

Identity Harmonization for Federated HPC, Grid and Cloud Services

Benjamin Ertl, Uros Stevanovic, Arsen Hayrapetyan, Bas Wegh, Marcus Hardt
Steinbuch Centre for Computing (SCC)
Karlsruhe Institute of Technology (KIT)
Karlsruhe, Germany
{benjamin.ertl,uros.stevanovic,arsen.hayrapetyan,bas.wegh,marcus.hardt}@kit.edu

Abstract—With the increasing acceptance of multiple authentication mechanisms, federated infrastructures need to provide means of keeping consistency between multiple user identities. Although the current authentication and authorization infrastructures are designed to support multiple ways of authentication (SAML, OpenID Connect, X.509), they are missing unified protocols and interfaces to harmonize multiple user identities. This article introduces the concept of identity harmonization for federated cloud services. Our approach is based on the standardized System for Cross-domain Identity Management (SCIM) protocol. We add the support for account linking and per-service verification. Furthermore, the concept is put into context of currently existing federated infrastructures and is exemplified within a federated e-infrastructure currently developed in the course of the INDIGO-Datacloud project. The concept is evaluated in the INDIGO testbed in terms of deployability, scalability, provisioning and deprovisioning of user accounts, as well as maintenance and integration effort.

Keywords—Authentication, Authorization, Cloud Computing, Federated Services, Grid Computing, Identity Management, HPC

I. INTRODUCTION

The use of federated identities is an established practice in various research communities, e.g. high-energy physics, social science and humanities, bioinformatics, fusion energy. It offers numerous advantages and convenience to organizations and their users [1]. Federated Identity Management (FIM) relies on an arrangement between multiple organizations that lets subscribers use the same identification data to obtain access to the secured resources of all organizations in the federation [2]. Using federated identity, however, introduces potential privacy and security issues, and also comes with architectural challenges [3]. With the development and incorporation of new authentication mechanisms, e.g. OpenID Connect (OIDC) [4], into existing federated environments, current issues could be more easily resolved. Traditional Authentication and Authorization Infrastructures (AAI) in scientific research communities are usually based around centrally managed user accounts and groups. This does not scale well with the increasing amount of users and with changes of the users' affiliations [5].

Extending federated environments like HPC-, Grid-, Cloud- and other e-infrastructures with new authentication mechanisms requires a consistent way of keeping multiple user identities harmonized.

Researchers may want to access data of an experiment conducted with a federated user identity using also their institutional identity. Today this is only supported to a limited extent if at all within current federated environments.

Current e-infrastructures are predominantly based on X.509 [7] authentication. Security Assertion Markup Language (SAML) based installations [6] are currently being developed and deployed e.g. in the EGI context. They are used with Attribute Authorities (AA) or attribute certificates to enrich authenticated user identities with attributes that are required for access control at the services [5].

The integration of novel authentication mechanisms leads to users having multiple identities in existing environments, which requires a unified approach across different authentication protocols to harmonize these identities. This is especially the case if local accounts and federated identities can not be synchronized or local systems require specific access credentials. Further, the demand to support changes of users' affiliations can not always be supported by current approaches. For example attribute certificates lack the required flexibility for changing user attributes and can not always support users authenticated via different protocols.

In this paper we propose a novel approach for harmonizing user identities based on the standardized System for Cross-domain Identity Management (SCIM) protocol [8]. The concept adds support for account linking and user attribute changes, introducing a per-service verification mechanism.

In the context of our approach, a user identity always refers to a human being, which however can have multiple identities. For example user Alice is participating in two research groups G1 and G2 at different research institutes using different authentication mechanisms, e.g. SAML and OpenID Connect. Additionally to the user's main identity Alice, two secondary identities AliceG1 and AliceG2 are established. The user's information are encapsulated in accounts, sometimes referred to as profiles [9].

