CSAP

COMMON SECURITY ARCHITECTURE
for PRODUCTION

RECOMMENDED PRACTICES FOR THE DEPLOYMENT OF ZERO TRUST IN MEDIA PRODUCTION

Contents

This document is intended as a guide for companies developing or implementing products, solutions, or services for the future of media creation. No effort is made by Motion Picture Laboratories, Inc. to obligate any market participant to adhere to the recommendations in this document. Whether to adopt these recommendations in whole or in part is left to the discretion of individual market participants, using independent business judgment. Each MovieLabs member company shall decide independently the extent to which it will utilize, or require adherence to, these recommendations. All questions on member company adoption or implementation must be directed independently to each member company.

1 Introduction

The MovieLabs 2030 Vision says that cloud services are common resources shared across users engaged in a production, be that the production company, the studio, VFX houses, finishing houses, small specialist providers and individual contributors. That presents a very different security challenge from production infrastructure that is owned[1] by one entity and most users are its employees.

Securing cloud infrastructure using a traditional security perimeter is somewhere between complex and impossible especially when the cloud infrastructure is used for the workflows of the 2030 Vision.

> *The enterprise perimeter is no longer a location; it is a set of dynamic edge capabilities delivered when needed as a service from the cloud.*
>
> The Future of Network Security Is in the Cloud
> Neil MacDonald et al, Gartner, August 30, 2019

At the same time, traditional security perimeters are failing even when used to protect private data centers. The premise of perimeter security is that good actors are on the inside and bad actors are on the outside but if the perimeter is breached, that premise is false. With cloud production, there is no inside and outside.

The cybersecurity landscape is shifting with increasing speed from perimeter security to zero trust security architectures. Zero trust is a security architecture far more suited for protecting cloud resources than perimeter security. It starts with the belief that nothing should be implicitly trusted either inside or outside of any security perimeter. Instead, the rule is to verify anything and everything trying to connect before authorizing access.

US National Institute of Standards and Technology (NIST) Special Publication 800-207, Zero Trust Architecture[2] presents seven tenets of the zero-trust architecture. While we are largely quoting from SP 800-207, MovieLabs have substituted certain terms for the CSAP equivalent and in each case the original term is in the footnote. The seven NIST tenants are:

1. All data sources and computing services are considered resources.
2. All communication is secured regardless of network location. Network location alone does not imply trust.

[1] We use the term “owned” to include “leased” and “rented.”
[2] National Institute of Standards and Technology, Zero Trust Architecture, NIST Special Publication 800-207, Abstract: https://csrc.nist.gov/publications/detail/sp/800-207/final, PDF: https://doi.org/10.6028/NIST.SP.800-207.

3. Access to individual enterprise resources is granted on a per-session basis. Trust in the requester is evaluated before the access is granted. Access should also be granted with the least privileges needed to complete the task.
4. Access to resources is determined by *authorization*[3] policy - including the observable state of client identity, application/service, and the requesting *resource*[4] – and may include other behavioral and environmental attributes.
5. The enterprise monitors and measures the integrity and security posture of all owned and associated *resources*. No *resource* is inherently trusted. The enterprise evaluates the security posture of the *resource* when evaluating a resource request.
6. All resource authentication and authorization are dynamic and strictly enforced before access is allowed. This is a constant cycle of obtaining access, scanning, and assessing threats, adapting, and continually reevaluating trust in ongoing communication.
7. The enterprise collects as much information as possible about the current state of *resources*, network infrastructure, and communications and uses it to improve its security posture. This data can also be used to provide context for access requests from subjects.

The MovieLabs Common Security Architecture for Production (CSAP) is a zero trust architecture designed to secure content production, protecting assets and the integrity of workflows directly and not by protecting the infrastructure the workflow runs on.

This document presents recommended practices for the deployment of zero trust security in media production, and specifically in a way that that supports the CSAP Zero Trust Foundation (ZTF). The ZTF is a zero trust implementation as might be used in an enterprise adopting zero trust. It is not media production specific, but it does require inclusion of specific functionality in the NIST zero trust architecture that other deployments may not.

[3] CSAP term, in SP 800-207 it is “dynamic policy.”
[4] CSAP term, in SP 800-207 it is “asset”, see Vocabulary section of this document.

