



CSAP

COMMON SECURITY ARCHITECTURE
for PRODUCTION

VERSION 1.2

PART 5A:

IMPLEMENTATION CONSIDERATIONS –
STARTING OUT



Contents

- 1 Introduction 1
 - 1.1 How to Use Part 5 1
 - 1.2 Terms 2
 - 1.3 Visual Language Security Icons 3
 - 1.4 Choice of Examples 3
 - 1.5 References 3
 - 1.5.1 MovieLabs Publications 3
 - 1.5.2 Publications from Government Organizations 4
- 2 CSAP Recap 5
 - 2.1 Rules and Policies 5
 - 2.2 CSAP Components 5
 - 2.3 Security Levels 6
 - 2.3.1 Level 100 7
 - 2.3.2 Level 200 7
 - 2.3.3 Level 300 8
- 3 The Path to CSAP Implementation 9
 - 3.1.1 Getting to Zero-Trust 10
 - 3.1.2 Getting to Level 100 11
 - 3.1.3 Getting to Levels 200 and 300 12
- 4 Trust 14
 - 4.1 Trust 14
 - 4.2 Authorization 15
 - 4.3 Analytics 16
 - 4.4 Trust inference 16
 - 4.4.1 BeyondCorp’s Trust Tiers 17

© 2022 Motion Picture Laboratories, Inc.

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

CSAP v1.2 is presented in six parts:

Part 1: Architecture Description the main architecture document.

Part 2: Interfaces describes the possible interfaces between the modules in a canonical form.

Part 3: Security Levels presents a metric-based approach to scaling security.

Part 4: Securing Software-Defined Workflow discusses how the security architecture can be applied to software-defined workflows that are managed using a service bus.

Part 5: Implementation Considerations is broken into sub-documents (5A, 5B, and 5C), which cover different aspects of CSAP implementation. This part is a new addition to CSAP version 1.2.

Part 6: Policy Description describes how policies and rules are defined. This part has not been published as of December 2022.

It is assumed that the reader is familiar with the previously published parts of CSAP, 1 to 4, and we do not reiterate the concepts described in those parts.

1.1 How to Use Part 5

In creating the CSAP architecture, the designers wanted the result to be implementable using existing technologies and with the least development possible. Part 5 provides a perspective on the way the CSAP designers ensured those goals were met. However, the descriptions of implementation approaches may not represent the optimal approach and are not detailed enough to serve as implementation guides.

Part 5 is broken into sub-parts, each of which covers a particular aspect of implementation. This initial release as an addition to CSAP v1.2 consists of three parts that cover some aspects of implementing CSAP, by no means all of them. We expect to release additional parts as developers gain experience in implementing CSAP.

- In Part 5A (this document) subtitled “Starting Out,” Section 2 “CSAP Recap” and Section 3 “The Path to CSAP Implementation” provide general guidance and set the stage. Section 4 “Trust” looks at the meaning of trust, which is a core concept of CSAP, and of zero-trust architectures in general.
- In Part 5B subtitled “CSAP Core,” Section 2 “Identity and the Authentication Service” and Section 3 “Authorization and Authorization Rule Distribution Services” discuss implementation considerations for CSAP core security components. Section 4 “The User Experience” is a lesson in a way to create a good user experience.
- In Part 5C: subtitled “Approaches,” Section 2 “The Network” covers ways in which networks may be used to support CSAP functions. Section 3 “Access Controls” discuss ways access to assets and resources can be controlled. Section 4 “End to End Security” looks at ways that the CSAP architecture can be used to facilitate end-to-end security on untrusted infrastructure.

1.2 Terms

Authentication is the security mechanism used to validate an entity's identity by a trusted authority. The entity might be a user, a service, a device, an application, etc.

Authorization is the security mechanism used by a trusted authority to determine whether an entity can perform an action.

An *Asset* is the broad term we use to mean any data and metadata that is part of the process of media creation including image data, sound data, and metadata. *This is the media definition of the word "asset," and not the definition used in cybersecurity where the word asset means any data, device, or other component (hardware or software) that supports information-related activities.*

The use of *Context* is the normal definition of that word, the setting, circumstances, or environment of an event. The MovieLabs Ontology for Media Creation has a specific and different meaning for *context*, which is not used in this document.

Policies are the abstract representation of what is to be authorized.¹

Rules are the actionable representation of a policy.

A *Device* is a piece of infrastructure in the form of a physical or virtual system that serves as a platform for the execution of software.

Mutual authentication occurs when each entity that is part of forming a trust relationship can authenticate all the others. (As will be discussed later, in this context a user and their system may each be an entity.)

mTLS (mutual TLS) is a form of TLS with additional steps that authenticate the client to the server. In TLS, the server is authenticated to the client, but out-of-band authentication is required to authenticate the client to the server, e.g., using managed device services.

