streaming video alliance

TECHNICAL DOCUMENT

Forensic Watermarking Implementation Considerations for Streaming Media

Created and Approved by the Streaming Video Alliance
July 19, 2018

WORKING GROUP:
Privacy and Protection

GROUP CHAIRS:
Brian Stevenson, Nokia
Mike Wilkinson, NBCUniversal

1. Contributors

The following people contributed to this document:

- Niels Thorwirth (Verimatrix)
- Erik Hietbrink (Irdeto)
- Jaap Haitsma (NexGuard Labs France S.A.S.)
- Gwenaël Doërr (ContentArmor S.A.S.)
- Glenn Deen (Comcast)
- Mike Wilkinson (NBCUniversal)
- Robin Wilson (BAMTECH Media)
- Brian Stevenson (Nokia UK Ltd.)
- Christopher White (Friend MTS Limited)

2. Abstract

The piracy of online video content is a significant issue for content owners and distributors. This document explains the technology of watermarking and how it can be employed, in various methods, to secure video against theft.

The Streaming Video Alliance (the Alliance) is an industry forum open to all companies from all sectors of the online video value chain. The Alliance focuses on the ecosystem, architecture and best practices needed to support the future of online video.

Membership is comprised of industry leaders from the entire online video ecosystem, including content providers, service providers, commercial CDNs and streaming video technology providers.

Notice:

This document has been created by the Streaming Video Alliance. It is offered to the Alliance Membership solely as a basis for agreement and is not a binding proposal on the companies listed as resources above. The Alliance reserves the rights to at any time add, amend or withdraw statements contained herein. Nothing in this document is in any way binding on the Alliance or any of its members. The user's attention is called to the possibility that implementation of the Alliance agreement contained herein may require the use of inventions covered by the patent rights held by third parties. By publication of this Alliance document, the Alliance makes no representation or warranty whatsoever, whether expressed or implied, that implementation of the specification will not infringe any third party rights, nor does the Alliance make any representation or warranty whatsoever, whether expressed or implied, with respect to any claim that has been or may be asserted by any third party, the validity of any patent rights related to any such claim, or the extent to which a license to use any such rights may or may not be available or the terms hereof.

©2018 Streaming Video Alliance

This document and translation of it may be copied and furnished to others, and derivative works that comment on or otherwise explain it or assisting its implementation may be prepared, copied, published and distributed, in whole or in part, without restriction other than the following, (1) the above copyright notice and this paragraph must be included on all such copies and derivative works, and (2) this document itself may not be modified in any way, such as by removing the copyright notice or references to the Alliance, except as needed for the purpose of developing Alliance Specifications.

By downloading, copying, or using this document in any manner, the user consents to the terms and conditions of this notice. Unless the terms and conditions of this notice are breached by the user, the limited permissions granted above are perpetual and will not be revoked by the Alliance or its successors or assigns.

This document and the information contained herein is provided on an "AS IS" bases and THE ALLIANCE DISCLAIMS ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY, TITLE OR FITNESS FOR A PARTICULAR PURPOSE.

streaming video alliance

Table of Contents

3. Overview

The protection of commercial video content from piracy is obviously important to content owners and is typically a requirement for distributors in the licensing terms. Watermarking is a useful tool to combat theft as it permits tracing stolen copies back to the source of the leak. Improvements in watermarking technology have advanced to the point that it is now practical to embed watermarks containing distributor identifiers, content identifiers, and subscriber identifiers for nearly all types of content including live content, delivered to individual subscribers.

The implementation details to embed, and later extract, watermarks vary by watermarking technology provider. They share fundamental similarities to embed data that can later be extracted by a detector. This document examines the use of watermarks for forensic identification of the sources of pirated content. Specifically, how different schemes can integrate with media distribution workflows to embed marks and how extractions can be employed in source leak identification processes.

Various watermark embedding approaches, such as server-side and client-side, are examined in this document along with the considerations and challenges of each approach. Likewise, watermarking in both on-demand and live streamed distribution scenarios is explored. For completeness, this document also discusses watermarking in other parts of the distribution chain including review copies and digital cinema.

This is not a marketing document to promote any specific approach. Instead, it aims to educate readers on different watermarking approaches to help them identify the one that best fits their specific media workflows and business needs.

The reader is referred to the companion document on watermarking from the UHD Forum[1] for the technical details behind watermarking for streaming as this document does not explore the technologies behind how marks are embedded or protected in video

[1] Ultra HD Forum: Phase A Guidelines, April 24, 2017,Revision: 1.3 available at https://ultrahdforum.org/wp-content/uploads/Ultra-HD-Forum-Guidelines-v1.3.pdf

4. Definitions

Referenced Definitions

For the definition of the following terms, see the security chapter of the UHDF companion document (see section 13 *Additional Information* Sources):

Forensic Watermarking

Set of Variants

Variant

Variant Sequence Generator

Variant Sequence

Watermark Identifier

Unique Definitions

Operator Mark—A watermark payload intended for distribution to, and identification of, a unique operator's content. Also known as a Distributor or Distribution Mark.

Session Mark—A watermark payload intended for distribution to, and identification of, a unique playback session. This identifier may note a unique session, subscriber, or playback device. If the mark is used to identify an individual subscriber, it is also known as a Subscriber Mark.

Mitigation—The action of eliminating the distribution of pirated content. This may include, but is not limited to, legal action and disablement of the source stream.

Transaction Database—Collection of all unique embedded payloads with associated information about the transaction.

Actor Definitions

Content Owner—The owner of the intellectual property rights for the content being distributed.

Operator—A video provider delivering content from multiple Content Owners. The Operator is obligated to deliver content in compliance with the Content Owners' distribution agreements.

Watermarking Supplier—A provider of watermark embedding and extraction technologies that may run on the site of the operator or content owner.

Piracy Monitoring—A service for the discovery of pirated content. This service may or may not be run by the Watermarking Supplier, however, the service is dependent upon the Watermarking Supplier's technology. In the context of forensic watermarking, Piracy Monitoring will trigger the extraction of the mark.

Pirate Source—A website or service designed for the illicit or illegal distribution of content.

Client Device—The end-user consumption device of the Operator's network. This device may be Consumer Provisioned Equipment (CPE) or Customer Owned and Managed (COAM).

