SMPTE Meeting Presentation

Is The Future of Content Protection Cloud(y)?

Diehl Eric

Technicolor, 1 Rue Jeanne d'Arc, 92130 Issy Les Moulineaux, France,
eric.diehl@technicolor.com.

Written for presentation at the

SMPTE 2014 Annual Technical Conference & Exhibition

Abstract. *As the adoption of cloud computing continues to progress and becomes a critical element in the production and post-production processes of the movie industry, new strategies for protecting content will have to emerge. The lack of sufficient content protection may hinder the adoption of cloud in the industry, which could deprive the community of the many operational and financial benefits that this new approach to delivering technology services can offer.*

This paper defines four types of trust assumptions and their relative strength in the different cloud deployment models. It addresses the major threats associated with cloud -- such as data breaches, account hijacking, denial of services, or malicious insiders -- with a particular focus on content protection.

The cloud increases exposure to risk. The paper explores two use cases: digital delivery and processing. They demonstrate that the clever use of hybrid cloud and new advanced technologies, associated to proper security, may allow secure deployment in the cloud without sacrificing content protection.

Keywords. Security, cloud, hybrid cloud, encryption, watermark, digital screener

Introduction

Cloud computing and cloud storage offer many advantages over traditional computing and storage. These benefits are both operational (rapid elasticity, scalability, and broad access) and financial (metered services, on-demand service, and resource pooling). With such assets, cloud is on a fast path to become a commodity (Carr, 2008). As adoption of cloud computing continues to progress, it will become a critical element in the production and post-production processes of the media industry.

Content protection is an essential requirement for production and post-production. The lack of sufficient content protection may hinder the adoption of cloud in the media industry, which could deprive the community of the many operational and financial benefits that this new approach to delivering technology services can offer. New strategies and secure architectures will have to emerge.

This paper explores whether acceptable content protection is currently achievable in the cloud. The first section briefly introduces the different cloud architectures with a rough estimation of their respective exposure to risk. The second section addresses the major threats associated with cloud. Most of these threats are generic for any cloud application. Nevertheless, content protection poses some specific challenges that this section examines. The last section describes two use cases: secure delivery and processing in the cloud. It analyzes the content protection practices applicable to private clouds and hybrid clouds.

This paper uses a writing convention. When the paper does not explicitly use Alice, Bob and Eve characters, it assumes that the attacker is a female character whereas the attacked person is a male character.[1] This convention allows easy identification in the same sentence of who is who, and helps to distinguish clearly the attacker from her victim.

Cloud architectures and trust

Definitions

There are several cloud deployment models. The National Institute of Standards and Technologies (NIST) (Mell and Grance, 2011) and the International Telecommunication Union (ITU) (Anon, 2014c) define four deployment models:

- *Private cloud* is for the exclusive use by a single organization. The organization – or a third party – may own, manage and operate the infrastructure. However, only principals belonging to the organization share the pool of resources. The infrastructure may be on or off premises of the organization. For instance, a private cloud can be designed and hosted by the organization's IT team or can be provisioned from a cloud service provider.
- *Community cloud* is for the exclusive use by a community of organizations that share a common concern or goal. One or more organizations of the community can own, manage or operate the infrastructure, or a third party can undertake those functions. Regardless of the arrangement, only principals belonging to the community may share the pool resources. The infrastructure may be on or off the premises of the community members.
- *Public cloud* is for the open use of infrastructure by general public. A third-party cloud service provider owns, manages and operates the infrastructure. Entities may share the resources. The infrastructure is on the premises of the cloud service provider.
- *Hybrid cloud* is the composition of two or more of the previous deployment models.

[1] The reader should not look for any stealthy meaning or interpretation in this distribution of roles.

All four deployment models support the three NIST-defined service models: Infrastructure as a Service (IaaS), Platform as a Service (PaaS) and Software as a Service (SaaS).[2]

This paper will focus on the SaaS model in which:

- A service provider (which could be a content distributor or post-production house) manages a cloud-based service for its users. The service provider designs the offered service and controls it.
- A user uses this service. The user may be an employee an employee of the service provider or may be an external worker.
- The service provider may use its private cloud or an infrastructure owned by another cloud provider.
- The service may use software solutions provided by a cloud partner.

Trust

Trust is the cornerstone of security. There are many definitions of trust. This paper uses Roger Clarke's definition (Clarke, 2002):

> *Trust is confident reliance by one party on the behavior of other parties.*

In other words, it indicates a belief that other parties will be reliable, and that they are worthy of trust. Trust is the foundation of any security system. If the foundations of a building are weak, then the building may collapse. If the foundations of a building are sound, then the building will be robust and survive many incidents. The same is true for trust. Without a good trust model, security system will inevitably collapse. Therefore, it is key to know and to understand the trust assumptions associated with different technology architectures.

Secure architectures rely on many types of trust assumptions. This paper categorizes them in four large clusters.

- *Trust in the hardware.* This category encompasses all assumptions about the physical environment. The usual assumptions are:
    - That the hardware is reliable and genuine – meaning that it has not been modified by an attacker. For instance, the hardware device should not have been the subject of a successful supply chain attack (Anon, 2012), or maliciously being refurbished.
    - That the hardware is located in a secure area where access is restricted to authorized staff. Physical security often requires fences, guards, surveillance cameras and access control measures -- such as badges.
    - That the hardware is in a known and specific geographical location. Often, the actual geographical location of hardware can determine whether or not the system is in compliance with certain national regulations -- such as the European Directive on personal data.
- *Trust in the Operating System (OS) and Virtual Machines (VM).* It is critical to ensure that the OS and VM are genuine, properly patched, and not infected or otherwise compromised. In the cloud environment, particular attention must be focused on effective, secure virtualization. When assessing virtualized environments, one of the most-important security considerations is to ensure the total isolation between VMs operating on the same machine. Once this trust established, it is also important to

[2] ITU defines additional models called cloud service categories. New cloud service categories are Compute as a Service, Communications as a Service, Data storage as a Service, or Network as a Service.

ensure that trust has been established for the different software components and final applications that are in the virtualized environment.

