Network Working Group R. Copeland, Ed.
Internet-Draft Institut Mines Telecom-Telecom Sud Paris
Intended status: Informational K. Corre
Expires: March 30, 2017 Orange Labs
I. Friese
Deutsche Telekom AG
S. El Jaouhari
Institut Mines Telecom-Telecom Bretagne
September 26, 2016

Requirements for Trust and Privacy in WebRTC Peer-to-peer Authentication
draft-copeland-rtcweb-p2p-idp-auth-00

Abstract

This document studies the relationships of WebRTC communication users with their web Calling Services (CS) and their Identity Providers (IdPs), in order to identify requirements for IdP based peer-to-peer authentication. This study focuses in particular on issues of privacy, security and trust that are raised by the introduction of the IdP into the WebRTC call model, and by a different browser-based calling paradigm, compared with Mobile networks or traditional VoIP systems. The document lists privacy and trust scenarios for WebRTC authentication for individuals as well as organizations. This contribution is proposed to the RTCWEB working group.

Status of This Memo

Copyright Notice

Table of Contents

1. Introduction   3
2. Terminology   4
3. Call Context Aspects   6
  3.1. Existing Protocols and Drafts   6
  3.2. WebRTC Architecture Components   7
  3.3. WebRTC Call Models   7
    3.3.1. Single-CS Call Model   7
    3.3.2. Dual-CS Call Model   7
  3.4. Service Types and their Trust Models   9
  3.5. Destination Categories 10
4. Architecture Vulnerabilities 10
  4.1. Untrusted Parties 10
  4.2. Network Messages Across Multiple Parties 10
  4.3. Confidentiality of Call Logs 11
  4.4. Dependency of Identifiers 12
  4.5. IdP Selection Issues 12
5. Identity Privacy 12
  5.1. Desirable and Undesirable Identity Privacy 13
  5.2. Undetectable Calling 13
  5.3. Pseudonymous Calling 13
  5.4. Unlinkable Calling 14
  5.5. Potential Methods of Identity Protection 14
    5.5.1. Sensitive User Information 14
    5.5.2. Proposed Surrogated identities with Pseudonyms 14
6. Trust Relationships 15
  6.1. Three-Way Trust: User-CSP-IdP 15
  6.2. Choice Indicates Trust 16
  6.3. User trust in Calling Services 16
  6.4. User trust in Identity Providers 17
  6.5. Communication Service trust in Identity Providers 17
  6.6. Trusted Identities for non-Browser Interworking 18

7. Use Cases for Privacy Requirements 18
7.1. Anonymous Caller Connecting to Call-Centers 18
7.2. Call Center Worker's Privacy 19
7.3. Online Gaming Calling by Pseudonyms 19
7.4. Hosted Enterprise WebRTC Conferencing Service 20
7.5. Variable Trust modes for Employee's Calls 20
7.6. Employee using untrusted WebRTC service 20
7.7. WebRTC service Interworking with SIP users 21
8. Requirements Summary 21
8.1. Anonymity 21
8.2. Unlinkability 22
8.3. Independent IdP 22
8.4. User Information Confidentiality 23
8.5. Calling Services 23
8.6. Usability and Notification 23
9. Conclusions 24
10. Contributors 24
11. Acknowledgements 25
12. References 25
12.1. Normative references 25
12.2. Informative references 26
Authors' Addresses 27

1. Introduction

This study provides requirements for supporting identity privacy in peer-to-peer (P2P) WebRTC calling services. While the WebRTC specification aims to manage the media flow, it does not include procedures for call initiation and privacy setting. However, WebRTC architecture supports peer authentication that is performed separately from the calling service, decoupling user authentication from the granting of service resources. The authentication is performed by a third party Identity Provider (IdP), while service resources are granted by each Calling Services Provider (CSP).

WebRTC calling creates a different paradigm from 'traditional' VoIP services or Telecom, because it is browser-based. All that is needed for a website to support WebRTC calling is to download a client to the device to access the user-agent JavaScript APIs. This simplicity will encourage many websites to add calling facilities to their 'shop-window'. In turn, such easy click-to-link services will entice occasional website visitors to initiate 'opportunistic' calls, often using unknown, maybe untrusted calling services, giving little thought to the risk of leaking personal information. Hence, this paradigm increases risks of abuse of user data, identity theft and commercial exploitation. This necessitates establishing appropriate privacy protection, even without user explicit input of preferences. It is possible to achieve this by attaching privacy rules to the

IdP’s maintained user identity, so that the IdP will support anonymity where required.

An independent identity should prevent the undesirable lock-in effect of ’service-bound’ identities and allow for a single identity, with its linked pseudonyms identifiers (with same credentials), to be used instead of numerous identifiers and separate passwords. Users should be able to specify particular privacy rules that are applied during authentication across all services, or for a particular service type, since the privacy rules are attached to the identifier, not to the service.

The current peer authentication procedure provides some flexibility for the choice of IdP, but does not allow users to determine privacy requirements for different circumstances. There are no means of evaluating trust models and required privacy between calling parties, their CSs and IdPs. The call model (single or dual CS model) affects the level of trust in the other parties. Users may trust their chosen IdPs and CS, but the same cannot be assumed for CS and IdPs of their communication partners. The service type and the means of activating it also influence the trust level, e.g. a social media contact may be more trusted than a call to an unknown website. Similarly, the type of destination (e.g. public organizations versus unfamiliar websites) also impacts the trust level. Such destinations have their own privacy requirements that need to be negotiated, e.g. callers’ traceability may be mandated in order to avoid nuisance calls, but at the same time, unlinkability is required for employees when they respond to public enquiries.

