Exploring Privacy Implications
in OAuth Deployments

Srivathsan G. Morkonda, Paul C. van Oorschot and Sonia Chiasson

School of Computer Science, Carleton University

Ottawa, ON, Canada

Email: srivathsan.morkonda@carleton.ca, paulv@scs.carleton.ca, chiasson@scs.carleton.ca

Abstract—**Single sign-on authentication systems such as OAuth 2.0 are widely used in web services. They allow users to use accounts registered with major identity providers such as Google and Facebook to login on multiple services (relying parties). These services can both identify users and access a subset of the user's data stored with the provider. We empirically investigate the end-user privacy implications of OAuth 2.0 implementations in relying parties most visited around the world. We collect data on the use of OAuth-based logins in the Alexa Top 500 sites per country for five countries. We categorize user data made available by four identity providers (Google, Facebook, Apple and LinkedIn) and evaluate popular services accessing user data from the SSO platforms of these providers. Many services allow users to choose from multiple login options (with different identity providers). Our results reveal that services request different categories and amounts of personal data from different providers, with at least one choice undeniably more privacy-intrusive. These privacy choices (and their privacy implications) are highly invisible to users. Based on our analysis, we also identify areas which could improve user privacy and help users make informed decisions.**

I. INTRODUCTION

An increasing number of web applications encourage users to log in to their services in exchange for a personalised experience. However, this comes with a usability problem for users having to administer a growing list of account credentials, e.g., to choose unique and strong passwords for accounts on each service. Managing large sets of credentials is a difficult task for users and can result in insecure practises such as reusing passwords and choosing weak passwords [34]. Many web authentication schemes have been proposed to improve usability and convenience. Federated single sign-on (SSO) schemes involve a trust relationship between an *identity provider* (IdP) and one or more other-party services (*relying parties* or RPs) that allow users to identify themselves on the service using login credentials registered with the IdP. The OAuth 2.0 protocol is an example of a federated SSO scheme used to establish trust in identity-related interactions between an IdP and an RP.

Instead of requiring users to create new credentials, RPs can use the OAuth 2.0 framework [20] to outsource user identification to an IdP with whom users are likely to have an existing account. As part of this transaction, the RP can also request access to additional user personal data stored with the IdP. Major identity providers (e.g., Facebook, Google, Microsoft) expose web APIs to grant RPs controllable access to protected user data stored on their platforms. With the user's

permission, an IdP allows the RP access to one or more user-data attributes, which in some cases includes sensitive user data such as emails, contact information and documents in personal cloud storage—raising privacy concerns. Such user data allows the RP to extend functionality and personalise web content to the user. It also reduces implementation costs for the RP since they are outsourcing login-related tasks, including key management and credential verification, to the identity provider.

Prior work (e.g., [10] [16]) involving OAuth-based SSO gives focus to security issues including leaks of user data from relying parties to other third-party actors due to implementation flaws. In contrast, we focus on privacy consequences for users of OAuth-based SSO that arise when RPs gain access to user data stored on IdP sites. In this study, we evaluate privacy practises in popular services using OAuth for SSO. Our contributions include:

- *OAuthScope*, a tool to extract OAuth protocol request parameters from sites supporting major SSO providers (Google, Facebook, Apple and LinkedIn).
- An empirical study of OAuth-based logins in the Alexa Top 500 sites for five countries. The study reveals considerable variation in how relying parties implement the different SSO options, in the data made available by identity providers, and in apparent trends across countries.
- An explication of how user choices, typically made without full information, can result in release of considerably different amounts of user data and have privacy implications.

Based on our analysis, we further discuss possible modifications to OAuth which could lead to improved end-user privacy.

Note: while some enterprise SSO providers (e.g., Microsoft) are popular among users, these are primarily used in closed systems using enterprise accounts. Since we target websites involving personal user accounts, we do not include these providers in our study.

The next section provides background on OAuth 2.0 framework. Section III introduces the OAuthScope tool. Section IV presents the empirical study. Section V provides our classification of user data available in four identity providers. Results of the empirical study are reported in Section VI. Section VII presents our analysis of privacy implications to SSO users. In Section VIII, we propose stakeholder-based changes to
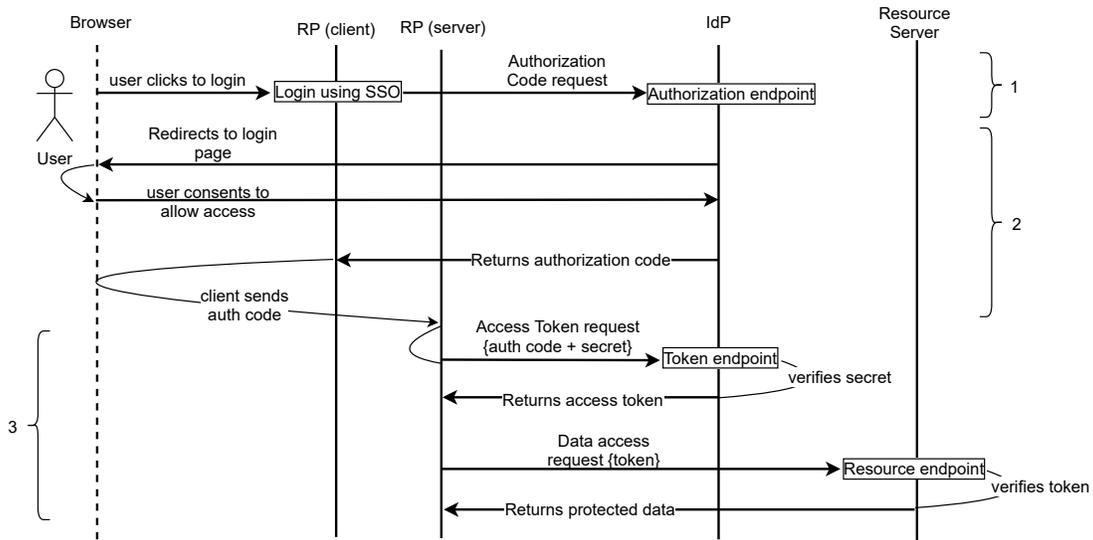
Fig. 1. Procedure for OAuth 2.0 Authorization Code flow (based on [20]). Diagram for the Implicit flow is in Appendix A.

improve user privacy. Related work is described in Section IX and concluding remarks in Section X.

II. OAuth 2.0 Framework (Background)

OAuth 2.0 [20] is a web resource authorization protocol popular in client-server deployments worldwide for granting applications access to protected resources without sharing the user's credentials. For example, a website prompts a user to log in and optionally allow access to specified user data by relying on their Google account rather than creating new credentials on that website. The website does not gain access to the user's Google credentials, but instead gains access to a subset of the user's data and trusts Google's verification of the user's credentials. The user is able to use the same Google account to log in on other websites that support Google as a SSO option. User data is referred as *protected resources* and the user is called as the *resource owner* (RO) in OAuth 2.0 specification.

The OAuth 2.0 protocol is composed of several "grant types" that define how credentials are granted to RPs. The procedure for obtaining access to protected resources is defined by *OAuth flows* and each flow is designed to serve different use cases (and provide different levels of security). Since our study involves OAuth 2.0 use in web applications, we describe the steps involved in two flows commonly used in these applications.

A. Authorization Code Flow (Server-side flow)

In order for an RP to use OAuth 2.0, it must first register itself with an IdP to obtain a `client_id` (used for identification of the RP in requests). Additionally, in the case of confidential clients (RPs with the ability to securely store secrets), an associated `client_secret` is issued by the IdP. This allows an IdP to authenticate requests from an RP. We provide an overview of the three primary steps (labelled in Fig. 1) involved in the *authorization code flow* [20].

1) **RP request to IdP**: The resource owner triggers the flow by clicking a login element, sending a request constructed by the RP to the IdP's *authorization endpoint*. The RP programs the request to add several parameters (as query components in the request URI) specifying the access request. To indicate this flow type, the RP must set the value of `response_type` to `code`. The request also includes the parameters `scope` and `redirect_uri`. The `scope` parameter lists one or more protected resources (e.g., name, email) the RP is requesting to access. The allowed values for this parameter are specified by individual IdPs based on the resources made available by the IdP. The `redirect_uri` is the endpoint the RP is requesting the IdP to redirect the RO, at the end of the flow. In order to redirect to the intended party, the OAuth 2.0 specification [20] requires the IdP to check if the URI specified in the request matches with the value registered by the RP.

2) **IdP request to RO**: After receiving the RP's request, the IdP redirects the RO to a login page on its domain to login with the IdP account. If the RO is already logged in, the IdP displays a prompt to confirm resource access to the RP. In this process, the IdP provides the RO with a list of user-data attributes the RP is requesting to access and optionally provides the ability to deny RP the access to one or all of the requested attributes. If the RO approves the request, the IdP issues an authorization `code` and redirects the RO back to the RP at the verified redirect endpoint. The `code` is included as a query component to the redirection URI. The standard does not specify what should happen if the RO denies access, leaving it to the IdP's discretion.