- *Trust in the administrators.* A computer system is only as secure as the administrator that oversees it (Anon, 2014a). Malicious insiders are a serious threat.
- *Trust in the access controls.* Only trusted, authenticated principals should be granted access to the system. Principals encompass individuals, services and computers. Trusted access is a foundational assumption for a perimeter defense system. In this configuration, the system executes behind the secure walls of the IT defense.

Table 1 provides a subjective evaluation of the strength of these four categories of trust assumptions for different deployments. The rating is subjective because it assumes that security has been properly implemented for each deployment. Unfortunately, this assumption is not always true. In other words, depending on the threat model, a proper secure implementation in the public cloud may be more trustworthy than a weak implementation in a data center. Nevertheless, this paper assumes that the implementations comply with the best practices.

Table 1: Robustness of trust assumptions

	Data Center	Private cloud	Community cloud	Public cloud
Trust in Hardware	✸✸✸	✸✸✸	✸✸	✸✸
Trust in OS/VM	✸✸✸	✸✸✸	✸✸	✸✸
Trust in Admin	✸✸✸	✸✸✸	✸✸	✸
Trust in Access	✸✸✸	✸✸✸	✸✸	✸

In this document, data center means an IT infrastructure that does not use cloud technologies (virtualization, multi-tenancy, elasticity…) and is hosted in a secure area. The data center is the traditional environment for most mature organization. The system runs within a secure data center owned by the organization and managed by the organization's administrators. Physical and logical access is only granted to principals approved by the organization. From the point of view of security, this configuration is optimal. Unfortunately, it lacks some key benefits offered by cloud solutions – such as rapid elasticity.

The private cloud turns the traditional data center into cloud deployment. The private cloud runs within secure data centers owned (or at least controlled) by the organization, managed by the organization's administrators. Access is only granted to principals approved by the organization. Nevertheless, some researchers believe that private cloud can theoretically be less secure than a traditional data center because cloud environments are more complex than traditional data center environments (Haimes et al., 2014). Indeed, complexity increases the surface attack. Nevertheless, in this paper, we attribute to the private clouds the same trust levels associated with data centers.

With community cloud, trust starts to dilute. For the sake of simplicity, this section assumes that Companies A and B share a community cloud. Company A manages the cloud infrastructure. Company B has to trust that Company A will do a proper job of securely managing hardware, OS and other critical enterprise system components. Company B must also trust that Company A's administrators are trustworthy. Furthermore, each company has to trust that the principals authorized by the other company are trustworthy. If an attacker compromises a principal of Company B, she may use it to attack Company A's assets. Usually, the purpose of community cloud is to serve a common goal. Thus, these trust assumptions may be reasonably reliable -- although they are weaker than the trust assumptions for private cloud.

With public cloud, the dilution of trust is even greater. For hardware trust and OS trust, under the assumption that the cloud service provider is serious, the strength of these assumptions may be equivalent those associated with community cloud. However, one issue that can be difficult to resolve with public cloud solutions revolves around the location of the provider's processing units. Many organizations may be uncomfortable with having sensitive data stored in foreign countries. Also, the management of the infrastructure is under the full control of the cloud service provider. Most cloud providers do proper screening when hiring their administrators. Nevertheless, the cloud service customer has no way to assess the trustworthiness of the administrators of the cloud service provider. Access to key systems is outside of the control of cloud service users. Unknown principals may share the same CPU, the same physical memory and hard drives. Thus, trust for administration and access may be weak.

When a service provider uses a public or community cloud, it delegates some of its security-related roles to the cloud provider. As such, the security assumptions theoretically become more fragile because the service provider loses some control. Hybrid cloud may reinforce some of these assumptions by carefully defining the distribution of tasks between private cloud environments that offer strong assumptions and public cloud resources that present weaker assumptions. (The section on "Use cases" will show some examples.)

Threats

Generic threats

Security requires understanding threats. In 2013, the Cloud Security Alliance listed the nine top threats associated with cloud computing (Los et al., 2013):

1. Data breaches. An attacker may gain access to data. The usual technical solutions are authenticated access control and encrypting data at rest. Encryption of data at rest is the best solution as the attacker could gain access to clear data. Nevertheless, encryption requires careful management of cryptographic keys.
2. Data loss. Losses may result from an attack, a technical problem, or a consequence of a catastrophic event. The loss may also come from the loss or destruction of cryptographic material used to protect against data breaches. Ransomware is a new threat that now falls into this category. New ransomware can encrypt files on networked directories (Leyden, 2013), or attack network attached storage (NAS) infrastructures (Pott, 2014). Thus, this threat may extend to online backups.
3. Account hijacking. This threat is not specific to cloud – or exacerbated by cloud. Strong, secure identity management can mitigate account hijacking. Unfortunately, user authentication is not sufficient as some sophisticated attacks may use session hijacking which occurs after successful authentication.
4. Insecure application program interfaces (APIs).This threat must be particularly kept in mind by system designers. It is not necessarily unique to the cloud, but a cloud infrastructure undoubtedly exacerbates this threat.
5. Denial of Service (DoS). Since cloud operations are often remote, they are prone to DoS attacks. The primary type of attack is network access flooding via Distributed DoS attacks (DDoS) or amplified DDoS attacks. This threat is unlikely for private clouds that operate within the local area network. However, this threat is highly likely to be unleashed on public cloud environments. On the other hand, it should be noted that public clouds are designed to expect huge amounts of traffic. As a result, public cloud

providers may be better prepared to counter this type of attack than a typical data center.