The accumulated record of calls is a precious commodity in the monetized web space, but many users wish to exercise better control of what is divulged. To solve this, it is argued that by using context-based privacy to obscure certain details or to present surrogate identities, the calling service logs will contain only filtered information, in accordance with users’ wishes.

Hence, this study proposes requirements for user-controlled, multi-purpose, identity, with service-independent authentication by IdPs, which is protected not only by preferences and negotiated privacy rules, but also by detected context-based privacy settings.

2. Terminology

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119 [RFC2119]

Privacy and data minimization terms, such as Anonymity, Unlinkability, Undetectability, and Pseudonymity, follow the definitions by Pfitzmann et al. [TerminologyForPrivacy], but refer to inter-party interactions, not user-to-server, where the 'sender' (Caller) and the Destination (called party) are separate entities from their own services and endpoint clients. Table 1 contains further terms that convey a particular meaning or an extension meaning in this document.

Terms Definitions

Term	Description
Called-party, Destination or Callee	A target destination, as nominated by the caller, who is an individual called party, an organization representative or an intelligent object, with a WebRTC valid identity.
Caller or Sender	An individual, an organization representative or an intelligent object with a valid identity who is initiating a WebRTC call or session by using a WebRTC-enabled browser.
Calling Service (CS)	The WebRTC based service that is used to connect calling parties. This service may provide calling features and group calling management, as well as users' preferences and group privacy policies. A CS is assumed to be server-side software from a CS Provider, with clients downloaded to the users' endpoints.
IdP (Identity provider)	IdP (Identity Provider) creates, maintains, and manages identity information for entities (users, services, or systems) and provides authentication to other service providers. It is a trusted third party that can be relied upon by users and servers when they establish a dialog that must be authenticated.
IdP Proxy	A client (user agent) constructed by each IdP for its own operation, and is downloaded to users' endpoints when they wish to verify a given 'assertion' for a user.
Unlinkability or Untraceability	While the definition in [TerminologyForPrivacy] refers to unlinkability of subject to a message or a particular attribute, in this document it is defined as the inability to link calling parties to their routable address for calling back.
Privacy Rules	Rules containing parameters for service delivery that are compiled from preferences of the

	involved individuals, groups, or institutions, to determine when, how, and to what extent information is divulged.
Surrogated Identities	Identities that identify a real person uniquely through their association with a universal unique surrogate key (a database indexing key that is system-generated as an artificial primary key value, independent of any attribute). Such identities need not have their own credentials, and can use either meaningful pseudonyms or meaningless (anonymous) identifiers.

Table 1

3. Call Context Aspects

This section describes the dependency of privacy settings on the call model (single/dual CS), service type and destination category, in the light of the P2P authentication procedure.

3.1. Existing Protocols and Drafts

WebRTC calls connect two browsers that are exchanging media and data, as presented in the "WebRTC Overview" [I-D.ietf-rtcweb-overview]. Prior to the media connection, the calling parties may authenticate each other, in a possibly reciprocal peer-to-peer authentication process, as described in "WebRTC Security Architecture" [I-D.ietf-rtcweb-security-arch].

The RFC "Privacy Considerations for Internet Protocols" [RFC6973] offers guidance for privacy considerations in Internet protocols. This RFC identifies privacy threats and threat mitigation solutions which includes data minimization.

The Internet draft "Security Considerations" [I-D.ietf-rtcweb-security] for WebRTC identifies two threats to privacy in WebRTC. These are the correlation of anonymous calls through persistent identifiers such as DTLS certificates, or the risk of browser fingerprinting through the WebRTC API itself.

Examples of WebRTC binding to specific identity protocols are given in "WebRTC Security Architecture", such as OAuth 2.0 [RFC6749] (or OpenID Connect [OIDC]). However, these protocols do not address privacy and identity management for peer-to-peer authentication.

3.2. WebRTC Architecture Components

Current WebRTC communications rely on standard browser APIs (getUserMedia) that execute at both endpoints to enable media streaming, and WebRTC APIs (RTCPeerConnection and RTCDatachannel) that execute at the endpoints to manage the flow. The endpoints also execute a CS client, which drives the initial call setup.

To enable decoupling of the CS from the identity authentication process, WebRTC security architecture [I-D.ietf-rtcweb-security-arch] proposes that the WebRTC PeerConnection component interacts directly with the IdP, to initiate and manage the authentication process. It negotiates Session Description Protocol (SDP) with the other party by sending an SDP offer and accepting, rejecting or offering a counter offer.

Both parties set up their own PeerConnection instances and download an IdP proxy from their own IdP. Each IdP authenticates its user’s and returns an identity assertion containing the identity and session key fingerprint as claims. When a party receives an SDP offer or answer containing an identity assertion, that party also downloads the IdP proxy from the other party’s IdP. This proxy is then used to verify the received identity assertion. Each IdP Proxy only connects to its own IdP server. This architecture is presented in Figure 1.

3.3. WebRTC Call Models

3.3.1. Single-CS Call Model

In a single-CS model, which is the dominant model in the current Internet Voice services, only one calling service is involved, where both the caller and the called party are registered to the same service. The service manages the subscribing users’ privacy policy via a set of options, which are then reconciled for the call. These services utilize ’service-bound’ identities, where users must log on with the service-specific credentials. By contrast, the WebRTC Single-CS call model permits users to select their own identities, and perform user-to-user mutual authentication, even though both users are logged on to the same service.

3.3.2. Dual-CS Call Model

While early implementations of WebRTC calling between individuals go no further than the single-CS model, it is expected that the dual-CS model will become more common if WebRTC services are adopted in business, especially for small-to-medium enterprises. In dual-CS model, the calling parties log on to two different CSs, with their own sets of privacy rules and security policies, so CSs must discover the respective privacy/security requirements and negotiate an acceptable set of rules for both parties per session.