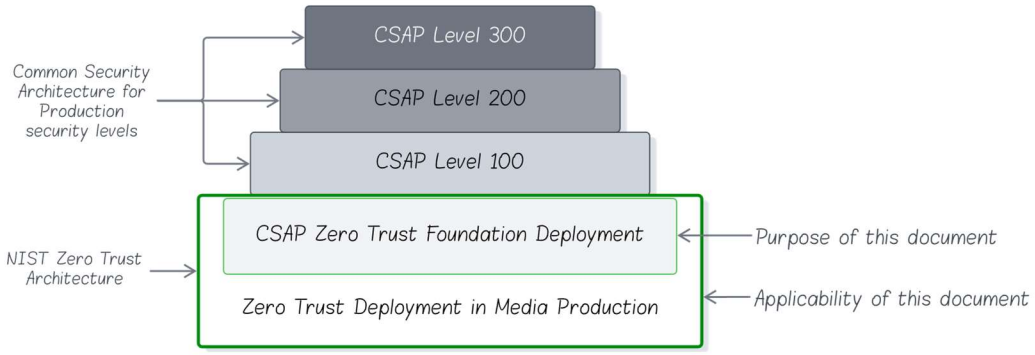

Figure 1 The scope of this document

The required functionality is primarily the logical separation of authentication and authorization, the use of security policies for authorization and identifiable policy enforcement points (PEP). The PEPs may be existing security components of the infrastructure.

- The CSAP ZTF is described in detail in the CSAP documentation, Part 5A.
- The CSAP ZTF does not require a rigorous implementation of NIST's Zero Trust Architecture.

CSAP level 100 is implemented by adding additional functionality to the CSAP ZTF.

1.1 Scope

Zero trust is a philosophy of security management. It improves security because of the way security is managed. It does not replace the need to secure systems. MovieLabs has published recommended practices applicable to the broader subject of cloud or hybrid cloud security and they are to be found in MovieLabs Enhanced Content Protection for Production (ECPP)[5].

This document does not include the recommended practices in ECPP, and those recommended practices should be followed in addition to the recommended practices in this document.

This document does not explain zero trust, nor does it offer a background to the recommended practices. We have endeavored to supply a rich set of external references both in section 1.4 and throughout the document. The reader should be aware that since "zero trust" has become a marketing term of art, not all products described as "zero trust" follow the recommended practices in this document and therefore buyers should do their own investigations to ensure products and services match the recommendations in this document.

[5] https://movielabs.com/prodtech/security/ML_ECPP_v1.0.zip

Some recommended practices are by reference to external sources. We chose this approach because the external source is more authoritative on the topic.

1.2 Formatting

Recommended practice is formatted as bold text in a numbered indented paragraph.

1. **This is an example recommended practice.**

1.3 Terminology

We use the words *must, should* and *consider* in this way:

- *Must* means a recommended practice that is to be followed.
- *Should* means a recommended practice that is optional.
- *Consider* means a practice that does not meet the threshold of being a recommended practice but is nonetheless relevant. For example, if it is snowing outside, *consider* wearing gloves.

There is often a choice of granularity in the implementation of security measures. The appropriate granularity of security controls is, in part, a matter of risk management. For example:

- *Each* user must be authenticated individually. Coarser granularity (e.g., shared accounts) is not acceptable because it undermines the security model.
- *Every* user must be authorized to access a compute resource or an asset. This can be done at an appropriate level of granularity, meaning the authorization is for the individual, the group they belong to or the role they have been assigned and that level of granularity can be chosen as appropriate.

The text endeavors to be clear when there is no choice of granularity. The word *each* is used to mean the maximum granularity is required and the word *every* means that the granularity can be coarser.

In some cases, a recommended practice includes the term *as appropriate* or *feasible.*

As appropriate is a qualifier that means that the recommended practice is to be interpreted in the context of the implementation choices, risk management, organizational security policies, etc. We have endeavored to be explicit as to the conditions for the application of “as appropriate” recommended practices.

The term *feasible* is a qualifier that acknowledges that there are circumstances where the recommended practice cannot be implemented for technical or cost reasons. This cannot be quantized in any absolute terms and is a matter for risk management. Where something is not feasible, any residual risk should be considered.

1.4 Reference and Informative Material

We have collected a series of references for you. These include specifications as well as guidance and informative material.

There are also references to specific topics embedded in the rest of the document.

While we have endeavored to be vendor neutral, vendors of security products are a source of excellent literature on the matters covered here, however any reference to a vendor's documentation is not an endorsement of that vendor.

