

oauth  
Internet-Draft  
Intended status: Best Current Practice  
Expires: 23 January 2026

P. Kasselmann  
SPIRL  
D. Fett  
Authlete  
F. Skokan  
Okta  
22 July 2025

Cross-Device Flows: Security Best Current Practice  
draft-ietf-oauth-cross-device-security-11

## Abstract

This document describes threats against cross-device flows along with practical mitigations, protocol selection guidance, and a summary of formal analysis results identified as relevant to the security of cross-device flows. It serves as a security guide to system designers, architects, product managers, security specialists, fraud analysts and engineers implementing cross-device flows.

## Discussion Venues

This note is to be removed before publishing as an RFC.

Discussion of this document takes place on the Web Authorization Protocol Working Group mailing list ([oauth@ietf.org](mailto:oauth@ietf.org)), which is archived at <https://mailarchive.ietf.org/arch/browse/oauth/>.

Source for this draft and an issue tracker can be found at <https://github.com/oauth-wg/oauth-cross-device-security>.

## Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <https://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on 23 January 2026.

## Copyright Notice

Copyright (c) 2025 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Revised BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Revised BSD License.

## Table of Contents

1. Introduction . . . . .	4
1.1. Cross-Device Authorization . . . . .	5
1.2. Cross-Device Session Transfer . . . . .	6
1.3. Defending Against Cross-Device Attacks . . . . .	6
1.4. Conventions and Terminology . . . . .	7
2. Best Practices . . . . .	7
3. Cross-Device Flow Patterns . . . . .	8
3.1. Cross-Device Authorization . . . . .	8
3.1.1. User-Transferred Session Data Pattern . . . . .	9
3.1.2. Backchannel-Transferred Session Pattern . . . . .	10
3.1.3. User-Transferred Authorization Data Pattern . . . . .	12
3.2. Cross-Device Session Transfer . . . . .	13
3.2.1. Cross-Device Session Transfer Pattern . . . . .	13
3.3. Examples of Cross-Device Flows . . . . .	15
3.3.1. Example A1: Authorize Access to a Video Streaming Service (User-Transferred Session Data Pattern) . . . . .	15
3.3.2. Example A2: Authorize Access to Productivity Services (User-Transferred Session Data Pattern) . . . . .	15
3.3.3. Example A3: Authorize Use of a Bike Sharing Scheme (User-Transferred Session Data Pattern) . . . . .	15
3.3.4. Example A4: Authorize a Financial Transaction (Backchannel-Transferred Session Pattern) . . . . .	15
3.3.5. Example A5: Add a Device to a Network (Cross-Device Session Transfer Pattern) . . . . .	16
3.3.6. Example A6: Remote Onboarding (User-Transferred Session Data Pattern) . . . . .	16
3.3.7. Example A7: Application Bootstrap (Cross-Device Session Transfer Pattern) . . . . .	16
3.3.8. Example A8: Access a Productivity Application (User-Transferred Authorization Data Pattern) . . . . .	17
3.3.9. Example A9: Administer a System (Backchannel-Transferred Session Pattern) . . . . .	17

4.	Cross-Device Flow Exploits . . . . .	17
4.1.	Cross-Device Authorization Flow Exploits . . . . .	17
4.1.1.	User-Transferred Session Data Pattern Exploits . . . . .	18
4.1.2.	Backchannel-Transferred Session Pattern Exploits . . . . .	20
4.1.3.	User-Transferred Authorization Data Pattern Exploits . . . . .	22
4.2.	Cross-Device Session Transfer Exploits . . . . .	24
4.3.	Examples of Cross-Device Flow Exploits . . . . .	26
4.3.1.	Example B1: Illicit Access to a Video Streaming Service (User-Transferred Session Data Pattern) . . . . .	26
4.3.2.	Example B2: Illicit Access to Productivity Services (User-Transferred Session Data Pattern) . . . . .	27
4.3.3.	Example B3: Illicit Access to Physical Assets (User-Transferred Session Data Pattern) . . . . .	27
4.3.4.	Example B4.1: Illicit Transaction Authorization (Backchannel-Transferred Session Pattern) . . . . .	28
4.3.5.	Example B4.2: Fake Helpdesk (Backchannel-Transferred Session Pattern) . . . . .	28
4.3.6.	Example B5: Illicit Network Join (Cross-Device Session Transfer Pattern) . . . . .	28
4.3.7.	Example B6: Illicit Onboarding (User-Transferred Session Data Pattern) . . . . .	29
4.3.8.	Example B7: Illicit Application Bootstrap (Cross-Device Session Transfer Pattern) . . . . .	29
4.3.9.	Example B8: Account Takeover (User-Transferred Session Data Pattern) . . . . .	29
4.3.10.	Example B9: Illicit Access to Administration Capabilities Through Consent Request Overload (Backchannel-Transferred Session Pattern) . . . . .	30
4.3.11.	Out of Scope . . . . .	30
5.	Cross-Device Protocols and Standards . . . . .	30
6.	Mitigating Against Cross-Device Flow Attacks . . . . .	32
6.1.	Practical Mitigations . . . . .	32
6.1.1.	Establish Proximity . . . . .	33
6.1.2.	Short Lived/Timebound QR or User Codes . . . . .	35
6.1.3.	One-Time or Limited Use Codes . . . . .	35
6.1.4.	Unique Codes . . . . .	36
6.1.5.	Content Filtering . . . . .	36
6.1.6.	Detect and Remediate . . . . .	37
6.1.7.	Trusted Devices . . . . .	37
6.1.8.	Trusted Networks . . . . .	38
6.1.9.	Limited Scopes . . . . .	39
6.1.10.	Short Lived Tokens . . . . .	39
6.1.11.	Rate Limits . . . . .	39
6.1.12.	Sender-Constrained Tokens . . . . .	40
6.1.13.	User Education . . . . .	40
6.1.14.	User Experience . . . . .	41
6.1.15.	Authenticate-then-Initiate . . . . .	42

6.1.16. Request Initiation Verification . . . . .	42
6.1.17. Request Binding with Out-of-Band Data . . . . .	43
6.1.18. Practical Mitigation Summary . . . . .	43
6.2. Protocol Selection . . . . .	45
6.2.1. IETF OAuth 2.0 Device Authorization Grant RFC8628: .	45
6.2.2. OpenID Foundation Client Initiated Back-Channel Authentication (CIBA): . . . . .	46
6.2.3. FIDO2/WebAuthn . . . . .	47
6.2.4. Protocol Selection Summary . . . . .	48
6.3. Foundational Pillars . . . . .	49
7. Conclusion . . . . .	51
8. Contributors . . . . .	51
9. References . . . . .	51
9.1. Normative References . . . . .	51
9.2. Informative References . . . . .	53
Appendix A. Document History . . . . .	55
Authors' Addresses . . . . .	58

## 1. Introduction

Protocol flows that span multiple end-user devices are in widespread use today. These flows are often referred to as cross-device flows. A common example is a user that uses their mobile phone to scan a QR code from their smart TV, giving an app on the TV access to their video streaming service. Besides QR codes, other mechanisms are often used such as PIN codes that the user has to enter on one of the devices, or push notifications to a mobile app that the user has to approve.

In all cases, it is up to the user to decide whether to grant authorization or not. However, the QR code or PIN are transferred via an unauthorized channel, leaving it up to the user to decide in which context an authorization is requested. This may be exploited by attackers to gain unauthorized access to a user's resources.

To accommodate the various nuances of cross-device flows, this document distinguished between cases where the cross-device flow is used to authorize access to a resource (cross-device authorization flows) and cases where the cross-device flow is used to transfer an existing session (cross-device session transfer flows).

### 1.1. Cross-Device Authorization

Cross-device authorization flows enable a user to initiate an authorization flow on one device (the Consumption Device) and then use a second, personally trusted, device (Authorization Device) to authorize the Consumption Device to access a resource (e.g., access to a service). The Device Authorization Grant [RFC8628] and Client-Initiated Backchannel Authentication [CIBA] are two examples of popular cross-device authorization flows.

In these flows, the Consumption Device and the Authorization Device are not directly connected and there is no technical mechanisms for the Authorization Device and Consumption Device to establish mutual authentication. It is left to the user to decide whether the source of the authorization request (the Consumption Device) should be trusted before they scan a QR code, enter a user code, or accept an authorization request pushed to their Authorization Device. The transfer of the authorization request and context between the Consumption Device and Authorization device is done over an unauthenticated channel. The only mitigation against this unauthenticated channel is the user's judgement.

Cross-Device Consent Phishing (CDCP) attacks exploit the unauthenticated channel between the Consumption Device and Authorization Device using social engineering techniques to gain unauthorized access to the user's data. Several publications have emerged in the public domain ([Exploit1], [Exploit2], [Exploit3], [Exploit4], [Exploit5], [Exploit6]), describing how the unauthenticated channel can be exploited using social engineering techniques borrowed from phishing. Unlike traditional phishing attacks, these attacks don't harvest credentials. Instead, they skip the step of collecting credentials by persuading users to grant authorization using their Authorization Devices.

Once the user grants authorization, the attacker has access to the user's resources and in some cases is able to collect access and refresh tokens. Once in possession of the access and refresh tokens, the attacker may use these tokens to execute lateral attacks and gain additional access, or monetize the tokens by selling them. These attacks are effective even when multi-factor authentication is deployed, since the attacker's aim is not to capture and replay the credentials, but rather to persuade the user to grant authorization.

## 1.2. Cross-Device Session Transfer

Session transfer flows enable a user to transfer access to a service or network from a device on which the user is already authenticated to a second device such as a mobile phone. In these flows, the user is authenticated and then authorizes the session transfer on one device, referred to as the Authorization Device (e.g., a personal computer, web portal or application), and transfers the session to the device where they will continue to consume the session, referred to as the Consumption Device (e.g., a mobile phone or portable device).

The session may be transferred by showing the user a session transfer code on the Authorization Device, which is then entered on the Consumption Device. This flow may be streamlined by rendering the session transfer code as a QR code on the Authorization Device and scanned by the Consumption Device.

The session transfer preserves state information, including authentication state, at the second device to avoid additional configuration and optimize the user experience. These flows are often used to add new devices to a network, onboard customers to a mobile application, or provision new credentials (e.g., [OpenID.SIOPV2]).

In these cross-device session transfer flows, the channel between the Authorization Device and the Consumption Device is unauthenticated.

Cross-Device Session Phishing (CDSP) attacks exploit the unauthenticated channel between the Authorization Device and Consumption Device by using social engineering techniques to convince the user to send the session transfer code to the attacker. These attacks borrow techniques from traditional phishing attacks, but instead of collecting passwords, attackers collect session transfer codes and other artefacts that allow them to setup a session and then use it to access a user's data.

## 1.3. Defending Against Cross-Device Attacks

This document provides guidance to implementers to defend against Cross-Device Consent Phishing and Cross-Device Session Phishing attacks. This guidance includes:

1. Practical mitigations for susceptible protocols (Section 6.1).
2. Protocol selection guidance to avoid using susceptible protocols (Section 6.2).

3. Results from formal analysis of susceptible protocols (Section 6.3).

#### 1.4. Conventions and Terminology

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

This specification uses the terms "access token", "refresh token", "authorization server", "resource server", "authorization endpoint", "authorization request", and "client" defined by The OAuth 2.0 Authorization Framework [RFC6749].

#### 2. Best Practices

This section describes the set of security mechanisms and measures to secure cross-device protocols against Cross-Device Consent Phishing and Cross-Device Session Phishing attacks that the OAuth working group considers best practices at the time of writing.

1. Implementers MUST perform a risk assessment before implementing cross-device flows, weighing the risks from Cross-Device Consent Phishing and Cross-Device Session Phishing attacks against benefits for users.
2. Implementers SHOULD avoid cross-device flows if risks cannot be sufficiently mitigated.
3. Implementers SHOULD follow the guidance provided in Section 6.2 for protocol selection.
4. Implementers MUST implement practical mitigations as listed in Section 6.1 that are appropriate for the use case, architecture, and selected protocols.
5. Implementers SHOULD implement proximity checks as defined in Section 6.1.1 if possible.

These best practices apply to the Device Authorization Grant ([RFC8628]) as well as other cross-device protocols such as the Client Initiated Backchannel Authentication [CIBA], Self-Issued OpenID Provider v2 [OpenID.SIOPV2], OpenID for Verifiable Presentations [OpenID.VP], the Pre-Authorized Code Flow in ([OpenID.VCI]) and other cross-device protocols that rely on the user to authenticate the channel between devices.

Section 3 provides details about susceptible protocols and Section 4 provides attack descriptions. Section 6.1 provides details about the security mechanisms and mitigations, (protocol-selection) provides protocol selection guidance and Section 6.3 provides details from formal analysis of protocols that apply to cross device flows.

### 3. Cross-Device Flow Patterns

Cross-device flows allow a user to start a flow on one device (e.g., a smart TV) and then transfer the session to continue it on a second device (e.g., a mobile phone). The second device may be used to access the service that was running on the first device, or to perform an action such as authenticating or granting authorization before potentially passing control back to the first device.

These flows typically involve using a mobile phone to scan a QR code or enter a user code displayed on the first device (e.g., Smart TV, Kiosk, Personal Computer or other electronic devices.).

#### 3.1. Cross-Device Authorization

In a cross-device authorization flow, a user attempts to access a service on one device, referred to as the Consumption Device, (e.g., a smart TV) and then uses a second device, referred to as the Authorization Device (e.g., a smartphone), to authorize access to a resource (e.g., access to a streaming service) on the Consumption Device.