3) **Token exchange**: From its server-side application (which is considered more secure than a channel from

the RO's browser), the RP uses the `client_secret` and the authorization `code` to exchange for an *access token*. After verifying the `code` and RP's `secret`, the IdP issues an access token allowing the RP to access RO's resources within the IdP. Access tokens can vary in format (chosen by each IdP), but the purpose is to represent authorization information (e.g., allowed scope values, token expiry date). A commonly used self-contained format is the JSON Web Token (JWT) [21] that protects the integrity of information within the token.

B. Implicit Flow (Client-side flow)

The *implicit* grant type is simpler than the authorization code flow in that it allows the RP to directly obtain the access token without exchanging an authorization code. It is useful for browser-based (JavaScript) apps that lack the ability to securely store secrets. The RP initiates the implicit flow by sending an access request to the IdP with the parameter `response_type=token`. The IdP prompts the user with a login screen and, if access is granted, an access token is issued and included in the URI fragment of the redirect endpoint. The redirect URI is verified by the IdP to ensure it matches the value registered by the RP. Since the access token is returned in the redirection URI, it is included in the user's browsing history and, therefore, the security of the access token relies on the security of the user's system. A malicious application with access to the user's browser could misuse the access token. Additionally, third-party scripts running within the RP site will be able to access the token.

The OAuth 2.0 implicit flow was originally designed when browsers restricted apps to make requests only to its own domain [28]. This browser restriction prevented browser-based apps from using the authorization code grant since it requires sending a HTTP POST request to the IdP's authorization endpoint, which in many apps (not owned by the IdP) is different from the RP. The implicit flow in OAuth 2.0 offered a workaround that avoids using the POST request and includes the access token directly in the redirection URI. Though implicit flow provides the needed functionality, this is not recommended from a security perspective due to the risks associated with storing credentials in URIs. Modern browsers now support Cross-Origin Resource Sharing (CORS) that enables a website to request resources from other permitted origins, removing the need for this workaround. The challenge still remains for browser-based applications to securely store secrets, and the recommended flow for such situations is the authorization code flow with Proof Key for Code Exchange (PCKE) [30].

C. Authentication (OpenID Connect)

OAuth 2.0 can be adapted to allow an IdP to authenticate users to the RP. The OpenID Connect 1.0 [31] (OIDC) specification is designed for authentication and is built upon the OAuth 2.0 protocol. OIDC introduces a special value (`openid`) to the `scope` parameter for specifying intent to authenticate. The OIDC specification also extends OAuth 2.0's `response_type` parameter to define additional flows. If this parameter includes the value `id_token`, the OpenID Provider (OP) will issue an *ID token* to the RP after a successful authentication, encoded as a JWT [21] and containing key-value pairs (*Claims*) about the user's identity. *Claims* are digitally signed using JSON Web Signature (JWS). Since many RPs use both OIDC and OAuth 2.0, we refer to *OpenID Providers* as IdPs in our study. OIDC allows any combination of `id_token`, OAuth 2.0 values `code` and `token` for the `response_type` parameter [31]. Each combination refers to a *hybrid flow* that defines the values (access token, authorization code and ID token) included in IdP response and the endpoint (authorization and token) issuing the values.

Although the OIDC standard is built on top of OAuth 2.0 specification, some identity providers instead use custom modifications of OAuth 2.0 to provide authentication capabilities. In Section V, we briefly describe these modifications as part of our analysis of the four identity providers.

D. Refresh Tokens

An OAuth 2.0 access token is issued for a limited lifetime and once it expires, the user is required to approve access again in order for a new access token to be issued to the RP. A common access token type is "bearer" token that allows use by any party in possession of the token. An unauthorized party in possession of such tokens can use it to access protected resources. While long-lived access tokens improve user experience by reducing the frequency of required logins, the longer life increases the risk of token leaks. The recommended approach is to combine the use of a access token with a *refresh token*, a string issued by the IdP that allows the RP to extend existing access without involving the user to approve access again. The RP can extend its access by exchanging the refresh token for a new access token with a scope (identical or lesser than the previously issued access token) determined by the IdP. During this exchange, the IdP validates the refresh token before issuing a renewed access token and optionally, a new refresh token. Due to the sensitivity of refresh tokens, the specification restricts its use to only server-side flows such as the authorization code flow [20].

III. THE OAUTHSCOPE TOOL

We designed and built OAuthScope, a Java web tool that scans and extracts OAuth 2.0 protocol-related parameters from authorization requests made by relying parties to identity providers. We provide a list of URLs as input, and OAuth-Scope visits each URL in a headless browser setup from a local server. It scans each site to locate OAuth-based SSO requests to any of the four providers in our study (i.e., Google, Facebook, Apple, and LinkedIn). For each SSO option available on the RP, OAuthScope simulates a mouse-click to trigger an OAuth 2.0 flow, which initiates an authorization request by the RP to the IdP's authorization endpoint. OAuthScope captures the parameters included in the request, including the flow type and `scope` parameter that specifies the protected

resources for which the application intends to obtain access. OAuthScope uses *Selenium WebDriver* [32], an open-source browser automation framework built primarily for automated in-browser testing of web applications. OAuthScope uses the *WebDriver* framework for identifying HTML elements on an RP site and simulating web page navigation to extract OAuth 2.0 protocol data from supported IdP sites. We summarize the primary steps performed by OAuthScope as follows:

1) **Identify login element**: After loading a website's landing page, OAuthScope searches for potential login elements based on a set of predefined match criteria and simulates a user click if an element is found. Most RPs display login forms in either a new page or a new iframe within the landing page. We consider both cases and switch contexts as necessary. Our tool captures a screenshot of the login page, for use in manual verification of correctness.

2) **Identify SSO elements**: On the new page (or iframe), it performs a search for HTML elements that potentially lead to a login page of one of the four IdPs in our study. OAuthScope identifies SSO elements based on element texts and HTML tag values commonly found in RPs pointing to IdPs. After obtaining potential SSO login elements, our tool simulates clicks on each element and checks if the resulting page's URL matches with a list of predefined endpoints for each IdP in our study. We built this list to contain OAuth 2.0 endpoints published on each IdP's developer documentation pages.

3) **Extract OAuth 2.0 parameters**: The OAuth 2.0 framework specifies that protocol parameters should be added as query components of the request URI (an example is listed in Appendix A) initiated by an RP. OAuthScope extracts these parameters encoded in the link as key-value pairs. Finally, the parameters and login screenshots are persisted to a database for further analysis.

We note that the current version of OAuthScope does not automatically ensure identification of all SSO elements in all websites. A manual inspection of the web pages is still necessary to ensure completeness of the collected data. To facilitate this process, our tool captures screenshots of the login pages and of the landing page for any site where our tool did not find a login element, and stores the screenshots to a database. We use these images to identify any SSO options missed by the automated scan and manually feed the SSO login links to OAuthScope for extraction of the protocol parameters. Since our goal in this study is to evaluate privacy implications of top sites using OAuth-based SSO, we did not focus our efforts on building a tool that can fully automate the data extraction process. However, we did iteratively improve our match criteria based on strings and tags found in RP sites during our data collection process. This gradually reduced the need for manual involvement in data collection. OAuthScope includes a web front-end (included in Appendix A) for displaying and analysing the collected dataset.

IV. EMPIRICAL STUDY: OVERVIEW

We conducted a study of SSO systems in popular websites that implement the OAuth 2.0 framework for identifying users and accessing user data stored with different identity providers. In an initial exploration, we manually scanned the Alexa Top 500 sites [3] in a first target country and found that three major identity providers (Google, Facebook and Apple) are predominantly supported as SSO options. In addition to these providers which primarily contain personal data of users, we also included LinkedIn as an IdP given its popularity as a platform for sharing professional data.

For each of these four IdPs, we collected OAuth 2.0 protocol-related data from the Alexa Top 500 sites in five countries: Australia, Canada, Germany, India and the United States. We describe our data collection procedure below.

A. Research Questions

The goal of this study is to understand the privacy implications for users opting to use OAuth-based SSO to log in to the top websites. To achieve this, we pursue the following research questions:

RQ1: What categories of user data do relying parties request from SSO providers? (Sec. VI-B, VI-C)

RQ2: How prevalent is each SSO scheme in popular relying parties across the five countries? (Sec. VI-D)

RQ3: If a relying party supports multiple SSO options, how do they differ in terms of requested user-data attributes? (Sec. VI-C)

B. Data Collection

To collect accurate representation of websites as served to local users in each country, we use a VPN service to connect to a server in the country when scanning and collecting data from sites. For each country included in our study, we first manually visit each of the Alexa Top 500 sites in the country and identify websites that support at least one of our four chosen IdPs. We then run each filtered site through OAuthScope to extract the OAuth 2.0 parameters for each of these IdPs supported by the website. As a cross-check, we noted the IdPs supported by each website during our initial manual scan and verified that OAuthScope collected all available data from each website. For any omissions, we manually obtained the IdP links and passed it to OAuthScope for extraction of the protocol parameters. In total, this provided us with a dataset consisting of details from 815 RPs (Australia: 174; Canada: 159; Germany: 126; India: 172; US: 184). Our dataset consists of all RPs from each country that use at least one of the four IdPs; if a site appeared in the top 500 list of more than one country, it is included each time so that we could make direct comparisons across countries.

