

A First Look at Digital Rights Management Systems for Secure Mobile Content Delivery

Amir Rafi
Information Security Group
Royal Holloway, University of London
Egham, Surrey, United Kingdom
amir.rafi.2019@live.rhul.ac.uk

Carlton Shepherd
School of Computing
Newcastle University
Newcastle-upon-Tyne, United Kingdom
carlton.shepherd@ncl.ac.uk

Konstantinos Markantonakis
Information Security Group
Royal Holloway, University of London
Egham, Surrey, United Kingdom
k.markantonakis@rhul.ac.uk

Abstract—Digital rights management (DRM) solutions aim to prevent the copying or distribution of copyrighted material. On mobile devices, a variety of DRM technologies have become widely deployed. However, a detailed security study comparing their internal workings, and their strengths and weaknesses, remains missing in the existing literature. In this paper, we present the first detailed security analysis of mobile DRM systems, addressing the modern paradigm of cloud-based content delivery followed by major platforms, such as Netflix, Disney+, and Amazon Prime. We extensively analyse the security of three widely used DRM solutions—Google Widevine, Apple FairPlay, and Microsoft PlayReady—deployed on billions of devices worldwide. We then consolidate their features and capabilities, deriving common features and security properties for their evaluation. Furthermore, we identify some design-level shortcomings that render them vulnerable to emerging attacks within the state of the art, including micro-architectural side-channel vulnerabilities and an absence of post-quantum security. Lastly, we propose mitigations and suggest future directions of research.

Index Terms—Digital rights management (DRM), media streaming, trusted execution environments, mobile security

I. INTRODUCTION

Video streaming consumes over half of the global Internet bandwidth [1]. The market share of over-the-top (OTT) media services—those offered directly to viewers via the Internet, bypassing cable, broadcast, and satellite television—is projected to reach USD 139 billion in 2028 [2]. Current trends in OTT broadcasting indicate the migration of workflows to the cloud, from production and packaging to delivery [3]. The OTT model, however, poses some significant security challenges. Media content is received and displayed by devices that are under the full control of customers, i.e. mobile phones and tablets. Cloud-based content providers, such as Netflix and Disney+, rely on DRM to prevent copyright infringements and unauthorised consumption of content. A DRM system restricts content usage by delivering encrypted content to authorised devices and controlling its decryption; legitimate users are thus potential adversaries in a typical DRM threat model.

The current DRM landscape consists of proprietary solutions causing a fragmentation in the industry; content providers are thereby compelled to implement multiple DRM systems

This work has been supported as part of the EPSRC Centre for Doctoral Training in Cyber Security for the Everyday at Royal Holloway, University of London (EP/S021817/1).

to support multiple platforms. Today’s mobile platforms are shipped with proprietary DRM systems, which are integrated into the application or operating system itself, such as Android (Widevine), iOS (FairPlay), and Microsoft Windows (PlayReady). Android and iOS have a global OS market share of 72.11% and 27.22% respectively for mobiles; 47.54% and 52.38% respectively for tablets; and Windows OS is used on 76.31% of desktops [4]. Widevine, for instance, is available on 5 billion devices worldwide [5].

A. Contributions

Since 1998, the Digital Millennium Copyright Act (DMCA) has hindered security researchers from analysing DRM systems. Recent literature describes a number of exploits, reverse engineering of code obfuscation techniques used for protecting low quality content, and analysis of OTT apps that rely on DRM protection. As a result, despite their widespread use, the understanding of various DRM technologies remains highly fragmented. This paper aims to bridge this gap. We systematically evaluate the security of mainstream DRM systems using an approach and criteria that analyses such systems holistically. Such an evaluation remains forthcoming in the existing literature. Our contributions can thus be summarised as follows:

- To the best of our knowledge, we present the first domain-specific security evaluation of mobile-based DRM systems, considering the security of both the content and user devices.
- Based on a comprehensive analysis of public domain material, we identify probable attack vectors and present a comparative security analysis (§V) of Widevine, FairPlay and PlayReady DRM systems.
- We discuss mitigations (§VI) for security vulnerabilities discovered during our analysis and suggest future directions of research (§VII).

B. Scope

A DRM system can be deployed on user platforms as a separate application, as a part of the web browser, or as a part of the OS, with optional hardware protection. From a deployment perspective, DRM solutions may be found in the form of an obfuscated library (see Widevine L3 [6]) or,

increasingly, within a trusted execution environment (TEE) that prevents kernel-mode adversaries from accessing relevant assets, such as content decryption keys. Broadly speaking, the most sophisticated DRM implementations are integrated into the OS, leveraging kernel-level security mechanisms, and hardware-assisted protections offered by TEEs.

The scope of this security analysis is limited to DRM deployments implemented as a part of the OS and using TEEs for hardware-level protection. In this work, we focus on DRM systems implemented on mobile platforms, comprising smartphones, tablets, and laptop devices possessed by viewers of media content (i.e. the general public). DRM solutions are notoriously closed-source in nature, which restricts a complete understanding of their characteristics. To conduct this study, we thus rely on publicly available materials pertaining to Widevine, FairPlay and PlayReady to analyse their security. We take a security-oriented approach to this work; issues relating to privacy and content watermarking are deemed outside the scope of this work.

II. BACKGROUND & RELATED WORK

The DMCA was signed in to the United States law in 1998 to discourage content piracy, and amended as Title 17 §1201 [7] prohibiting the circumvention measures for protecting copyrighted material. Although piracy continued to be conducted, the DMCA restricted security researchers from analysing DRM systems [8]. In the 2000s, DMCA takedowns were used against security researchers who identified issues with the High Bandwidth Digital Content Protection (HDCP) scheme developed by the Intel Corporation [9], [10]. By 2016, the DMCA anti-circumvention claims had been raised in nearly 50 reported cases in the United States federal courts [11]. The exemptions to DMCA were expanded by the Copyright office and the Library of Congress in 2018, allowing security researchers greater freedom to analyse DRM systems.

In recent years (>2012), the growth of OTT mobile streaming has led to new generation of DRM systems, prompting researchers to analyse their claimed security assurances. The MITRE CVE database lists 53 CVE records covering different security issues in Widevine and PlayReady since 2014.¹ In 2013, Wang et al. [14] introduced an automated technique to remove DRM protection from content protected by PlayReady. This technique tracks the flow of content between memory buffers in real-time during playback, identifies where the content is decrypted, and copies the decrypted content, which is then reconstructed into a new media file without any DRM protection. The authors demonstrated its effectiveness against Amazon Instant Video, Hulu, Netflix, and Spotify.

D’Orazio and Choo [15] proposed a DRM adversary model in 2016 to formalise potential attackers’ capabilities and to evaluate DRM protection on iOS devices. They highlighted vulnerabilities resulting from a lack of public-key pinning and server-side verification. Since 2013, a technique known as

¹There are no CVE records for FairPlay, although work has identified some significant security weaknesses [12], [13].

