NISO RP 27-2019

Recommended Practices for Improved Access to Institutionally-Provided Information Resources

Results from the Resource Access in the 21st Century (RA21) Project

A Recommended Practice of the
National Information Standards Organization and the STM Association

About NISO Recommended Practices

A NISO Recommended Practice is a recommended "best practice" or "guideline" for methods, materials, or practices in order to give guidance to the user. Such documents usually represent a leading edge, exceptional model, or proven industry practice. All elements of Recommended Practices are discretionary and may be used as stated or modified by the user to meet specific needs.

This recommended practice may be revised or withdrawn at any time. For current information on the status of this publication contact the NISO office or visit the NISO website (www.niso.org).

Published by
National Information Standards Organization (NISO)
3600 Clipper Mill Road, Suite 302
Baltimore, MD 21211
www.niso.org

ISBN: 978-1-937522-99-5

Contents

Foreword

About this Recommended Practice

This document details the findings from the Resource Access in the 21st Century initiative. It provides recommendations for using federated identity as an access model and improving the federated authentication user experience. The recommendations describe privacy preserving methods for access to scholarly content, and specific federated authentication workflow guidance to Service Providers (SPs), such as publishers, research infrastructure providers, and Identity Providers (IdPs) including Libraries and institutional Identity and Access Management systems, as well as any other groups involved in user engagement, such as Identity Federation Operators.

NISO Information Policy & Analysis Topic Committee Members

The Information Policy & Analysis Topic Committee had the following members at the time it approved this Recommended Practice:

Anne Campbell
EBSCO Information Systems

Michael Habib
Clarivate Analytics

Athena Hoeppner *(co-chair)*
University of Central Florida

Moon Kim
Ohio State University

Betty Landesman
Independent

Jack Maness *(co-chair)*
University of Denver

Stuart Maxwell
Scholarly IQ

Siôn Romaine
University of Washington

Christine Stamison
NERL

Kimberly Steinle
Duke University Press

Gavin Swanson
Cambridge University Press

RA21 Initiative Participants

The following individuals served on the Steering and Outreach committees, and on the User Experience and the Security and Privacy working groups, all of which developed this Recommended Practice:

RA21 Steering Committee

Chris Shillum *(co-chair)*
Elsevier

Ralph Youngen *(co-chair)*
American Chemical Society

Dan Ayala
ProQuest

Laird Barrett
Springer Nature

Peter Brantley
University of California, Davis

Todd Carpenter
NISO

Heather Flanagan
RA21 Academic Pilot Coordinator

Dave Flynn
EBSCO Information Services

Ann Gabriel
Elsevier

Gerry Grenier
IEEE

Don Hamparian
OCLC

Josh Howlett
JISC

Leif Johansson
SUNET

Phil Leahy
OpenAthens

Tim Lloyd *(co-chair, RA21 Outreach and Communications Committee)*
LibLynx

Helen Malone
GSK

Jay Neill
Wiley

Eefke Smit
International Association of STM Publishers

Jenny Walker
RA21 Corporate Pilot Coordinator

Julia Wallace
RA21 Project Director

Rich Wenger *(co-chair, RA21 Outreach and Communications Committee)*
Massachusetts Institute of Technology

Ann West
InCommon

RA21 User Experience Working Group

Serena Rosenhan *(Chair)*
ProQuest

Laird Barrett
Springer Nature

Rachel Bock
Wiley

Judy Chen
American Chemical Society

Jim Gavin
Wolters Kluwer

Matt Kleiderman
Copyright Clearance Center

Anna Rouben
ProQuest

Casey Schwartz
IEEE Xplore

Gergely Szabo
Elsevier

RA21 Security and Privacy Working Group

Reginald Zamora *(Chair)*
Wiley

Dan Ayala
ProQuest

Elias Balafoutis
Atypon

Laird Barrett
Springer Nature

Todd Carpenter
NISO

Diane Cogan
Ringgold

Meltem Dincer
Wiley

Paul Dixon
LibLynx

Ken Ferris
Taylor and Francis

Sari Frances
IEEE

Joe Greene
CAS

Lisa Janicke Hinchliffe
University of Illinois

Phil Leahy
OpenAthens

Scott McCarthy
ProQuest

Richard Northover
Elsevier

David Orr
OpenAthens

Christian Pruvost
Elsevier

Peter Reid
Bath Spa University

Andy Sanford
EBSCO

Chris Shillum
Elsevier

Will Simpson
ORCID

Adam Snook
OpenAthens

Heather Ruland Staines
Hypothes.is

Jos Westerbeke
University of Rotterdam

Ralph Youngen
American Chemical Society

RA21 Outreach and Communications Committee

Tim Lloyd *(Co-chair)*
Liblynx

Michelle Brewer
Wolters Kluwer

Rich Wenger *(Co-chair)*
Massachusetts Institute of Technology

Todd Carpenter
NISO

Jane Charlton
OpenAthens

John Felts
Coastal Carolina University

Sari Frances
IEEE

Ann Gabriel
Elsevier

Don Hamparian
OCLC

Robert Kelshian
American University

Laurence Lockton
University of Bath

Judy Luther
Informed Strategies

Matt McKay
International Association of STM Publishers

Jean Shipman
Elsevier

Emily Singley
Boston College

Raoul Teeuwen
SURFnet

Tracy Tolliver
University of Illinois

Lauren Tulloch
Copyright Clearance Center

Keith Webster
Carnegie Mellon University

Acknowledgements

The RA21 Initiative wishes to acknowledge those outside the formal project leadership who contributed to this effort, in particular all participants in the various pilot programs, working groups, extended user experience testing participants, and advisory committee members.

Trademarks, Services Marks

Section 1. Introduction

For several years, scholars have expressed increasing frustration with obtaining access to institutionally-provided information resources given changing work habits and the expectation of always-on connectivity from any location, at any time, from any device. Traditionally, the predominant method for authorizing access to scholarly information resources has been through IP address recognition. In the era where the ability to be "online" was solely facilitated through a physical connection to one's institutional network, IP address recognition worked well. However, given today's geographically dispersed workforce and variety of online connection options, institutional IP address recognition is no longer a satisfactory means for authorizing access to scholarly information resources. Today's researchers face a confusing diversity of options to facilitate remote access to scholarly information resources, such as VPN servers, proxy servers, various access code and registration schemes, and third-party software solutions. These difficulties in navigating today's remote access solutions impede research, frustrate users, and may encourage fully entitled users to resort to illicit, pirate websites. These sites in turn compromise the trustworthiness of the scholarly record and pose broad information security risks to institutions.

Several events in 2016 heightened industry awareness of the remote access problem. CNI's 2016 Report on the Authentication and Authorization Survey[1] revealed that even though IP address recognition was the most widely used method for authorizing access to scholarly content, more than half of the survey respondents reported having attribute-based authorization systems such as Shibboleth in use. The Copyright Clearance Center hosted a summit at the prompting of members of the Pharma Documentation Ring to call attention to problems with IP address recognition in the corporate space[2]. Finally, the International Association of STM Publishers (STM) and the National Information Standards Organization (NISO) brought together stakeholders from the publishing, library, software, and identity communities to form a broad initiative called Resource Access in the 21st Century (RA21) to address the issue[3].

RA21 conducted an assessment of the remote access capabilities in existence toward the end of 2016, and quickly surmised that SAML-based federated authentication held the most promise for providing a robust, scalable solution for remote access to scholarly content. Several factors influenced this assessment, including:

1) How IP address recognition can often allow the end user to remain anonymous to the service/content provider, and it is important to the library and researcher communities to preserve this option. SAML authentication protocol has the capability to preserve the privacy as well.
2) SAML being the only protocol that currently supports the concepts of federation and multilateral trust. Through established processes, entities can coordinate policies of trust and share configuration information through multilateral, as opposed to bilateral arrangements which quickly run into scaling difficulties.

3) The wide deployment of SAML-based federated authentication technology across research institutions and the scholarly publishing industry. More than 5,000 academic institutions and 11,000 service providers worldwide are members of academic identity federations and have SAML-based technology, such as Shibboleth, already deployed. In addition, most major publishers provided support for SAML, and a survey issued by RA21 revealed that many corporations used SAML-compliant identity management systems to manage their own workforces.

RA21 further discovered that despite the fundamental suitability and widespread deployment of SAML-based technologies, few users chose to leverage SAML federated authentication as a remote access solution. Several publishers reported findings from focus groups and usage analysis that indicated the reason why: when faced with the challenge of trying to navigate publisher websites to authenticate via SAML, users abandon those efforts and seek alternate means for accessing information resources. In particular, users find the process for locating their home institution (their Identity Provider, or IdP) to be onerous; users are often faced with disparate and complex implementations across publisher sites that frequently employ inconsistent terminology and visual elements.

RA21 quickly determined that for SAML-based federated authentication to become more widely used, these significant issues in user experience must be addressed. Creation of a streamlined user experience for federated authentication therefore became the central goal of the RA21 initiative. An improved user experience was prototyped and thoroughly tested with practicing researchers in both academic and corporate settings, incorporating iterative improvements based upon findings from those user studies. The technical feasibility of implementing the user experience was validated by pilots conducted by RA21 participants. Information about the methodologies followed are included in the Appendix.

This document details RA21's findings and provides recommendations for improving the federated authentication user experience. The recommendations offer guidance to Service Providers (SPs) (e.g., publishers, research collaboration platforms, and research infrastructure providers), IdPs (e.g., academic institutions or corporations), as well as any other groups involved in user engagement such as Identity Federation Operators.

The following elements emerged as the key components of an improved solution:

1. A **common UI element** (e.g., a button) that SPs may add to their sites to invite users to authenticate with a federated identity or initiate the IdP discovery process.
2. An **improved, search-based IdP discovery experience** which makes use of enhanced IdP metadata to enable reliable selection of the appropriate IdP using institution name or email domain.
3. A **cross-domain IdP persistence service** which enables a user's previous choice of IdP to be remembered by their browser through web sites that use the IdP persistence service, thus decreasing the frequency with which the user has to choose their IdP.

1.1. Guiding Principles

RA21 developed the following Guiding Principles to structure the overarching effort:

1. The user experience for researchers will be as seamless as possible, intuitive and consistent across varied systems, and meet evolving user expectations.
2. The solution will work effectively regardless of the researcher's starting point, physical location, or preferred device.
3. The solution will be consistent with emerging privacy regulations, will avoid requiring researchers to create yet another ID, and will achieve an optimal balance between security and usability.
4. The system will achieve end-to-end traceability, providing a robust, widely adopted mechanism for detecting fraud that occurs at institutions, vendor systems, and publishing platforms.
5. The IdP will not be burdened with administrative work or expenses related to implementation and maintenance. The implementation plan should allow for gradual transition and account for different levels of technical and organizational maturity in participating institutions.

These principles led to specific requirements in the areas of privacy, security, user experience, and governance.

1.1.1. Privacy Principles

An area of concern for many when it comes to using federated authentication to access resources is whether this access model impacts the privacy of the user. A SAML-based workflow has the potential to share arbitrary information that the IdP holds about the user with the SP. When compared with IP authorization models, in which limited information is transmitted about the user, this is an understandable concern, although in some circumstances IP addresses themselves are considered to be personally identifying.

SAML authentication has the capability to effectively preserve the privacy of the end user through the exchange of anonymous assertions such as the user's membership of an authorized user community. Identity federations have long-standing mechanisms in place to set and enforce policies, including those regarding privacy, through the federation membership agreements which all participants must sign. In addition, norms have been established for certain classes of resources through the definition of entity categories and attribute bundles. Finally, end users may be given control over what information is released by their IdP through the use of attribute release consent user interfaces.

To ensure application of these mechanisms in the use cases addressed by RA21, RA21 endorses the guidance provided in the GEANT Data Protection Code of Conduct[4].

> "The Data protection Code of Conduct describes an approach to meet the requirements of the EU Data Protection Directive in federated identity management.

> The Data Protection Code of Conduct defines behavioral rules for SPs which want to receive user attributes from the Identity Providers managed by the Home Organisations." -- Introduction to the Data Protection Code of Conduct

Further, we propose that the code of conduct be adopted globally for all users regardless of the location of their home institution or the location of the service they are accessing.

