RA21

WAYF Cloud and P3W
Security & Privacy Recommendations

July 2018

Security & Privacy Working Group

Elias Balafoutis (Atypon)
Reggie Zamora (Wiley)
Lisa Janicke Hinchliffe (University of Illinois)
Sari Frances (IEEE)
Scott McCarthy (Proquest)
Dan Ayala (Proquest)
Meltem Dincer (Wiley)
Ken Ferris (Taylor and Francis)
Peter Reid (Bath Spa University)
Todd Carpenter (NISO)
Christian Pruvost (Elsevier)
Andy Sanford (EBSCO)
Ralph Youngen (ACS)
Heather Ruland Staines (Hypothes.is)
Adam Snook (OpenAthens)
Richard Northover (Elsevier)
Chris Shillum (Elsevier)
Joe Greene (CAS)
Jos Westerbeke (University of Rotterdam)
Will Simpson (ORCID)
Paul Dixon (LibLynx)
Phil Leahy (OpenAthens)
David Orr (OpenAthens)
Diane Cogan (Ringgold)
Laird Barrett (Springer Nature)

Executive Summary

I. Introduction

Publishers, libraries, and consumers have all understood that authorizing access to content based on IP address no longer works in today's distributed world. The RA21 project hopes to resolve some of the fundamental issues that create barriers to moving to federated identity in place of IP address authentication by looking at some of the products and services available in the identity discovery space today, and determining best practice for future implementations going forward. www.ra21.org

II. Objective

This document offers a purely technical analysis of the security and privacy risks associated with the Cloud WAYF and P3W pilots. The social and policy-related aspects of privacy are not within the scope of this analysis. The desired outcome is to provide recommendations tailored to mitigate risks identified for each pilot.

III. Analysis

Security threat analysis was performed for each pilot based on the threat model included in the appendix. Risks were estimated based on the threat and impact levels for the given scenarios, and countermeasures were identified. Recommendations for each pilot are based on the countermeasures. All security threats identified were deemed low priority and risks can be mitigated by applying standard security and data protection practices.

For data privacy risks, a data protection impact assessment ("DPIA") was performed compliant with GDPR requirements to determine if "high risks" were involved. The first step involved using UK's Information Commissioner's Office (https://ico.org.uk/) GDPR screening checklist to determine if a data privacy impact assessment is required. Since the pilots met three of the required criteria, a DPIA was performed. Results from the DPIA indicated that there were no "high risks," and the same results were used as the basis for the privacy risk recommendations.

IV. Conclusion

There are no significant risks which prevent the WAYF Cloud and P3W pilots from moving forward. Residual risks from both a security and privacy perspective are LOW. The nature of the data involved is low value, i.e., not directly or easily attributable to any natural person, and appropriate safeguards are in place to mitigate confidentiality concerns.

Below is a summary matrix of recommendations per pilot:

S&P RECOMMENDATIONS	Cloud WAYF	P3W
Privacy Policy/Opt-In	√	√
Data Protection Impact Analysis	√	√
Data Retention Policy	√	
Denial of Service Protection	√	√
Browser Security (https + access controls)	√	√
Database/Data Protection	√	
Server hardening	√	√
Security Code Scanning	√	√
Vulnerability Scanning/Penetration Testing	√	√
API Security	√	
Audit Logging	√	
Security Monitoring	√	
Incident Response Plan	√	
High Availability Infrastructure	√	√
Anti-Virus Software	√	√
GDPR Compliance	√	

Overview of the Pilots

Both pilots were established to test the emerging user experience for federated access to scholarly content that is being developed by the RA21 initiative. There are two critical steps to this user experience:

1) Identity provider discovery (allowing the user to search for his/her home institution in order to perform a login).
2) Identity provider persistence (storing the user's identity provider choice to avoid repeating the discovery step).

The two pilots tested different technical architectures for the implementation of identity provider persistence. In addition, the P3W pilot also prototyped a means of providing a central identity provider discovery service. The WAYF Cloud pilot assumed that identity provider discovery would happen on each individual publisher's site and did not prototype a central discovery service.

WAYF Cloud

Following is a very short summary of the WAYF Cloud pilot design. Please refer to the WAYF Cloud architecture document for a detailed description. (https://drive.google.com/file/d/1XI3dqWHLFLerpJypKyTRhsrZchysF556/view?usp=sharing)

APIs

The following picture illustrates at a high level the data flow between the different elements of the WAYF Cloud architecture:

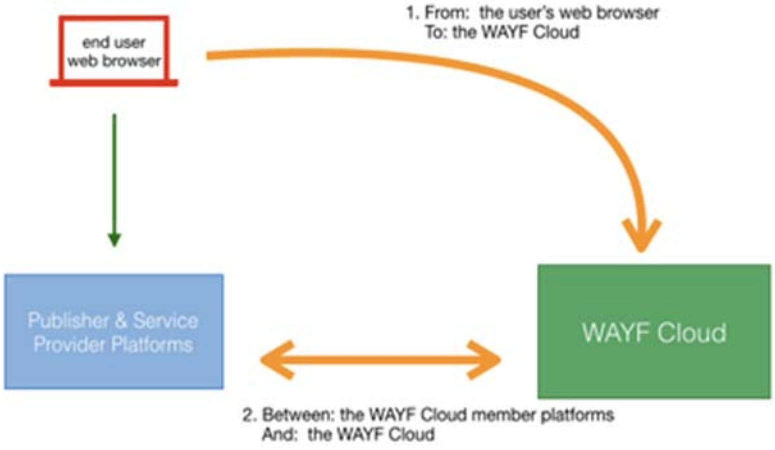
end user
web browser
1. From: the user's web browser
To: the WAYF Cloud
Publisher & Service Provider Platforms
WAYF Cloud
2. Between: the WAYF Cloud member platforms
And: the WAYF Cloud