We use the terms for the structure of the organizations creating a creative work that are the common parlance of Hollywood, but they map directly to the terms used elsewhere.

The Studio is the entity that owns the rights to the creative work, is responsible for funding production and has a say (usually creative) in the production process. This is the same role as a commissioning broadcaster or a network (in the way that the term is used when referring to US linear broadcasters.) View the word "Studio" as a shorthand construct for anything that fits the definition.

Depending on the context, the term *The Production* is used to mean either:

- The entity responsible for carrying out production. This may be a production company, an organization set up to produce one creative work or a department or business unit that is part of the studio.

¹ Or, in the very specific case of Global Security Policies, what is to be denied. Global Security Policies are the only place where a "deny" construct is needed since everywhere else, CSAP is deny by default.

Or

- The complete process of producing the creative work.

Vendors are companies that provide services to the production. They may also be called production service providers. Examples are a VFX company, a transportation company, and a cloud infrastructure provider.

Please do not assume that our use of these terms means that CSAP is only for Hollywood studios or only for motion picture production. CSAP is for all types of media production including scripted and unscripted television.

1.3 Visual Language Security Icons

The shapes and icons used in the diagrams in this document are part of the MovieLabs Visual Language. Rather than add a key for the security icons to each diagram, we include it here.



Further information on the MovieLabs Visual Language can be found here at www.movielabs.com/production-technology/visual-language-for-media-creation/

1.4 Choice of Examples

In this document we present examples using commercial products and services. The choice of technology vendor, for example a cloud provider, is in no way an endorsement and does not imply that one product is better than another. In fact, where possible, we have steered away from examples where there is only one vendor.

We also discuss Google's BeyondCorp, and when we do, unless we say otherwise, we are talking about the zero-trust security solution Google implemented in their own offices and about which they have published a range of papers. These papers are cited frequently in the literature. Consider this to be an academic reference.

1.5 References

1.5.1 MovieLabs Publications

[The Evolution of Production Workflows](#), MovieLabs, 2020



[The Common Security Architecture for Production](#), Parts 1 to 4, MovieLabs.

1.5.2 Publications from Government Organizations

Zero Trust Architecture, NIST Special Publication 800-207, <https://doi.org/10.6028/NIST.SP.800-207> .

2 CSAP Recap

Before we discuss implementation of CSAP, it is timely to review the components of CSAP and some of the concepts embodied in it. We recommend that anyone who has read previous versions of Part 1: Architecture review version 1.2, since the update to the architecture has changed the process of creating and distributing authorization rules.

2.1 Rules and Policies²

CSAP uses the terms policies and rules in this way:

- *Policies* are the abstract representation of what is to be authorized. They are what is requested by workflow management.
- *Rules* are the actionable representation of a policy. They are what is sent to the Policy Enforcement Points.

The CSAP designers anticipate that, in most cases, rules would be specific to the point of application.³

2.2 CSAP Components

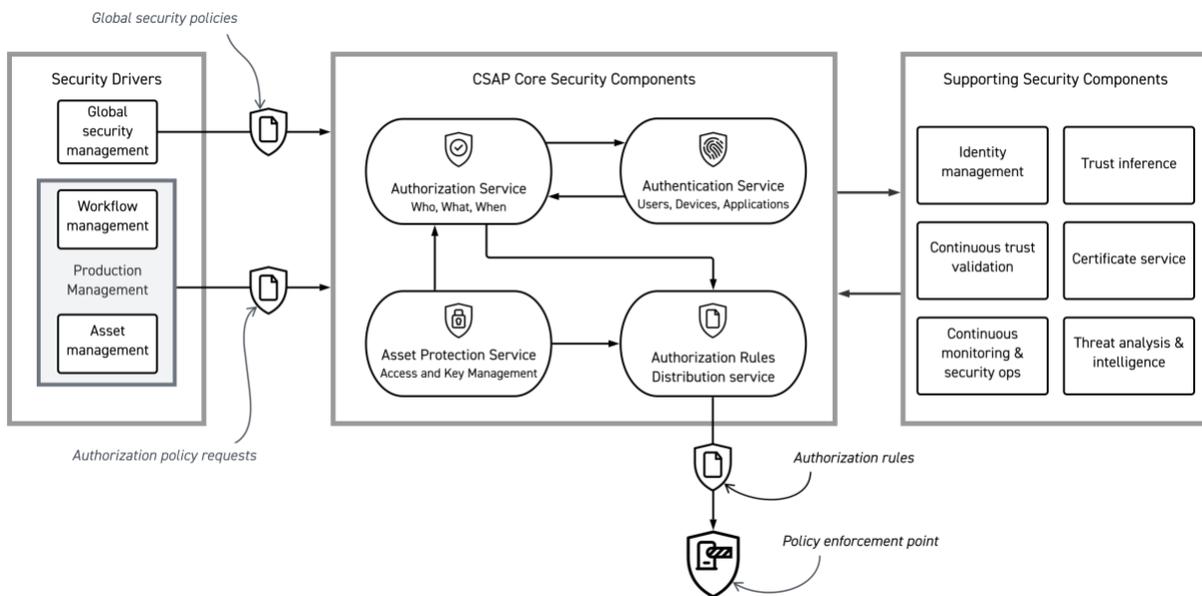


Figure 2-1 CSAP components

Four groupings of components are shown in:

