

ESPreSSO: Establishing Suggested Practices Regarding Single Sign-On

Public Comment Period: May 23 – June 22, 2011

*A Recommended Practice of the
National Information Standards Organization*

Abstract: ESPReSSO explores practical solutions for improving the success of SSO authentication technologies for providing a seamless experience for the user and makes recommendations for promoting the adoption of one or more of these solutions to make the access improvements a reality.



Published by:
NISO, Baltimore, Maryland, U.S.A.

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This recommended practice may be revised or withdrawn at any time.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)
One North Charles Street, Suite 1905
Baltimore, MD 21201
www.niso.org

Copyright © 2011 by the National Information Standards Organization

All rights reserved under International and Pan-American Copyright Conventions. For noncommercial purposes only, this publication may be reproduced or transmitted in any form or by any means without prior permission in writing from the publisher, provided it is reproduced accurately, the source of the material is identified, and the NISO copyright status is acknowledged. For permission to photocopy or use material electronically from NISO RP-11-201x, ESPreSSO: Establishing Suggested Practices Regarding Single Sign-On, please access www.copyright.com or contact the Copyright Clearance Center, Inc. (CCC) at 222 Rosewood Drive, Danvers, MA 01923, 978-750-8400. CCC is a not-for-profit organization that provides licenses and registration for a variety of users. All inquiries regarding translations into other languages or commercial reproduction or distribution should be addressed to: NISO, One North Charles Street, Suite 1905, Baltimore, MD 21201.

ISBN (13): [to be added at publication]

Table of Contents

Foreword	v
Part 1: Introduction	1
1.1 Purpose and Scope	1
1.2 Terms and Definitions.....	1
Part 2: Why Is It Time to Act?	4
2.1 Overview of Issues.....	4
2.2 Library Community.....	4
2.4 End User Community.....	5
Part 3: Traditional Approaches to Controlling Access to Licensed Resources.....	6
3.1 The Evolution of Authentication Requirements.....	6
3.2 The Evolution of Access Control	7
3.2.1 Client Machine IP Address and Client Organization VPN Services	7
3.2.2 Proxy Servers	9
3.2.3 Userids/Passwords for a Service Provider Site.....	10
3.2.4 Federated Login (Authentication)	11
Part 4: ESPReSSO Recommendations	16
4.1 Overview.....	16
4.2 Use Cases	17
4.3 Summary of Recommendations	18
4.4 Recommendations to Service Providers.....	20
4.4.1 Service Provider Open Page.....	20
4.4.2 Service Provider Identity Discovery Page.....	21
4.4.3 Service Provider Protected Page.....	22
4.4.4 Attribute-Based Authorization	23
4.5 Recommendations to Libraries / Institutions	23
4.5.1 Institution Login Page	23
4.5.2 Institution Menu Page.....	24
4.6 Role of a Proxy Server in Supporting a Hybrid Environment.....	26
4.7 Rewriting OpenURLs	26
4.8 Appropriate Use of Branding	26
4.9 Additional Functionality	27
4.9.1 Pseudonymous Access.....	27
4.9.2 User Consent to Attribute Release.....	27
Part 5: Content Discovery Services.....	29
5.1 Content Discovery Services	29
5.1.1 Overview of Federated Search	29
5.1.2 Overview of Web-Scale Discovery Services.....	30
5.2 Existing Authentication with Discovery Services.....	31
5.3 Recommendations for Authentication in a Discovery Search Environment	31
Appendix A Description of Functions in Current Authentication Environments	32
Bibliography.....	34

Figures

Figure 1: Use case #1 scenario17
Figure 2: Use case #2 scenario17
Figure 3: Use case #3 scenario17
Figure 4: Use case #4 scenario18
Figure 5: Mock-up of Identity Discovery page using recommendations.....22
Figure 6: Mock-up of Institutional Login Page using recommendations24
Figure 7: Example of Institution Menu Page.....25
Figure 8: Federated search.....29
Figure 9: Web-scale discovery search30
Figure 10: Functional components of current authentication environments.....32

Foreword

About This Recommended Practice

In 2009, NISO launched a new Chair's Initiative—a project of the chair of NISO's Board of Directors, focusing on a specific issue that would benefit from study and the development of a recommended practice or standard. Oliver Pesch, Chair of NISO's Board of Directors at the time, chose the issue of standardizing seamless, item-level linking through single sign-on (SSO) authentication technologies in a networked information environment.

Accessing information in a networked environment has been a reality for most library user communities for over a decade. Recent years have seen an explosion in this type of usage. With the advent of hosted aggregated full-text databases and the proliferation of e-journals and e-books, users' searches for information often take them to a number of different online hosts and platforms as part of a single transaction. When those information resources are commercial products, each platform traditionally required the user to be authenticated and authorized. Service providers (SPs) have used two approaches to this issue: 1) ensuring that the requesting IP address is within a range assigned to the license holder, and 2) issuing userids and passwords to users. In the latter case, the user may have a different identity on each platform.

As usage habits and technology have evolved, these traditional methods no longer work well. With the growing complexity of licensing situations and network design, along with the increased usage from mobile devices, campuses have outgrown the capabilities of these approaches. Using and managing this environment has become confusing and error-prone for browser users, licensee organizations, and service providers. The problems caused by having to manage multiple identities have led to the development of so-called "Single Sign-On" (SSO) authentication technologies, including proprietary technologies such as Athens and formal open standards such as SAML (security assertion markup language). With these technologies, the user authenticates once and can then access all compliant content platforms using the same identity. (The user would typically be authenticated by the organization holding the licenses.) More importantly, these technologies have been designed so the user would encounter only one login event while traversing a multitude of in-sourced and outsourced service providers. In addition, with the SSO technologies' the user does not have to be using a device attached to the license holder's network; they can be anywhere in the world. Simplifying the user experience has become more important as organizations have outsourced more and more of their supporting business functions (not just to licensed content).

A bridge is needed to address today's hybrid environment and move all parties towards a longer-term effective SSO solution. The ESPReSSO Recommended Practice document recommends practical solutions for improving the success of existing SSO authentication technologies to provide a seamless experience for the user. Specifically, ESPReSSO recommends best practices related to selection of authentication method and transparent flow between the service provider (SP) site and the identity provider (IdP) site during authentication.

Recommendations to service providers include the preferred location for login links and input boxes, standard approaches for guiding users to a desired authentication method, where local branding information could be inserted on a webpage, as well as approaches for handling automatic logins.

Recommendations for libraries/institutions include display of the login page, branding of the login page, use of a menu page with all available content listed that transfers with automatic login to the selected service provider, and appropriate passing of parameters to the service provider that authenticate the user.

Additional recommendations are made about methods that provide tradeoffs between privacy and advanced functions. Specific recommendations in federated search and web-scale discovery environments are made that will lead all parties from the current environment to a longer-term recommendation to use the Shibboleth authentication model.

ESPreSSO did not invent any new technology or protocols. Instead, ESPReSSO aims to promote the adoption of best practices that make access improvements a reality by using existing technologies while preparing for the future.

Discovery to Delivery Topic Committee

NISO's Discovery to Delivery (D2D) Topic Committee had the following members at the time it approved this Recommended Practice:

[to be added by NISO after final approval]

ESPreSSO Working Group Members

The following individuals served on the NISO ESpReSSO Working Group that developed and approved this Recommended Practice:

Steven Carmody (Co-chair)
Brown University

Ted Koppel
Auto-Graphics, Inc.

Frank Cervone
Purdue University Calumet

Ulrike Lengwenat
Springer

Pete Ciuffetti
CredoReference

Peter Noerr
MuseGlobal, Inc.

Andy Dale
OCLC, Inc.

Lyn Norris
Eduserv

Kristine Ferry
University of California, Irvine

Kishor Patel
ProQuest

Andy Ingham
University of North Carolina, Chapel Hill

Heather Staines
Springer

Harry Kaplanian (Co-chair)
Serials Solutions, Inc.

Pieter van Lierop
Infor Library and Information Solutions

David Kennedy
Johns Hopkins University

Foster Zhang
Johns Hopkins University

Acknowledgments

The ESpReSSO Working Group would like to offer a special thanks to the following individuals for their assistance:

Adam Chandler (Working Group Observer)
Cornell University

Oliver Pesch (Working Group Observer; Project Proposal Author)
EBSCO Information Services

Rob Walsh (original Working Group member)
EnvisionWare, Inc.

For input regarding publisher, aggregator, and platform experience with implementing single sign-on, we thank the following:

American Institute of Physics: Paul DeCillis

Cambridge University Press: Chris Fell

EBSCO Information Services: Sarah Buck and Heather Klusendorf

Elsevier: Chris Shillum and Ale DeVries

HighWire Press: John Sack

H. W. Wilson: Ronald Miller

IEEE: Gerry Grenier

Institute of Physics: Laura Shaw

Ithaka/JSTOR: Matthew Callow and Brian Larsen

MetaPress: Matthew Wren and Tiffany Rich

Nature Publishing Group: Amanda Ward

Oxford University Press: Claire Dowbekin

Semantico: Colin Caveney and Richard Padley

Taylor and Francis: Margaret Walsh and Rosa Perez

Wiley-Blackwell: Caroline Rothaug

Part 1: Introduction

1.1 Purpose and Scope

Making the Single Sign-On (SSO) environment work better and smarter will certainly help increase the success of users getting to the content to which they are entitled. Over the last several years many of the larger service providers (SPs) have implemented SSO technologies. However, it is probably fair to say that many content hosts have not implemented these technologies. Library users are required to operate in an environment that includes a mix of authentication technologies with internet protocol (IP) authentication being the most common. An effective solution needs to address this hybrid environment and, at the very least, take into consideration the needs of IP authentication and proxy servers and how they interoperate with SSO authentication technologies.

The ESPReSSO Recommended Practice document recommends practical solutions and a path forward for improving the success of SSO authentication technologies for providing a seamless experience for the user. It further aims to promote the adoption by campuses and service providers of a family of solutions to make the access improvements a reality. This initiative did not invent any new technology or protocols. Rather, it has developed a set of “best practice” recommendations surrounding the use of existing technologies.

The ESPReSSO Working Group was primarily concerned with the situation where an organization (a company, a campus, a public library, etc.) acquires a license to access specific content that is delivered via the web, and where the browser user is a member of the group authorized to access that content. The working group did not address the situation where an individual would obtain a license for his or her own personal use and then use a social account to authenticate themselves to the service provider. Service providers are reporting that users are not currently requesting this functionality. In addition, supporting this approach requires as much work for the publisher in managing userids and passwords within their site as it does for the licensee organization.

Best practices for user experience on mobile devices are rapidly evolving. Consequently, this report avoids recommendations for screen layout and use on mobile devices. However, the flows described in later sections will work on mobile devices.

1.2 Terms and Definitions

The following terms, as used in this recommended practice, have the meanings indicated. See also [Appendix A](#), which contains definitions and descriptions of functional components found in an authentication environment.

<u>Term</u>	<u>Definition</u>
service provider (SP)	A website offering services and content to browser users. Often, part of the site is available for open browsing and part of the site is access controlled, requiring proof that the user is authorized for access under an existing contract. Many publishers or content providers are offering licensed content via the Internet.
identity provider (IdP)	A user's home organization. The user typically has a strong relationship with this organization, and it can make up-to-date and accurate assertions about the user.
identity discovery service	A webpage presented by service providers to which users are redirected for authentication after selecting their home organization. This page was previously called a WAYF (Where Are You From) Service.
deep link	A URL pointing to a specific resource (e.g., an article) at a service provider site.