- SP access to the WAYF Cloud API is protected via client ID/secret pair
- Access to the WAYF Cloud admin interface is protected via username/password
- WAYF Cloud Widget requests to the WAYF Cloud API are pre-authorized via the WAYF Cloud server (REST) API
- User access to the WAYF Cloud UI is authorized using a cookie

Data Model

The following picture illustrates the WAYF Cloud data model:

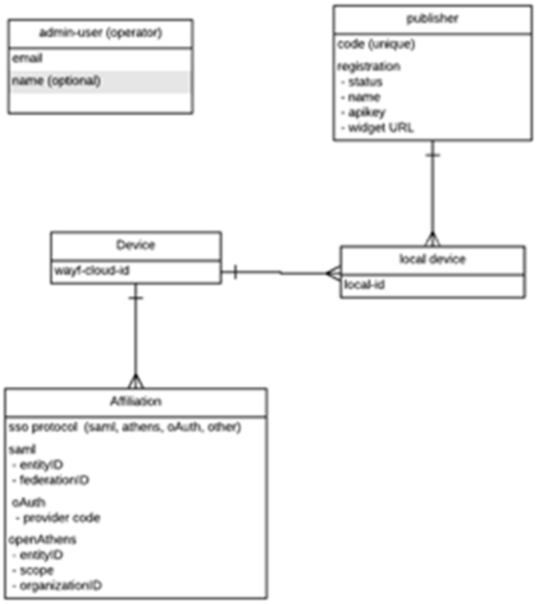

The data exchanged between Service Providers and the WAYF Cloud is a subset of this data model, which looks as follows:

```
POST: https://wayf-cloud-sandbox.literatumonline.com/1/device/<wayf-local-id>/history/idp
{□
  "entityId": "sample-saml-entity-id",□
  "federationId": "sample-saml-federation-id",□
  "type" : "SAML"
□}
```

P3W

Below is a picture that illustrates at a high level the data flow between different elements of the P3W Architecture:

P3W Data Flows

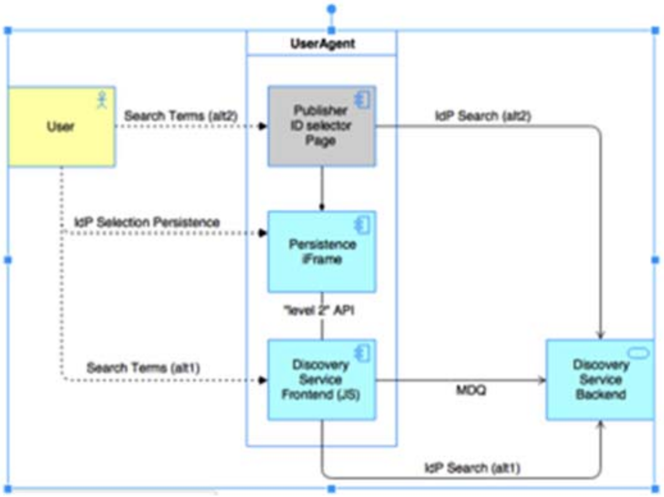

P3W supports the following levels of implementation:

Level 1: Centrally hosted central discovery service. The simplest implementation is that of a completely external discovery service. A service provider would point to a single, external URL.
Level 2: Local instantiation of a central discovery service. Rather than have the entirety of the identity discovery service be external to the service provider domain, a copy of the metadata in the discovery service exists within the service provider's domain.
Level 3: Login happens on SP site by integrating the rendered HTML to the discovery stage. Additional integration is possible by a service provider through integration of the identity discovery service UX into the SPs website.
Level 4: Rendering all the UI and interrogating the JSON object provided by a central service. For sites that have complex requirements around identity (such as the need to support both local accounts and federated accounts within the same interface) and significant IT resources, full integration is definitely possible.

Please refer to https://docs.google.com/spreadsheets/d/157v2A5MR1HgR88YZ8M-xZ0jMvwKNJ5Tau_GmBo_mBFA/edit?usp=sharing for further details on the P3W pilot.

Security Threat Assessment

Threat assessment was performed for the WAYF Cloud and P3W using Microsoft's STRIDE Model Classification to validate coverage of well-known threat vectors. The model represents possible security threats based on the proposed architecture (see Appendix A). For each theoretical threat, the damage and the attack effort are analyzed. The results are documented in tables that describe the threat and corresponding countermeasure. Recommendations for the pilots are based on the countermeasures identified.