“FairPlay MITM (Man-In-The-Middle)” [12] exploited vulnerabilities in FairPlay to propagate pirate iOS apps. In 2016, Xiao [16] reported how FairPlay MITM could be used to spread malware on iOS devices.

Against Widevine, a vulnerability [17] was discovered in 2016 that allowed access to a local cache containing decrypted version of protected content. This vulnerability was tested against Amazon Prime and Netflix, although other DRM systems may also be vulnerable. In 2018, Chau et al. [13] analysed 141 content delivery apps on the Android platform and discovered vulnerabilities that not only compromised the DRM-protected content but also put users’ security and privacy at risk. The vulnerabilities were caused by insecure policy delivery and bootstrapping, client-side policy enforcement, and reusing encryption keys; the whole Amazon Music collection of 40 million songs was found to be encrypted with the same key. The work evaluated the apps that *use* DRM systems for content protection, but not the DRM systems themselves.

In 2019, Buchanan [18] used differential fault analysis to recover keys for protected content using Widevine level 3 (L3) [19] (see §III-A). Most GitHub repositories containing this proof of concept code were removed by Google’s DMCA take-down [20] in 2020. More recently, Zhao [21] broke in to the Widevine L1 security level on the Qualcomm TEE, an Android TrustZone implementation. The L1 is the most secure implementation of Widevine, which uses TEEs on user devices to protect content and encryption keys, and to securely decrypt content. Zhao recovered the Widevine keybox from the trusted storage of QTEE by breaking the ASLR (Address Space Layout Randomization).

Based on an analysis of OTT apps including Netflix, Hulu and Showtime, Patat et al. [22] revealed that most apps use the same cryptographic keys for a particular media item across all users. Moreover, these apps support vulnerable legacy devices, which can be exploited to compromise the cryptographic keys. In 2022, they used a vulnerable legacy device to recover the Widevine L3 root of trust (CVE-2021-0639) by bypassing the obfuscation used to conceal cryptographic keys in software [23]. The authors exploited design decisions for supporting legacy devices; however, they only compromised low-quality content made available by Widevine L3. Patat et al. [22] raised privacy concerns over the device identifiers in the Widevine protocol, which can be used for third-party user tracking. Interestingly, their responsible disclosure revealed that Netflix was unaware that it does not protect audio tracks. The work investigated OTT apps and their compliance with Widevine guidelines on Android devices; it did not directly analyse the security of Widevine.

III. SECURITY DESIGN OF DRM SYSTEMS

A. Security Architecture of Google Widevine

Widevine is the primary DRM implementation for Android devices, which is deployed as a part of the operating system and it can also be deployed as a browser- or application-based implementation. It is used by various content providers, including Google Play, Amazon Prime, Netflix, YouTube, BBC,

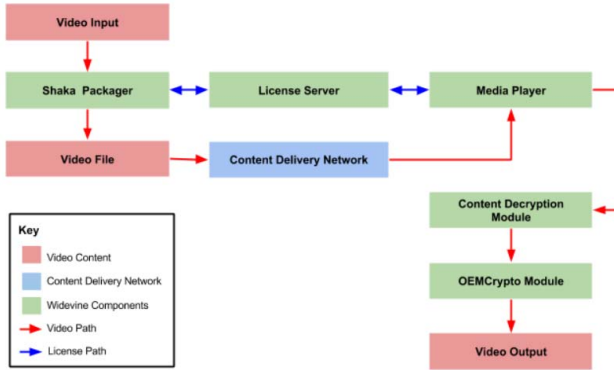


Fig. 1. Widevine DRM architecture [6].

Hulu, Disney+ and Spotify [24], [25]. Its main components comprise the Shaka packager, the Widevine license server, a supported media player, the Content Decryption Module and the OEMCrypto module [6]. Figures 1 and 2 demonstrate how the components interact together to deliver, decrypt, and playback content on user devices. We describe each area as follows.

1) *Security Levels*: Widevine uses three different security levels: L1 to L3. L1 is deemed the most secure and uses the TEE for both cryptography and media processing. The L2 security level uses the TEE for cryptography but not media processing, returning decrypted content to video decoder in the non-secure world for display. The L3 security level is the least secure and uses only software-based protection on devices that do not have a TEE [6]. All Widevine security levels require signed system images to be loaded on devices by a secure bootloader. The L1 and L2 security levels require the Widevine components implemented in the TEE to be included in the bootloader chain of trust; L1 also requires the components used for the secure video path to be included [26]. The L1 and L2 security levels also require the Widevine Keybox and the OEMCrypto API to be implemented on the device for key management and content decryption. The Keybox is encrypted with a unique factory-provisioned and device-specific key [26].

2) *Shaka Packager*: This is Widevine’s open-source packager that uses the Dynamic Adaptive Streaming over HTTP (DASH) protocol [6], [27] for optimised streaming over a dynamic bandwidth; HTTP Live Streaming (HLS) [28] is also supported. The DASH protocol converts content into fragmented MP4 files containing multiple versions of segments with the same length but different resolutions and bitrates. A media presentation description (MPD), or a manifest file, describes the different resolutions and bitrates for the content. The media player can use the manifest to request segments of different resolutions depending on the available bandwidth.

3) *Widevine Encryption*: The Shaka packager encrypts the media segments during packaging. Widevine supports the CENC (Common Encryption) [29] and recommends the ISOBMFF (ISO Base Media File Format) standard; FairPlay Streaming encryption [30] and the WEBM format are also

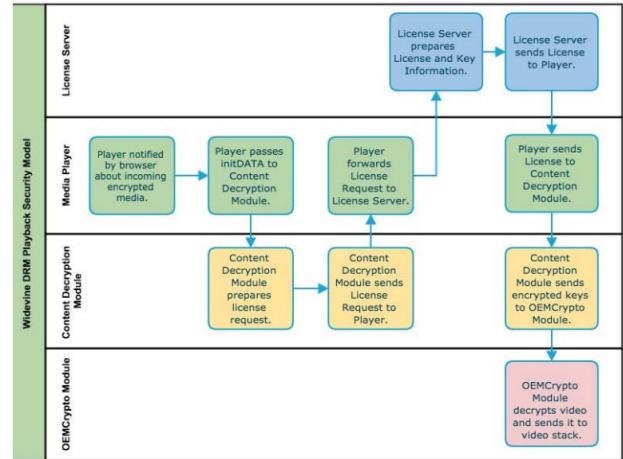


Fig. 2. Widevine playback security model [6].

supported. Widevine uses the Encrypted Media Extensions (EME) [31] and Media Source Extensions (MSE) [32] to allow playback of encrypted media.

4) *License Server*: Widevine provides a cloud-based license server that uses a request-response protocol over HTTPS to issue content licenses to the Content Decryption Module.