More detail on how the Data Protection Code of Conduct should be implemented in different scenarios is discussed below in the Recommendations.

1.1.2. Security Principles

Coupled with the concept of user privacy is the goal of information security. In addition to the concept of data minimization enshrined in the Data Protection Code of Practice, the process by which information is requested and stored must also support the confidentiality, integrity, and availability of the service.

To that end, a thorough analysis of the technologies and architecture proposed for RA21 was carried out according to the STRIDE model[5], a best practice for security evaluations. This analysis is publicly available on the RA21 website: WAYF Cloud and P3W Security & Privacy Recommendations

All proposed technologies and architectures reviewed during this evaluation demonstrated a strong security profile. Minimal information was requested or stored, resulting in a very low risk environment.[6,7]

1.1.3. User Experience Principles

To achieve an authentication experience that is "as seamless as possible, intuitive and consistent across varied systems" RA21 adopted guiding user experience principles focused on removing friction and reducing cognitive load* at every opportunity in the user workflow. Specific guiding principles include:

- Reducing the number of steps required for federated authentication.
- Limiting the choices presented to the user at a given time (buttons, input fields, links, options in lists, etc.).
- Helping users find the next click in the access workflow by creating clear and simple calls to action (CTAs).
- Using best practice feedback patterns, such as typeahead and suggestions, to give the user confidence to proceed.

* The amount of mental resources required to operate a system (Whitenton) https://www.nngroup.com/articles/minimize-cognitive-load/

1.1.4. Governance Principles

In order to ensure consideration of a wide set of viewpoints and support broad stakeholder engagement, the RA21 initiative employed an open approach to governance. A Steering Committee was established with membership from all stakeholder groups; a number of working groups were established with open participation from anyone in the community with the willingness and experience to participate, and many stakeholder update meetings and numerous conference presentations were conducted through the life of the project. All final outputs are freely available on the RA21 website.

1.2. Terms and Definitions

Authentication	The process of verifying the identity of a user, process or device, often, but by no means exclusively, through the use of a username and password.
Authorization	The process of verifying against a set of access controls whether an account is authorized to access a given service or resource.
Entity Category[8]	Entity categories group federation entities that share common criteria, such as SPs that are operated for the purpose of supporting research and scholarship interaction. The intent is that all entities in a given entity category are obliged to conform to the characteristics set out in the definition of that category, thus allowing IdPs means of classifying SPs without having to research each one.
Federated Authentication	The mechanism by which an identity provider, such as a home organization, indicates to one of more service providers that the user has been authenticated and may be authorized by the service provider to access relevant resources.
Federated Identity	A digital identity which is asserted by one system (an identity provider) which may be consumed by other systems (service providers) by means of federated authentication.
Federation	A federation is an association of organizations that agree to exchange information as appropriate about their users and resources in order to enable collaborations and transactions such as user authentication.
Identity Provider (IdP)	An organization that manages digital identities and issues authentication assertions and potentially other attributes to Service Providers.

Identity Provider (IdP) Persistence	The storage and re-use of a previous IdP choice made during an identity provider discovery process.
IP address-based Authorization	A method where an SP and a home organization have agreed that every request coming from a range of network/Internet Protocol (IP) addresses associated with the home organization should be authorized for the services provided by the SP.
Multilateral Federated Authentication	A federated authentication workflow within a multilateral identity federation.
Multilateral Identity Federation	A form of identity federation where a trusted third-party registers and publishes all entity metadata to all members, preventing the need for bilateral agreements between an IdP and SP.
Security Assertion Markup Language (SAML)[9]	An open standard for exchanging authentication and authorization data between entities, thus enabling single sign on (SSO) and federated authentication and authorization workflows. Many interoperable open source and commercial implementations of SAML are available.
Service Provider (SP)	An organization that makes online resources available to users based in part on information, in particular authentication assertions, from IdPs.
Single Sign On (SSO)	The ability of a user to access multiple discrete systems or sets of resources with a single set of access credentials. This is often a precursor to supporting Federated Authentication.
Web Storage[10]	Where web applications can store data locally within the user's browser. Before HTML5, application data had to be stored in cookies, included in every server request. Use of web storage (sometimes known as browser local storage) prevents sending the data to server with each server calls (which is what cookies do).

Section 2. Recommendations

RA21 has the following recommendations as outlined in the subsequent subsections:

2.1. Adopt Multilateral Federated Authentication

The foundational recommendation of the RA21 project is to strongly encourage the use of multilateral federated authentication for all inter-organizational collaboration and access management. Such technologies offer the benefits of:

- Decoupling of access management from the physical network configuration of organizations, simplifying administration, reflecting the reality of a mobile workforce and anticipating the growing adoption of cloud-based secure Internet Service Providers.
- Offering users the convenience of Single Sign On, obviating the need to establish and remember separate credentials for every service they wish to personalize.
- Offering a privacy-preserving method of authentication and authorization.
- Creating a uniform mechanism for access to resources which works in the same way whether or not the user is connected to the institutional network, thus improving user understanding and adoption.
- Providing for more granular access control by the institution's Identity Provider over that of IP address recognition should it be needed.

The predominant technology used in multilateral Federated Authentication today is SAML, thus these recommendations focus on SAML-based platforms. However, in the future, a different underlying technology may prevail such as OpenID Connect[11]. The RA21 recommendations are expected to remain applicable to any future shift in underlying technology.

These recommendations do not dictate the use of any particular software platform, it should be noted that there are a variety of platforms available which support SAML-based Federated Authentication, including open source, free, and paid services.

2.1.1. Employ appropriate authentication mechanisms for specific use cases

The use of federated access does not imply the exclusive use of login credentials (usernames and passwords) as the authentication mechanism. In merely implies that responsibility for authenticating the user is delegated by the Service Provider to the Identity Provider.

Identity Providers are free to use a variety of different authentication methods for different classes of users and different use cases. Whilst login credentials are most often used for individuals affiliated with an institution, other authentication mechanisms are likely to be better suited to other use cases, for example:

- Many agreements for access to scholarly information resources provide for onsite access to members of the public (known as walk-ins). In this case, the institution may choose to use device- or network-based authentication scheme such certificates installed on library workstations or the network perimeter, or locally-managed IP address validation to support these users. Another alternative would be to provide guest accounts or one-time use codes on the local authentication system.
- In clinical care settings where rapid access to information is often necessary, institutions may choose to leverage existing authentication mechanisms, for example proximity cards to manage access to these resources, or again provision workstations using certificates of locally managed IP address verification.

Further work is required with the specific communities concerned to test and validate the most appropriate authentication mechanism for each use case.

2.2. Establish Multilateral Identity Federations Where They Do Not Exist

In both corporate and academic settings, organizations are using SAML-enabled tools to support local and, increasingly, point-to-point bilateral Federated Authentication. In academia, many institutions take this a step further by participating in a SAML-based identity federation which enables mutual trust to be established among a set of collaborating organizations.

RA21 recommends that Identity Federations are established and adopted by communities where they currently do not exist, such as corporate consumers of information resources.

Multilateral identity federations support broader and more rapid adoption of federated authentication by establishing common policies and standards, establishing mutual trust

without the need for bilateral agreements and reducing the need for point-to-point configuration through the distribution of common metadata.

In addition, while common at the institution level, the use of SAML-based technologies for authentication and access control does not always permeate to all departments at an institution, for example the campus library at academic institutions. RA21 recommends that the groups responsible for operating IdPs within institutions do more to educate their communities about the benefits of Federated Authentication and do more to promote its broad adoption.

2.3. Ensure that Privacy is Preserved while Enabling Convenient SSO and Granular Authorization

The GEANT Data Protection Code of Conduct enshrines the principles of Legal compliance, Purpose limitation and Data minimization with regards to user's personal information received by SPs in federated authentication transactions.

To ensure application of these mechanisms in the use cases addressed by RA21, RA21 endorses the guidance provided in the GEANT Data Protection Code of Conduct[12]. While this endorsement focuses on v1 of the Data Protection Code of Conduct as the currently approved version, SPs should continue to follow guidance as it evolves in approved successors. Further, we propose that the code of conduct be adopted globally for all users of RA21 regardless of the location of their home institution or the location of the service they are accessing.

Identity federations have long-standing mechanisms in place to set and enforce policies, including those regarding privacy, through the federation membership agreements which all participants must sign. In addition, norms have been established for certain classes of resources through the definition of entity categories and attribute bundles.

For the research collaboration use case, for example the use of tools and services such as wikis, blogs, project and grant management tools that require some personal information about users to work effectively, the Research and Scholarship entity category has already been established. RA21 proposes that this Entity Category should not be used for access to scholarly information resources. Instead, RA21 recommends that new entity categories are established to cover these use cases as follows:

Use Case	Description	Purpose	Attributes
1a. Access to a scholarly information resource that doesn't provide	Unless the SP has a specific, contractual agreement with an IdP, the IdP should only send	Assert that the user is a member of the institution's authorized user community for the resources being accessed.	Anonymous assertion (e.g., eduPersonEntitlement[13])

personalization by a user from an academic institution†	anonymous attributes to the SP.		
1b. Access to a scholarly information resource that provides personalization by a user at an academic institutions*	Unless the SP has a specific, contractual agreement with an IdP, the IdP should only send anonymous and pseudonymous attributes to the SP.	Assert that the user is a member of the institution's authorized user community for the resources being accessed.	Anonymous assertion (e.g., eduPersonEntitlement)
		Enable SSO to any personalized features the resource may offer using institutional credentials.	Pseudonymous pairwise user identifier (e.g., eduPersonTargetedID‡, pairwise-id[14]
2. Access to scholarly information resources by users at a corporation	In addition to the attributes provided by academic libraries, an additional attribute may be sent to the SP to support specific granular usage analysis or charge back requirements that the IdP may have.	Assert that the user is a member of the institution's authorized user community for the resources being accessed.	Anonymous assertion (e.g., eduPersonEntitlement)
		Enable SSO to any personalized features the resource may offer using institutional credentials.	Pseudonymous pairwise user identifier, (e.g., eduPersonTargetedID, pairwise-id)
		Enable granular usage measurement and charge back of costs at the IdPs request.	A repeating, opaque reporting group code meaningful only to the IdP (note: a standardized attribute is being established for this purpose).

Service providers that do support personalization should allow users to create local accounts which are linked to their institutional SSO credentials using the supplied pseudonymous pairwise identifier, thus obviating the need to maintain separate sign-in

† Some SPs that support optional personalization, i.e., they offer some degree of functionality to anonymous users while other features are available only to users who choose to create a personalized account. SPs in this category must provide the anonymous feature set to users who authenticate and only provide an anonymous entitlement attribute.

‡ Note that while the eduPerson specification includes a number of generic identifier attribute types such as eduPersonTargetedID, it is increasingly common for individual security protocols such as OpenID Connect and SAML to define their own "standard" subject identifiers and related functionality. In some cases (e.g., SAML) this material has been explicitly informed by, and is a reaction to, problems or limitations arising from the application of the eduPerson-defined identifiers to federated authentication. It is advisable to defer to a particular protocol's specifications to understand what constitutes best practice in that particular context.

credentials for each service. Pseudonymous identifiers should not be used by SPs for other purposes without explicit user consent.

If services require information from users in addition to the attributes released by their IdP in order to create local accounts, then this should be gathered via an explicit, opt-in registration process.

In the case that the IdP sends more attributes than defined by these entity categories, the SP must not use, collect, or store that data in any way, under any circumstances.

Further, RA21 also recommends that before personally identifying attributes, including pseudonymous identifiers, are sent by an IdP to an SP, the IdP obtains consent from the user via a consent-informed attribute release process.

2.4. Improve the User Experience of Identity Provider Discovery

The technological capabilities of a multilateral federated authentication workflow are powerful, but the implementation across SPs is inconsistent and often difficult for the user to follow. Building best practices around the user experience is one of the key outcomes of RA21.