<u>Term</u>	<u>Definition</u>
federated login (authentication)	An approach to authentication where a service provider site redirects browser users to their home organization for authentication. After a successful authentication, the browser user is returned to the service provider accompanied by attribute assertions describing his/her rights at the service provider. As used in this report, the term includes the authorization process performed at the SP site.
federated search (metasearch; Web scale discovery search)	A portal that allows a user to enter search terms and then uses the terms to search across a wide variety of potential sources. It allows patrons unfamiliar with where content can be found to discover all relevant licensed content available in the library.
IP-based authentication	A method where a service provider and a licensee organization have agreed that every request coming from a range of network/Internet Protocol (IP) addresses associated with the licensee organization should be granted access to the licensed resources.
Just-in-Time (JIT) authentication	A process at a service provider's website that will initiate the login process when an anonymous user attempts to access controlled content. For instance, JIT will often be triggered if an anonymous user attempts to access a deep link.
licensee organization	The institution that signed a contract with a service provider on behalf of the users who are affiliated in some way with the institution. The licensee organization then becomes an identity provider for those users. This report will sometimes use the word campus to refer to a licensee organization.
link resolver	A service run by the licensee organization that will accept an OpenURL syntax (e.g., from a site with abstracts) and map that to a URL at a site where the licensee organization has access to a licensed copy of the full text of the article
personally identifiable information (PII)	Information which makes it possible to identify a specific individual. A person's name is considered PII (although it may not be sufficient to identify a single individual). This report mentions situations where an identity provider might release attributes which contain PII to a service provider.
proxy server	A computer system that acts as an intermediary between a user making a request and the requested service. Licensee organizations operate proxy servers in order to provide their users with access to licensed resources when the users are not directly connected to the licensee organization's network (e.g., they are at home or traveling). Users authenticate to the proxy, which is connected to the licensee organization's network, and the proxy then allows them access to the licensed resources "through" the proxy server.
security assertion markup language (SAML)	A standards-based approach to web single sign-on (SSO) authentication. Many interoperable open source and commercial implementations of SAML are available.

<u>Term</u>	<u>Definition</u>
virtual private network (VPN)	A secure method of connecting to an organization's protected network from a remote location through a special server that authenticates the remote user to the network using credentials provided by the organization. To service providers, the user then appears to be coming from an IP address on the licensee organization's network, even though the user is physically remote from that network.
WAYFless URL	A URL created by a licensee organization and pointing to a service provider that includes information about the licensee organization. Clicking a WAYFless URL allows a user to bypass the Identity Discovery process.
web single-sign on (SSO) authentication	A framework for authenticating users across a family of services. The user authenticates once and all succeeding authentication requests are handled transparently by the service without requesting anything further from the user.

Part 2: Why Is It Time to Act?

2.1 Overview of Issues

The current mélange of approaches to authentication and authorization has resulted in nightmarish operational situations for licensee organizations, for service providers, and for users. Unfortunately, there is no silver bullet that will quickly return us to a set of simple options. We are in a phase of transition to new models for authentication and authorization; it is likely that—for the short term—all parties will have to work in a hybrid environment where a variety of mechanisms need to be supported. However, because of the complexity this adds to the user experience, all organizations are strongly encouraged to transition to the newer models as soon as possible.

The complexity arises for several reasons:

- The internet world has evolved to provide users with many more options. Users can follow different paths, traversing multiple websites, in order to enter a publisher's site. They can have different starting points (e.g., a link found on the open web, a search on the open web, directly accessing a publisher front page, a redirect from a library navigation page/menu page, a redirect via an OpenURL to a deep-link URL within the publisher site, a bookmark, e-mailed hyperlinks, etc.). The result is that users arrive at many different points on the publisher's site. It has been difficult to create a consistent, coherent user experience amidst all this variety.
- Users may experience multiple authentication mechanisms, depending on how they enter a publisher's site. Sometimes, the user's physical location could affect the browser flows and authentication mechanisms. Within the publisher's site, the user might—innocently—navigate from a public page to a protected page, and thus unwittingly trigger authentication.
- Service providers generally have to present and support multiple authentication mechanisms. They have to construct and present a usable authentication graphical user interface (GUI) that somehow combines multiple methods into an interface that can be used successfully by people with a low familiarity with technical concepts.
- Campuses have deployed various approaches over the years; some of them require users to be able to use, handle, and manipulate proxy-prefixed URLs that are incomprehensible to the average person.

The hybrid environment will likely include both IP and web Single Sign-On (SSO) authentication, and may include other options, as well. Licensee organizations will likely have contracts with some organizations that support SSO and some that still only support IP authentication. Service providers will likely have to work with licensee organizations that want to use web SSO and some that are still relying on IP authentication. However, it is clear that requirements have evolved such that IP authentication is now a poor fit with user requirements, and that, going forward, web SSO implementations will be more secure, easier to user, and more functional for service providers, licensee organizations, and users.

Each of these three communities is impacted in negative ways when confronted with the hybrid environment as described in the following sections.

2.2 Library Community

Patron demand for remote access to content via computer or mobile device has become the norm rather than the exception. Keeping track of numerous usernames and passwords for the multitude of publisher sites and platforms is unmanageable. Ensuring the privacy of user searches, reading history, and behavior has long been a principle of library operations. Librarians are insisting that the traditional levels of privacy remain as an option for the digital user. Many users may be willing to trade some of their privacy for a higher level of service. However, a set of users and situations remain where user privacy is an absolute requirement. Libraries must provide patrons with an efficient, seamless way to access content and to search across this content from multiple sources without continually being challenged for credentials, or having to change the steps they follow as a

function of their physical location. For libraries in countries where national federations have evolved to advocate for SSO, data challenges remain. However, for libraries in locations without government-funded federations, the obstacles are greater, as there is no legal mandate requiring content providers to offer such a service.

2.3 Publisher Community

As licenses increase in their complexity, customers may participate in numerous agreements, allowing varying degrees of access at an institutional, consortial, departmental, or other level. Keeping track of which affiliated users have access to what content becomes more challenging all the time.

At the same time, customer demands for privacy concerning their users' personal details and online search behavior have grown at an even quicker pace. Implementing single sign-on that allows institutions to vouch for their authorized users puts the responsibility and the authority for access back into the hands of the institution, allowing content providers to concentrate on the content rather than the access details. Spurred to action by support for single sign-on amongst European federations, service providers and content providers have labored to meet the varying requirements, including certification, interface adaptation, required attributes, and more. Demands for SSO implementation come now from around the globe and include corporate and government, as well as academia and libraries.

Streamlining the process and the variety of approaches has become essential to service providers who must also ensure that their branding remains visible to end users. The current variability in the end user experience creates a high level of confusion and results in users giving up rather than being able to complete their tasks. A decrease in usage may be used as evidence for cancelation by librarians, so this situation poses a threat to publisher business models.

2.4 End User Community

Researchers and students have access to content through a variety of channels; however, if access is from outside of the university's IP range, a multitude of usernames and passwords might be required. When seeking access to a secured resource, a researcher should be able to identify easily whether the publisher/aggregator supports SSO. He or she should be able to navigate to the institutional login page, to identify the appropriate federation and institution, and, once authenticated, to return to the secured resource with minimal disruption. All stages of this process should be identified and branded so that the request for credentials is not misinterpreted as phishing or malware. As an authorized user of a resource licensed by his/her institution, the researcher should be able to gain access in a manner that is as seamless and painless as possible.

Working in the hybrid environment is particularly confusing for users. Various studies of user interface issues have consistently focused on usability and consistency issues. Browser users are visiting multiple sites, but are seeking some consistency in how they identify themselves to service provider (SP) sites. In addition, in the Federated Login model, users can be redirected from the SP site to the identity provider (IdP) site and back to the SP site; throughout this process, users are seeking visual feedback indicating that the sequence of pages they are seeing are related to their goal of obtaining access to a particular SP site.

Part 3: Traditional Approaches to Controlling Access to Licensed Resources

3.1 The Evolution of Authentication Requirements

Organized communities, such as campuses, public libraries, and corporations have been accessing licensed content via the Internet since the late 1990s. Today, the library at a typical research university or community library may have hundreds of licenses in place for use by its patrons. Remote access to this content via computer or mobile device has become the norm rather than the exception.

An initial approach to license enforcement was to equate the licensee organization with its assigned network addresses, and to assume that any request coming from within that range of addresses was coming from a user covered by the organization's license—a method called IP-based authentication. This turned out to be relatively easy to implement for both service providers and licensee organizations. All users were anonymous to the publisher—protecting user privacy—and all users within the authorized IP range were covered by the same license.

As licenses became more complex, so did authentication. Different slices of a campus community may have different licensed privileges at a publisher site. For example, medical school faculty and students may be granted online access to the full text of certain articles, while non-medical school community members would have to visit the campus library and use a printed version. Faculty, staff, and students may have access to one set of services at a publisher site; alumni may have access to a different set. Someone who is an alumnus as well as a staff member could have access to both sets. This trend toward user categorization to support different license terms was very difficult to support with IP-based authentication.

At the same time, there has been a steadily increasing trend of users accessing licensed resources while away from the campus. Faculty may be working from home or while traveling. Students may be living off campus and working from their apartments. In such instances, their IP address would be outside of the allowable ranges contained in the IP authentication list. Campuses deployed two approaches to this problem: proxy servers and virtual private network (VPN) services. The proxy server would be on the campus network; users would access it, authenticate themselves, and then access resources “through” it; the proxy server would scan the HTML received from the service provider and rewrite links to point to itself before delivering the stream to the browser user. Clicking a link would take the browser user back through the proxy server and on to the resource. VPN services would allow a remote user to connect to a special server on the campus network, authenticate themselves, and then connect to remote services but appear to be coming from the campus network.

Two other changes were also occurring on campuses. The first of these was that campuses began to issue credentials to many people outside their core community of students, faculty, and staff (the group typically covered by the contracts with vendors). This created a significantly larger group of people who could successfully authenticate, including applicants, alumni, and/or other affiliates. Having a proxy server authenticating a user was no longer sufficient; an additional authorization step was required to determine if this person was eligible under the contract to access the resource. Currently, many campuses lack the infrastructure to be able to make this second step feasible.

There has been a strong tradition of preserving personal privacy in the use of library resources and a sustained interest in preserving privacy with the use of online resources. While on some levels this seems easy (the browser never tells the publisher who the users are), remaining anonymous on the Internet has turned out to be much harder than expected.¹ Many usability features designed to allow users to save searches, mark favorite articles or journals, and those which analyze user behavior to recommend additional articles end up compromising personal privacy in unexpected ways. To address this concern, libraries have tried to use their contracts with service providers to enforce personal privacy requirements—with varying degrees of success (e.g., constraining what service providers can do to identify returning users, and what service providers can do

¹ See: Larkin, Erik. “Browser Fingerprints: A Big Privacy Threat.” *PCWorld*, March 26, 2010. Available at: http://www.pcworld.com/article/192648/browser_fingerprints_a_big_privacy_threat.html

with any information they collect about users and usage patterns). Such requirements can also add to the service providers' burden of supporting even more diverse authentication set-ups.

Among the user community, there is a wide spectrum of views on the importance of personal privacy. Many people are willing to share some personally identifiable information (PII) with a publisher in return for a higher level of service. For example, the publisher may allow users with credentials associated with some PII to save search requests from one session to the next. Users who share their e-mail address with the publisher can receive monthly newsletters describing new content in areas related to past searches. In such an environment, PII becomes the currency for higher service levels. Thus, there may not be a universal correct answer to the trade-off between privacy and service—each individual may have to choose what feels comfortable. Interestingly, an individual's perspective on privacy may depend on the situation; a medical researcher may want personalization when doing queries related to work, but want to remain anonymous when doing queries related to his own health issues.

Another evolutionary change involves the integrated use of multiple service providers' platforms. Where previously a user would be searching and accessing only a single service provider's holdings, current usage patterns involve transfers from one site to another. A user might begin searching at a site offering abstracts and then be redirected to a link resolver for full text options, and then further linked to a site containing the full text of an article. With users often beginning their searches with Google, content providers have scrambled to ensure that licensed content is visible in both Google web searches and Google Scholar searches. Libraries with link resolvers often make their holdings available to Google through the Library Links program. While Library Links shows patrons what is available to them "at their library," such information may not always be visible or apparent.