5) *Content Decryption Module (CDM)*: This is a proprietary component of Widevine on user devices, which creates encrypted license requests that are relayed to the license server by the media player application. The CDM receives encrypted license responses from the license server, through the media player, and uses the OEMCrypto module to decrypt them.

6) *OEMCrypto Module*: Decrypts the license responses to extract the encryption keys, which are used to decrypt the encrypted media content within the TEE of the user’s device. OEMCrypto is implemented by the system-on-chip vendor for Widevine L1, or by Widevine (Google, Inc.) for L3.

7) *Widevine Playback Security*: The Shaka packager produces encrypted media segments and a corresponding manifest. The encryption keys are sent to the license server and the encrypted media segments are streamed to the media player on user devices through the Content Delivery Network (CDN). The application receives a manifest and encrypted content from the CDN, extracts the initialisation data (InitData), and sends it to the media player. The media player forwards the InitData to the CDM to create a license request; the player sends the (encrypted) licence request produced by the CDM to the license server, which returns an (encrypted) license response back to the CDM. The license server authenticates the media player before responding to any requests. Next, the CDM sends the license response and media to the OEMCrypto module in the TEE to decrypt the content; OEMCrypto sends the decrypted content to the video stack for display.

B. Security Architecture of Apple FairPlay

Apple FairPlay, implemented in Apple iOS and tvOS, is based on the FairPlay Streaming (FPS) specification [30].

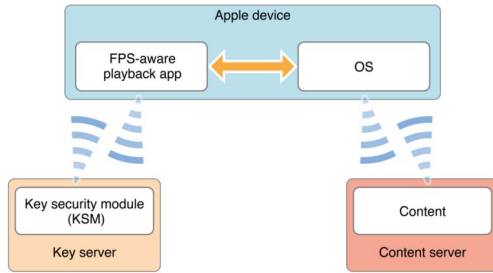


Fig. 3. Interactions between FPS and content provider's servers [33].

FPS provides secure delivery of encryption keys to devices and allows secure playback of encrypted content, which is provided using the HTTP Live Streaming (HLS) protocol [28]. We note that the Safari browser implementation on macOS uses FPS with an EME interface [33].

FPS requires content providers to implement their own content server, key (license) server and a player application; these must meet FPS requirements to be compatible with the FPS framework [33]. The content providers also need a public key infrastructure and public key certificates for RSA keys that are used to encrypt the license requests [33]. FPS authenticates the key server but requires content providers to source their own mechanisms for the player application to be authenticated by the key server [33]. FPS does not include content encryption and packaging mechanisms. Content providers are required to use the HLS streaming protocol for content delivery, and AES-128 CBC mode for encryption; video content is encrypted on a per frame basis, using a unique initialization vector (IV) for each encryption, and the audio content is fully encrypted on a per sample basis. Content providers handle the encryption keys and the IVs, as well as their associations with the content [33]. Figure 3 shows interactions between FPS on user devices and the content provider's key and content servers. FPS is integrated into the device OS and creates a server playback context (SPC) to request content encryption keys from the key server. The key server responds with the content key context (CKC). FPS extracts the encryption keys from the CKC and decrypts the content for display [33].

1) *FairPlay Playback Security*: Figure 4 shows the sequence of data flow during content playback. The player application sends a playback request to FPS on the device, which refers to the content manifest playlist on the content server and obtains the encrypted content. The FPS requests the player application for the encryption key, which responds with a request to create the SPC. FPS verifies the content provider's public key certificate and encrypts the SPC with the provider's RSA public key. FPS also uses a challenge-response mechanism and an undisclosed "symmetric algorithm" [33] to authenticate the key server. The player application creates a session with the key server and sends the encrypted SPC along with the content ID to the key server.

The key server decrypts the SPC using its RSA private key and creates a CKC. The CKC contains the IV, the encryption key and its expiration time, which is based on the

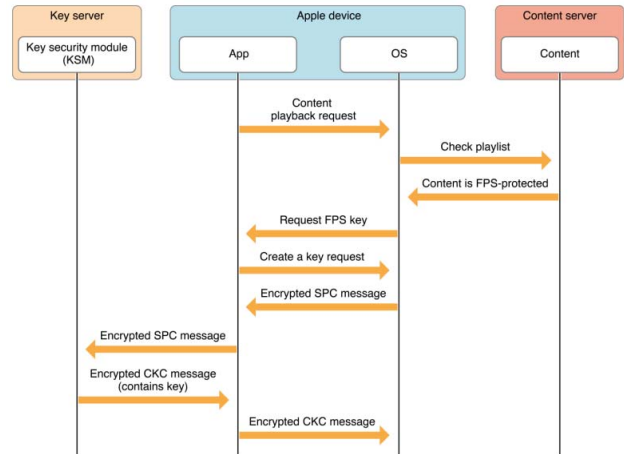


Fig. 4. Sequence of data flow in FPS [33].

player application's time reference instead of the server's time location. The key server encrypts the CKC with the session key contained in the SPC and sends it to the FPS on the device through the player application [33].

The FPS creates a play context and uses the iOS kernel to decrypt the CKC using the session key. The content encryption key and the IV are extracted from the CKC, and content is decrypted on a per frame basis. The decrypted content is sent directly to the decoder by the FPS and HDCP is enforced to display the decoded content on device display [33].

2) *Server Playback Context (SPC)*: The SPC created by FPS contains various security elements, including:

- A random AES-128 session key for single SPC use; a new session key is required even if another SPC contains the same content ID.
- An anti-replay seed to prevent replay of key server responses.
- A secure version of the content ID for server verification.
- Optional integrity verification to protect session key exchange.
- Player application certificate and a unique and anonymous device identifier; these are used by the key server to restrict the number of simultaneous playbacks for a user account.
- Player application's time reference, which is used by the key server to manage expiration of the encryption keys.
- Versions of security mechanisms on the key server that are supported by the player application.

3) *Key Expiration*: FPS uses two modes of key expiration: *video rental* and *secure lease*, which can be used together or independently. FPS does not decrypt the content and declines playback if the encryption key has expired for a video rental. For a mid-playback expiration, the current playback is allowed to continue and the expiration is enforced for the next playback session. The key server restricts the number of simultaneous playbacks by a single user account for a secure lease. The key server uses the anonymous device identifier to allocate a

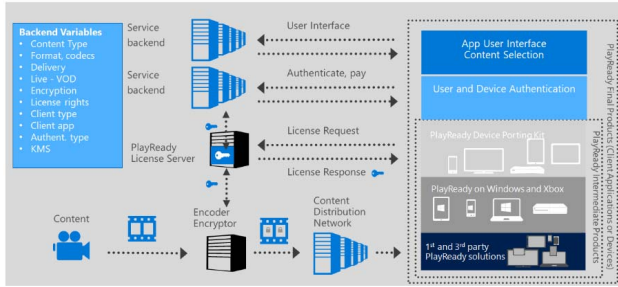


Fig. 5. Interactions between PlayReady client and servers [34].

slot to the device; if the lease is not renewed by the player application before its expiration then the slot is deallocated and FPS stops content playback on the device.