The user experience associated with a research workflow often involves the use of a search tool that provides the user with links to multiple SPs (e.g., a literature search nearly often produces links to documents found on multiple publisher websites). As users follow such links, their first interaction with each SP is in an unauthenticated state, whereby the user must invoke an IdP discovery process (a.k.a., "Where Are You From") to search for their home institution.

This federated authentication pattern is known as SP-initiated authentication, because the user must initiate the authentication process from the SP. An alternative pattern, known as IdP-initiated authentication also exists but is far less common in scholarly research. IdP-initiated federation requires the user to first authenticate with an institutional portal, from which the user can then seamlessly access an SP using a specially constructed URL (a "WAYFless URL") that bypasses the need for IdP discovery on the SP. RA21 endorses the use of WAYFless URLs wherever applicable, but the recommendations made in this document apply to the more common authentication pattern that requires IdP discovery.

Solving poor IdP discovery processes requires SPs to look at and understand the user journey holistically, specifically recognizing that a single user often has multiple interactions with different SPs in a research session. They are typically not interacting with only one SP or having just one experience.

More detail on each of the user experience recommendations follows, but to summarize, RA21 recommends the following best practices for SPs:

1. Implement a new call to action (e.g., a button) for federated authentication on all relevant SP pages.
2. Present all access choices to the user in close visual proximity and in hierarchical order.
3. Place the primary call to action button in a location that does not require the user to scroll.
4. Display the user's previous choice of IdP.

The best practice recommendations in this document are intended to minimize cognitive load and friction for the user and are based on findings of multiple user studies conducted with both academic and corporate users[§].

2.4.1. Overview of the Expected User Experience

This section describes the expected experience for a user in two scenarios: 1) a user who has not previously authenticated and whose institution (the IdP) is not known. 2) a user whose institution is known (remembered or pre-defined). Scenarios 2b and 2c also describe how the experience may be further streamlined when a dynamic version of the call to action is used.

2.4.1.1. Scenario 1: Identity Provider is Not Known

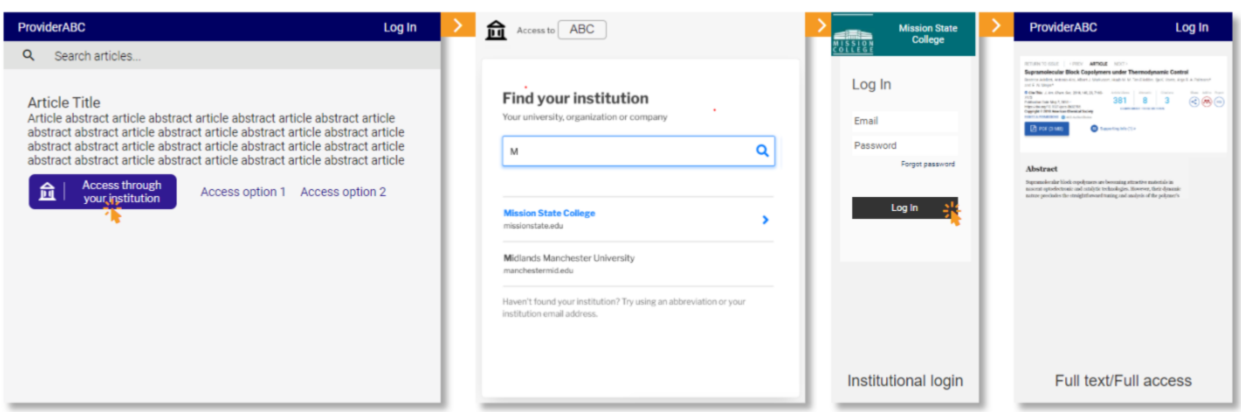

Figure 1. *Initial experience; institution is not yet known*

1. Users navigate to an article/other resource page on an SP site. They click the familiar "Access through your institution" button which opens an IdP discovery page.
2. On the IdP discovery page, users search for and select their identity-providing institution.

---

§ Usability studies were conducted over the two-year project to develop and validate the design and best practice recommendations in this document. See Appendix for a summary of key findings or RA21 User Research Methods and Findings for a detailed report including research methods and corresponding findings.

3. On the institution login page, users authenticate using their institutional credentials and access the full article/other resource, if available through their institution.

2.4.1.2. Scenario 2a: Identity Provider is Known (Static Button)

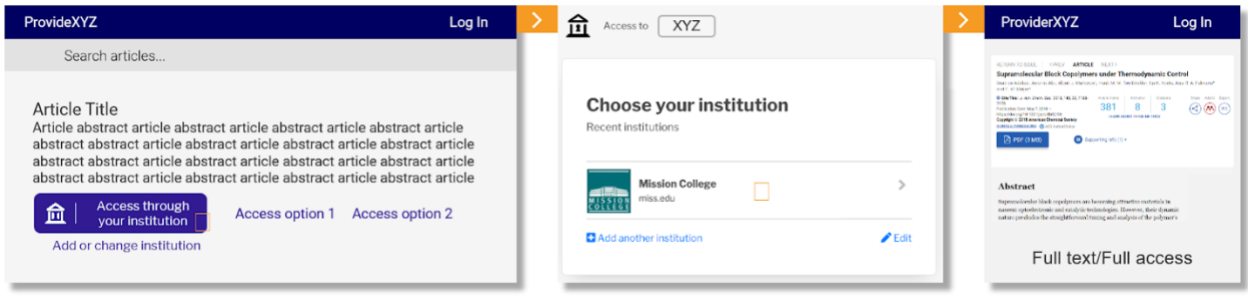

Figure 2a. *Subsequent experience; institution is remembered in IdP discovery*

1. Users navigate to an article/other resource page on a SP site. They click the "Access through your institution" button which opens an IdP discovery page.
2. On the IdP discovery page, users see their previously used institution(s) and select the IdP they wish to use; they don't need to search for their institution again. Users also have the opportunity to find another institution (e.g., if they have multiple IdPs) or to remove or change the previously used institution.
3. If users still have an active session with their IdP, they should be able to by-pass the institution login step and gain immediate access to the full article/other resource, if available through their institution. In this scenario, when users access a second SP page, that SP uses the known institution to direct authentication calls to the correct IdP, without requiring the user to re-identify their institution to the new provider.
4. If users no longer have an active session with their IdP, they will be redirected to their institution login page and asked to re-authenticate.

2.4.1.3. Scenario 2b: Identity Provider is Known (Dynamic Button - Recommended)

Figure 2b. *Subsequent experience; institution is remembered and shown in call to action*

1. Users navigate to an article/other resource page on an SP site. They see the dynamic button with the label: "Access through <your institution>" where <your institution> is the name of their previously used IdP. Note: Users will also have the opportunity to remove institutions or find another institution. (See Section 2.4.3.2. for more detail on changing or finding another institution.)
2. If users still have an active session with their IdP, they should be able to skip the institution login page and gain immediate access to the full article/other source, if available through their institution.
3. If users no longer have an active session with their IdP, they will be redirected to their institution login page and asked to re-authenticate.

2.4.1.4. Scenario 2c: Identity Provider is Pre-Defined (Dynamic Button – Recommended for Managed Environments)

To ensure users always have the known or "remembered" experience, identity providers could use one of the following approaches:

1. For a company or site with centralized desktop/laptop management policies and tools, it should be possible to push an update to managed desktops/laptops that would pre-populate the company or site name to be used in the federated access button.
2. An alternate option is to make a specially crafted link available to site users through trusted IT channels. Clicking on the link would pre-populate the "Access through your institution" call to action with the identity producer name.
3.

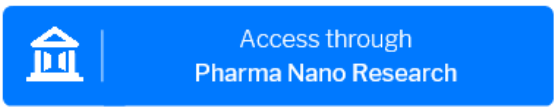

Figure 2c. *Example of pre-defined Identity Provider in call to action for managed environments*

2.4.2. Detailed Design Recommendations

2.4.2.1. Adding a Call to Action (Access Button) to Service Provider Pages

When adding the call to action button for federated authentication to SP pages, the following best practices are recommended.

✓ DO	✕ DON'T
✓Group all access options together and show the fewest possible number of choices. Users can follow the right path faster when they can scan and differentiate available options at once. Access options might include: Access through your institution, Log-in (local account), Rent, or Purchase.	✕ Spread access options across different areas of the page. (Note: it is still appropriate to offer options such as SP login in traditional locations such as the page header.) ✕ Give users a false sense of access by providing a PDF icon, Download, or View Article option on the article page when users must first authenticate before getting the full text.
✓ Clearly identify and communicate a primary access method. Users success improves when they can clearly identify a single primary action. In many cases, federated access is the user's best means of successfully accessing content and should be offered as the primary call to action. Other access options such as local account log in, rent or purchase, may be presented, but should not be weighted equally. For example, the primary access option should be a button and other choices might be links.	✕ Give secondary access options (e.g., Purchase) equivalent visual weight as the "Access through your institution" button.
✓ Place grouped access options in the first view of the page. Getting full access to content is a top goal for users of SP pages. Users can rapidly identify the right path when the primary	✕ Place "Access through your institution" in a location that requires the user to scroll to find it.

action is visible on the first page view. Whenever possible, avoid making the user scroll to locate access options. (Note: it is still appropriate to display options such as SP login in traditional locations such as the page header.) Design responsive and mobile experiences with this goal in mind. Best practice considerations include: • Identify the main view ports used by your user base • Considerations for mobile and responsive experiences: Identify the main view ports for your product and place the CTAs to that first view. This is still a recommended best practice for mobile or responsive, smaller view ports.	✕ Place "Access through your institution" on an overlay or separate page.

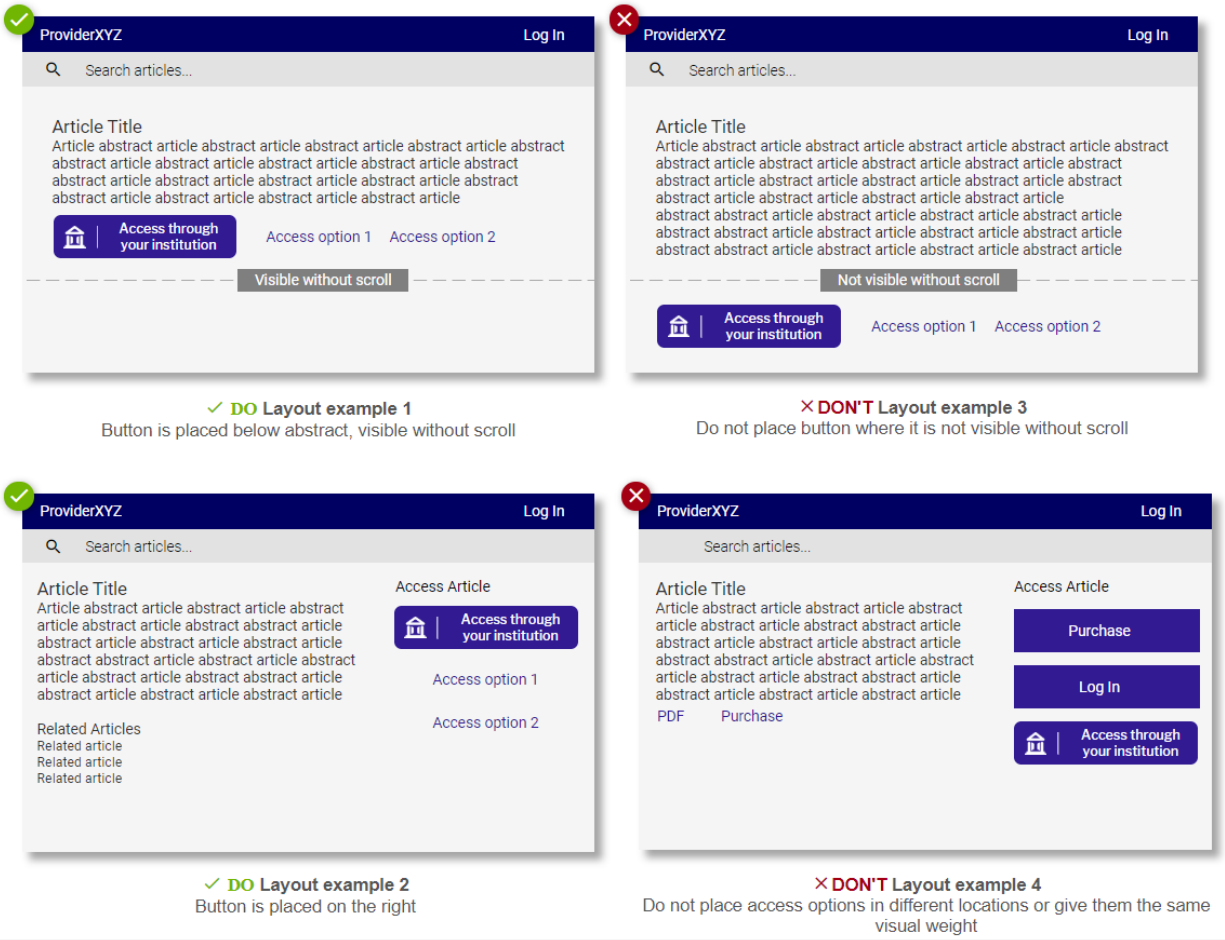

Figure 3a. *DOs and DON'Ts when adding the access button to a Service Provider page*