Acronyms

ABR	Adaptive Bit Rate
ACR	Automated Content Recognition
API	Application Programming Interface

AVC	Advanced Video Coding
AR	Aspect Ratio
CAS	Conditional Access System
CDN	Content Delivery Network
cDVR	Cloud Digital Video Recorder
COAM	Customer Owned and Managed (device)
CPE	Consumer Provisioned Equipment
DA-AD	Digital-to-Analog Analog-to-Digital (conversion)
DRM	Digital Rights Management
ERW	Early Release Window
HD	High Definition (1080p and below until SD)
HDR	High Dynamic Range
HDCP	High-bandwidth Digital Content Protection
HDMI	High Definition Media Interface
HEVC	High Efficiency Video Coding
HFR	High Frame Rate
ID	Identifier
ISP	Internet Service Provider
OS	Operating System
PII	Personally Identifiable Information
SaaS	Software as a Service
SD	Standard Definition (570p and below)
SDR	Standard Dynamic Range
TEE	Trusted Execution Environment
TSTV	Time-shifted Television
UGC	User Generated Content
UHD	Ultra-High Definition (4K and below until UHD)
VBR	Variable Bit Rate
VOD	Video on Demand

5. Introduction

Streamed video, both live and on-demand, is growing in popularity with viewers. Thanks to a wide array of networks and consumer devices, it is becoming accessible anywhere and at any time across a multitude of devices. Video also comes in many forms, from user generated content (UGC) videos on social media outlets to ultra-high quality and high-definition professional content, including world class sporting events and first run movies delivered by licensed Internet streaming services.

Although the Internet serves as a great delivery mechanism for streamed video, it is also a popular distribution channel for pirated content. To safeguard revenues both Content Owners and Operators have interest to protect licensed video content from piracy. In the licensing terms, Content Protection is also a typical requirement from Content Owners to Operators

The ability to prevent piracy, take down illegal content, and act against illegal sources are key objectives of Content Protection. Meeting these objectives requires the use of a variety of technologies including watermarking and others such as Digital Rights Management (DRM), fingerprinting, and cryptography.

This document focuses on the application of watermarks in online streamed video delivery, while references to information sources covering the other techniques can be found in Section 13, *Additional Information Sources.*

6. Technology Introduction and Background

Watermarking is the technique of embedding data into either the audio[2] or video portions of an asset that can be reliably extracted, even if the asset has been modified. In the event of modification, the watermark is designed to travel along with the asset without itself being modified.

Purpose of Forensic Watermarking

In the context of this document Forensic Watermarking is intended to provide a means to identify the source of leaked content at the distributor level, at the more granular device level, or even the subscriber level. When leaked content is found on piracy sites, it can be analyzed for any embedded marks which, when extracted, can provide chain of custody through to device/subscriber identification depending on the marks present. This information can then be employed by investigative teams to locate the leak source and act to stop future piracy.

Application Areas

Forensic Watermarking will provide anti-piracy functionality, however, there are many uses where watermarking can add value in the content delivery chain:

Digital Cinema—Digital Cinema offers the best quality and earliest availability to video content. This makes it highly susceptible to theft. The quality is so good it is possible to create a high-quality artifact, even if recorded in the theatre. Digital Cinema Watermarks do not apply to individual users, but can be used to identify the theater, date, and time where the content was screened.

Screeners—Pre-release content for press and other viewers associated with the studio can also represent a source of illegal distribution. This is a very critical security breech as screeners are normally provided access to the content prior to the start of distribution windows. Watermarks can be used to identify the viewer that received the screening copy.

Physical Media—In this case, watermarking is attached to the physical media such as a Blu-ray disc, that later can be used to determine the playback device used for capture and illegal distribution.

(Premium) Video on Demand (VoD)—The VoD use case is like Screeners except that the scale of the distribution is much wider. Forensic watermarks are typically used to identify subscribers who redistribute content. The closer the VoD distribution is from theatrical release the more likely forensic watermarking is to be mandated.

Live—Piracy of premium live events raises unique challenges for watermarking. The overall latency of adding a watermark to a live stream is of paramount importance to the end user. The streamed version of the content is usually behind the broadcast version. In some cases, like sport events, it is critical for the watermarking process not to further impair delivery time. Another issue is that the value for individual live events declines very quickly. The objective is to shut down illegal retransmissions as soon as possible. Forensic identification needs to be fast to reduce the end-to-end delay from scanning online piracy services, extracting the watermark from illegal content, and the issuing of take down notices or switching off the illegal retransmission at the source.

VoD and Live applications using video distribution over IP networks is a focus of this document. The use of a watermark can be used to determine the network and/or service provider responsible for delivering a piece of content and the user that accessed the content. This use-case has increasing relevance as distribution of valuable content through IP networks is becoming closer to the release dates, making the

[2] In this document emphasis is placed on video watermarking as it is more commonly used for forensic tracking, although audio can also serve as a carrier for forensic information for a variety of use cases.

content more valuable. Since the quality of the IP source video is increasing, the pirate can still provide a very high-quality copy, even after several generations.

Application Layers

Many watermarking systems permit layered watermarks where the mark can be inserted at different stages of the content packaging and distribution workflows without interfering or replacing one another. Such layered watermarks can be used to recreate the entire chain of custody should pirated content be found on sharing sites like BitTorrent, social media sites such as YouTube or Facebook, as well as Kodi[3] streaming devices.

Each identifier embedded at various stages throughout the production and distribution lifecycle can inform the Content Owner about all post-production entities that handled the content, the Operator that distributed the content, and even the end-user that consumed it. This allows the content owner to pinpoint the actual source of the agent responsible for leaking the content. Identifiers may be applied using technologies from different Watermarking Suppliers.

Related Technologies

Conditional Access (CA) / Digital Rights Management (DRM):

CA and DRM schemes can restrict content access to only authorized users who meet a set of conditions.

CA systems were initially designed for cable and satellite networks to switch on and switch off access to mainly linear programming based on a user's subscription or package.

DRM extends the CA concept and augments content encryption with an array of rights and conditionals based on a user's subscription, service tier, time and location, and even device. DRM allows a rights holder to grant extremely granular rights for the same piece of content, allowing one user to perhaps play the content only one time while granting another user access in perpetuity.

CA and DRM are a valuable means to prevent unauthorized access to secure content. While they are both very effective in their own rights, they can do little to isolate pirated content or identify the wrongdoers if that content is stolen and made freely available. With watermarking, unlike content encryption, the malicious user is not able to verify successful removal of the protection and only receives a notification about a possible infringement after the illegal use. In that way, a user intent on pirating content is never safe when attempting to distribute the content and emphasis is on deterrence and investigation rather than direct prevention.