Another example of new usage is federated search, where the user goes to a portal site and enters search terms, which are forwarded by the portal to the set of licensed resources that the user is authorized to access. The search is run across multiple content platforms and the results are returned to the portal, which then displays an aggregation of the results. Clicking a link accesses the content at the relevant service provider. Federated search has a unique set of challenges in that it performs searches of licensed content acting as an agent for the user. The user interacts with the federated search portal and not directly with the content provider's site when issuing queries and viewing search results. Proper operation requires the library, the federated search provider, and the content provider to all be integrated and correctly configured; otherwise the situation of multiple authentication pop-ups can result.

Web-scale discovery services are the latest vendor offering to simplify the search experience for the user. The process is similar to federated search: the user goes to a portal, enters search terms, is presented with a summary of the results, and can click a link to see the full text. However, the implementation of web-scale discovery is very different from the former federated search implementations, and consequently the user may encounter authentication at different steps in the workflow.

Many college faculty and students possess multiple affiliations. They may be associated with a campus with licenses, and may also be a member of one or more professional organizations that also give them privileges at certain service providers. In the near future, users may expect that their access rights to a service provider's content would be a combination of all their privileges.

A growing amount of usage is moving from traditional computers to mobile devices. At the low end, this might be a mobile phone with a display. But, it also includes mid-range devices (e.g., netbooks and tablet computers) addressing mobility needs and providing much smaller screens. It is widely recognized that these smaller devices often require a different presentation to the user. In addition, because these are considered to be personal devices, authentication mechanisms are appearing that are tied to possessing the device, rather than the traditional possessing of a secret device-independent login, i.e., a password.

3.2 The Evolution of Access Control

3.2.1 Client Machine IP Address and Client Organization VPN Services

As discussed in the previous section, license compliance has traditionally been enforced using access controls based on the IP address of the client machine making the request. The assumptions of this approach are:

- 1) The license belongs to an organization with a computer network. The organization supplies (or maintains at the service provider) the range(s) of IP addresses assigned to its networks.
- 2) Only the licensee organization has computers attached to that network.
- 3) All members of the organization have the same privileges and all are treated as anonymous users.
- 4) Only members of the licensee organization will be using computers attached to the organization's network. Therefore anyone using a computer attached to the organization's network is recognized as covered by the organization's license because the IP address is within the organization's ranges.
- 5) Non-affiliates of the organization would be contractually allowed to use the resources *so long as* they did so from a machine physically present on the premises, such as a computer within the library.

Some organizations also implement virtual private network (VPN) services. This allows home computers to connect to a special server on the organization's network and authenticate to the network (using credentials provided by the organization). To SPs, the user then appears to be coming from an IP address on the organization's network, even though the user is physically remote from that network.² Using VPN services usually requires installing specialized software on the remote machine and using a specialized process to start the VPN access.

The general perception has been that IP-based authentication guarantees anonymity for the browser user. However, the use of persistent browser cookies and browser fingerprints has started to reduce a user's ability to remain completely anonymous.

Advantages

- IP-based access control is usually relatively easy for both the service provider and the licensee organization to configure and manage. While the set of IP addresses is maintained at the service provider site, the SP may delegate to the licensee organization the authority to manage the list.
- If the user is using a machine attached to one of the configured networks, access to the service provider is easy and transparent.
- Privacy and anonymity of individual users can easily be maintained.

Disadvantages

- If the user is using a machine that is *not* attached to one of the networks with the allowable IP ranges, then he/she will be unable to obtain access to the desired resource. With the growing use of mobile devices, this is now a common occurrence.
- On campuses, machines directly connected to the campus network were sometimes accidentally misconfigured to become open proxies (see 3.2.2, with the result that anyone in the world connecting through that proxy could access the licensed resources. This clearly violates the assumption that a machine coming from a configured IP address can only be used by a person covered by the license contract.
- In situations involving abuse of the publisher's site, it can be difficult to identify the responsible individual who used the organization's network for access.
- Some international institutions connect to the Internet through centrally managed networks with so-called "anonymous proxies" that obscure the IP of the institution originating the request.
- Network topology does not always successfully map to an authorized subscriber. Some institutions, such as regional school districts, get internet access from centrally managed networks and share an IP range with other organizations; the individual institutions are not segmented by IP address. This inability to segment the user base by IP would also be the case when a department of a larger institution was the licensee and not the entire institution.

² At some institutions, a VPN only intercedes in communications with sites *within* that organization's domain; in those cases, a remote service provider would see the remote user's IP address and therefore deny access.

- Since all users are anonymous to the service provider, it becomes impossible to segment different communities within an organization and identify the individuals within those segments so that different content can be licensed to those segments.
- If a user is offsite and using an institution-provided VPN service to obtain an IP address that a service provider would recognize, then there is a need to install and configure the VPN software on the user's computer. Some users find this process to be confusing and interruption prone and the support issues for this method have proved to be challenging. Consequently, a number of sites are now moving away from providing VPN services, resulting in offsite users with no way of using IP authentication.
- Devices using mobile 3G access, even when onsite, will bypass the established IP addresses and thus have no access to licensed content.
- With the coming transition to Internet Protocol version 6 (IPv6), which allows data packets to have more than one IP address, authentication management using IP address is going to become harder and more confusing.

Summary

Use of mobile communications, need for remote access, changes in network management methods, and desire to segment the community for different content have all contributed to significantly reducing the utility of the machine- or VPN-based IP authentication approach.

3.2.2 Proxy Servers

Licensee organizations deploy proxy servers for a number of security-related purposes and as an alternative to VPN services. One benefit of proxy servers is the ability to provide their members with access to the organization's networked resources, including licensed resources, when they are not directly connected to the licensee organization's network. Users authenticate to the proxy and the proxy then allows them access to the licensed resources "through" the proxy. Because the proxy sits on the licensee network, it is authenticated by the service provider using IP-based access control (or sometimes via referring URL). The proxy server receives the stream of data from the service provider, scans it and even rewrites portions of it, and then sends it on to the browser user. The proxy can rewrite links embedded in the page so that clicking the link would return the browser user to the proxy site and through it to the real destination. This link rewriting is necessary because, if left intact, the original links would take the browser user outside the proxy and directly to the provider, thus resulting in possible denial of access when access was no longer coming from the organization's network.

Central to the use of a proxy server is the understanding that the organization is ensuring that only users covered by the license agreement are allowed to pass through the proxy to the service provider. Open proxy servers, which contain no authentication or authorization steps, are a clear violation of the contract.

A proxy server typically authenticates users using the organization's standard authentication system. The licensee organization is responsible for ensuring that only users authorized under the contract are allowed to use the proxy. Once logged in, a user session is maintained by the proxy; details of the user and the session are unknown to the service provider. Maintaining a user session means that proxies can provide SSO across all of the services they front. Once a session is active, the user will be able to access all proxied target resources without further authentication or authorization.

Many proxy server implementations also offer organizations the ability to control which people and groups are allowed to access what resources. A user attempting to access a resource without the proper permissions would be stopped by the proxy. Of course, this requires that staff at the licensee organization maintain group memberships and the required access rules.

EZproxy is a commercial proxy server implementation that is used by many campuses. Later sections of this report specifically mention it as an example of a proxy server.

Advantages

- Proxies can successfully remap an HTML data stream most of the time. Consequently they are in widespread use by campuses.

- Most configuration changes are confined to a single controllable place—the library’s proxy server. Entitlement setup and checks occur across all content providers in that one location.
- All search and content-viewing traffic routed through the proxy can be logged for later use in reporting.
- For most service providers, supporting access via a proxy is the same as supporting access control via desktop machine IP authentication. No additional work is required.
- To the user, access via proxy looks like and generally behaves like SSO; the experience can be transparent.
- Most proxy servers no longer require installation of software onto personal machines as some older implementations did.

Disadvantages

- The campus must configure and manage the proxy server for every set of licensed content.
- Proxy servers scan and edit the stream of data coming from the publisher site before they pass it to the browser user. When the HTML page content gets very complicated, includes scripts, or has deeply embedded links, proxies may miss the data string that must be edited or incorrectly edit other strings, thereby rendering the page returned to the user as useless. If the service HTML pages contain absolute hyperlinks or vendor-to-vendor hyperlinks, they also might not be correctly recognized or edited by the proxy (resulting in the bypass situation described in the next bullet).
- There are several ways that a user can bypass the proxy by following a URL directly to a content provider. The user will be outside of a proxy-authenticated session and will be denied access to any licensed content. Users frequently encounter this situation when using web search engines to discover content; the proxy does not participate when users arrive from the open web, and the user is denied access to content even though licensed by their organization. Articles and blogs may reference citations or URLs that are in a protected content area; clicking on those will not route the user through the proxy. Users who know that content to which they’ve been denied access is licensed by their organization will be confused, while unaware users will assume they have no access. There is rarely a service-side discovery process that will rescue this class of users and lead them back to their proxy login screen. Usually, the only option is for the user to start their query over from a library-managed page that is coded with proxy syntax-preserving, starting-point URLs.
- Running a proxy server requires server resources and technical configuration skills that are sometimes unavailable at smaller institutions. (This is a disadvantage of SAML IdPs too (see 3.2.4), but there are a greater number of available outsourcing options for SAML than there are for proxy servers).
- Instructors adding deep links to course home pages must write them using “proxy syntax” to ensure they will work. Most people are unaware of this requirement and/or find proxy syntax to be very confusing. Even a tech-savvy instructor may not be aware that content he is linking to through a transparent proxy server is not available on the open web.

Summary

Proxies generally function well, but are now failing in an increasing number of situations, particularly when they are unable to successfully parse some HTML data streams and the growing number of situations where users arrive at a resource without going through the proxy, and consequently are denied access.

3.2.3 Userids/Passwords for a Service Provider Site

Some service providers issue and maintain credential pairs of userid/password. The provider might issue a small pool of credentials or even a single credential for an entire organization. The library would then distribute credentials to researchers with a need to access the resources provided by that publisher.

Another option allows individual users covered by an organizational license to create their own userids/passwords that are then registered with the service provider. Having an individual account provides these individuals with additional functionality (e.g., saving searches from one session to the next).

Advantages

- Someone possessing the userid/password pair can access the service while offsite.
- When people have individual userids, the publisher can provide added value by customizing aspects of the service to the individual and providing mechanisms to recognize a user when s/he returns and to remember what was done during previous sessions.

Disadvantages

- Identity management is not a core competency of most service providers. This management is further complicated because they are working with a set of remote users and do not have any mechanisms to automatically provision users and privileges. Processes for identity management tend to be highly manual and thus labor intensive and costly.
- The user is required to take an additional step (another login) to access each collection of licensed content.
- Each data source and publisher could potentially result in a different userid/password pair, which tends to be frustrating to users.
- Potentially, a login would need to be created for *every* patron of the library for the publisher content site, which could result in a very large pool of accounts. (The alternative is shared userids and passwords, which reduces accountability.)
- Users' frustration with the extra logins could encourage them to use sources other than the library's licensed content for research (e.g., default to what they can get with a Google search). This not only increases the cost/use of licensed content for the organization, it can also result in users missing out on important information for their research.
- Passwords are easily shared by users with other people outside the licensee organization, contributing to unauthorized use of the publisher's content. Users who don't properly secure their userid/password information may inadvertently make them accessible by others. (This is happening much less frequently now, since that same password often controls access to a host of other services for the user who is more conscious about protecting and not sharing such passwords.)
- The organization is required to take suitable steps to ensure that the credentials are not shared with individuals outside the organization and did not leak to the broader internet. Other than user education, this process tends to be a reactive one of dealing with infractions as they occur, which can be very time consuming.
- There is no auto-provisioning or, more significantly, de-provisioning. A userid/password would continue to work after someone leaves the community unless either the library or the publisher deactivated it. Deactivations require the library to be notified of everyone who leaves, which rarely occurs on a systematic and timely basis.
- To ensure all the proper controls are in place and to manage all the user activation and deactivation requires constant maintenance by both the library and the publisher.