Cross-device authorization flows have several benefits, including:

- \* Authorization on devices with limited input capabilities: End-users can authorize devices with limited input capabilities to access content (e.g., smart TVs, digital whiteboards, printers or similarly constrained devices).
- \* Secure authentication on shared or public devices: End-users can perform authentication and authorization using a personally trusted device, without risk of disclosing their credentials to a public or shared device.
- \* Ubiquitous multi-factor authentication: Enables a user to use multi-factor authentication, independent of the device on which the service is being accessed (e.g., a kiosk, smart TV or shared Personal Computer).
- \* Convenience of a single, portable, credential store: Users can keep all their credentials in a mobile wallet or mobile phone that they already carry with them.

There are three cross-device flow patterns for transferring the authorization request between the Consumption Device to the Authorization Device.

- \* **\*User-Transferred Session Data Pattern:**\* In the first pattern, the user initiates the authorization process with the authorization server by copying information from the Consumption Device to the Authorization Device, before authorizing an action. By transferring the data from the Consumption Device to the Authorization Device, the user transfers the authorization session. For example the user may read a code displayed on the Consumption Device and enter it on the Authorization Device, or they may scan a QR code displayed on the Consumption Device with the Authorization Device. The Device Authorization Grant ([RFC8628]) is an example of a cross-device flow that follow this pattern.
- \* **\*Backchannel-Transferred Session Pattern:**\* In the second pattern, the OAuth client on the Consumption Device is responsible for transferring the session and initiating authorization on the Authorization Device via a backchannel with the Authorization Server. For example the user may attempt an online purchase on a Consumption Device (e.g., a personal computer) and receive an authorization request on their Authentication Device (e.g., mobile phone). The Client Initiated Backchannel Authentication [CIBA] is an example of a cross-device flow that follow this pattern.
- \* **\*User-Transferred Authorization Data Pattern:**\* In the third pattern, the OAuth client on the Consumption Device triggers the authorization request via a backchannel with the Authorization Server. Authorization data (e.g., a 6 digit authorization code) is displayed on the Authorization Device, which the user transfers to Consumption Device (e.g., by manually entering it). For example the user may attempt to access data in an enterprise application and receive a 6 digit authorization code on their Authentication Device (e.g., mobile phone) that they enter on Consumption Device.

#### 3.1.1. User-Transferred Session Data Pattern

The Device Authorization Grant ([RFC8628]) is an example of a cross-device flow that relies on the user copying information from the Consumption Device to the Authorization Device by either entering data manually or scanning a QR code. The figure below shows a typical example of this flow:

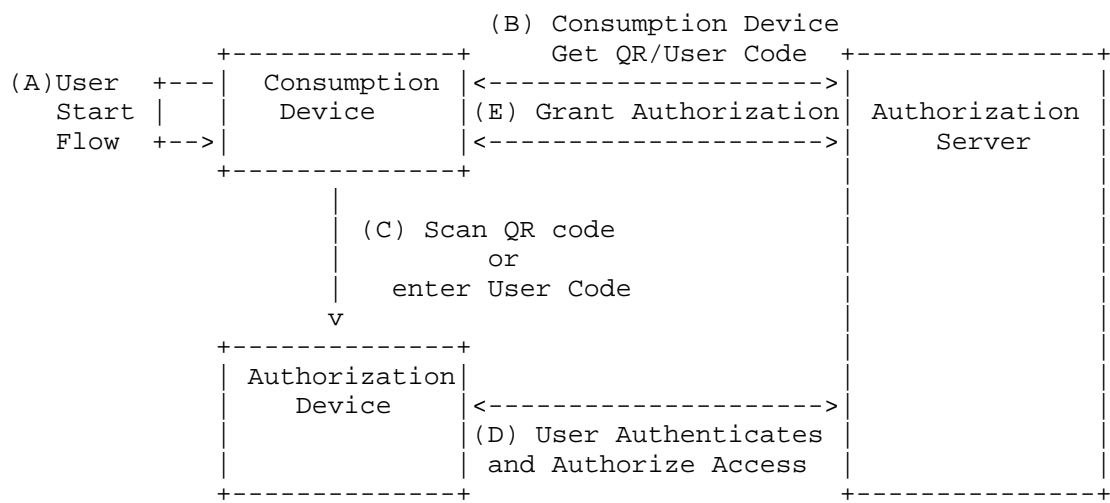


Figure: User-Transferred Session Data Pattern

- \* (A) The user takes an action on the Consumption Device by starting a purchase, adding a device to a network or connecting a service to the Consumption Device.
- \* (B) The Consumption Device retrieves a QR code or user code from an Authorization Server.
- \* (C) The QR code or user code is displayed on the Consumption Device where the user scans the QR code or enters the user code on the Authorization Device.
- \* (D) If the user is unauthenticated, they authenticate to the Authorization Server before granting authorization.
- \* (E) The Authorization Server issues tokens or grants authorization to the Consumption Device to access the user's resources.

3.1.2. Backchannel-Transferred Session Pattern

The Client Initiated Backchannel Authentication [CIBA] transfers the session on the backchannel with the Authorization Server to request authorization on the Authorization Device. The figure below shows an example of this flow.

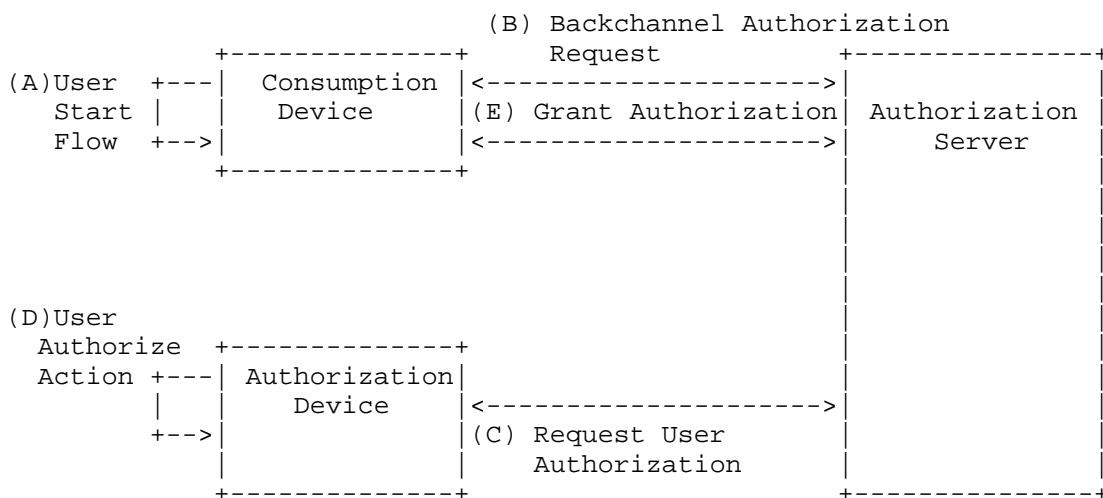


Figure: Backchannel-Transferred Session Pattern

- \* (A) The user takes an action on the Consumption Device by starting a purchase, adding a device to a network or connecting a service to the Consumption Device.
- \* (B) The client on the Consumption Device requests user authorization on the backchannel from the Authorization Server.
- \* (C) The Authorization Server requests the authorization from the user on the user's Authorization Device.
- \* (D) If the user is unauthenticated, they authenticate to the Authorization Server before using their device to grant authorization.
- \* (E) The Authorization Server issues tokens or grants authorization to the Consumption Device to access the user's resources.

The Authorization Server may use a variety of mechanisms to request user authorization, including a push notification to a dedicated app on a mobile phone, or sending a text message with a link to an endpoint where the user can authenticate and authorize an action.

### 3.1.3. User-Transferred Authorization Data Pattern

Examples of the user-transferred authorization data pattern include flows in which the Consumption Device requests the Authorization Server to send authorization data (e.g., a 6 digit authorization code in a text message, e-mail or mobile application) to the Authorization Device. Once the Authorization Device receives the authorization data, the user enters it on the Consumption Device. The Consumption Device sends the authorization data back to the Authorization Server for validation before gaining access to the user's resources. The figure below shows an example of this flow.

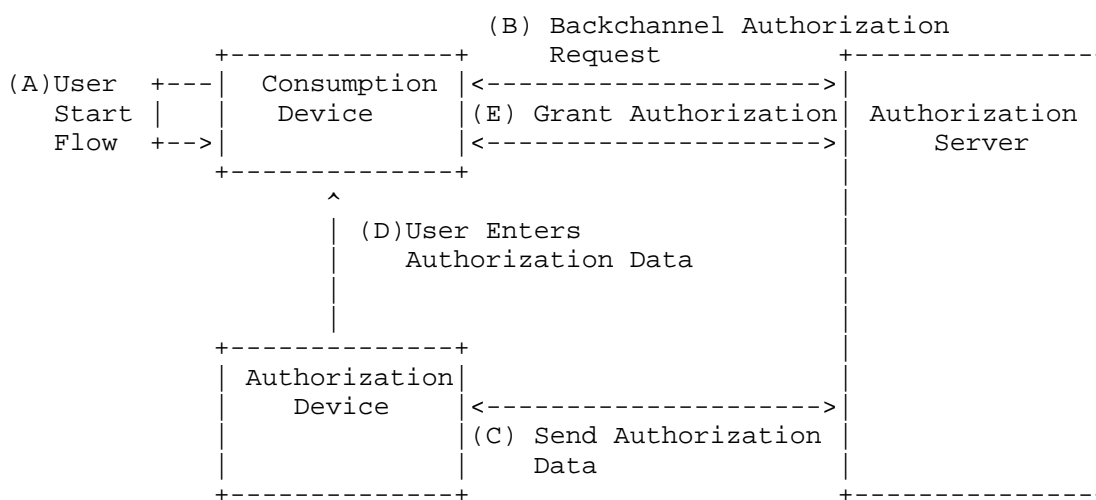


Figure: User-Transferred Authorization Data Pattern

- \* (A) The user takes an action on the Consumption Device by starting a purchase, adding a device to a network or connecting a service to the Consumption Device.
- \* (B) The client on the Consumption Device requests user authorization on the backchannel from the Authorization Server.
- \* (C) The Authorization Server sends authorization data (e.g., a 6 digit authorization code) to the Authorization Device. Examples of mechanisms that may be used to distribute the authorization data include text messages, email or a mobile application.
- \* (D) The user enters the authorization data (e.g., the 6 digit authorization code) on the Consumption Device.

- \* (E) The Authorization Server issues tokens or grants authorization to the Consumption Device to access the user's resources.

The Authorization Server may choose to authenticate the user before sending the authorization data.

### 3.2. Cross-Device Session Transfer

Session transfer flows enable a user to transfer access to a service or network from a device on which the user is already authenticated to a second device such as a mobile phone. In these flows, the user is authenticated and then authorizes the session transfer on one device, referred to as the Authorization Device (e.g., a personal computer, web portal or application), and transfers the session to the device where they will continue to consume the session, referred to as the Consumption Device (e.g., a mobile phone or portable device).

The session transfer preserves state information, including authentication state, at the second device to avoid additional configuration and optimize the user experience. These flows are often used to add new devices to a network, onboard customers to a mobile application, or provision new credentials (e.g., [OpenID.SIOPV2]).

#### 3.2.1. Cross-Device Session Transfer Pattern

In this flow, the user is authenticated and starts the flow by authorizing the transfer of the session on the Authorization Device. The Authorization Device requests a session transfer code that may be rendered as a QR code on the Authorization Device. When the user scans the QR code or enters it on the Consumption Device where they would like the session to continue, the Consumption Device presents it to the Authorization Server. The Authorization Server then transfers the session to the Consumption Device. This may include transferring authentication and authorization state to optimize the user experience. This type of flow is used, for example, for adding new devices to networks, bootstrapping new applications, or provisioning new credentials. The Pre-Authorized Code Flow in ([OpenID.VCI]) is an instance of using this pattern to provision a new credential. The figure below shows a typical flow.