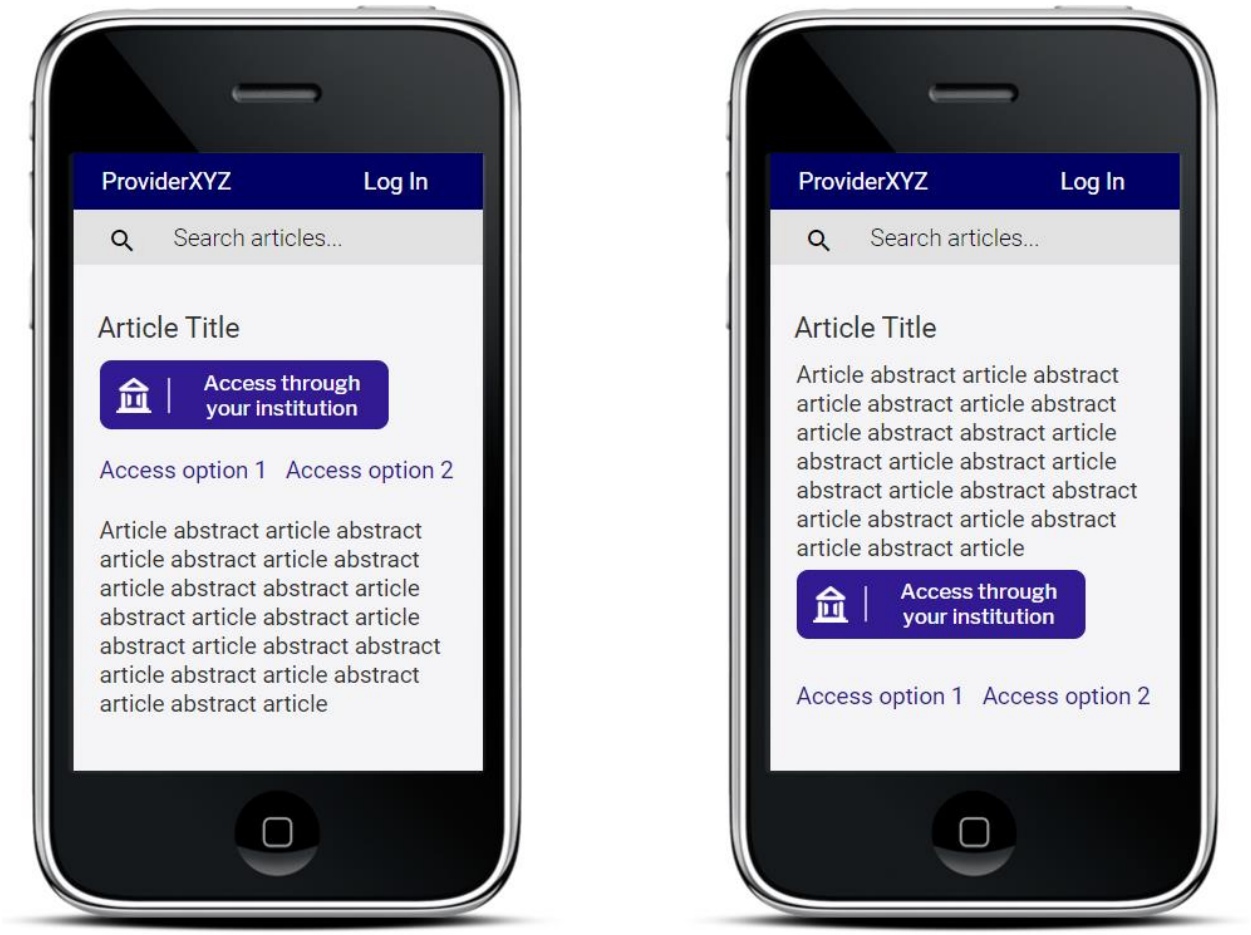

Figure 3b. *Apply similar best practice recommendations when adding the access button to a mobile or responsive Service Provider page*

2.4.2.2. Access Button Guidelines

When adding the call to action button for federated authentication ("Access through . . . ") to SP pages, the following best practices are recommended.

Figure 4. *Federated access button*

✓ DO	✕ DON'T
✓ Use a consistent presentation of the "Access through your institution" button. Repetition of the language and visual styling of the access button including icon, text, button, contrast and behavior will significantly reduce friction and build recognition and confidence for users. The specific labeling and design of the call to action were informed by best practice models and multiple rounds of iterative user testing. (See Appendix for additional detail.)	✕ Change "Access through your institution" text label. ✕ Use low contrast for the text and button background. ✕ Use a link for "Access through your institution" instead of a button when it is the primary call to action for access.

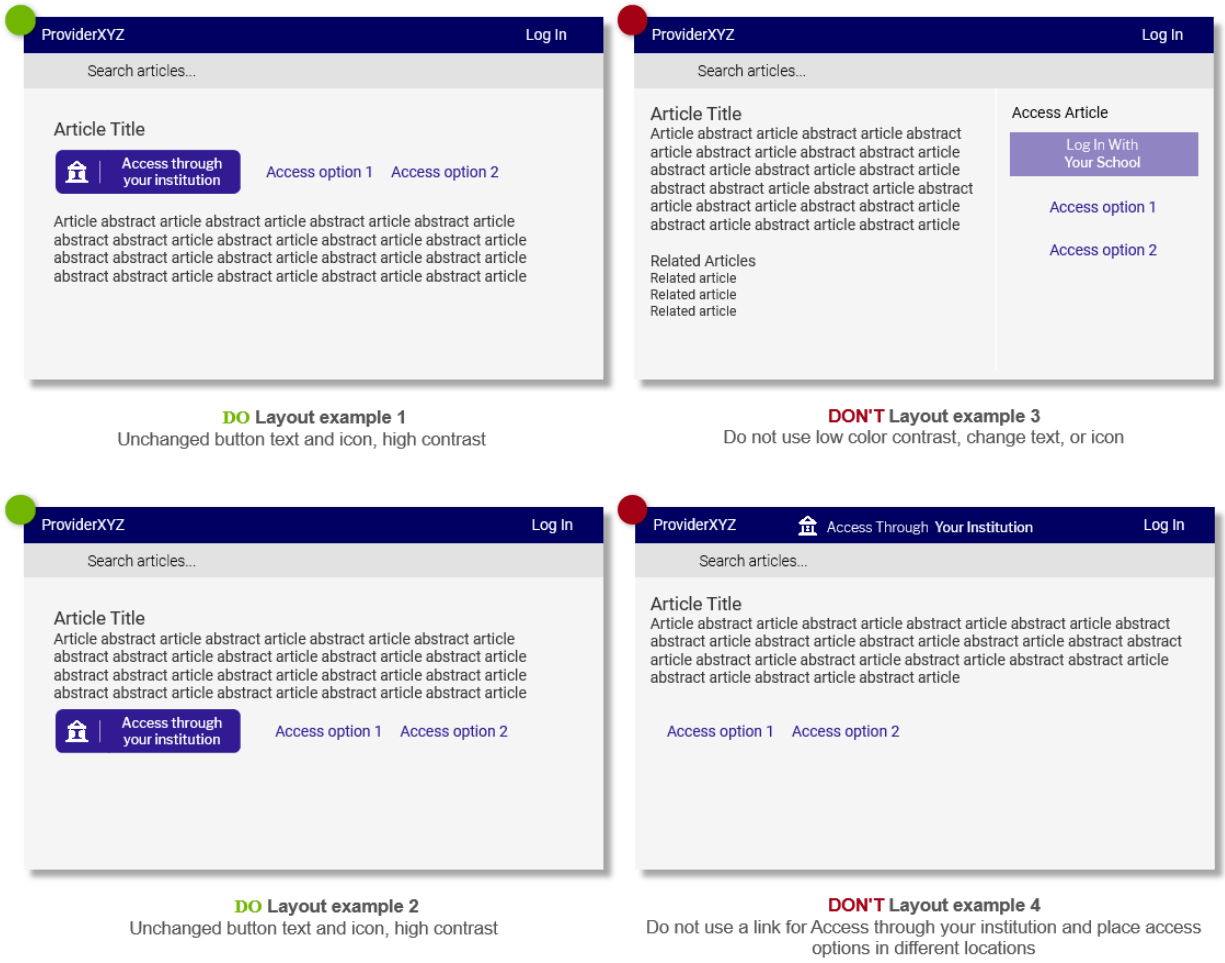

Figure 5. *DOs and DON'Ts when implementing the access button*

2.4.3. Implementation Recommendations for the Visual Elements of Identity Provider Call to Action

This section discusses options for achieving the above recommendations through different levels of implementation complexity for the SP. This section only discusses implementation considerations. Detailed implementation specifications and instructions will be provided in a separate document.

2.4.3.1. Static Access Button

The simplest implementation is for the SP to place a button matching the guidelines on its pages and link that button to an IdP discovery service, either a centralized service offered by a federation or other organization, or a customized service developed locally by the SP. Please refer to 2.4.1.1. Scenario 1: Identity Provider is Not Known.

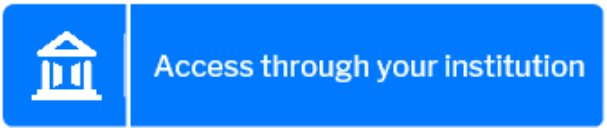

Figure 6. *Example of static button implementation*

Implementation requirements[**]:

- Button design specifications, logo file
- IdP discovery service link.

Expected experience:

- Users recognize and click on the "Access through your institution" button on the SP site.
- Users are taken to an IdP discovery service where they look up and select their institution. That institution is remembered for the subsequent uses.
- Users are taken to their familiar institutional login page, authenticate as they normally do, and are then taken to the full version of the SP content (if that content is provided by their chosen institution).
- On subsequent clicks of "Access Through Your Institution", users will see a screen asking them to confirm the remembered institution (see Figure 14 in Section 2.4.4.1.)

2.4.3.2. Dynamic Access Button

A dynamic button is provided by the Cross-domain IdP Persistence Service (see Section 2.5.). This allows the remembered institution to be displayed as the call to action on an SP page, which saves the user from going through the steps of finding their institution every time they visit a new SP. Users are only presented with the "Access through your institution" wording

[**] Detailed specifications will be provided in separate implementation documents.

in the "cold-start state," when the institution is not known (e.g., browser cache has been cleared, new/different device, etc.).

Returning users are presented with a button that is populated with their most recently used institution (e.g., "Access through Southfield University").

- The target of the button in the "cold start" state is an IdP discovery service (centrally offered, or locally offered by the SP).
- The target of button in the remembered state is the remembered IdP's login page.

Please refer to Scenario 2b: Institution is Known (Dynamic Button).