V. API ANALYSIS OF IDENTITY PROVIDERS

User data available to third-party services through OAuth-backed APIs vary with each identity provider. We first compare IdPs before evaluating the RPs that request permissions via these APIs. The diverse native services provided by IdPs lead

Data category	Google	Facebook	Apple	LinkedIn
Basic	email *profile* openid	email *public_profile*	email name	r_emailaddress name profilePicture headline
Identity	user.birthday.read user.addresses.read* user.gender.read* user.phonenumbers.read*	user_birthday user_hometown user_gender user_age_range instagram_graph_user_profile*		address birthDate phoneNumbers backgroundPicture
Personal	userinfo.profile photoslibrary* fitness* tasks*	user_location user_photos user_videos instagram_graph_user_media*		geoLocation
Behavioural	games* user.organization.read*	user_likes user_posts user_link		organizations positions educations projects certifications skills volunteeringInterests volunteeringExperiences
Sensitive	contacts drive gmail documents* spreadsheets* youtube*	user_friends		websites industryName courses testScores summary

to differences in the types of user data available within each IdP's SSO platforms. These differences make it difficult to objectively compare permissions across multiple IdPs at different granularity. To assist with this comparison, we scanned OAuth-backed data attributes relating to personal user data for each IdP and categorised the attributes based on how the information could be used (or misused). We grouped items into five data categories: *basic*, *identity*, *personal*, *behavioural*, and *sensitive*, as shown in Table I.

Categories: Data in the *basic* category includes the attributes we consider least privacy-invasive (e.g., name, email, profile image) shared with RPs through OAuth APIs. Here, name refers to the user's full name or profile username, and in the case of Apple SSO, the user is allowed to choose a custom name to be shared with the specific RP being accessed. Data attributes that facilitate identification of specific users such as user's birthday, gender, mobile number and street address are grouped into the *identity* category. These are security-critical user data that, in the hands of adversaries can be misused to impersonate users (e.g., obtain a new mobile SIM). The *personal* category includes data that enable RPs to access the user's personal data including photos, videos and current location. While such data pertains to users, it could also involve other secondary individuals who are part of the data being shared (e.g., a friend in the user's videos). We include data attributes (e.g., social content liked by users) that could exhibit a user's personality/habits as *behavioural* data. User's *behavioural* data can possibly be used by online trackers and other scripts to study users. Finally, we include user data

such as emails, contact lists and documents that could contain personal information under the *sensitive* category.

We note that some user data can be associated with multiple categories. For example, a passport copy uploaded to cloud storage can be both *sensitive-* and *identity*-related. In such cases, we include the attribute in the category we deemed more relevant. Our goal is to provide an overview of the types of user data accessed by RPs, and not a mutually-exclusive categorisation of OAuth-backed APIs. Table I groups relevant attributes from each provider into the five categories. We selected an initial list of attributes relevant to user data by reviewing OAuth documentation pages available on each IdP platform. Then, we refined this list to include all attributes requested by the top 500 sites in each of the five countries. In Fig. 6, we include screenshots of UI shown to SSO users of the four IdP. We now discuss the four IdP in the context of data categories made available through their OAuth-backed APIs.

A. Google OAuth API

Google's *OAuth 2.0 Scopes* document [18] categorizes data attributes into APIs based on the service where each is normally used within Google's ecosystem. For example, the Gmail API groups attributes relevant to sending and receiving emails in a Gmail inbox. This is useful for third-party email clients (e.g., Mozilla Thunderbird) that externally manage users' emails. Personal user information such as birthday, gender, and contact information are exposed through the People API. Other Google APIs that we found to be commonly used by RPs include the Calendar API (used for managing
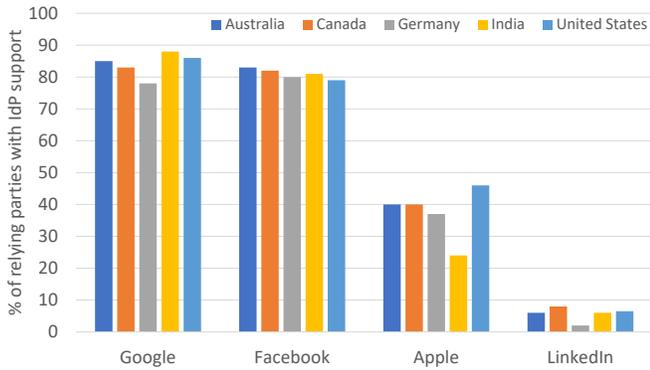
Fig. 2. Percentage of RPs per country in our dataset that support each IdP.
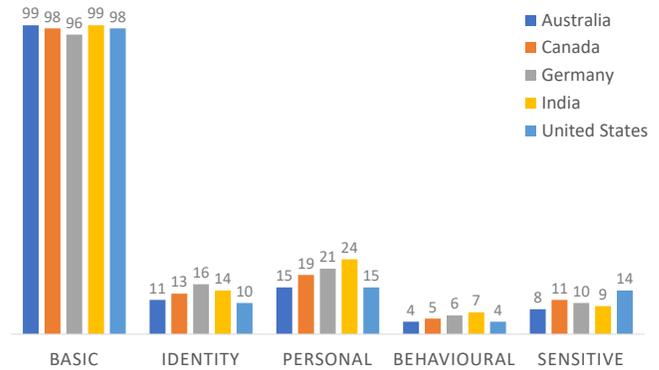

Fig. 3. Percentage of RPs per country in our dataset requesting permissions per category (as defined in Table I). Some *basic* attributes are required by default by IdPs, partially explaining high number of requests for this category.

Google calendars) and the Drive API (used for accessing the user's Google Drive storage). Although not relevant to our study, Google makes available several other attributes to RPs through OAuth.

In addition to attributes found in our dataset, we included attributes related to personal user data from all Google APIs listed in the Scope document [18]. Google's OAuth 2.0 APIs classify a subset of their attributes as belonging to *sensitive* (different from *sensitive* category defined in our study) and *restricted* scopes. Most RP applications requesting access to any of these attributes must go through a verification process reviewed by Google prior to use [19].

Google's OAuth 2.0 platform includes two types of profile-related attributes. The `profile` [17] attribute under *basic* category shown in Table I includes only the user's name, email and public profile image (and is included by default in response to requests), whereas `userinfo.profile` includes all publicly available information from the user's Google profile and must be explicitly requested. Outside of the SSO platform, profile information can be accessed through Google services such as Hangouts, Maps, YouTube and Play[1].

Authentication: Google's OAuth 2.0 platform supports the OpenID Connect attribute, `openid`, which allows RPs to identify users as specified by the OpenID Connect 1.0 specification [31]. Including the `openid` attribute in the authentication request allows the RP to obtain from Google an ID token containing fields that assert the user's identity (e.g., username, token's issuer, token expiry timestamp). The fields are encoded as key-value pairs in the data portion of a JWT (JSON Web Token) [21] object with a digital signature signed by Google.

Supported flows: Google's SSO platform uses both OAuth 2.0 and OIDC specifications to support standard *flows*, including the implicit and authorization code flows. RPs specify the flow using the `response_type` parameter, as discussed in Section II. Additionally, the value `permission` can be used for the parameter to specify use of implicit flow[2].

[1]https://support.google.com/accounts/answer/6304920

[2]https://developers.google.com/identity/sign-in/web/reference

B. *Facebook OAuth API*

Facebook's Graph API [12] is the platform for third-party applications to interact with a user's Facebook data. Applications implementing SSO with Facebook use the *Facebook Login* interface to identify users. RPs can gain read-access to a wide range of data attributes from a user's Facebook profile. *Facebook Login* provides social media-specific *behavioural* data (e.g., Facebook pages likes and social posts created by the user) not available through other IdPs. We find that almost all of these attributes (as marked in Table I) are requested by a site in our empirical study.

Other permissions available in *Facebook Login* allow applications to view, create, edit and delete content on user-administered Facebook Pages. This is also useful for business management applications to integrate business-owned Facebook Pages in promoting products and services. Facebook also provides Instagram Permissions for RPs (e.g., a service managing user's digital photo library) to access a user's Instagram content [13]. Although these permissions are available for access, we focus our analysis on user-data attributes (listed on Table I) accessed by RPs in the top 500 sites.

All SSO requests (including those that do not explicitly request the `public_profile` attribute) to *Facebook Login* require the user to allow the RP access to default fields (Facebook name and profile picture) [13]. RPs requesting permissions other than `public_profile` (user's name and profile picture), `email` and `pages_show_list` (list of Facebook Pages managed by the user) are required to undergo a Facebook App Review [11].

Authentication: Although *Facebook Login* does not support the OIDC specification, it allows RPs to authenticate users without requesting access to user data. Authentication can be done using *Facebook Login*'s *signed request* which is a signed base64url encoded JSON object containing a user ID issued by Facebook.

Supported flows: *Facebook Login* supports the standard OAuth 2.0 authorization code flow (`code`) and the implicit flow (`token`). RPs can additionally include the `signed_request` parameter to obtain the user's ID and the

granted_scopes parameter to obtain a list of permissions approved by the user.