4) *Offline Playback*: The key server allows offline playback by indicating the persistent nature of the encryption key in the CKC. As such, FPS creates a persistent play context that is linked to the iOS device, which is stored in the device's file system with the expiration time of the encryption key.

5) *Rollback Prevention*: The SPC includes the key server versions supported by the player application. FPS uses the latest compatible version of security mechanisms available on the platform to protect encryption keys. An undisclosed "protection mechanism" [33] is used to prevent roll-back attacks to an older version of the security mechanisms.

6) *AirPlay Streaming*: AirPlay streaming from an iOS device to Apple TV uses a streamer context that supports secure streaming of encrypted content and CKC over AirPlay. According to Apple's documentation [33], FPS over AirPlay uses the same level of protection as FPS on iOS devices for the delivery of encryption keys and IVs.

C. Security Architecture of Microsoft PlayReady

PlayReady is widely deployed on Windows-based platforms, such as workstations and laptops, first developed in 2007 and currently integrated into Windows 10. PlayReady clients can be integrated into the OS or in the hardware. They can be implemented in media players on desktops, and as applications on mobile devices and other consumer devices, including smart TVs, set-top boxes and network receivers [34]. PlayReady serves more as an end-to-end framework than the DRM solutions discussed hitherto, which provide more rigid requirements. Indeed, it does not provide any mechanisms for content encryption, user authentication, and key management; content providers are required to implement these features. However, it does employ a key seed, a random 30-byte value that content providers can use to derive encryption keys or they can use their own key management system. PlayReady supports MPEG-DASH, HLS and Microsoft Smooth Streaming protocols, and CENC Common Encryption using either AES-CBC or AES-CTR modes of operation [35]. Content providers can use the PlayReady Server Software Development Kit (SDK) to build various servers that are compatible with PlayReady clients under a flexible model. These servers in-

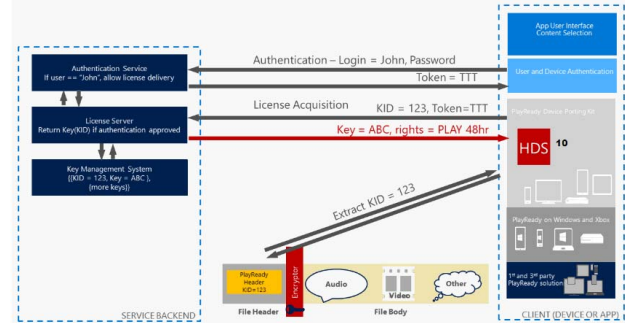


Fig. 6. Acquiring a PlayReady license [34].

clude content packaging servers, web servers, license servers, metering servers, and domain controllers.

1) *PlayReady Playback Security*: Interactions between PlayReady-enabled clients and servers are illustrated in Figure 5. Here, encrypted content is distributed through a CDN for streaming to clients. The packaging server includes the KeyID for the content encryption key in the content header, according to the PlayReady header specification.

The client streams the selected content from the CDN after browsing through the user interface. The client extracts the KeyID from the content header and obtains an authentication token from the authentication server. Figure 6 details the process flow to obtain a license from the license server. The client sends a license request to the license server including the KeyID and the authentication token. The license server verifies the authentication token with the authentication service; retrieves the encryption key from the key management system; creates a license containing the encryption key and content usage rights and restrictions; encrypts the license using the client's public key or the client domain public key; digitally signs the license with its own private signing key; and sends the signed license to the client. Multiple licenses and encryption keys can be sent in a single transaction and a license can also be granted for a domain instead of a single client [34], [36].

The PlayReady client verifies the License server's digital signature and certificate expiration, and then decrypts the license using its own private key [36]. The client decrypts the content using the encryption keys from the license and it enforces the license policy. PlayReady client implementations are expected to follow PlayReady's compliance and robustness rules, which govern how clients decrypt the content and require the clients to correctly process and apply the restrictions in the license [34].

2) *Offline Playback*: The license server can specify a particular license to be persistent or non-persistent. Persistent licenses are stored in non-volatile memory, such as the hard drive, and non-persistent licenses are stored in volatile memory so they are discarded when the current session ends [34], [37]. The HDS (Hashed Data Store), shown in figure 6, is a client-side license store provided by PlayReady.

3) *PlayReady Security Levels*: The security level of a PlayReady client is included in the client certificate that is embedded in the client during manufacturing [38]. A copy of the client certificate is included in the license request and the license server can select suitable licenses and license properties depending on the client security level. Moreover, the license server sets a minimum required security level for each license; the security level of a client must exceed this threshold to be able to use the license. The PlayReady SL150 security level is the least secure and it is only meant to be used for development and testing of clients as the content and encryption keys are not protected. The SL2000 security level is for commercial client implementations in hardened devices and applications, which use software-based protection mechanisms. The SL3000 security level offers the highest security and it is designed for commercial implementations on hardened devices that use hardware-based protection mechanisms, utilising the TEE for core functionalities.

D. Discussion

In Table I, we summarise the features provided by Widevine, FairPlay and PlayReady. We observe that none of these technologies provide all of the stated features based on publicly available information, placing an onus on content providers to implement any shortcomings. For instance, we note that no solution provides user authentication, while metering facilities are provided only by PlayReady. Conversely, all three DRM systems provide creation of license requests, license encryption, content decryption in the TEE of the device, and secure key storage for offline playback. License request encryption and a secure display are provided by Widevine and FairPlay, but not PlayReady. Player application, license server and content packaging are provided by Widevine and PlayReady but not FairPlay; PlayReady does not provide the actual mechanisms, it has the client compliance and robustness rules for player applications, and a server SDK to develop packaging and license servers. Content encryption is only provided by Widevine using the Shaka packager. Only FairPlay authenticates the license server and anonymously limits simultaneous playback sessions. Features unique to PlayReady include a key seed for key generation, and server SDK to develop web servers, domain-based license servers and metering servers.

IV. DRM EVALUATION METHODOLOGY AND CRITERIA

In this section, we distil the features and security properties of cloud-based DRM systems for mobile devices.²

A. DRM Security Features

Based on the previous sections, we propose a general feature set for classifying OTT DRM systems for cloud-based content

²We note that one possible approach is to employ a threat modelling methodology, such as STRIDE [39] or NIST SP 800-154 [40]. These approaches identify vulnerabilities and potential threats using data flow diagrams and attack trees; threats are then ranked and appropriate mitigations are determined. Generic threat modelling methodologies, however, do not specifically cover the features necessary for constructing secure DRM systems, thereby warranting a domain-specific evaluation.