Automated content recognition (ACR)

This approach to content protection analyzes the unique features of an audio or video asset and compares these against "reference" fingerprints stored in a database used for content identification. One of the key characteristics of fingerprinting is that it works off the unique characteristics of the content itself instead of modifying the content.

For example, a fingerprint used to verify content uploaded to a site for user-generated content such as YouTube) is recognized as being the same as previously analyzed copyrighted content. If this is the case, the content may be rejected, or the Content Owner may be alerted.

When watermarks are used to trace the source of piracy, ACR is routinely confused with watermarking. This is the reason that the term "forensic watermarking" is preferred today.

[3] For more information about Kodi, visit https://en.wikipedia.org/wiki/Kodi_(software)

Automated content recognition is also often called “fingerprinting,”. Since this term is sometimes used to denote visible on-screen messages, the term ACR is recommended.

7. Integration Options

There are two main approaches to Forensic Watermarking, routinely referred to as *one-step watermarking* and *two-step watermarking*. These two alternate strategies have been described in detail in the guidelines document published by the Ultra HD Forum (see Section 13 *Additional Information Sources*).

In this Section, we will focus on key aspects prior to integration. While all combinations are in use, selecting one-step vs. two-step watermarking may vary from one deployment to the other and are typically driven by considerations to optimize performance, control over device base, and implementation efforts.

One-Step Watermarking

This single step watermarking approach typically requires access to the uncompressed video where the embedding algorithm can identify locations that allow for robust and invisible modifications.

Integration in Client Devices

The uncompressed video (also called baseband or YUV format) may be accessible during decoding in a Client Device. In this case, care should be taken that the access is controlled and secured. One method to establish security is to implement the watermark embedding algorithm in the hardware itself, thereby making it part of the secure video pipeline. The watermark execution may be controlled with parameters such as an "ON/OFF" switch, a watermark embedding strength, and/or an embedded payload. These parameters need to be transmitted securely and handled by the Client Devices. A trusted execution environment (TEE) can be helpful to accomplish this.

The advantages of this client integration are:

1. Seamless support for broadcast content where content delivery is the same for all users,
2. Support for live content since there is no delay introduced or preparation required,
3. Independence of the codec and container (since embedding occurs after decoding).

The information to be embedded may be taken from the Client Device or provided from the head-end for a specific session.

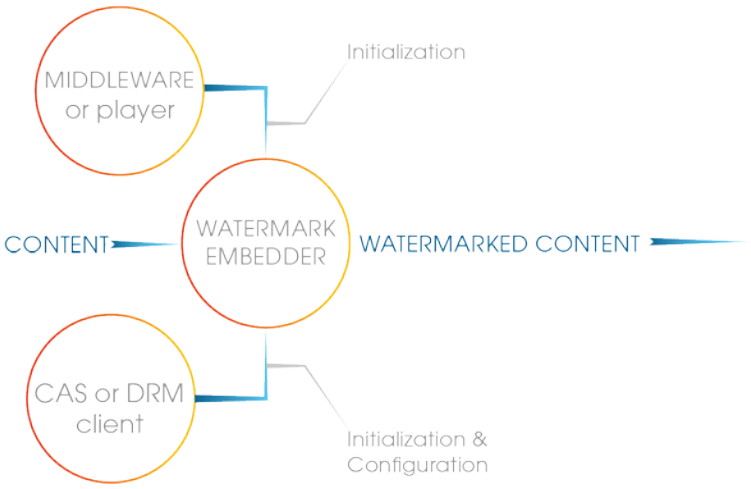

Figure 1: One-Step watermark performed on the client side.

Integration in the Head-End

One-step watermarking is usually employed for client-side embedding. However, it can also be integrated within a video encoder as well. This application requires a new encode for each individualized file which limits the scale to a small number of uniquely marked copies. As such, it is used less frequently and is most suitable for smaller batches in applications such as post-production or screeners.

Two-Step Watermarking

As its name suggests, two-step watermarking essentially breaks the entire watermark embedding process down to two sequential steps.

The first step creates Variants. These Variants are pre-watermarked variations of segments (i.e., chunks or fragments) of the video bitstream that can be used interchangeably without affecting the viewing experience while still providing means to perform Forensic Watermarking. As depicted in Figure 2, the Variants are encoded video data and may be generated pre- or post-encoding. The Variants are then forwarded along with the video where the second step is applied. The length of a Variant differs between Watermarking Suppliers, ranging from a couple of bytes to a whole video segment in ABR video networks.

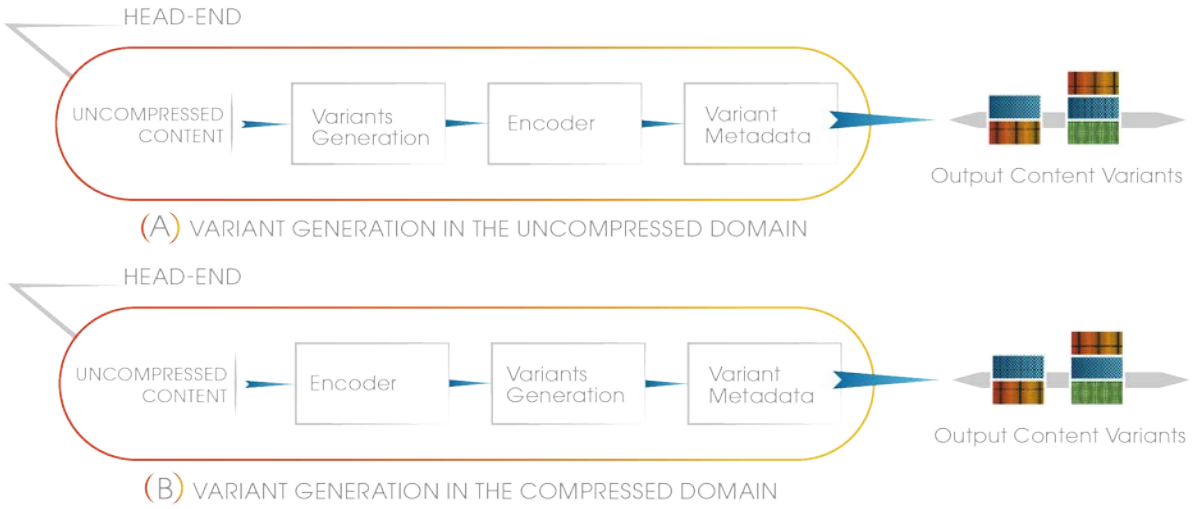

Figure 2: Two breeds of two-step watermarking systems

The second step is responsible for creating a serialized video bitstream that uniquely identifies the recipient to whom it is delivered. Using the Watermark Identifier (i.e., the payload) that is to be embedded, the serialization engine selects a unique sequence of Variants and then assembles the segments of video accordingly. The sequence of Variants encodes the Watermark Identifier with sufficient redundancy to establish reliable forensic evidence.