² In CSAP v1.2 Part 1, the architecture was changed from a view where the Policy Enforcement Point interpreted the abstract representation to one that may be infrastructure specific.

³ For example, cloud storage on a particular provider

Grouping	Explanation	Implementation examples
Core security components	Services providing: <ul style="list-style-type: none"> • Authorization service • Authentication service • Asset protection service • Authorization rule distribution service (ARDS) 	<ul style="list-style-type: none"> • Created specifically for CSAP • Commercial solutions customized for CSAP • Aggregation of solutions/services
Supporting security components	Services used by CSAP core components	<ul style="list-style-type: none"> • Commercial solutions configured for CSAP
Workflow management	External to CSAP, this is whatever is driving the workflow and is the source of authorization policy requests	<ul style="list-style-type: none"> • Software-defined workflow manager • Post-production house scheduling tool
Policy enforcement points	The point where CSAP interacts with the infrastructure, applications, workflow, etc.	<ul style="list-style-type: none"> • Service mesh proxy • Endpoint security • SaaS API • Storage ACLs • Asset encryption/decryption

Policy enforcement points are CSAP’s data plane, and they interact with the workflow’s control plane.

Note: this diagram does not address how the CSAP services are themselves protected. For example, the authorization service needs to only accept authorization policy requests from authenticated and authorized workflow management, and only accept global security policies from authenticated and authorized global security management.

2.3 Security Levels

Part 3 of the CSAP specifications describes three security levels, 100, 200 and 300, as a way of illustrating how CSAP security can be scaled according to the risk tolerance of a production.

The security levels are characterized by the availability of core and supporting security components, and the functionality/capabilities that are required to be available. Not all areas of a production will be at the same security level and not all technologies and services need to be capable of level 200 or level 300. For example:

- In a TV series, the security level for the season opener and the season finale might be set higher than for mid-season episodes.
- Certain scenes of a sensitive nature in a motion picture might warrant a higher level of security than has been set for the rest of the production.

Lastly, security can also be scaled, by which we mean the level of security can be turned up and down, by how the components are used. For example, if all access controls on the cloud storage are set according to group membership, it is likely that some members of the group will have more access

privileges than their immediate task requires. CSAP does not set rules for how components are configured.

The tables below list the necessary and beneficial components. Changes from the previous level are bolded.

2.3.1 Level 100

Necessary Core Components	Beneficial Core Components
<ul style="list-style-type: none"> • Authentication service • Authorization service <ul style="list-style-type: none"> ○ System wide capability for long lifetime authorization rules • Asset protection <ul style="list-style-type: none"> ○ Access controls • Authorization rules distribution system (ARDS) 	<ul style="list-style-type: none"> • Authorization service <ul style="list-style-type: none"> ○ Localized capability for medium and short lifetime authorization rules • Asset protection <ul style="list-style-type: none"> ○ Local asset encryption capability ○ End-to-end asset encryption capability

Necessary Supporting Components	Beneficial Supporting Components
<ul style="list-style-type: none"> • Identity management • Certificate service 	<ul style="list-style-type: none"> • Trust inference • Continuous trust evaluation • Continuous monitoring and security operations (CMSO) • Threat analysis and intelligence

2.3.2 Level 200

Necessary Core Components	Beneficial Core Components
<ul style="list-style-type: none"> • Authentication service • Authorization service <ul style="list-style-type: none"> ○ System wide capability for long lifetime authorization rules ○ Localized capability for medium lifetime authorization rules • Asset protection <ul style="list-style-type: none"> ○ Access controls ○ Local asset encryption capability • Authorization rules distribution system (ARDS) 	<ul style="list-style-type: none"> • Authorization service <ul style="list-style-type: none"> ○ Localized capability for short lifetime authorization rules • Asset protection <ul style="list-style-type: none"> ○ End-to-end asset encryption capability

Necessary Supporting Components	Beneficial Supporting Components
<ul style="list-style-type: none"> • Identity management • Certificate service 	<ul style="list-style-type: none"> • Trust inference • Continuous trust evaluation • Continuous monitoring and security operations (CMSO) CMSO • Threat analysis and intelligence