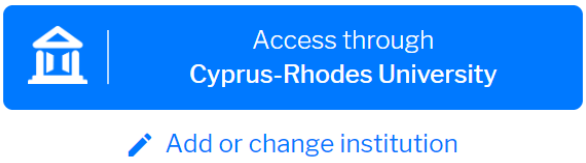

Figure 7. *Dynamic access button: the most recently used institution is displayed as the call to action on the SP page*

With the Dynamic access button, users also have the option to remove a remembered institution or to use more than one institution. This is useful for users with multiple or changing affiliations.

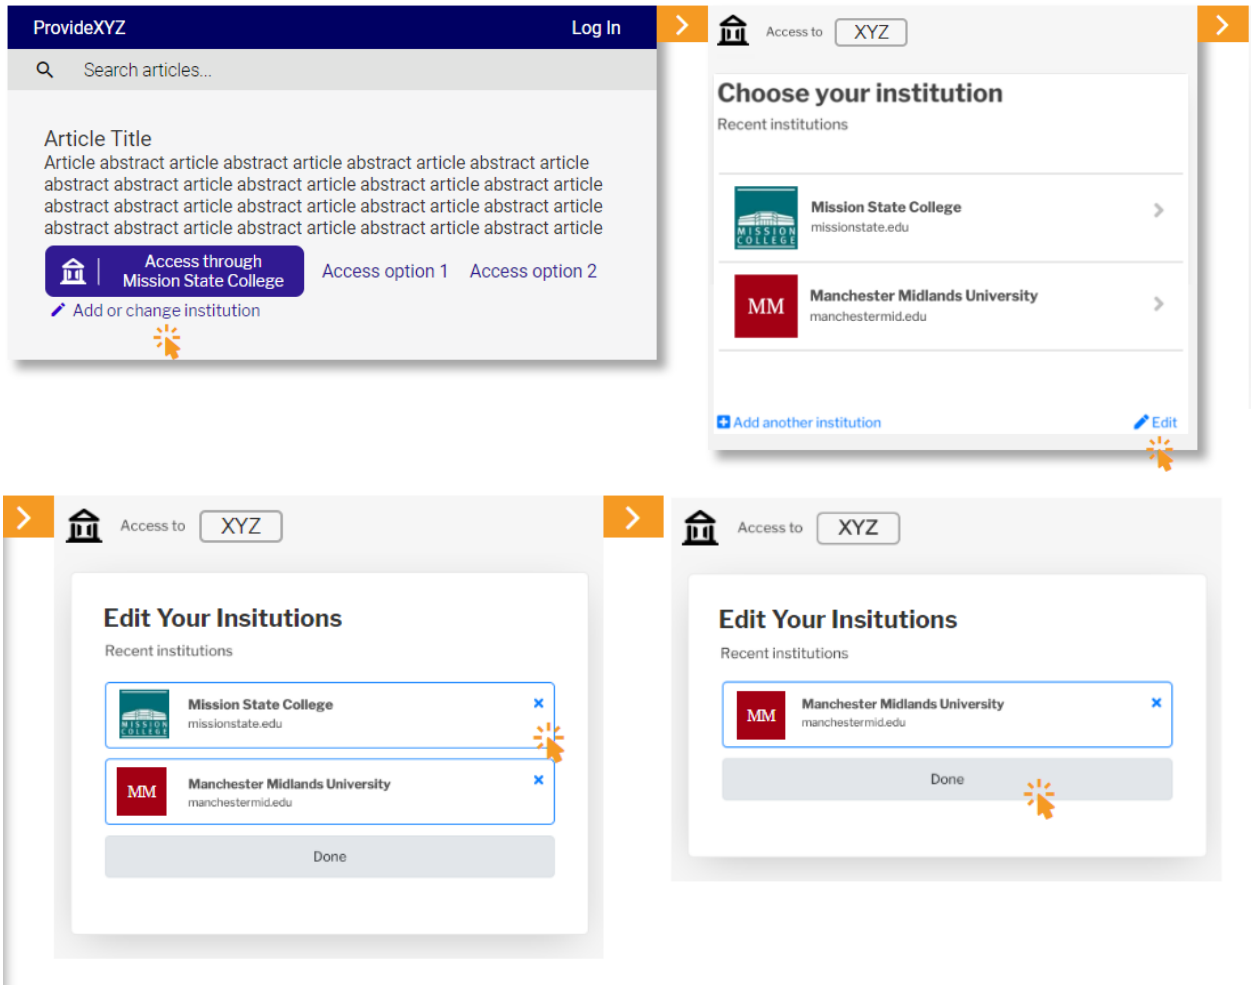

Figure 8. *Dynamic access button allows user to remove a previously used institution or to use more than one institution.*

Implementation requirements ††:

- Single line of JavaScript to create the i-frame button, including any parameters to pass to the button to control language and presentation.
- For managed environments, where desktops and laptops are provided by the institution, there will also be a mechanism allowing IT to systematically install the site's login location, so the identity provider name displays in the call to action without the user have to go through the IdP discovery flow.
-

Expected experience:

- On SP pages users are presented with the "Access through your institution" when there is no known institution (e.g., first time visiting any SP, browser cache has been cleared, new/different device, user has disallowed JavaScript, etc.)

†† Detailed specifications will be provided in separate implementation documents.

- Returning users are presented with the button populated with their most recently used institution (e.g., "Access through Miami University"). This saves the user from going through the steps of finding their institution every time they visit a new SP.
- IdP information should be stored in local browser storage after IdP search, but prior to visiting the IdP login page.
- When a user can authenticate but is not authorized to view the desired full text, provide clear messaging of the reasons the full text is not available.

2.4.4. Implementation Recommendations for the Visual Elements of Identity Provider Discovery

2.4.4.1. Default Identity Provider Discovery Service

It is recommended that a central, cross-federation IdP discovery service is made available for SPs to integrate with, based on a common metadata source such as eduGAIN. This provides a low-effort implementation route for SPs who do not need, or do not have the resources, to develop their own custom IdP discovery service.

In order to offer a good user experience, all IdP discovery services should ensure that it makes use of clear metadata and standardized logos, as described in Section 2.5.

2.4.4.2. Service Provider Custom Identity Provider Discovery Service

SPs may also wish to create their own customized IdP discovery service in order to incorporate their own local metadata from bilateral IdP relationships and/or integrate it with other authentication models (such as local accounts). For example, the SP can present a single point of entry for the user to choose their institution, and then offer an appropriate response depending on the authentication mechanism used by that institution.

Those doing so are encouraged to follow the best practices for IdP discovery outlined below.

2.4.4.3. Streamlined Institution Search

IdP discovery search interfaces should adhere to the following recommendations when presenting the initial search UI.

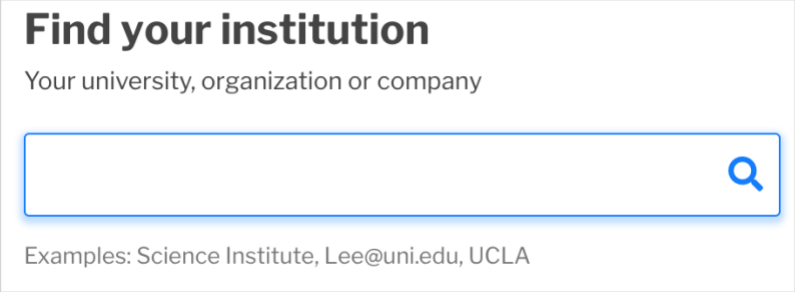

Figure 9. *IdP discovery service search field.*

- Provide clear instruction on what to search for by labeling the search box: "Find your institution". Including descriptive text below the heading is also recommended.
- Provide labeling that is available to assistive technology. Users need to be aware of control labels, headings, tip, and other content using screen readers. We recommend complying with the most current version of WCAG Accessibility guidelines at the AA level: WCAG 2.1 or its subsequent versions. [15]
- On page load, bring keyboard focus into the search field so that users can start typing and searching without additional hand movements or clicks.
- Provide type-ahead in the search field; users expect to see results when they type.
- Provide support for searching institution abbreviations, e.g., UCLA.
- Providing support for deriving institution from entered email domain is optional.
- User testing showed that most users are unlikely to enter their email address as a means of identifying their IdP.‡‡ However, email address is included as an option in the IdP discovery process because in some situations (for example mergers and acquisitions of corporations) institution name is ambiguous.
    - Email addresses provided for IdP look up should not be retained; only the domain portion of the email address should be used to conduct the search.

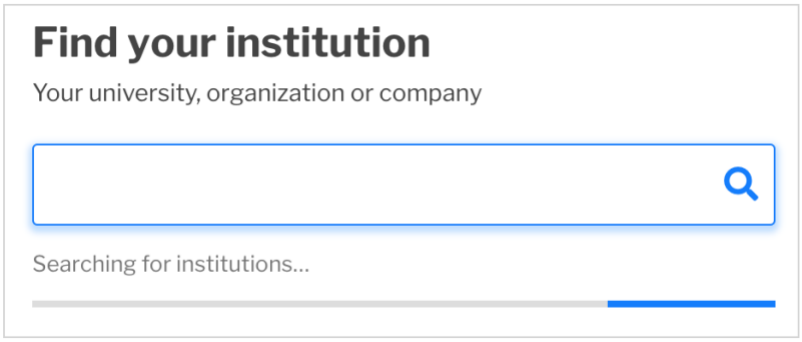

Figure 10. *Continuous progress indicator.*

- Provide a continuous progress indicator (Figure 9) if search results are not displayed immediately so that users know that something is happening.

‡‡ See Section A.3 for a summary of the user research conducted and compiled as part of the RA21 initiative.

- Provide progress updates that are available to assistive technology. Users of screen readers need to be aware of dynamically changing content.

2.4.4.4. Search Results Display

IdP discovery search interfaces should adhere to the following recommendations when displaying the results of institution searches.

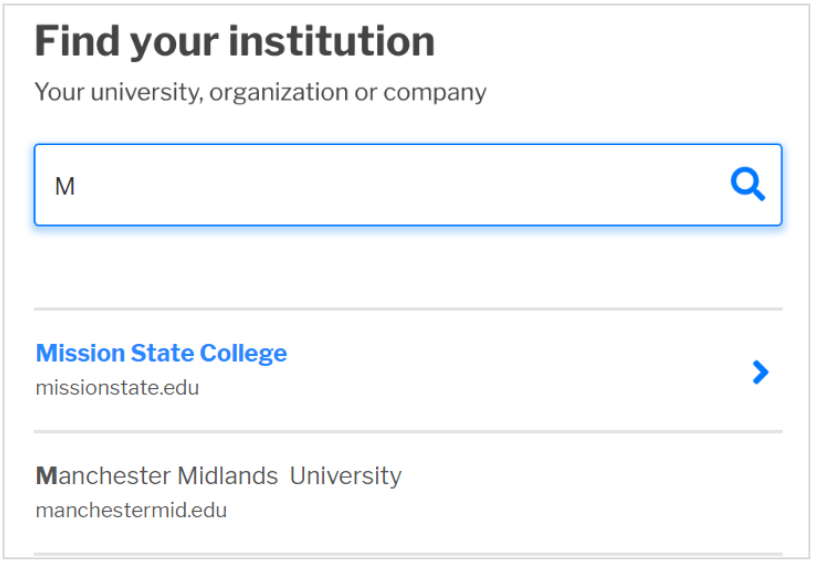

Figure 11. *Displaying search results*

- Limit the number of displayed search results to what will fit in the visible frame. Users should not have to scroll through a list. If the number of matches from type ahead is too large to have reasonable confidence of a relevant match displaying near the top of the list (e.g., greater than 10), wait for users to type more characters before displaying the matches.
- Display institution domain, in addition to the institution name, to show users that they will be taken to a different site.
- Show institutional icons, unless they impact performance. Search results without icons tested effectively and users were able to find their institutions.
- Support accessibility by providing full keyboard support to navigate to the search result and select it. Provide a visible "on focus" style for all elements so that users know when elements are in focus. Provide information about the number of search results to assistive technology. Users need to be able to learn about dynamically changing results using screen reader. (e.g., "Five institutions found matching New York. Use Up and Down arrows to move through results.")