STRIDE Threat Model Classification

STRIDE is a classification scheme for characterizing known threats according to the kinds of exploit that are used (or motivation of the attacker). The STRIDE acronym is formed from the first letter of each of the following categories:

Spoofing Identity "Identity spoofing" is a key risk for applications that have many users but provide a single execution context at the application and database level. In particular, users should not be able to become any other user or assume the attributes of another user.

Tampering with Data Users can potentially change data delivered to them, return it, and thereby potentially manipulate client-side validation, GET and POST results, cookies, HTTP headers, and so forth. The application should not send data to the user, such as interest rates or periods, which are obtainable only from within the application itself. The application should also carefully check data received from the user and validate that it is sane and applicable before storing or using it.

Repudiation Users may dispute transactions if there is insufficient auditing or recordkeeping of their activity. For example, if a user says, "But I didn't transfer any money to this external account!", and you cannot track his/her activities through the application, then it is extremely likely that the transaction will have to be written off as a loss.

Therefore, consider if the application requires non-repudiation controls, such as web access logs, audit trails at each tier, or the same user context from top to bottom. Preferably, the application should run with the user's privileges, not more, but this may not be possible with many off-the-shelf application frameworks.

Information Disclosure Users are rightfully wary of submitting private details to a system. If it is possible for an attacker to publicly reveal user data at large, whether anonymously or as an authorized user, there will be an immediate loss of confidence and a substantial period of reputation loss. Therefore, applications must include strong controls to prevent user ID tampering and abuse, particularly if they use a single context to run the entire application.

Also, consider if the user's web browser may leak information. Some web browsers may ignore the no-caching directives in HTTP headers or handle them incorrectly. In a corresponding fashion, every secure application has a responsibility to minimize the amount of information stored by the web browser, just in case it leaks or leaves information behind, which can be used by an attacker to learn details about the application or the user, or to potentially become that user.

Finally, in implementing persistent values, keep in mind that the use of hidden fields is insecure by nature. Such storage should not be relied on to secure sensitive information or to provide adequate personal privacy safeguards.

Denial of Service Application designers should be aware that their applications may be subject to a denial of service attack. Therefore, the use of expensive resources such as large files, complex calculations, heavy-duty searches, or long queries should be reserved for authenticated and authorized users and not available to anonymous users.

For applications that do not have this luxury, every facet of the application should be engineered to perform as little work as possible, to use fast and few database queries, to avoid exposing large files or unique links per user, in order to prevent simple denial of service attacks.

Elevation of Privilege If an application provides distinct user and administrative roles, then it is vital to ensure that the user cannot elevate his/her role to a higher-privilege one. In particular, simply not displaying privileged role links is insufficient. Instead, all actions should be gated through an authorization matrix to ensure that only the permitted roles can access privileged functionality.

WAYF Cloud

Threats Identified	**Risk Priority**	**S**	**T**	**R**	**I**	**D**	**E**
1. User's data can be accessed by unauthorized users with physical access to the user's device	4	S1	T1	R1	I1	D1	E1
2. Lost or broken API Keys can be used by an unauthorized SP to access the WAYF Cloud API	4	S2	T2	R2	I2	D2	E2
3. The data exchanged between the WAYF Cloud server and SP platforms is intercepted	4	S3	T3	R3	I3	D3	E3
4. Denial of Service - WAYF Cloud Server	4	S4	T4	R4	I4	D4	E4
5. The WAYF Cloud server can be completely compromised by the attacker		S5	T5	R5	I5	D5	E5
6. Elevation of privileges which gives access to WAYF Cloud data	4	S6	T6	R6	I6	D6	E6
6. Repudiation of actions taken by user	4	S7	T7	R7	I7	D7	E7

STRIDE Threat	**Mitigation**
S1	Require opt-in before cookie is stored per privacy policy
T2	Apply API security best practices
T5	Infrastructure hardening and security monitoring
I3	Use secure protocols such as HTTPS
D4	DDOS services subscription
E6	Install anti-virus/anti-malware software
R7	Implement audit logging

P3W

RA21 Pilot - P3W Threats	Risk Priority	S	T	R	I	D	E
1. Data stored in the user's local storage can be manipulated by unauthorized users that have physical access to the user's device.	4	S1	T1	R1	I1	D1	E1
2. User's data stored in the local storage of the device can be accessed by unauthorized users	4	S2	T2	R2	I2	D2	E2
3. Access to the user's local storage by unauthorized SPs.	4	S3	T3	R3	I3	D3	E3
4. Denial of Service – Directory Service	4	S4	T4	R4	I4	D4	E4
5. The P3W server can be compromised by the attacker.	4	S5	T5	R5	I5	D5	E5
6. Elevation of privileges which gives access to P3W data	4	S6	T6	R6	I6	D6	E6

Threat	Mitigation
I1	Provide privacy policy
I2	Require user Opt-in before any data is stored
T3	Privacy policy
T5	Hardening of Infrastructure, Monitoring
D4	DDOS services subscription
E6	Install anti-virus/anti-malware software

WAYF Cloud

Client Threats

Risk	Priority 4
Threat	Disclosure User's data stored in the WAYF Cloud can be accessed by unauthorized users that have physical access to the user's device.