2.3.3 Level 300

Necessary Core Security Components	Beneficial Core Security Components
<ul style="list-style-type: none"> • Authentication service • Authorization service <ul style="list-style-type: none"> ○ System wide capability for long lifetime authorization rules ○ System wide capability for medium lifetime authorization rules ○ Localized capability for short lifetime authorization rules • Asset protection <ul style="list-style-type: none"> ○ Access controls ○ Local asset encryption capability ○ End-to-end asset encryption capability • Authorization rules distribution system (ARDS) 	

Necessary Supporting Security Components	Beneficial Supporting Security Components
<ul style="list-style-type: none"> • Identity management • Trust inference • Continuous trust evaluation • Certificate service • Continuous monitoring and security operations (CMSO) • Threat analysis and intelligence 	

3 The Path to CSAP Implementation

If you are reading this document, you may be thinking to yourself that CSAP sounds like a good idea, but where to start?

The good news is that CSAP is not an architecture that stands apart from the trends in security. CSAP is aligned with a major shift in the cybersecurity landscape, the move to the advanced protection of zero-trust security architectures. Simply put, CSAP is zero-trust architecture applied to the security of content production.

CSAP is a shift from a location-centric security model (e.g., perimeter) to a data and workflow-centric security model that is a better fit for media production in the cloud. It supports fine-grained security controls between participants, systems, assets, and workflows. This makes perfect sense in a world like the world of the MovieLabs 2030 Vision where either there is no centralized point of control or it is constantly moving.

Current security solutions are extrinsic, they are bolted onto or around the workflow systems. CSAP is security by design and, properly implemented, the security is intrinsic to the workflow. Apart from being a superior approach to security, this is one way that CSAP meets the 2030 Vision principle that security must not get in the way of the creative process.

In CSAP we, conceptually, make a distinction between authentication and authorization and the tenets of CSAP are:

1. Nothing can join any workflow unless it has been authenticated.
2. Nothing can join a specific workflow unless it has been authorized.

This means that everything must be authenticated and authorized before it can join a specific workflow. However, we describe it as we have done because authentication and authorization may not, in fact it is likely that they will not, originate from the same place. This isn't new to CSAP, it happens today. For example, a CG artist is hired and immediately given a task. When they try to log into their workstation, they discover that they are not yet in the corporate SSO system. Another example is that mid-way through a task an editor leaves the production and, as part of off boarding their account in the identity management system is locked.

Authentication means that something has been added to an identity management system used by the production's authentication service. The identity management system might be completely controlled by the production, it might be an SSO used across a larger organization, or it might be federated identity management with multiple owners. For example, someone working on a production may already be in the SSO identity management system of an organization (for example the studio) that the production uses. In that case, the organization would control who or what was in their identity management system.

On the other hand, authorization does come from the workflow management whether when a workflow is set up or as it proceeds. CSAP is workflow driven and does not require administrator intervention to

authorize activity such access to assets. As we showed in our examples above, it is possible for the workflow management to authorize something to be part of a workflow that cannot be authenticated.

There is no “trust, but verify,”⁴ no trusted networks (at least, not at the hardware level) where something is trusted because of the port it is connected to, and so on.

As we noted earlier, CSAP is zero-trust security applied to production. Thus, the journey from conventional or perimeter security to CSAP can be divided into two steps. The first step is the adoption of a zero-trust security model which is foundational to CSAP. The second step is to add the CSAP components and functions required by level 100.

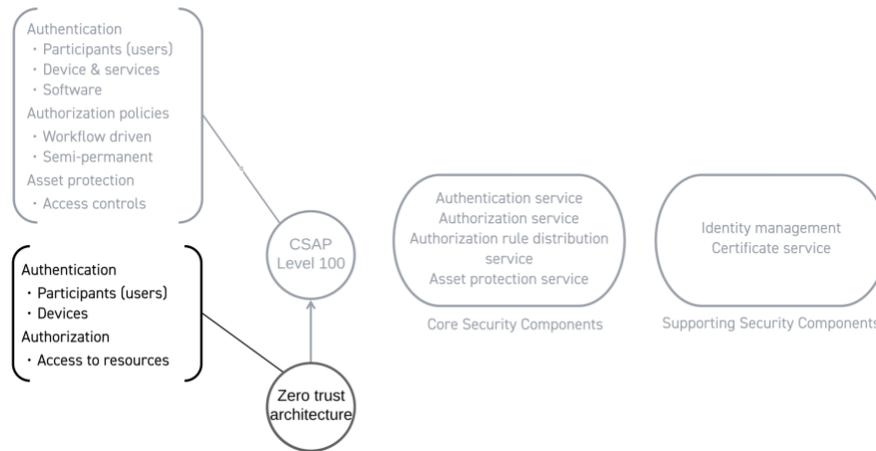


Figure 3-1 Zero-trust is foundational to CSAP

Zero-trust is primarily a change in security philosophy. It doesn’t necessarily require new technology although it may require services that have not been deployed in a particular instance of a traditionally managed security environment.