C. Apple OAuth API

Apple introduced the *Sign in with Apple* framework in 2019 as a privacy-friendly alternative for SSO users wanting to use third-party web and mobile applications without disclosing their data. In addition to libraries for Apple devices, it includes a JavaScript library compatible with web applications in Windows or Android operating systems. *Sign in with Apple* allows RPs to authenticate SSO users and to optionally request access to the user's name and email through the scope parameter. When requested, *Sign in with Apple* allows users to either share their original email address or create an anonymous email address enabled by Apple's Private Email Relay Service. This service generates an anonymous email address unique to the user-RP pair and routes all email correspondence between RP and user through this email, hiding the user's real email from RP. *Sign in with Apple* also helps RPs distinguish real users from bots through a boolean-value real user indicator [5]. Applications on Apple's App Store using a third-party SSO service (e.g., Facebook, Google) are now required to also offer *Sign in with Apple* as an SSO option[3].

Authentication: Unlike other IdPs, *Sign in with Apple* is primarily used for authentication with limited user data (real or anonymous) available to RPs. Although the documentation [5] lacks explicit mention of the OpenID standard, it closely follows OpenID conventions discussed in Section II-C. A successful authentication returns an ID token in a JWT object, allowing the RP to identify the user.

Supported flows: The accepted values for the response_type parameter in *Sign in with Apple* include either code (for authorization code flow), id_token (for authentication) or both. It does not support the implicit flow.

D. LinkedIn OAuth API

LinkedIn is a widely used platform for sharing professional and personal user data. The *Sign In with LinkedIn* platform allows users to authenticate and authorize profile access to third-party applications. In addition to r_emailaddress (user's email address), *Sign In with LinkedIn* provides three scope parameters related to a user's LinkedIn profile. Each parameter groups several attributes, providing less granular access to user data compared to other IdPs. Applications can specify the r_liteprofile (or r_basicprofile used in older version of the API) scope to request access to the user's full name, profile picture (including image meta data) and profile headline. The r_fullprofile scope additionally includes all the other fields (listed in Table I) entered to the user's LinkedIn profile [23]. For each SSO request, *Sign In with LinkedIn* provides a unique user ID that allows RPs to authenticate users. If a request includes the w_member_social scope, the user can authorize the RP to create new LinkedIn posts on their behalf.

[3]https://developer.apple.com/news/?id=09122019b

IdP	Scope	Australia	Canada	Germany	India	US
	email	99	97	98	99	97
G	profile*	79	78	69	78	81
O	openid	55	54	53	64	62
O	userinfo.profile	11	11	17	17	10
G	user.birthday.read	2	1	2	1	1
L	contacts	1	1	1	2	1
E	gmail	0	0	1	0	1
	drive	0	0	0	1	0
	email	91	87	90	89	89
	public_profile*	29	36	36	45	38
F	user_birthday	9	12	14	12	8
A	user_friends	7	10	9	6	13
C	user_location	3	8	6	6	3
E	user_hometown	3	6	4	3	2
B	user_likes	2	4	3	3	3
O	user_gender	2	6	6	0	2
O	user_photos	1	2	2	2	3
K	user_link	1	2	2	2	0
	user_posts	1	1	1	1	1
	user_age_range	1	3	0	1	0
	user_videos	1	0	1	1	1
APPLE	email	100	100	94	100	99
	name	94	95	87	95	92
LINKEDIN	r_emailaddress	90	85	100	100	83
	r_liteprofile	90	85	100	80	75
	r_basicprofile	10	0	0	20	8
	r_fullprofile	10	0	0	10	8

Fig. 4. The percentage of RPs per IdP in our dataset that requests a particular scope attribute. Blue cells represent attributes from the *basic* category. Darker cells indicate a higher percentage of RPs making a given request. * identifies default scope attributes that the IdP requires users to share with RP.

Authentication: In LinkedIn's documentation for *Sign In with LinkedIn*, we did not find explicit mention of the OpenID standard or information on whether relying parties could request only the member ID, a unique identifier specific to the RP-user pair. However, RPs can authenticate users using the OAuth flows supported by the platform. After a successful login, the RP obtains an OAuth 2.0 access token which includes the user's LinkedIn member ID.

Supported flows: *Sign In with LinkedIn* defines two types of *consents* (analogous to OAuth flows): *member authorization* and *application authorization*. We concern ourselves only with *member authorization*, where *Sign In with LinkedIn* requires the use of the authorization code flow from OAuth 2.0. Conversely, *application authorization* uses OAuth 2.0's *client credentials flow* for systems requiring machine-to-machine authorization, without user involvement [20], and is outside the scope of this study.

TABLE II
OAuth 2.0 and related OpenID Connect (OIDC) flows used by RPs per country. N = total number of sites (among top 500) offering the IdP per country (number of RPs); n = number of RPs using the given flow; % = percentage of RPs using a given flow (i.e., $\% = n/N * 100$). Hybrid flows represent multi-valued response types described in Section II-C.

IdP	OAuth 2.0/OIDC Flow	Response Type	Australia			Canada			Germany			India			US		
			N	n	%	N	n	%	N	n	%	N	n	%	N	n	%
Google	Authorization code	code	148	95	64	132	85	64	98	74	76	151	83	55	159	95	60
	Implicit	token	↑	4	3	↑	4	3	↑	0	-	↑	3	2	↑	2	1
		id_token		2	1		1	1		0	-		0	-		2	1
		id_token token		34	23		35	27		17	17		55	36		50	31
	Hybrid	code id_token		0	-		0	-		0	-		0	-		0	-
		code token		0	-		0	-		0	-		0	-		1	1
		code id_token token		13	9		7	5		7	7		10	7		9	6
Facebook	Authorization code	code	147	102	69	134	90	67	103	73	71	144	82	57	148	84	57
	Implicit	token	↑	4	3	↑	5	4	↑	0	-	↑	2	1	↑	3	2
		token signed_request		41	28		39	29		30	29		60	42		61	41
Apple	Authorization code	code	70	34	49	64	30	47	47	30	64	42	16	38	85	35	41
	Hybrid	code id_token	↑	36	51	↑	34	53	↑	17	36	↑	26	62	↑	50	59
LinkedIn	Authorization code	code	9	9	100	11	11	100	3	3	100	10	10	100	11	11	100

VI. Empirical Results

In this section, we report the findings of our empirical study on the use of OAuth 2.0 to access user data in popular SSO services.

A. Distribution of Providers

Our results show that Google and Facebook are consistently the most popular SSO options in top websites across all five countries. Apple is currently the third most popular option, possibly due to its relatively recent introduction of *Sign in with Apple* in 2019. As shown in Fig. 2, *Sign in with Apple* is less popular in India, which is consistent with Apple's lack of popularity in India [33]. However, recent requirements (discussed in V-C) for apps on Apple's App Store could lead to an increase in its use with RPs.

B. Comparing requested data across countries

Fig. 4 provides a sorted list of the most requested attributes for each IdP. The majority of the RPs request one or more *basic* attributes that help identify users (e.g., to display the user's name and profile picture on the RP site). Although relatively similar patterns emerge across countries, we do note some variation on particular attributes. For example, Google's `userinfo.profile` is requested more frequently in Germany and India, while LinkedIn's `r_basicprofile` and `r_fullprofile` are requested most frequently in India, followed by Australia, and not requested at all in Canada and Germany. These variations could be a result of an IdP's popularity with the users in a given country. Perceived sensitivity of specific user data could also vary across countries, and consequently lead to different privacy implications for users.

During our analysis, we also observed that RPs use different web designs and make available different SSO options in different countries. For example, Rakuten.com (a popular e-commerce marketplace) supports three SSO options (Google, Facebook and Apple) for US users and requests read access to the user's emails when signed on with Google SSO. However, Rakuten.ca (for Canadian users) only supports Facebook and Apple SSO options. Interestingly, Rakuten.de (for German users) does not support any SSO options. We did not have any Rakuten sites specific to users in Australia or India in our dataset.

C. Comparing requested data across providers

Although each IdP exposes different user data, we notice differences in which attributes RPs requested for similar data categories across the providers. For example, 13% of RPs in the US request access to users' friends list with the Facebook SSO, but only 1% request the contacts list with the Google SSO. Fig. 5 links RPs in the US sites with the user data they request from each IdP. For readability, this chart does not include the *basic* attributes, and instead focuses on RPs that request more sensitive non-*basic* attributes. From Fig. 5, we see that the US airbnb.com and tripadvisor.com request more non-*basic* attributes than other RPs (towards the center of the graph, each with 5 attributes from Facebook).