Impact	Low The attacker can view the IdPs that the user has successfully signed in to in the past.
Effort	Medium Physical access to the user’s device is required. Since no user data is stored in the user’s device, the attacker needs access to the cookie stored in the device in order to access the user data by visiting an SP.
Countermeasures	- Require user opt-in once before the cookie is stored per privacy policy

Server Threats

Risk	Priority 4
Threat	Masquerading, Manipulation, Insertion, Destruction, Disclosure Lost or broken API Keys can be used by an unauthorized SP to access the WAYF Cloud API.
Impact	Low Worst case impact is the degradation of the user experience to “first time use” and the disclosure of a list of IdPs (anonymously).
Effort	High Brute force attack, or theft of API keys. In addition to the API Keys, the attacker needs also get the SP specific device IDs in order to make any use of the API.
Countermeasures	- Apply API security best practices: - Blacklist API clients after consecutive authentication failures - Implement process for API key refresh - Protect user data (encryption/salting)

Risk	Priority 4
Threat	Eavesdropping, Disclosure

Impact	Low The data exchanged between the WAYF Cloud server and SP platforms is of low risk/value. The attacker can get access to IdP list, but the data subject cannot be identified.
Effort	High Intercept network traffic between the WAYF Cloud server and SP platforms.
Countermeasures	- Use HTTPS

Risk	Priority 4
Threat	Denial of Service
Impact	Low UX degrades to the ”first time” experience. The ability of user to access institutional resources is not affected.
Effort	High Denial of Service (DoS) attacks are typically performed by professionals.
Countermeasures	- Implement DoS protection measures - Use high availability deployment

Risk	Priority 4
Threat	Masquerading, Manipulation, Discloser, Insertion, Destruction□ The WAYF Cloud server can be completely compromised by the attacker.
Impact	Low UX degrades to the ”first time” experience. The ability of user to access institutional resources is not affected. Relationship between IdPs and users is disclosed. The users cannot be identified.
Effort	High

	Exploit infrastructure weaknesses to take full control of the WAYF Cloud domain.
Countermeasures	- Hardening of the infrastructure - Monitoring

Threat	Elevation of Privilege
Impact	Low The ability of user to access institutional resources is not affected. Relationship between IdPs and users is disclosed. The users cannot be identified.
Effort	Medium Elevation of privilege requires obtaining the privileged credentials by first compromising a non-privileged account then running hacking tools to further exploit the access.
Countermeasure	Install anti-virus/anti-malware software to detect hacking tools
Risk	Priority 4
Threat	Medium
Impact	User may claim that they did not take action that was performed, i.e., entered credentials, email address, accessed data, etc.
Effort	Low
Countermeasures	- Implement audit logging

Summary & Recommendations

The WAYF Cloud has a larger attack surface on the server side, due to the server API and centralized storage. However, the effort required to perform an attack is high, provided that standard security practices are followed.

The WAYF Cloud is a passive element in the process of user authentication and authorization and its availability doesn't affect the ability of users to authenticate at a Service Provider. Given the nature of the data involved, the damage of a successful attack is small as long as sufficient data protection measures are taken.

In a worst-case scenario, damage is limited to the degradation of user experience and to the disclosure of Identity Providers and their association with anonymous entities.

All threats are low priority (Priority 4) and the risks can be mitigated by applying standard security and data protection practices.

Security Recommendations

1. All browser traffic should use secured protocols, such as https, to prevent unauthorized access and to preserve confidentiality.
2. Database used for WAYF Cloud should be encrypted to safeguard information against unauthorized access and data breaches. (Already performed.)
3. Servers used to support the WAYF Cloud should be hardened.
4. Code used for WAYF Cloud should be scanned for vulnerabilities (OWASP Top 10, etc.).[1]
5. A vulnerability scan and/or penetration test should be conducted to validate security of the WAYF Cloud system.[1]
6. Any access to the WAYF Cloud API should be secured.
7. Audit logging should be enabled to log all activities related to the security of the systems and any direct data access (outside of API).
8. Administrative access to the WAYF Cloud systems should be strictly limited.[1]
9. Transfer of any credentials during the creation of a new publisher/registration approval should be secured.
10. Security monitoring should be in place to prevent/detect breaches.
11. An incident response plan should be in place.

[1] Security Best Practices Recommendations

P3W

Client Threats

Risk	Priority 4
Threat	Insertion, Destruction, Manipulation Data stored in the user’s local storage can be manipulated by unauthorized users that have physical access to the user’s device.
Impact	Low Degradation of the UX to the “first time” experience on the event that the user’s local storage contents have been manipulated and/or deleted.
Effort	Medium Physical access to the user’s device is required.
Countermeasures	Privacy policy

Risk	Priority 4
Threat	Disclosure User’s data stored in the local storage of the device can be accessed by unauthorized users.
Impact	Low The attacker can view the IdPs that the user has successfully used in to sign in in the past.
Effort	Medium Physical access to the user’s device is required.
Countermeasures	- Require user opt-in once before any data is stored

Risk	Priority 4
Threat	Disclosure, Destruction, Manipulation Access to the user’s local storage by unauthorized SPs
Impact	Low

	1. Degradation of the UX to the “first time” experience on the event that the user’s IdP cache has been altered. 2. The attacker can get access to the user’s IdP list. The data subject cannot be identified solely by the data stored in the local storage.
Effort	Low 1. Client integration of P3W JS cannot be secured to restrict access to SPs. 2. XSS attacks to non-secured SPs can result to the attacker getting access to the shared storage.
Countermeasures	- Use P3W Implementation Levels 1 to 3 - Require user opt-in before any data is stored in the local storage

Server Threats

The server threats noted below have been included in the P3W section of this report since that pilot prototyped a central discovery service. However, these server threats would be relevant to any implementation of a central discovery service regardless of the architecture used for identity provider persistence.