TABLE I
FEATURE FULFILMENT OF WIDEVINE, FAIRPLAY, AND PLAYREADY.

Feature	Widevine	FairPlay	PlayReady
Key generation	—	—	Key seed
Packaging	Shaka packager	—	PlayReady server SDK
Encryption	Shaka packager	—	—
User authentication	—	—	—
License server authentication	—	RSA-based protocol	—
Create license request	Widevine CDM	FPS framework	PlayReady client
License request encryption	Encrypted	Encrypted with key server's RSA public key	—
License server	Widevine license server	—	Server SDK
License encryption	Encrypted	Session key and anti-replay	Client/domain public key, digitally signed
Content management	—	—	—
Content storage	—	—	PlayReady server SDK
Player application	Widevine media player	—	Client compliance rules
TEE component	OEMCrypto module	FPS framework	PlayReady client
Display	EME/MSE	HDCP enforced	—
Offline playback	Widevine key-box	iOS file system	HDS (Hashed Data Store)
Limit playback sessions	—	Anonymous device ID	—
Domain license	—	—	Server SDK
Metering	—	—	Server SDK

—: Specific measures not known from publicly available documentation.

delivery to mobile devices. The set assumes an end-to-end workflow of security processes and types of vulnerabilities specific to cloud-based mobile DRM systems, comprising 21 key security aspects listed under the following categories (also given in Table II). The following sections describe the various security aspects, numbered [SP1]–[SP21].

1) *Key management*: In mobile DRM systems, cryptographic keys and IVs must be generated securely, and keys must also be protected in transit (to the packaging and license servers) and storage (in a key database). The packaging server uses the keys to encrypt content, while the license server embeds the keys in an encrypted license which is transported to the TEE, where it is used to support content decryption.

[SP1] **Key generation**: a secure random number generator (RNG), or a key seed, is required to provide unpredictable values for generating keys and IVs with sufficient entropy.

[SP2] **Pre-license key transit**: secure channels are required to transfer keys to the packaging and license servers.

[SP3] **License transit**: keys are securely embedded in an encrypted and authenticated license.

[SP4] **Key database**: keys are securely held in a key database before encryption by the license server.

[SP5] **On-device key storage**: keys should be stored in TEE-

TABLE II
SECURITY PROPERTIES OF CLOUD-BASED MOBILE DRM SYSTEMS.

Category	Subcategories		Property
Key management	Key generation		[SP1]
	Key in transit	Pre-encryption	[SP2]
		Post-encryption	[SP3]
	Key storage		[SP4]
	Key disposal		[SP5]
Data encryption	Content encryption	Standardisation	[SP7]
		Re-keying	[SP8]
	Key encryption	Session security	[SP9]
		Public-key crypto.	[SP10]
		Pre-encryption	[SP11]
Content management	Content storage	Post-encryption	[SP12]
		Pre-encryption	[SP13]
	Content in transit	Post-encryption	[SP14]
		Post-decryption	[SP15]
		Manifest security	[SP16]
Authentication	User authentication	[SP17]	
	Server authentication	[SP18]	
DDoS	DDoS mitigations	[SP19]	
TEE	Secure TEE usage	[SP20]	
Quantum security	Post-quantum mitigations	[SP21]	

reserved memory regions on user devices.

[SP6] **Key disposal:** keys should be disposed securely to avoid content decryption once the license has expired.

2) *Encryption:* DRM systems commonly employ symmetric keys for content encryption, which are protected using key encryption keys. Public key cryptography is used for license encryption, containing content encryption keys, and for exchanging (symmetric) session keys.

[SP7] **Suitable encryption algorithms:** Standard symmetric algorithms are required for secure and efficient content encryption, e.g. 128-bit AES.

[SP8] **Content re-keying:** refers to encrypting different parts of the content with different keys to limit the impact of a key compromise, and to allow different license rules to be applied to each part.

[SP9] **Session key management:** session keys must be generated, exchanged, and stored securely.

[SP10] **Public key management:** public-key pairs must be generated and managed securely using appropriate standards and modules, e.g. NIST FIPS 140-3.

3) *Content management:* Content and the manifest files require secure storage and transit. Content is transferred from the content database to the packaging server to be encrypted before it is stored in a content server and cached by the CDN. The encrypted content is streamed to user devices, decrypted in the TEE, and sent to the device display for viewing.

[SP11] **Pre-packaged content storage:** Content requires secure storage in the content database to prevent attackers from obtaining plaintext content before DRM protection is applied.

[SP12] **Packaged content storage:** Packaging servers generally apply encryption during content packaging.

[SP13] **Pre-packaged content transit:** content requires secure transfer from the content database to the pack-

aging server where it is encrypted.

[SP14] **Packaged content transit:** Packaged content is encrypted for transit.

[SP15] **Secure video path:** decrypted content is displayed directly from the TEE to the display device over a trusted path, e.g. to thwart screen scraping attacks.

[SP16] **Secure manifest:** manifest files, containing URLs for different content segments, must be transmitted securely. A MITM attack, where the CDN is impersonated, could potentially replace these URLs with malicious ones.

4) *Authentication:* Mutual authentication between users and the servers is required to ensure only authorised users can access the content, and to prevent false server attacks, where an attacker-controlled server impersonates a legitimate server to launch attacks on user devices.

[SP17] **User authentication:** users should be authenticated by the CDN and the license server.

[SP18] **Server authentication:** CDN and license servers should be authenticated by user devices.

5) *DDoS:* Compromise the availability of DRM-protected content by consuming server and network resources.

[SP19] **DDoS vulnerabilities:** CDN and license servers could potentially be vulnerable to DDoS attacks.

6) *TEEs:* DRM protection relies on the security provided by TEEs on user devices. Vulnerabilities in TEEs could be leveraged to compromise DRM-protected content.

[SP20] **TEE attack resistance:** TEEs share hardware resources, such as the processor and memory, with the native OS on user devices. A TEE must be able to thwart relevant attacks, such as side-channel and fault injection attacks.

7) *Quantum security:* Suitable post-quantum algorithms must be used to maintain the security of DRM systems against quantum adversaries, which can undermine public- and symmetric-key cryptographic algorithms.

[SP21] **Post-quantum security:** the outcomes of NIST's Post-quantum Cryptography (PQC) standardisation process must be implemented.

V. SECURITY ANALYSIS OF DRM TECHNOLOGIES

In this section, we consider the highest security levels of Widevine, FairPlay and PlayReady, e.g. Widevine L1 and PlayReady SL3000 levels, which utilise hardware-assisted mechanisms for content protection. Widevine, FairPlay and PlayReady provide a subset of the security features (§III-D) required for a complete DRM implementation; content providers must implement the remaining features securely (§VI-B) for the DRM protection to be effective. We analyse each technology in turn below, and compare them in Table III.