3.1.1 Getting to Zero-Trust

Describing the step from traditional security to zero-trust is not something this document will address. It all depends on where you start and, more importantly, there are many authoritative sources that can provide guidance and the technology for the transition to a zero-trust security architecture. Here is a short (and incomplete) list:

Transitioning to modern access architecture with Zero Trust, Microsoft,
<https://www.microsoft.com/en-us/insidetrack/transitioning-to-modern-access-architecture-with-zero-trust>

⁴ The phrase means that someone is trusted to do something, but their compliance will be monitored/verified. See https://en.wikipedia.org/wiki/Trust,_but_verify

A unified and proven Zero Trust system with BeyondCorp and BeyondProd, Google Cloud, <https://cloud.google.com/blog/products/identity-security/applying-zero-trust-to-user-access-and-production-services>

Zero Trust on AWS. Advancing your security model with a Zero Trust approach, AWS, <https://aws.amazon.com/security/zero-trust/> (this is set of resources including videos)

Zero Trust. A revolutionary approach to Cyber or just another buzz word?, Deloitte, <https://www2.deloitte.com/content/dam/Deloitte/de/Documents/risk/deloitte-cyber-zero-trust.pdf>

5 Steps to Zero Trust. A simple guide to deploying Zero Trust networks, Palo Alto Networks. <https://start.paloaltonetworks.com/5-steps-to-zero-trust.html> (registration required)

Cisco Zero Trust Architecture Guide, Cisco, <https://www.cisco.com/c/en/us/solutions/collateral/enterprise/design-zone-security/zt-arch-guide.html>

IBM Security Zero Trust Acceleration Services. Accelerate zero trust adoption, IBM Security, <https://www.ibm.com/downloads/cas/DBG68MKM>

Getting Started with Zero Trust Access Management: Trust Begins with Secure Identity, Okta, <https://www.okta.com/resources/whitepaper-getting-started-with-zero-trust-access-management-wbs/> (registration required)

3.1.2 Getting to Level 100

The step from zero-trust to level 100 requires implementing the necessary functionality which of course means that the associated core and supporting components must be available.

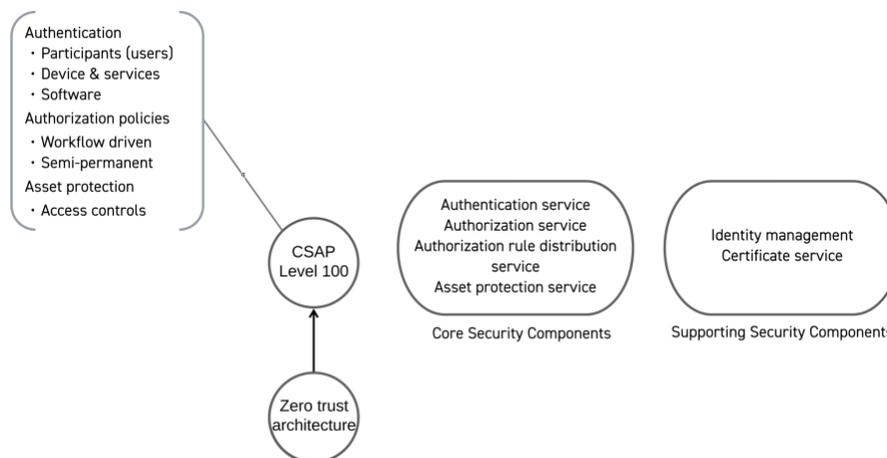


Figure 3-2 Transitioning from zero-trust to CSAP level 100

The core components may be specific services implemented for CSAP or they may be, to a greater or lesser extent, embodied in existing services.

The authorization service, asset protection service, ARDS, and possibly the authentication service are new services that can be run at the production or vendor levels. The service providers create PEPs and can define the actions/templates used by their PEP when those templates are used to create service specific authentication rules. It is the authorization policy request that provides the data to the authorization service that it used to transform those templates into authorization rules.

Here is where we see CSAP’s central enhancement of zero-trust: security, and in particular authorization, is controlled by the production workflow.

3.1.3 Getting to Levels 200 and 300

Whether security needs to meet the CSAP level 200 or 300 requirements is a matter for the production to decide. We anticipate that, initially at least, only parts of a production and certain services and technologies will need to be at level 200 or 300.