The proposed approach allows Alice to link both identities AliceG1 and AliceG2 so that the identities are associated with the same local user at both research institutes. In general, the Identity Harmonization Service proposed receives the identity information and interacts with the local user management to perform the identity linking.

We evaluate the concept from the perspective of an infras-

structure and service provider according to the deployability, scalability, provisioning and deprovisioning of user accounts, maintenance and integration effort.

The INDIGO-Datacloud [10] project has been selected as an example architecture, because it supports authentication from home Identity Providers (IdPs) (e.g. institutional, SAML-based) and social-IdPs (e.g. Google, OpenID Connect based) alike.

The remainder of the paper is organized as follows: Section II provides the background of the proposed method while Section III describes the concept in detail. Section IV describes the exemplification of the concept within the INDIGO-Datacloud project. Section V presents the practical implementation in the INDIGO -Datacloud prototype testbed at the Karlsruhe Institute of Technology (KIT) and the evaluation, test results and their assessment. At the end in Section VI the conclusions and outlooks are provided.

II. BACKGROUND

A typical AAI infrastructure of a federated community can be described as following. Users can be represented by different identity providers and different technologies, including certificate authorities and social IdPs (user representation level). These user representations can be enriched with attributes from additional attribute authorities (attribute enrichment level). To get access to end-services, either direct or via token translation, access will be granted to the authenticated and authorized user. The attribute enrichment can either occur at the end-service level directly, via a translation layer (token translation level) or at the Service Provider level (SP/IdP Proxy). SP/IdP Proxies are most commonly supported by current e-infrastructures to support SAML authentication, as outlined in Fig. 1.

The Security Assertion Markup Language (SAML) is an XML-based framework that allows identity and security in-

formation to be shared across security domains [6]. SAML is one of the most widely adopted industry standards for Single Sign-On (SSO) for cloud and enterprise applications [11]. The outdated versions 1.0 and 1.1 have been specified by the Organization for the Advancement of Structured Information Standards (OASIS), where the current version 2.0 is the convergence of three standards: SAML 1.1, Liberty Alliance Identity Federation Framework (ID-FF) and Shibboleth, a single sign-on infrastructure standard and reference implementation [11].

In the SAML Web SSO profile, the user requesting an access to an end-service is redirected by the end-service provider to an identity provider, which is usually chosen by the user. Beforehand, an initial level of trust must be established between SP and IdP. The user has to authenticate at the IdP which issues an authentication assertion to the SP. The SP verifies the assertion and makes the authorization decision based on the verified user's identity. A simplified view of the SAML authentication architecture is illustrated in Fig. 1.

However, a major drawback of this approach is that it requires the user always to login via a web browser. Although this issue has been addressed with the SAML v2.0 Enhanced Client or Proxy Profile (ECP) [12], which defines a SAML SSO profile over HTTP that allows clients to contact the user's IdP directly without requiring redirection by the SP, the ECP solution has not been widely adopted. Today more flexible and lightweight protocols e.g. OpenID Connect exist. Therefore, more and more e-infrastructures also allow different non-SAML based identity providers with different authentication mechanisms.

OpenID Connect 1.0 [4] is a simple identity layer on top of the OAuth 2.0 protocol [13]. It allows clients to verify the identity of an authenticated user, by introducing an authorization server. Furthermore, it allows obtaining profile information about the user in an interoperable and RESTful manner.

Extending federated environments like cloud- and e-infrastructures to make use of new authentication mechanisms increases the flexibility and user acceptance of such infrastructures. Enabling users to combine multiple identities, potentially spanning different technologies, has not been addressed in current architectures. Therefore, we propose a novel approach for harmonizing user identities. Our implementation is based on the System for Cross-domain Identity Management (SCIM) protocol [8]. It defines a RESTful (Representative State Transfer) protocol for identity management. It is specified as an HTTP-based protocol with JavaScript Object Notation (JSON) based user, groups and general resources representation. The supported authentication and authorization mechanisms for the SCIM RESTful APIs are all standard HTTP authentication and authorization schemes, for example bearer tokens and basic authentication. In industry, SCIM is among others, supported by Google, Ping, and Salesforce [11].