Figure 1 shows the calling parties (A and B) using different services (CS-A and CS-B) and different IdPs (IdP-1 and IdP-2). It shows the mutual P2P authentication that is performed by each side symmetrically. A gets its own identity assertion from IdP-1 and verifies B via a downloaded IdP-2 proxy. B gets its own assertion from IdP-2 and verifies A by the downloaded IdP-1 proxy. Each proxy runs in its own sandbox to protect it from interference. The mechanism for interaction between calling Services can be, for example, SIP or XMPP.

```
                    +----------+           +----------+
                    | Caller’s |Unspecified| Called   |
    +----------+    | Service  | Protocol  | Service  |    +----------+
    | IdP-1    |    |  CS-A    |<--------->|  CS-B    |    | IdP-2    |
    |          |    | (caller) |(SIP,      | (callee) |    |          |
    +-v------^-+    +----^-----+  XMPP,...)+--------^-+    +--^-----v-+
      |      |           /                          \         |     |
      |      |          /                            \        |     |
      |      |         /                              \       |     |
      | +----|----------------+             +-----------------|---+ |
      | |    |                |    Media    |                 |   | |
      | |    | Browser A      |<------------>|    Browser B    |   | |
      | | +-------+ +-------+ |             | +-------+ +-------+ | |
      | | | IdP-1 | | IdP-2 | |             | | IdP-1 | | IdP-2 | | |
      | | | Proxy | | Proxy | |             | | Proxy | | Proxy | | |
      | | |sandbox| |sandbox| |             | |sandbox| |sandbox| | |
      | | +-------+ +---^---+ |             | +----^--+ +-------+ | |
      | +---------------|-----+             +------|--------------+ |
      |                 |                          |                |
      |                 |                          |                |
      +-----------------|--------------------------+                |
       B verifies A     |                                           |
                        +-------------------------------------------+
                                                  A Verifies B
```

Figure 1: Dual CS Calling Model

Calling parties may connect to unknown parties who are using unfamiliar services across the globe, and are authenticated by IdPs that the caller may not know or even be aware of. In such cases, many users would prefer to prevent unnecessary data disclosure.

3.4. Service Types and their Trust Models

In order to understand the privacy requirements for WebRTC architecture, it is important to classify calling services by the manner of initiating the sessions, choosing destinations (called parties), and consulting the other party CS. The following is a classification of web calling service types that may have varying privacy requirements, due to different trust models:

a. Contact-List-Service: The service enables users to maintain their contact list, and make or receive calls from them. This service type is used by social media, VoIP OTT, and internal enterprise call servers.

b. Click-to-Link Service: Browsing users activate this service type when they click on a website link to talk to anyone at that destination, not to an individual person. The 'linked' party in the visited website uses its own CS in a single-CS call model. This service type is common for shopping websites and customers' enquiries, which today are mostly chat services only. The caller may not trust such services, and users often wish for anonymity to be preserved. On the other hand, unlinkability is often required for the responding individuals at the website.

c. Negotiated Service: This type of service, which is common for business-to-business, allows both calling parties to have their own privacy rules. In a single-CS call model, the service reconciles the differences, but in a dual-CS call model, the calling services should conduct a dialogue resulting in negotiated privacy rules per call.

d. Interworking Service: WebRTC calling services should enable connecting to other types of services, using temporary or 'interworking' special identities. This service type is used, for example, for interworking with SIP servers and telephony. Although users cannot authenticate each other, the interworking depends on one-sided IdP authentication.

e. Conferencing Multi-Party Service: Services that connect several parties in one call need different mechanisms to mix the media, such as a bridge, router/mixer or a matrix, but they also need to reconcile privacy needs of several parties. This is often a single-CS service between subscribing participants, but it may also support multiple calling services, where each call leg is managed by the caller's own CS. Peer authentication could still be performed between the conferencing shared identity and each participant.

These service types influence the required privacy, since they are correlated to particular trust models, e.g. contact-List calling permits full information exchange, but Clink-to-Link sets the strictest privacy level.

3.5. Destination Categories

The destination category influences the level of security and privacy that the calling parties require. The destination may be an untrusted web site in click-to-link service type that users prefer to connect while preserving anonymity and/or unlinkability, but such a destination may also be a completely trusted, often used website. Businesses often wish to apply different confidentiality rules for internal or external calls, i.e. distinguishing between classes of destinations. Hence, privacy rules should be determined per destination category, by: the addressing mode (click-to-link or contact list); the domain (government etc.); or previously logged calls. In the absence of explicit user preference, context-based privacy level can be set by the CS according to the destination category. The IdP can also perform such a service, based on the call log and the target domain names.

4. Architecture Vulnerabilities

This section considers the new paradigm that webRTC creates, which gives rise to different trust models between the parties and other stakeholders.

4.1. Untrusted Parties

The decoupling of the authentication from the service logic, both in networked functions and in business entities, increases vulnerability, but also the variety of trust models that are needed to support permissive and intuitive web calling patterns. This environment encourages users to take more risks and launch interactions without careful considerations of the information that may leak to various parties. In additions, not only a casually invoked service may receive user data, but also the IdP and the CS of the other party, who may not be trusted.

4.2. Network Messages Across Multiple Parties

WebRTC architecture relies on the IdP to perform the authentication independently from the call initiation process. This entails further network based interactions and more parties gaining knowledge of them. The additional network messaging between more parties increase the risk of Man-in-the-Middle (MITM) attacks. The number of involved parties in person-to-person calls rises in peer authentication, depending on the call model. A single-CS call model with service-bound identification involves only one CS with two parties. In the dual-CS call model with independent IdPs, there are two calling parties, two service providers and two IdPs, i.e. double the number of involved parties. This increases the risk of leaking information and abuse of call history, as well as MitM attacks.