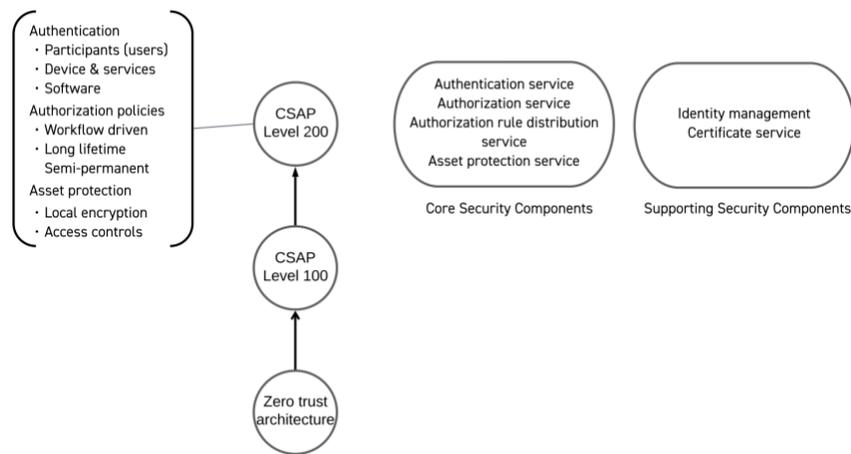


Figure 3-3 Progression from level 100 to level 200

As we move up the CSAP levels we primarily see that there is more granularity in the authorization rules, a shift from assets protected by access controls to assets encrypted individually in a way that is not a property of the storage infrastructure, and more capabilities are end-to-end leading to the end-to-end asset encryption in level 300.

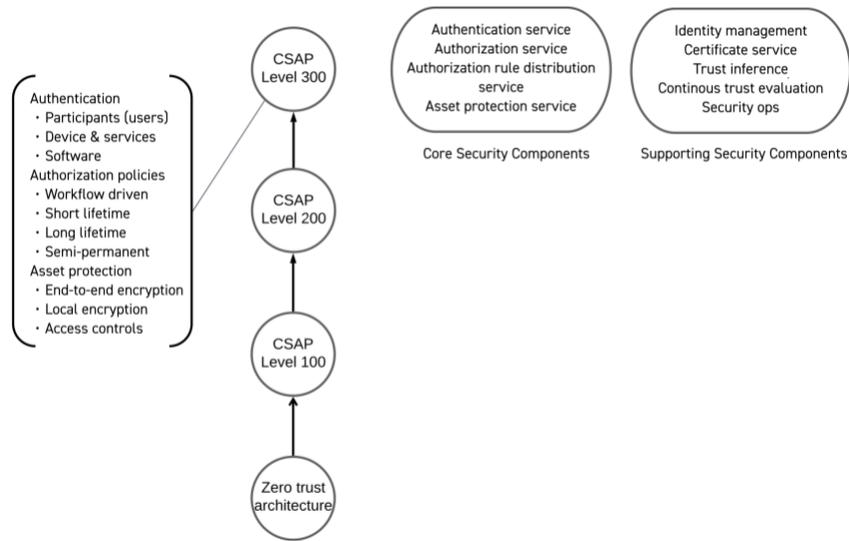


Figure 3-4 Progression from level 200 to level 300

Level 300 also requires the use of additional supporting security components. Some of these serve other security components. For example, continuous monitoring serves trust inference while trust inference itself is required to improve the user experience that may otherwise be negatively affected by the more granular security of level 300.

4 Trust

CSAP is a Zero Trust architecture which means we must first have a common understanding of what trust means.

Mayer, Davis, and Schoorman (1995)⁵ define trust as the willingness of a party to be vulnerable to the actions of another party based on the expectation that the other party will perform a particular action important to the trustor, irrespective of the ability to monitor or control that other party. This is an excellent definition for our purposes because it hints at the consequences of trusting something that is not trustworthy.

There are two factors in creating a trust relationship:

1. Determining whether an entity can be trusted
2. Determining where something claiming to be a trusted entity is indeed that entity and not an impostor

The first of these is outside of the scope of CSAP. It is a decision that is based on factors that vary from one organization to another, from one situation to another, and requires some form of risk assessment.

The second of those is the fundamental role of identity management. It is also important to avoid implicit⁶ trust. For example, a trusted user's device should not be trusted just because that trusted user is using it.

Since trust is central to any zero-trust architecture including CSAP, robust identity management is a prerequisite.

4.1 Trust

If I say I trust you, I probably don't mean I trust you to do anything. I might trust a cardiac surgeon to perform heart surgery, but that doesn't mean I'm going to let them do brain surgery on me. Trust has boundaries.

⁵ Mayer, R.C., Davis, J.H., Schoorman, F.D., 1995. An integrative model of organizational trust. *Academy of Management Review* 20 (3), 709–734.

⁶ This is not the same as CSAP's Trust Inference where many factors such as location can be used to determine a trust level.