Risk	Priority 3
Threat	Denial of Service
Impact	High The ability of users to sign in to their IdP and access institutional resources depends on the availability of the discovery service (affects implementation Levels 1, 2, and 3).
Effort	High Denial of Service (DoS) attacks are typically performed by professionals, but they are also possible at low cost for non-experts.
Countermeasures	- Implement DoS protection measures - Use high availability deployment - Use the Level 4 implementation option

Risk	Priority 3

Threat	Masquerading, Manipulation, Disclosure, Insertion, Destruction□ The discovery service can be compromised by the attacker.
Impact	High Hijacking of the discovery service can be exploited to perform user password phishing by directing users to fake IdPs.
Effort	High Exploit infrastructure weaknesses to take full control of the P3W domain.
Countermeasures	- Hardening of the infrastructure - Monitoring

Summary & Recommendations

P3W has lower attack surface on the server side as it doesn’t utilize server APIs or central storage.

Its role in the process of user authentication and authorization, is that of an external SAML discovery service (except for integration Level 4).

The availability of the discovery service can affect the ability of users to sign in at the SP using their institutional IdP and ultimately get access to institutional resources. In addition, the compromise of the discovery service can be used to perform password phishing attacks by directing users to fake IdP pages. This risk is not, however, inherent in the IdP persistence solution, whichever architecture is selected.

The P3W Implementation Level 4 utilizes the browser local storage as common storage across different SP domains. Client-side cross-domain storage is generally considered insecure and its restricted by the web browser. Furthermore, the security of user’s data stored in such storage is that of the least secure Service Provider (“SP”) which is integrated with P3W. This is because Cross Site Scripting attacks at less secured SPs can be used by an attacker to get access to the user’s local storage which is shared across SPs

The threats for P3W are priority 3 and 4. The risks can be mitigated by using a highly available and secure infrastructure. Due to the low impact of service loss (users will be required to login again), no mitigation is required to address any availability concerns.

Following is a more detailed list of recommendations:

Security Recommendations

1. Access for the Discovery Service Frontend (JavaScript) should be limited to its own data.
2. All browser traffic should use secured protocols such as https to prevent unauthorized access and to preserve confidentiality.
3. Servers used to support the Discovery Services Backend should be hardened.[3]
4. Code used for Discovery Services should be scanned for vulnerabilities (OWASP Top 10, etc.).[1]
5. A vulnerability scan and/or penetration test should be conducted to validate security of the Discovery Services server(s). [1]
6. Install anti-virus/anti-malware software

Privacy Analysis

A GDPR based Data Protection Impact Analysis ("DPIA") was performed to determine privacy risks for the WAYF Cloud and P3W pilots (see Appendix C). Based on the nature of data being collected, privacy risks are low:

a. Data is randomized and not directly attributable to a natural person
b. No "high-risk" personal data such as biometric, health, or criminal records are processed
c. Data is encrypted and salted.[1].

Consequently, when reviewing the data elements, security controls, and mitigations in place, there are no residual high risks for either pilot (please see details in Appendix C).

Summary & Recommendations
The risks to personal data involved in both pilots are low based on the data protection impact analysis performed. P3W does not use centralized storage, hence the majority of GDPR requirements are not applicable.

P3W Privacy Recommendations

1. An opt-in option shall be provided with details of data being collected and how it is shared between Publishers, Service Providers, and the WAYF Cloud.
2. A privacy impact assessment should be conducted to ensure that the intended purpose for the data being collected is justified against its intended purpose.

Due to the use of central storage, the WAYF Cloud needs to take measures to protect personal user data according the GDPR requirements. Such measures are included in the design of the WAYF Cloud pilot.

WAYF Cloud Privacy Recommendations

- An opt-in option shall be provided with details of data being collected and how it is shared between Publishers, Service Providers, and the WAYF Cloud.
- A data retention policy should be created to cover the entire data lifecycle, from identifying the useful life of data collected to its secure destruction requirements.
- Access to the data collected should be limited to individuals who have the “need to know” based on their job function.
- A privacy impact assessment should be conducted to ensure that the intended purpose for the data being collected is justified against its intended purpose.
- Since data transfers will be global and subject to GDPR, controls should be put in place to account for GDPR and similar regulations which will require securing, reporting, extracting, deletion, etc. of data.

Appendix

A. Approach to Threat Modelling & Analysis

Overview

This section describes the threat modelling and risk assessment performed. It provides an explanation of the terms *threat* and *threatened property* in an attempt to help answer the questions of what needs to be protected from what threat. It further presents the method used to classify threats and to weigh the risk that a threat will lead to an attack.