2.4.4.5. Discovery Interface Error Handling

IdP discovery service search interfaces should handle errors as follows:

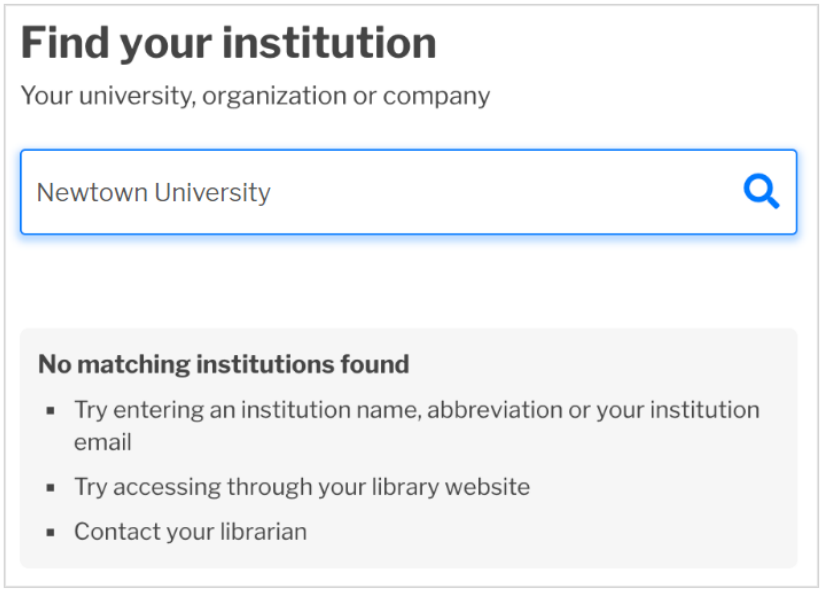

Figure 12. *Institutions are not found.*

- When there are no matches, provide a helpful message instructing users on next steps that may lead to success.
- Accessibility: Assistive technology needs to be aware of the message. Users need to be able to learn that no matches are found using screen reader.

2.4.4.6. Displaying Remembered Institutions

IdP discovery service search interfaces should handle remembered institutions as follows:

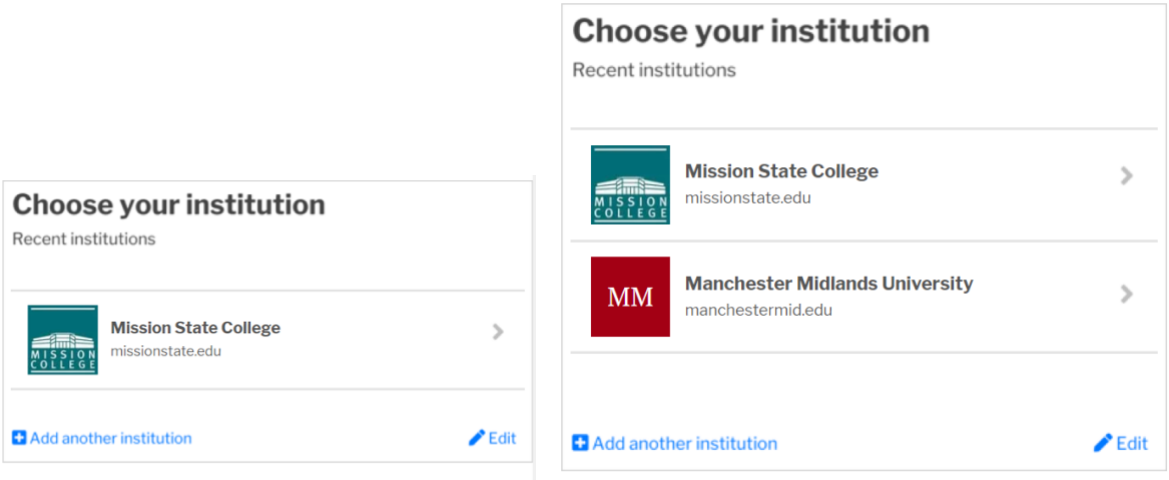

Figure 13*. Remembered institutions (single and multiple)*

- Display institution names with institution icons if available.
- When IdP institution is known for a user coming to the IdP discovery service, instead of presenting the user with Figure 8: “Find your institution”, display that institution name (with relevant metadata as shown) and ask the user to confirm using that institution or choose another institution.
- Note: a recommended alternative is to display the remembered institution in the SP page CTA (see Section 2.4.3.2. Dynamic Button).
- Accessibility: Provide labeling for all content to be accessible to assistive technology. Provide full keyboard support to navigate to the search result and select it. Provide a visible “on focus” style for all elements so that users know when elements are in focus.

2.5. Establish a Cross-domain Identity Provider Persistence Service

A Cross-domain IdP Persistence Service, available for use by all participating SPs, is essential for the implementation of the user experience envisioned by the RA21 initiative. The Where Are You From (WAYF) process, in which a user searches for and selects one's home institution, is an essential function for federated authentication. However, it is a multi-step process and even the best WAYF implementations yield a suboptimal user experience. Further exacerbating this problem for scholarly communications is the fact that a user will encounter this WAYF process multiple times, on multiple SP websites, during a typical single browser session.

The IdP Persistence Service envisioned by RA21 will greatly improve this user experience. This service will allow users to select their institution once and have that same choice presented to them later, either when the user returns to access resources from the same SP or when visiting a different participating SP, so that the user is not required to go through the full IdP discovery workflow again. Provided the user's IdP session stays active, the persisted choice of IdP provides for one-click access to resources when users move from resource to resource and across SPs. Users will be transparently authenticated to new SPs as needed, avoiding any disruption in the user’s train of thought.

RA21 recommends the establishment of this Cross-domain IdP Persistence Service adopting the technical approach from the RA21 Privacy Preserving Persistent WAYF (P3W) pilot. The technical approach from the P3W pilot persists a user's IdP choice in web storage. This preserves user privacy, and also minimizes any information security risks by ensuring there is no centralized store of user preferences. The user is free to delete their remembered choices at any time.

Access to the Cross-domain IdP Persistence Service is handled through a JavaScript API served from a single, trusted domain; this API would be available to IdP discovery Interfaces (Section 2.4.4.6.) and is built into the dynamic version of the access button (Section 2.4.3.2.). SPs that desire to take advantage of this service, but who will not be

using the default access button or discovery service, may develop their own integration with the API but care must be taken to analyze the security aspects of any such implementation.

In the event that the Cross-domain IdP Persistence Service is not available, or not usable by the user's browser, the system will gracefully fall back by allowing the user to select make a fresh IdP choice for every session.

The Cross-domain IdP Persistence Service must be governed in an open and transparent manner, and any architectural changes to the service must be evaluated for any possible security and privacy implications.

2.5.1. Establish an Option to Allow IdP Choices to be Pre-populated in Managed Environments

During the RA21 corporate pilot, the ability to pre-populate a Web browser with the user's identity provider was tested. Pre-populating the user's identity provider means that a user would never have to search for her home institution, and that the RA21 button would automatically appear with the name of the user's institution. The RA21 corporate pilot tested two ways of accomplishing this pre-population.

For a site that uses a desktop/laptop management tool, it should be possible for central IT to push an update to managed desktops/laptops that would pre-populate the RA21 button. This would likely be an option for corporations that have centralized control over their desktops/laptops. The second option is for a specially crafted link to be created which could be sent to the site's users via email. Clicking on such a link would pre-populate the RA21 button.

2.6. Improve Metadata Quality and Apply Consistent Standards

IdPs and federation operators can help--or hinder--the IdP discovery experience for the end user in a variety of ways. The SAML specification offers a variety of metadata fields that are used to automatically share information among member IdPs and SPs, such as the organizational logos, the sector in which the organization does business, where in the world the organization is located, and more. In particular, an extension of the SAML specification called the Metadata Extensions for Login and Discovery User Interface (MDUI)[16] provides a clear and consistent way for IdPs to format and share this information automatically within their federation or with other partners.

Currently, many IdPs do not provide appropriate information in MDUI attributes, either because they do not provide this information at all, or because the federation to which they belong does not support the necessary attributes. This has a significant negative impact on the user experience of IdP discovery interfaces. RA21 recommends that all federations and IdP implementations support keywords and description MDUI elements. As support for

MDUI elements becomes available, all IdP administrators populate them and IdP discovery interfaces support appropriate use of these fields.

Two fields in particular offer enormous value in improving the IdP discovery process: DisplayName and Logo.

2.6.1. DisplayName

DisplayName is the name of the institution shown to the end user when they interact with an IdP discovery service. Ideally, DisplayName should be provided both in English and the local languages of the IdP. By including it in the MDUI information, IdP discovery services will be able to efficiently collect the necessary information to present to the user avoiding the need for federation or IdP-specific processing.

2.6.2. Logo

The Logo field is a field that includes information on the organization's logo, including the URL where the image may be found and its size in pixels. Most federations have specific guidelines around the height and width requirements for a logo, to help make the user experience consistent. In addition to the full logo, a smaller icon is also useful. The full logo and the icon, as separate images, may be used in different scenarios. For example, the icon may be used in IdP discovery service search results, while the logo may be used to show the user which IdP they have used previously.

Logos should meet the following basic criteria:

- The logo needs to be able to display on any background.
- Choose the best version of your logo to fit into a small square. Your institution's branding guidelines will include appropriate options.
- The logo must be scalable (e.g., an SVG image).

Each federation currently has their own requirements and guidance regarding logos. RA21 strongly recommends that further work is done within the identity community to standardize logo recommendations across all federations.

2.6.3. Keywords and Descriptions

The Keywords element allows IdPs to help a user locate the correct IdP provider using familiar acronyms or alternative names (e.g., "MIT" for "Massachusetts Institute of Technology", "Caltech" for "California Institute of Technology"). Keywords should be in both English and the native language(s) of the IdP.

The Description element is another field that helps the user understand who and what a service is for. For an IdP, the goal is to specify a brief, localized description of the community supported by this IdP in human language fit for an end user. If an organization has multiple IdPs for different departments, schools, or for testing purposes, this description will help a user identify which IdP they should choose when offered a selection.

RA21 recommends that all federations support keywords and description MDUI elements, that all IdPs populate them, and IdP discovery interfaces support appropriate use of these fields.

2.7. Set Session Timeout Periods Contextually Based on the Type of Resource Being Accessed and Institutional Risk Management Policy

IdP session timeouts are set by IdPs and determine how often a user has to explicitly reauthenticate with their IdP as they move from resource to resource. In general, the user experience will be better with longer timeouts whilst information security considerations may point to shorter timeouts.

Session management (in general) is a control which attempts to mitigate unauthorized access risks. The level of risk is dependent on various factors. For example, physical location is one key factor: If the computer is located inside a steel locked room where access is only provided by registering via a security guard, a picture is taken, and access is only allowed if you successfully perform biometric authentication via retina scan and a PIN code (like in the movies), then unauthorized access risk is low. However, if the computer is publicly accessible (cafeteria, study hall, etc.), it presents a higher risk. Type of data being accessed and level of access are other factors. Session management is a balance between security and usability based on what is being protected. If the computer is in public but the data/system being protected is low risk (not sensitive, personal or confidential), then sessions should be extended. However, if the login is to a privileged account (account administrator, root access, etc.) or the data or system is not low risk (e.g., it is sensitive, personal, or confidential) sessions should be highly restricted (15-30 minutes).

In environments where the risk is considered fairly low by the IdP, the IdP can help improve the user experience by adjusting single sign on (SSO) and authentication timeouts. In the Shibboleth software[§§], these timeouts are set to 30 minutes for SSO and 60 minutes for authentication--other IdP software may have different timeouts but are likely still configurable.

In a low-risk environment such as access to scholarly information resources, RA21 recommends that session timeouts are mapped to a typical users' work period in that environment (e.g., 10 hours). This will result in users having to login only once per business day and having a seamless experience across all SPs. Login accounts that have elevated privileges (account/system administrators, root, etc.) should be more highly restricted (15-30 minutes) because of the level of access and associated risks involved.