Widevine. We observe that Widevine poses security risks due to the way in which manifest files could be used as an attack vector. Content encryption keys in Widevine are at a low risk during license transit, as the license is encrypted and sent over HTTPS. Risks to the encryption algorithm and

packaged content, both in storage and in transit, are low as the Shaka packager uses CENC encryption to encrypt the content. The risk level of TEE-related vulnerabilities is low due to the specialist equipment and skills required for these attacks. The risk from post-quantum attacks is present, albeit low given the absence of large-scale quantum computers. New post-quantum public key algorithms will be required to replace the current ones, and the current symmetric key lengths will not offer sufficient protection from quantum attacks.

FairPlay. Like Widevine, FairPlay poses a high risk to user devices due to insecure manifest files. FairPlay requires the license to be encrypted with a unique AES-128 session key; generated by the FairPlay client; protected by the license server’s public key; and also protected from replay attacks by an anti-replay seed. Therefore, the security risk to the license during its transit is low. Content providers are required to encrypt the packaged content using AES-128 CBC mode and a unique IV for each encryption. Unlike other TEEs, Apple’s Secure Enclave Processor (SEP) [41] does not share hardware resources with host OS on the device and protects against some side-channel attacks, including attacks based on caching differences. It is unclear whether the SEP is used by FairPlay, although we recommend that it should be used for additional hardware-assisted protection. We note that FairPlay is also vulnerable to quantum adversaries.

PlayReady. The PlayReady manifest files also pose a high risk to user devices. The license is encrypted with the PlayReady client public key or the client domain public key, and it is digitally signed by the license server; hence, the license in transit has a low risk level. The PlayReady client verifies the digital signatures and decrypts the license using its private key; the level of risk to the public-private key pairs is low. PlayReady supports CENC Common Encryption with AES-CBC and AES-CTR for content encryption, resulting in a low risk to packaged content in storage and in transit. There is a risk to PlayReady from TEE and quantum adversaries.

VI. MITIGATIONS

Following the previous analysis, this section proposes mitigations for the identified issues.

A. Malicious Manifest Attacks

There is a high risk of potential attacks on user devices from insecure manifest files used in Widevine, FairPlay and PlayReady. The manifest file is sent as plaintext, and it contains DRM client initialization data, KeyIDs, and download locations of content segments for a content stream. Manifest files are sent as plaintext to allow DRM clients to stream relevant content segments defined in the manifest, and to initialize a license request using the KeyIDs contained in the manifest. Moreover, an encrypted manifest can be replaced with a malicious one by an attacker without requiring decryption. Data origin authentication of the manifest, as opposed to confidentiality, is required to mitigate these attacks. Content providers do not share long-term secrets with DRM clients to allow symmetric cryptography, such as MACs (Message

TABLE III
SECURITY COMPARISON OF WIDEVINE, FAIRPLAY, AND PLAYREADY.

Security Property	Technology		
	Widevine	FairPlay	PlayReady
[SP1] Key generation	*	*	*
[SP2] Pre-license key transit	*	*	*
[SP3] License transit	✓,▼	✓,▼	✓,▼
[SP4] Key database	✚	✚	✚
[SP5] Key storage on device	*	*	*
[SP6] Key disposal	✚	*	*
[SP7] Suitable encryption algorithms	✓,▼	✓,*	✓*
[SP8] Content re-keying	✚	✚	✚
[SP9] Session key management	*	✓,▼	✚
[SP10] Public key management	*	✚	✓,▼
[SP11] Pre-packaged content storage	✚	✚	✚
[SP12] Packaged content storage	✓,▼	✓,▼	✓,▼
[SP13] Pre-packaged content transit	✚	✚	✚
[SP14] Packaged content transit	✓,▼	✓,▼	✓,▼
[SP15] Secure video path	☆	*	*
[SP16] Secure manifest	☆,▲	☆,▲	☆,▲
[SP17] User authentication	✚	✚	✚
[SP18] Server authentication	✚	*	✚
[SP19] DDoS security	*	✚	✚
[SP20] TEE security	✓,▼	*	✓,▼
[SP21] Quantum security	☆,▲	☆,▲	☆,▲

✓: Provided natively. ✚: Not provided. ☆: Somewhat provided. *: Insufficient information available. ▼: Contains security weaknesses of low concern. ▲: Weaknesses of high concern.

Authentication Codes), to be used. Instead, content providers can digitally sign the manifest using a private signing key to provide non-repudiation, a stronger notion of data origin authentication, which may be accompanied by mutual authentication and TEE-based attestation (e.g. see [42]); DRM clients can use the content provider’s public verification key to verify the digital signatures before processing the manifest.

B. Secure Configuration

The overall security of a DRM system depends on its complete configuration; Widevine, FairPlay and PlayReady provide a subset of the features required for DRM protection (§V), leaving the remaining features to be implemented by content providers. An insecure implementation of these features could potentially introduce vulnerabilities in a DRM system. DRM systems use encryption to protect content and the content encryption keys. A secure implementation of standards-based encryption with correct key management should be used to ensure confidentiality of the content and the keys.

C. Authenticated Encryption

Similar to the malicious-manifest attacks (§VI-A), the encrypted content segments and licenses can be replaced by malicious payloads or malware in a MITM attack. Data origin authentication of the content segments and licenses is required to prevent these attacks. Confidentiality of content segments and licenses is already provided by the packaging and license servers, which can be combined with data origin authentication protocols. There are several modes of operation for authenticated encryption that securely combine data origin authentication and confidentiality into a single primitive, such as CCM mode, EAX mode, OCB mode and GCM mode.

Encryption of content segments by the packaging server, and licenses by the license server, should be replaced with authenticated encryption to protect the content and licenses in transit, and to protect user devices from MITM attacks described above.

D. Trusted Execution Environments

DRM systems use TEEs in user devices to securely decrypt and store content, and to manage cryptographic keys. In recent years, TEEs have been subjected to a myriad of side-channel attacks, particularly micro-architectural ones, as well as physical fault injection attacks. Some side-channel attacks, such as attacks based on cache-timing differences, are caused by hardware (processor and memory) sharing between the TEE and the host OS on user devices; these can be mitigated by using an isolated processor that is not shared with the main OS. Using an isolated processor for TEE could prevent cache-based and transient execution attacks that exploit cache memory sharing between the ‘normal’ and the ‘secure’ worlds to leak sensitive TEE data [43]. However, it does not prevent fault injection and other physical attacks, like those based on electromagnetic and power analysis. For this, tamper-resistant packaging is a widely suggested countermeasure [41].

E. Post-quantum Mitigations

Quantum computers are known to fatally undermine traditional public key algorithms, e.g. RSA and ECDSA, which has prompted the NIST PQC standardisation competition [44]. To make DRM systems resistant to post-quantum attacks, TEEs should be adapted to the new PQC standards. The various security services provided by TEEs, such as authenticated boot, remote attestation and key management, lack widespread resistance to quantum adversaries, as do the supporting protocols for license distribution in existing DRM systems. As such, we strongly urge developers to expedite the transition to post-quantum algorithms for license distribution and authentication in particular.