1.4.1 MovieLabs Publications

MovieLabs has produced a video explaining zero trust and why it is the security architecture of choice for protecting cloud production on both public and private cloud infrastructure:

> Zero Trust and Protecting Cloud Production (video), https://movielabs.com/zero-trust-and-protecting-cloud-production/

The MovieLabs Security Blog series is an introduction to the concepts behind zero trust and CSAP. At the time of publication, three posts have been published and there are more to come (https://movielabs.com/2030-vision-blog/). In order, the three posts are:

1. "Can I Trust You?" https://movielabs.com/can-i-trust-you/
2. "I Don't Trust You, You Don't Trust Me, Now What?" https://movielabs.com/i-dont-trust-you-you-dont-trust-me-now-what/
3. "Am I Authorized to do That?" https://movielabs.com/am-i-authorized-to-do-that/

The Common Security Architecture for Production (CSAP), a 5-part set of documents that describes the architecture and discusses implementation options.
https://movielabs.com/production-technology/production-security/

Enhanced Content Protection for Production (ECPP), recommended practices for the transition from on premises to a hybrid or full cloud infrastructure,
https://movielabs.com/prodtech/security/ML_ECPP_v1.0.zip

Ontology for Media Creation, https://movielabs.com/production-technology/ontology-for-media-creation/

1.4.2 Publications from US and UK Government Agencies

Zero Trust Architecture, US National Institute of Standards and Technology (NIST) Special Publication 800-207, https://doi.org/10.6028/NIST.SP.800-207

Zero Trust and Trusted Identity Management, US National Security Telecommunications Advisory Committee (NSTAC), https://www.cisa.gov/sites/default/files/publications/Final%20Draft%20NSTAC%20Report%20to%20the%20President%20on%20Zero%20Trust%20and%20Trusted%20Identity%20Management.pdf

Advancing Zero Trust Maturity Throughout the Network and Environment Pillar, US National Security Agency (NSA) Cybersecurity Information Sheet (CIS) https://media.defense.gov/2024/mar/05/2003405462/-1/-1/0/csi-zero-trust-network-environment-pillar.pdf

Zero Trust Architecture Design Principles, UK National Cyber Security Centre, https://www.ncsc.gov.uk/collection/zero-trust-architecture

Network Architectures, UK National Cyber Security Centre, (this article describes how zero trust improves security for remote access). https://www.ncsc.gov.uk/collection/device-security-guidance/infrastructure/network-architectures#section_2

1.5 Vocabulary

Even if you are familiar with the vocabulary of cybersecurity, we recommend reviewing this subsection as it defines what we mean by the terms of art used in this document.

Access control is a security technique that regulates who or what can view or use resources in a computing environment.

Access controls are, for the purposes of this document, the means of enforcing access control that are integrated into a resource such as storage.

System/workflow means that which CSAP is protecting other than assets and is used when a more specific term like server is not appropriate or too restrictive.

> For example, a recommended practice is "**everything must be authenticated before it can join the system/workflow**" which might encompass more specific recommended practices such as:
>
> - All users must be authenticated before they can access a service.

- All compute resources must be authenticated before they communicate with another.[6]
- Workflow participants must be authenticated.

Resource: any part of the infrastructure including services that is used by the system/workflow.

Network: a physical or virtual data communications network (a LAN, an IP network, a VLAN, etc.).

authorization policy (lower case): zero trust security policies, called dynamic policies in SP 800-207. In this context, that would mean they are in form appropriate for media production, and generally that means a form that is the same or like a CSAP Authorization Policy.

Authorization Policy (capitalized): collectively CSAP Authorization Policy and CSAP Authorization Rule, as defined in the CSAP Architecture Part 1.

Organizational security policies document required security measures to be taken by an organization's staff.

An organizational security policy is not the same as a "policy" as that term is used in this document (except where the term "organizational security policies" is used), CSAP documents, NIST SP 800-207, and generally in the context of zero trust. In zero trust, a "policy" is a machine executable description of what is authorized such as activity and access to resources and assets.

Asset: an image, video, sound, CG model, metadata, etc. that is typically contained in a file. This is the standard use of the word in media production. Assets is a defined term as part of the MovieLabs Ontology for Media Creation.

> **Note:** NIST SP 800-207 and other security documents use the word "asset" to refer to what this document calls a "resource."

Protect surface: the thing that must be protected, it is also the thing that attackers are after. They are defined by:

- Data that needs to be protected
- Applications that consume sensitive information
- Resources that are the most sensitive
- Services that can be exploited to disrupt operations

[6] This is an example, not a recommended practice. If it were a recommended practice, we would have explained what "communicate" means it the particulate context.