A well-known example of two-step watermarking is A/B watermarking for ABR video content. In this example, a Variant has the same size as the video segments used by the delivery network and there are two variants per video segment, Version A and Version B. The Client Device can use playlist manipulation or segment selection at the CDN edge to only receive a collection of video segments encoded with its Watermark Identifier.

First Step Integration in the Head-End

By design, Variant Generation must be integrated with the head-end.

The transmission of the Variant metadata along with the video content inherently induces an overhead that is directly related to the granularity of the Variants. For instance, in ABR A/B watermarking, video segments are duplicated which doubles storage on the Origin Server and notably reduces the caching efficiency of the CDN. This overhead is "internal" and not transferred to the egress of the CDN, which is still equal to the non-watermarked scenario.

The integration of the first step in the head-end may be more or less intimate with encoding. To avoid introducing dependencies with a specific encoder vendor, one could aim at making the Variant generation process separate from the encoding equipment. On the other hand, this may introduce delays due to the video transport layer. As an example, in A/B watermarking, the generation of the variants may result in a whole segment delay. This may not be acceptable for some application use cases such as live transmission. These cases may need a tighter integration with encoder vendors.

Second Step Integration

The rationale for splitting the watermarking in two steps is to offload the computational burden in the first step to make the second step as simple and scalable as possible. This lightweight serialization agent generates a unique Sequence of Variants that encodes the desired Watermark Identifier. Therefore, it can be integrated at any point in the video delivery network where an identifier is available.

Serialization in the Head-end

The second step can be placed at the head-end, like the origin server in an OTT scenario. All segment requests are escalated to the origin server that is then in charge of producing a serialized segment containing a unique Sequence of Variants that encodes part of the Watermark Identifier associated to the entity querying the segment. In A/B watermarking, it amounts to returning either version A or version B of the requested segment. In such an integration, all segment requests reach the origin server so there is no cache. While it precludes large scale deployments, such integration may be appropriate in some low volume cases such as screeners, hospitality window, or mezzanine content distribution.

ABR Playlist Embedding

As a first solution to achieve good scalability with A/B watermarking, playlist manipulation through the manifest has been introduced to restore some caching capabilities of the video delivery network. The key idea is to deliver serialized playlists to only declare segments that encode the Watermark Identifier of the recipient. Regardless of the network conditions, the client can only request segments that encode the Watermark Identifier associated to them. The advantage of this approach is that the A and B versions of the segment are cached in the video delivery network and requests may not escalate to the origin server if the desired segment is already present in cache which would reduce latency. For security reasons, the segment URIs need to be obfuscated to prevent malicious users from manipulating a playlist themselves. The intelligence to decide which segment to deliver lies in the central manifest manipulation component.

ABR Segment Selection on the Edge

A/B segment selection is when the selection is performed between prepared segments of A/B variants just in time for each client segment request. Here, the serialization step of a two-step watermarking system is performed in the edge servers.

In this case, all recipients receive the same playlists. The video delivery system still needs to deliver a unique sequence of A and B segments to every individual. This can be achieved by making the edge servers select either A or B version of the segment to be returned when they receive a segment request. When the version of the segment has been selected, the edge server can then query the cache possibly to the origin server which delivers it to the recipient. The edge server then has intelligence to make the

decision on what segment to deliver. This solution has the same CDN caching properties as the playlist embedding for the content along with the additional benefit that the playlist can also be cached.

This application may be required in cases where byte-range indexed playlists are used, or the playlist is templated (such as with Smooth Streaming and DASH VoD profiles) and segments cannot be addressed individually.

The serialization effort during playlist delivery is eliminated and the same playlist can be used for all streams. However, the edge needs to apply logic to decide for each requested segment. This logic includes identification of the information to be embedded, a decision for segment selection, and the delivery of the corresponding segment.

The fact that all recipients receive the same playlist provides an extra layer of security against comparison and manipulation of playlists before the content is downloaded. It is recommended to make use of https-based segment URIs or other strategies to avoid local ISPs from further caching the seemingly "common" content segments after they leave the CDN edge, thus destroying the A/B serialization pattern.

Segment Modification on the Edge

Some edge servers can perform late repacketization operations on the segment content itself. This allows for segments to not only be selected on the edge allowing choice between A/B, but preceding assembly of these segments even after they have been requested. The idea is to use a single video transport protocol between the origin server and edge servers to optimize caching capabilities and perform repacketization operations (container and scrambling) at the edge to deliver segments in the desired format. Such repacketization implies that the encoded video buffers are available in cleartext at some point. For two-step watermarking systems that use Variants at a finer granularity than the whole segment, it provides an opportunity to perform the serialization step at the edge. In this case, the transmission of Variants metadata alongside the video can be significantly lower than the ratio inherent to A/B watermarking where storage is doubled, and cache is doubled for A/B segments. This approach is not limited to ABR, it can apply to any progressive download and employs. a common playlist for all subscribers.

Serialization in the Client Device

Finally, the serialization step of two-step watermarking can be integrated in the client device. As mentioned earlier, this process is very lightweight and may not require extra hardware. It does imply, though, that the Variant metadata travels along with the video from the head-end down to the device. To keep control over the bandwidth overhead, such integration usually requires having Variants of the finest possible granularity.

When the serialization process operates on the cleartext encoded bitstream, its integration within a TEE may be recommended for security reasons. Alternately, the different Variants may be encrypted with different keys and each device provisioned with a unique set of keys that provide access to a unique Sequence of Variants. Such crypto-binding of the access to Variants has been standardized for ISO BMFF[4] and for Blu-ray discs[5].

[4] For this standard, see ISO/IEC 23001-12:2015, Information technology -- MPEG systems technologies - Part 12: Sample Variants in the ISO base media file format

[5] H. Jin, J. Lotspiech, and S. Nusser, "Traitor Tracing for Prerecorded and Recordable Media", ACM DRM 2004.