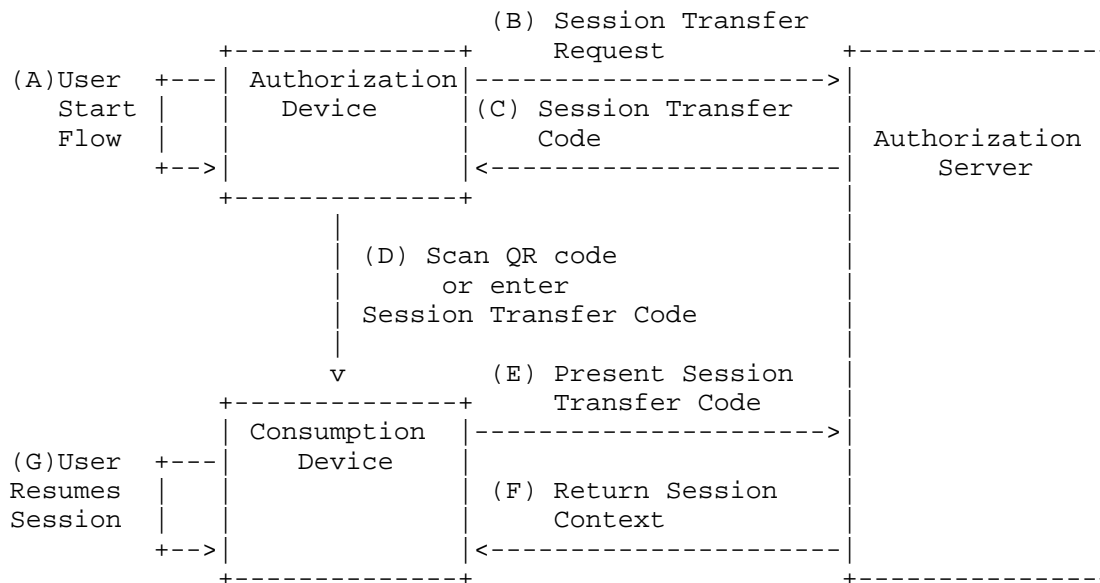


Figure: Cross-Dvice Session Transfer Pattern

- \* (A) The user is authenticated on the Authorization Device and authorizes the transfer of the session to the Consumption device.
- \* (B) The client on the Authorization Device requests a session transfer code from the Authorization Server.
- \* (C) The Authorization Server responds with a session transfer code, which may be rendered as a QR code on the Authorization Device.
- \* (D) The user scans the QR code with the Consumption Device (e.g., their mobile phone), or enters the session transfer code on the target Consumption Device.
- \* (E) The client on the Consumption Device presents the session transfer code to the Authorization Server.
- \* (F) The Authorization Server verifies the session transfer code and retrieves the session context information needed to resume the session on the Consumption Device.
- \* (G) The user resumes the session and is able to access the information on the Consumption Device that they authorized on the Authorization Device.

### 3.3. Examples of Cross-Device Flows

Examples of cross-device flow scenarios include:

#### 3.3.1. Example A1: Authorize Access to a Video Streaming Service (User-Transferred Session Data Pattern)

An end-user sets up a new smart TV and wants to connect it to their favorite streaming service. The streaming service displays a QR code on the TV that the user scans with their mobile phone. The user is redirected to the streaming service provider's web page and asked to enter their credentials to authorize the smart TV to access the streaming service. The user enters their credentials and grants authorization, after which the streaming service is available on the smart TV.

#### 3.3.2. Example A2: Authorize Access to Productivity Services (User-Transferred Session Data Pattern)

An employee wants to access their files on an interactive whiteboard in a conference room. The interactive whiteboard displays a URL and a code. The user enters the URL on their personal computer and is prompted for the code. Once they enter the code, the user is asked to authenticate and authorize the interactive whiteboard to access their files. The user enters their credentials and authorizes the transaction and the interactive whiteboard retrieves their files and allows the user to interact with the content.

#### 3.3.3. Example A3: Authorize Use of a Bike Sharing Scheme (User-Transferred Session Data Pattern)

An end-user wants to rent a bicycle from a bike sharing scheme. The bicycles are locked in bicycle racks on sidewalks throughout a city. To unlock and use a bicycle, the user scans a QR code on the bicycle using their mobile phone. Scanning the QR code redirects the user to the bicycle sharing scheme's authorization page where the user authenticates and authorizes payment for renting the bicycle. Once authorized, the bicycle sharing service unlocks the bicycle, allowing the user to use it to cycle around the city.

#### 3.3.4. Example A4: Authorize a Financial Transaction (Backchannel-Transferred Session Pattern)

An end-user makes an online purchase. Before completing the purchase, they get a notification on their mobile phone, asking them to authorize the transaction. The user opens their app and authenticates to the service before authorizing the transaction.

### 3.3.5. Example A5: Add a Device to a Network (Cross-Device Session Transfer Pattern)

An employee is issued with a personal computer that is already joined to a network. The employee wants to add their mobile phone to the network to allow it to access corporate data and services (e.g., files and e-mail). The employee is logged-in on the personal computer where they initiate the process of adding their mobile phone to the network. The personal computer displays a QR code which authorizes the user to join their mobile phone to the network. The employee scans the QR code with their mobile phone and the mobile phone is joined to the network. The employee can start accessing corporate data and services on their mobile device.

### 3.3.6. Example A6: Remote Onboarding (User-Transferred Session Data Pattern)

A new employee is directed to an onboarding portal to provide additional information to confirm their identity on their first day with their new employer. Before activating the employee's account, the onboarding portal requests that the employee present a government issued ID, proof of a background check and proof of their qualifications. The onboarding portal displays a QR code, which the user scans with their mobile phone. Scanning the QR code invokes the employee's digital wallet on their mobile phone, and the employee is asked to present digital versions of an identity document (e.g., a driving license), proof of a background check by an identity verifier, and proof of their qualifications. The employee authorizes the release of the credentials and after completing the onboarding process, their account is activated.

### 3.3.7. Example A7: Application Bootstrap (Cross-Device Session Transfer Pattern)

An employee is signed into an application on their personal computer and wants to bootstrap the mobile application on their mobile phone. The employee initiates the cross-device flow and is shown a QR code in their application. The employee launches the mobile application on their phone and scans the QR code which results in the user being signed into the application on the mobile phone.

### 3.3.8. Example A8: Access a Productivity Application (User-Transferred Authorization Data Pattern)

A user is accessing a Computer Aid Design (CAD) application. When accessing the application, authorization data in the form of a 6 digit authorization code is sent to the user's mobile phone. The user views the 6 digit authorization code on their phone and enters it in the CAD application, after which the CAD application displays the user's most recent designs.

### 3.3.9. Example A9: Administer a System (Backchannel-Transferred Session Pattern)

A network administrator wants to access an administration portal used to configure network assets and deploy new applications. When attempting to access the service, the network administrator receives a notification in an app on their mobile device, requesting them to confirm access to the portal. The network administrator approves the request on their mobile phone and is granted access to the portal.

## 4. Cross-Device Flow Exploits

Attackers exploit the absence of an authenticated channel between the two devices used in a cross-device flow by using social engineering techniques typically used in phishing attacks.

In cross-device authorization flows the attacker uses these social engineering techniques by changing the context in which the authorization request is presented to convince the user to grant authorization when they shouldn't. These attacks are also known as Cross-Device Consent Phishing (CDCP) attacks.

In cross-device session transfer flows the attacker uses these social engineering techniques to convince the user to initiate a session transfer and send them a session transfer code. Once the attacker is in possession of this session transfer code, they present it to the Authorization Server to transfer the session and access the users resources. These attacks are referred to as Cross-Device Session Phishing (CDSP) attacks.

### 4.1. Cross-Device Authorization Flow Exploits

Attackers exploit cross-device authorization flows by initiating an authorization flow on the Consumption Device and then use social engineering techniques to change the context in which the request is presented to the user in order to convince them to grant authorization on the Authorization Device. The attacker is able to change the context of the authorization request because the channel

between the Consumption Device and the Authorization Device is unauthenticated. These attacks are also known as Cross-Device Consent Phishing (CDCP) attacks.

#### 4.1.1.1. User-Transferred Session Data Pattern Exploits

A common action in cross-device flows is to present the user with a QR code or a user code on the Consumption Device (e.g., smart TV) which is then scanned or entered on the Authorization Device (the mobile phone). When the user scans the code or copies the user code, they do so without any proof that the QR code or user code is being displayed in the place or context intended by the service provider. It is up to the user to decide whether they should trust the QR code or user code. In effect the user is asked to compensate for the absence of an authenticated channel between the Consumption Device (e.g., smart TV) and the Authorization Device (e.g., the mobile phone).

Attackers exploit this absence of an authenticated channel between the two devices by obtaining QR codes or user codes (e.g., by initiating the authorization flows). They then use social engineering techniques to change the context in which authorization is requested to convince end-users to scan the QR code or enter it on their Authorization Device (e.g., mobile phone). Once the end-user performs the authorization on the mobile device, the attacker who initiated the authentication or authorization request obtains access to the users resources. The figure below shows an example of such an attack.

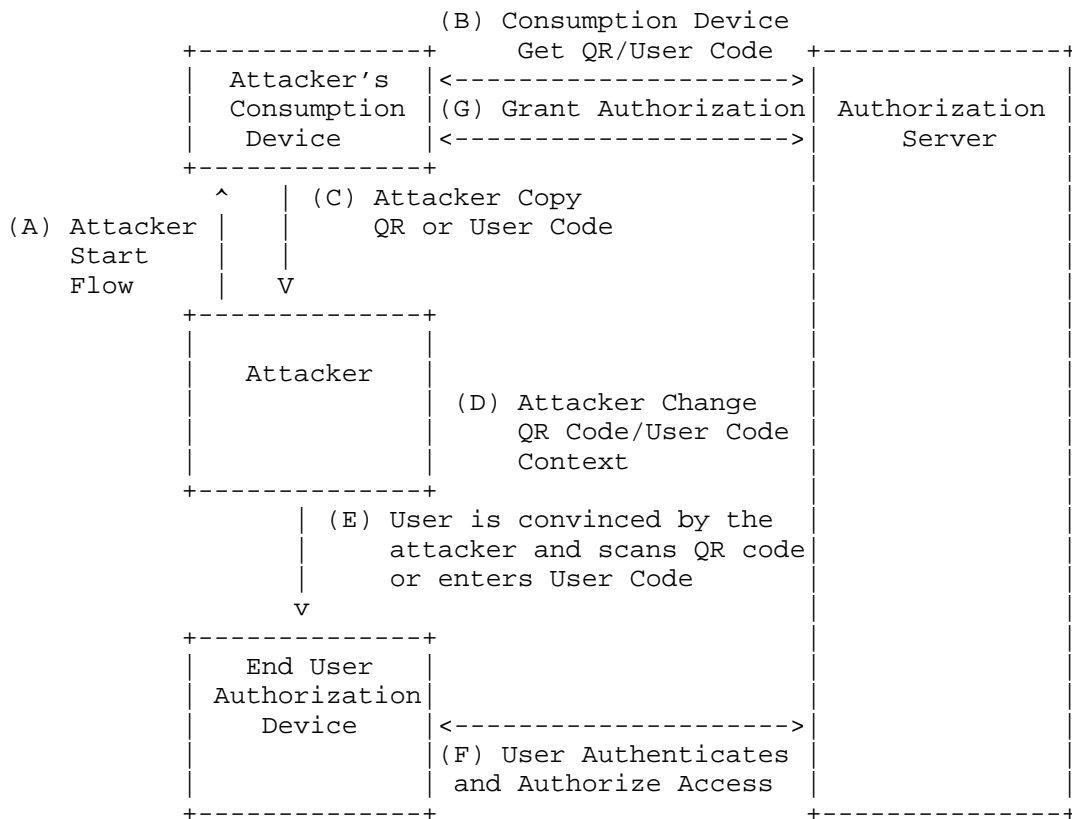


Figure: User-Transferred Session Data Pattern Exploits

- \* (A) The attacker initiates the protocol on the Consumption Device (or mimicks the Consumption Device) by starting a purchase, adding a device to a network or connecting a service to the Consumption Device.
- \* (B) The Consumption Device retrieves a QR code or user code from an Authorization Server.
- \* (C) The attacker copies the QR code or user code.

- \* (D) The attacker changes the context in which the QR code or user code is displayed in such a way that the user is likely to scan the QR code or use the user code when completing the authorization. For example, the attacker could craft an e-mail that includes the user code or QR code and send it to the user. The e-mail might encourage the user to scan the QR code or enter the user code by suggesting that doing so would grant them a reward through a loyalty program or prevent the loss of their data.
- \* (E) The QR code or user code is displayed to the user in a context chosen by the attacker. The user is convinced by the attacker's effort and scans the QR code or enters the user code on the Authorization Device.
- \* (F) The user authenticates to the Authorization Server before granting authorization.
- \* (G) The Authorization Server issues tokens or grants authorization to the Consumption Device, which is under the attacker's control, to access the user's resources. The attacker gains access to the resources and any authorization artifacts (like access and refresh tokens) which may be used in future exploits.

#### 4.1.2. Backchannel-Transferred Session Pattern Exploits

In the backchannel-transferred session pattern, the client requests the authorization server to authenticate the user and obtain authorization for an action. This may happen as a result of user interaction with the Consumption Device, but may also be triggered without the users direct interaction with the Consumption Device, resulting in an authorization request presented to the user without context of why or who triggered the request.

Attackers exploit this lack of context by using social engineering techniques to prime the user for an authorization request and thereby convince them to granting authorization. The social engineering techniques range in sophistication from messages misrepresenting the reason for receiving an authorization request, to triggering a large volume of requests at an inconvenient time for the user, in the hope that the user will grant authorization to make the requests stop. The figure below shows an example of such an attack.