[§§] https://wiki.shibboleth.net/confluence/display/SHIB2/IdPAuthnSession

2.8. Establish Security Incident Reporting Frameworks

The recommendation to adopt federated authentication implies that the onus for detecting and preventing hacking and abuse of resources by malicious actors has to be shared between IdPs and SPs. It is therefore important to adopt a mechanism that allows information about suspected security incidents to be shared between SPs and IdPs without compromising legitimate user's privacy.

With the ever-increasing threat of cybersecurity attacks, it is important to understand that security incidents threaten not only resources provided by service providers, but also systems operated by IdPs themselves, some of which may contain sensitive data such as HR information, patient health records or sensitive intellectual property. As an example, in March 2018, the FBI charged nine Iranian citizens with an elaborate and massive hacking campaign intended to steal academic data and intellectual property. According to the FBI report, the massive spear phishing attack targeted more than 100,000 accounts of professors around the world and successfully compromised approximately 8,000 of those accounts.

Working together, SPs and IdPs can provide an effective means of detecting and limiting the impact of compromised login credentials without compromising user privacy. Today's widespread use of IP address recognition for authorizing access to scholarly information resources means that when suspicious activity is detected on an SP's site, SPs are unable to provide IdP institutions with sufficient forensic information to trace and block the attack. This often leaves SPs with no choice but to block entire IP ranges, inappropriately impacting legitimate users.

SAML federated authentication provides enhanced end-to-end traceability and more targeted forensic data to assist in detecting compromised accounts, while maintaining privacy. Each federated authentication user session establishes a unique, ephemeral SAML session ID between the IdP and the SP. This SAML session ID conveys no identity information about the user to the SP. However, in the case where a SP detects potentially suspicious activity, the SP can report the SAML session ID to the IdP for further investigation. The IdP can then quickly determine the compromised user account using the SAML session ID and work directly with the user to determine whether or not the account has been compromised and take appropriate protective measures.

The manner in which SPs and IdPs communicate security incidents has been formalized by several academic federation operators. A notification framework called Sirtfi[17] (Security Incident Response Trust Framework for Federated Identity) has been developed for this purpose. RA21 recommends the adoption of Sirtfi, joining others in the community, such as ORCID, who notes: "To ensure an effective security incident response, Service Providers and Identity Providers need to coordinate efforts. Sirtfi provides a component necessary for trust in identity federations, and the community would benefit from widespread adoption."

2.9. Leverage Existing or Establish New Interfederation Services for SP and IdP Interoperability

Multilateral federated authentication relies on the exchange of a common set of metadata that describes characteristics of Service Providers and Identity Providers. In order to enable seamless SAML access to resources across globally, RA21 recommends the establishment of global interfederation services for metadata exchange across sectors.

This includes developing and promoting best practices for the exchange of global metadata maximize interoperability between SPs and IdPs. Identity federations participating in interfederation should publish the metadata for all SPs and IdPs with the interfederation service to avoid the creation of "walled gardens" (i.e. to avoid situations where certain SPs can only federate with certain IdPs). A working example of such a service and related community development is the eduGAIN Interfederation Service run by GÉANT for the Research and Education sector. RA21 will leverage eduGAIN service for federating with that community.

RA21 also endorses the formalization and adoption of scalable metadata management and distribution services such as pyFF or the Shibboleth Metadata Aggregator, and per-entity metadata via the MDQ protocol. These services are needed for the scalable operation of federations moving forward.

Section 3. Future Work Items

RA21 recommends that the community establishes a follow-on collaborative structure to oversee the implementation, operation and governance of the Cross-domain IdP Persistence Service and central IdP discovery service proposed in this document. This structure should follow the principles of broad stakeholder participation and open, transparent governance which RA21 followed.

In addition, the following work items should be considered by this or other community groups for further standardization and formalization:

- The user interface for user-controlled consent for attribute release.
- The user interface for increased transparency about information use, including:
    - informing users that the IDP selection will be remembered as well as instructions on how to manage their remembered institutions.
    - indicating that email addresses entered will not be retained; they are solely for matching on the domain.
- The user interface for the browser storage policy notification to support storing the user choice in a way to be accessible across domains and compatible with all browsers.
- Guidance on the use of hints (e.g., IP address ranges, geolocation) to improve IdP discovery options.
- The attribute specification for granular usage reporting.
- The establishment of entity categories and standard attribute bundles for the use cases outlined Section 2.3.
    - Refer to the work underway with the Federated Identity Management for Libraries effort.
- Guidance on increasing the visibility of the institution providing access to resources on SP sites, for example by aligning on a common UI for display of institutional name and logos on SP pages and streamlining administration through the use of standardized metadata elements.
- Development of best practice guidance for logos.
- Establish an open and transparent governance model for any operational service developed from the findings of the RA21 initiative.
- More granular control by the SP regarding what IdPs are displayed, even when using the Cross-domain IdP Persistence Service.
- The best authentication mechanisms to use for library “walk in” users
- The best authentication mechanisms to use in clinical care settings
- Finalization of the mechanics of integration with the Cross-domain IdP Persistence Service for SPs who are not using the default access button.
- Development recommendations on a standard authentication model for machine-to-machine access to Service Provider APIs.
- Report from the conclusion of the Hospital/Clinical Access Working Group.

Appendix

A.1 Reference Architecture

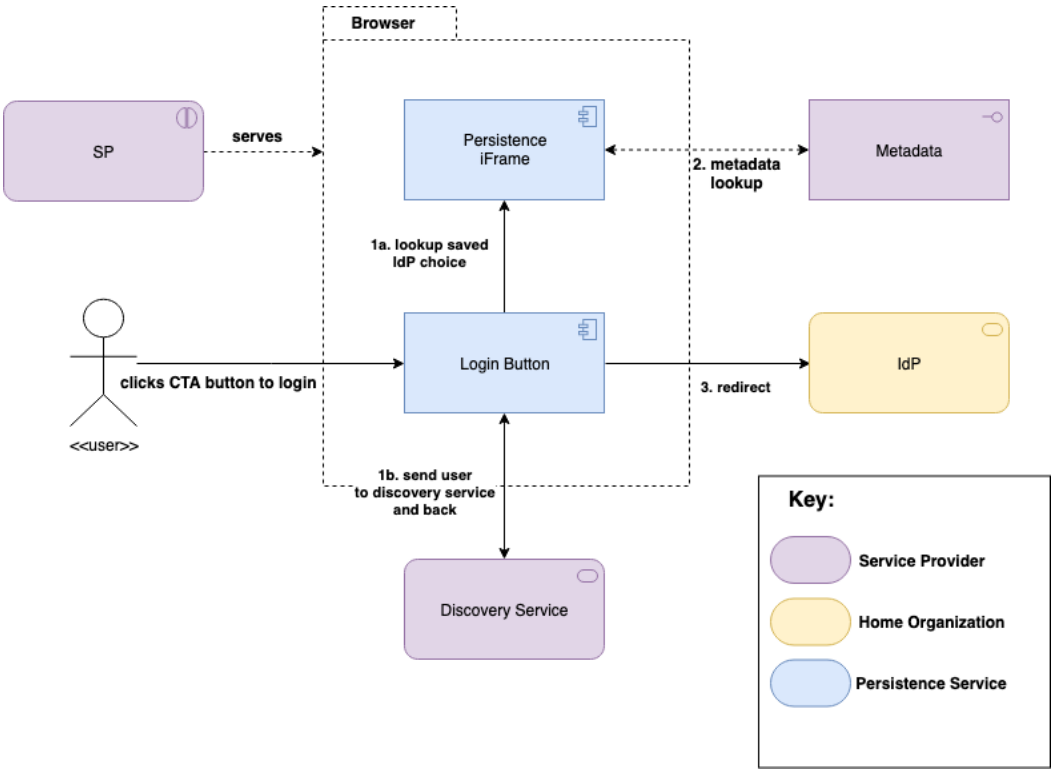

Figure 13. *Reference Architecture for IdP discovery and Cross-domain IdP Persistence Services*

In this architecture diagram, as the user clicks on the CTA button on the SPs website, the associated JavaScript looks initially in web storage for previous IdP choice. Depending on the age of the information, it may refresh the data for the IdP (such as its target URL) from metadata store, and then will redirect the user to the IdP. (Note that the SP does not see the choice stored in the browser itself). If there is no choice stored (or the user requests to select another option), the user is redirected to an initial IdP discovery service that interacts with the Cross-domain IdP Persistence Service to present a searchable list of IdPs to the user and store that choice in the user's browser.

A.2 Pilot Technologies

The ultimate goal of RA21 was to develop a set of recommendations based on potential implementations. Three pilots went forward to test the functionality and limitations of federated authentication:

- Corporate Pilot
- Academic Pilot - WAYF Cloud
- Academic Pilot - Privacy Preserving Persistent WAYF

At the conclusion of the pilot phase, three reports were created: a report on the findings of the Corporate Pilot, a high-level review of the two academic pilots, and an in-depth security and privacy evaluation of the academic pilots. After reviewing the three pilots, it was determined by the RA21 Steering Committee that the group should advance a recommendation based upon the Privacy Preserving Persistent WAYF pilot, which then became the basis for this Recommendation, The WAYF Cloud pilot technology was not advanced and is provided below for reference. The corporate pilot focused more on user experience and testing the idea of establishing a corporate federation; the feedback from that pilot has been incorporated into the user experience guidance and in other aspects of the guidance for the use of federated identity.

A.2.1 Corporate Pilot

The Corporate Pilot was the kick-off for RA21, as the pilot began in 2016 and was merged into the larger effort. Participants in this pilot included members of the Pharmaceutical Documentation Ring (P-D-R), including AbbVie, BASF, Glaxo Smith Kline (GSK), Novartis, and Roche, and several publishers, including the American Chemical Society (ACS), Elsevier, Springer Nature, and Wiley.

The three goals for this pilot were:

- Improved user login experience at the publisher sites
- Provision for granular usage statistics reporting
- Ability to easily set up and maintain Single Sign On with multiple publishers.

During the first phase of the pilot in 2017, the initial User Experience (UX) development was tested with end users at the P-D-R companies. Key findings from the phase one (2017) testing included:

- Equal support for the use of institutional name and personal email address for identification at the publisher site.
- Privacy concerns raised around use of email address.
- Confusion identified around variety of names for an institution.
- Individual user registration seen as being more valuable for frequent users but could be a privacy issue for some.

The SAML protocol was tested with two publishers (Elsevier; Springer Nature) and additional attributes identified that would be required for department billing, differentiating between employee types and for granular usage reporting.

The second phase of the corporate pilot (January to June 2018) focused on further UX testing by the P-D-R pilot participant companies, the specification for granular usage reporting and the exploration of options for the set up and maintenance of a P-D-R-specific federation.

A.2.2 Academic Pilot - WAYF Cloud

The WAYF Cloud pilot focused on the use of a cloud service, rather than local browser storage to facilitate the exchange of data between publisher platforms to simplify IdP discovery and user login. This was also an open source software package[***], developed by Atypon.

The WAYF Cloud service assumes that IdP discovery is handled by the SP. The focus of this service is to persist that choice such that any participating SP can use the WAYF Cloud to present the user's choice back to the user. The WAYF Cloud does this by using both JavaScript and back-end API calls to store a mapping between each SP's local unique device identifier and a common global unique device identifier in a central database. The central database then maps the global shared device IDs to the IdPs the user has successfully logged into.

Note that the WAYF Cloud architecture does not store username or passwords or other personally identifying information. The data trail will exist during the life of the session when the user is in Incognito mode, and in the browser until the cache is cleared.

A.2.3 Academic Pilot - Privacy Preserving Persistent WAYF (P3W)

The Privacy Preserving Persistent WAYF (P3W) pilot focused on the validation of the use of SAML-based federated authentication technologies to provide seamless access to scholarly resources for authorized users at participating institutions while protecting the user's privacy. A key aspect of the technology was the need for a lightweight central service that would enable the storage of user-specific data within the user's local browser. The pilot was built on an open source software package developed by the Swedish national research and education network organization, SUNET, called pyFF[†††] (Python Federation Feeder).