You may wish to explore these sources to find out more about the protect surface concept:

> Define a Protect Surface to Massively Reduce Your Attack Surface, by John Kindervag, September 2018. https://www.paloaltonetworks.com/blog/2018/09/define-protect-surface-massively-reduce-attack-surface/
>
> Defining the Zero Trust Protect Surface, Cloud Security Alliance Zero Trust Working Group, March 2024. https://cloudsecurityalliance.org/artifacts/defining-the-zero-trust-protect-surface (free registration may be required)

Note that a protect surface is *not* the same as an attack surface although an Internet search will return references to both.

Microsegmentation "divides a network into small, discrete sections, each of which has its own security policies and is accessed separately. The goal of microsegmentation is to increase security by confining threats and breaches to the compromised segment, without impacting the rest of the network." Definition by Palo Alto Networks.

> Palo Alto Networks offers an introduction to microsegmentation at https://www.paloaltonetworks.com/cyberpedia/what-is-microsegmentation.

2 Fundamental Recommended Practices

1. **Implement the recommended practices in the MovieLabs Enhanced Content Protection for Production (ECPP)**
2. **Every entity must be authenticated and authorized before it can access a resource or service.**
3. **The deployment should assume that the system/workflow security is in a constant state of breach.**[7]
4. **The deployment should have logically separate origins of authentication and authorization such that it is possible to manage authentication and authorization independently.**

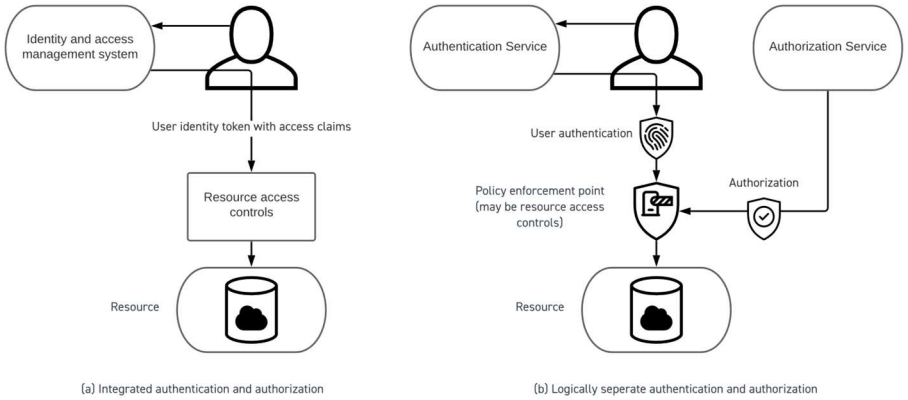

Figure 2 Logical separation of authentication and authorization

5. **Deployment must be "deny by default." No activity should be permitted unless it is explicitly authorized.**

[7] This statement may sound harsh but in most circumstances it's impossible to guarantee there aren't any intruders in a network. The assumption is a core principle of zero trust. An implication of not making this assumption is implied trust of everything on the network which is contrary to the principles of zero trust and experience.

3 Deployment Recommended Practices

The recommended practices in this section apply when zero trust is being deployed.

6. **Protect surfaces must be defined. Each protect surface should be as small as possible.**
7. **Every workflow should be mapped to understand the transactions between tasks, infrastructure, assets, and participants that are part of that workflow.**
8. **The zero trust architecture should be tailored to protect surfaces. Controls should be placed as close as possible to protect surfaces.**
9. **The system must use authorization policies and they should be stored, for example, in identifiable components, such that the current set of authorization policies can be identified.**
10. **All traffic should be inspected. All traffic and transactions should be logged, and an audit trail created by a comprehensive analytics and monitoring systems.**
11. **Analytics should be used to determine the efficacy of authorization policies and to identify normal and abnormal activity.**
12. **Layer 7, the application layer, controls should be used for access to resources and assets wherever feasible.[8] Reliance on security at layers 1-4 may not address threats as effectively.**
13. **The security system should take measures to ensure that authorization policies conform to organizational security policies.**

[8] This is consistent with the goal of CSAP to protect assets and the integrity of workflows rather than the infrastructure.