Our concept is exemplified within the INDIGO-Datacloud project, which is an EU project funded under the Horizon-2020 programme with Grant Nr 653549. The main objective of the

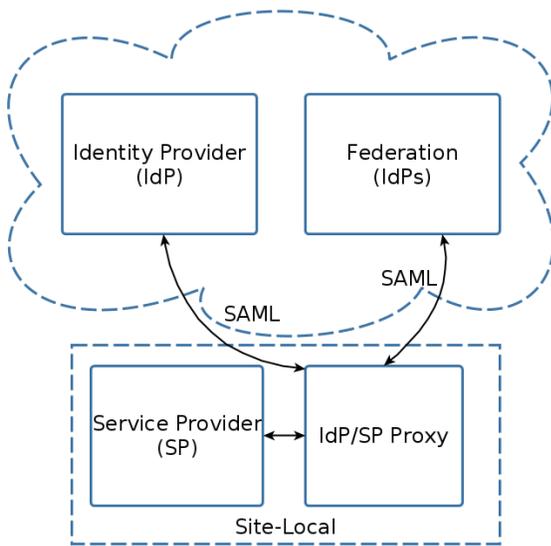


Fig. 1. Simplified SAML authentication architecture

project is to develop an open-source platform for scientific computing and data storage, deployed on public and private cloud infrastructures.

III. CONCEPT

Our concept introduces extensions to the current authentication infrastructures as indicated in Fig. 2, specific usage of the SCIM data format as seen in Listing 2 and an algorithm for user identity harmonization, see Listing 1.

The basic idea of the concept is to have a global access management service that release sets of linked user identities with user and group claims in the SCIM data format. It is provided as an aggregated user information object, where user and group claims are attributes that can be used to identify and map a primary user identity to additional, site-local and service specific user identities. An identity harmonization service at a given site receives this information and creates or updates local user accounts. According to the provided information, user and group claims can be verified with the proposed algorithm, Listing 2.

A simplified view of an authentication architecture including identity harmonization is outlined in Fig. 2. Multiple identity providers (SAML/ OIDC) are connected with a central Identity and Access Management service (IAM). The IAM is one of many global authentication services for the federated infrastructures, managing federated user identities and allowing the users to link together different identities provided by the supported identity providers. This functionality is already supported by the access management services in production, for example by the WSO2 Identity Server [14] or the OpenAM project [15].

Users can authenticate to the IAM via different authentication protocols, such as SAML **c.1** or OpenID Connect **c.2** or X.509 (not shown), but in principle there is no limit with regard to the choice of authentication scheme.

Linking different identities at the IAM by the user will result in an aggregated user identity information object with user and group claims, as exemplified in Listing 2. The IAM communicates with the site-local Identity Harmonization services via the SCIM protocol, see Fig. 2 connection **b**, forwarding the aggregated user information. Several models for ensuring that SCIM updates reach sites reliably and timely are possible, however, the choice of push, pull or pub-sub is an implementation detail.

After receiving the aggregated user information at the Identity Harmonization service (IdH), the service accesses the site-local User Management service (UM). IdH then modifies existing user accounts to satisfy the user and group claims of multiple user identities in the aggregated user information, see Fig. 2 connection **a**. Depending on site policies IdH could also be used to create missing user accounts, if needed.

Fig. 3 illustrates the typical flow for a user that wants to link two user accounts together in order to access a local service managed by the local user management service. Following this approach, users retain the privileges of both accounts, assuming that the local user management service allows the creation

of users authenticated at an external identity provider. This means that all users authenticated at the access management service can request access to the local end-services.