8. Watermarking Payload for Forensic Evidence

Identification

Viewer identification is critical in the process of forensic analysis. To make such identification valuable, the watermark should contain as much information as possible while also respecting the constraints of limited payload in the content and privacy. One possibility to accomplish this is to store a short device ID. If the number that can be embedded in the watermark is the maximum expected population size and does not contain additional information or coding, it is most efficient. The number that is stored in the watermark should be traceable to a viewer or household when found in pirated content, but not publicly traceable by entities other than the operator providing content to the consumer.

Transaction Database is another approach that enables precise identification while still maintaining privacy. This approach assigns a unique number for each transaction that is stored with information about the client, the content, a timestamp, IP address, device type and any other relevant details. The transaction number on its own does not reveal any personally identifiable information (PII), but the presence and access to the database is required to maintain the use of the watermark. If a transaction number is used as an embedding element, the Transaction Database is the only location that can identify the transaction once marked content is found. The database needs to be maintained for the duration of possible use of the watermark and is required for any legal actions.

Since the database does contain PII, it should be treated with similar security levels to an operator's comparable information like billing and operational information.

Granularity

The Transaction Database can contain any information that may be relevant to forensic investigation. The more information a transaction identifier links to, the higher the resulting level of forensic evidence.

In other cases, however, the device type identification may be sufficient. In this instance, a viewer will not be able to be identified except by a device type or configuration. This may be helpful to understand vulnerabilities to content protection exploited by a specific device, OS, application version, etc. In this case, it would be used to identify the technical source of a leak rather than a specific end user.

To understand the general timing of a content leak, this may be combined with watermark information relating to a time interval.

Formatting Payload to Discourage Collusion Attacks

A collusion attack is an approach where an adversary aiming to make a mark unreadable is combining several copies of content. Some methods to foil this attack are to spread payload bits in different domains, and to apply error correction codes that identify the mark after collision. In general, anti-collusion techniques are more effective with more payload bits retrieved from the content.

Formatting the payload with this attack in mind can help to spoil it. Based on the "marking assumption" that bits in common to combined videos will survive any averaging approach, those are maximized by using the same bits to signal information that is likely the same between attackers. This may be a device type if the copies leaked due to the same device weakness, a geographic location if the attackers are close together, a timestamp if the copies are made in a similar time, etc.

In addition, it may make sense to reduce the number of transactions assigned to an individual user by repeating the same transaction number for a certain time interval. In this example, a user will receive the same payload for any access of the same content for the duration several weeks which prevents collusion between copies of a single user in a short interval.

Marking & Detection Workflows

Content Owner's View

The following process represents a common workflow for the Operator Mark (Figure 3):

1) The Content Owner prepares a piece of content for distribution. Within this asset a unique per-operator identifier is watermarked.
2) The watermarked content is delivered to the Operator.
3) The watermark ID is registered with the Content Owner's Watermarking Supplier system on-premises or remotely.
4) The Transaction Database is shared with the Content Owner's Piracy Monitoring Service.
5) The Operator processes the content and inserts a unique session-based watermark into the stream.
6) The identifier associated with the watermark is registered into the watermarking supplier database. This may also include content information for non-blind extraction.
7) The watermarked content is streamed via the content consumption client. During this consumption, the content is distributed to a Pirate Source.
8) The Content Owner's Piracy Monitoring Service identifies the pirated content and extracts the operator watermark to determine the origin. The operator may independently run a Monitoring Service.
9) For mitigation, the Content Owner could defer to the Operator to act on closing a leak.
10) The detected content is provided to the Operator for forensic analysis.
11) The detected content is provided to the Watermarking Supplier for watermark extraction.
12) The watermark ID is returned to the Operator.
13) The Operator mitigates the breach based on content owner policy. This could be done by providing policies to Monitoring Service. See Section 8.4 for Incidence Management details.

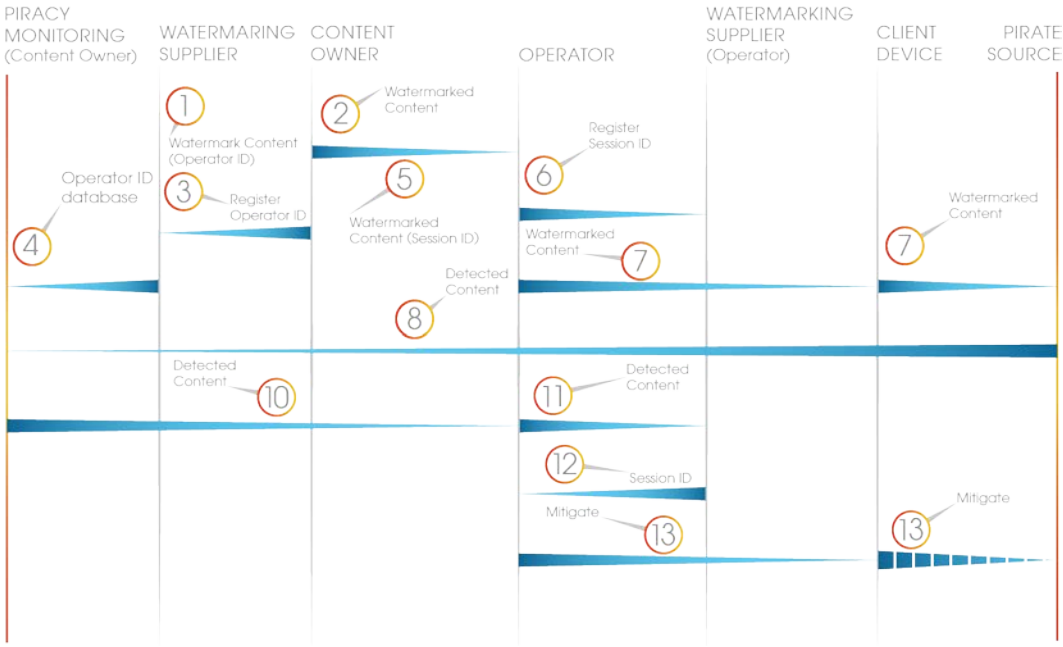

Figure 3: Watermarking Workflow, Content Provider Perspective

OTT Streaming Service View

The following process represents a common workflow to manage the Session Mark (Figure 4):

1) The Operator acquires a piece of content for distribution from the Content Provider.
2) The Operator processes the content for ABR delivery and inserts a unique session-based watermark into the stream.
3) The identifier associated with the watermark is registered into the Watermarking Supplier database.
4) The Transaction Database is shared with the Operator's Detection Service.
5) The watermarked content is streamed via the content consumption Client Device. During this consumption, the content is distributed to a Pirate Source.
6) The Operator's Piracy Detection Service identifies the content on the Pirate Service and extracts the session watermark to determine the origin.
7) The Operator's Piracy Detection Service informs the Operator of the breach.
8) The Operator informs the Content Owner of the breach.
9) The Operator's Piracy Detection Service mitigates the breach based on the Content Owner's policy. This could be done by providing policies to Monitoring Service. See Section 10.4 for Incidence Management details.