Both Facebook and Google SSOs are supported by tripadvisor.com. When a user logs in with Facebook, tripadvisor.com requests access to user's hometown, location (as listed in profile), list of likes (of Facebook pages), friends list and photos. However, the request only includes the *basic* attributes
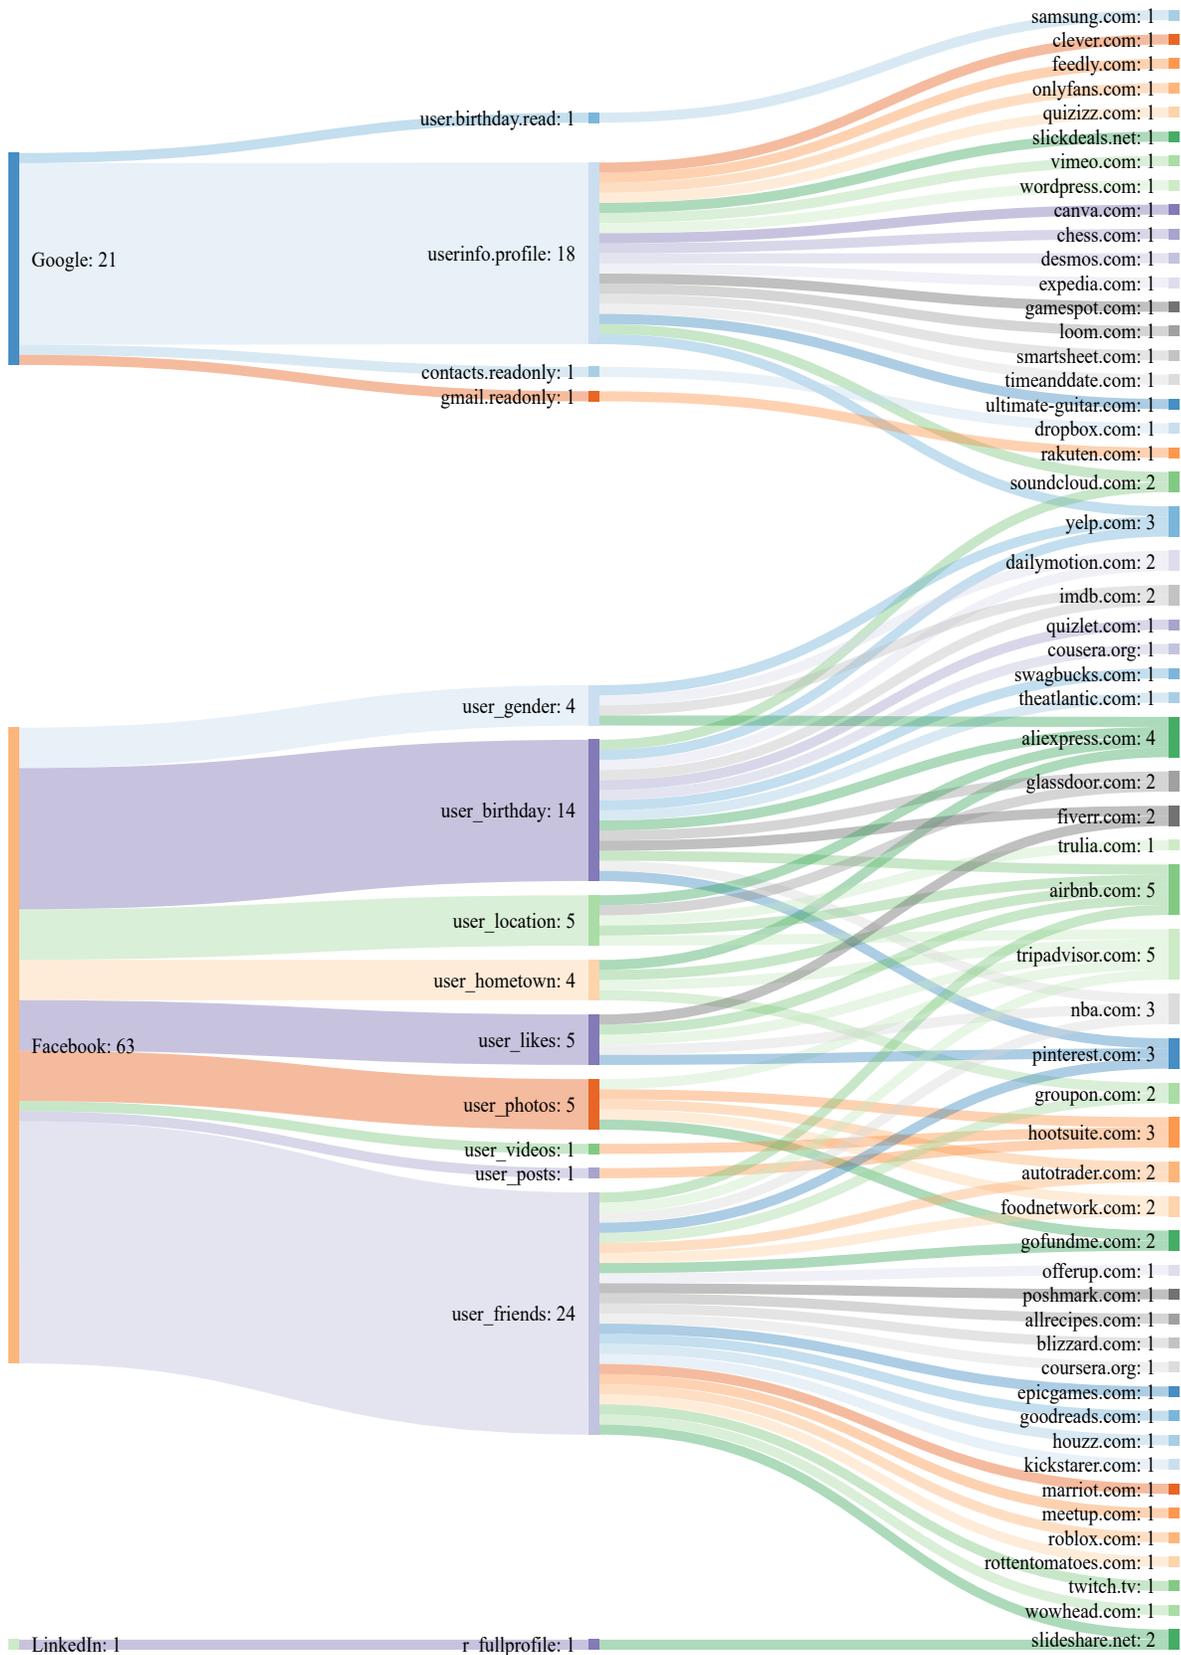
Fig. 5. Data attributes (from Table I) requested by relying parties in the top 500 US sites. For readability, this diagram excludes the *basic* category. Apple is not included in this chart since it only supports *basic* attributes.
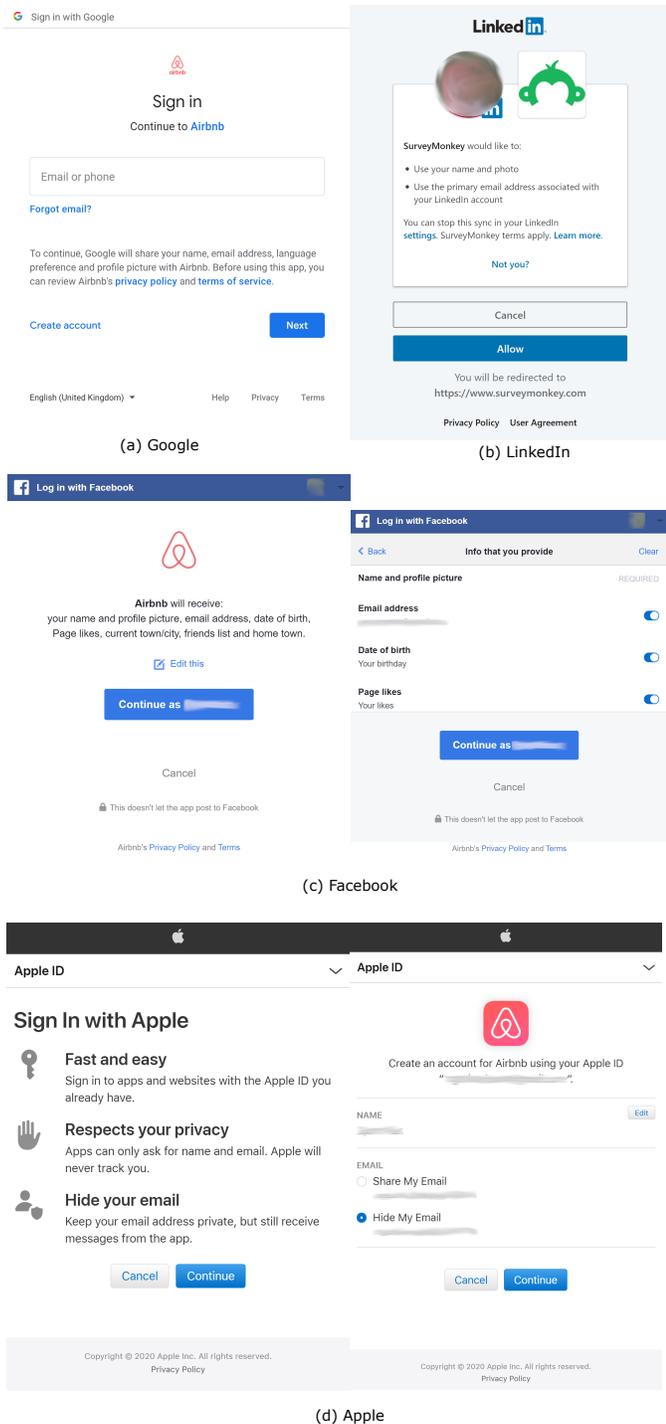
Fig. 6. UI of SSO login forms on IdP sites. (a) Google and (b) LinkedIn do not allow users to alter fine-grained permissions granted to RPs. (c) Facebook allows users to selectively opt-out of non-*default* permissions requested by an RP. (d) Apple allows users to use a substitute name and anonymous email with an RP. In cases of (b), (c) and (d), the IdP presents its default login dialog before showing the SSO screens if the user is not already logged in, as discussed inline. Personal details are greyed out in these images.