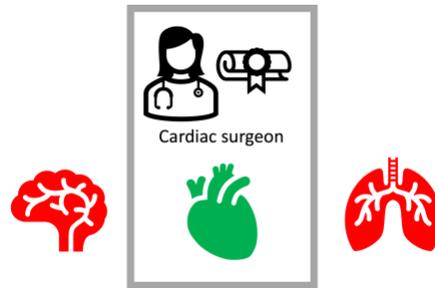


Figure 4-1 A trust boundary

In CSAP, trust boundaries are set by authorization. I verify you are the cardiac surgeon you claim to be⁷ (authentication), and I’m going to let you do heart surgery (authorization.)

4.2 Authorization

Media production workflows are complicated and dynamic and what needs to be authorized can be complicated and dynamic. There are many vendors, large and small, different departments, contract employees, studio employees, etc., all performing workflows that have a limited duration or change frequently. To authorize activity in this constant change is why CSAP has authorization rules that can be created on-demand as workflows progress. Simplistically, a CSAP authorization rule says who can do what using what resources to which assets and when it can be done.

CSAP’s security is variable. Level 300 supports constantly changing authorization rules so that the principle of least privilege⁸ can be applied as literally as required, including temporarily, but its policies can also be derived from the privileges that accompany user authentication in an IAM (identity and access management) system. CSAP does not require the security level to be consistent across the production.

So, yes, implementing CSAP authorization rules at the granularity possible in the architecture may not be easy right now, in part because it requires tight integration with currently nascent workflow management systems. However, there is no reason to start at the deep end. Authorization rules can be created and applied using existing access management systems. That supports the shift in the approach to security but draws on what is already in place.

All that having been said, separating out authentication and authorization is an essential part of the CSAP architecture, but it is conceptual, and it does not preclude authentication and authorization being implemented using an appropriate IAM system that handles both.

⁷ For example, verifying a hospital ID. Verifying the doctor’s diploma is part of the decision to trust the doctor, but it is not part of the authentication.

⁸ The security principle of least privilege means that the privileges given to a user should be only those that are required to conduct the immediate task and no more.

4.3 Analytics

Analytics is a level 300 requirement.

Zero-trust requires pervasive instrumentation.

The starting point in zero-trust networking is the assumption that the network has been breached. The starting point in perimeter security is that network breaches can be prevented.

In perimeter security, activity analytics are looking for abnormal behavior. It looks for activity that indicates that the network has been breached. Those analytics might be using fingerprints associated with malicious activity. If a fingerprint of malicious activity is the movement of large amounts of data, that's not going work in a production environment where moving large amounts of data is normal. Furthermore, malicious data movement can be buried in the noise floor of legitimate activity.

Analytics in a zero-trust network needs to do more: it needs to learn what is normal. That has two purposes: activity that the analytics determines is normal might not need the same amount of scrutiny; and identifying normal behavior is a key part of trust inference.

This approach does present challenges. Creating a picture of normal behavior while assuming that the network is in a constant state of breach means avoiding polluting the picture of normal behavior with unauthorized behavior assumed to be present on the network. While a thesis on how to solve that dilemma is, fortunately, beyond the scope of this document, we posit that the answer lies in the fact that production workflows are known.

4.4 Trust inference

Trust inference is a level 300 requirement.

Note: We do not anticipate anyone will implement trust inference from scratch. This section describes how it can be implemented but the goal is to show considerations in evaluating trust inference services.

Trust inference and its sibling, continuous trust validation, are assessing whether behavior is normal. It is a property of a zero-trust architecture that it is assumed that the network is in a constant state of breach. Normal behavior of devices and users is learned from past behavior and trust inference is combining that with rules that say, for example, access must never come from a particular set of locations. Conceptually, trust inference uses a trust rating based on a comparison of the context of the sign-in attempt (for example, IP address, device, time of day, etc.) and behavior learned about the user or device.

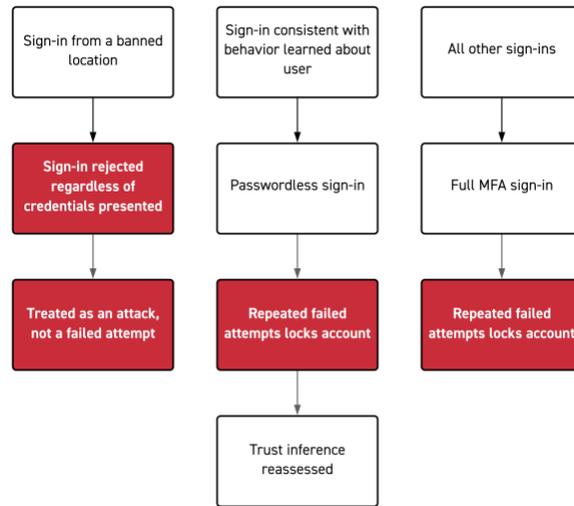


Figure 4-2 Trust inference at work

In our example, Figure 4-2, the circumstances of the sign-in attempt determine the level of credentials required but, to be very clear, the decision to use a reduced level of credential requirements is based on behavior learned about the user and is not solely a factor of where they are logging in from or the device they use.