Definitions

Threats

In Information Security, a threat is a *potential event that, when it turns to an actual event,* an attack*, it may cause an incident that can harm an organization or system (ISO27001).*

Attacks can be classified as inside or external and may be active or passive. Active attacks result in the alteration of system resources and affect its operation. Passive attacks make use of information in the system but do not affect its operation. Inside attacks occur when legitimate users inside a security domain behave in unauthorized or unintended ways. External attacks may be carried outside of a security domain. Generally, insider attacks pose a harder set of security problems, because the attacker is trusted to some extent.

While there is no single formal list of threat definitions in Information Security, following is a list of the top common threats, in alphabetical order, collected from various sources.

Denial of Service	Denial of service covers actions and events that prevent information processing systems from providing agreed levels of service to authorized users.
Destruction	The unauthorized deletion of information.
Disclosure	To get knowledge (accidentally/intentionally) of information, which is not available by authorization.
Insertion	The unauthorized introduction of information.

Interception, Eavesdropping	The observation of user data during a communication by an unauthorized user.
Manipulation	The unauthorized modification of information.
Masquerade	The pretense by an entity to be a different entity. The unauthorized impersonation of an authorized user or common entity by discovering and using his authentication credentials.
Replay	The recording and subsequent replay of a communication at some later time.
Repudiation	The false denial that an entity sent (or created) something.

Threatened Properties

A successful attack on a target always implies a loss or reduction of a property. Following is a list of threatened properties which are considered during the threat assessment.

Authenticity	The property that the claimed identity of an entity is correct.
Availability	The property of being accessible and usable upon demand of an authorized entity [ISO 7498-2].
Confidentiality	The property that information is not made available or disclosed to unauthorized individuals, entities, or processes [ISO 7498-2].
Integrity	The property that information has not been altered or destroyed in an unauthorized manner (cf. [ISO 7498-2]).

Liability	The property that occurs upon agreed arrangements or conventions.
Privacy	The right of individuals to control or influence what information related to them may be collected and stored and by whom that information may be disclosed.
Trustworthiness	The property that the real system state coincides with the assumptions of the subjects using the system.

Threat Assessment

Threat assessment is the process of the evaluation of the system or a feature in order to identify threats and/or the threatened properties and determine the associated *risk*.

Risk is often expressed as a combination of two factors: probability and consequences. It asks two basic questions: What is the probability that a particular information security event will occur in the future? And what consequences would this event produce or what impact would it have if it actually occurred? (ISO27001)

For the purposes of this document we use the terms *Attack Effort* and *Damage Factor* as the two factors whose combination determines the risk of a threat. Attack Effort and Damage Factors are determined as follows:

Attack Effort	
Low	The target system does not provide any measures or only provides weak measures to counter the attack caused by the threat. A successful attack is easy to perform with low effort even for an unskilled attacker without any additional technical equipment.
Medium	The target system provides some measures to counter the threat. A successful attack is feasible for skilled attackers with medium effort.
High	An attacker with up-to-date knowledge needs to navigate strong technical difficulties with high effort for a successful attack. A successful attack is carried out by professionals.

Risk Factor	
Low	Negligible damage
Medium	Non-critical damage
High	Critical system functionality affected

The risk factor, which characterizes the supposed danger imposed by a threat, is then determined as follows:

Risk Factor	Low Impact	Medium Impact	High Impact
Low Effort Attack	Priority 3	Priority 2	Priority 1
Medium Effort Attack	Priority 4	Priority 3	Priority 2
High Effort Attack	Priority 4	Priority 4	Priority 3

Priority 1 and 2 represent credible dangers. It should be assumed that if no adequate security countermeasures are provided, the user/platform may be at risk. Priority 2 may still cause significant damage dependent on the concrete scenario.

Priority 3 identifies threats for which the level of effort and expense is proportional to the level of damage, e.g., to cause high damage, a high effort is necessary and with a low level of effort, only low damages are achievable. In the latter case, these threats might be still annoying if the attack is performed multiple times.

Priority 4 threats require the attacker to incur a high level of effort and expense to cause, at worst, a medium level of damage. As a result, these threats are classified as non-important risks.

The results of the threat assessment and risk effort determination are documented in tables, as shown in the example below. In particular, for each threat found during this analysis one such table is used to document and classify the risk. The table describes the threat and provides the resulting risk factor. A more detailed description regarding the reasoning of the damage and effort factor is also given. Security countermeasures are provided as recommendation.

Risk	Values: 1 to 4 based on the risk factor assessment
Threat	The name of the threat or its description
Impact	The Damage Factor (low, medium, high) along with a description, as appropriate, of the impact of a successful attack
Effort	The Effort Factor (low, medium, high) along with short explanation of how the attack can be performed
Countermeasures	Description of the actions needed to prevent the risk

B. Approach to Data Protection Impact Analysis (DPIA)

The GDPR regulation was used to baseline data protection requirements for each pilot. Below are the steps performed to determine recommendations for each pilot:

1. Using UK's Information Commissioner's Office ("ICO") screening questionnaire, assess if a DPIA is needed based on pilot's data model.
2. Perform a DPIA as required.
3. Use DPIA results for recommendations.

ICO Screening Questionnaire – DPIA Required

Based on meeting three of the criteria below, the WAYF Cloud and P3W pilots require a DPIA. Under the GDPR, a DPIA is needed if the system:

✗ uses systematic and extensive profiling with significant effects;
✗ processes special category or criminal offence data on a large scale; or
✗ systematically monitors publicly accessible places on a large scale.