(`email`, `openid`, and `profile`) with Google SSO, and thus these do not appear in Fig. 5.

To offer a more complete example, we provide Fig. 7. Here, we identified a subset of RPs that request 2 or more non-*basic* attributes (from Fig. 5) and also include the *basic* attributes requested from other IdPs. Out of the 17 RPs shown in Fig. 7, 14 support an alternate SSO login requesting only the *basic* attributes, suggesting that, in these cases, some user options are less privacy-invasive than others. For the remaining RPs in the figure, nba.com and gofundme.com support only Facebook as a SSO option; slideshare.net supports two SSOs but neither option uses only *basic* attributes.

D. Use of OAuth 2.0 and related OIDC Flows

Table II summarizes the different OAuth 2.0/OpenID Connect flows used by RPs in our dataset. Of particular concern is that a significant number of RPs use the less secure implicit flow, especially in India (38% with Google SSO, 43% with Facebook SSO) and US (33% with Google SSO and 43% with Facebook SSO). RPs using implicit flow receive access tokens in the user's browser directly from the token endpoint. Apple and LinkedIn do not support the implicit flows on their SSO platforms, thus forcing all RPs to use more secure options. As shown in Table II, both Google and Facebook SSOs support both the client-side and server-side flows. From our analysis of countries, we also note that fewer RPs in Germany use the less secure implicit flows compared to other countries.

VII. PRIVACY IMPLICATIONS

In this section, we discuss privacy implications from our evaluation of OAuth 2.0 use in popular SSO services.

A. Impact of interface design

Fig. 6 shows the design of SSO user interfaces (UI) presented to the user when they choose to login using SSO. Usable SSO interfaces are essential in informing users about the permissions requested by an RP. Designs should be as simple as possible and provide users with a path-of-least-resistance to securely complete authentication tasks [37, §9.8]. IdPs must obtain informed consent during SSO workflows, which theoretically provides users with some control over their privacy. However, earlier research on Android's permission system, which tackles a similar issue, has shown that permission warnings are ineffective in informing users about the risks associated with allowing access to applications [14]. Android permission warnings focus on resource access but lack useful information for users to understand the associated benefits and risks [29]. Even when users read permission warnings, they are unaware of risks and simply trust the marketplaces to have reviewed the hosted applications [22]. We draw parallels to SSO permission requests and, through our inspection, observe that the existing SSO UI designs similarly lack useful information for users to convey risks associated with sharing personal data. They also do not convey the value provided to the user from granting RPs access to their data.
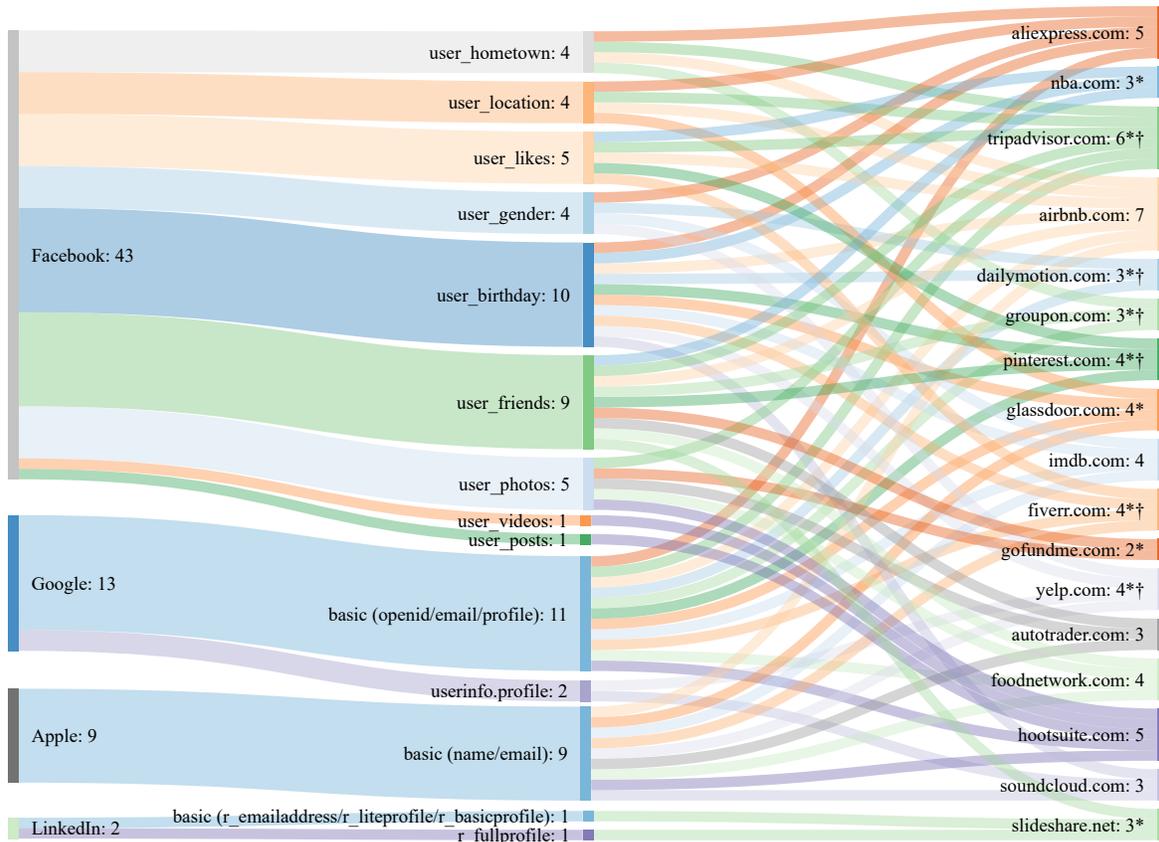
Fig. 7. Comparison of all data attributes requested (including from the *basic* category) by a subset of US RPs. The included RPs request at least 2 non-*basic* permissions. For readability, we exclude redundant connections to the *basic* category for RPs that already request one or more non-*basic* attributes with an IdP, since such requests include *default* fields. RPs using client-side flows in OAuth 2.0 are shown with *(Facebook) and †(Google). The main takeaways are explained inline.

Information about requested permissions is presented differently by each IdP (see Fig. 6). If the user is not already logged in, some IdPs (i.e., LinkedIn, Facebook and Apple) present their default login screens before showing SSO screens. This means that redirected users can view the requested permissions only after they complete authentication with the IdP. This may lead to users granting permissions they would have chosen to deny if the information was presented before they logged in. While some IdPs allow users to edit the requested permissions (shown in Fig. 6), others require users to grant all the requested permissions. Given the differences on the amount of private data requested among available SSO options, we highlight the importance of presenting the requested permissions to users prior to their decision to login with a specific provider.

B. Implications due to Implicit Flow

The OAuth 2.0 implicit flow was created due to past browser restrictions limiting websites to making requests only within its own domain [28]. This prevented JavaScript-based apps from using the authorization code flow since it involves making requests to the IdP domain, which is always different when the RP and IdP are different entities. As discussed in Section II-B, modern browsers support cross-origin requests, allowing the use of more secure flows. Our analysis shows that many RPs still use the implicit flow that returns access tokens in the redirection URL. This is a security concern since URLs are persisted in users' browsing history and it increases the attack surface for access token leakage. Access tokens are a type of "bearer" tokens, and any party in possession can use it to access protected resources without involving the user or RP. Access tokens provide access to specific resources for a limited duration. Users may not be aware of the risks associated with access token leaks, especially when the RP uses client-side flows. The potential damage to users increases when the access token's scope allows excessive access to sensitive user data. RPs can reduce the attack surface by requesting minimum access and using secure flows.

C. Offline data leaks

RPs and IdPs should clearly indicate to users the purpose of requested permissions prior to issuing access tokens. OAuth enables an RP to improve usability for users through customization based on user-data attributes from an IdP. Once granted access, the RP is able to download and persist user data for further processing. As mentioned in Section V, many IdPs review RP applications that request access to sensitive user data. IdPs also provide an interface for users to revoke previously granted access to an RP, invalidating all access

tokens issued to the RP. However, this does not prevent a rogue RP from misusing any user data already accessed. Since user data is processed on RP applications (not controlled by IdP), it is not possible for the user or IdP to be aware of any misuse of accessed data.

Without proper security measures, even a well-intentioned RP can be vulnerable to data breaches that increase the attack surface for users and their data. Although using OAuth ensures that the user's passwords are safe from an attack on RP, leaked access tokens can equally cause damage by allowing attackers to access user data [9]. It can be challenging for users to track attacks on RPs, and understanding the implications requires a mental model of the SSO system that many users lack [36]. Another challenge for users involves access to decommissioned accounts. RPs identify users by trusting the IdP's verification of user credentials. If a user has stopped using an RP or de-linked their IdP account with RP, it may not be possible for the user to later correspond with the RP (e.g., to demand deletion of personal data).

VIII. TOWARD PRIVACY-FRIENDLY OAUTH SSO

Having uncovered examples of services offering a variety of SSO IdP alternatives but with major differences—largely invisible to users—in the categories and amounts of personal user data accessed, we now ask: How might *privacy*, and the *transparency* and *accountability* of RPs, be improved? To this end, we suggest changes below—as much to encourage discussion and exploration of such protocol and interface changes, as for any merit in the specific suggestions per se. Our discussion briefly notes the impact on and roles of four stakeholders: Users, IdPs, RPs, and OAuth Specification authors.