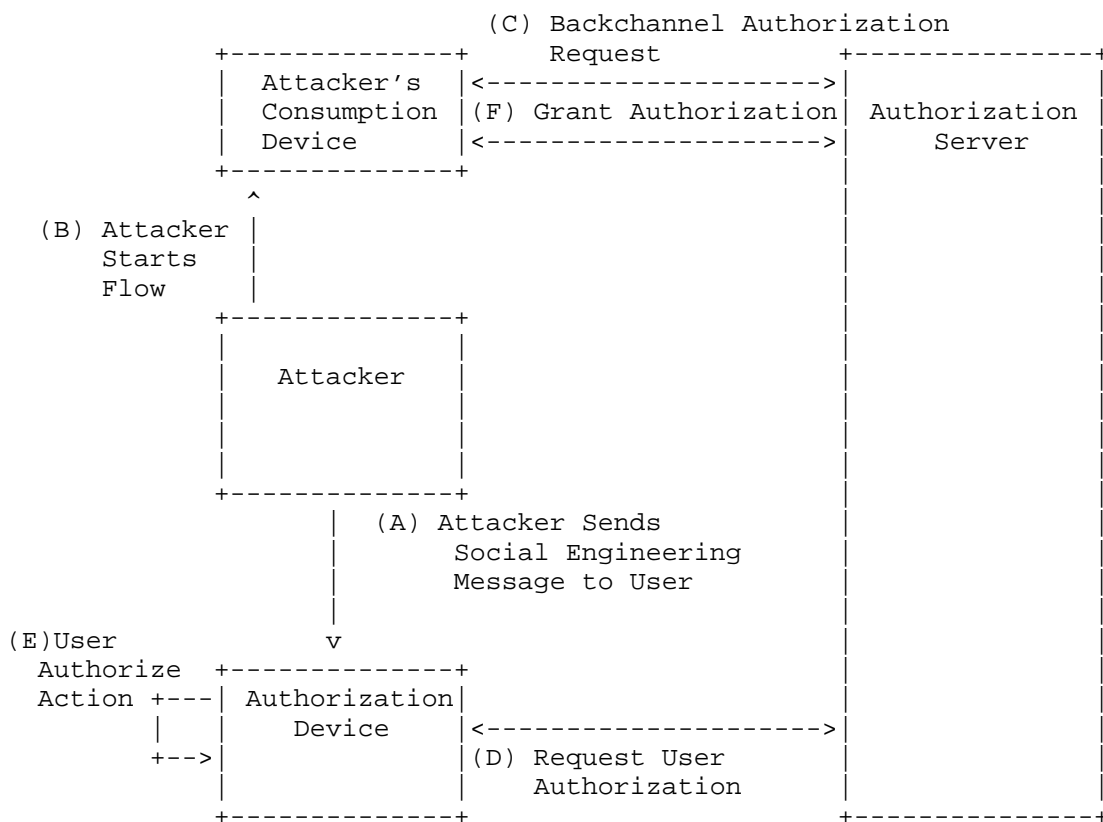


Figure: Backchannel-Transferred Session Pattern Exploits

- \* (A) The attacker sends a social engineering message to prepare the user for the upcoming authorization (optional).
- \* (B) The attacker initiates the protocol on the Consumption Device (or by mimicking the Consumption Device) by starting a purchase, adding a device to a network or accessing a service on the Consumption Device.
- \* (C) The client on the Consumption Device requests user authorization on the backchannel from the Authorization Server.
- \* (D) The Authorization Server requests the authorization from the user on the user's device.
- \* (E) The user authenticates to the authorization server before granting authorization on their device.

- \* (G) The Authorization Server issues tokens or grants authorization to the Consumption Device, which is under the attacker's control. The attacker gains access to the user's resources and possibly any authorization artifacts like access and refresh tokens.

#### 4.1.3. User-Transferred Authorization Data Pattern Exploits

In cross-device flows that follow the user-transferred authorization data pattern, the client on the Consumption Device initiates the authorization request, but the user still has to transfer the authorization data to the Consumption Device. The authorization data may take different forms, including a numerical value such as a 6 digit authorization code. The authorization request may happen as a result of user interaction with the Consumption Device, but may also be triggered without the user's direct interaction with the Consumption Device.

Attackers exploit the user-transferred authorization data pattern by combining the social engineering techniques used to set context for users and convincing users to providing them with authorization data sent to their Authorization Devices (e.g., mobile phones). These attacks are very similar to phishing attacks, except that the attacker also has the ability to trigger the authorization request to be sent to the user directly by the Authorization Server.

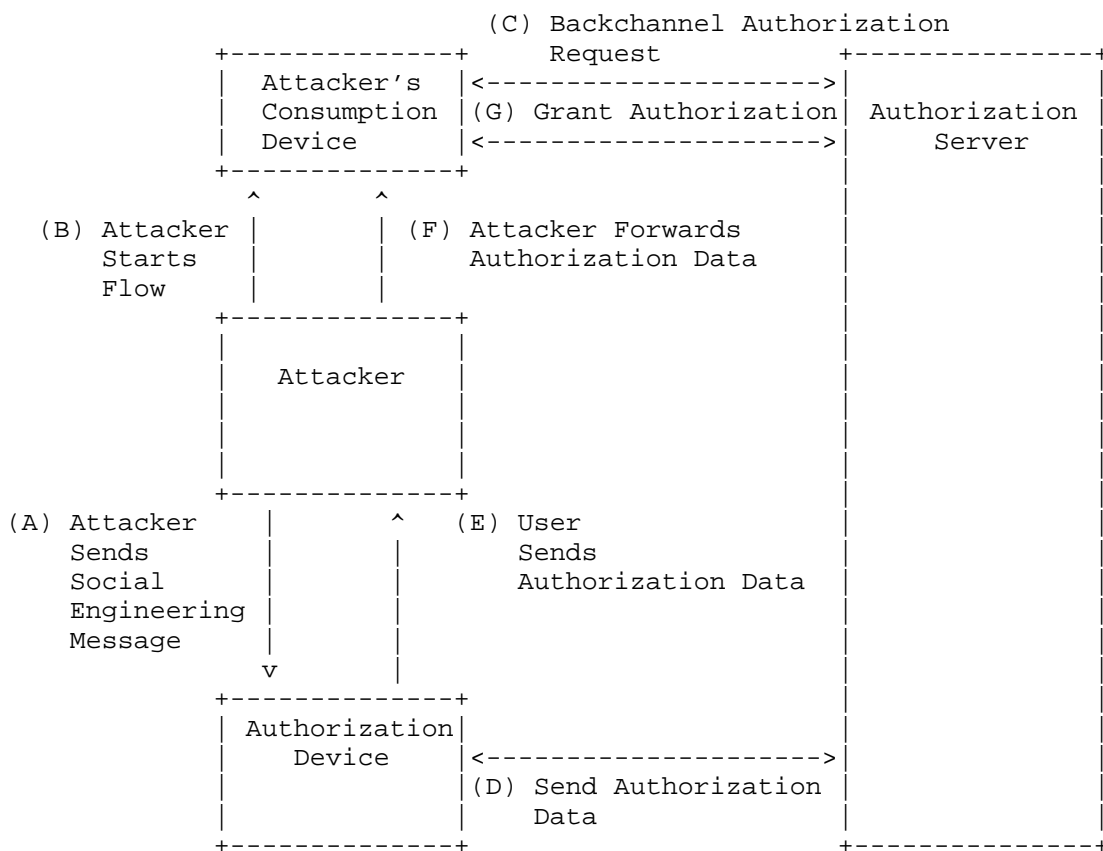


Figure: User-Transferred Authorization Data Pattern Exploits

- \* (A) The attacker sends a social engineering message to prime the user for the authorization request they are about to receive, including instructions on what to do with the authorization data once they receive it.
- \* (B) The attacker initiates the protocol on the Consumption Device (or by mimicking the Consumption Device) by starting a purchase, adding a device to a network or accessing a service on the Consumption Device.
- \* (C) The client on the Consumption Device requests user authorization on the backchannel from the Authorization Server.

- \* (D) The Authorization Server sends authorization data (e.g., a 6 digit authorization code) to the user's Authorization Device (the authorization data may be presented as a QR code, or text message).
- \* (E) The user is convinced by the social engineering message received in step (A) and forwards the authorization data (e.g., a 6 digit authorization code) to the attacker.
- \* (F) The attacker enters the authorization data (e.g., a 6 digit authorization code) on the Consumption Device.
- \* (G) The Authorization Server grants authorization and issues access and refresh tokens to the Consumption Device, which is under the attacker's control. On completion of the exploit, the attacker gains access to the user's resources.

The unauthenticated channel may also be exploited in variations of the above scenario where the user (as opposed to the attacker) initiates the flow and is then convinced using social engineering techniques into sending the authorization data (e.g., a 6 digit authorization code) to the attacker. In these flows, the user is already authenticated and they request authorization data to transfer a session or obtain some other privilege such as joining a device to a network. The authorization data may be represented as a QR code or text string (e.g., 6 digit authorization code). The attacker then proceeds to exploit the unauthenticated channel by using social engineering techniques to convince the user to send the QR code or user code to the attacker. The attacker then use the authorization data to obtain the privileges that would have been assigned to the user.

#### 4.2. Cross-Device Session Transfer Exploits

Attackers exploit cross-device session transfer flows by using social engineering techniques typically used in phishing attacks to convince the user to authorize the transfer of a session and then send the session transfer code or QR code to the attacker. The absence of an authenticated channel between these two devices enables the attacker to use the session transfer code on their own device to obtain access to the session and access the users data. These attacks are referred to as Cross-Device Session Phishing (CDSP) attacks.

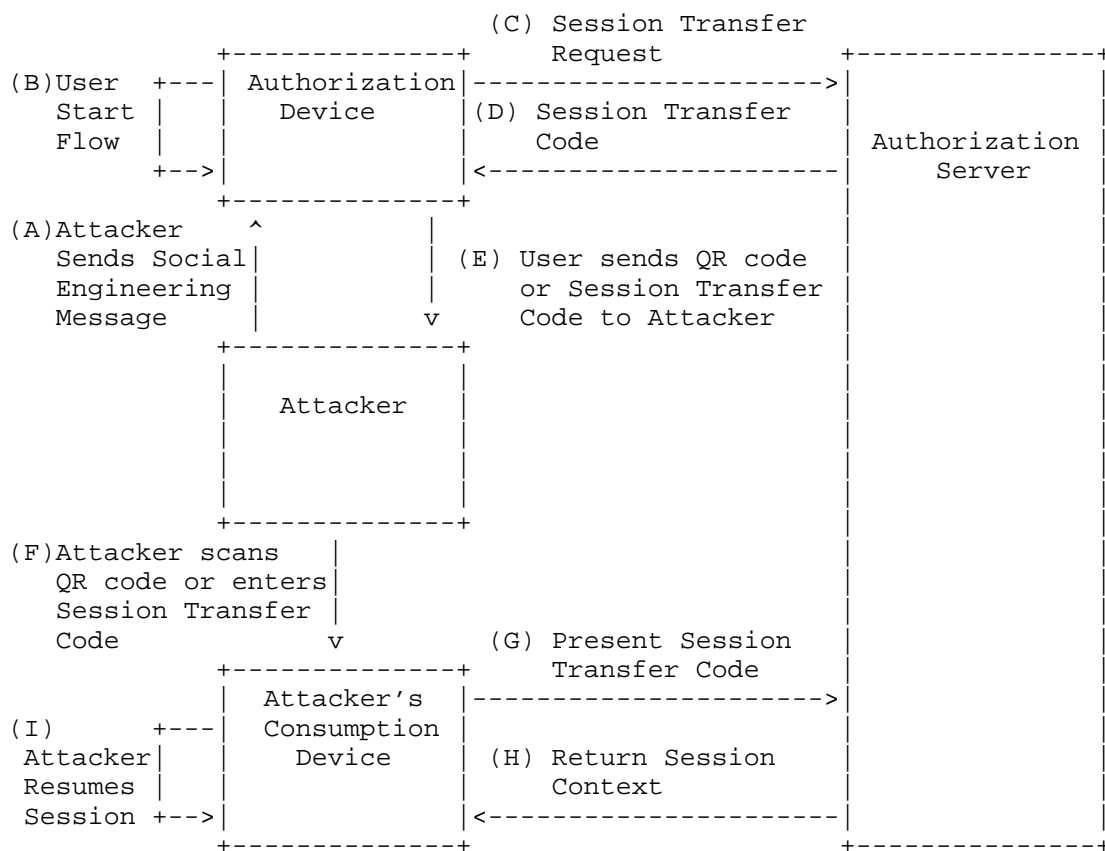


Figure: Cross-Device Session Transfer Pattern Exploit

- \* (A) The attacker sends a social engineering message to that convinces the user that they should authorize a session transfer including instructions on what to do with the QR code or session transfer code once they receive it.
- \* (B) The user is authenticated on their Authorization Device and authorizes the transfer of the session to the Consumption Device.
- \* (C) The client on the Authorization Device requests a session transfer code from the Authorization Server.
- \* (D) The Authorization Server responds with a session transfer code, which may be rendered as a QR code on the Authorization Device.

- \* (E) The user sends the QR code or session transfer code to the attacker, following the instructions they received in step (A).
- \* (F) Once the attacker receives the QR code, they scan it or enter it on their own Consumption Device.
- \* (G) The client on the Consumption Device presents the session transfer code to the Authorization Server.
- \* (H) The Authorization Server verifies the session transfer code and retrieves the session context information needed to resume the session on the Consumption Device.
- \* (I) The attacker resumes the session on their own Consumption Device and is able to access the information that the user authorized on their Authorization Device in step (B).

#### 4.3. Examples of Cross-Device Flow Exploits

The following examples illustrate these attacks in practical settings and show how the unauthenticated channel is exploited by attackers who can copy the QR codes and user codes, change the context in which they are presented using social engineering techniques and mislead end-users into granting consent to avail of services, access data and make payments.

##### 4.3.1. Example B1: Illicit Access to a Video Streaming Service (User-Transferred Session Data Pattern)

An attacker obtains a smart TV and attempts to access an online streaming service. The smart TV obtains a QR code from the streaming service authorization server and displays it on screen. The attacker copies the QR code and embeds it in an e-mail that is sent to a large number of recipients. The e-mail contains a message stating that the streaming service wants to thank them for their loyal support and by scanning the QR code, they will be able to add a bonus device to their account for no charge. One of the recipients open the e-mail and scan the QR code to claim the loyalty reward. The user performs multi-factor authentication, and when asked if they want a new device to be added to their account, they authorize the action. The attacker's device is now authorized to access the content and obtains an access and refresh token. The access token allows the attacker to access content and the refresh token allows the attacker to obtain fresh tokens whenever the access token expires.

