BNL-221976-2021-CPPJ

Federated User Account Management

M. Karasawa

Submitted to the CHEP2019 Conference
to be held at Adelaide
November 04 - 08, 2019

November 2020

Physics Department

Brookhaven National Laboratory

U.S. Department of Energy

USDOE Office of Science (SC), High Energy Physics (HEP) (SC-25)

DISCLAIMER

Federated User Account Management

Mizuki Karasawa[1,*] and *John* Hover[1,**]

[1]Brookhaven National Laboratory, Upton, N.Y. U.S.A

Abstract. BNL SDCC (Scientific Data and Computing Center) recently deployed a centralized identity management solution to support Single Sign On (SSO) authentication across multiple IT systems. The system supports federated login access via CILogon and InCommon and multi-factor authentication (MFA) to meet security standards for various application and services such as Jupyterhub / Invenio that are provided to the SDCC user community. CoManage (Cloud) and FreeIPA / Keycloak (local) are utilized to provided complex authorization for authenticated users. This talk will focus on technical overviews and strategies to tackle the challenges/obstacles in our facility.

1 Introduction

The Scientific Data and Computing Center (SDCC) is the main scientific computing center at Brookhaven National Laboratory (BNL). The SDCC manages computing for large collaborations like the experiments at the Relativistic Heavy Ion Collider at BNL and the ATLAS experiment at the Large Hadron Collider (LHC) at CERN. Recently, it has started to support other communities including researchers at the National Synchrotron Light Source II (NSLS-II) and the Center for Functional Nanomaterials (CFN), both at BNL, the Belle II experiment at KEK in Japan, as well as a number of smaller groups. As part of its mission, the SDCC has deployed numerous web services like Jupyter, Invenio, Mattermost, and Gitea that require authentication (AuthN) and authorization (AuthZ) services.[7–10] The AuthN and AuthZ services are required to ensure that only authorized users are accessing the service.

From the user's perspective, the need to authenticate when accessing each service can be a major irritation with the proliferation of web services at the SDCC. This is further exacerbated by the fact that the user typically access services hosted at multiple data centers, each with a different set of accounts and passwords. From the facility side, managing user accounts is a growing burden as the number of users and supported groups increases. To alleviate these problems, the SDCC has deployed a Keycloak based single sign on (SSO) system and is using it with CiLogon's CoManage collaborative management platform and the InCommon Federation. [1–4]

2 Keycloak

Keycloak is an open source identity management system that is the upstream code base for Red Hat's SSO solution. It was chosen by the SDCC for its integration with FreeIPA, its

*e-mail: mizuki@bnl.gov
**e-mail: jhover@bnl.gov

support for OpenID Connect (OIDC), OAuth2, and SAML2, its ability to broker identities, and its support for multi-factor authentication (MFA).[1, 5, 11, 12] Integration with FreeIPA allowed the Keycloak Id provider (IdP) to access local authentication and authorization information from SDCC's FreeIPA server, which is used to manage SDCC UNIX accounts. Multi-factor authentication support, in the form of time based one time passwords (OTP) was important as cybersecurity policy requires MFA for specific types of services, including interactive access. OIDC/OAuth2 support was necessary to allow the SDCC to use commercial identity providers like Google and Facebook while SAML2 was needed to enable participation in academic federations like InCommon and eduGAIN.[4, 6]

3 Single Sign On

Single sign on for local SDCC users was accomplished by connecting the Keycloak IdP (identity provider) to the SDCC FreeIPA server. The latter provides the user's identity (user's UNIX account name) and authentication method (user's Kerberos), as well as authorization information.[14] SDCC web services are protected by a service provider (SP) that utilized authentication information from the SDCC Keycloak IdP. When a user accesses a SDCC service for the first time, they are redirected to the SDCC IdP for authentication. Once authenticated the user can access any SDCC web service, for which he/she is authorized, without re-authenticating. Authorization information, in the form of group association, can be provided through Keycloak, from FreeIPA, from databases within Keycloak or from other external sources like CoManage. Figure 1 shows the relationship between the user, IdP, FreeIPA, and protected web services.

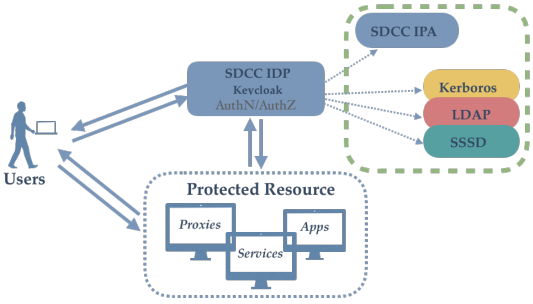

Figure 1. Single Sign On for SDCC users for local web services

4 Federation

Federation allows users to access SDCC web services without authenticating to the SDCC IdP. With federation, SDCC service providers can be configured to trust external identity providers, with IdPs in the InCommon Federation and major commercial identity providers like Google and Facebook as the likely choices. When accessing an SDCC web service, the user would be directed to an SDCC approved Id provider for authentication. The relationship between SDCC SPs and external IdPs is shown in figure 2. The drawback of this architecture is that each SDCC SPs must be configured to work with the external IdPs, either individually, or through a federation like InCommon.