F. DRM Standardisation

The fragmentation in the current DRM landscape and lack of standardisation is a source of potential vulnerabilities in DRM systems. The DRM systems analysed in this work use different formats, protocols and APIs. The complexity of the ecosystem, with different implementations using numerous components, may inadvertently lead to insecure configurations. Additionally, various standardised formats and APIs for DRM systems have been developed, including CENC [29], CMAF (Common Media Application Format) [45], CPIX (Content Protection Information Exchange) [46] and EME APIs [31]. This paper has shown that some of these standards are not used by any of the DRM systems, such as the CPIX, a key exchange data format for confidential and authenticated key delivery in DRM systems. Further DRM standardisation could potentially improve interoperability, reduce fragmentation and complexity, and minimise vulnerabilities.

VII. FUTURE DIRECTIONS

The centralised servers in cloud-based DRM can be a single point of failure and vulnerable to attacks. In a centralised DRM model, content owners do not have direct access to consumers and must sell their content to large content providers and media platforms. Furthermore, the copyright and license transaction information in current DRM systems lacks transparency, and authentication is inefficient, requiring multiple interactions. To overcome these challenges, Zhang and Zhao [47] have proposed a decentralised DRM model with collective maintenance using a blockchain-based license structure. Licenses are issued automatically using smart contracts, offering reliability of license transactions, and the license transaction information is recorded on the blockchain to make it transparent and secure. The authorisation of licenses is scalable and does not need interactive communication between content owners and consumers. Content owners can set flexible pricing rules for content usage and deliver DRM-protected content directly to the consumers.

Rana and Mishra [48] have proposed an authenticated key distribution framework for IoT enabled DRM systems, allowing authorised users to consume content through smart devices. The framework uses biometric data and three-factor authentication to verify users; allows secure and anonymous communication; but does not preserve the secrecy of consumer preferences. Data centres consume a vast amount of power to meet consumer demand for low-latency access to content. Rana et al. [49] have discussed the security of authenticated DRM key distribution frameworks, which support efficient communication, computing and mutual authentication to conserve energy consumption. Lastly, we draw attention to Li et al. [50], who explored using federated cloud and edge computing for cost-efficient streaming, and developing load distribution mechanisms to reduce latency.

VIII. CONCLUSION

In this paper, we presented the first security evaluation of three major DRM systems under the OTT content delivery model with respect to mobile devices. The security analysis identified attack vectors in all aspects of key management, content management, encryption and access control in each of the DRM systems, drawing from parallel advancements in the state of the art (e.g. TEE micro-architectural vulnerabilities). We showed how all three DRM systems are vulnerable to post-quantum attacks, and issues pertaining to manifest security. Moreover, TEEs are vulnerable to a range of side-channel and physical fault injection attacks that could compromise DRM-protected content. However, it is worth being cognisant of the high levels of expertise and specialist equipment needed to mount such attacks successfully. Another source of potential vulnerabilities is the complexity and fragmentation in the DRM landscape, which could be remedied by appropriate standardisation efforts to reduce the fragmentation and complexity of the ecosystem. Some of the current trends and future directions of research relevant to cloud-based DRM include a decentralised DRM model based on blockchain

technology; IoT-enabled DRM systems; improving the power and computational efficiency of DRM data centres; and using federated cloud and edge computing with load distribution mechanisms for cost-efficient and low latency streaming.