To accomplish this, the user logs in at the IAM with both accounts, following the preferred authentication mechanisms, i.e. SAML and OIDC in the given example. If successfully authenticated, the user can link both accounts together. This account linking triggers a harmonize user request at the IdH using the aggregated user information. The aggregated user information consists of user and group claims from both user identities. These claims are verified at the IdH against the local UM, e.g. a directory service or a local user database. The proposed algorithm for the harmonization step is described in Listing 1.

In the harmonization step, the Identity Harmonization service processes the aggregated user information in the specified format, see Listing 2, and returns the site-local user managed by the local User Management service (UM). The algorithm first retrieves or creates the local user account and default group identified by the primary user identity and then makes the connection between the user account and group (1-3). After the default user account has been established locally, additional group and user claims are verified and the default user account is updated accordingly (4-9).

The execution of the algorithm for different user and group claims results in the mappings outlined below, where the verification of users and groups of an authenticated user will always result in the retrieval or creation of local user and group accounts.

- One user identity and no groups provided
 - ◊ Verify the user against the local user management service and add a verified user to the default group

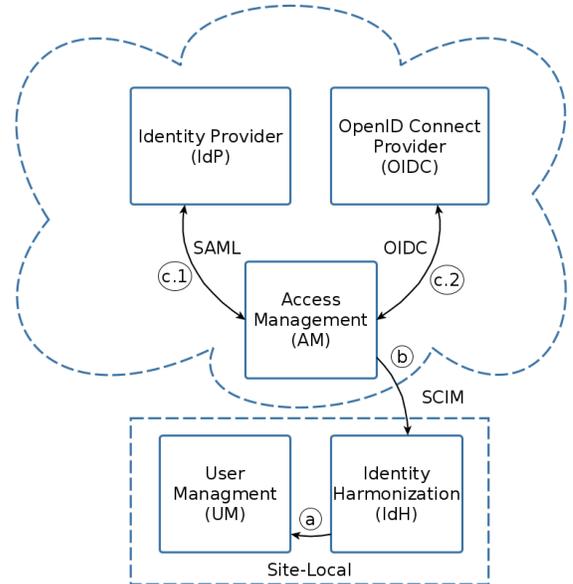


Fig. 2. Simplified authentication architecture with identity harmonization

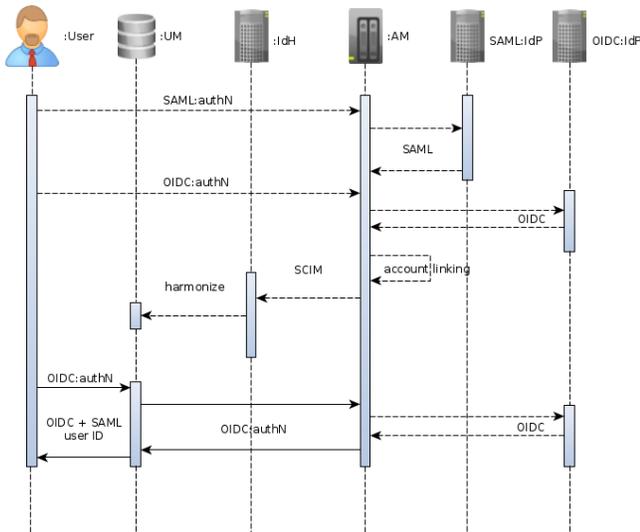


Fig. 3. Simplified authentication flow with identity harmonization

- One user identity and one or multiple groups provided
 - ◊ Verify the user and group claims against the local user management service and add a verified user to verified groups
- Multiple user identities and no groups provided
 - ◊ Verify user claims against the local user management service, map identities to the preferred and verified primary user identity and add a verified primary user to the default group
- Multiple user identities and one or multiple groups provided
 - ◊ Verify user claims against the local user management service, map identities to the preferred and verified primary user identity and add verified primary user to verified groups

Providing the aggregated user information in the SCIM data format can be achieved in various ways. For example