Implementing trust inference and continuous trust validation requires a set of rules describing the circumstances when sign-in attempts should always be rejected and a system that learns normal user and device behavior, and a processing engine that combines the two sets of data. Learning behavior is not confined to individual users or devices. It can also be role or group behavior. For example, someone performing the duties of an editor would have similar behavior to someone else performing the duties of an editor on the same production.

Continuous trust validation also requires information about a user’s or device’s current activity. See the next section for a description of how BeyondCorp trust inference works, particularly for devices.

4.4.1 BeyondCorp’s Trust Tiers

By looking at how BeyondCorp⁹ manages trust we get key pointers as to how trust inference and continuous trust evaluation can be implemented, the fact that trust does not need to evaluate to yes or no. The level of trust assigned can be a scale and, for example, some access is curtailed if the trust value is below a threshold.

The fundamental components of BeyondCorp are the Trust Inferer, device inventory services, access control engine, access policy, gateways, and resources.

⁹ <https://cloud.google.com/beyondcorp>

Osborn, McWilliams, Beyer and Saltonstall¹⁰ describe the BeyondCorp *Trust Inferer* which is a system that continuously analyzes and annotates device state. The system sets the maximum trust tier accessible by the device and assigns the VLAN to be used by the device on the corporate network. These data are recorded in the Device Inventory Service. Reevaluations are triggered either by state changes or by a failure to receive updates from a device.

BeyondCorp's trust tiers are assigned to each device by the Trust Inferer.

As a device is allowed to access more sensitive data, we require more frequent tests of user presence on the device, so the more we trust a given device, the shorter-lived its credentials. Therefore, limiting a device's trust tier to the minimum access requirement it needs means that its user is minimally interrupted. We may require installation of the latest operating system update within a few business days to retain a high trust tier, whereas devices on lower trust tiers may have slightly more relaxed timelines.

Trust inference for a device comes from two sets of data:

- Observed data generated from data collected such as the last time a security scan was performed on a device, the OS version and patch level, any installed software.
- Prescribed data which is manually assigned and might include the assigned owner of the device, users and groups allowed to access the device, DNS and DHCP assignments and explicit access to particular VLANs.

Data is analyzed from a variety of disparate sources to identify where data conflicts. There is no single or small number of systems that are regarded as the source of truth.

¹⁰ BeyondCorp: Design to Deployment at Google, Barclay Osborn, Justin McWilliams, Betsy Beyer Max Saltonstall, ;login:, ;, vol. 41 (2016), pp. 28-34, <https://scholar.google.com/scholar?lr&ie=UTF-8&oe=UTF-8&q=BeyondCorp%3A+Design+to+Deployment+at+Google+Osborn+McWilliams+Beyer+Saltonstall>

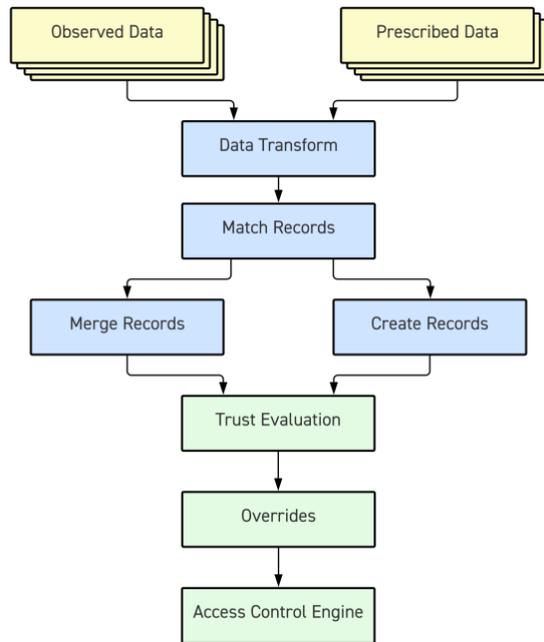


Figure 4-3 The BeyondCorp data processing pipeline (Source: Google)

Here the trust evaluation overrides the access control engine, and this is one way to implement not only initial trust inference but also CSAP’s continuous trust evaluation.

Trust tiers could also be used to assist the CSAP authorization service. As well as assigning a trust tier to each device, each resource is associated with a minimum trust tier required for access. To access a given resource, a device’s trust tier assignment must be equal to or greater than the resource’s minimum trust tier requirement.