Summary

This approach is time-consuming for both library and publisher with an end result that is still a rather insecure authentication method. It is not scalable in today's complex environment. Users also tend to be frustrated resulting in less use of the licensed content.

3.2.4 Federated Login (Authentication)

3.2.4.1 Generic Federated Login

A growing number of service providers have implemented support for federated authentication. In this model the user authenticates once at his/her home organization's identity provider (IdP) site. Subsequent accesses to protected sites, e.g., a publisher's licensed content site, are redirected to the user's IdP. The user's IdP recognizes that the user already has a session and "asserts" to the publisher the privileges associated with this

user. The publisher's site uses the assertion to authenticate the user to the licensed content without requiring the user to re-enter credentials. The user does a single login to the IdP and the system handles subsequent login requests to other service providers in a transparent fashion. This new model provides:

- a) a simplified user experience across a wide range of services because of SSO, and
- b) more effective authentication and authorization of users.

In this manner, federated login supports the model of web single sign-on (SSO) authentication.

This method was first implemented using the proprietary Athens protocol (see 3.2.4.2). More recently, the open Security Assertion Markup Language (SAML) standard has become the preferred way to implement this approach. The IdP would assert a unique opaque identifier to the SP along with a representation of this user's permissions. Some IdPs support maintaining and asserting a unique persistent identifier associated with each user and service provider pair. A service provider could use this value to identify a returning user (without knowing the real world identity) and allow saved searches across sessions and personalization of the service. The IdP can also assert personally identifying information (e.g., a name), although in the library scenarios PII is rarely passed on by the IdP. Generally, the identifier includes assertions about which services the user is authorized to access. The service provider's system then makes an access control decision, based on the information supplied by the IdP.

Campuses and businesses are adopting this model to access a wide variety of in-sourced and out-sourced services. A growing number of cloud-based services supporting the business and instructional needs of campuses offer SAML-based access. Many of the websites supporting the collaborative work of "Virtual Organizations" (research projects) that draw their membership from multiple campuses now rely on SAML-based authentication.

Over the last several years, several protocols and implementations have been used to support the concept of federated authentication. With some protocols, the licensee organization's system (or its agent's system) makes assertions about the user—including, for instance, whether the user is covered by the institution's license and what privileges they have. Sometimes, this additional information allows the service provider to provide a higher level of service. The service provider decides which of the federated protocols to support (e.g., Athens and/or Shibboleth, a SAML-based protocol), based on customer requirements and their comfort level with the security and level of assurance supported by each protocol.

Advantages

- The home organization can utilize the identity management framework and processes that are already in place to track who enters and leaves its community and manage access to other services (e.g., e-mail). Consequently it is already well positioned to authenticate its users and manage their access to licensed resources. As a result, security, privacy, and manageability of license enforcement are enhanced.
- Because the home organization asserts the privileges associated with each individual user, different users can have different privileges and/or have access to different licenses. This supports the community segmentation model and provides the needed flexibility to implement it.
- There is no need for individuals' credentials to be stored or managed by the service providers, relieving them of an often onerous burden. The campus is the obvious place to maintain this information, which must be tracked for various other purposes (including course registration, access to the physical campus and more), rather than attempting to do so at potentially hundreds of content and service providers who work with the campus.
- A browser user can be anywhere in the world (a likely scenario these days) and still be safely and securely authenticated and authorized to the licensed content.
- This model supports web single sign-on (SSO), which provides a simpler user experience across multiple service providers.
- The number of passwords that users must remember is reduced to the single password for the home organization's IdP site. With SSO the user only logs in once per session, regardless of which service providers s/he connects to.

- The home organization no longer needs to support and operate proxy servers or a VPN service for access to licensed or protected content.
- The use of additional optional attributes could provide individual users with a higher level of service.
- The worldwide higher education/research community has standardized attribute definitions; consequently, SP implementations can be standardized.
- Logs at the home organization might provide beneficial usage information at a level of detail that is shielded from the SP.

Disadvantages

- The library and the service provider have to both support the same federated authentication method. Since there is not as yet a single accepted method that everyone is adopting, the library and user could end up with a hybrid environment where some content still requires a separate login or proxy.
- The service provider might have to support multiple federated authentication methods to satisfy different customers.
- As usage has evolved, there are many different ways that a browser user can find his/her way to a publisher site. Federated authentication must be able to authenticate the user in all of these scenarios. Existing implementations cannot always accommodate every possible set-up or navigation path.
- The user interface of current implementations of discovery services (see 5.1) vary tremendously from one service provider to the next (e.g., placement and terminology of the login request) and can thus be confusing to users. A recent focus group study by JISC identified many of these issues, and developed a set of recommendations³.
- The user does have to log in. With IP-based authentication, an on-campus user does not have to log in at all.

Summary

Currently, this methodology comes the closest to offering single sign-on authentication as envisioned by the ESPReSSO working group. However, some changes or improvements are needed to fulfill the group's full vision. Many of the recommendations in this report are targeted at improving the usability of this process.

3.2.4.2 Athens

Athens is an early version of federated authentication centrally that was funded by the UK Joint Information Systems Committee (JISC) for a number of years as the preferred solution for UK higher and further education and is now widely adopted in the UK for both academic and health organizations. It achieved SSO with a proprietary protocol exchange between a vendor-side client library and a centrally run Athens authentication system.

Athens initially acted as an identity provider, allowing an organization to provision identities for their users into the central Athens directory. Organizations had tools to de-provision users when they left the community.

Users initiate an Athens authentication by following a login link from the vendor to the Athens authentication point. If they have an active session, they are immediately granted access; otherwise, a vendor-customized login screen at Athens asks the users to log in with their Athens credentials. Athens also provides a simple portal called "My Athens," which allows the user to move from service to service by connecting with the supported services via menu URLs pointing to the registered entry points.

SPs have tools to integrate the Athens library with their authentication system. Athens now offers organizations the choice of integration with their own IdP or to use the Athens directory. It can also act as a fee-based SAML identity provider (see 3.2.4.3) for institutions not wanting to run their own IdPs. Athens now offers SAML protocols for both IdPs and SPs.

³ Smith, Rhys. *Publisher Interface Study*. Cardiff University and JISC, September, 3 2009. [Note: PDF version has the title of *Service Provider Interface Study, v 0.2.*] Available at: <http://sites.google.com/site/publisherinterfacestudy/>

Advantages

All of the advantages of the generic federated authentication described in 3.2.4.1 apply in addition to the following.

- A single authentication mechanism is used by all the service providers. The user experience is consistent from one site to the next.
- Athens uses organizational attributes rather than IP addresses, which offers more sophisticated subscriber identification than the network topology can offer. In particular, it offers granularity of different subscribers within a single IP address range.
- Athens provides a central support service providing SPs with standardized support.

Disadvantages

All of the disadvantages of the generic federated authentication described in 3.2.4.1 apply in addition to the following.

- Many service providers have implemented Athens but similar support has not been shown by the growing number of cloud-based services and sites supporting collaborative work within the higher education environment.

Summary

While Athens was the first implementation of federated login, it is not seeing adoption by the broad set of service providers beyond licensed content that many campuses now use.

3.2.4.3 SAML

Security Assertion Markup Language (SAML) is a standards-based web SSO protocol promulgated by the OASIS organization. Many vendor and open-source implementations of this standard are readily available. These implementations use metadata exchange that allows SPs and IdPs to know about each other. The metadata is managed by federations of cooperating institutions and authorized vendors. More than 30 higher education/research federations exist around the world to jointly implement SAML⁴. These federations have generally been organized around national boundaries. In the SAML model, service providers run software that “speaks” SAML. Subscribing institutions run (IdP) software that manages the login process and makes assertions describing the user.

There are a variety of ways to initiate a SAML session, depending on whether the user starts at the SP login page or the institutional menu. Library menus typically use WAYFless URLs to connect to an SP. These URLs indicate which institution the browser comes from. This allows the vendor to redirect the user back to the correct IdP if a valid session does not currently exist. If the correct IdP is not known, service providers use either built-in identity discovery services that allow the user to select their home institution, or they redirect users to federation-based, centralized discovery services.

Deep page linking is possible in a SAML-based session, but depending on how the SP-implemented protection scheme works, the user may have to sidestep through the discovery and authentication process before seeing the initially-requested page. Oftentimes, deep-linking is implemented by starting with an OpenURL. The licensee organization’s link resolver can often be configured to redirect the user to the local IdP, and then on to the SP.

The open source Shibboleth implementation of SAML has seen widespread adoption by the higher education and research communities, and is also used by many commercial service providers. There are also several vendor implementations of the SAML standard.

The higher education community has created Federations⁵ as an organizing approach because this model makes processes easier, safer, simpler, and cheaper; it also requires less coding and provides more flexibility.

⁴ See <http://www.terena.org/activities/refeds/> for more information about Research and Education Federations (refeds). See <http://refeds.org/resources.html> for a map showing identity federations that currently work with refeds.

⁵ See http://refeds.org/resources_list.html for a list of the extant federations.

Advantages

All of the advantages of the generic federated authentication described in 3.2.4.1 apply in addition to the following.

- SAML is a mature protocol; there is significant real world experience using it in a wide variety of situations.
- There is widespread adoption of SAML within the higher education/research community, which is a significant customer of library licensed resources.
- A wide variety of interoperable open source and vendor SAML implementations exist.
- In addition to controlling access to licensed resources, many campuses use SAML-based SSO to control access to instructional and administrative applications in order to provide a consistent user experience across a wide range of applications.
- There are standard, widely adopted approaches to common situations. For instance, a standard mechanism and attribute syntax for asserting privileges associated with individual users is available to IdPs and SPs.

Disadvantages

- Managing a SAML environment may require a level of technical skill that some organizations do not possess. (Outsourcing options exist for organizations without the requisite skill set.)

Summary

Going forward, standards-based SAML would appear to be the best choice for implementing a federated login model. Library licensed content can be integrated with the same authentication method that is becoming the preferred protocol in the higher education and research community.

Part 4: ESPreSSO Recommendations

4.1 Overview

These recommendations are intended to define a path forward from the current access control mechanisms (IP address and SP-managed userids/passwords)—which are increasingly problematic—to the next generation approach (federated login) that promises to be more secure, more flexible, and provide better functionality and user experience. This section defines and describes specific recommendations to service providers and licensee organizations designed to provide both types of organizations with a manageable, scalable approach to access control that meets the rapidly evolving, increasingly complex requirements of today's users. The recommendations describe a hybrid environment containing older authentication approaches and newer approaches that are in the early stages of implementation. This is necessary to support the transition to the newer authentication environment.

The majority of the recommendations refer to the newer authentication approaches and are intended to provide a consistent user experience at multiple service provider sites. The recommendations found in this report draw on several years of experience and a variety of approaches. Many of these recommendations were called out in the JISC-sponsored focus group study, published in September 2009. This ESPreSSO recommended practice builds on that study and presents a set of recommendations to both IdP and SP sites. The recommendations specifically address typical browser flows, the sequence of pages presented to users, page layout, what information to include in each of those pages, consistent GUI elements, and additional features and functionality to provide users with added value. The recommendations are intended to:

- Provide users with a consistent experience across a multitude of sites and situations.
- Reduce user confusion and aborted sessions during the identity discovery/login process by using a consistent set of visual elements as the user is transferred between sites in order to reinforce the “this is normal and expected” aspect of the experience.
- Be straightforward and simple to implement for both IdPs and SPs.