4 Operational Recommended Practices

4.1 Authentication

14. **Each Participant[9] must be individually authenticated before it is allowed to join or access a workflow/system.**

15. **All devices, systems and services should be authenticated before being used in any workflow. Consider extending authentication to software applications.**

16. **Interactions, for example secure connections, between two devices, systems or services must be mutually authenticated at the same time.**

17. **Any certificate used for authentication, such as device certificates, must be validated against a trusted certificate authority before being accepted. Validation must include checks for revocation.**

18. **A participant or device must not be trusted by virtue of how they are connected.**

4.2 Authorization Policies

19. **Management of authorization policies should be implemented in the simplest way that meets the security goals of the organization.**

20. **Consider automation of authorization policies as an appropriate means to meet the security goals of the organization.**

21. **All workflow activity should be authorized by an authorization policy.**

22. **Authorization policies should follow the principle of least privilege to the extent that is feasible and appropriate in each context.**

23. **Authorization policies should be valid for the shortest period needed to conduct a task or set of tasks where that is feasible and appropriate.[10]**

24. **Authorization policies should be revoked or set to expire when they are no longer needed.**

25. **Automation of authorization policy generation should be protected by zero trust methods, e.g., the source of policies should be authenticated.**

---

[9] For example, a User, but generally as "Participant" is defined in the MovieLabs Ontology for Media Creation.

[10] This adds a temporal dimension to the principle of least privilege.

5 Network Virtualization and Microsegmentation

Network virtualization and microsegmentation may be used to create protect surfaces. That is an implementation choice.

5.1 Microsegmentation

These recommended practices apply where a protect surface is created using a microsegment.

26. **The microsegment should include the minimum number of devices as is feasible and only those that need to communicate with each other.**

27. **Systems inside the microsegment must only communicate with anything outside of the microsegment through a defined security gateway or policy enforcement point.**

5.2 Mutually Authenticated Networks

A mutually authenticated network is a layer 2 (data link), or layer 3 (network layer) method of creating a network where nodes are mutually authenticated[11] before they can join the network. The following recommended practices apply if mutually authenticated networks are being used.

28. **All nodes must be authenticated and authorized before being allowed to join a mutually authenticated network.**

29. **Communication with nodes other than members of the mutually authenticated network must be prevented.**

5.2.1 Software Defined Perimeter (SDP)

SDP is a mutually authenticated network security architecture developed by the Cloud Security Alliance (CSA). The following recommended practice applies if an SDP is being used.

30. **Follow the recommended practices of the Cloud Security Alliance for the Software Defined Perimeter. https://cloudsecurityalliance.org/research/topics/software-defined-perimeter.**

5.2.2 Service Mesh

A service mesh is a dedicated infrastructure layer for handling service-to-service communication. Service meshes are well defined, but no single standard exists. The following recommended practices apply to service meshes.

[11] Exactly how mutual authentication is used depends on the system used.

31. **Services in the service mesh must each have their own service mesh proxy and must only communicate through the proxy.**

32. **Mutual authentication resulting in an encrypted connection must be established between service mesh proxies before further communication can take place.**

33. **Service mesh proxies must be authorized before they are allowed to connect to other service mesh proxies.**

5.3 Virtual Local Area Networks (VLANs)

The recommended practices in this subsection apply if VLANs are being used to deploy microsegmentation.

34. **Devices must be authenticated before communication with a VLAN is permitted.**

35. **VLAN tags must only be assigned to authenticated devices[12] that are authorized to join then VLAN.**

36. **VLAN tags must not be assigned based on parameters such as MAC address that can be spoofed.**

37. **Network components such as switches and routers must be configured such that VLANs are only accessible by authenticated and authorized devices.**

Correctly configuring VLANs requires a very good understanding of network equipment vendors' documentation.

5.4 Virtual Private Clouds

A Virtual Private Cloud is a configurable pool or set of shared resources allocated within a public cloud environment in a way that provides isolation between different organizations or groups of users. When using a virtual private the cloud, the recommended practice is:

38. **Follow the recommended VPC practices of your cloud provider.[13]**

[12] A common protocol for authentication of devices joining a networking is IEEE 802.1X. https://standards.ieee.org/ieee/802.1X/7345/ (This is not a free resource).

[13] For example, https://aws.amazon.com/vpc/ and https://cloud.google.com/vpc/

6 Access Controls

6.1 Resource Access Controls

When resource access controls that are part of the infrastructure are used:

39. **Resource access controls should be managed by authorization policies.**

40. **Resource access controls must be configured such that they deny all access except access authorized in current authorization policies.**

6.2 SaaS Access Controls

These recommended practices apply when using a SaaS service with its own identity and/or access management.

41. **A SaaS service's internal identity management should be integrated[14] with a system wide authentication system, for example the organization's SSO (single or same sign-on) or a federated identity management.**

42. **SaaS native access controls should be managed using authorization policies with an appropriate level of granularity.**

43. **If the native security controls in a SaaS do not meet the security goals of the organization (for example, they can only be set through an administration console), a discrete policy enforcement point should be deployed at all points of I/O to and from the SaaS service.**

[14] We are deliberately vague about what that means because there is more than one way to do it.