6. Malicious insiders. The previous section highlighted the importance of administrators' trustworthiness. For instance, Edward Snowden was a contractor of the National Security Agency (NSA) working as an administrator. Users and operators are also potential insiders.
7. Abuse of cloud services. This occurs when people use the cloud for nefarious actions -- such as password cracking, or flooding a service.
8. Insufficient due diligence. Jumping on the cloud bandwagon without enough preparation may be an issue. This threat is not specific to the cloud. It is true for any new paradigm.
9. Shared technology vulnerability. Because cloud environments often share hardware and software components, a single vulnerability can affect many players. A good example is the extensive use of common cryptographic library or HeartBleed (Anon, 2014b). If an attacker finds an exploit affecting one service provider, the same exploit may threaten other service providers. "Break Once, Break All" may be encountered too often. Software diversification is one line of defense.

Content protection

This section analyzes the relevance of these threats to content protection.

1. Data breaches. Content theft is the main threat that content protection strategies must address.
2. Data loss. Content loss is probably less an issue for content protection. As long as master material is safely backed up, it may be possible to restore lost content. Nevertheless, restoration of lost content may be costly and take a long time to recover.
3. Account hijacking. This is not particular to content protection.
4. Insecure APIs. This is not particular to content protection.
5. DoS. This is normally out of the scope of content protection. Nevertheless, DoS attacks can be combined with clever social engineering to open the door for other advanced attacks. Movie productions often operate under extremely time tight deadlines. Thus, DoSing a system may be a clever method to have an operator send out content without following the protected workflow protocols just to respect the delays. In so doing, content can be put at risk.
6. Malicious insiders. For years, they have been an issue for content protection (Byers et al., 2003). For instance, MPAA guidelines for site protection put a strong emphasis on tracking the movement of workers within the facility (Anon, 2013). Table 1 shows that trusting the administrators and trusting access in a private cloud is equivalent to the trust in a data center. Therefore, if best practices are implemented, high trust can be assumed when using private cloud. The trust is weaker in the case of public cloud. Encryption of content is a way to reduce this risk. Nevertheless, this is only true if the public cloud administrators (i.e. third party workers who are not screened and controlled by the cloud service customer), have no way of obtaining the decryption keys. The cloud service provider's administrators have access to the hardware and to the hypervisor. For instance, they may dump the memory of a given VM. With such memory snapshot, it is possible to extract confidential information such as passwords or private keys if they were not properly obfuscated (Rocha and Correia, 2011).

7. Abuse of cloud services. This does not concern content protection.
8. Insufficient due diligence. This does not primarily concern content protection. Nevertheless, there are currently no industry guidelines on how to protect content in the cloud. The current MPAA guidelines do not tackle this issue. Even some best practices are not suitable for cloud. Thus, implementing content protection strategies becomes an exercise of interpretation and interpolation of existing guidelines by the designers.
9. Shared technology vulnerability. For content protection, it is important to ensure cryptographic isolation between users sharing the same cloud service -- for instance, a digital screener portal. Whenever an attacker gains access to user A's content, she should not gain access to user B's content, even if they share the same volume. If an attacker compromises the access control list protection of the file system, cryptographic isolation should remain as a second line of defense.

CSA's list is not exhaustive. It lacks many threats. This part highlights some additional threats that are especially relevant to content protection.

10. Data integrity. It is paramount to put in place measures that prevent attackers from modifying content. As a backup, measures should be in place to detect and determine whether tampering has taken place. Unnoticed distribution of a tampered content may tarnish service provider's reputation and of its customers.
11. Media sanitization. Hard drives have to be disposed of at the end of their life. If the cloud provider does not securely erase or destroy them, attackers, or a second-hand buyer may find some treasures. In a private cloud, the service provider may ensure that its IT team uses proper sanitization methods (Kissel et al., 2012). This verification is not always possible by service provider in the case of public cloud.
    The same problem exists whenever a service provider resigns from a cloud provider, or even if the service provider deletes files rather than securely wiping them out. The service provider may not trust the cloud provider to clean storage that is not anymore allocated to a customer.
12. Weak implementation. Content security in the cloud is a new paradigm which requires new security postures. The media industry is at the foot of the learning curve. Without clear guidelines, implementers will not know whether the implemented security measures are sufficient. Current MPAA guidelines are not adapted to cloud implementation (and especially for public cloud) (Anon, 2013). Some cloud service providers have announced that they are MPAA compliant. Unfortunately, their scope of compliance is very superficial. These statements of compliance are often mainly about certain elements of the infrastructure. Service providers will have to implement much more measures to provide users in the media industry with proper assurances. As security is only as strong as its weakest link, new media industry-specific guidelines dedicated to cloud computing are needed.

Some interesting attack vectors

The cloud introduces new risks or attacks. This section highlights some of them and analyzes their likelihood for the different cloud deployment models.

The main risk with public cloud comes from multi-tenancy. With virtualization, different principals may share the same hardware server. In the public cloud, these principals may belong to different organizations. In other words, Alice may share the same hardware server with Bob. Unfortunately, Alice does not know Bob and does not trust Bob. Indeed, she may even be unaware that she is sharing resources with Bob. If Bob is malicious, then he may attempt to gain access to Alice's information. There are many potential vectors.

- Co-residence on servers. In 2009, Thomas Ristempart et al. attempted to launch a malicious VM on the same server than the victim (Ristenpart et al., 2009). Once able to co-reside with the target, they mounted different side channel attacks. Then, by measuring parameters such as cache usage, or estimated traffic rates, they inferred some valuable information. Since 2009, much progress have been made in malicious exploitation of co-residence. Side channel attacks are an active domain of research. Co-residence is not a risk in the private cloud as only applications trusted by the service provider will operate. It is a risk in the public cloud.

- Malicious VM; if an attacker can replace a trusted VM with a forged VM, then she has a nice tool. Using covert communication, such as timing channels and storage channels, the forged VM may stealthily communicate confidential information to a remote VM in the public cloud. It is often difficult to distinguish malicious VMs from official ones. The exchange could even occur in official cloud images repositories (Wei et al., 2009).