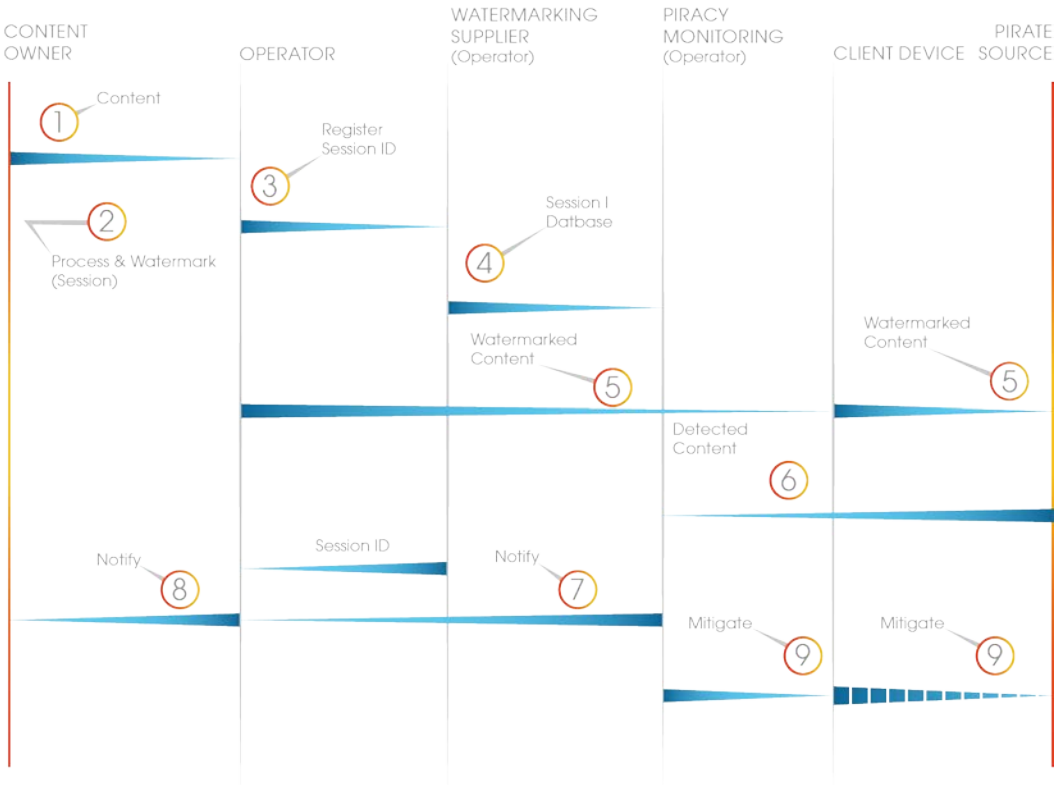

Figure 4: Watermarking Workflow, Service Perspective

9. Privacy Consideration for Watermarking

Helpful privacy guidelines are listed by the Center for Democracy & Technology [CDT][6]. These guidelines provide considerations to maintain basic privacy rights for consumers of forensically marked content.

[6] https://cdt.org/insight/privacy-principles-for-digital-watermarking

10. Survivability Considerations

Watermarking survivability requirements may vary depending on what should be protected, such as quality level or release window. While the maximum survivability and robustness is desired, for some use cases the amount of modifications and possible visible impact should be minimized. Priorities for different use cases can help to find suitable tradeoffs. Different requirements can be identified whether the protected value is of higher quality content only, or if it is more important to identify the piracy origin under several content degradations or within immediate time window.

Public Requirements

While the requirements vary between content owners, with different applications, and over time, there are some published requirements that help to provide orientation or serve as a template.

Digital Cinema Requirements

The requirements for forensic watermarking within Digital Cinema requires both an audio and video watermark with a 35-bit payload. These two watermarks make it robust against a list of common video and audio processing transformations. These have been described in detail in section 9.4.6.1 of the Digital Cinema specification V1.2[7] .

Studio Requirements (MovieLabs)

MovieLabs Specification for Enhanced Content Protection, v1.1[8] has the following requirements around security of the watermarking solution:

- The system shall have the ability to securely forensically mark video at the server and/or client to recover information necessary to address breaches.
- The watermarking shall be robust against corruption of the forensic information.
- The watermark shall be inserted on the server or on the client such that the valid insertion is guaranteed during playback even if the device and its secrets are compromised.

Protecting Quality (UHD over HD)

The protection of the content quality is most important for some applications. This is especially true if the content is available thru unprotected channels in a lower quality version. As recording technology improves in resolution and quality, new distribution and storage technology is created, and new security is applied. This happens while existing distribution channels with traditional quality may have exploitable weaknesses. In these instances, watermarking may be used to protect the highest quality version while other quality levels are not as important. In addition, the mark may not have to survive strong degradation. Relevant attacks include:

- 4k recompression with VBR, 20 mbps, AVC/HEVC
- HDR removal or modification
- Downscale one level e.g. 4k to H1080p (or HD to 720p)
- Cropping to 20 min in time
- Fps conversion between Film/PAL/NTSC and HFR
- Targeted Attacks:

---

[7] http://dcimovies.com/specification/DCI_DCSS_v12_with_errata_2012-1010.pdf

[8] http://movielabs.com/ngvideo/MovieLabs%20Specification%20for%20Enhanced%20Content%20Protection%20v1.1.pdf

- Geometric or temporal de-synchronization
- Collusion of 2 streams
- Add / reduce spatial noise, add / reduce flicker, sharping, saturation, rotate

Protecting Release Windows

Revenue loss is a large concern for content owners and distributors. Due to this, both the quality of the content as well as the earlier-release content must be protected from compromise.