4.3. Confidentiality of Call Logs

Currently, both IdPs and calling services acquire user knowledge, which is a) contextual; b) historical; and c) subscription-based. This information can be static (user personal details, additional email identities) or dynamic (calling patterns, frequent destinations, preferred services/websites, and more). Every time a new call is set up, the exchanged information can provide means of tracking user activity or linking back to the parties. The accumulation of such information reveals trends, habits, preferences and behavior patterns that are highly valuable to traders and marketeers. Exploiting this data is often the only monetization methods available to web service providers, but it is a cause for concern for users who regard it as compromised privacy.

The IdP gains knowledge of personal details (to enhance user profile), and associated identities (to enhance identity resilience), which users may be reluctant to share with numerous websites. In addition, the IdP can log authentication requests of every CS using its identity service, while the CS only has visibility of calls made to and from that service. Since users’ calling activities are now spread over a number of web communications services, an IdP will have wider perspective on the user’s intelligence.

IdP is not able to know much about the involved CS, as the Peer-Connection method interfaces between the endpoint’s client and the IdP Proxy directly, not via a CSP server. On the other hand, the CS has better understanding of the call context in the preamble before initiating a call request.

As the IdP Proxies are deployed on both calling parties’devices, the IdPs can log identity verification requests for incoming as well as outgoing calls. It may be possible for an IdP to track call history for a particular destination user who is using another IdP, even if an pseudonymous identity is given, and perhaps eventually link the records with the real identity. However, this risk only arises with habitual pseudonymous calling to the same destination.

4.4. Dependency of Identifiers

The aim of providing completely independent and portable identities is not easy to achieve. While the IdP-generated identity (with an IdP domain) is not related to any specific calling service, i.e. portable between services, it is still dependent on the IdP. If users wish to keep their well-known published identifiers when moving to another IdP, they need identities that ensure both CS and IdP independence. This requires users to use a global user identifier, such as a Universally Unique Identifier (UUID) [RFC4122] that would acts as a unifying key for all identities. This globally unique identifier could link several identifiers, both IdP-generated and CS-generated.

4.5. IdP Selection Issues

In WebRTC, each CS client in the device is responsible for setting up the authentication requests for its own party. The CS client decides what form of authentication to apply, i.e. peer authentication, server-side Single Sign-On, or service-specific authentication. This means that the CS controls the selection process, and may restrict the choices of IdP to choose from, or even prevent an IdP to be involved.

Current WebRTC specifications define two options for the CS to select an IdP for an identity assertion request:

- o If the setIdentityProvider() method has been called by the CS, the provided IdP will be used.

- o If the setIdentityProvider() method has NOT been called, the browser may use a pre-configured IdP.

Pre-configuring an IdP via the browser means that yet another party - the browser vendor - is a stakeholder in the WebRTC call initiation. It is argued that currently, users do not have sufficient control on the selection of the IdP with these facilities.

5. Identity Privacy

This section sheds a new light on the privacy requirements for undetectable, pseudonymous, and unlinkable callers that arise from the webRTC peer calling. Although these terms do exist, the associated privacy requirements have not been previously identified.

5.1. Desirable and Undesirable Identity Privacy

Many websites may be content to receive enquiries from anonymous callers, because this may generate impulse-buying, so they only request user details before completing a transaction. Such websites also require some privacy rules themselves, to protect specific personnel serving at a call center. Web calling encourages opportunistic calling by users who are merely visiting the websites, where users identities are ’incognito’, i.e. their status is ’undetectable’ or ’unobservable’. In certain circumstances, calls from undetectable identities should always be supported, e.g. calls to emergency services that are passed through without any authentication. While undetectable status is passive, in other cases callers may specifically wish to withhold their personal details for a variety of legitimate reasons, e.g. to avoid revealing interests in sensitive material or avoid personal embarrassment. In such cases, users can choose to use assumed names (pseudonyms). Similarly, there are good reasons to support the requirement of unlinkability that prevents tracing back previous calls, e.g. to avoid traders chasing business.

The contrasting requirements to prevent anonymity should also be considered, in order to prevent abuse, e.g. for nuisance calls or malicious disruption of service. The solution to block all callers who are not on the personal contact list may suffice for individual users, but this is too restrictive for a business.

Therefore, since privacy protection is both desirable and undesirable depending on context and point of view, privacy rules need finer granularity, so that they can be applied judiciously, according to context and circumstances. Hence, the requirements are for WebRTC authentication to support different states of user privacy and anonymity: Undetectable (undisclosed identity), Pseudonymous (false or fictitious identity) or unlinkable (untraceable address/path).

5.2. Undetectable Calling

Undetectable calling may be initiated without logging to any CS, while the user is unknown. When a call is made without logging to the CS, a call request may be processed without any authentication.

5.3. Pseudonymous Calling

Pseudonymous calling means that the username identifier is replaced with an assumed name that hides the user’s identity. The authentication can be performed by an IdP who is aware of the pseudonym owner. Hence, the other parties can be reassured that the identity is verified. Such authentication may not be sufficient for monetized transaction and non-repudiation, but is considered acceptable for web calling.

5.4. Unlinkable Calling

Unlinkability prevents other parties from calling back, or from tracing the user's cyber activities, such as visited websites and calling patterns. Unlinkable callers seek to hide the originating website or redirected services. Unlinkability is also both desirable and undesirable, depending on the context. Preventing linkability is often needed to protect individual employees who respond to enquiries from the public. Conversely, linkability is highly desirable by emergency and health services, to locate incapacitated callers in distress.

5.5. Potential Methods of Identity Protection

5.5.1. Sensitive User Information

User information that may be subject to privacy includes:

- o Forename and last names, which are often incorporated in the username part of the identifier.