- Rootkits and hypervisor. The primary security assumption of cloud is that the hypervisor enforces total isolation between VMs. If the hypervisor is corrupted, then the security may be annihilated (Wang and Jiang, 2010). Rootkit attacks similar to old Blue Pill (Rutkowska, 2006) have broken this protection. Weak implementation of running applications with buffer overflow attacks may allow this type of malware to succeed. This threat is valid for all cloud deployment models because a malevolent user may inject the malicious payload. Users are not necessarily trustworthy, even in private cloud deployment.

Use cases

Securely delivering content

Digital delivery of screeners is a promising use case for cloud deployment. Already some commercial services use the cloud. Digital Rights Management (DRM) is mature and can be easily ported to the cloud (Diehl, 2012). Nevertheless, the most technically challenging aspect of digital screeners is the requirement to provide an invisible forensics mark unique to the viewer or to the session. In the remainder of this document, watermarking means an imperceptible (invisible or inaudible) watermark. This paper does not address visible watermark or burn-in. When referring to payload, it means the information carried by the watermark.

The first generation of video watermarking technologies operated in the clear baseband domain. Operating in the baseband domain implies lengthy calculation on clear content. Figure 1 describes a simplified process for secure delivery of content. The clear content is first watermarked for Alice. Then the watermarked content is encrypted. The encryption may be performed by using a DRM packager (such as PlayReady) or by using a point-to-point encryption mechanism, such as SSL/TLS or HLS.

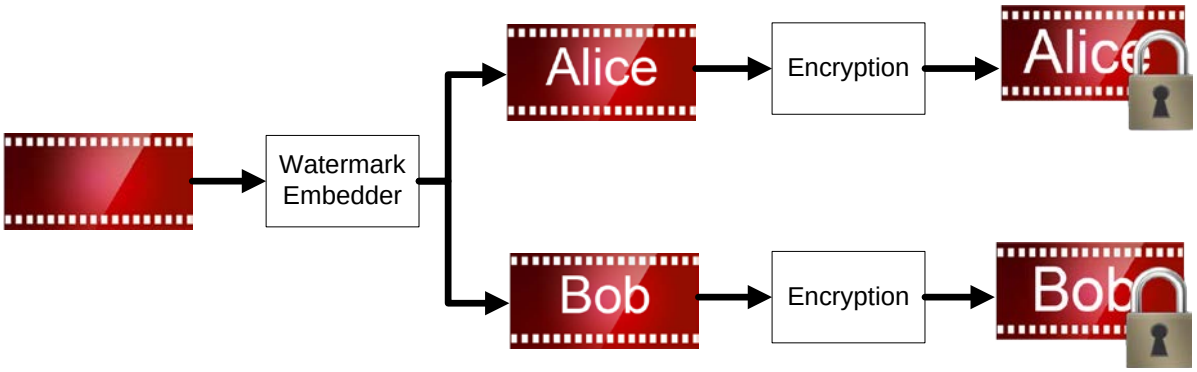

Figure 1: Watermarking in the clear at head end

This architecture has three issues from the point of view of content protection.

1. The most obvious issue is that content is in the clear until delivery time. The least secure point of the workflow is when content is in the clear at rest. If the system operates in a private cloud, then the risk may be acceptable. Unfortunately, at least from a financial perspective, delivery is better suited in the public cloud than in the private cloud. If the system has to be in the public cloud, content should be encrypted when it is at rest. This encryption means that the described architecture should evolve with a decryption prior to watermarking. In this case, the data path from decryption to ultimate encryption should be protected to the maximum levels possible. For instance, ensuring that other tenants do not share the same physical servers would limit the risk of side channel attacks. Unfortunately, this additional protection increases the delay between the request for a piece of content and its actual delivery. From a user experience perspective, this delay should be minimal.
2. The second issue is about the non-repudiability of the watermark. For the watermark to be non-repudiable, it is paramount to prove that Eve can access Alice watermarked content only from Alice's rendering device. This architecture does not allow this demonstration. Within the workflow, watermarked content is in the clear. Eve may have stolen Alice watermarked content just before encryption. For instance, Eve could be an insider. This weakness is due to the embedding in the clear rather than to cloud implementation.
3. The third issue is that the scheme assumes that the watermark embedding cannot be bypassed. If embedding were bypassed, non-watermarked content could be delivered. Traceability would be lost. The workflow may enforce mandatory processes. However, this enforcement would require additional mechanisms.

Over the past few years, a new generation of video watermark technologies has appeared. The principal characteristic is that watermark embedding occurs in the compressed domain (Robert et al., 2014). Embedding is faster and requires less computation power than the previous generation. Operating in the compressed domain reduces the delay introduced by watermark. Thus, it enhances the user experience. Unfortunately, it does not enhance security compared to the previous scheme.

Some of the watermark technologies of this new generation support watermarking in the encrypted compressed domain. Figure 2 describes a simplified process for securely delivering content in this context. The clear content is first pre-processed. Indeed, this new generation of watermark -- sometimes called two-step watermark -- prepares the content to be watermarked. This preparation may occur offline and requires massive calculation. The actual embedding occurs later and can be done on the fly. Once the content is preprocessed, it is encrypted. The

optimal solution is to use DRM. At delivery time, the encrypted content is personalized for the recipient. The watermark embedder adds its payload in the encrypted bytes without decrypting content. The encrypted watermarked content can then be delivered to the user.

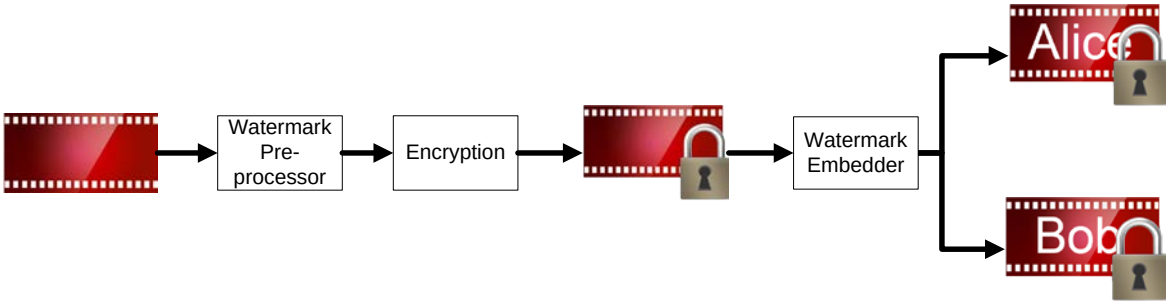

Figure 2: Watermarking in the encrypted domain

This architecture partly solves the previous issues.