Listing 1 Harmonize with local user

```

Input: SCIM:User{User:Claims, Group:Claims}
Output: SCIM:LocalUser
1: user ← createOrGetDefaultLocalUser(User)
2: group ← createOrGetDefaultLocalGroup
3: addToGroup(user, group)
4: localGroups ← verifyGroupClaims(Group:Claims)
5: for localGroup in localGroups do
6:   addToGroup(user, localGroup)
7: end for
8: localUser ← verifyUserClaims(User:Claims)
9: user ← fromLocalUser(localUser)
10: LocalUser ← toScim(user)
11: return LocalUser

```

listing all the user identities in a multi-valued list of complex objects e.g. according to the `api:messages:2.0>ListResponse` and `schemas:core:2.0>User` schemas of the `urn:ietf:params:scim` schema, which defines all schema attributes. This format is similar to the defined query operations in the SCIM protocol.

Another approach could utilize the existing attributes of the `schemas:core:2.0>User` object in a specified predefinition and add additional user and group claims in the complex *meta* attribute of the `schemas:core:2.0>User` schema. For our proposal, the latter approach has been implemented since it provides a more compressed user information object, see Listing 2.

In the proposed data format, the `userName` attribute always maps to the `uid` of the primary user identity and the `externalId` attribute to the unique `id` of the access management service. If the user has linked different accounts together, the `userName` attribute is set to the `uid` which the user declares as the primary user's `id`.

All group claims are listed in the `groups` list in the `User` object, predefining the `display` attribute to map always to the group's common name (`cn` for POSIX groups).

Listing 2 Example SCIM user schema with extended metadata

```

1: {
2:   "schemas":["urn:ietf:params:scim:schemas:core:2.0:User"],
3:   "userName":"Alice",
4:   "externalId":"2abf56d2-9bdc-47d2-a6c7-074c24717879",
5:   "name":{
6:     "formatted":"Alice A.",
7:     "familyName":"A.",
8:     "givenName":"Alice"
9:   },
10:  "meta":{
11:    "uid":"AliceG1",
12:    "uid":"AliceG2"
13:  },
14:  "groups":[
15:    {
16:      "display":"G1"
17:    },
18:    {
19:      "display":"G2"
20:    }
21:  ]
22: }

```

The introduced concept results in a user identity that is harmonized with the local user management service and allows a local per-service authorization decision. It still requires an initially established trust relation between access management service, in the conventional SAML or OIDC concept mapped to the SP/IdP Proxy or Relying Party, and the user management service. However, this can now be achieved in a more dynamic

and scalable process than with currently existing authentication and authorization concepts.

The identity harmonization service also needs write permissions to the local user management service for creating and modifying users. This however can also take place in isolated and controlled environments, e.g. separate directory server partitions or user databases, to keep the user identities enabled for harmonization apart from local users and groups that should be excluded from the process.

IV. INDIGO-DATA CLOUD EXEMPLIFICATION

The proposed solution for consistent user identities in cloud environments has been realised in the INDIGO-Datacloud project as part of the authentication and authorization infrastructure. The first release of the project is scheduled for July 2016 and the second release for March 2017. The releases will contain the Platform as a service (PaaS) and Infrastructure as a service (IaaS) layer components, toolkits, APIs and services of the INDIGO open-source platform for scientific computing and data storage. As of this writing the project is in its beta phase for the evaluation of available components and services.

Fig. 4(a) outlines the envisioned INDIGO AAI, grouped according to the general AAI view in the federated infrastructures context, showing **Layer 1: User Representation**, **Layer 2: Attribute Enrichment**, **Layer 3: Translation**, **Layer 4: End-Services**.

In the INDIGO AAI, the identity harmonization service is part of the proxy and translation services (**Layer 3**) that interact with the authentication and extended attributes services (**Layer 1 and 2**), but also with the service providers and end-services (**Layer 4**).

For the beta release of the platform, the identity harmonization service has been implemented and evaluated as a proof of concept service that will be further developed and evaluated according to the scheduled public releases of the INDIGO project.