- o Domain name of associated organization, which often incorporates the organization name in the @domain part.

- o Message path and IP address, which are revealed in the SDP (Session Description Protocol).

5.5.2. Proposed Surrogated identities with Pseudonyms

Using pseudonyms avoids undesirable disclosure of the identity and/or incidental private information. However, pseudonyms should still be associated by a common key to the real user. This is achieved by a fully independent identifier that acts as a 'surrogate' key, i.e. an indexing key that is not based on any meaningful personal details, as described in [SurrogateKeys] for indexing data. Such surrogate keys provide stability, because they are not affected by users changing circumstances (e.g. married names) and personal attributes (e.g. changing employer domain name).

The surrogate key for identities is a Global User ID (GUID) that uniquely identifies a particular user. GUIDs can be generated by a number of known algorithms. They may inject user-specific variable attributes as 'salt' to the otherwise random number generator, to ensure uniqueness. It is proposed that this GUID/surrogate key will link the IdP-based identity, service-bound identities and pseudonym identities, at the discretion of the user. All the linked identities (i.e. the 'surrogate identities') can share the same credentials that the IdP has verified. Service-bound identities from a variety of calling services that do have their own credentials (usually just passwords) can also link to the surrogate key, thus benefiting from deeper user verification and from the SSO effect that such arrangement brings.

6. Trust Relationships

This section analyzes discernible trust models which are proposed as the basis of setting up appropriate privacy levels.

6.1. Three-Way Trust: User-CSP-IdP

The current WebRTC security architecture only assumes that users trust their CSP, or that an IdP is used in P2P authentication if the CS is untrusted. Trust in the IdP is only considered regarding its verified, https origin. In this model, the browser constitutes the trust anchor. However, this simple trust model does not describe trust in the privacy context, nor the difference between a user's own IdP and the other IdP.

Users need to trust other involved actors, i.e. CSs and IdPs, to manage their privacy and provide solid identity claims. The CS in turn may need to trust both IdPs regarding user authentication and identity claims. However, IdPs do not require particular trust relations with the CS, as they merely provide a service, without risk to themselves. Figure 2 and subsequent sections details this trust model.

```
                            +---------+
        +------------------->   CSP   |
        |                   +---------+
    +---------+              |
    |  Alice  |              |
    +---------+              +-------+
     |  |                    |       |
     |  |            +-------v-+  +--v------+
     |  +------------>  IdP A  |  |  IdP B  |
     |               +---------+  +-^-------+
     |                              |
     +------------------------------+
```

Figure 2: Trust relationships between communication setup actors

6.2. Choice Indicates Trust

The purpose of establishing trust is to make a decision in situations where an action with an inherent risk depends on informed judgement. Explicit choice can be interpreted as a measure of trust. Users' chosen IdPs are more trusted than mandated IdPs imposed by the communication services, though enterprise mandated IdPs are trusted by virtue of the enterprise selection. Similarly, calling services that are casually invoked by click-to-Link in a visited website are less trustworthy than those which the user has registered to.

Trust is also assumed towards a call-party that is known to the user, but there is no implied trust level for the calling services and IdPs of that party. Trust may be associated with the type of the call destination, which can be categorized as:

- o Fully trusted parties (government, public organization, known business)
- o Inner-circle of a social network or family members
- o Uncertain (not-in-contact-list, no information)
- o Untrusted (known as unreliable).

6.3. User trust in Calling Services

Traditionally, CSP had access to full user profile information and accumulated call history, but user habits now favor using different services according to context, so the CS has only their own usage records. While some CS may enjoy greater trust, in other cases, users do not wish to share or even store their call history. These preferences are usually agreed when users register with the CS, but it is up to the CS to respect them.

Since the CS manages the call signaling, it is well placed to intercept the peer-to-peer media stream that otherwise is deemed private. Even if the caller's CS can be trusted not to do so, the CS of the other party is not so well trusted. Trust in the CS also varies between single-CS and dual-CS Call Model. In a single-CS call model both parties use the same service to communicate, however this does not guarantee that they have the same trust models and the same privacy requirements. In a dual-CS call model, the other party's CS merits even less trust, as it may not even be known (depending on user-interface implementations). Hence, variable precautions and privacy negotiations are necessary according to the context and the involved parties. In cases where the CS is untrusted, enforcing authentication by an independent IdP ensures that the exchange of media key is on a third-party path (the identity path) between the authenticated users. Therefore, the CS, who is not in control of this path, cannot mount a MitM attack. However, the CS, not the user, determines whether an IdP is to be used, and users have no means of ensuring that they are protected by the IdP authentication.

6.4. User trust in Identity Providers

The IdP, more than the CS, is the custodian of user intelligence; hence it must have trust relationship with users that subscribe to its services. It is assumed that such services include storing and linking service-bound identities, to allow for flexible means of authentication of related identifiers. Using an IdP makes it easier for users to control privacy, since a single agreement with their chosen IdP is simpler than managing numerous web services, some of which are use only rarely.

Users may have more than one IdP, perhaps different IdPs for special purposes, e.g. most commonly, a separate IdP from an employer. They may choose an IdP that they do not fully trust for private activities that they wish to keep separate. In such cases, the user can limit what personal details are disclosed to the IdP, but the IdP will still know of any authentication request to this identity.

Trusting the IdP of the other party is a more difficult issue, since this IdP may not be even known. The P2P authentication procedure ensures that the IdP origin is not circumvented, but there are no ways of assessing the strength and veracity of the origin statement.

Currently, called users have no way of controlling the downloading of an ’alien’ IdP Proxy (of the other party) to their device, since this is performed automatically by the CS client at the behest of the caller. Hence, both IdP proxies are subject to the same sandbox restrictions, although they have different trust models.