The attacker scales up the attack by emulating a new smart TV, obtaining multiple QR codes and widening the audience it sends the QR code to. Whenever a recipient scans the QR code and authorizes the addition of a new device, the attacker obtains an access and refresh token, which they sell for a profit.

#### 4.3.2. Example B2: Illicit Access to Productivity Services (User-Transferred Session Data Pattern)

An attacker emulates an enterprise application (e.g., an interactive whiteboard) and initiates a cross-device flow by requesting a user code and URL from the authorization server. The attacker obtains a list of potential victims and sends an e-mail informing users that their files will be deleted within 24 hours if they don't follow the link, enter the user code and authenticate. The e-mail reminds them that this is the third time that they have been notified and their last opportunity to prevent deletion of their work files. One or more employees respond by following the URL, entering the code and performing multi-factor authentication. Throughout the authentication experience, the user is interacting with a trusted user experience, re-enforcing the legitimacy of the request. Once these employees authorized access, the attacker obtains access and refresh tokens from the authorization server and uses it to access the users' files, perform lateral attacks to obtain access to other information and continuously refresh the session by requesting new access tokens. These tokens may be exfiltrated and sold to third parties.

#### 4.3.3. Example B3: Illicit Access to Physical Assets (User-Transferred Session Data Pattern)

An attacker copies a QR code from a bicycle locked in a bicycle rack in a city, prints it on a label and places the label on a bicycle at the other end of the bicycle rack. A customer approaches the bicycle that contains the replicated QR code and scans the code and authenticates before authorizing payment for renting the bicycle. The bicycle rack unlocks the bicycle containing the original QR code and the attacker removes the bicycle before cycling down the street while the customer is left frustrated that the bicycle they were trying to use is not being unlocked [NYC.Bike]. The customer proceeds to unlock another bicycle and lodges a complaint with the bicycle renting company.

#### 4.3.4. Example B4.1: Illicit Transaction Authorization (Backchannel-Transferred Session Pattern)

An attacker obtains a list of user identifiers for a financial institution and triggers a transaction request for each of the users on the list. The financial institution's authorization server sends push notifications to each of the users, requesting authorization of a transaction. The vast majority of users ignore the request to authorize the transaction, but a small percentage grants authorization by approving the transaction.

#### 4.3.5. Example B4.2: Fake Helpdesk (Backchannel-Transferred Session Pattern)

An attacker obtains the contact information for a user and contacts them, pretending to be a representative of the user's financial institution. The attacker informs the user that there were a number of fraudulent transactions against their account and asks them to review these transactions by approving or rejecting them. The attacker then triggers a sequence of transactions. The user receives an authorization request for each transaction and declines them as they do not recognize them. The attacker then informs the user that they need to close the users account and transfer all the funds to a new account to prevent further fraudulent transactions. The user receives another authorization request which they approve, or provide additional authorization information to the attacker which enables the attacker to complete their attack and defraud the user.

#### 4.3.6. Example B5: Illicit Network Join (Cross-Device Session Transfer Pattern)

An attacker creates a message to all employees of a company, claiming to be from a trusted technology provider investigating a suspected security breach. They ask employees to send them the QR code typically used to join a new device to the network, along with detailed steps on how to obtain the QR code. The employee, eager to assist, initiates the process to add a new mobile device to the network. They authenticate to the network and obtain a QR code. They send the QR code to the attacker. The attacker scans the QR code and adds their own device to the network. They use this device access as an entry point and perform lateral moves to obtain additional privileges and access to restricted resources.

#### 4.3.7. Example B6: Illicit Onboarding (User-Transferred Session Data Pattern)

An attacker initiates an employee onboarding flow and obtains a QR code from the onboarding portal to invoke a digital wallet and present a verifiable credential attesting to a new employee's identity. The attacker obtains a list of potential new employees and sends an e-mail informing them that it is time to present proof of their background check or government issued ID. The new employee scans the QR code, invokes their digital wallet and presents their credentials. Once the credentials are presented, the employee's account is activated. The employee portal accessed by the attacker to obtain the QR code displays a message to the attacker with instructions on how to access their account.

#### 4.3.8. Example B7: Illicit Application Bootstrap (Cross-Device Session Transfer Pattern)

An attacker creates a message to all employees of a company, claiming to be from the company's IT service provider. They claim that they are trying to resolve an application performance issue and ask employees to send them the QR code typically used to transfer a session. The employee, eager to assist, initiates the process to transfer a session. They authenticate and obtain a QR code and then send the QR code to the attacker. The attacker scans the QR code with their mobile phone and access the users data and resources.

#### 4.3.9. Example B8: Account Takeover (User-Transferred Session Data Pattern)

An attacker wants to use some website which requires presentation of a verifiable credential for authentication. The attacker creates a phishing website which will in real time capture log-in QR Codes from the original website and present these to the user. The attacker tries to get the user to use the phishing website by sending messages by e-mail or other messaging technologies. The user scans the QR code on the phishing website, invokes their digital wallet and presents their credentials. Once the credentials are presented, the original session from the attackers device is authorized with the user's credentials.

#### 4.3.10. Example B9: Illicit Access to Administration Capabilities Through Consent Request Overload (Backchannel-Transferred Session Pattern)

An attacker attempts to access an administration portal repeatedly, generating a stream of authorization requests to the network administrator. The attempts are timed to occur while the administrator is asleep. The administrator is woken by the incoming requests on their phone, and, in an attempt to stop the notifications, they accidentally approve access and the attacker gains access to the portal.

#### 4.3.11. Out of Scope

In all of the attack scenarios listed above, a user is misled or exploited. For other attacks, where the user is willingly colluding with the attacker, the threat model, security implications and potential mitigations are very different. For example, a cooperating user can bypass software mitigations on their device, share access to hardware tokens with the attacker, and install additional devices to forward radio signals to circumvent proximity checks.

This document only considers scenarios where a user does not collude with an attacker.

### 5. Cross-Device Protocols and Standards

Cross-device flows that are subject to the attacks described earlier typically share the following characteristics:

1. The attacker can initiate the flow and manipulate the context of an authorization request. E.g., the attacker can obtain a QR code or user code, or can request an authentication/authorization decision from the user.
2. The interaction between the Consumption Device and Authorization Device is unauthenticated. I.e., it is left to the user to decide if the QR code, user code, or authentication request is being presented in a legitimate context.

A number of protocols that have been standardized, or are in the process of being standardized that share these characteristics include:

- \* \*IETF OAuth 2.0 Device Authorization Grant ([RFC8628]):\* A standard to enable authorization on devices with constrained input capabilities (smart TVs, printers, kiosks). In this protocol, the user code or QR code is displayed on the Consumption Device and entered on a second device (e.g., a mobile phone).
- \* \*Open ID Foundation Client Initiated Back-Channel Authentication (CIBA) [CIBA]:\* A standard developed in the OpenID Foundation that allows a device or service (e.g., a personal computer, smart TV, Kiosk) to request the OpenID Provider to initiate an authentication flow if it knows a valid identifier for the user. The user completes the authentication flow using a second device (e.g., a mobile phone). In this flow the user does not scan a QR code or obtain a user code from the Consumption Device, but is instead contacted by the OpenID Provider to complete the authentication using a push notification, e-mail, text message or any other suitable mechanism.
- \* \*OpenID for Verifiable Credential Protocol Suite (Issuance, Presentation):\* The OpenID for Verifiable Credentials enables cross-device scenarios by allowing users to scan QR codes to retrieve credentials (Issuance - see [OpenID.VCI]) or present credentials (Presentation - see [OpenID.VP]). The QR code is presented on a device that initiates the flow.
- \* \*Self-Issued OpenID Provider v2 (SIOP V2):\* A standard that allows end-user to present self-attested or third party attested attributes when used with OpenID for Verifiable Credential protocols. The user scans a QR code presented by the relying party to initiate the flow.

Cross-device protocols SHOULD not be used for same-device scenarios. If the Consumption Device and Authorization Device are the same device, protocols like OpenID Connect Core [OpenID.Core] and OAuth 2.0 Authorization Code Grant as defined in [RFC6749] are more appropriate. If a protocol supports both same-device and cross-device modes (e.g., [OpenID.SIOPV2]), the cross-device mode SHOULD not be used for same-device scenarios. An authorization server MAY choose to block cross-device protocols used in same-device scenarios if it detects that the same device is used. An authorization server may use techniques such as device fingerprinting, network address or other techniques to detect if a cross-device protocol is being used on the same device. If an implementor decides to use a cross-device protocol or a protocol with a cross-device mode in a same-device scenario, the mitigations recommended in this document SHOULD be implemented to reduce the risks that the unauthenticated channel is exploited.

## 6. Mitigating Against Cross-Device Flow Attacks

The unauthenticated channel between the Consumption Device and the Authorization Device allows attackers to change the context in which the authorization request is presented to the user. This shifts responsibility of authenticating the channel between the two devices to the end-user. End-users have "expertise elsewhere", are typically not security experts, and don't understand the protocols and systems they interact with. As a result, end-users are poorly equipped to authenticate the channel between the two devices. Mitigations should focus on:

1. Minimizing reliance on the user to make decisions to authenticate the channel.
2. Providing better information with which to make decisions to authenticate the channel.
3. Recovering from incorrect channel authentication decisions by users.

To achieve the above outcomes, mitigating against Cross-Device Consent Phishing attacks require a three-pronged approach:

1. Reduce risks of deployed protocols with practical mitigations.
2. Adopt or develop protocols that are less susceptible to these attacks where possible.
3. Provide analytical tools to assess vulnerabilities and effectiveness of mitigations.

### 6.1. Practical Mitigations

A number of protocols that enable cross-device flows that are susceptible to Cross-Device Consent Phishing attacks are already deployed. The security profile of these protocols can be improved through practical mitigations that provide defense in depth that either:

1. Prevents the attack from being initiated.
2. Disrupts the attack once it is initiated.
3. Remediate or reduces the impact if the attack succeeds.

It is RECOMMENDED that one or more of the mitigations are applied whenever implementing a cross-device flow. Every mitigation provides an additional layer of security that makes it harder to initiate the attack, disrupts attacks in progress or reduces the impact of a successful attack.

#### 6.1.1. Establish Proximity

The unauthenticated channel between the Consumption Device and Authorization Device allows attackers to obtain a QR code or user code in one location and display it in another location. Consequently, proximity-enforced cross-device flows are more resistant to Cross-Device Consent Phishing attacks than proximity-less cross-device flows. Establishing proximity between the location of the Consumption Device and the Authorization Device limits an attacker's ability to launch attacks by sending the user or QR codes to large numbers of users that are geographically distributed. There are a couple of ways to establish proximity:

- \* **Physical connectivity:** This is a good indicator of proximity, but requires specific ports, cables and hardware and may be challenging from a user experience perspective or may not be possible in certain settings (e.g., when USB ports are blocked or removed for security purposes). Physical connectivity may be better suited to dedicated hardware like FIDO devices that can be used with protocols that are resistant to the exploits described in this document. The authorization server does not directly establish proximity, but relies on the deployed system and selected protocols and devices to establish proximity.
- \* **Wireless proximity:** Near Field Communications (NFC), Bluetooth Low Energy (BLE), and Ultra Wideband (UWB) services can be used to prove proximity between the two devices. NFC technology is widely deployed in mobile phones as part of payment solutions, but NFC readers are less widely deployed. BLE presents another alternative for establishing proximity, but may present user experience challenges when setting up. UWB standards such as IEEE 802.15.4 and the IEEE 802.15.4z-2020 Amendment 1 enable secure ranging between devices and allow devices to establish proximity relative to each other [IEEE802154]. The authorization server does not directly establish proximity, but relies on the deployed system and selected protocols and devices to establish proximity using these wireless proximity protocols.
- \* **Shared network:** Device proximity can be inferred by verifying that both devices are on the same network. This check may be performed by the authorization server by comparing the network addresses of the device where the code is displayed (Consumption Device) with

that of the Authorization Device. Alternatively the check can be performed on the device, provided that the network address is available. This could be achieved if the authorization server encodes the Consumption Device's network address in the QR code and uses a digital signature to prevent tampering with the code. This does require the wallet to be aware of the countermeasure and effectively enforce it. Note that it is common for a Consumption Device (e.g., a TV) to use a Wi-Fi connection while the Authorization Device (e.g., a phone) uses a mobile network. Though physically in close proximity, they don't share a network, so other proximity checks are needed.

- \* Geolocation: Proximity can be established by comparing geolocation information derived from global navigation satellite-system (GNSS) co-ordinates or geolocation lookup of IP addresses and comparing proximity. Geolocation based on GNSS may vary in accuracy depending on the users location, and when mapped to national or regional boundaries may show a Consumption and Authorization Device in different locations if those devices are close to a border. Since relative position is more important than absolute location, implementations should consider relative location to both devices rather than absolute location when determining proximity. Geolocation based on IP addresses may be inaccurate along regional or national borders due to overlapping coverage by different network providers from the respective regions. This may result in the Consumption device being mapped to one region, while the Authorization Device may be on another network from another provider and mapped to another region. These inaccuracies may require restrictions to be at a more granular level (e.g., same city, country, region or continent). Similar to the shared network checks, these checks may be performed by the authorization server or on the users device, provided that the information encoded in a QR code is integrity protected using a digital signature.