1. Content is in the clear during the preparation phase. Watermark preparation may be part of the ingest process. Once encrypted, content is secure. A hybrid cloud is suitable. Preparation and encryption occur in the private cloud, thus benefiting from its higher security and trust status. Once content is encrypted, it is safe to post in the public cloud. Encrypted content is useless to an attacker without the corresponding decryption keys. Furthermore, encryption also partially protects the integrity of content. Even if an attacker gets access to stored content, she cannot alter genuine content. Most attempts would end up in wrecked content that could not be decrypted. The attacker would just have generated a DoS attack. Nevertheless, as the attacker would have access to content, she would have reached the same result by just deleting stored content. Watermarking can occur in the public cloud without weakening the trust model. Deletion would be an easier attack.
2. Non-repudiability of watermark is stronger in this scenario. Since the clear watermarked content exists only after decryption, Eve needs Alice's decryption keys to access Alice watermarked content. This decryption is only possible with the license delivered to Alice's device or from Alice device.[3] In all other locations of the workflow, the watermarked content is always encrypted. Therefore, if Alice's watermarked material is found in the wild, there are only three potential scenarios:
    - Alice posted the clear content after camcording or screen scrapping it.
    - Alice collaborated with Eve and passed her credentials to Eve, who posted the pirated content.
    - Alice's machine was compromised by Eve. Forensics analysis of her device could confirm (or not) this hypothesis.
3. As for the previous architecture, if the workflow is designed correctly, it is reasonable to assume that all content will be watermarked. Even if embedding occurs in the public cloud, we may expect that it will be difficult for an attacker to modify the implemented workflow.

Figure 3 describes a real example of such hybrid deployment. The schematic focuses on the content protection aspects of the deployment. Microsoft PlayReady™ provides content protection for DRM and Technicolor ContentArmor™ Video Watermark ensures forensics marking.

[3] We assume that the DRM protecting the license is robust.