6.5. Communication Service trust in Identity Providers

The CS may rely on a third-party IdP to authenticate users when they log in, and link the given identity with its internal user account. In such cases, the CSP must trust the IdP regarding the authentication strength and the validity of the provided profile information.

In most cases, the CSP provides a set of preferred IdPs for users to choose from, through SSO implementations on the website and usage of the setIdentityProvider function. However, users could also select an IdP, e.g. with OIDC discovery. In dual-CS call model, the CS could receive a SDP message containing an assertion from an unknown

IdP. The verification of the assertion could be performed using the browser's default IdP, with the CS only receiving a confirmation that the identity is authenticated. In these cases, the CS has only low trust level in the IdP, while IdPs that have been vetted by the CS are higly trusted.

6.6. Trusted Identities for non-Browser Interworking

WebRTC browser-based calling services may need to communicate with users on non-browser services, including users of existing SIP servers. The interworking should not be only at the level of signaling and applications, but also at the authentication stage. For a mobile network, as specified in 3GPP TS 23.228, mutual authentication is not possible, but the WebRTC identities, which have been authenticated by an IdP or a CS, are linked to the allocated SIP identities. The 3GPP WebRTC-SIP Client at the device enables it to contact the local network proxy.

7. Use Cases for Privacy Requirements

While [I-D.ietf-rtcweb-use-cases-and-requirements] provides use cases for webRTC media, in this study, use cases demonstrate the calling context scenarios that require different privacy settings, which enhance the examples in [I-D.cazeaux-rtcweb-oauth-identity].

7.1. Anonymous Caller Connecting to Call-Centers

Alice is surfing on websites of several insurance or healthcare companies and wants to discuss matters of some sensitivity. She clicks on links within these websites, in order to talk to their experts. Alice is concerned with her privacy and prefers to remain anonymous, especially towards her employer. Some websites treat her identity as undetectable, since she has not logged in to the service, but they allow such callers to visit. For websites that require authentication, she will use a pseudonym and authenticate to her personal IdP, to avoid her employer's IdP becoming aware of which websites she calls. This use case demonstrates a single-CS call model with the 'Link-to-Call' Service type. The privacy requirements demonstrates undetectability, pseudonymity and unlinkability for the caller. The alternative for Alice is to create unrelated identities for each website, but this is much more laborious. Using an independent IdP with surrogate pseudonyms, Alice can rely on the same credentials, while reassuring the destination websites that she is properly authenticated and is not making nuisance calls.

7.2. Call Center Worker's Privacy

Bob is a member of a company's product support group, working in a customer support center. The company presents clickable icons on the website that connect visitors to the right expert. Bob answers Alice's call, but when Alice calls again, she cannot contact Bob directly, and her call is answered by another group member. This use case represents single-CS Click-to-Link service model from the called destination point of view. Once Alice indicates that she would like to talk, the called destination invokes its own calling service, so the roles are reversed: the destination is the caller and Alice becomes the called party. This enables the destination group members (Bob) to remain unlinkable vis-a-vis Alice. On the other hand, Alice is an undetectable identity, since she has not logged into the service, so her call request uses the browser default IdP to retrieve a surrogate pseudonym identity, but she is still traceable in terms of IP address and path. Hence, this example also shows that unlinkability is not necessarily attached to all other anonymity states, e.g. detectability.

7.3. Online Gaming Calling by Pseudonyms

Alice is playing poker on a gaming website. Alice is a customer, with an account which the gaming business administrator has verified (e.g. via a credit card). Alice wishes to communicate with other players through a voice channel provided by the gaming facility. She is registered on the gaming site under her chosen pseudonym, which is all she wants to reveal to the other players. The calling service verifies the users and their accounts through a server-side IdP validation, so the identities with their pseudonyms are strictly service-bound. This use case demonstrates single-CS /Contact-List call model, where the calls are placed between two registered and logged-on users of the same service. The model is the same as traditional OTT VoIP, where the CS manages the users' identities. Callers and called-parties implicitly rely on the calling service to provide anonymity, where the anonymity set is determined by the number of players.

It should be possible for the gaming CS to permit P2P authentication and independent IdPs, but the gaming host may still require authorization of user accounts (if money changes hands). The IdP authentication benefits the CS, since the IdP's identification is coupled with deeper user verification.

7.4. Hosted Enterprise WebRTC Conferencing Service

Alice is working for a corporation that provides her with a comprehensive web-based communication suite of internal and external conferencing, which is hosted by a Service Provider. The conferencing Service Provider uses the mandatory corporate IdP to authenticate the employees. Alice calls Bob, who works for a partnering company, and is logged on to his own company’s CS. Alice uses Bob’s identity from his own CS, which is recorded in her contact list. Although Alice’s company does not insist on confidentiality, Bob’s company does, so Bob’s calling service demands that all the conferencing participants use the security level that matches Bob’s. This use case demonstrates dual-CS negotiated call model, where the parties have their own preferences determined by their respective organizations. The service providers need to agree on a common set of privacy rules. Although an IdP is mandated by Alice’s organization, external calling parties may not have the same IdP, hence external callers should be authenticated in the mutual peer-to-peer authentication process.

7.5. Variable Trust modes for Employee’s Calls

Employees can have different requirements of privacy depending on type of calls and types of destinations. Alice is a Sales representative calling Bob (a potential customer), to conduct a consumer survey and she wishes to remain untraceable (unlinkable) and unrecognised (pseudonymous). However, when she calls her colleague, Charlie, to discuss invoices, she would like Charlie to call her back. Thus, Alice needs unlinkability when calling Bob, but full linkability when calling Charlie. This use case shows privacy decision by context, according to destination type. Rules may be defined per class of destinations (e.g. internal-colleague, external-corporate, or external-personal). The privacy rules may be executed by the IdP, but other rules may be executed by the CS, hence discrepancies may occur, for example, when the destination-based privacy rule conflicts with corporate policies for this customer. Hence, this example also shows the need for privacy policy negotiation and reconciliation.