The P3W pilot supported two models of integration. Level 1, the most basic integration model assumes that the SP — generally a publisher in the RA21 use case — wants to completely externalize its federated IdP discovery services. As such, it would use a common URL to point to a central IdP discovery service that would allow the user to choose among a list of possible IdPs, and then record that choice in a user's browser so that it is available for future sessions.

[***] https://github.com/Atypon-OpenSource/wayf-cloud

[†††] https://github.com/IdentityPython/pyFF

In the more advanced level 2 scenario, the SP would use a local IdP discovery service that could include any local accounts hosted by the SP and use the Cross-domain IdP Persistence Service only to store the user's choice of IdP in the browser. This would be accomplished by calling an API within the browser provided by JavaScript hosted by the Cross-domain IdP Persistence Service on a trusted domain so that participating SPs can all access the same shared set of remembered IdPs.

Note that the P3W architecture only supports storing the user's choice (or choices) of IdP in their browser, no usernames, passwords or other personally identifying information is stored. If a user uses a private browsing mode, any choices made will not be stored after that browser windows is closed.

A.3 User Experience Design Rationale and Research

A.3.1 Common Design Questions

1. **Why use a button?** For most users, institutional access is the best bet for accessing the full text, so the call to action for institutional access needs to be easy to find and easy to choose. In most cases, institutional access should be presented as the primary or preferred path for a user, as opposed to an equivalent option. (It should only be presented as a secondary option if it truly is not the preferred access path.) Alternative presentations of the call to action (versions with and icon and link) were tested and were significantly harder for users to recognize. This design also follows best practice/ familiar conventions for similar experiences users encounter in everyday digital experiences (e.g., button signaling option to check out with PayPal).

2. **Why have a symbol or icon?** The goal of the design has always been to create a visual cue that instantly signals to the user that their institutional relationship is what will get them access. The institution symbol aides in quick recognition across different publisher sites. Along with the label, it reinforces the relationship (the institution) that is most likely to gain them access. With that association, the user knows what to expect and can confidently complete the subsequent steps with minimal friction.

3. **Why are words like "log in" or "sign in" not used?** Several versions of the call to action text were tested including "log in through your institution" and "Get full text". While users generally understood the labels, they also introduced ambiguity or incorrect expectations, because users possess multiple login credentials and they are uncertain which ones are being requested. "Access through your institution" was clear and unambiguous and set expectations that were aligned with what users then experienced once they clicked the button.

4. **When the identity provider is known (e.g., Scenario 2b) why does the user need to click a button to trigger the authentication function? Why can't authentication happen automatically without the user needing to click? This**

would be closer to the IP authentication experience. This is not a recommended experience for three reasons: 1- if the IdP is unavailable for some reason, the user will be left in an unrecoverable error situation; 2 - the user needs the ability to change identity providers because they use multiple providers or because the last used provider may have been selected in error. 3 Asking users to verify the relationship that provides them access is a common best practice for increased transparency and user confidence.

5. **How will display of different interface languages be accommodated?**
   The dynamic button will recognize a language parameter and display appropriate text.

6. **How should long institution names be handled?**
   Long names will be truncated to fit the available width of the button or the IdP discovery service results panel. The dynamic button is designed to be responsive to the available display space and additional handling rules will be developed as the service is built out.

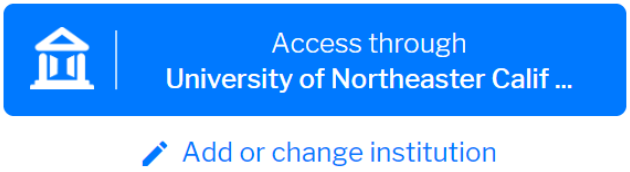

Figure 14. *Examples of truncation for long institution names.*

7. **How will accessibility standards be met?** The central IdP discovery service interface will be WCAG 2.1 AA compliant. The call to action button on service provider pages should be implemented using standard web technology that allows communication with assistive technology. Specifically:

   - Button label and text should be available to screen readers and have high color contrast.
   - Users must be able to navigate to and activate the button using only a keyboard.
   - The button should to be placed in a logical location; assistive technology users should be able to reach the button with a logical sequence of steps.

A.3.2 User Research Insights

The recommendations in this document are informed by best practice research as well as multiple usability studies, interviews with and surveys of representative users.

The goal of this research was (1) to understand the source of frustration and challenges users encounter when they are presented with barriers to access full content in the midst

of their research process and (2) to test different solutions for removing and minimizing those challenges and provide an informed optimum user experience and recommendations.

A.3.2.1 User Needs and Frustrations

Remote Access. A majority of researchers do not have a good understanding of how remote scholarly content access works. Users expressed a desire for an easy remote access solution to scholarly information:

> "*This is always a tricky thing to do, to be able to get access to online journals when you are off campus which I commonly do."*
>
> *"We have monthly study groups and they are off campus."*
>
> *"Very applicable to me and very glad that people are continuing to work on providing these features."*

Access tasks disrupt research workflow. Users express frustration when they are interrupted in their research flow and are forced to shift their attention to figuring out how to get access to full text. The number of steps required to complete the process (often 6 or more clicks) and the uncertainty of the outcome is off-putting to many users, leading them to abandon the attempt. Even users who understand federated access and are motivated to complete the process, still often experience a lot of difficulty getting access.

> "*It never works 100% of the time, sometimes the system just doesn't work.*"
>
> *"A number of times when I click on the institutional logins, there are a number of options, I think Shibboleth is one where not every institution is listed there. So, although that is the most intuitive button on a website, typically when I click on it, I find that my institution is not listed among those with that option. Typically, the best option for me is login through my institution externally and then go back to the article link and just click PDF."*

A.3.2.2 Key Findings and Insights

This section provides a summary of key findings and insights from multiple user studies that informed the design of the user experience and recommendations in Section 2.4.

Using a consistent visual cue reduces friction and cognitive load.

User testing showed that users start recognizing the pattern (text, button, icon) as a primary call to action quickly. Users expect the familiar button to have the same

functionality from site to site. This reduces the demands on users' memory, removes hesitations and allows users to remain in their flow.

User testing also showed that none of the users had a strong feeling about the color of the button. They have stronger feelings about the location of the button and their ability to get access.

> *"I actually read what the words say. I am not too concerned about the color. I would say if it was a consistent color, it would make no difference to me across journals because I think I would still read what the button is telling me to do."*

> *"I do like something that helps it* [the call to action] *stand out, doesn't matter if it's the same or different colors or if there are even colors. The words help, the icon helps, color helps but what it comes down to its just the ability to access it through the university subscription."*

In the same vein, adding a previously selected institution name to the call to action, increased users' understanding and trust of the button.

Placement, grouping and hierarchy of access options are essential for recognition.

Users took >20 seconds longer to locate the "Access through your institution" button when they had to scroll down to find it and when it was separated from other access options.

A/B testing revealed that users were 4 times more likely to choose the wrong access option when all access choices are presented as equivalent.

Placement, or the location of the primary call to action, is the most important factor for recognition.

User testing showed that when the primary call to action is placed in the initial page view and (1) on the top right of the page or (2) between the title of the article and the abstract, users are able to easily and quickly spot it and both placements have similar effectiveness. However, when the primary call to action is placed lower on the page, it is significantly less effective.

Offering access options on an overlay or page layer is not desirable.

Some sites place their access options in an overlay or separate page that is presented when the user clicks on a pdf, full text or Log in link. Users expressed that they prefer to see the institutional access options directly on the page, as opposed to presented in this way.

> *"It encourages you to know you can access the article as you have an institution to log in to. Also, some people will be looking for that option immediately and it makes it easier to find."*

> Seeing the access button on the page *"cuts out a step and it tells me that I automatically have to be logged into my school account or have an offline account so that I don't waste time trying to view the article when it only gives me the abstract."*
>
> *"Say you are a new student and you didn't know that you could login through your institution on that website. It will help them recognize that they have that option to access full text, oh I can do that!"*

User research also showed that, presenting access options in a layer added 25 seconds to the average task time over presenting options directly on the page. While successfully identifying the correct access option in a layer is high (88%), the median total task time from landing on the abstract page to clicking on the primary call to action was 48 seconds (33 seconds for identifying the initial call to action and 16 seconds for selecting "Access through your institution").

A PDF icon or label distracts users from other access options.

When a PDF icon or label is present, a majority of users do not explore alternative calls to action on the page. In testing, users were annoyed with the experience when they were presented with a PDF icon but don't have access to the PDF, particularly if they are presented with a no access error message.

> *"It was confusing since it showed PDF link even though I didn't have access yet. The last two templates were more intuitive because the* [PDF] *links didn't appear until I was signed in."*

Entering institution name is preferred over entering institutional email address.

For IdP discovery, users more strongly identified with entering an institution name than with email address or domain when searching for institution.

> *"I associate access with my institution. I don't associate myself having the access individually. The relationship is with my institution, not me as an individual."*
> *"I have experience doing this elsewhere--I look for the institution name first, then I go to my university website to enter my email and password."*

Endnotes

[1] Cliff Lynch, *Report on the CNI Authentication and Authorization Survey 2016*, Coalition for Networked Information, August 2016, https://www.cni.org/wp-content/uploads/2016/08/CNI-AuthenticationSurveyReport.2016.pdf

[2] Judy Luther, *Universal Resource Access: Finding a Solution*, Copyright Clearance Center, August 2016, https://www.informedstrategies.com/wp-content/uploads/2015/10/CCC_Universal_Resource_Access_Finding_a_Solution.pdf

[3] Meltem Dincer, Chris Shillum, *RA21: Resource Access in the 21st Century*, presented at the International Association of STM Publishers meeting in London, December 2016, https://www.stm-assoc.org/2016_12_11_RA21_2016_Dec_8_016_Outreach_Meeting_CSMD.pptx

[4] *Data Protection Code of Conduct Home*, https://wiki.refeds.org/display/CODE/Data+Protection+Code+of+Conduct+Home

[5] "STRIDE (security)," *Wikipedia*, last modified 11 March 2019, https://en.wikipedia.org/wiki/STRIDE_(security)

[6] "TheIdentitySelector," *GitHub*, June 2019, https://github.com/TheIdentitySelector/

[7] "Atypon-OPenSource/wayf-cloud," *GitHub*, April 2018, https://github.com/Atypon-OpenSource/wayf-cloud/

[8] "Entity-Categories Home," *REFEDS wiki*, REFEDS, last modified February 2018, https://wiki.refeds.org/display/ENT/Entity-Categories+Home.

[9] "Security Assertion Markup Language," *Wikipedia*, last modified 4 June 2019, https://en.wikipedia.org/wiki/Security_Assertion_Markup_Language

[10] *Web Storage (Second Edition)*. World Wide Web Consortium. 19 April 2016. https://www.w3.org/TR/2016/REC-webstorage-20160419/

[11] *Welcome to OpenID Connect*, OpenID Foundation, https://openid.net/connect/

[12] *Data Protection Code of Conduct Home*, https://wiki.refeds.org/display/CODE/Data+Protection+Code+of+Conduct+Home

[13] *eduPerson Object Class Specification (201602)*, Internet2 Middleware Architecture Committee for Education, Directory Working Group (MACE-Dir), 9 March 2016, https://wiki.refeds.org/x/KgCuAg

[14] *SAML V2.0 Subject Identifier Attributes Profile Version 1.0.* Edited by Scott Cantor. 16 January 2019. OASIS Committee Specification 01. https://docs.oasis-open.org/security/saml-subject-id-attr/v1.0/cs01/saml-subject-id-attr-v1.0-cs01.html. Latest version: https://docs.oasis-open.org/security/saml-subject-id-attr/v1.0/saml-subject-id-attr-v1.0.html.

[15] *Web Content Accessibility Guidelines (WCAG) Overview*, Web Accessibility Initiative. June 2018, https://www.w3.org/WAI/standards-guidelines/wcag/

[16] http://docs.oasis-open.org/security/saml/Post2.0/sstc-saml-metadata-ui/v1.0/sstc-saml-metadata-ui-v1.0.html

[17] *SIRFTI,* REFEDS, https://refeds.org/sirtfi