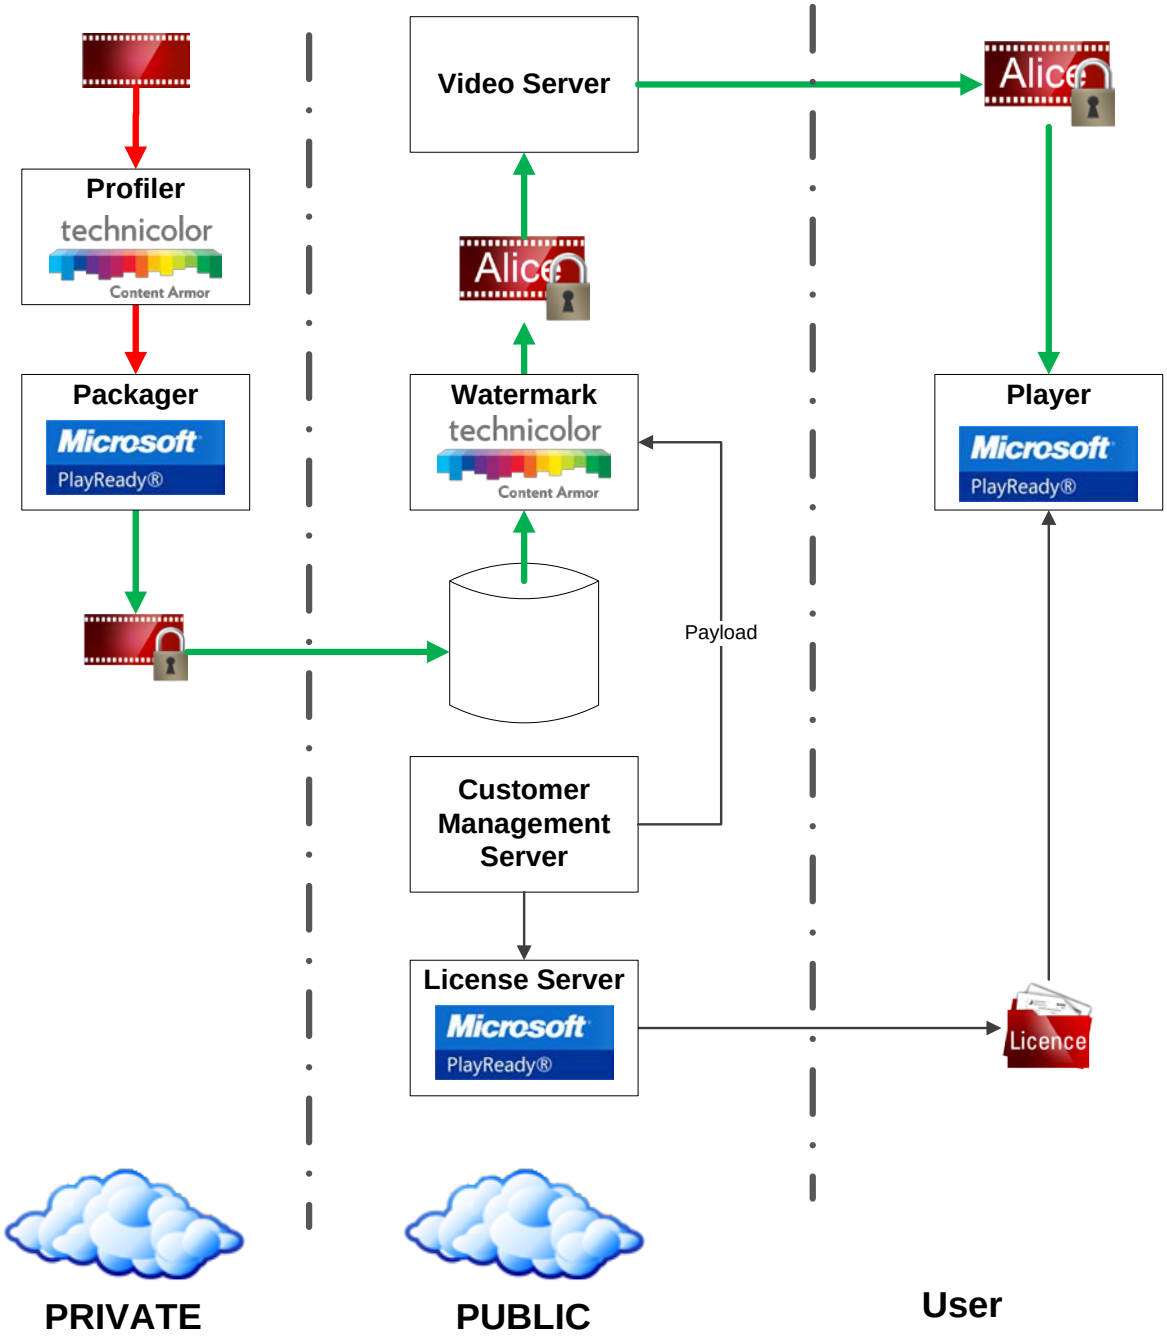

Figure 3: E-screener with watermark at head end

The ingest phase occurs in a private cloud. The clear H264 content is passed to the ContentArmor profiler. The profiler analyzes the input content and extracts watermark forensics metadata (WEM). This operation occurs only once for each piece of content. PlayReady's packager encrypts the H264 content. These operations may take some time and usually have no real time constraints.

The PlayReady-protected content and WEM are pushed on to a public cloud storage resource. The following operations will occur in the public cloud. When a user requests a digital awards screener, the Customer Management Server verifies whether the user is authorized to watch

this screener. If authorized, the corresponding content is extracted from the data storage and passed to ContentArmor embedder. The Customer Management Server defines the payload to embed in this instance of content. The payload may be unique to the user, may be unique to the session, may contain a time stamp, or any other suitable configuration. Using the WEM, the ContentArmor embedder inserts the payload defined by the Customer Management Server into the PlayReady protected content.

This operation is extremely fast. It can even be done during the transfer from the storage to the video server, thus introducing no noticeable delay. The video server streams PlayReady protected content to the player. The PlayReady license server delivers a PlayReady license to the player, defining the usage rights associated with this screener.

The trust model of this architecture is robust. It mainly relies on three hypothesis:

1. Operations in the private cloud are secure.
2. PlayReady is secure.
3. ContentArmor Video Watermark is secure.

We will analyze the robustness of these hypotheses.

Clear content is available only in the private cloud. Thus, if the first hypothesis holds, attackers have not access to clear content during processing. According to Table 1, hypothesis 1 is valid.

In the public cloud, and during transfer from the private cloud to the public cloud, and from a video server to the player, the e-screener is always protected by PlayReady. As the content owner selected PlayReady as the DRM for this solution, this means that the content owner trusted this solution. Therefore, hypothesis 2 is valid.

Once decrypted, clear watermarked content can be recorded either by camcording or screen scrapping. ContentArmor Video Watermark offers robust protection against these recording means. Thus, hypothesis 3 is valid.

Therefore, this architecture enables the distribution of pre-theatrical release e-screeners without compromising the security level while still securing the benefits and advantages of using public cloud resources.

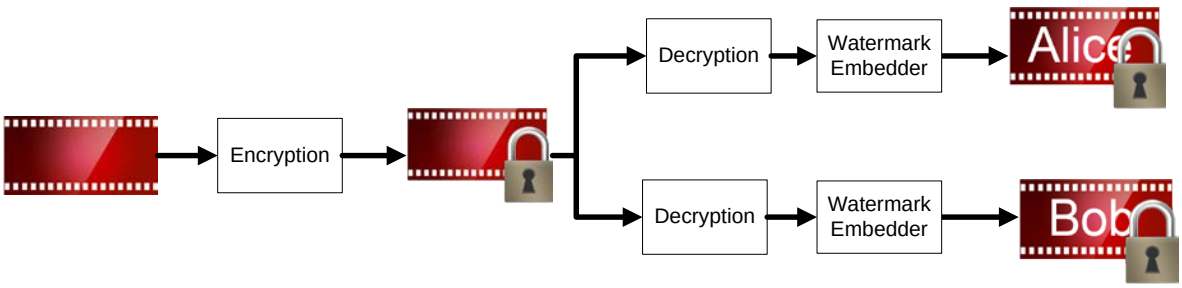

Figure 4: Watermarking at player

For the sake of completion, it is important to address another potential design. Watermarking may occur at the rendering point rather than at the head end. Figure 4 illustrates this architecture. The clear content is encrypted. The optimal solution is to use DRM. At rendering time, the encrypted content is personalized for the recipient. The watermark embedder adds its payload in the clear content -- or before decryption if the watermark technology supports this feature. The advantage of this architecture is that the distribution may benefit from a Content Distribution Network (CDN) since all recipients receive the same piece of encrypted content.

Previous solutions cannot benefit from the caching feature of CDNs because each recipient receives his personalized instance of content.

This architecture partly solves the two first issues.

1. Content is in the clear during ingest phase. Once encrypted, content is secure. Then, it is safe to post it in the public cloud. This scenario presents the same advantage as the previous architecture.
2. Non-repudiability is similar to the previous architecture. Only Alice can access Alice-watermarked content.
3. Since embedding occurs in the user's platform, it is not safe to assume that watermark embedding cannot be by-passed. Indeed, the attacker may entirely control the user's platform and modify the application running on the user's platform. Nevertheless, some watermark technologies offer implementations that enforce the need for the content to go through the watermark embedder. If the content by-passes the embedder, then this piece of content cannot be rendered.[4]

Processing content

Transforming content in the cloud would allow users to benefit from on demand self-service, rapid elasticity and scalability. Processing has two aspects: content storage and content transformation.

