdPs.

Figure 2 shows how the technologies interact to create a new IdP.

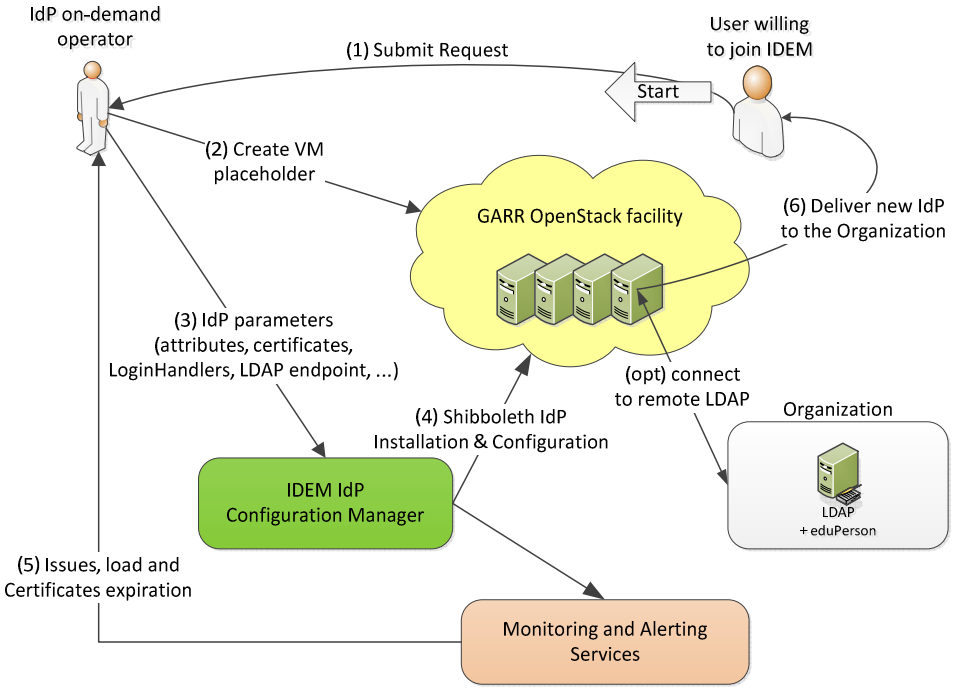

Figure 2 – OpenStack, Puppet and monitoring interactions during IdP creation.

Another advantage of the "IdP in the cloud" approach is that the IdPs have an inherently harmonized set of attributes and metadata, fulfilling since the beginning IDEM and eduGAIN recommendations. In addition, the metadata are automatically enriched as prescribed by REFEDS Discovery guidelines.

3. Early adopters and learned lessons

"IdP in the cloud" opened to beta users in June 2013. Three biomedical organizations have been chosen to assess the service. Their feedback on the issues with federated resources, misconfigurations and usability allowed us to improve the service. The DevOps approach showed its full potential by enabling the release of fixes to all the IdPs automatically and transparently in a Continuous Delivery way.

Resiliency and live-migration have prevented interruptions of both the IdPs and the service itself when minor infrastructure issues occurred during the beta period.

The collected feedback indicates appreciation of the Commercial-Off-The-Shelf Software (COTS) approach: the IdPs are provided and maintained at a marginal cost, updates and fixes are more frequent. The overall quality of the federation is higher as compliancy and reliability are assessed regularly by the IDEM operators, stressing in particular security review.

4. Conclusion and next steps

This paper discusses how GARR and the IDEM Federation have designed and implemented a new cloud service to lower the complexity for the organizations to join a federation.

For 2014 we plan to offer every Italian research hospital an IdP in the cloud. Other communities, like the Digital Cultural Heritage community, are already showing interest in the service.

On the technical perspective, IDEM and GARR will continue the effort to bring other AAI entities into the cloud, aiming at the creation of a whole Federation-as-a-Service. In this direction GARR is collaborating with other NRENs in the frameworks of GN3+ and ELCIRA projects.

References

[1] IDEM *Federation for Authentication and Authorization*. https://www.idem.garr.it (2009)

[2] Ragouzis, N. et al. *Security Assertion Markup Language (SAML) V2.0 Technical Overview*. OASIS Committee Draft. Document ID sstc-saml-tech-overview-2.0-cd-02 (2008)

[3] Morgan, R. L., Cantor, S., Carmody, S., Hoehn, W., and Klingenstein, K. *Federated security: The Shibboleth approach*. EDUCAUSE Quarterly, 27(4):12–17 (2004)

[4] simpleSAMLphp http://simplesamlphp.org/

[5] eduGAIN *Inter-federation of national identity systems across Europe*. http://www.edugain.org (2011)

[6] REFEDS *The voice of research and education identity federations*, https://refeds.org/index.html (2009)

[7] TERENA *A Study on Authentication and Authorisation Platforms For Scientific Resources in Europe*, https://confluence.terena.org/display/aaastudy/AAA+Study+Home+Page (2012)

[8] Edwards D. *What is DevOps*, http://dev2ops.org/2010/02/what-is-devops/ (2010)

[9] OpenStack *Open source software for building private and public clouds*, http://www.openstack.org/ (2012)

[10] GlusterFS *GlusterFS the open source, distributed scalable file system*, http://www.gluster.org/ (2005)

[11] Puppet *The Puppet installation and configuration tool*. https://puppetlabs.com (2005)

Short Biographies

Fabio Farina, PhD. He joined GARR in 2010 after some years at CERN working on the CMS Grid on distributed data analysis. Currently Fabio works on cloud services, focusing on storage and Future Internet test beds. He designed the architecture for “IdP in the cloud” and coordinated the initiation of the project.

Andrea Biancini. He has developed his career in different private companies where he was responsible for managing IT projects, planning and governance and project portfolio management. Always interested in the human dimension, he leaded the execution phase for "IdP in the cloud" and in the meanwhile got a degree in Psychology.

Maria Laura Mantovani, MSc. Networks and information security expert, for 20 years worked at University of Modena and Reggio Emilia as network architect and subsequently as chief security officer. In this role she built the academic identity management system. She joined GARR in 2009 and here is the coordinator of IDEM, the Italian identity federation for R&E.

Marco Malavolti, BSc. is a developer and he joined GARR in January 2013. He got a bachelor in Computer Science at University of Bologna working on AAI entities and Metadata Registry tools. He develops the code for IdP in the Cloud.

Pasquale Mandato. He is a system support engineer and an OpenStack expert. He took care of the IaaS and security aspects for “IdP in the cloud”.

Cristiano Valli. He coordinates the system support unit at GARR. Cristiano designed and rolled out the distributed infrastructure hosting the “IdP in the cloud” service.

Luca Prete, BSc. He is network engineer. He’s doing a two-year internship with GARR R&D on Software Defined Networking. He is interested in networking, cloud and virtualization technologies.

Sabrina Tomassini, MSc. As a senior network engineer, she is part of the Network Planning and Engineering team that collects the user community requests for connectivity and plans technical solutions to implement them. She is currently following initiatives connected to Ministry of Health and Ministry of Culture. She joined GARR in 2007 and she is also involved in GARR member and partner relations.