C1: *When registering with an IdP, an RP could be required to provide descriptions justifying each OAuth user-data attribute they plan to request from users.*

This might introduce *value labels* (cf. *privacy nutrition labels*[4]) and convey intended uses and benefits (if any) to users.

C2: *Each RP could disclose to the user before they select an IdP whether one IdP choice will access more user data than others and provide an explanation of potential user benefits.*

C3: *The RP could display information justifying each OAuth user-data attribute requested, before asking a user to select an IdP.*

Such information might be conveyed via a second-level RP user interface. Recall that OAuth next redirects users to the IdP.

C4: *The IdP could re-iterate the justification (e.g., value labels) before asking users to authorize release of RP-requested data attributes.*

C5: *IdPs could enforce exclusive use of server-side flows (i.e., disallow implicit flow) for any RP request involving access to sensitive user data.*

Here "sensitive" is as defined in Table I. Our motivation is that widely scoped access tokens create greater risks, and client-side flows increase the attack surface (see Section II-B). The OAuth spec could mandate the above, as well as the following.

C6: *The OAuth spec could allow optional scope parameters distinct from those denoted mandatory for RP operation.*

Rather than the present all-or-none scenario (forcing users to accept all, or abandon the IdP for a given RP visit), privacy-conscious users could opt-out of selected *optional* parameters.

The above changes could allow audits or privacy compliance checks (by IdP or third parties), and support informed choices by privacy-conscious users. Over time, this could result in RP-IdP pairs following the privacy best practice of requesting only *need-to-know* data, and privacy-friendly IdPs gaining a "preferred-IdP" status based on their user interface quality, and auditing of RP compliance with new IdP-to-RP guidelines for how RPs should display *value labels* to users. We suggest two further items that may help privacy-aware users make informed decisions.

C7: *A third-party tool could be built to shed light on the consequences of different IdP (OAuth attribute) choices.*[5]

C8: *Reputation-based or data-driven community efforts could provide privacy ratings for SSO options at popular RPs.*

Many of these suggestions may of course face numerous hurdles, one being that the agendas of commercial RPs and IdPs are often not aligned with those of privacy-conscious users. However, we argue that awareness and discussion of technical possibilities are important steps towards supporting user privacy, and shedding light on any RP *dark patterns* [27].

IX. RELATED WORK

Zhou et al. [40] built SSOScan, a black-box testing tool to automatically scan the top 20,000 US sites and found at least one serious vulnerability due to implementation flaws in 345 of the 1660 sites that supported Facebook SSO in 2014. Drakonakis et al. [10] built an auditing framework for evaluating web applications for implementation flaws related to authentication and authorization, including applications that support SSO logins. They simulate user interactions to automatically create accounts and login to 25K websites to find that 9,324 domains are vulnerable to leaking sensitive user data to unauthorized parties. Their approach to simulate user interactions and automatically obtain SSO protocol-related information is similar to ours but as noted in our introduction, instead of privacy leaks to unauthorized parties, we evaluate privacy implications in websites that are explicitly granted (although users might not be aware) access to the user's personal data protected by SSO providers. Mainka et al. [24] design a testing framework to investigate malicious IdPs in OpenID implementations and identify four novel attack classes affecting 11 of 16 systems tested. Fett et al. [15] pursue formal analysis of the OAuth 2.0 standard and proofs of security properties for all OAuth 2.0 flow types. Chen et al. [8] evaluate

---

[4]https://cups.cs.cmu.edu/privacyLabel/

[5]Perhaps analogous to AppCensus for mobile app permissions ([1]; Reardon et al. [29]).

the use of OAuth in mobile applications and found 89 of 149 applications incorrectly implemented OAuth, thus making them vulnerable to attacks. In a 2012 field study of popular SSO systems, Wang et al. [38] analysed SSO web traffic through the browser and identified 8 serious flaws in popular RPs and IdPs, allowing attackers to impersonate the victim user.

Mainka et al. [25] analyse the OpenID Connect protocol and identify security flaws similar to vulnerabilities found in other SSO protocols. They implement a fully-automated evaluation tool to identify implementation flaws in OpenID Connect libraries. Bai et al. [6] provide a tool to automatically identify security vulnerabilities in implementations of web authentication protocols including OAuth-based SSO. Yang et al. [39] propose an OAuth 2.0 security testing framework and automatically evaluate four IdPs (Facebook, Sina, Renren and Tencent Weibo) and 500 top-ranked web apps in US and China. Their empirical study reveals web apps that lack TLS protection for OAuth sessions leading to novel exploits.

Addressing challenges related to user awareness, AppCensus [1] (cf. [29]) uses dynamic analysis to reveal privacy implications of granting data access to Android apps. More recently, Apple introduced *privacy labels* [4] to highlight privacy practises to users of iOS apps. Narayanan et al. [27] discuss *dark pattern* designs in online services used to influence less-informed users into choices not in their best interest. Mathur et al. [26] investigate ~11K shopping websites and find 1,818 instances of *dark patterns* designed to increase user purchases. Felt et al.'s [14] user studies evaluate the effectiveness of Android permissions and find 20 of 24 participants were unaware or did not look at permission warnings. Unlike mobile apps where users are given only one set of permissions, SSO users often have the choice (although hidden) to login to a given RP with a less privacy-intrusive alternative so user awareness could significantly impact decisions.

Sun et al. [36] empirically found users hesitant to adopt OpenID due to a lack of understanding and to concerns over releasing personal information. Many users held the misconception that their IdP credentials were shared with the RP. In 2012, Sun et al. [35] also evaluated OAuth 2.0 implementations by three major IdPs (Facebook, Microsoft, Google) and explored 96 RP sites supporting the Facebook SSO. Results revealed several implementation decisions causing security concerns, including possible access token theft. Privacy implications discussed herein complement their work on security implications from identified vulnerabilities. Bonneau et al. [7] surveyed 35 password-replacement schemes and found that compared to other schemes, federated SSO systems offer more benefits across various usability, deployability and security properties. Alaca et al. [2] propose a framework to evaluate 14 web SSO schemes, including OAuth 2.0 and OpenID, and compare various properties including privacy benefits. They identify defining characteristics for each scheme and highlight priorities for stakeholders.

X. Concluding Remarks

OAuth-based systems provide many benefits such as flexibility and convenience to SSO users. Services using OAuth benefit from reduced development costs related to outsourcing identity management. When an RP supports multiple SSO logins, users must commit to an SSO option (and in many cases, complete the authentication) before finding what user data will be requested by the RP. This design means that users never find out what data would be requested by other SSO options, and consequently, are not fully informed about available choices on the RP site. Our empirical results reveal privacy practises where popular RPs request vastly different amounts of user data from different IdPs, with at least one option unquestionably more privacy-intrusive than others, similar to dark patterns found in website designs [26] [27]. SSO users are likely to make privacy decisions not in their best interest, due to the lack of information on available choices.

When granting RPs access to user data, users are not given information on the duration of the access. This lack of information, combined with an RP's ability to extend previously granted access without additional user involvement (Section II-D), poses ongoing danger to user privacy. Further research is needed to mitigate risks related to allowing such continued access by RPs.

To the best of our knowledge, we offer the first in-depth analysis of OAuth-based SSO with a primary focus on user privacy as opposed to security. Based on the empirical work facilitated by our novel OAuthScope tool, we identify 8 areas to improve the privacy of OAuth. These improvements to OAuth's architecture will require effort and cooperation between IdPs, RPs, and specification authors. We hope that greater awareness by technically-savvy users and privacy enthusiasts, of the privacy implications identified through our work (Section VII), may result in further attention to privacy violations, further community-based monitoring, and a more privacy-friendly OAuth-based SSO ecosystem.

Acknowledgment

References

[1] Appcensus. https://www.appcensus.io/, 2020.
[2] F. Alaca and P. C. van Oorschot. Comparative Analysis and Framework Evaluating Web Single Sign-on Systems. *ACM Computing Surveys*, 53(5):1–34, 2020.
[3] Alexa. The top 500 sites on the web. https://www.alexa.com/topsites, 2021.
[4] Apple. App privacy labels now live on the App Store. https://developer. apple.com/news/?id=3wann9gh, 2020.
[5] Apple. Sign in with Apple. https://developer.apple.com/documentation/ sign_in_with_apple, 2021.
[6] G. Bai, J. Lei, G. Meng, S. S. Venkatraman, P. Saxena, J. Sun, Y. Liu, and J. S. Dong. AuthScan: Automatic Extraction of Web Authentication Protocols from Implementations. In *NDSS*, 2013.

[7] J. Bonneau, C. Herley, P. C. van Oorschot, and F. Stajano. The Quest to Replace Passwords: A Framework for Comparative Evaluation of Web Authentication Schemes. In *IEEE Symp. Security and Privacy*, pages 553–567, 2012.

[8] E. Y. Chen, Y. Pei, S. Chen, Y. Tian, R. Kotcher, and P. Tague. OAuth Demystified for Mobile Application Developers. In *ACM CCS*, pages 892–903, 2014.