A typical use-case for an INDIGO user is to link the global INDIGO user identity with a local user identity, to access a local service managed by the local user management service with the INDIGO user as the primary user but with the privileges of both accounts.

First, the user needs to register or login as an INDIGO user at the INDIGO Identity and Access Management service (IAM), see Identity Management in Fig 4(a), with the preferred identity provider. At the initial registration, the INDIGO user identity information is propagated via the IAM to the utilized federated cloud environments, where each environment has an instance of the Identity Harmonization service (IdH) connected to the site-local User Management service (UM).

In the next step, the logged-in user authenticates additionally against the site specific home-IdP and links the INDIGO identity together with the home-IdP identity. After successfully linking the identities together, the IAM triggers a harmonization user request at the Identity Harmonization service with the aggregated user information following the proposed SCIM protocol and data format.

When the harmonization service successfully harmonized the global INDIGO user identity with the user's local identity, the user can then login with the INDIGO user identity at the site-local end-service. This might require an additional token translation step at the Site Token Service if the chosen authentication mechanism is not natively supported. The logged-in INDIGO user will have the same privileges as the linked local user at the provided end-services.

V. IMPLEMENTATION AND EVALUATION

The IdH service has been implemented as a RESTful web-service utilizing the Java application framework Spring [16]. The webservice provides an interface for the access management service to trigger the harmonization process for user identities provided as an aggregated SCIM user information.

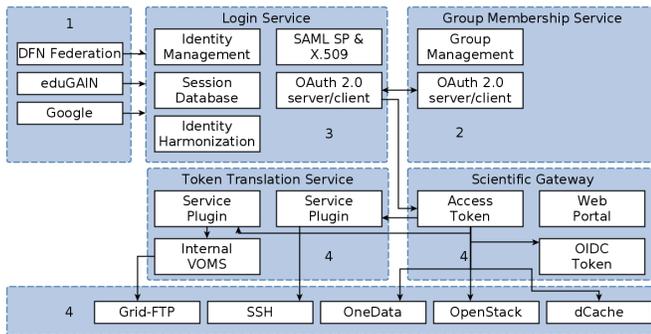
For the evaluation of the IdH service we deployed a local testbed according to the simplified authentication and authorization architecture depicted in Fig. 1. The set-up is composed of following components as outlined in Fig. 4(b).

- a SAML-based IdP
- an OIDC-based test authentication provider
- an instance of the LDAP-Facade [17] as an access management service
- the prototype implementation of the Identity Harmonization service
- an LDAP-based directory service for the local user management
- a test SSH end-service that authenticates users against the Lightweight Directory Access Protocol (LDAP) server

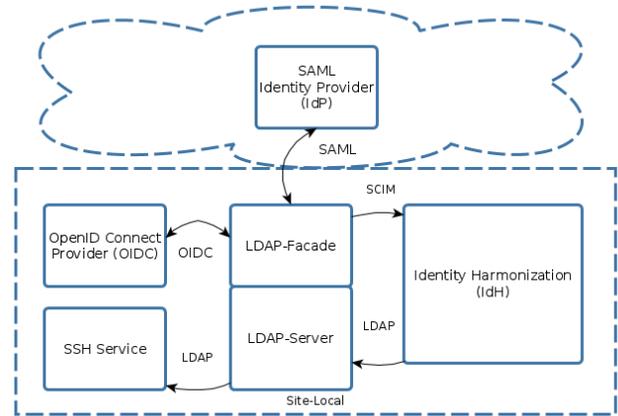
The LDAP-Facade, a service to seamlessly join non web-based services into the SAML federations, has been developed as part of the bwIDM project that was finalized in 2013 [18]. It is deployed in productive environments to provide a federated access to storage resources at the Karlsruhe Institute of Technology.