In these cases, the mark needs to be robust against the degradations above as well as:

- Heavy compression to HD and SD formats
- Camcording in both high and low quality
- Image modifications like: rotation, cropping, adding mattes, mirroring
- Frame Rate changes
- Aspect ratio changes
- Editing portions of content
- Decimating frames

Protecting Live Content

For the protection of live content, it is important to identify the leaking source quickly to disable the stream. While quality and security requirements are important, they are secondary. The watermark needs to survive the degradation with the additional requirement of fast extraction with a short duration of content.

Content Piracy Use Cases

Scenarios for unauthorized capturing or re-streaming include:

1. Straight digital capture of the stream using screen grabbing software or capture from an HDMI signal with HDCP protection removed. Content captured in this manner would make detecting of the watermark much more straightforward. Detection would need to be robust enough to withstand any heavy compression and possible changes in frame rate, scale, and aspect ratio. HDR content may be down converted to SDR. Also, in this scenario, the image could possibly be reversed if the pirates thought it would eliminate the ability to trace it back to them.
2. When a straight capture of the stream is not feasible, the most likely method of content capture is by camcording the content with a mobile phone or a high-quality camcorder. Testing shows that very high-quality video can be captured. Quality of pirated content in this manner can vary greatly depending on factors such as lighting in the room, positioning of the camcorder, and expertise of the person recording the content. For this type of piracy, the detection should be straightforward as there would be little movement of the image during the recording, but there might still be degradations such as image cropping and rotational issues. Watermarking in this case still needs to be robust to frame rate changes and images being rotated since it is tough to get the recording device perfectly aligned and synchronized. The more difficult detection is when there is no control of the lighting in the room and the camcorder may just be sitting on something or being held. This will introduce rotation and movement in the recording as well as a washed-out or very dark recording of the image.

In the first scenario, expectation is to detect very quickly as the content should be taken in or close to the state that it would have been watermarked. The second scenario would require more effort to detect, especially if the utilized watermark relies on spatial detection. In this case, the investigator would need to align the pirated file with the original file from a time perspective, then do a continuous

temporal alignment depending on how much movement is in the image. In a blind detection scenario, the detector should automatically account for the temporal and spatial issues as part of its functionality.

Understanding Robustness Trade-Offs

Watermark robustness generally describes the survivability of the mark after content transformations. It is a complex parameter since it has dependencies on many different factors that may include:

- the nature and configuration of the transformation,
- the characteristics of the underlying content,
- the length of the content, and
- the amount of embedded information.

Even for a common transformation like compression, the robustness depends on details such as codec, codec configuration, and implementation making it difficult to define robustness boundaries. There is a common trade-off that can help the configuration of watermarking approaches depending on the goal of protecting the release window like content value decreasing rapidly over time, or quality level like content in lower quality having a lower protection threshold. For instance, in live sports, the quick evaluation, tracing, and blocking of piracy is of the highest importance, while the security against targeted manipulation does not need to protect against complex attacks with long execution times. As a result, the mark should be strong enough to be read quickly.

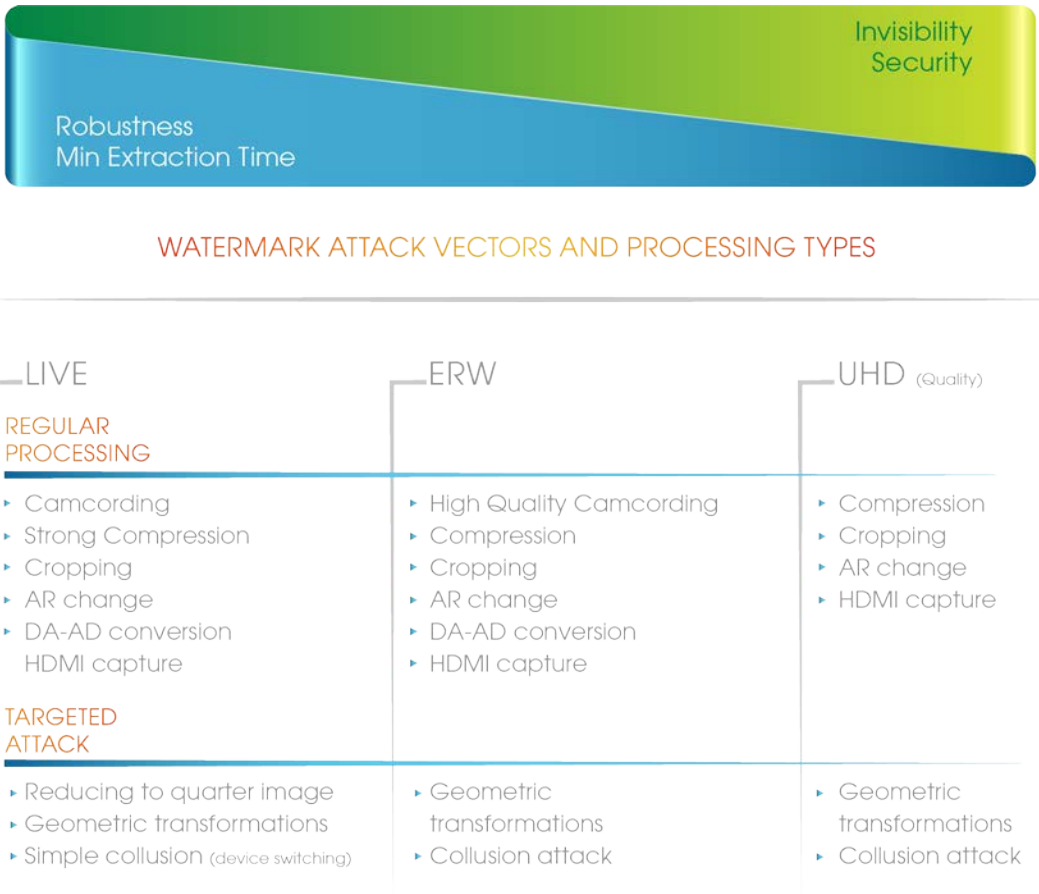

Figure 5: Watermark attack vectors and processing types

11. Performance Considerations

Extra processing is required to be performed during content delivery or playback. This will individualize the content with a unique watermark. When doing so, it is critical to minimize any integration complexity that might impact performance or introduce playback latency.

One-Step Marking

During embedding in the baseband domain, the content is modified during playback. The critical item for performance is marking during content processing. If this occurs in the client playback device, it is in parallel to content decryption, decoding, and display. Depending on the implementation, resources may be shared with these functions or dedicated hardware blocks may perform those elements without using the general CPU.