Specifically, this report recommends:

- 1) Service providers should continue to support multiple authentication options for the near future. We are in a transition period and must accommodate that some organizations will move faster than others.
- 2) Service providers and licensee organizations should move quickly to reduce reliance on IP-based access control. As described in 3.2.1, there are many security issues with this approach and it is no longer adequate in today's rapidly ubiquitous computing environment.
- 3) Service providers and licensee organizations should move as fast as possible to deprecate userids/passwords that are validated at the service provider site.
- 4) Service providers and licensee organizations should move as quickly as possible to implement and use standards-based federated authentication.
- 5) Service providers and licensee organizations should improve the federated authentication user experience by implementing the recommended federated access user experience that will ensure:
 - a) SPs adopt standard placement/wording of the login link on all the public pages on their site.
 - b) SPs present a standard approach (discovery) for guiding the user to the desired authentication method.
 - c) IdPs create a consistent experience as the user moves from SP to IdP to SP.
 - d) SP and IdP web designers insert branding at appropriate places in the flow to provide visual feedback that the flow is progressing as expected.
 - e) SPs offer a single URL point of access for IP authentication and federated login. SPs should not require use of a new separate URL entry point in order to use federated login.

- 6) Service providers and licensee organizations should move to take advantage of the additional features and functionality supported by the federated authentication model as described in Section 4.9.

4.2 Use Cases

The most common use cases for the recommended federated login approach are listed below. They reference the functional components shown in Figures 1 through 4; Appendix A describes each function in more detail.

- 1) Browser user goes directly to an SP Open Page, clicks a login link, is redirected to the SP Login page, selects a Home Organization, is redirected back to their Institution Login Page, authenticates successfully, is redirected back to the SP Open Page, and is granted various privileges.

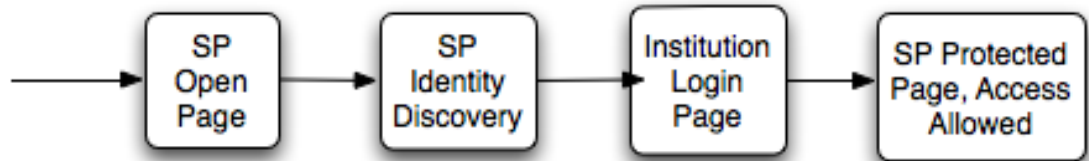


Figure 1: Use case #1 scenario

- 2) Browser user goes to their Institution Menu Page maintained by their home institution, finds the entry for the desired service, clicks a link, is redirected to their IDP, optionally authenticates (if an authenticated session is not already in place), and is redirected to the SP site.

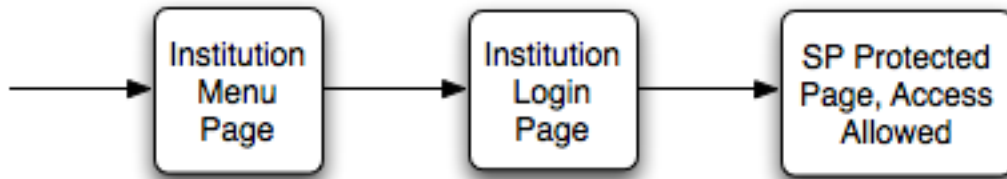


Figure 2: Use case #2 scenario

- 3) Browser user searches an open Provider site with abstracts, clicks an OpenURL link, is redirected to a link resolver run by their home organization (or a contracted third party) which creates a deep link URL, is redirected to their IDP, optionally authenticates (if an authenticated session is not already in place), and is redirected (via the deep link) to a full text article at an SP site.

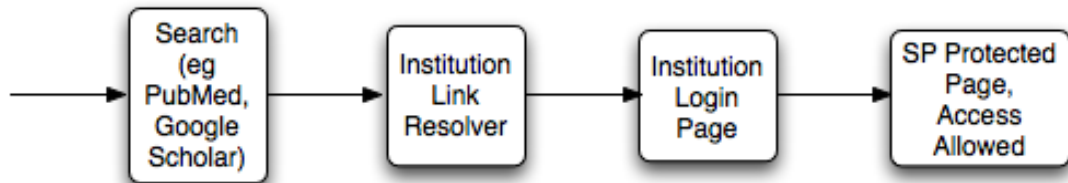


Figure 3: Use case #3 scenario

- 4) Browser user does a search on the open web, clicks a link, and is taken to a deep link at an SP site. Generally, in this scenario (especially in the likely event that the researcher is not within an appropriate IP range), the user will be prompted to authenticate as in the first user case. Successful authentication requires that the institution has configured this SP.

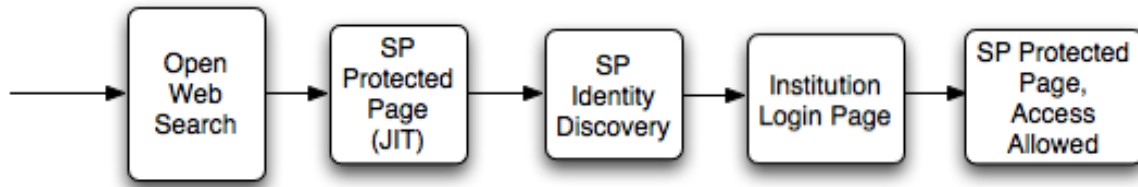


Figure 4: Use case #4 scenario

Appendix A describes the various functions that users may encounter during these use cases in more detail. In many implementations and deployments, these functions are combined into a single component in order to simplify operational issues.

4.3 Summary of Recommendations

This section provides a summary of the changes that all service providers *and* licensee organizations should implement. Successful implementation and achieving a positive user experience require that both types of organizations make some changes.

Table 1 provides, in summary form, the recommendations for each of the functional components. Sections 4.4 through 4.9 provide further details on each of these recommendations for SP and IdP sites.

Table 1: Summary of Recommendations

Functional Component	Recommendation
Service Provider Open Page (see 4.4.1)	<ul style="list-style-type: none"> ▪ SP open pages should display to anonymous users a login link positioned and labeled in the recommended/preferred fashion. This link should lead to an Identity Discovery Page (see 4.4.2) embedded within the SP site. ▪ SP open pages should support IP-based access control.
Service Provider Identity Discovery Page (see 4.4.2)	<ul style="list-style-type: none"> ▪ SPs should provide, within their site, an identity discovery page that offers unauthenticated anonymous users—who require authorization in order to view protected content—with a process for identifying their home organization (i.e., the subscribing institution). <ul style="list-style-type: none"> – As a transition vehicle, this page should allow one or more non-federated mechanisms as well as federated login. – This page should provide an auto-complete search box for finding subscribing institutions by name. – This page should remember the user’s institution selection from previous visits through the use of a cookie and offer a direct or automatic link to the institution login page. (The institution’s logo—obtained from metadata—may be useful for this purpose.) The institution login page is controlled by the selected home organization, not the SP. – This page should <i>not</i> prompt for a username/password since users may not realize that they are not yet at their institution login page. (Users who know they have to use a vendor-provided username and password should select that login preference by following a link to a

Functional Component	Recommendation
	<p>vendor login form.⁶) The GUI flow design of the discovery process should reduce the likelihood of this confusion by offering a vendor login form only to users who explicitly request login via a vendor-provided username and password.</p> <ul style="list-style-type: none"> ▪ SPs may also consider additional helpful discovery features for this page such as: <ul style="list-style-type: none"> – Geo-detection of the user’s IP address to suggest nearby institutions – Searching for institutions by city or postal code – Links to legacy login methods such as Athens or vendor-provided login forms – Ability to change the user’s preferred institution (if a previous selection was remembered via a browser cookie)
<p>Service Provider Protected Page (see 4.4.3)</p>	<ul style="list-style-type: none"> ▪ The SP should continue to first use automatic login techniques when receiving requests for protected pages before determining if SSO authentication is needed for the user. ▪ The SP should, when receiving requests for protected pages that are forwarded to the Identity Discovery Page, return the user to the originally requested page after successful authentication.
<p>Identity Provider/Licensee Organization/ Institution Login Page (see 4.5.1)</p>	<ul style="list-style-type: none"> ▪ The licensee organization should use recognizable institutional branding so users know which credentials to provide. ▪ In addition, the licensee organization should include branding from the publisher site to provide users with visual feedback that the institution login page is a normal part of the flow. (See section 4.8.)
<p>Identity Provider/Licensee Organization Institution Menu Page (see 4.5.2)</p>	<ul style="list-style-type: none"> ▪ If a licensee organization maintains a page listing the databases and journals that have been licensed, along with links to these resources, then the menu should support users even when they are not on the local institution network (and thus ineligible for IP-address based authentication). ▪ WAYFless URLs may simplify the flow for users associated with the institution that provided the menu.
<p>Proxy Server in a Hybrid Environment (see 4.6)</p>	<ul style="list-style-type: none"> ▪ Use an Institutional Switch in conjunction with the proxy server. ▪ Provide the URL of the target SP within the URL query string sent to the proxy server (which has been configured to know about the relevant IdP) rather than providing the appropriate IdP information within the URL query string sent to the target SP.
<p>Rewriting OpenURLs (see 4.7)</p>	<ul style="list-style-type: none"> ▪ Have the link resolver configured to produce URLs that will redirect the user to the Session Initiation endpoint at the SP.

⁶ It has been demonstrated that users cannot always distinguish who is requesting the username and password when presented with a login form. The user might mistakenly type their institutional username and password into a vendor login form. Additionally, presenting a vendor login form to an unsuspecting user could be construed as “phishing” by some institutions.

Functional Component	Recommendation
Appropriate SP and IdP Use of Branding (see 4.8)	<ul style="list-style-type: none"> ▪ SP should insert branding within the Identity Discovery Page and the Institution Login Page. ▪ Licensee organizations should insert branding in an SP’s Identity Discovery Page, if the user has a pre-selected IdP, and within the Institution Login Page.
Additional Functionality (see 4.9 and 4.4.4)	<ul style="list-style-type: none"> ▪ SPs should support pseudonymous access, e.g. with the use of the eduPersonTargetedID tag. ▪ Users should be presented with a “User Consent to Attribute Release” before the institution releases personally identifiable information to the SP.

4.4 Recommendations to Service Providers

4.4.1 Service Provider Open Page

An Open Page at the SP site refers to the provider’s main webpage (i.e., the primary landing spot for users) and to all its publicly available pages. From a license enforcement point of view, the goal for this page is to provide unauthenticated (anonymous) users who have arrived at the page with an easy and consistent way to prove their privileges so the SP can determine which resources they are authorized to see.

Recommendations:

- 1) Present a login link in the upper right hand corner of the interface.

This is, at least anecdotally, the most common location for such a link and should therefore be a familiar place for users to look. While presenting a link instead of a login form on the initial page does require one more click from the user, it allows for utilizing an entire screen’s worth of real estate via the discovery mechanism described in 4.8.

The login link provides a way to proactively authenticate when:

- auto-determination via IP authentication is not possible,
- the user wishes to override the auto-determination via IP authentication (e.g., the user is onsite at another institution), or
- a currently active SSO session has not already been established.

- 2) The button/link should say Login.

This is a term understood by nearly all users, as opposed to things like “Shibboleth,” “Federation,” “Federated,” “Athens,” and other more technical terms often used in this area.

- 3) The Open Page should present a specific indication (when relevant) that the user has an active session.

This session could be a result of:

- auto-determination via IP authentication, or
- previous SSO login during the same session.

This indication can be accommodated with simple text such as: “Access to XYZ provided via Sample University.”

- 4) The login link should be made available in conjunction with just in time (JIT) authentication that intercedes at the point of need.

The login link is clicked by the user, whereas JIT authentication activates automatically when the user attempts access to a protected item. This allows for the most convenient method to be used based on the situation or on user preference.

- 5) Both the login link and the JIT process should present the user with the same unified, identity discovery mechanism (see 4.4.2).