7.6. Employee using untrusted WebRTC service

Alice is an employee making use of a WebRTC service that she considers to be untrusted, in order to communicate some important messages to Bob, while she is out of the office. Alice registers to an untrusted CS with her corporate identity. Bob can ’discover’ Alice’s address when he logs into the same untrusted service, because her corporate identity was linked with the CS service-bound identity, when she registered. This use case is an example of single-CS call model, using an untrusted calling service in combination with a trusted IdP. Mutual peer authentication can take place, with each party authenticating the CS based identity via surrogate or explicit corporate identities. Alice wants to be sure that her trusted corporate IdP is used, in order to minimize risks of an MitM attack by the CS, and ensure that the media flow is confidential. Currently, the decision to use an IdP, or a particular user-chosen IdP, is in the hands of the CS, so the possible attacker is responsible for setting the protection! Hence, it is very important that users gain the power to proactively protect their communication by opting to use IdP authentication.

7.7. WebRTC service Interworking with SIP users

Alice uses a social media website to connect to her friends, and her contact list includes people with mobile numbers only. Alice can initiate and receive web calls from her mobile, to connect with mobile phones. Users receiving calls from Alice will see Alice’s phone number displayed, or her IdP identity. Alice can also call Bob’s mobile number from her laptop using a social media service. To connect to Bob, Alice is authenticated by her social media service provider (her CS), who also provides her with a SIP identity that is linked to her other identities. Her SIP identity [RFC3261] is authenticated by the mobile service provider, who had provided a pool of SIP identities to the social media calling service. This use case demonstrates dual-CS Call Model with the interworking service type, using server-side authentication.

8. Requirements Summary

This section lists the new requirements, as discussed above. These requirements call for greater user’s autonomy, greater transparency, and greater variety of trust models that affect the level of divulged information.

8.1. Anonymity

1. It should be possible to set different anonymity rules by standard service types, call models and destination categories.

2. Personal information must not leak via identity assertions.

3. IdP should facilitate pseudonimity via surrogate linked identities.

8.2. Unlinkability

1. It should be possible to set different unlinkability rules by standard service types, call models and destination categories.

2. Callers should be able to request and enforce unlinkability with respect of a called party, separately from other anonymity states.

3. Called destinations should be able to refuse unlinkability requests (e.g. to avoid nuisance calls), while respecting pseudonimity.

4. Unlinkability (e.g. via "origin" in the SDP) should be subject to pre-defined policy, whether that policy is CS-based or IdP-based. Currently, such policies are not transparent to users.

5. Non-disclosure of organization domains is a type of unlinkability, as well as anonymity.

8.3. Independent IdP

1. Users should be able to choose an IdP independently from any calling service, though some services will still mandate or restrict the choice of IdP. In particular, authentication by a trusted IdP must be an option for users who activate untrusted services.

2. IdPs must not lock-in users through non-portable identifiers.

3. Users should be able to create and link surrogate identities and pseudonyms to a globally unique identifier that is portable between IdPs.

4. Users should be able to associate service-bound identities with their independent identity (albeit with distinctive assertion tokens), thus achieving Single-Sign-On.

5. Browsers should allow users to set their chosen default IdPs and log in on browser start-up. Currently browsers select their own factory-set IdPs.

6. User-chosen IdPs should be able to prompt users to log in. Currently, IdP proxies cannot open a dialogue with the user.

7. Users should be able to set IdP-based privacy rules for untrusted CS.

8.4. User Information Confidentiality

1. Linked identifiers and supporting personal verification data must be subject to users’ privacy preference.

2. Session setup messages, as well as identity assertions, should be protected to prevent tampering and eavesdropping.

3. Users’ affiliations with organizations should be subject to privacy preferences. Currently, corporate requirements are not addressed.

8.5. Calling Services

1. Users should be able to set privacy rules for untrusted CS or destinations websites acting as calling services, regardless of the service own parameters.

2. CS should be able to determine privacy parameters per organizations, for data confidentiality and anonymity.

3. Despite users’ choice of IdP, calling services should not be precluded from mandating their choice of IdPs, or offering a preferred IdP list.

4. There must be a transparent method of resolving conflicting privacy requirements arising from the respective CS options.

5. The original website that redirects to a calling service should not be named as ’origin’, if users wish to avoid divulging it.

6. To distinguish between withheld identity and undetected identity, the "origin" field should only provide a status indicator.

8.6. Usability and Notification

1. Users should be informed how their privacy will be handled by the calling service, and which identity and IdP are used.

2. It should be possible to request unlinkability and pseudonimity for a shared group identity, but allow members to maintain separate identities with personal privacy preferences.

3. The IdP must notify users (or CS clients) if it is not possible to support undetectable, pseudonymous or unlinkable calling. Currently, there are no IdP notifications to the user.

4. Users should be notified if a default IdP is assigned, and if other than their chosen IdP is assigned.

9. Conclusions

The web calling paradigm is transformed by browser-based calling facilities that are easily added to shop windows websites. However, this encourages opportunistic calling with increased risks from untrusted parties. The spread of such calling services means users have to maintain numerous identities/passwords, while established services lock-in users to the service-bound identity. Users need to manage their own structured identities, independently of any service. They also need to control their privacy preferences, and vary such preferences for high-risk connections. A solution is needed not only to allow users to control their privacy requirements according to web-calling context, but also to protect destination websites from abuse of anonymity.