The ICO also requires you to do a DPIA if you plan to:

✓ use new technologies;
✗ use profiling or special category data to decide on access to services;
✗ profile individuals on a large scale;
✗ process biometric data;
✗ process genetic data;
✓ match data or combine datasets from different sources;
✗ collect personal data from a source other than the individual without providing them with a privacy notice ("invisible processing");
✓ track individuals' location or behavior;

☒profile or target services at children; or
☒process data that might endanger the individual’s physical health or safety in the event of a security breach.

NOTE: The ICO also recommends as a good practice to perform a DPIA for any major new project involving the use of personal data even if there is “no specific indication of likely high risk.”

Data Protection Impact Analysis – Residual Risks: Low

A DPIA is a process to systematically analyze your processing and help you identify and minimize data protection risks. It must:

- describe the processing and your purposes;
- assess necessity and proportionality;
- identify and assess risks to individuals; and
- identify any measures to mitigate those risks and protect the data.

It does not have to eradicate the risk, but should help to minimize risks and consider if they are justified.

You must do a DPIA for processing that is likely to be high risk. But an effective DPIA can also bring broader compliance, financial, and reputational benefits, helping you demonstrate accountability more generally and building trust and engagement with individuals.

The DPIA resulted in a determination that residual risks are low (refer to Appendix C).

C. Data Protection Impact Analysis (DPIA)

Introduction

This document is a Data Protection Impact Assessment (DPIA) for evaluating the data privacy risks associated with the WAYF Cloud and P3W pilots. The DPIA is an analysis of expected processing activities related to assessments and covers details of the processing activity itself and an assessment of the risks associated with the processing activities including any measures that need to be taken to mitigate those risks.

This DPIA is based on a template by Questionmark (www.questionmark.com) but has been modified by our organization to reflect our use case(s).

Risks and measures

Risk methodology

This section considers the risk to natural persons—in this case, participants in the assessment process. Other risks that apply to the organization but do not impact privacy are out of scope. What is in scope are risks that could lead to physical, material, or non-material harm to the data subject, including any discrimination, damage to reputation, loss of confidentiality of data protected by professional secrecy, or any other significant economic or social disadvantage.

In our analysis, all risks are also associated with a probability:

- **Likely**. Strong (high) chance that the documented scenario could occur. High risks are going to occur from time to time; for example equipment failure in a situation where no redundancy is in place.
- **Possible**. Medium (neutral) chance that the documented scenario could occur. Between low and high.
- **Unlikely**. Scenario is unlikely—should not happen more often than once in a decade, or even more infrequently.

Risks are also associated with a severity:

- **Critical**. There is significant, real damage to a large number of data subjects; for example, a large-scale data breach.
- **Severe**. There is significant, real damage to one or a small number of data subjects, or more minor damage to a large number of data subjects.
- **Moderate**. Minor or procedural issue that does not lead to significant damage.

COMBINED PRIVACY RISK ANALYSIS

ID	Nature of Risk	Likelihood	Severity	Mitigation(s)
1	Pilot users are not aware of data being collected.	Likely	Moderate	A privacy policy will be provided.
2	JavaScript application that runs on user’s browser may access data/cookies it is not authorized for.	Unlikely	Moderate	Modern browsers have built-in security features that prevent data from being accessed by a different protocol, host, and port number.
3	Data is retained longer than necessary.	Possible	Moderate	A data retention policy will be provided.
4	Lack of access controls for data being provided to pilots.	Possible	Moderate	Access to data will be restricted based on a “need to know” basis.
5	Nature of data collected is not justified against its intended purpose.	Unlikely	Moderate	A privacy policy will be provided.
6	Data protection controls are not in place.	Possible	Moderate	Data is protected by encryption/hashing of the unique IDs.
7	Pilots do not comply with privacy regulations (GDPR).	Unlikely	Moderate	Reasonable compliance with GDPR will be assessed based on level of risk for the pilots.

PRIVACY RISK BY PILOT

Risk ID#	Issue	WAYF Cloud	P3W
1	Data Collection	Data is collected by the service providers and disclosed to the WAYF Cloud (and to other SPs via the WAYF Cloud).	Level 1–3: Data is collected by P3W. Level 4: Data is collected by the service providers and disclosed to P3W (and to other SPs via the P3W).
5	Nature of Data	Low risk to user rights and freedoms (IdP list). Includes unique IDs.	Low risk to user rights & freedoms (IdP list). Does not include unique IDs.
6	Data Protection Measures	User data is protected by the encryption/hashing of the unique IDs so that the data subject cannot be identified even when the data is combined with other IDs stored in other locations.	Not applicable, since no unique IDs are stored.
1	Minimization	Data collected is the minimum required and not used for additional purposes.	Data collected is the minimum required and not used for additional purposes.
1	User Consent	Required Recommendation: use dedicated opt-in box.	Required Recommendation: use dedicated opt-in box.
1	Consent Withdrawal	Required Can be supported by clearing browser cache or opt-out.	Required Can be supported by clearing browser cache or opt-out.
1	User Rights	As per GDPR Art. 11, user rights do not apply where the data subject cannot be identified. If the device cookie is additionally provided by the use, the user rights can be exercised by tools provided by the WAYF Cloud	Not applicable, since there is no central data storage.
7	Data Breach Notification	Not applicable as per GDPR Art. 33, as it is unlikely to result in a risk to the rights and freedoms of natural persons.	Not applicable, since there is no central data storage.
7	Data Retention (7)	Data retention policy shall be enforced by deleting data when no longer necessary.	Not applicable, since there is no central data storage.