4.4.2 Service Provider Identity Discovery Page

The Identity Discovery Page provides a discovery mechanism for the provider of the identity credentials. The user would arrive at this page after clicking on the login link available on all of the SP's open pages (see 4.4.1) or when attempting to access a protected page (at which point the SP would trigger JIT authentication and redirect the user to this page).

The following recommendations acknowledge that authentication at a protected service usually allows one or more non-SSO mechanisms as well as federated login. Accordingly, the recommendations in this section are to simplify the user experience in the event that any of the automated authentication checks failed.

The goal of these recommendations is to provide as much automated help as possible (e.g., internet location-detection information), to provide real-time suggestions (e.g., dynamic-input search box or display previous selections), and to present the user with a comprehensive palette of authentication options while still being concise with screen real estate. The authentication step will normally be considered an annoyance to the user, but by making a concerted effort to ease the process of getting the user on the right authentication path, the experience need not be more than a quick and minor inconvenience.

Recommendations:

- 1) SPs should deploy an integrated Identity Discovery service within their site.

This approach, in contrast to relying on a centrally managed Identity Discovery service, offers these benefits:

- It allows the SP to brand these pages, providing visual feedback to the user that this page is part of the normal flow.
- If the SP is a member of multiple federations, this approach presents the user with a one-level Discovery Service, rather than multiple-level (e.g., "Select your country first").
- It avoids presenting the user with *all* IdPs within a federation. (Some IdPs may not work if the SP requires non-standard user attributes for authorization or if there is no contract between the SP and an IdP.)
- It allows the service provider to gracefully handle non-Shibboleth institutional authentication sources (e.g., if an EZproxy prefix is known to the service provider).

- 2) The Identity Discovery mechanism should use the following GUI flow after selecting the login link or encountering the JIT authentication:

- a) Allow the user to select his/her Home Organization (detail on recommended GUI elements follows).
- b) Redirect the user to the Institution Login Page that is hosted by the selected Home Organization (see 4.5.1).
- c) Return the user to the originally desired page.

- 3) The Identity Discovery Page should take on certain UI elements (e.g., color and fonts) of the SP site.

Brand the page with the SP's logo to provide the users with visual feedback that they are still within the SP site.

Consistent UI elements between all sites allow users to feel comfortable even if they've never been to that particular SP site before.

- 4) Provide a dynamic-input search box that starts displaying at first character typed for finding

- authentication sources (usually IdPs).
- 5) Specify the authentication source within the search box in “institution name” format.
 - 6) Provide the capability for remembering and forgetting previous authentication source selections, preferably by using the logos for the home institution.
 - 7) Display one or more preferred IdPs as identified in the SP configuration. The preferred IdP should be represented by both an image button bearing the logo of the IdP and a textual link bearing the IdP’s name.
 - 8) Provide an option to display all partner IdPs.
 - 9) Provide a continue button, specifically labeled “Continue” to indicate that there are additional steps to perform.
 - 10) Enable hinting through the use of internet location-detection information to propose geographically proximate subscribing institutions.
 - 11) Include a “More Information” link that leads to information that explains exactly what is occurring and what the user should expect as s/he moves through the login process.

Figure 5 below illustrates a mock-up of an Identity Discovery Page that incorporates the recommendations.

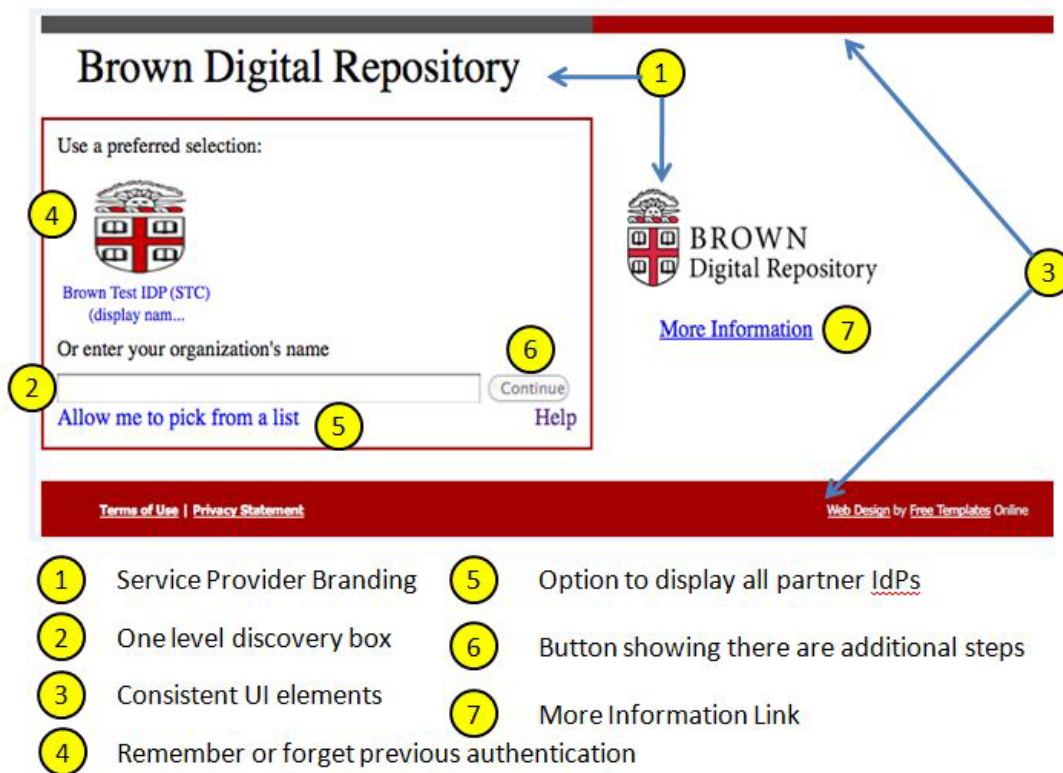


Figure 5: Mock-up of Identity Discovery page using recommendations

4.4.3 Service Provider Protected Page

The Service Provider Protected Page is a page secured for access by only authorized users of a licensee organization. To view content on a protected page, a user must be authenticated by one of the mechanisms described in this document. However, many users may be redirected to this protected page through a variety of methods—ranging from a Google search or use of a discovery tool, to linking in from a reference, DOI®, or

URL listed in another resource—that bypass the Service Provider Open Page and the Service Provider Identity Discovery Page.

Recommendations:

- 1) Requests for Protected Pages should continue to first use automatic login techniques before determining if SSO authentication is needed for the user. Even if that approach succeeds, a Login link should also be visible.
- 2) If an automatic login technique does not succeed, a JIT mechanism should be triggered to redirect the user to an Identity Discovery Page.
- 3) Requests for Protected Pages that are forwarded to the Identity Discovery Page should return the user to the originally requested page after successful authentication.

4.4.4 Attribute-Based Authorization

In the Federated Login Model, an IdP site makes assertions about the current browser user. These assertions contain named attributes and associated values. The higher education/research community has adopted a set of standard practices around attribute syntax and semantics. Virtually all campus and publisher SPs support these practices. This has greatly simplified the deployment and management of federated login for both campuses and SPs. These sites have been using the attributes defined in the eduPerson schema. Specifically, most sites use an entitlement value to represent that “this browser user is covered by the institutional license⁷.”

This model of using entitlement values would also allow a campus with multiple licenses granting different privileges to different sub-populations to easily assert which licenses are associated with each individual user.

The attribute model was designed to be extensible. See 4.9, *Recommendations for Additional Functionality*, for specific examples of how attributes can be used to add value to the user experience.

4.5 Recommendations to Libraries / Institutions

With the implementation of a federated login approach, the Institutional Login Page takes on a new importance. Users attempting to access an SP website supporting federated login will be redirected from the SP site, albeit briefly. To provide the user with a consistent experience, this section offers a number of recommendations regarding the use of Institutional Login Pages and the Institutional Menu Pages. The goal is to improve and simplify the user experience, providing visual feedback to encourage the user to continue.

4.5.1 Institution Login Page

With the Identity Discovery Service embedded in the service provider's website (see 4.4.2), the institution's login page will be the first time the user leaves the SP website. As the name implies, the Institution Login Page is used to authenticate the browser user. After a successful authentication, the browser user is returned to the SP site, accompanied by any attributes and values included with the authentication that describe the user's privileges at this particular SP site.

Recommendations:

- 1) The Institution Login Page should clearly identify the campus using logo images and other identifying marks appropriate to the institution.
- 2) The Institution Login Page should contain the logo, if possible, and the name and description of the relevant SP, to provide the user with visual feedback that this login page is a normal part of the flow. This information will most likely be retrieved from the federated authentication metadata describing the SP. This should help users understand that they are still involved in a process related to accessing that service.
- 3) In addition, the Institution Login Page should also include recognizable institutional branding so users

⁷ The MACE-Dir Working Group maintains registrations for the “entitlement: common-lib-terms” namespace of values that can be used to “indicate that the holder of the entitlement has access to resources under those contract terms.” For more information about common-lib-terms see <http://middleware.internet2.edu/urn-mace/urn-mace-dir-entitlement.html>

know the identity of the site where they will be authenticating.

Figure 6 below illustrates a mock-up of an Institution Login Page that incorporates the recommendations.

1 Service Provider Branding **2** Institutional Branding

Figure 6: Mock-up of Institutional Login Page using recommendations

4.5.2 Institution Menu Page

Most licensee organizations maintain webpages listing the SPs with which they have current contracts. Specifically in the library context, these Institution Menu Pages typically take the form of lists of electronic journals, e-books, and other electronic resources that are hosted externally. The listings may be arranged alphabetically and/or by subject and are often searchable. However they are presented, these listings are important not only because they indicate what content is available to affiliated users, but also because the links to the SPs that they provide can include crucial components in their URLs, such as a "proxy prefix" that initially routes all access through the appropriate proxy server.

Recommendations:

- 1) The entries on the Institution Menu Page should transfer the user first to the local identity provider (to authenticate the user) and then on to the appropriate SP site.
- 2) When transferring to the SP Session Initiation (SI) endpoint, pass as parameters:
 - a) the identity of the home organization (expressed in the manner required by the protocol that will be used), and
 - b) a target parameter describing where the user should be sent after successful authentication.

Figure 7 shows an example of an institution menu page.

The screenshot shows the Brown University Library website. At the top, there is a header with the library name and a search bar. Below the header is a navigation menu with links for 'Library A-Z', 'Off-Campus Access', 'Hours & Locations', and 'Contact'. The main content area is divided into two columns. The left column contains a 'RESEARCH' section with links like 'Ask a Librarian', 'WorldCat / easyBorrow', and 'eJournals A-Z'. Below this is a 'LIBRARY SERVICES' section with links like 'Borrowing', 'Online Forms', and 'Librarians by Subject'. The right column contains a paragraph about online journals, a search filter section for 'Find e-resources by title or identifying number', and a list of links to selected ejournal publisher's websites.

Brown University Library

SEARCH: Josiah WorldCat Articles Videos/DVDs Website

[Library A-Z](#) | [Off-Campus Access](#) | [Hours & Locations](#) | [Contact](#)

Ask a Librarian
 Josiah (Catalog)
 WorldCat / *easyBorrow*
 Databases A-Z
 eJournals A-Z
 eBooks
 Course Reserves (OCRA)

RESEARCH

Getting Started
 Guides: by Subject / Course
 Videos / DVDs
 Instructional Images
 Virtual Reference
 Collections / A-Z
 Citation Management
 Copyright & Fair Use
 Center for Digital Scholarship

LIBRARY SERVICES

Borrowing
 Online Forms
 Librarians by Subject
 Tools & Tutorials
 Giving to the Library
 Visiting
 About the Library

The Library provides access to 50,000 online journals and newspapers. Most are restricted to Brown affiliated users. [Activate your Account](#) to obtain a Brown username. If you have forgotten your username or password, call 863-7277 or the CIS Help Desk at 863-4357. Users accessing these resources from off-campus must use the [WebVPN](#), the [VPN client](#) or the [proxy server](#). If you are experiencing difficulty accessing a title, please send email to ejournals@brown.edu. Please note that [license agreements](#) govern how this material can be used by individuals.