REFERENCES

- [1] "The Global Internet Phenomena Report," Sandvine – Phenomena, 2022, <https://www.sandvine.com/phenomena>.
- [2] "Over The Top (OTT) Service Market Size," Fortune Business Insights, August 2022, <https://www.fortunebusinessinsights.com/industry-reports/over-the-top-services-market-100506>.
- [3] M. Fachot, "Cyber security — a priority for broadcasters and media companies," IEC (International Electrotechnical Commission) e-tech, 2019, <https://etech.iec.ch/issue/2019-01/cyber-security-a-priority-for-broadcasters-and-media-companies>.
- [4] "Operating System Market Share Worldwide," StatCounter Global Stats, 2022, <https://gs.statcounter.com/os-market-share>.
- [5] Google, Inc., "Widevine," <https://www.widevine.com>.
- [6] —, "Widevine DRM Architecture Overview - version 1.2," 2017, https://www.whymatematica.com/wp-content/uploads/2018/08/Widevine_DRM_Architecture_Overview.pdf.
- [7] "17 U.S. Code sec. 1201 - Circumvention of copyright protection systems," Legal Information Institute (LII), 1999, <https://www.law.cornell.edu/uscode/text/17/1201>.
- [8] "Digital Millennium Copyright Act (DMCA)," Electronic Frontier Foundation, <https://www.eff.org/issues/dmca>.
- [9] N. Ferguson, "Censorship in action: Why I don't publish my HDCP results," 2001, <https://web.archive.org/web/20120220014712/http://www.macfergus.com/niels/dmca/cia.html>.
- [10] S. Crosby, I. Goldberg, R. Johnson, D. Song, and D. Wagner, "A Cryptanalysis of the High-Bandwidth Digital Content Protection System," in *Security and Privacy in Digital Rights Management*. Springer, 2002.
- [11] "List of 1201 threats," Electronic Frontier Foundation, 2016, <https://www.eff.org/document/list-1201-threats>.
- [12] "Chinese app store offers pirated iOS apps without the need to jailbreak," ExtremeTech, 2013, <https://www.extremetech.com/mobile/153849-chinese-app-store-offers-pirated-ios-apps-without-the-need-to-jailbreak>.
- [13] S. Y. Chau, B. Wang, J. Wang, O. Chowdhury, A. Kate, and N. Li, "Why Johnny Can't Make Money With His Contents: Pitfalls of Designing and Implementing Content Delivery Apps," in *34th Annual Computer Security Applications Conference*. ACM, 2018, pp. 236–251.
- [14] R. Wang, Y. Shoshitaishvili, C. Kruegel, and G. Vigna, "Steal this movie: Automatically bypassing DRM protection in streaming media services," in *USENIX Security Symposium*, 2013, pp. 687–702.
- [15] C. D'Orazio and K.-K. R. Choo, "An adversary model to evaluate DRM protection of video contents on iOS devices," *Computers & Security*, vol. 56, pp. 94–110, 2016.
- [16] C. Xiao, "AceDeceiver: First iOS Trojan Exploiting Apple DRM Design Flaws to Infect Any iOS Device," Unit 42, 2016, <https://unit42.paloaltonetworks.com/acedeceiver-first-ios-trojan-exploiting-apple-drm-design-flaws-to-infect-any-ios-device/>.
- [17] R. Chirgwin, "Google's Widevine DRM doesn't quite manage," The Register, 2016, https://www.theregister.com/2016/06/28/google_widevine_drm/.
- [18] D. Buchanan, "Breaking Widevine L3 DRM," Twitter, <https://twitter.com/David3141593/status/1080606827384131590>.
- [19] "Google's Widevine L3 DRM cracked by a security researcher," Cyware Labs, 2019, <https://cyware.com/news/googles-widevine-l3-drm-cracked-by-a-security-researcher-eaed58ca>.
- [20] "DMCA, Widevine L3 decryptor," GitHub, 2020, <https://github.com/github/dmca/blob/7f665a41abb77dae9927cdc62c09edcec822d4cc/2020/11/2020-11-09-Google.md>.
- [21] Q. Zhao, "Wideshears: Investigating and Breaking Widevine on QTEE," <https://i.blackhat.com/asia-21/Thursday-Handouts/as-21-Zhao-Wideshears-Investigating-And-Breaking-Widevine-On-QTEE-wp.pdf>.
- [22] G. Patat, M. Sabt, and P.-A. Fouque, "WideLeak: How Over-the-Top Platforms Fail in Android," in *52nd IEEE/IFIP Int'l Conf. on Dependable Systems and Networks*, 2022, pp. 501–508.
- [23] —, "Exploring Widevine for Fun and Profit." [Online]. Available: <http://arxiv.org/abs/2204.09298>
- [24] "Google Widevine Digital Rights Management," Google Developers, 2021, <https://developers.google.com/widevine/drm/overview>.
- [25] D. Kim, "What is Google Widevine DRM and How It Works," PallyCon, 2019, <https://pallycon.com/blog/what-is-google-widevine-drm-and-how-does-it-work/>.
- [26] A. Lee, "Widevine DRM on devices - Getting started - version 1.5," 2018, https://web.archive.org/web/20190504082905/https://storage.googleapis.com/wvdocs/Widevine_DRM_Getting_Started_Devices.pdf.
- [27] T. Stockhammer, "Dynamic Adaptive Streaming over HTTP - Design Principles and Standards," 2010, https://www.w3.org/2010/11/web-and-tv/papers/webtv2_submission_64.pdf.
- [28] "HTTP Live Streaming," Apple Developer Documentation, https://developer.apple.com/documentation/http_live_streaming.
- [29] "ISO/IEC 23001-7:2016 - Information technology - MPEG systems technologies - Part 7: Common encryption in ISO base media file format files," 2016, <https://www.iso.org/standard/68042.html>.
- [30] "FairPlay Streaming - HTTP Live Streaming," Apple Developer, <https://developer.apple.com/streaming/fps/>.
- [31] D. Dorwin, J. Smith, M. Watson, and A. Bateman, "Encrypted Media Extensions," W3C Recommendation, 2017.
- [32] M. Wolentz and M. Watson, "Media Source Extensions™," W3C Working Draft, <https://www.w3.org/TR/media-source/>.
- [33] "FairPlay Streaming Overview," 2016, <https://developer.apple.com/streaming/fps/FairPlayStreamingOverview.pdf>.
- [34] M. Sasouvanh and R. L. Franc, "PlayReady Overview: Simple End to End System," Microsoft Documentation. [Online]. Available: <https://docs.microsoft.com/en-us/playready/overview/simple-end-to-end-system>
- [35] K. R. Vijayanagar, "What is Microsoft PlayReady DRM and How Does It Work?," OTTVerse, 2020, <https://ottverse.com/microsoft-playready-how-does-it-work/>.
- [36] M. Sasouvanh and D. Klopfenstein, "PlayReady Overview: License Acquisition," Microsoft Documentation, 2018, <https://docs.microsoft.com/en-us/playready/overview/license-acquisition>.
- [37] —, "PlayReady Overview: License Persistence," Microsoft Documentation, 2018, <https://docs.microsoft.com/en-us/playready/overview/license-persistence>.
- [38] —, "PlayReady Overview: Security Level," Microsoft Documentation, <https://docs.microsoft.com/en-us/playready/overview/security-level>.
- [39] "The STRIDE Threat Model," Microsoft documentation, 2009, [https://learn.microsoft.com/en-us/previous-versions/commerce-server/ee823878\(v=cs.20\)](https://learn.microsoft.com/en-us/previous-versions/commerce-server/ee823878(v=cs.20)).
- [40] M. Souppaya and K. Scarfone, "Draft NIST Special Publication 800-154: Guide to Data-Centric System Threat Modeling," https://csrc.nist.gov/files/pubs/sp/800/154/fpd/docs/sp800_154_draft.pdf.
- [41] C. Shepherd, K. Markantonakis, N. van Heijningen, D. Aboukassimi, C. Gaine, T. Heckmann, and D. Naccache, "Physical fault injection and side-channel attacks on mobile devices: A comprehensive analysis," *Computers & Security*, vol. 111, p. 102471, 2021.
- [42] C. Shepherd, R. N. Akram, and K. Markantonakis, "Remote credential management with mutual attestation for trusted execution environments," in *12th IFIP International Conference on Information Security Theory and Practice*. Springer, 2019, pp. 157–173.
- [43] W. Xiong and J. Szefer, "Survey of transient execution attacks and their mitigations," *ACM Computing Surveys*, vol. 54, no. 3, pp. 1–36, 2021.
- [44] "Post-Quantum Cryptography Standardization," NIST CSRC - Post-Quantum Cryptography (PQC), <https://csrc.nist.gov/Projects/post-quantum-cryptography/post-quantum-cryptography-standardization>.
- [45] "ISO/IEC 23000-19:2020 - Information technology - Multimedia application format (MPEG-A) - Part 19: Common media application format (CMAF) for segmented media."
- [46] "DASH-IF Implementation Guidelines: Content Protection Information Exchange Format (CPIX)," DASH Industry Forum, <https://dashif.org/docs/CPIX2.1/HTML/Index.html>.
- [47] Z. Zhang and L. Zhao, "A Design of Digital Rights Management Mechanism Based on Blockchain Technology," in *1st Int'l Conf. on Blockchain*. Springer-Verlag, pp. 32–46.
- [48] S. Rana and D. Mishra, "Secure and ubiquitous authenticated content distribution framework for IoT enabled DRM system," *Multimedia Tools and Applications*, vol. 79, no. 27, pp. 20319–20341, 2020.
- [49] S. Rana, M. Obaidat, D. Mishra, S. Mukhopadhyay, and B. Sadoun, "Computational Efficient Authenticated Digital Content Distribution Frameworks for DRM Systems: Review and Outlook," *IEEE Systems Journal*, 2020.
- [50] X. Li, M. Darwich, M. Bayoumi, and M. A. Salehi, "Cloud-Based Video Streaming Services: A Survey," *arXiv preprint arXiv:2011.14976*, 2020, <http://arxiv.org/abs/2011.14976>.