Depending on the risk profile and the threat model in which a system is operating, it MAY be necessary to use more than one mechanism to establish proximity to raise the bar for any potential attackers.

Note: There are scenarios that require that an authorization takes place in a different location than the one in which the transaction is authorized. For example, there may be a primary and secondary credit card holder and both can initiate transactions, but only the primary holder can authorize it. There is no guarantee that the primary and secondary holders are in the same location at the time of the authorization. In such cases, proximity (or lack of proximity) may be an indicator of risk and the system may deploy additional controls (e.g., transaction value limits, transaction velocity limits) or use the proximity information as input to a risk management system.

**\*Limitations:** Proximity mechanisms make it harder to perform Cross-Device Consent Phishing (CDCP) attacks. However, depending on how the proximity check is performed, an attacker may be able to circumvent the protection: The attacker can use a VPN to simulate a shared network or spoof a GNSS position. For example, the attacker can try to request the location of the end-user's Authorization Device through browser APIs and then simulate the same location on their Consumption Device using standard debugging features available on many platforms.

#### 6.1.2. Short Lived/Timebound QR or User Codes

The impact of an attack can be reduced by making QR or user codes short lived. If an attacker obtains a short lived code, the duration during which the unauthenticated channel can be exploited is reduced, potentially increasing the cost of a successful attack. This mitigation can be implemented on the authorization server without changes to other system components.

**\*Limitations:** There is a practical limit to how short a user code can be valid due to network latency and user experience limitations (time taken to enter a code, or incorrectly entering a code). More sophisticated Cross-Device Consent Phishing attacks counter the effectiveness of short lived codes by convincing a user to respond to a phishing e-mail and only request the QR or user code once the user clicks on the link in the phishing e-mail [Exploit6].

#### 6.1.3. One-Time or Limited Use Codes

By enforcing one-time use or limited use of user or QR codes, the authorization server can limit the impact of attacks where the same user code or QR code is sent to multiple victims. One-time use may be achieved by including a nonce or date-stamp in the user code or QR code which is validated by the authorization server when the user scans the QR code against a list of previously issued codes. This mitigation can be implemented on the authorization server without

changes to other system components.

**\*Limitations:** Enforcing one-time use may be difficult in large globally distributed systems with low latency requirements, in which case short lived tokens may be more practical. One-time use codes may also have an impact on the user experience. For example, a user may enter a code, but their session may be interrupted before the access request is completed. If the code is a one-time use code, they would need to restart the session and obtain a new code since they won't be allowed to enter the same code a second time. To avoid this, implementers MAY allow the same code to be presented a small number of times.

#### 6.1.4. Unique Codes

By issuing unique user or QR codes, an authorization server can detect if the same codes are being repeatedly submitted. This may be interpreted as anomalous behavior and the authorization server MAY choose to decline issuing access and refresh tokens if it detects the same codes being presented repeatedly. This may be achieved by maintaining a deny list that contains QR codes or user codes that were previously used. The authorization server MAY use a sliding window equal to the lifetime of a token if short lived/timebound tokens are used (see Section 6.1.2). This will limit the size of the deny list. This mitigation can be implemented on the authorization server without changes to other system components.

**\*Limitations:** Maintaining a deny list of previously redeemed codes, even for a sliding window, may have an impact on the latency of globally distributed systems. One alternative is to segment user codes by geography or region and maintain local deny lists.

#### 6.1.5. Content Filtering

Attackers exploit the unauthenticated channel by changing the context of the user code or QR code and then sending a message to a user (e-mail, text messaging, instant messaging or other communication mechanisms). By deploying content filtering (e.g., anti-spam filter), these messages can be blocked and prevented from reaching the end-users. It may be possible to fine-tune content filtering solutions to detect artefacts like QR codes or user codes that are included in a message that is sent to multiple recipients in the expectation that at least one of the recipients will be convinced by the message and grant authorization to access restricted resources.

**\*Limitations:** Some scenarios may require legitimate re-transmission of user, QR and authorization data (e.g., retries). To prevent the disruption of legitimate scenarios, content filters may use a

threshold and allow a limited number of messages with the same QR or user codes to be transmitted before interrupting the delivery of those messages. Content filtering may also be fragmented across multiple communications systems and communication channels (e-mail, text messaging, instant messaging or other communication mechanisms), making it harder to detect or interrupt attacks that are executed over multiple channels, unless there is a high degree of integration between content filtering systems.

#### 6.1.6. Detect and Remediate

The authorization server may be able to detect misuse of the codes due to repeated use as described in Section 6.1.4, as an input from a content filtering engine as described in Section 6.1.5, or through other mechanisms such as reports from end-users. If an authorization server determines that a user code or QR code is being used in an attack it may choose to invalidate all tokens issued in response to these codes and make that information available through a token introspection endpoint (see [RFC7662]). In addition it may notify resource servers to stop accepting these tokens or to terminate existing sessions associated with these tokens using Continuous Access Evaluation Protocol (CAEP) messages [CAEP] using the Shared Signals Framework (SSF) [SSF] framework or an equivalent notification system.

**\*Limitations:** Detection and remediation requires that resource servers are integrated with security eventing systems or token introspection services. This may not always be practical for existing systems and may need to be targeted to the most critical resource services in an environment.

#### 6.1.7. Trusted Devices

If an attacker is unable to initiate the protocol, they are unable to obtain a QR code or user code that can be leveraged for the attacks described in this document. By restricting the protocol to only be executed on devices trusted by the authorization server, it prevents attackers from using arbitrary devices, or by mimicking devices to initiate the protocol.

Authorization Servers MAY use different mechanisms to establish which devices it trusts. This includes limiting cross-device flows to specific device types such as interactive whiteboards or smart TVs, pre-registering devices with the authorization server or only allow cross-device flows on devices managed through device management systems. Device management systems may enforce policies that govern patching, version updates, on-device anti-malware deployment, revocation status and device location amongst others. Trusted devices MAY have their identities rooted in hardware (e.g., a TPM or equivalent technology).

By only allowing trusted devices to initiate cross-device flows, it requires the attacker to have access to such a device and maintain access in a way that does not result in the device's trust status from being revoked.

**\*Limitations:** An attacker may still be able to obtain access to a trusted device and use it to initiate authorization requests, making it necessary to apply additional controls and integrating with other threat detection and management systems that can detect suspicious behaviour such as repeated requests to initiate authorization or high volume of service activation on the same device.

#### 6.1.8. Trusted Networks

An attacker can be prevented from initiating a cross-device flow protocol by only allowing the protocol to be initiated on a trusted network or within a security perimeter (e.g., a corporate network). A trusted network may be defined as a set of IP addresses and joining the network is subject to security controls managed by the network operator, which may include only allowing trusted devices on the network, device management, user authentication and physical access policies and systems. By limiting protocol initiation to a specific network, the attacker needs to have access to a device on the network. This mitigation can be implemented on the authorization server without changes to other system components.

**\*Limitations:** Network level controls may not always be feasible, especially when dealing with consumer scenarios where the network may not be under control of the service provider. Even if it is possible to deploy network level controls, it SHOULD be used in conjunction with other controls outlined in this document to achieve defence in-depth.

#### 6.1.9. Limited Scopes

Authorization servers MAY choose to limit the scopes they include in access tokens issued through cross-device flows where the unauthenticated channel between two devices are susceptible to being exploited. Including limited scopes lessens the impact in case of a successful attack. The decision about which scopes are included may be further refined based on whether the protocol is initiated on a trusted device or the user's location relative to the location of the Consumption Device. This mitigation can be implemented on the authorization server without changes to other system components.

*\*Limitations:* Limiting scopes reduces the impact of a compromise, but does not avoid it. It SHOULD be used in conjunction with other mitigations described in this document.

#### 6.1.10. Short Lived Tokens

Another mitigation strategy includes limiting the life of the access and refresh tokens. The lifetime can be lengthened or shortened, depending on the user's location, the resources they are trying to access or whether they are using a trusted device. Short lived tokens do not prevent or disrupt the attack, but serve as a remedial mechanism in case the attack succeeded. This mitigation can be implemented on the authorization server without changes to other system components.

*\*Limitations:* Short lived tokens reduces the time window during which an attacker can benefit from a successful attack. This is most effective for access tokens. However, once an attacker obtains a refresh token, they can continue to request new access tokens, as well as refresh tokens. Forcing the expiry of refresh tokens may cause the user to re-authorize an action more frequently, which results in a negative user experience.

#### 6.1.11. Rate Limits

An attacker that engages in a scaled attack may need to request a large number of user codes (see exploit Section 4.3.1) or initiate a large number of authorization requests (see exploit Section 4.3.4) in a short period of time. An authorization server MAY apply rate limits to minimize the number of requests it would accept from a client in a limited time period.

**\*Limitations:** Rate limits are effective at slowing an attacker down and help to degrade scaled attacks, but do not prevent more targeted attacks that are executed with lower volumes and velocity. Therefore, it should be used along with other techniques to provide a defence-in-depth defence against cross-device attacks.

#### 6.1.12. Sender-Constrained Tokens

Sender-constrained tokens limit the impact of a successful attack by preventing the tokens from being moved from the device on which the attack was successfully executed. This makes attacks where an attacker gathers a large number of access and refresh tokens on a single device and then sells them for profit more difficult, since the attacker would also have to export the cryptographic keys used to sender-constrain the tokens or be able to access them and generate signatures for future use. If the attack is being executed on a trusted device to a device with anti-malware, any attempts to exfiltrate tokens or keys may be detected and the device's trust status may be changed. Using hardware keys for sender-constraining tokens will further reduce the ability of the attacker to move tokens to another device.

**\*Limitations:** Sender-constrained tokens, especially sender-constrained tokens that require proof-of-possession, raise the bar for executing the attack and profiting from exfiltrating tokens. Although a software proof-of-possession key is better than no proof-of-possession key, an attacker may still exfiltrate the software key. Hardware keys are harder to exfiltrate, but come with additional implementation complexity. An attacker that controls the Consumption Device may still be able to exercise the key, even if it is in hardware. Consequently the main protection derived from sender-constrained tokens is preventing tokens from being moved from the Consumption Device to another device, thereby making it harder sell stolen tokens and profit from the attack.

#### 6.1.13. User Education

Research shows that user education is effective in reducing the risk of phishing attacks [Baki2023]. The service provider MAY educate users on the risks of cross-device consent phishing and provide out-of-band reinforcement to the user on the context and conditions under which an authorization grant may be requested. For example, if the service provider does not send e-mails with QR codes requesting users to grant authorization, this may be reinforced in marketing messages and anti-fraud awareness campaigns. The service provider MAY also choose to reinforce these user education messages through in-app experiences. In [PCRS2023], it is proposed to advise users to verify the trustworthiness of the source of a QR code, for instance

by checking that the connection is protected through TLS or by verifying that the URL really belongs to the Authorization Server.

**\*Limitations:** Although user education helps to raise awareness and reduce the overall risk to users, it is insufficient on its own to mitigate cross-device consent phishing attacks. In particular, carefully designed phishing attacks can be practically indistinguishable from benign authorization flows even for well-trained users. User education SHOULD therefore be used in conjunction with other controls described in this document.

#### 6.1.14. User Experience

The user experience SHOULD preserve the context within which the protocols were initiated and communicate this clearly to the user when they are asked to authorize, authenticate or present a credential. In preserving the context, it should be clear to the user who invoked the flow, why it was invoked and what the consequence of completing the authorization, authentication or credential presentation is. The user experience SHOULD reinforce the message that unless the user initiated the authorization request, or was expecting it, they should decline the request.

This information MAY be communicated graphically or in a simple message (e.g., "It looks like you are trying to access your files on a digital whiteboard in your city center office. Click here to grant access to your files. If you are not trying to access your files, you should decline this request and notify the security department").

It SHOULD be clear to the user how to decline the request. To avoid accidental authorization grants, the "decline" option SHOULD be the default option or given similar prominence in the user experience as the "grant" option.

If the user uses an application on a mobile device to scan a QR code, the application MAY display information advising the user under which conditions they should expect to be asked to scan a QR code and under which circumstances they should never scan a QR code (e.g., display a message that the QR code will only be displayed on kiosks within trusted locations or on trusted websites hosted on a specific domain, and never in e-mail or other media and locations).

The user experience MAY include information to further educate the user on cross-device consent phishing attacks and reinforce the conditions under which authorization grants may be requested.

**\*Limitations:** Improvements to user experience on their own is unlikely to be sufficient and SHOULD be used in conjunction with other controls described in this document.

#### 6.1.15. Authenticate-then-Initiate

By requiring a user to authenticate on the Consumption Device with a phishing resistant authentication method before initiating a cross-device flow, the server can prevent an attacker from initiating a cross-device flow and obtaining QR codes or user codes. This prevents the attacker from obtaining a QR code or user code that they can use to mislead an unsuspecting user. This requires that the Consumption Device has sufficient input capabilities to support a phishing resistant authentication mechanism, which may in itself negate the need for a cross-device flow.

**\*Limitations:** Authenticating on the Consumption Device before starting a cross-device flow does not prevent the attacks described in Section 4.3.6 and Section 4.3.8 and it is RECOMMENDED that additional mitigations described in this document is used if the cross-device flows are used in scenarios such as Section 3.3.5 and Section 3.3.7.