Find e-resources by title or identifying number:
 Limit search to the following content types:
 All | Books | Journals | Other

Title begins with

Browse e-journals by title

0-9 A B C D E F G H I J K L M N O P Q R S T U V W X Y Z Other

Browse e-journals by subject

-- Please select a subject category --

If you already have a citation and want full text, try [Citation Linker](#)

Links to selected ejournal publisher's websites:

- [AGU Digital Library](#)
- [American Chemical Society Legacy Archives](#)
- [American Chemical Society Single Title Subscriptions](#)
- [American Chemical Society Web Editions](#)
- [American Economic Association Web](#)
- [American Physical Society Journals](#)
- [American Physical Society Publications](#)
- [American Phytopathological Society Journal Back Issues](#)
- [American Society of Agronomy](#)
- [American Society of Andrology](#)
- [American Statistical Association Publications](#)
- [Ammons Scientific Publications](#)
- [Austrian Academy of Sciences](#)
- [AVMA Journals Online](#)
- [Bioline International](#)
- [BioOne.1](#)
- [Brepols Journals Online](#)
- [Business Source Complete](#)
- [CAIRN Economics, Social and Political Science](#)
- [CAIRN Free Access Journals](#)
- [CAIRN General Collection](#)
- [Cambridge Journals Online](#)
- [Canadian Meteorological and Oceanographic Society \(CMOS\)](#)
- [Center For Digital Initiatives](#)
- [Charleston Advisor](#)
- [China Academic Journals - Economy & Management \(Series J\) - English](#)
- [CiNii: Open Access Journals](#)
- [Complete List of Russian, Ukrainian, CIS and Baltic Titles](#)
- [Computing Reviews](#)
- [Council for the Development of Social Science Research in Africa \(CODESRIA\) Department of Publications and Dissemination](#)
- [Culturesfrance publications](#)
- [Department of Mathematics at Princeton University Publications](#)
- [Digital Editions from Exact Editions](#)
- [DigiZeitschriften Journals](#)
- [Directory of Open Access Journals](#)
- [Duke Law Journals](#)

Figure 7: Example of Institution Menu Page

4.6 Role of a Proxy Server in Supporting a Hybrid Environment

For the next several years most campuses will likely be managing hybrid environments that include SPs supporting federated login along with some SPs still relying on IP-based authentication. Managing this environment will prove challenging. Librarians will want to reduce the number of places where they are maintaining configuration information about various SPs. Users will want persistent deep link URLs that will continue to work even if the campus moves from using IP authentication to federated login with the SP referenced by the URL.

The InCommon Federation Library Working Group is recommending⁸ the use of an Institutional Switch in conjunction with an Institutional Proxy Server to accommodate these two requirements.

By its very nature, the proxy server type of SSO does not require that any special URL be presented to the service provider as it works via configurations that have been coordinated between the institution and the vendor (either IP ranges or referring URLs) for authentication/authorization. Following that coordinated configuration setup, proxy SSO works by providing the URL of the target SP within the URL query string sent to the proxy server (which has been configured to know about the relevant IdP) rather than providing the appropriate IdP information within the URL query string sent to the target SP. For example:

`http://proxy.someinstitution.edu/login?url=http://www.sampleSP.com`

For proxy SSO users, it is more logical for the URL to go first to the institution's proxy server as opposed to the ultimate SP.

4.7 Rewriting OpenURLs

OpenURL link resolvers have become a critical component within a campus library services infrastructure. They are used to map references of specific resources (e.g., a journal article) to a specific deep link pointing to that resource at an SP with whom the campus has a license. OpenURLs are most often generated by sites holding databases with abstracts. Oftentimes, multiple SPs will hold copies of the same article. OpenURLs allow a site to redirect a browser user to an article's location at a site where the user will be authorized to access it.

A campus using federated login to provide browser users with an SSO experience should have its link resolver configured to produce URLs that will redirect the user to the Session Initiation (SI) endpoint at the SP. The link resolver already knows which campus the user is associated with. Sending the user to the SI endpoint will bypass the Identity Discovery process at the SP and allow for a more transparent flow for the user, with fewer annoying interruptions.

4.8 Appropriate Use of Branding

Branding—for both the SP and IdP sites—is important for conveying to the user that s/he is looking at a page that is part of the natural access and authentication flow to reach the desired resource. If users are presented with a page that offers no context and no hints about the SP or IdP identities, they may conclude that they have been sent down a dead end or be concerned about phishing and decide not to pursue their search.

Recommendations:

Service provider branding should be inserted:

- 1) Within the Identity Discovery Page that is part of the SP site (see 4.4.2).
- 2) Within the Institution Login Page (see 4.5.1) to reinforce to the users that the appearance of the login page is related to their efforts to access information at that service provider.

Institution branding should be inserted:

⁸ See: InCommon Best Practices [for library resource providers and libraries]. Last edited on December 22, 2009. Available at: <https://spaces.internet2.edu/display/inclibrary/Best+Practices>

- 1) In a Service Provider's Identity Discovery Page if the user's browser contains a cookie indicating a previously selected IdP (see 4.4.2).
- 2) Within the Institution Login Page (see 4.5.1), in order to reinforce:
 - the identity of the organization requesting the user's credentials, and
 - that access to the requested resources at the SP site will depend on being included under the institution's contract with the SP.

Figure 5 and Figure 6 illustrate the branding recommendations for the SP Identity Discovery Page and the Institution Login Page, respectively.

4.9 Additional Functionality

The recommendations above cover the basic functionality designed to meet the needs of most users, but individual libraries may require additional functionality. This section discusses two specific types of additional functionality that would be particularly useful to provide.

4.9.1 Pseudonymous Access

Culturally, access to online library-related SPs has been provided in the same tradition of access to brick and mortar libraries: a user can enter the site, search for and read resources, and leave the site without leaving any record of what s/he did. There may be an authorization step at the entry to the site (e.g., present an ID card or present some sort of authorization token to an online provider); however, no record is kept of the user's actions within the site. In recent years, however, some online sites have begun to use persistent cookies stored in the user's web browser to track the user's actions; laws and regulations governing this type of tracking vary from country to country.

With online resources, the lack of any memory of the user prevents provision of some advanced functionality such as saving searches from one session to the next or alerting services. One method of allowing these personalization features within a resource while still preserving most of the individual's privacy is pseudonymous access.

With such access, the SP's system would recognize that "This is User ABC returning to my site," but not have any means of linking the value ABC to the user's real identity. This approach also has the advantages of allowing the SP to track the actions of pseudonymous users (without knowing who they are) to provide such data as how many unique users are accessing their site or usage patterns. For many users, libraries, and service providers, these are acceptable tradeoffs.

SAML-based authentication provides an easy way for institutions and service providers to support pseudonymous access.

Specifically, this can be accommodated via eduPersonTargetedID, which "is a persistent, non-reassigned, privacy-preserving identifier designed to provide a service provider with a unique identifier for a logged in person while preserving the person's privacy. ...each eduPersonTargetedID is unique per person per service provider"⁹

4.9.2 User Consent to Attribute Release

Another method of trading off privacy for functionality is to offer users a "User Consent to Attribute Release" before the institution releases personally identifiable information, which allows users to retain control over which information is released to an SP. SAML-based authentication is an example of one that provides the option for the institution to provide additional attributes describing the user. The assumption is that users would obtain additional functionality as a result of providing additional information about themselves.

A user might choose, for instance, to release an e-mail address to an SP along with the targetedID value. The SP could then track the user's activities over some period of time by linking the activity to sessions associated with

⁹ eduPersonTargetedID [element]. Defined in: eduPerson Object Class Specification (200806). Internet2, June 30, 2008. Available at: <http://middleware.internet2.edu/eduperson/docs/internet2-mace-dir-eduperson-200806.html#eduPersonTargetedID>

that value of targetedID. At regular intervals, the SP might send an e-mail containing news that might be of interest to this user (e.g., recent acquisitions in industry areas where the user regularly does searches).

Clearly, there are privacy issues that the user should be aware of; e.g., if an institution-based e-mail address is used, the SP could probably use the e-mail address to obtain the user's real identity. For many users, however, this may be an acceptable tradeoff.

Part 5: Content Discovery Services

5.1 Content Discovery Services

When libraries started providing electronic catalogs, patrons had only one place to search. As publishers started the transformation from print to e-content, libraries were faced with setting up pages that listed the licensed databases and journals so students and patrons could find what they were looking for. Over time, these hosted lists of databases became un-manageable to keep current and users had to determine for themselves which database to search for specific material. In addition, users were getting familiar with the web search experience where a single search box provided access to everything available. Libraries required a tool that could solve both problems: a solution that could help users discover content wherever it might exist in or outside the library while using a single search box. Vendors responded with multiple solutions: Federated Search, Web Scale Discovery, and combinations of the two.

5.1.1 Overview of Federated Search

Federated search has become very popular in the library setting over the last 10 years as it allows patrons unfamiliar with where content can be found to discover all relevant content available through the library. Federated search allows the researcher to focus on the topic of research and not be concerned as to which are the best databases or resource collections for a particular query. It also allows the user to run one search query across many or all of the available resources simultaneously.

The user simply accesses their organization's federated search page and enters the search term(s) into the search text box to initiate a query. The federated search tool issues the query to all content providers the researcher has selected (or to some pre-selected combination of content resources) and combines and formats the result sets into a single unified result set that is presented to the researcher. When the researcher decides to view content presented by the citation in the search results, the user simply clicks on the visible URL, which directs the researcher's browser directly to the publisher's content presentation service of abstracts or full text, as illustrated in Figure 8.

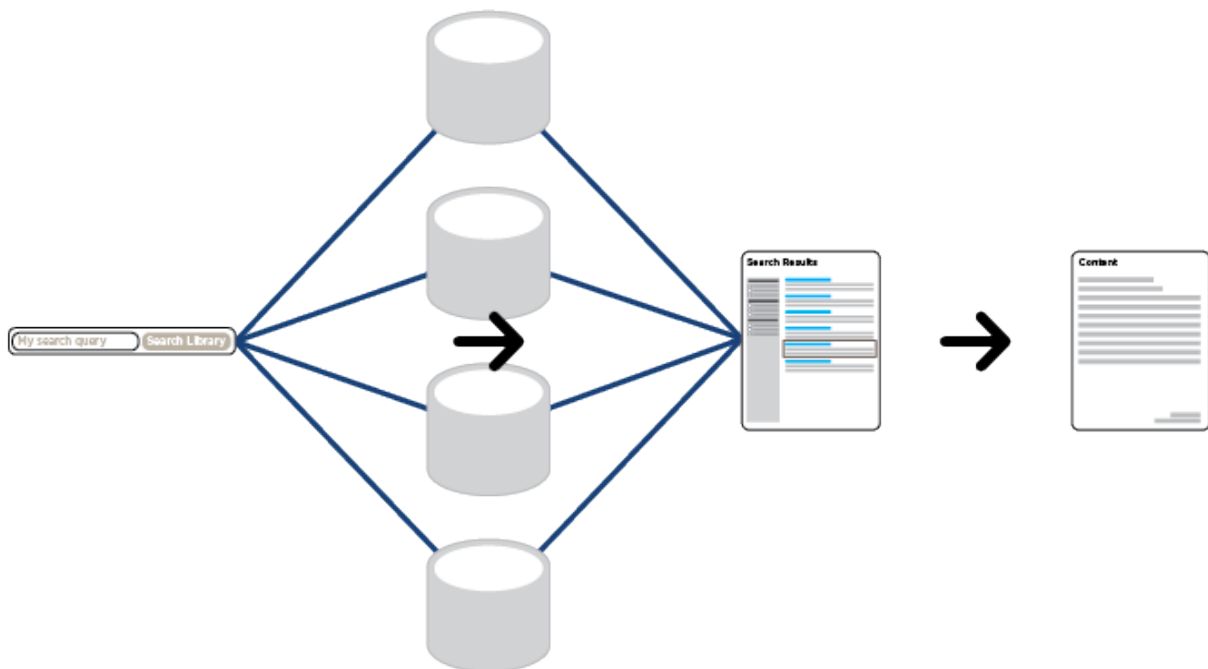


Figure 8: Federated search

Federated search can be a locally installed application, a remotely hosted service, or service provider managed service (such as software as a service or SAAS) that sends search queries to multiple but separate content

providers and retrieves the results to generate an integrated result set. Search queries can be sent to various search interfaces, including HTTP, Z39.50, and XML gateways. Search results are parsed as citations with varying amounts of metadata. Links are provided to access abstract and/or full-text content within the native interface of each content provider.