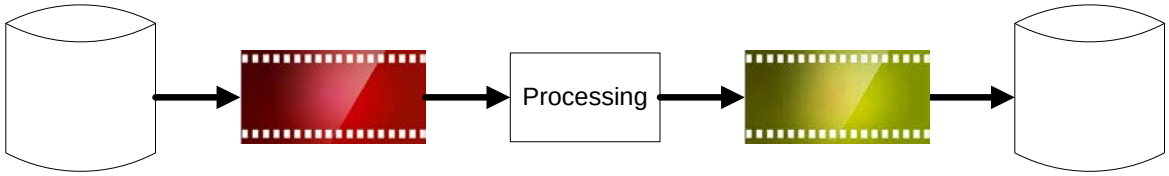

Figure 5: Processing

Figure 5 illustrates the simplified architecture. The piece of content is stored, transferred to an application that transforms it, and the transformed content is stored back. From the point of view of content protection, the major risk is theft. Content is in the clear at rest and during transfer before processing and after processing. This risk is low if all operations occur in a private cloud, but high in a public cloud.

If some elements of the workflow are in a public cloud, then encryption is most probably mandatory. Figure 6 provides the corresponding architecture. Content is encrypted at rest. Prior to processing, content is decrypted. Once processed, the transformed content is encrypted. Several implementations are possible. This section explores some of them.

[4] ContentArmor Video Watermark proposes such anti-bypass enforcement.

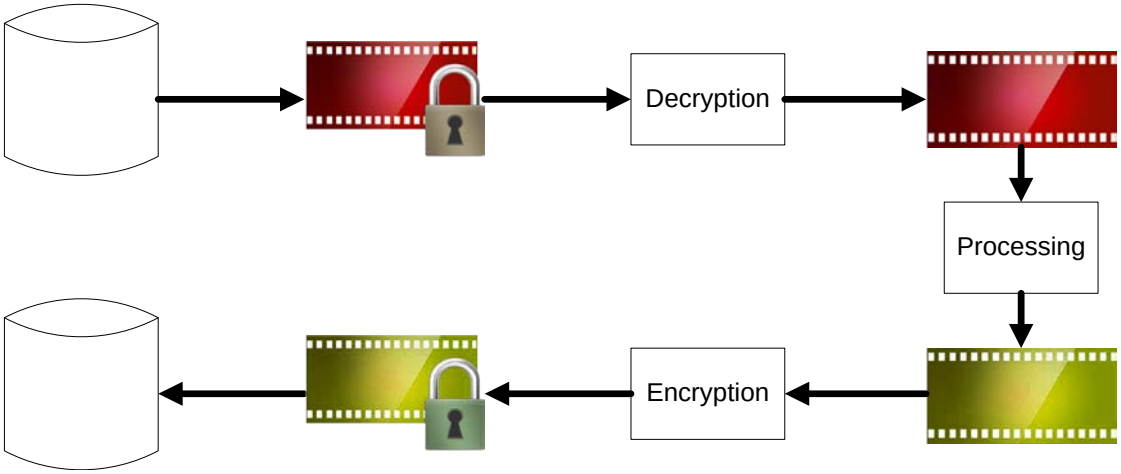

Figure 6: Processing and encrypted content at rest

In the first implementation, processing and storage occur in the public cloud. The following points offer security tricks to carefully fine tune this process:

1. Key management will occur in the public cloud. Special care should be applied to secure to the maximum extent possible. For instance, some cloud service providers propose cryptographic services using physical Hardware Secure Modules (HSM). HSM has the advantage of storing, the master keys in a secure fashion and of performing cryptographic operations without exposing secure keys.
2. Ideally, decryption and encryption are integrated within the processing application, thus limiting exposure of clear content.
3. If possible, all content should never be decrypted in one block. Partial decryption is preferred. Thus, no temporary file may contain all of the clear content – whether or not it has been transformed. Temporary files may not be wiped out properly.
4. The complexity of the design resides in the key management. It raises several questions, such as: How to ensure that only the right processing process can trigger decryption (Diehl et al., n.d.)? How to manage the keys used to encrypt the transformed content? Is the processed content encrypted with the same key as the original content?

In the second implementation, processing occurs in the private cloud whereas storage is in the public cloud. The second implementation should be more secure than the previous one.

1. Key management occurs in the private cloud. Thus, it is more secure. Nevertheless, special care is recommended as it handles secret material.
2. Decryption and encryption are performed in a private cloud, i.e., in a trusted execution zone.
3. Partial decryption is no longer preferred. Nevertheless, properly wiping out temporary files and memory buffers that may contain passwords and secret keys is part of the best practice.
4. The complexity of key management is the same in both implementations.

For the sake of completion, this section presents an additional architecture based on the use of homomorphic encryption. Homomorphic encryption is a kind of encryption that conserves operations. Thus, an operation applied to two encrypted data results in new data that once decrypted is the operation applied to both sets of clear data. Equation 1 illustrates this remarkable feature.

Equation 1: full homomorphic encryption

$$D\big(E(a)\ \mathrm{op}\ E(b)\big) = a\ \mathrm{op}\ b$$

This type of encryption opens opportunities in a cloud environment (Naehrig et al., 2011). It could enable secure transformation -- even in a public cloud. Figure 7 illustrates an architecture that would use homomorphic encryption. Because content is encrypted and stored at rest, the transformation would handle only encrypted data. This operation could securely operate in the public cloud.