#### 6.1.16. Request Initiation Verification

The user MAY be asked to confirm if they initiated an authentication or authorization request by sending a one-time password (OTP) or PIN to the user's Authorization Device and asking them to enter it on the Consumption Device to confirm the request. If the request was initiated without the users' consent, they would receive an OTP or PIN out of context which may raise suspicion for the user. In addition, they would not have information on where to enter the OTP or PIN. The user experience on the Authorization Device MAY reinforce the risk of receiving an out-of-context OTP or PIN and provide information to the user on how to report an unauthorized authentication or authorization request.

**\*Limitations:** The additional verification step may reduce the overall usability of the system as it is one more thing users need to do right. Attackers may combine traditional phishing attacks and target users who respond to those messages with an interactive attack that sets the expectation with the user that they will have to provide the OTP or PIN, in addition to granting authorization for the request.

#### 6.1.17. Request Binding with Out-of-Band Data

In the User-Transferred Session Data Pattern, users MAY enter out-of-band information on the Consumption Device to start the authorization process. The out-of-band data entered by the user MAY then be included in the QR code which is displayed on the Consumption Device. When the QR code is scanned by the Authorization Device, the out-of-band data is verified by the user or by the Authorization Device. The out-of-band data could be any attribute that the user or Authorization Device can retrieve during the authorization process. Examples include a serial number, one-time password or PIN, location or any other data that the user or the Authorization Device can recall or retrieve during the authorization process ([MPRCS2020], [PCRS2023]).

**\*Limitations:** A sophisticated attacker may include an additional step in their attack where they create a phishing attack that gathers the out-of-band data from the user before initiating the authorization request. The additional step could also have a negative impact on the usability level of the solution.

#### 6.1.18. Practical Mitigation Summary

The practical mitigations described in this section can prevent the attacks from being initiated, disrupt attacks once they start or reduce the impact or remediate an attack if it succeeds. When combining one or more of these mitigations the overall security profile of a cross-device flow improves significantly. The following table provides a summary view of these mitigations:

Mitigation	Prevent	Disrupt	Recover
Establish Proximity	X	X	
Short Lived/Timebound Codes		X	
One-Time or Limited Use Codes		X	
Unique Codes		X	
Content Filtering		X	
Detect and remediate			X
Trusted Devices	X		
Trusted Networks	X		
Limited Scopes			X
Short Lived Tokens			X
Rate Limits	X	X	
Sender-Constrained Tokens			X
User Education	X		
User Experience	X		
Authenticate-then-Initiate	X		
Request Initiation Verification		X	
Request Binding with Out-of-Band Data		X	

Table 1

Table: Practical Mitigation Summary

## 6.2. Protocol Selection

Some cross-device protocols are more susceptible to the exploits described in this document than others. In this section we will compare three different cross-device protocols in terms of their susceptibility to exploits focused on the unauthenticated channel, the prerequisites to implement and deploy them, along with guidance on when it is appropriate to use them.

### 6.2.1. IETF OAuth 2.0 Device Authorization Grant [RFC8628]:

#### 6.2.1.1. Description

A standard to enable authorization on devices with constrained input capabilities (smart TVs, printers, kiosks). In this protocol, the user code or QR code is displayed or made available on the Consumption Device (smart TV) and entered on a second device (e.g., a mobile phone).

#### 6.2.1.2. Susceptibility

There are several reports in the public domain outlining how the unauthenticated channel may be exploited to execute a Cross-Device Consent Phishing attack ([Exploit1], [Exploit2], [Exploit3], [Exploit4], [Exploit5], [Exploit6]).

#### 6.2.1.3. Device Capabilities

There are no assumptions in the protocol about underlying capabilities of the device, making it a "least common denominator" protocol that is expected to work on the broadest set of devices and environments.

#### 6.2.1.4. Mitigations

In addition to the security considerations section in the standard, it is RECOMMENDED that one or more of the mitigations outlined in this document be considered, especially mitigations that can help establish proximity or prevent attackers from obtaining QR or user codes.

#### 6.2.1.5. When to use

Only use this protocol if other cross-device protocols are not viable due to device or system constraints. Avoid using if the protected resources are sensitive, high value, or business critical. Always deploy additional mitigations like proximity or only allow with pre-registered devices. Do not use for same-device scenarios (e.g., if the Consumption Device and Authorization Device is the same device).

### 6.2.2. OpenID Foundation Client Initiated Back-Channel Authentication (CIBA):

#### 6.2.2.1. Description

Client Initiated Back-Channel Authentication (CIBA) [CIBA]: A standard developed in the OpenID Foundation that allows a device or service (e.g., a personal computer, smart TV, Kiosk) to request the OpenID Provider to initiate an authentication flow if it knows a valid identifier for the user. The user completes the authentication flow using a second device (e.g., a mobile phone). In this flow the user does not scan a QR code or obtain a user code from the Consumption Device, but is instead contacted by the OpenID Provider to complete the authentication using a push notification, e-mail, text message or any other suitable mechanism.

#### 6.2.2.2. Susceptibility

Less susceptible to unauthenticated channel attacks, but still vulnerable to attackers who know or can guess the user identifier and initiate an attack as described in Section 4.3.4.

#### 6.2.2.3. Device Capabilities

There is no requirement on the Consumption Device to support specific hardware. The Authorization Device must be registered/associated with the user and it must be possible for the Authorization Server to trigger an authorization on this device.

#### 6.2.2.4. Mitigations

In addition to the security considerations section in the standard, it is RECOMMENDED that one or more of the mitigations outlined in this document be considered, especially mitigations that can help establish proximity or prevent attackers from initiating authorization requests.

#### 6.2.2.5. When to Use

Use CIBA instead of Device Authorization Grant if it is possible for the Consumption Device to obtain a user identifier on the Consumption Device (e.g., through an input or selection mechanism) and if the Authorization Server can trigger an authorization on the Authorization Device. Do not use for same-device scenarios (e.g., if the Consumption Device and Authorization Device is the same device).

#### 6.2.3. FIDO2/WebAuthn

##### 6.2.3.1. Description

FIDO2/WebAuthn is a stack of standards developed in the FIDO Alliance and W3C respectively which allow for origin-bound, phishing-resistant user authentication using asymmetric cryptography that can be invoked from a web browser or native client. Version 2.2 of the FIDO Client to Authenticator Protocol (CTAP) supports a new cross-device authentication protocol, called "hybrid transports", which enables an external device, such as a phone or tablet, to be used as a roaming authenticator for signing into the primary device, such as a personal computer. This is commonly called FIDO Cross-Device Authentication (CDA). CTAP 2.2 hybrid transports is implemented by the client and authenticator platforms.

When a user wants to authenticate using their mobile device (authenticator) for the first time, they need to link their authenticator to their main device. This is done using a scan of a QR code. When the authenticator scans the QR code, the device sends an encrypted BLE advertisement containing keying material and a tunnel ID. The main device (CTAP client) and authenticator both establish connections to the web service, and the normal CTAP protocol exchange occurs.

If the user chooses to keep their authenticator linked with the main device, the QR code link step is not necessary for subsequent use. The user will receive a push notification on the authenticator.

##### 6.2.3.2. Susceptibility

The Cross-Device Authentication flow proves proximity by leveraging BLE advertisements for service establishment, significantly reducing the susceptibility to any of the exploits described in Examples 1-6.

#### 6.2.3.3. Device Capabilities

Both the Consumption Device and the authenticator require BLE support and access to the internet. The Consumption Device must support both the WebAuthn API [W3CWebAuthn] (or a platform-specific WebAuthn abstraction for native apps) and the FIDO Client to Authenticator Protocol (CTAP), specifically version 2.2 with hybrid transports [FIDOCTAP22]. The device serving as the FIDO authenticator must also support CTAP 2.2 or later to be used as a cross-device authenticator.

#### 6.2.3.4. Mitigations

FIDO Cross-Device Authentication (CDA) establishes proximity through the use of BLE, reducing the need for additional mitigations. An implementer MAY still choose to implement additional mitigation as described in this document.

#### 6.2.3.5. When to Use

FIDO2/WebAuthn SHOULD be used for cross-device authentication scenarios whenever the devices are capable of doing so and a suitable FIDO credential is not available on the Consumption Device. It MAY be used as an authentication method with the Authorization Code Grant [RFC6749] and PKCE [RFC7636], to grant authorization to a Consumption Device (e.g., smart TV or interactive whiteboard) using a device serving as the FIDO authenticator (e.g. a mobile phone) for authentication. This combination of FIDO2/WebAuthn and Authorization Code Flow with PKCE enables cross device authorization flows, without the risks posed by the Device Authorization Grant [RFC8628].

#### 6.2.4. Protocol Selection Summary

The FIDO Cross-Device Authentication (CDA) flow provides the best protection against attacks on the unauthenticated channel for cross device flows. It can be combined with OAuth 2.0 and OpenID Connect protocols for standards-based authorization and authentication flows. If FIDO2/WebAuthn support is not available, Client Initiated Backchannel Authentication (CIBA) provides an alternative, provided that there is a channel through which the authorization server can contact the end user. Examples of such a channel include device push notifications, e-mail or text messages which the user can access from their device. If CIBA is used, additional mitigations to enforce proximity and initiate transactions from trusted devices or trusted networks SHOULD be considered. The OAuth 2.0 Device Authorization Grant provides the most flexibility and has the lowest requirements on devices used, but it is RECOMMENDED that it is only used when additional mitigations are deployed to prevent attacks that exploit the unauthenticated channel between devices.

### 6.3. Foundational Pillars

Experience with web authorization and authentication protocols such as OAuth and OpenID Connect has shown that securing these protocols can be hard. The major reason for this is that the landscape in which they are operating - the web infrastructure with browsers, servers, and the underlying network - is complex, diverse, and ever-evolving.

As is the case with other kinds of protocols, it can be easy to overlook vulnerabilities in this environment. One way to reduce the chances of hidden security problems is to use mathematical-logical models to describe the protocols, their environments and their security goals, and then use these models to try to prove security. This approach is what is usually subsumed as "formal security analysis".

There are two major strengths of formal analysis: First, finding new vulnerabilities does not require creativity - i.e., new classes of attacks can be uncovered even if no one thought of these attacks before. In a faithful model, vulnerabilities become clear during the proof process or even earlier. Second, formal analysis can exclude the existence of any attacks within the boundaries of the model (e.g., the protocol layers modeled, the level of detail and functionalities covered, the assumed attacker capabilities, and the formalized security goals). As a downside, there is usually a gap between the model (which necessarily abstracts away from details) and implementations. In other words, implementations can introduce flaws where the model does not have any. Nonetheless, for protocol standards, formal analysis can help to ensure that the specification is secure when implemented correctly.

There are various different approaches to formal security analysis and each brings its own strengths and weaknesses. For example, models differ in the level of detail in which they can capture a protocol (granularity, expressiveness), in the kind of statements they can produce, and whether the proofs can be assisted by tools or have to be performed manually.

The following works have been identified as relevant to the analysis of cross-device flows:

- \* In "Formal analysis of self-issued OpenID providers" [Bauer2022], the protocol of [OpenID.SIOPV2] was analyzed using the Web Infrastructure Model (WIM). The WIM is specifically designed for the analysis of web authentication and authorization protocols. While it is a manual (pen-and-paper) model, it captures details of browsers and web interactions to a degree that is hard to match in

automated models. In previous works, previously unknown flaws in OAuth, OpenID Connect, and FAPI were discovered using the WIM. In the analysis of a cross-device SIOP V2 flow in [Bauer2022], the request replay attack already described in Section 13.3 of [OpenID.SIOPV2] was confirmed in the model. A mitigation was implemented based on a so-called Cross-Device Stub, essentially a component that serves to link the two devices before the protocol flow starts. This can be seen as an implementation of a trusted device relationship as described in Section 6.1.7. The mitigation was shown to be effective in the model.

- \* In "Security analysis of the Grant Negotiation and Authorization Protocol" [Helmschmidt2022], an analysis of a draft of the Grant Negotiation and Authorization Protocol (GNAP) [RFC9635] was performed using the Web Infrastructure Model. The same attack as in [Bauer2022] was found to apply to GNAP as well. In this case, a model of a "careful user" (see Section 6.1.13) was used to show that the attack can be prevented (at least in theory) by the user.
- \* In "The Good, the Bad and the (Not So) Ugly of Out-of-Band Authentication with eID Cards and Push Notifications: Design, Formal and Risk Analysis" [MPRCS2020], Pernpruner et al. formally analysed an authentication protocol relying on push notifications delivered to an out-of-band device to approve the authentication attempt on the primary device (Backchannel-Transferred Session Pattern, Section 3.1.2). The analysis was performed using the specification language ASLan++ and the model checker SATMC. According to the results of the analysis, they identified and defined the category of `_implicit attacks_`, which manage to deceive users into approving a malicious authentication attempt through social engineering techniques, thus not compromising all the authentication factors involved; these attacks are aligned with the definition of Cross-Device Consent Phishing (CDCP) attacks.
- \* In "An Automated Multi-Layered Methodology to Assist the Secure and Risk-Aware Design of Multi-Factor Authentication Protocols" [PCRS2023], Pernpruner et al. defined a multi-layered methodology to analyze multi-factor authentication protocols at different levels of granularity. They leveraged their methodology to formally analyze a protocol relying on a QR code that has to be scanned on a secondary device to approve the authentication attempt on the primary device (User-Transferred Session Data Pattern, Section 3.1.1). Given the results of the analysis, they proposed some practical mitigations either to prevent or reduce the risk of successful attacks, such as those described in Section 6.1.13, Section 6.1.16 and Section 6.1.17.