SDCC service providers can be isolated from the complexities of external federations by utilize Keycloak's identity brokering capabilities. In this configuration Keycloak acts as an

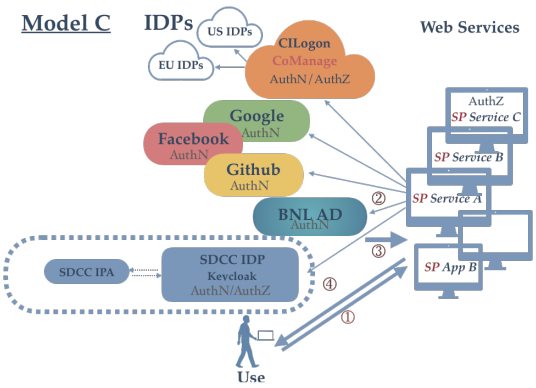

Figure 2. Configuration where SDCC Service Provides are federated with external IdP's or are part of a federation organization like InCommon

intermediary between the external IdP and the SDCC SP. Keycloak establishes the identity of the user for the SDCC service provider by interacting with the user's chosen identity provider. The identity broker effectively acts as an IdP for the SDCC service provider and acts as a service provider with respect to the external IdP. These relationships are shown in figure 3. As Keycloak interoperates with OIDC/OAuth2 and SAML2 based identity providers, the external IdP can be part of a SAML2 federation like InCommon, or an external OIDC provider like Google. This configuration is shown in figure 4.

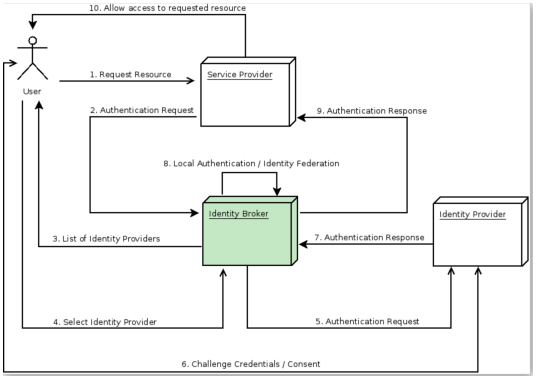

Figure 3. Interaction between an identity broker, service provider, identity provider and user when using identity brokering [15]

An additional simplication that has been applied to selected web services at the SDCC is shown in figure 5. In this configuration, the web services are not service providers, instead a web reverse proxy, sitting in front of the web services, acts as the service provider, gating access to the back end web services.

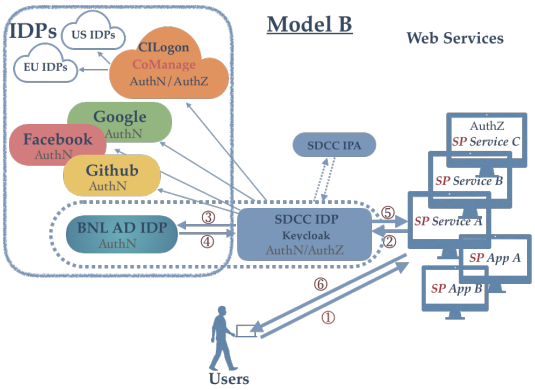

Figure 4. Federation using Keycloak's identity brokering capability

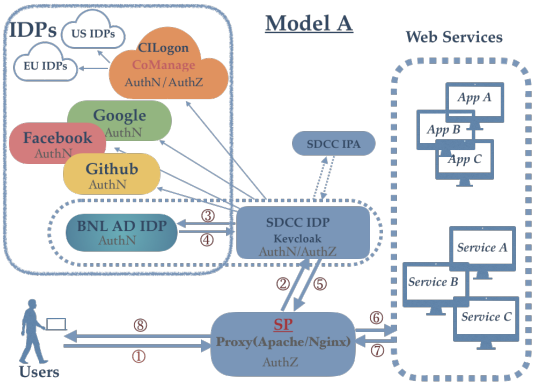

Figure 5. Web reverse proxy based service provider gating access to back end web applications.

5 Conclusion

Keycloak is a powerful solution, well suited in our facility, we are now looking into federating with InCommon and are considering hosting a CoManage instance to improve the authorization at the SDCC.

References

[1] Keycloak, https://www.keycloak.org
[2] CiLogon, https://www.cilogon.org
[3] CoManage,https://www.internet2.edu/products-services/trust-identity/comanage/
[4] InCommon, https://www.incommon.org
[5] FreeIPA, https://www.freeipa.org
[6] EduGain, https://www.edugain.org
[7] Project Jupyter, https://jupyter.org
[8] Invenio, https://invenio-software.org
[9] Mattermost, https://mattermost.com

[10] Gitea, https://gitea.io
[11] OpenID, https://openid.net
[12] OAuth, https://oauth.net
[13] SAML, https://wiki.oasis-open.org/security/FrontPage
[14] Kerberos Consortium, https://www.kerberos.org
[15] Red Hat, "Server Administration Guide: Red Hat Single Sign-On 7.4" [On line Documentation], Retrieved from https://access.redhat.com/documentation/en-us/red_hat_single_sign-on/7.4/html/server_administration_guide/identity_broker