Two-Step Marking

The separation of content modification and assembly creates integration points for both steps. During live scenarios and real-time preprocessing, the performance of the first step is relevant to minimize any delay during content preparation. Depending on the selection and application, individual Variants can be small and replace individual macroblocks or large, to replace several seconds of content. The additional content to create, distribute, and deliver may be a significant portion, double the size or more. The created content overhead may be a trade-off compared to the complexity of content processing before delivery[9].

Watermarking Technology Performance

Independent of the application mode, the actual watermarking technologies may vary in the ability to survive different degradations, security in the ability to withstand attacks, payload density in the ability to hide data in a given content period, and ease of extraction. While technology emphasis may differ, watermarking systems often allow for selection of these tradeoffs and can be configured for the given protection scenarios. For example, a larger payload may be selected despite increasing content duration or extraction time.

A high payload bitrate may further benefit the effectiveness of anti-collusion techniques. This may not only be a property of the watermarking technology itself but can also be dictated by the selected adaptive bitrate segment duration in the case of some two-step watermarking implementations.

[9] For more information on two-step watermarking, see the Ultra HD Forum: Phase A Guidelines, August 25, 2017,Revision: 1.4 available at https://ultrahdforum.org/wp-content/uploads/Ultra-HD-Forum-Guidelines-v1.4-final-for-release.pdf

12. Extraction Considerations

The extraction of the embedded information, also called readout or recognition, can be performed and initiated in different ways and may result in different outcomes. Depending on the application, technology, and implementation the following variations should be considered.

SaaS vs On-Premises

Extraction typically requires a computer process that analyzes the stream and video pixels to extract the embedded information. This could be performed using hardware or software provided by the watermarking supplier. Having the extraction application on site allows control over the process by the operator but does require the ability, training, and staff to run the extraction.

Alternatively, the extraction may be provided as a service (SaaS) where content is uploaded by the operator and extraction results are returned. Depending on the performance of the extraction device and upload speed, the time for extraction may increase in this method.

The security of the overall solution is another item to consider. For example, if the extraction logic is compromised through a system intrusion and algorithms are used to detect and expose watermarked content, it is possible that pirates could reverse engineer the extraction process and reveal how content is watermarked. Alternatively, the pirates could launch an “oracle attack” where they employ the leaked detection software as a black box to learn how to degrade the content just enough to prevent successful watermark payload extraction.

Blind / Informed / Non-Blind

Extraction of the mark may be aided with information about the content that it was embedded in. Watermarking approaches are called *non-blind* if they require the original for extraction or *informed* if some information is required. Otherwise they are called blind.

Blind approaches are easiest to deploy and use since the original content is not required. This method is most often used when there are challenges in acquiring the source content. However, non-blind approaches may increase the robustness of the mark as well as security of the solution, since the original, unmarked content is something that a pirate will not have access to.

Different scenarios of Creating on Operator Mark

There are often two stages of marking, as illustrated in the workflow diagrams, to identify the view in Section 8. A first level is used to identify an operator or distribution channel and a second level is used to identify the end user or session. The Operator Mark is useful to identify the relevant database for lookup of the session, embedding parameters, and possible content information for informed or non-blind embedding. It is required if several operators distribute the same content and there are no visible clues that reveal that information such as a broadcast bug or operator logo. Different marks should not interfere with each other if marking technologies differ in their technical embedding approach. Simple tests can be performed to verify specific technology combinations considered for deployment. These tests should include different levels of watermarks, applications using different strength, embedding and possibly removal, and include operator marks, applied by content owners.

Incident Management

After the mark has been identified and the source of the leak is located using the operator and possibly end user mark, different actions may be taken by the Content Owner or Operator to use this information:

1. **Intelligence**—the information is recorded to gain knowledge of content abuse and information about timing, location, and frequency. This may help to inform the Content Owner's decisions on future content releases as well as inform law makers to improve anti-piracy regulations. Cross-correlating this information with other records such as from billing systems can often help identify patterns and associated suspicious transactions.
2. **Service limitations**—an Operator may initiate an immediate take-down of the infringing stream. Users suspected of piracy may face limited access to content that is otherwise prone to piracy such as longer delay windows or suspension of their account.
3. **Legal actions**—the evidence from the extraction is used as proof to initiate measures against illegal content use. Measures can be addressed towards both pirate sites or end users and are typically initiated by Content Owners. This may serve as a deterrent to widespread piracy.

Some Content Owners also make use of services that identify copyright-infringing content on the Internet based off manual identification or content fingerprinting. Then they use automatic take-down tools or send take-down notices to relevant parties, including ISPs and CDNs. This may harm evidence collection when a stream distribution point is shut down before enough content is recorded for a complete watermark payload extraction if this is used in combination with session-based watermarking. Although the pirate viewing experience is disrupted, without source identification the stream may return quickly at an alternative distribution point. An Operator should establish a proper balance with the Content Owner.

13. Additional Information Sources

Digital Cinema System Specification by Digital Cinema Initiatives contains useful information on watermarking and robustness as used in digital cinemas. The resource can be found at:

http://dcimovies.com/specification/DCI_DCSS_v12_with_errata_2012-1010.pdf

Movielabs Specifications for Next Generation of Video and Enhanced Content Protection has information about watermarking as well as other content protection methods. The resource can be found at:

http://www.movielabs.com/ngvideo/

Phase A Guidelines by Ultra HD Forum covers Forensic Watermarking in Section 7.2 and is a useful technical companion to this document. The resource can be found at:

https://ultrahdforum.org/wp-content/uploads/Ultra-HD-Forum-Guidelines-v1.4-final-for-release.pdf

14. Illustrations

15. About the Streaming Video Alliance

Comprised of members from across the video ecosystem, the Streaming Video Alliance is a global association that works to solve critical streaming video challenges in an effort to improve end-user experience and adoption. The organization focuses on three main activities: first is to educate the industry on challenges, technologies, and trends through informative, publicly-available resources such as whitepapers, articles, and e-books; second is to foster collaboration among different video ecosystem players through working groups, quarterly meetings, and conferences; third is to define solutions for streaming video challenges by producing specifications, best practices, and other technical documentation. For more information, please visit www.streamingvideoalliance.org.

Streaming Video Alliance
5177 Brandin Court
Fremont, CA 94538 USA
(510) 492-4000
streamingvideoalliance.org

© 2018 Streaming Video Alliance.