## 7. Conclusion

Cross-device flows enable authorization on devices with limited input capabilities, allow for secure authentication when using public or shared devices, provide a path towards multi-factor authentication and, provide the convenience of a single, portable credential store.

The popularity of cross-device flows attracted the attention of attackers that exploit the unauthenticated channel between the Consumption Device and Authorization Device using techniques commonly used in phishing attacks. These Cross-Device Consent Phishing (CDCP) attacks allow attackers to obtain access and refresh tokens, rather than authentication credentials, resulting in access to resources even if the user used multi-factor authentication.

To address these attacks, we propose a three pronged approach that includes the deployment of practical mitigations to safeguard protocols that are already deployed, provide guidance on when to use different protocols, including protocols that are not susceptible to these attacks, and the introduction of formal methods to evaluate the impact of mitigations and find additional issues.

## 8. Contributors

The authors would like to thank Tim Cappalli, Nick Ludwig, Adrian Frei, Nikhil Reddy Boreddy, Bjorn Hjelm, Joseph Heenan, Brian Campbell, Damien Bowden, Kristina Yasuda, Tim W端rtele, Karsten Meyer zu Selhausen, Maryam Mehrnezhad, Marco Pernpruner, Giada Sciarretta, Dean H. Saxe, Roy Williams, Dan Moore and others (please let us know, if you've been mistakenly omitted) for their valuable input, feedback and general support of this work.

## 9. References

### 9.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/rfc/rfc2119>>.
- [RFC6749] Hardt, D., Ed., "The OAuth 2.0 Authorization Framework", RFC 6749, DOI 10.17487/RFC6749, October 2012, <<https://www.rfc-editor.org/rfc/rfc6749>>.

- [RFC7636] Sakimura, N., Ed., Bradley, J., and N. Agarwal, "Proof Key for Code Exchange by OAuth Public Clients", RFC 7636, DOI 10.17487/RFC7636, September 2015, <<https://www.rfc-editor.org/rfc/rfc7636>>.
- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/rfc/rfc8174>>.
- [RFC8628] Denniss, W., Bradley, J., Jones, M., and H. Tschofenig, "OAuth 2.0 Device Authorization Grant", RFC 8628, DOI 10.17487/RFC8628, August 2019, <<https://www.rfc-editor.org/rfc/rfc8628>>.
- [RFC7662] Richer, J., Ed., "OAuth 2.0 Token Introspection", RFC 7662, DOI 10.17487/RFC7662, October 2015, <<https://www.rfc-editor.org/rfc/rfc7662>>.
- [CIBA] Rodriguez, G. F., Walter, F., Nennker, A., Tonge, D., and B. Campbell, "OpenID Connect Client-Initiated Backchannel Authentication Flow - Core 1.0", September 2021, <[https://openid.net/specs/openid-client-initiated-backchannel-authentication-core-1\\_0.html](https://openid.net/specs/openid-client-initiated-backchannel-authentication-core-1_0.html)>.
- [CAEP] Tulshibagwale, A. and T. Cappalli, "OpenID Continuous Access Evaluation Profile 1.0 - draft 01", June 2021, <[https://openid.net/specs/openid-caep-specification-1\\_0-01.html](https://openid.net/specs/openid-caep-specification-1_0-01.html)>.
- [SSF] Tulshibagwale, A., Cappalli, T., Scurtescu, M., Backman, A., and J. Bradley, "OpenID Shared Signals and Events Framework Specification 1.0", June 2021, <[https://openid.net/specs/openid-sse-framework-1\\_0-01.html](https://openid.net/specs/openid-sse-framework-1_0-01.html)>.
- [W3CWebAuthn]  
Cappalli, T., Jones, M.B., Kumar, A., Lundberg, E., and M. Miller, "Web Authentication: An API for accessing Public Key Credentials Level 3", January 2025, <<https://www.w3.org/TR/2025/WD-webauthn-3-20250127/>>.
- [FIDOCTAP22]  
Bradley, J., Jones, M.B., Kumar, A., Lindemann, R., Verrept, S., and D. Waite, "Client to Authenticator Protocol (CTAP)", February 2025, <<https://fidoalliance.org/specs/fido-v2.2-ps-20250228/fido-client-to-authenticator-protocol-v2.2-ps-20250228.html>>.

[IEEE802154]

Institute of Electrical and Electronics Engineers, "IEEE Std 802.15.4-2020: IEEE Standard for Low-Rate Wireless Networks", {"name" => "IEEE", "value" => "802.15.4-2020"} , 2020, <[https://standards.ieee.org/standard/802\\_15\\_4-2020.html](https://standards.ieee.org/standard/802_15_4-2020.html)>.

## 9.2. Informative References

- [RFC9635] Richer, J., Ed. and F. Imbault, "Grant Negotiation and Authorization Protocol (GNAP)", RFC 9635, DOI 10.17487/RFC9635, October 2024, <<https://www.rfc-editor.org/rfc/rfc9635>>.
- [Exploit1] Cooke, B., "The Art of the Device Code Phish", July 2021, <<https://0xboku.com/2021/07/12/ArtOfDeviceCodePhish.html>>.
- [Exploit2] Min, D., "Microsoft 365 OAuth Device Code Flow and Phishing", August 2021, <<https://www.optiv.com/insights/source-zero/blog/microsoft-365-oauth-device-code-flow-and-phishing>>.
- [Exploit3] Syynimaa, N., "Introducing a new phishing technique for compromising Office 365 accounts", October 2020, <<https://o365blog.com/post/phishing/#new-phishing-technique-device-code-authentication>>.
- [Exploit4] Hwong, J., "New Phishing Attacks Exploiting OAuth Authentication Flows (DEFCON 29)", August 2021, <<https://www.youtube.com/watch?v=9slRYvpKHp4>>.
- [Exploit5] Secureworks Counter Threat Unit (CTU), "OAuth's Device Code Flow Abused in Phishing Attacks", August 2021, <<https://www.secureworks.com/blog/oauths-device-code-flow-abused-in-phishing-attacks>>.
- [Exploit6] Talebzadeh, K. and N. Romsdahl, "SquarePhish: Advanced phishing tool combines QR codes and OAuth 2.0 device code flow", August 2022, <<https://www.helpnetsecurity.com/2022/08/11/squarephish-video/>>.
- [NYC.Bike] Byrne, K. J., "Citi Bikes being swiped by joyriding scammers who have cracked the QR code", August 2021, <<https://nypost.com/2021/08/07/citi-bikes-being-swiped-by-joyriding-scammers-who-have-cracked-the-qr-code/>>.

## [OpenID.SIOPV2]

Yasuda, K., Jones, M. B., and T. Lodderstedt, "Self-Issued OpenID Provider v2", November 2022, <[https://bitbucket.org/openid/connect/src/master/openid-connect-self-issued-v2/openid-connect-self-issued-v2-1\\_0.md](https://bitbucket.org/openid/connect/src/master/openid-connect-self-issued-v2/openid-connect-self-issued-v2-1_0.md)>.

## [OpenID.VP]

Terbu, O., Lodderstedt, T., Yasuda, K., and T. Looker, "OpenID for Verifiable Credential Presentations", November 2023, <[https://openid.net/specs/openid-4-verifiable-presentations-1\\_0.html](https://openid.net/specs/openid-4-verifiable-presentations-1_0.html)>.

## [OpenID.VCI]

Lodderstedt, T., Yasuda, K., and T. Looker, "OpenID for Verifiable Credential Issuance", October 2023, <[https://openid.net/specs/openid-4-verifiable-credential-issuance-1\\_0.html](https://openid.net/specs/openid-4-verifiable-credential-issuance-1_0.html)>.

## [Bauer2022]

Bauer, C., "Formal analysis of self-issued OpenID providers", 2022, <<https://elib.uni-stuttgart.de/handle/11682/12417>>.

## [OpenID.Core]

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

## [PCRS2023]

Pernpruner, M., Carbone, R., Sciarretta, G., and S. Ranise, "An Automated Multi-Layered Methodology to Assist the Secure and Risk-Aware Design of Multi-Factor Authentication Protocols, IEEE Transactions on Dependable and Secure Computing (TDSC)", 2023, <<https://doi.org/10.1109/TDSC.2023.3296210>>.

## [MPRCS2020]

Pernpruner, M., Carbone, R., Ranise, S., and G. Sciarretta, "The Good, the Bad and the (Not So) Ugly of Out-of-Band Authentication with eID Cards and Push Notifications: Design, Formal and Risk Analysis, Proceedings of the Tenth ACM Conference on Data and Application Security and Privacy, Pages 223234, Association for Computing Machinery", 2020, <<https://doi.org/10.1145/3374664.3375727>>.

[Helmschmidt2022]

Helmschmidt, F., "Security analysis of the Grant Negotiation and Authorization Protocol", 2022, <<https://elib.uni-stuttgart.de/handle/11682/12220>>.

[Baki2023] Baki, S. and R. M. Verma, "Sixteen Years of Phishing User Studies: What Have We Learned?", IEEE Transactions on Dependable and Secure Computing, Volume 20, Number 2, Pages 12001212", 2023, <<https://doi.org/10.1109/TDSC.2022.3151103>>.

#### Appendix A. Document History

[[ To be removed from the final specification ]]

-latest

- \* Fixed malformed labels
- \* Clarified common use case for when phone and TVs do not use the same network.
- \* Clarified role of authorization server in establishing proximity.
- \* Clarified which mitigations can be implented by the authorization server only.
- \* Add Dan Moore to acknowledgements

-10

- \* Shepherd feedback: Describe unauthenticated channel.
- \* Shepherd feedback: Seperate normative and informative references.
- \* Shepherd feedback: Update FIDO/WebAuthn references

-09

- \* Affiliation change to allow publication to Datatracker.
- \* No content changes - re-published to avoid expiry while waiting on shepherd review.

-08

- \* Editorial updates.

-07

- \* Clarification of FIDO\WebAuthn section.
- \* Updated language in section on FIDO to allow for use of FIDO keys on consumption devices.
- \* Clarified origin of QR Code.
- \* Editorial updates
- \* Updated examples to be consistent.
- \* Made diagram description clearer.
- \* Added CTAP 2.2 Draft.
- \* Added additional guidance on geolocation inaccuracies.
- \* Added Roy Williams to acknowledgements
- \* Clarified that authorization servers can detect
- \* Consistent use of "smart TV"
- \* Fixed references

-06

- \* Corrected typos.

-05

- \* Added section to provide actionable guidance to implementers on how to use this document.
- \* Expanded section on formal analysis to include completed research projects.
- \* Added reference to OpenID for Verifiable Presentations.

-04

- \* Corrected formatting issue that prevented history from showing correctly.

-03

- \* Introduced normative SHOULD, RECOMMENDED and MAY when applied to actions the Authorization Server, Resource Server or Client may implement.
- \* Added User Education as a standalone mitigation.
- \* Added Maryam Mehrnezhad, Marco Pernpruner and Giada Sciarretta to the contributors list.
- \* Added Request Binding with Out-of-Band Data as an additional mitigation.
- \* Adopted the OpenID Foundation terminology from [CIBA] and changed Initiating Device to Consumption Device
- \* Added Fake Helpdesk and Consent Request Overload examples
- \* Replaced "Authenticated Flow" mitigation name with "Authenticate-then-Intitiate"
- \* Added Cross-Device Session Transfer pattern

-02

- \* Fixed typos and grammar edits
- \* Capitalised Initiating Device and Authorization Device
- \* Introduced Cross-Device Consent Phishing as a label for the types of attacks described in this document.
- \* Updated labels for different types of flows (User-Transferred Session Data Pattern, Backchannel-Transferred Session Pattern, User-Transferred Authorization Data Pattern)
- \* Adopted consistent use of hyphenation in using "cross-device"
- \* Consistent use of "Authorization Device"
- \* Update Reference to Secure Signals Framework to reflect name change from Secure Signals and Events
- \* Described difference between proximity enforced and proximity-less cross-device flows
- \* General editorial pass

-01

- \* Added additional diagrams and descriptions to distinguish between different cross-device flow patterns.
  - \* Added short description on limitations of each mitigation.
  - \* Added acknowledgement of additional contributors.
  - \* Fixed document history format.
- 00 (Working Group Draft)
- \* Initial WG revision (content unchanged from draft-kasselman-cross-device-security-03)
- 03 draft-kasselman-cross-device-security
- \* Minor edits and typos
- 02 draft-kasselman-cross-device-security
- \* Minor edits and typos
  - \* Upload as draft-ietf-oauth-cross-device-security-best-practice-02
- 01 draft-kasselman-cross-device-security
- \* Updated draft based on feedback from version circulated to OAuth working group
  - \* Upload as draft-ietf-oauth-cross-device-security-best-practice-01
- 00 draft-kasselman-cross-device-security
- \* Initial draft adopted from document circulated to the OAuth Security Workshop Slack Channel
  - \* Upload as draft-ietf-oauth-cross-device-security-best-practice-00

#### Authors' Addresses

Pieter Kasselman  
SPIRL  
Email: pieter@spirl.com

Daniel Fett  
Authlete  
Email: mail@danielfett.de

Filip Skokan  
Okta  
Email: panva.ip@gmail.com