10. Contributors

This document is the result of a very fruitful work within the reThink project, with also the following contributors:

I. Tariq Javed
Institut Mines Telecom-Telecom Sud Paris
9 rue C.Fourier
Evry 91011
France
Email: ibrahim_tariq.javed@telecom-sudparis.eu

N. Crespi
Institut Mines Telecom-Telecom Sud Paris
9 rue C.Fourier
Evry 91011
France
Email: noel.crespi@mines-telecom.fr

A. Bouabdallah
Institut Mines Telecom-Telecom Bretagne
2 rue de la Chataigneraie
Cesson Sevigne 35576
France
Email: ahmed.bouabdallah@telecom-bretagne.eu

S. Becot
Orange Labs
4, rue du Clos Courtel
Cesson Sevigne 35510

France
Email: simon.becot@orange.com

J.-M. Crom
Orange Labs
4, rue du Clos Courtel
Cesson Sevigne 35510
France
Email: jeanmichel.crom@orange.com

11. Acknowledgements

This document was inspired by a draft authored by Victoria Beltran, Orange (Email: vicbelma@gmail.com); Emmanuel Bertin, Orange (Email: emmanuel.bertin@orange.com); Stephane Cazeaux, Orange (Email: stephane.cazeaux@orange.com).

Thank you to Ricardo Chaves (INESC-ID) who provided us with his interesting feedback and comments.

This work has received funding from the European Union’s Horizon 2020 research and innovation program under grant agreement No 645342, project reTHINK.

12. References

12.1. Normative references

[I-D.ietf-rtcweb-overview]
Alvestrand, H., "Overview: Real Time Protocols for Browser-based Applications", draft-ietf-rtcweb-overview-15 (work in progress), January 2016.

[I-D.ietf-rtcweb-security-arch]
Rescorla, E., "WebRTC Security Architecture", draft-ietf-rtcweb-security-arch-12 (work in progress), June 2016.

[I-D.ietf-rtcweb-use-cases-and-requirements]
Holmberg, C., Hakansson, S., and G. Eriksson, "Web Real-Time Communication Use-cases and Requirements", draft-ietf-rtcweb-use-cases-and-requirements-16 (work in progress), January 2015.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <http://www.rfc-editor.org/info/rfc2119>.

[RFC3261] Rosenberg, J., Schulzrinne, H., Camarillo, G., Johnston, A., Peterson, J., Sparks, R., Handley, M., and E. Schooler, "SIP: Session Initiation Protocol", RFC 3261, DOI 10.17487/RFC3261, June 2002, <http://www.rfc-editor.org/info/rfc3261>.

[RFC4122] Leach, P., Mealling, M., and R. Salz, "A Universally Unique IDentifier (UUID) URN Namespace", RFC 4122, DOI 10.17487/RFC4122, July 2005, <http://www.rfc-editor.org/info/rfc4122>.

[RFC4949] Shirey, R., "Internet Security Glossary, Version 2", FYI 36, RFC 4949, DOI 10.17487/RFC4949, August 2007, <http://www.rfc-editor.org/info/rfc4949>.

[RFC6749] Hardt, D., Ed., "The OAuth 2.0 Authorization Framework", RFC 6749, DOI 10.17487/RFC6749, October 2012, <http://www.rfc-editor.org/info/rfc6749>.

[RFC6973] Cooper, A., Tschofenig, H., Aboba, B., Peterson, J., Morris, J., Hansen, M., and R. Smith, "Privacy Considerations for Internet Protocols", RFC 6973, DOI 10.17487/RFC6973, July 2013, <http://www.rfc-editor.org/info/rfc6973>.

12.2. Informative references

[I-D.cazeaux-rtcweb-oauth-identity]
Beltran, V., Bertin, E., and S. Cazeaux, "Additional Use-cases and Requirements for WebRTC Identity Architecture", draft-cazeaux-rtcweb-oauth-identity-00 (work in progress), March 2015.

[I-D.ietf-rtcweb-security]
Rescorla, E., "Security Considerations for WebRTC", draft-ietf-rtcweb-security-08 (work in progress), February 2015.

[JeffSayreHenryStory]
Sayre, J. and H. Story, "The WebID Protocol & Browsers", May 2011, <https://www.w3.org/2011/identity-ws/papers/idbrowser2011_submission_22/webid.html>.

[OIDC] "OpenID Connect Core 1.0 incorporating errata set 1", <https://openid.net/specs/openid-connect-registration-1_0.html>.

[SurrogateKeys]
Carter, B., "Intelligent Versus Surrogate Keys", 1997, <http://www.bcarter.com/intsurr1.htm>.

[TerminologyForPrivacy]
Pfitzmann, A., Hansen, M., and H. Tschofenig , "A terminology for privacy by data minimization: Anonymity, Unlinkability, Undetectability, Unobservability,Pseudonymity, and Identity Management - V0.34", August 2010.

[WebRTC] Bergkvist, A., Burnett, D., Jennings, C., and A. Narayanan, "WebRTC 1.0: Real-time Communication Between Browsers. World Wide Web Consortium WD WD-webrtc-20120821.", August 2012.

Authors' Addresses

Rebecca Copeland (editor)
Institut Mines Telecom-Telecom Sud Paris
9 rue C.Fourier
Evry 91011
France

Email: rebecca.copeland@coreviewpoint.com

Kevin Corre
Orange Labs
4, rue du Clos Courtel
Cesson Sevigne 35510
France

Email: kevin1.corre@orange.com

Ingo Friese
Deutsche Telekom AG
Winterfeldtstr. 21
Berlin 10781
Germany

Email: Ingo.Friese@telekom.de

Saad El Jaouhari
Institut Mines Telecom-Telecom Bretagne
2 rue de la Chataigneraie
Cesson Sevigne 35576
France

Email: saad.eljaouhari@telecom-bretagne.eu