[9] C. Cimpanu. Hackers stole GitHub and GitLab OAuth tokens from Git analytics form Waydev. https://www.zdnet.com/article/hackers-stole-github-and-gitlab-oauth-tokens-from-git-analytics-firm-waydev/, 2020.

[10] K. Drakonakis, S. Ioannidis, and J. Polakis. The Cookie Hunter: Automated Black-box Auditing for Web Authentication and Authorization Flaws. In *ACM CCS*, 2020.

[11] Facebook. App Review. https://developers.facebook.com/docs/app-review, 2021.

[12] Facebook. Graph API. https://developers.facebook.com/docs/graph-api/, 2021.

[13] Facebook. Permissions Reference. https://developers.facebook.com/docs/permissions/reference/, 2021.

[14] A. P. Felt, E. Ha, S. Egelman, A. Haney, E. Chin, and D. Wagner. Android Permissions: User Attention, Comprehension, and Behavior. In *SOUPS*, pages 1–14, 2012.

[15] D. Fett, R. Küsters, and G. Schmitz. A Comprehensive Formal Security Analysis of OAuth 2.0. In *ACM CCS*, pages 1204–1215, 2016.

[16] M. Ghasemisharif, A. Ramesh, S. Checkoway, C. Kanich, and J. Polakis. O Single Sign-Off, Where Art Thou? An Empirical Analysis of Single Sign-On Account Hijacking and Session Management on the Web. In *USENIX Security*, pages 1475–1492, 2018.

[17] Google. Google API for Authentication. https://developers.google.com/identity/sign-in/web/reference#users, 2021.

[18] Google. OAuth 2.0 Scopes for Google APIs. https://developers.google.com/identity/protocols/oauth2/scopes, 2021.

[19] Google. OAuth API verification FAQs. https://support.google.com/cloud/answer/9110914, 2021.

[20] D. Hardt. RFC 6749: The OAuth 2.0 Authorization Framework. https://tools.ietf.org/html/rfc6749, 2012.

[21] M. B. Jones, J. Bradley, and N. Sakimura. RFC 7519: JSON Web Token (JWT). https://tools.ietf.org/html/rfc7519, 2015.

[22] P. G. Kelley, S. Consolvo, L. F. Cranor, J. Jung, N. Sadeh, and D. Wetherall. A Conundrum of Permissions: Installing Applications on an Android Smartphone. In *Financial Cryptography and Data Security*, pages 68–79. Springer, 2012.

[23] LinkedIn. Sign In with LinkedIn. https://docs.microsoft.com/en-us/linkedin/consumer/integrations/self-serve/sign-in-with-linkedin, 2018.

[24] C. Mainka, V. Mladenov, and J. Schwenk. Do Not Trust Me: Using Malicious IdPs for Analyzing and Attacking Single Sign-on. In *IEEE EuroS&P*, pages 321–336, 2016.

[25] C. Mainka, V. Mladenov, J. Schwenk, and T. Wich. SoK: Single Sign-On Security — An Evaluation of OpenID Connect. In *IEEE EuroS&P*, pages 251–266, 2017.

[26] A. Mathur, G. Acar, M. J. Friedman, E. Lucherini, J. Mayer, M. Chetty, and A. Narayanan. Dark Patterns at Scale: Findings from a Crawl of 11K Shopping Websites. *ACM Human-Computer Interaction*, 3(CSCW):1–32, 2019.

[27] A. Narayanan, A. Mathur, M. Chetty, and M. Kshirsagar. Dark Patterns: Past, Present, and Future. *ACM Queue*, 18(2):67–92, 2020.

[28] A. Parecki. Is the OAuth 2.0 Implicit Flow Dead? https://developer.okta.com/blog/2019/05/01/is-the-oauth-implicit-flow-dead, 2019.

[29] J. Reardon, A. Feal, P. Wijesekera, A. E. B. On, N. Vallina-Rodriguez, and S. Egelman. 50 Ways to Leak Your Data: An Exploration of Apps' Circumvention of the Android Permissions System. In *USENIX Security*, pages 603–620, 2019.

[30] N. Sakimura, J. Bradley, and N. Agarwal. RFC 7636: Proof Key for Code Exchange by OAuth Public Clients. https://tools.ietf.org/html/rfc7636, 2015.

[31] N. Sakimura, J. Bradley, M. B. Jones, B. de Medeiros, and C. Mortimore. OpenID Connect Core 1.0. https://openid.net/specs/openid-connect-core-1_0.html, 2014.

[32] Selenium. Selenium WebDriver. https://www.selenium.dev/documentation/en/webdriver/, 2021.

[33] M. Singh. Why Apple sells just 2.5% of India's smartphones. https://www.cnbc.com/2018/01/29/why-apple-sells-just-2-point-5-percent-of-indias-smartphones.html, 2019.

[34] E. Stobert and R. Biddle. The Password Life Cycle: User Behaviour in Managing Passwords. In *SOUPS*, pages 243–255, 2014.

[35] S.-T. Sun and K. Beznosov. The Devil is in the (Implementation) Details: An Empirical Analysis of OAuth SSO Systems. In *ACM CCS*, pages 378–390, 2012.

[36] S.-T. Sun, E. Pospisil, I. Muslukhov, N. Dindar, K. Hawkey, and K. Beznosov. What Makes Users Refuse Web Single Sign-On? An Empirical Investigation of OpenID. In *SOUPS*, pages 1–20, 2011.

[37] P. C. van Oorschot. *Computer Security and the Internet: Tools and Jewels*. Springer Nature, 2020.

[38] R. Wang, S. Chen, and X. Wang. Signing Me onto Your Accounts through Facebook and Google: A Traffic-Guided Security Study of Commercially Deployed Single-Sign-On Web Services. In *IEEE Symp. Security and Privacy*, pages 365–379, 2012.

[39] R. Yang, G. Li, W. C. Lau, K. Zhang, and P. Hu. Model-based Security Testing: An Empirical Study on OAuth 2.0 Implementations. In *AsiaCCS*, pages 651–662, 2016.

[40] Y. Zhou and D. Evans. SSOScan: Automated Testing of Web Applications for Single Sign-On Vulnerabilities. In *USENIX Security*, pages 495–510, 2014.

APPENDIX

We provide an example authorization request discussed in Section III. RPs specify OAuth 2.0 parameters in authorization requests to the IdP. A brief description is included for each parameter in the request [20].

```
HTTP GET /authorizationEndpoint?
response_type=code
&scope=email%20profile
&redirect_uri=https%3A%2F%2Fclient%2Ecom%2Fcb
&client_id=lp4qazfnh1
&state=hnz3krb2mn
```

authorizationEndpoint: endpoint URI used by the RP for sending authorization requests to the IdP.

response_type: specifies the OAuth flow type the RP intends to use with IdP.

scope: a list of resources requested for access by the RP.

redirect_uri: user is redirected to this endpoint after completing interactions with the IdP. For security reasons, this value must match the endpoint registered with the IdP during RP's app registration.

client_id: a unique string issued to RP during registration.

state: a unique (non-guessable) string generated by RP and included in the authorization request. The IdP returns the value when redirecting the user back to RP. To mitigate cross-side request forgery attacks, it must be ensured that the returned value is equal to value included in the initial request.

As an extension to background provided in Section II, Fig 8 lists the process for the OAuth 2.0 implicit flow. Since the access token is returned to the RP in the redirection URI, it is vulnerable to token thefts from the user's browser.

Fig. 9 is a screenshot of OAuthScope described in Section III and lists identified OAuth 2.0 parameters for each RP. This UI is used for analysis of data collected by OAuthScope.
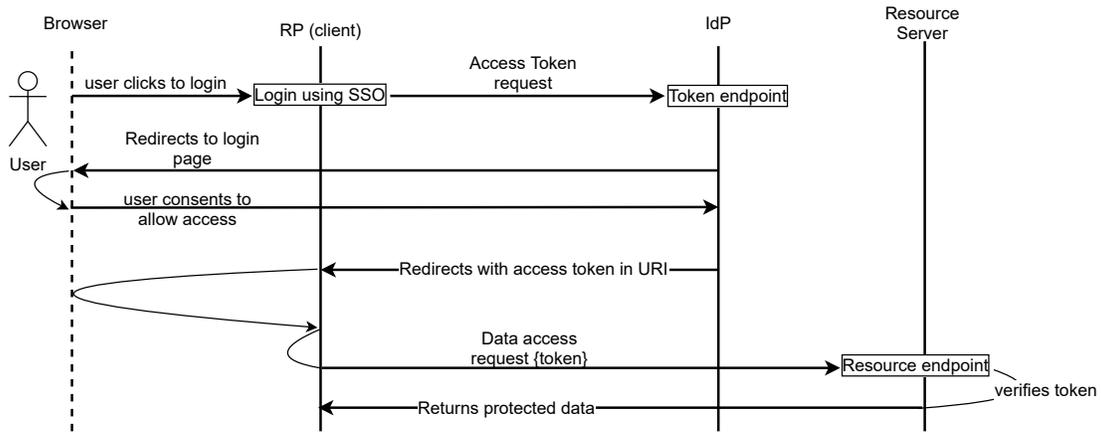
Fig. 8. Procedure for the OAuth 2.0 Implicit flow (derived from [20]).

Fig. 9. Screenshot of OAuthScope tool listing OAuth 2.0 parameters included in authorization requests from top US sites.