Federated search has a unique set of challenges in that it performs searches of licensed content acting as an agent for the user. The user interacts with the federated search portal and not directly with the content provider's site when issuing queries and viewing search results. Proper operation requires the content provider, the federated search provider, and the library to all be correctly configured. Federated search supports various authentication methods to ensure that only authorized users are able to send searches and view returned citations. Many libraries use a proxy server in conjunction with a federated search service to authorize user requests to access content provider resources. In addition, in the federated search scenario, there are additional issues of trust between the federated search portal and content providers as well as how the browser user's permissions are represented by the portal to the content providers.

5.1.2 Overview of Web-Scale Discovery Services

Web-scale discovery services have for many libraries started to replace federated search as a mechanism for allowing people to easily discover content. Web-scale services allow patrons unfamiliar with where content can be found to discover all relevant content available in the library. In essence, web-scale discovery outwardly provides the same feature set as federated search, but internally does the job in a very different manner.

A web-scale discovery service can be a locally installed application, remotely hosted service, or service provider managed service (software as a service or SAAS) that queries a single index that represents multiple but separate content providers and retrieves the results as an integrated result set. Unlike federated search, the queries are never sent to the actual provider of the content. Search results are parsed as citations with varying amounts of metadata. And like federated search, links are provided to access abstract and full-text content within the native interface of each content provider. Web-scale search offers the researcher three key benefits over federated search:

- Instantaneous search results
- Consistency in relevance ranking of results through normalization into a single unified index
- Complete results unlike federated search that returns subsets of results from many providers (Federated search must typically limit the number of results returned from a large result set due to performance issues.)

To use web-scale discovery, the user simply accesses the organization's Search page, enters a search term into the search text box, and initiates a query. The search engine searches the single index and ranks the results that are then presented to the researcher. To view content presented by a citation in the search results, the user clicks on the visible URL that directs the researcher's browser directly to the publisher's content presentation service, as shown in Figure 9.



Figure 9: Web-scale discovery search

Since web-scale discovery services don't send queries to service providers, no loads are created on the provider's equipment. Because of this, libraries are free to allow researchers to search at will without logging into the search tool or library website. In essence, anyone on the Web can use the search page to determine what exists in a library's collection. The only time that a researcher must use some type of authentication service is when actually attempting to view content hosted at the provider's site.

Web-scale search performs searches of licensed content acting as an agent for the user. The user interacts with the search portal and not directly with the content provider's site when issuing queries and viewing search

results. Proper operation requires the content provider and the library to be correctly configured so content can be viewed. Note that this is distinctly different from federated search in that queries are never sent to the content provider. Login should only occur once when the researcher attempts to view content from a provider for the first time. Many libraries use a proxy server in conjunction with the web-scale search service to authorize user requests to access content provider resources.

5.2 Existing Authentication with Discovery Services

Discovery services use the same authentication methods as described in section 3.2 with the same advantages and disadvantages.

5.3 Recommendations for Authentication in a Discovery Search Environment

With content discovery services, the user is authenticating to a portal; the portal may access backend services acting as the user's agent. However, the user does not authenticate to the backend services. Consequently, the recommendations for authentication are the same as those found in Section 4.

Recommendations:

- 1) Federated search services should require a single login to access search if the user is remote to eliminate excessive traffic from users that are not members of the library. From within the library, login should only happen during the first request for content. After the initial login has occurred, no more requests for login should occur.
- 2) Web-scale discovery services should not require a user to login to search. Login should only occur during the first request for content. After the initial login has occurred, no more requests for login should occur.

Appendix A Description of Functions in Current Authentication Environments

Figure 10 depicts the typical functions that users may encounter while using today’s authentication environment. Each of these components is described in this Appendix.

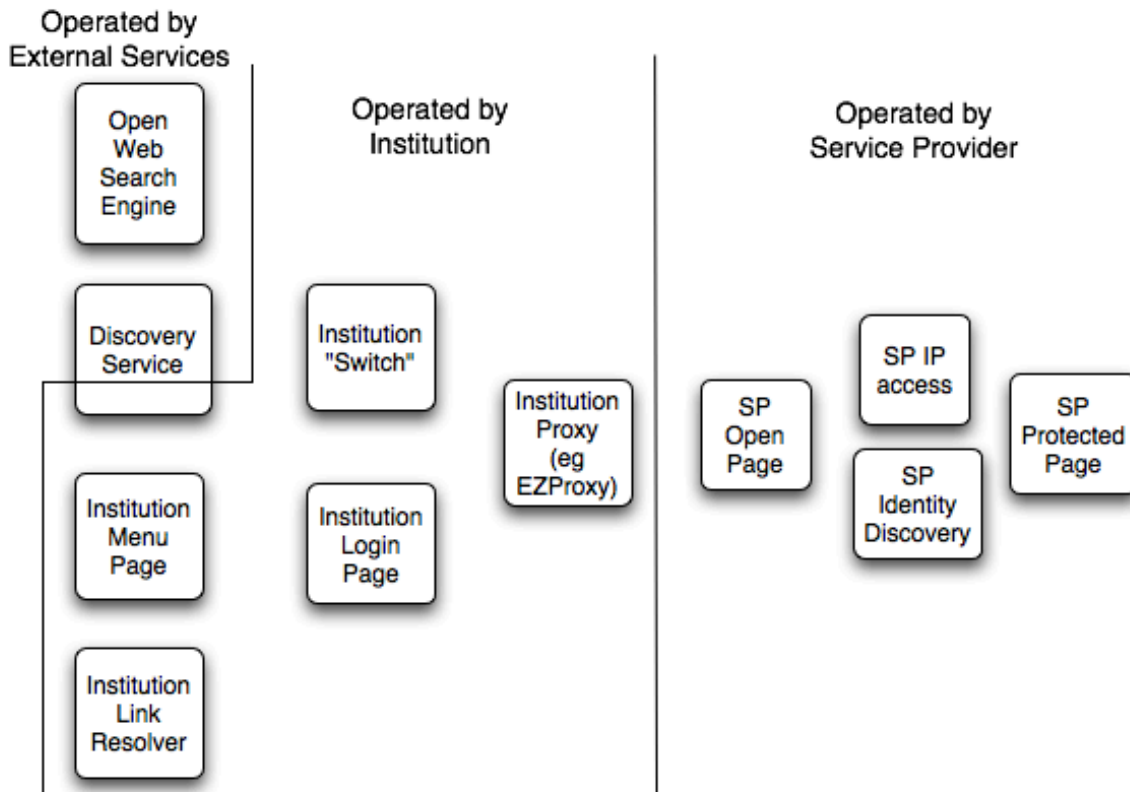


Figure 10: Functional components of current authentication environments

Functional Component

Description

Open Web Search Engine

Any of the search engines currently available on the World Wide Web. Some of these search engines index the content held by commercial service providers in protected areas, e.g., journal articles, and will generate search results containing deep links to those articles. Some search engines, such as Google Scholar, offer search results that use OpenURL syntax and can point to the link resolver at each scholar’s home campus to provide access to licensed content.

Discovery Service

Term that collectively describes for the federated search and web-scale discovery services described in Section 5.1.

<u>Functional Component</u>	<u>Description</u>
Institution Menu Page	A webpage maintained by the licensee organization that lists the databases and journals that have been licensed, along with links to these resources.
Institution Link Resolver	A service run by a licensee organization that accepts incoming URLs encoded in OpenURL syntax and maps those references to deep link URLs at a content provider holding the referenced resource and with whom the licensee organization has a contract allowing access to that resource.
Institution Login Page	A webpage presented by the licensee organization's identity provider service (IdP) that collects information used to authenticate the browser user.
Institution Switch	<p>A service run by the licensee organization which knows which service providers offer SSO-based access and which continue to rely on IP-authentication. Based on that information, the Switch redirects a browser user either: 1) directly to the service provider (SSO), or 2) to the licensee organization's proxy server.</p> <p>In alternative #2, if the user's IP address is within the ranges that the service provider will recognize, the Switch may forward the user transparently to the service provider.</p>
Institutional Proxy	A service at a licensee organization that authenticates users not directly connected to the licensee organization's network (e.g. they are at home, or traveling) and then authorizes the user to access licensed resources "through" the proxy server.
Service Provider Open Page	A webpage presented to the user by the service provider that can be accessed by anonymous users. An SP's landing page usually falls into this category.
Service Provider IP Access	A mechanism used by an SP uses to obtain a browser user's IP address (or the apparent address provided by a proxy server), and determine whether that address is within one of the ranges used by a customer licensee organization. If it is, the browser user is granted appropriate access.
Service Provider Identity Discovery	A webpage presented by service providers that allows users to select their home organization and be redirected there for authentication. This page was previously called a WAYF (Where Are You From) Service.
Service Provider Protected Page	A webpage presented by a service provider that can only be accessed by browser users who have been granted permission to see that page through some authentication process. The contract between the SP and licensee organization usually specifies which pages and content members of the licensee organization are authorized to view.

Bibliography

ANSI/NISO Z39.50-2003 (R2009), Information Retrieval: Application Service Definition & Protocol Specification. Baltimore, MD: National Information Standards Organization, 2003. Available at: <http://www.niso.org/standards/z39-50-2003/>

ANSI/NISO Z39.88-2004 (R2010), The OpenURL Framework for Context-Sensitive Services. Baltimore, MD: National Information Standards Organization, 2004. Available at: <http://www.niso.org/standards/z39-88-2004/>

eduPerson & eduOrg Object Classes [website]. Ann Arbor, MI: Internet2. Available at: <http://middleware.internet2.edu/eduperson/>

EZproxy® authentication and access software [website]. Dublin, OH: OCLC. Available at: <http://www.oclc.org/ezproxy/>

InCommon Library Services Collaboration [website]. Ann Arbor, MI: Internet2. Available at: <https://spaces.internet2.edu/display/inclibrary/InC-Library>

Klingenstein, Nate, and Scott Cantor. *Understanding Shibboleth*. Ann Arbor, MI: Internet2, March 10, 2011. Available at: <https://wiki.shibboleth.net/confluence/display/SHIB2/UnderstandingShibboleth>

OASIS Security Services (SAML) TC [website]. Burlington, MA: Organization for the Advancement of Structured Information Standards (OASIS). Available at: http://www.oasis-open.org/committees/tc_home.php?wg_abbrev=security

OpenAthens Access and Identity Management [website]. Bath, UK: Eduserv. Available at: <http://www.eduserv.org.uk/aim>

SAML XML.org: Online community for the Security Assertion Markup Language (SAML) OASIS Standard [website]. Burlington, MA: Organization for the Advancement of Structured Information Standards (OASIS). Available at: <http://saml.xml.org/>

Security Assertion Markup Language (SAML) v2.0. Burlington, MA: OASIS Security Services TC, March 2005. Available at: <http://www.oasis-open.org/standards#samlv2.0>

Shibboleth® [website]. Ann Arbor, MI: Internet2. Available at: <http://www.shibboleth.net/>