To demonstrate and verify the use-cases of harmonized identities, we enabled the test SSH service for the LDAP authentication against the deployed LDAP server. The LDAP server redirects the login request to the LDAP-Facade which in turn authenticates the user either against the SAML home-IdP or OIDC-based test authentication provider. Having both user accounts linked together at the LDAP-Facade, triggers the identity harmonization at the local user management service (the LDAP server in our testbed). This ensures that the user will always be logged in with the preferred user account information; in the case of the test SSH service, POSIX user id number (uidNumber) and POSIX primary group id number (gidNumber), and the associated POSIX groups (memberUid) as user and group claims.

Our evaluation shows that the proposed solution can be deployed in existing infrastructures without major adaptations, facilitated by the compliance of the IdH to standardized protocols and data formats. The programmatic extensions made to the LDAP-Facade are compatible to the SCIM standard specification and the connection between the IdH and the



(a) INDIGO authentication and authorization infrastructure



(b) Testbed authentication and authorization infrastructure

Fig. 4. Authentication and authorization infrastructure for evaluation

user management service is based on the standardized LDAP protocol.

Our evaluation criteria to measure the deployability have been the number of new components that need to be deployed, the number of existing components that need to be changed and the portability (OS and software dependencies) of the service. For our testbed the IdH service reaches a minimum count on all three criteria.

All supported harmonization use-case scenarios are successfully verified and demonstrated, providing consistent user identities and user privileges at the tested end-services.

The supported use-cases tested have been the following.

- Harmonize two SAML authentication based accounts
- Harmonize one SAML and one OIDC authentication based accounts
- Harmonize two ODIC authentication based account

Moreover, the proposed concept is scalable both in terms of supporting additional identity providers using standardized authentication protocols like SAML and OpenID Connect, and in terms of performance, due to the utilization of the lightweight SCIM protocol and RESTful IdH interface implementation.

Additional identity providers could be integrated in our tests without complications. This has been performed for our own test OIDC provider, the INDIGO IAM and Google OIDC.

Our performance measurements showed that the service scales vertically with the underlying server and horizontally with the number of provided servers with standard practices use for scaling web services, for example increasing the number of application processes and load balancing.

The identity harmonization service can run as a standalone Spring Boot application [19] on a separate server or on existing servers hosting other components, for example the LDAP-server, which keeps the additional administration and maintenance overhead low.

For our testbed we used one server hosting the LDAP-Faced, the LDAP server and the IdH service and an other server providing the LDAP-enabled SSH service. Due to this set-up,

the administrative effort for testing and debugging the service has been very low.

Concerning provisioning and deprovisioning of user accounts, the concept demonstrates a flexible approach via the newly introduced Identity Harmonization (IdH) service in collaboration with the Identity and Access Management service (IAM). The provisioning and deprovisioning processes are directly controlled at the IAM service by linking or unlinking user identities that represent different user accounts respectively. These actions are controlled by the users themselves, which need to authenticate at the different IdPs individually. As part of the harmonization step, the IdH service synchronizes the aggregated user information with the respective User Management (UM) service then.

VI. CONCLUSIONS AND FURTHER WORK

In this article we have described a novel concept for harmonizing multiple user identities for federated service in HPC-, Grid- and Cloud-environments. The concept introduces a new identity harmonization service that communicates via standard protocols with the local user management and the access management for federated environments. The communication to the access management service is based on the SCIM protocol with a specific, predefined usage of the data format. This maintains a high level of compatibility with the standard. The communication to the local user management service has been implemented and demonstrated using the LDAP protocol which is supported natively by numerous end-services, reducing the integration effort into current environments.

The concept also provides an algorithm for harmonizing and linking multiple user identities belonging to a single user. Identities are linked by the user at the access management service with the site-local user identities at the user management service. The algorithm verifies user identity and group membership claims expressed in the defined SCIM data format and modifies local user identities accordingly.

The introduced concept has been further exemplified through adaptation and integration in the scope of the INDIGO-Datacloud project. As the proof of concept implementation in the INDIGO-Datacloud project is still in progress, more site-specific user management services will be integrated and supported by the identity harmonization service.