Project Details

The table below sets out key information about the project:

Key information		
(a)	Data controller	RA21.ORG
(b)	Purpose of project	Resource Access for the 21st Century (RA21) is a joint STM and NISO initiative aimed at optimizing protocols across key stakeholder groups, with a goal of facilitating a seamless user experience for consumers of scientific communication. In addition, this comprehensive initiative is working to solve long standing, complex, and broadly distributed challenges in the areas of network security and user privacy. Community conversations and consensus building to engage all stakeholders is currently underway in order to explore potential alternatives to IP-authentication, and to build momentum toward testing alternatives among researcher, customer, vendor, and publisher partners.
(c)	Context, scope and background	A simple and secure access infrastructure requires involvement of a diversity of players and therefore a diverse approach. Key stakeholders will explore pathways to move beyond IP-recognition as the primary authentication system. The RA21 Taskforce will not build a specific technical solution or an industry-wide authentication platform; rather its objectives are to: • Recommend new solutions for access strategies beyond IP recognition practices. • Explain the standard measures that publishers, libraries and end-users should undertake for better protocols and security. • Test and improve solutions by organizing pilots in a variety of environments for the creation of best practice recommendations.
(d)	Data subjects (called “Participants” within this document)	Assessments are delivered to the following types of data subjects: • *Students* • *Publishers* • *Researchers* • *Customers* • *Other Partners and Institutions*

(e)	Types of personal data	The main personal data captured is identification data for the data subject including: • email address (P3W) • local-ID • group-ID
(f)	Special categories of data	Special/sensitive categories of data are not captured per the GDPR.
(g)	Who will be able to see and have access to data collected	The following roles will have access to assessment results: • WAYF Cloud • P3W developers and administrators

Assets, Including Processors and Sub-processors

This section of the DPIA contains a list of the assets through which personal data processing takes place, both internal and external to the organization.

1. **Internal IT**
   Personal data necessarily flows through the internal computer systems of the RA21 members.

2. **Server Environment**
   Personal data also flows through servers used for the WAYF Cloud and P3W pilots.

3. **User's Personal Computer**
   Data is primarily entered using users'/participants' personal computers. Some cookies may be stored locally on these computers.

Personal Data Captured

This section describes at a high level the types of personal information captured on participants. The following general information may be captured on participants:

- Email address (P3W)
- Local-ID
- Group-ID

Necessity and Proportionality – GDPR Article 6(1)

General Reasons Why the RA21 Pilots are Beneficial

RA21's mission is to align and simplify pathways to subscribed content across participating scientific platforms. RA21 will address the common problems users face when interacting with multiple and varied information protocols.

Legitimate Interests for RA21

Legitimate Interests for the pilots include Researcher requirements for:

- Seamless access to subscribed resources, from any device, from any location, from any starting point
- A consistent, intuitive user experience across resources
- Increased security for management of personal data
- Streamlined text and data mining

Resource Provider objectives:

- Provide individualized and differentiated access for better reporting to governing bodies and clients
- Offer personalized services to accelerate insight and discovery
- Ensure the integrity of content on both institutional and commercial platforms

Customer responsibilities:

- Minimize the administrative burden of providing access to authorized user communities
- Maximize the use of the resources purchased
- Protect the privacy of user communities and advocate for their security

(www.ra21.org)

3. Conclusion

Summary

In previous parts of this DPIA, we have:

a) Described the project and given a functional overview
b) Described the personal data that has been captured
c) Identified the purpose of the processing and the legitimate interests in conducting such processing
d) Identified the risks to privacy of data subjects and the mitigations in place for them

We now need to consider whether residual high risk remains and whether there is a need to consult the supervisory authority.

Residual Risk - LOW

Under the GDPR, it is necessary to consult the supervisory authority prior to processing where a DPIA indicates that the processing would result in a high risk in the absence of measures taken by the controller to mitigate the risk. Essentially, if the residual risk after mitigations taken remains high, then the supervisory authority needs to be consulted.

The Article 29 Working Party guidance on DPIAs gives examples of unacceptable high residual risk:

a) Instances where the data subjects may encounter significant, or even irreversible, consequences, which they may not overcome (e.g.: an illegitimate access to data leading to a threat on the life of the data subjects, a layoff, a financial jeopardy)
b) When it seems obvious that the risk will occur (e.g.: by not being able to reduce the number of people accessing the data because of its sharing, use or distribution modes, or when a well-known vulnerability is not patched)

For the Cloud WAYF and P3W pilots, the nature data being collected is low risk.

- Data is randomized and not directly attributable to a natural person.
- No ”high risk” personal data such as biometric, health, or criminal records is processed.
- Data is encrypted and salted.[2]

[2] In cryptography, a salt is randomized data added to a hash function to make the password harder to crack or mathematically guess.