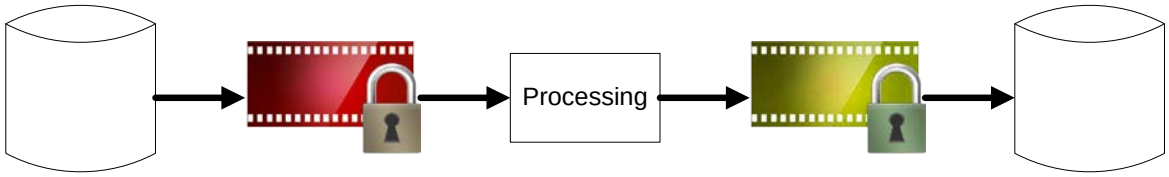

Figure 7: Processing with homomorphic encryption

Unfortunately, current homomorphic cryptosystems are not practical for complex calculations such as those associated with typical video production. Experts estimate that several decades will be needed to design full homomorphic encryption cryptosystems that will handle large calculations efficiently. For instance, a google search on encrypted request would today take a trillion times longer than the current search time (Greenberg, 2009). Even for basic operations such as addition, homomorphic encryption is far slower than normal encryption and requires much longer keys.

Conclusion

Cloud computing is changing the world. This wave of change will soon impact the media production and post-production sector. Content protection is one of the major challenges to solve in order for the media industry to adopt – and benefit from – cloud computing. Cloud dilutes the current trust model. Service providers outsource many security responsibilities to cloud providers and their cloud partners. This delegation is a problem since the service provider remains the ultimate responsible for liability in the event of a content leak.

Hybrid cloud is a smart approach that mitigates the risks associated to cloud. Private cloud handles content while content is in its most vulnerable state; public cloud handles content once it is reasonably protected. This hybrid deployment offers a balanced trust model. New developments in technologies, such a watermarking in the encrypted domain, will facilitate this transition towards cloud.

Nevertheless, cloud computing is just in its infancy. Once in the cloud, content will be under fire. Pre-theatrical content is more monetizable than nude pictures of celebrities. It will attract high-profile, well-funded attackers. Currently, the cloud industry is at the beginning of its learning curve. So are the attackers. The endless arm race is starting. Content owners and technology providers should quickly work together to issue best practices for content security in the cloud. As the industry and the attackers climb the learning curve, these best practices will need regular updates and take into account new technologies and new attacks.

References

Anon, "10 Immutable Laws of Security," Microsoft TechNet, available at: http://technet.microsoft.com/library/cc722487.aspx. 2014a.

Anon, "Content Security Best Practices: Common Guidelines V2.1," Jan. 2013.

Anon, "Cyber Supply Chain Risks, Strategies and Best Practices," in Priorities for America's Preparedness: Best Practices from the Private Sector. U.S. Resilience Project, 2012.

Anon; "The heartbleed bug;" available at: heartbleed.com. 2014b

Anon, "Y.3500 - Cloud computing ─Overview and Vocabulary," ITU, Aug. 2014c.

S. Byers, L.Cranor, D. Korman, P. McDaniel, and E. Cronin, "Analysis of security vulnerabilities in the movie production and distribution process," *Proceedings of the 3rd ACM Workshop on Digital Rights Management,* pp. 1–12 2003.

N. Carr, "The Big Switch: Rewiring the World, From Edison to Google," W. W. Norton & Company, 2008.

R. Clarke, "Trust in the Context of e-Business," Internet Law Bulletin, 4(5), Feb. 2002.

Diehl, E. "Securing Digital Video". Springer Berlin / Heidelberg. 2012.

E. Diehl, A. Durand, and S. Onno, US Patent No. 20100146298 A1, Nov. 26, 2008, "Method and system for processing digital content according to a workflow".

A. Greenberg, "IBM's Blindfolded Calculator," Forbes, Jun 2009.

Y.Y. Haimes, B.M. Horowitz, Z. Guo, E. Andrijcic, and J. Bogdanor, "Assessing Systemic Risk to Cloud Computing Technology as Complex Interconnected Systems of Systems," Jan. 2014.

R. Kissel, M. Scholl, S. Skolochenko, and X. Li, "Special for Publication 800-88: Guidelines for Media Sanitization Rev 1," Sep. 2012.

J. Leyden, "Fiendish CryptoLocker ransomware: Whatever you do, don't PAY," The Register, Oct. 2013.

R. Los, D. Shackleford, and B. Sullivan, "The Notorious Nine: Cloud Computing Top Threats in 2013," Feb. 2013.

P. Mell, and T. Grance, "The NIST Definition of Cloud Computing," NIST, Sep. 2011.

M. Naehrig, K. Lauter, and V. Vaikuntanathan, "Can Homomorphic Encryption Be Practical," *Proceedings of the 3rd ACM Workshop on Cloud Computing Security Workshop CCSW '11*, pp. 113–124 2011.

T. Pott, "Ransomware attack hits Synology's NAS boxen," The Register, Aug. 2014.

T. Ristenpart, E. Tromer, H. Shacham, and S. Savage, "Hey, You, Get off of My Cloud: Exploring Information Leakage in Third-party Compute Clouds," *Proceedings of the 16th ACM Conference on Computer and Communications Security. CCS '09*, pp. 199–212 2009.

A. Robert, O. Alvarez, and G. Doerr, "Adjusting bit-stream video watermarking systems to cope with HTTP adaptive streaming transmission," *IEEE International Conference on Acoustics, Speech and Signal Processing ICASSP '14*, pp. 7416–7419 May 2014.

F. Rocha, and M. Correia, "Lucy in the sky without diamonds: Stealing confidential data in the cloud," *IEEE/IFIP 41st International Conference on Dependable Systems and Networks Workshops DSN-W '11*, pp. 129–134 Jun. 2011.

J. Rutkowska, "Subverting VistaTM kernel for fun and profit," Black Hat Briefings, 2006.

Z. Wang, and X. Jiang, "HyperSafe: A Lightweight Approach to Provide Lifetime Hypervisor Control-Flow Integrity," *IEEE Symposium on Security and Privacy SP '10*, pp. 380–395 May 2010.

J.Wei, X. Zhang, G. Ammons, V.Bala, and P. Ning, "Managing Security of Virtual Machine Images in a Cloud Environment," *Proceedings of the 2009 ACM Workshop on Cloud Computing Security. CCSW '09*, pp. 91–96 2009.