The proof of concept implementation in the KIT testbed demonstrated that the solution can be deployed and integrated into productive authentication and authorization infrastructures without major changes to existing environments. The proposed solution is scalable in terms of number of supported identity providers and performance.

As the concept and implementation evolves, an in-depth testing of the reliability and performance in stress use-cases will be conducted.

Future work will deal with the definition of an integrated user schema within the current SCIM data format for consistent data formats and interoperability and will include additional methods for verification of user and group claims, e.g. attribute signing.

ACKNOWLEDGMENT

The authors want to acknowledge the support of the INDIGO-Datacloud project (grant number 653549) and the AARC project (grant number 653965), funded by the European Commissions Horizon 2020 Framework Programme.

REFERENCES

- [1] D. W. Chadwick, "Federated identity management," in *Foundations of security analysis and design V*. Springer, 2009, pp. 96–120.
- [2] D. Broeder, R. Wartel, B. Jones, P. Kershaw, D. Kelsey, S. Lüders, A. Lyall, T. Nyrönen, and H. J. Weyer, "Federated identity management for research collaborations," Tech. Rep., 2012.
- [3] E. Maler and D. Reed, "The venn of identity: Options and issues in federated identity management," *IEEE Security & Privacy*, no. 2, pp. 16–23, 2008.
- [4] N. Sakimura, J. Bradley, M. Jones, B. de Medeiros, and C. Mortimore, "OpenID connect core 1.0," *The OpenID Foundation*, p. S3, 2014.
- [5] M. Hardt, A. Hayrapetyan, P. Millar, and S. Memon, "Combining the X. 509 and the SAML Federated Identity Management Systems," in *Recent Trends in Computer Networks and Distributed Systems Security*. Springer, 2014, pp. 404–415.
- [6] P. Hallam-Baker, "Security assertions markup language," *May*, vol. 14, pp. 1–24, 2001.
- [7] R. I.-T. X.509, "X.509-The Directory: Public-key and attribute certificate frameworks," *International Telecommunication Union*, 2012.
- [8] E. K. LI, P. Hunt, B. Khasnabish, A. Nadalin, and Z. Zeltsan, "System for Cross-domain Identity Management: Definitions, Overview, Concepts, and Requirements (RFC7642)."
- [9] M. Benantar, *Access control systems: security, identity management and trust models*. Springer Science & Business Media, 2006.
- [10] INDIGO. (2015) INtegrating Distributed data Infrastructures for Global ExplOitation. [Online]. Available: <https://www.indigo-datacloud.eu/>
- [11] T. Q. Thanh, S. Covaci, B. Ertl, and P. Zampognano, "An Integrated Access Control Service Enabler for Cloud Applications," in *Future Network Systems and Security*. Springer, 2015, pp. 101–112.
- [12] S. Cantor, "SAML V2. 0 Enhanced Client or Proxy Profile Version 2.0," *OASIS Working Draft OASIS. sstc-saml-ecpv2. 0-wd04*, 2011.
- [13] D. Hardt, "The OAuth 2.0 Authorization Framework (RFC6749)."
- [14] WSO2. (2016) WSO2 Identity Server. [Online]. Available: <http://wso2.com/products/identity-server/>
- [15] OpenAM. (2016) OpenAM Project. [Online]. Available: <http://openam.forgerock.org/>
- [16] Spring. (2016) Spring Framework. [Online]. Available: <http://spring.io/>
- [17] J. Köhler, M. Simon, M. Nussbaumer, and H. Hartenstein, "Federating HPC Access via SAML: Towards a Plug-and-Play Solution," in *Supercomputing*. Springer, 2013, pp. 462–473.
- [18] bwIDM. (2013) Federated Identity Management for Institutions of Higher Education in the State of Baden-Wuerttemberg. [Online]. Available: <https://www.bwidm.de/>
- [19] Spring. (2015) Spring boot. [Online]. Available: <http://projects.spring.io/spring-boot/>