ENHANCED CONTENT PROTECTION FOR PRODUCTION (ECPP)

RECOMMENDED SECURITY PRACTICES FOR THE USE OF CLOUD SERVICES IN MEDIA CREATION

VERSION 1.0

Contents

1 Introduction

The transition to production in the cloud is challenging the perimeter security models that have been the foundation of on-premises security. Once cloud services of any form are introduced, the perimeter security model breaks down because of the very nature of cloud services: for example, the cloud service provider controls the hardware and the network[1]. You can't airgap the cloud.

Cloud security is very different from on-premises security.

MovieLabs, as part of its work on the evolution of content creation and its move to the cloud,[2] has developed the MovieLabs Common Security Architecture for Production (CSAP) upon the premise that a different approach is required for securing production in the cloud where the cloud is a resource shared across everyone working on a production. CSAP is equally applicable to the cloud resources that are not shared and on-premises infrastructure.

This document seeks to provide a series of recommended practices to provide guidance in creating and managing cloud security. These recommended practices are not intended to be static, but rather to evolve as the available technology evolves. Although the applicability of recommended practices may vary by situation, MovieLabs recognizes that most of these recommended practices will have broad and strong studio-wide support in most contexts involving the use of cloud resources for the production of motion picture and television content.

In this document the word *production* means the entire production process including pre-production, principal photography, VFX, post-production, sound, mastering, etc.

Each studio will determine individually which practices are required of their suppliers. Those providing production servers and technology should consult with individual studios to determine which recommended practices are required.

Note: the links in this document were correct at the time that it was written but the material referenced, particularly in the cloud provider documentation, may not be the most current information. The reader should confirm the most recent information especially before acting on any of the references.

[1] The cloud providers' shared responsibility models are addressed in a later section of this document.

[2] https://movielabs.com/production-technology/

2 Overview

2.1 Scope

This document offers security recommended practices for the use of public cloud[3] resources either as an extension of on-premises infrastructure or as the sole infrastructure of motion picture and television content development, production, and post-production (collectively referred to in this document as *production*). They apply to use cases where the vendor is in control of the configuration of the cloud infrastructure and security within the boundaries of customer responsibility as described in the cloud provider's shared responsibility model (see Section 5) and the users are their employees or contractors.

This document limits itself to covering only the "cloud deltas" or practices that are new or significantly different for cloud. It does not include practices that should already be in place for productions that use local and on-site infrastructure. Therefore, general IT and network security are out of scope, as are remote work, software development, human resources, site security, and processes for security management. These are all handled in other, well-documented industry practices and are not issues specific to cloud.

This document does not address privacy. Data protection regulations are out of scope. However, cloud providers have plenty of guidance of how to, for example, comply with the GDPR.[4]

2.2 Intended Uses

The recommended practices are intentionally high level. They are intended to be used as one input among others to provide guidance to those who are developing or deploying production workflows and services in the cloud or who are defining more specific and assessable requirements for vendors or for use in assessment programs.

They are not a set of controls and are not intended to be used directly for compliance assessment or for certification. And they are not a set of requirements. Security requirements are a matter for agreement between the customer and vendor.

2.3 Audience

This document is intended primarily for those responsible for ensuring security at:

1. Production services vendors who provide services for TV and motion picture production whether that be large or small organizations or individual contributors. It is their responsibility to provide a secure environment for the work in hand. For brevity, we refer to that group as *vendors.* Examples of members of this group would be post-production companies, VFX houses, sound post services companies and freelance individual contributors.

[3] Meaning a cloud infrastructure available to the general public.

[4] General Data Protection Regulation. https://gdpr-info.eu/

2. Providers of XaaS[5] services used by the vendors defined above. Examples of members of this group would be providers of cloud rendering, and application vendors who are now offering application-in-a-box cloud services.
3. Studios, productions, and others who want to ensure the security of their vendors and service providers, and of their internally developed workflows.

A broader audience would also include those designing, building, deploying, assessing, or configuring production workflows in the cloud.

This document inevitably contains a substantial amount of technical discussion especially from mid-way through Section 2 onward. To fully understand it, the reader needs to be familiar with IT technology such as networks and servers, how cloud services are used to provide functionality for the production process and the technical challenges of doing so. Readers not familiar with information security concepts such as Single Sign On (SSO), access controls in a shared multi-user environment, and how data is protected through encryption might need to research specific aspects as they read this document.

A companion document of a less technical nature will be published.

2.4 Recommended practice labels

2.4.1 Baseline practices

Some of the recommended practices are marked "**Baseline**". These are the "top 5" recommended practices for that section; however, this is from the perspective of protecting your customer's data. There are recommended practices, such as those around response and recovery, that are more to do with you protecting your business.

2.4.2 Foundational components

This document addresses the security of the public cloud portion of hybrid cloud (the part of the infrastructure that is outside the vendor's on-premises security perimeter).

However, there is a set of security components that are required for both cloud and on-premises security. In this document, we refer to those security components as foundational. Another way of expressing that is that foundational means "things you should be doing already." Where a recommended practice has text labelled "foundational," it means that foundational components are required to meet that recommended practice or the recommended practices in general.

If you are going straight to the cloud, you will not, of course, have the foundational components and will need to implement them.

2.5 Terms and definitions

Production is used, as is the common usage of the industry, to mean either:

[5] XaaS is short for Everything-as-a-Service and sometimes Anything-as-a-Service.

- The process of the creation of TV and motion picture content including pre-production, principal photography, VFX, post-production, sound, mastering, etc.

or

- The entity, infrastructure, and people responsible for producing the content.

2.5.1 The NIST definition of cloud computing

We use these definitions of cloud computing from NIST Special Publication 800-145

2.5.1.1 Essential Characteristics:

On-demand self-service	A consumer can unilaterally provision computing capabilities, such as server time and network storage, as needed automatically without requiring human interaction with each service provider.
Broad network access	Capabilities are available over the network and accessed through standard mechanisms that promote use by heterogeneous thin or thick client platforms (e.g., mobile phones, tablets, laptops, and workstations).
Resource pooling.	The provider's computing resources are pooled to serve multiple consumers using a multi-tenant model, with different physical and virtual resources dynamically assigned and reassigned according to consumer demand. There is a sense of location independence in that the customer generally has no control or knowledge over the exact location of the provided resources but may be able to specify location at a higher level of abstraction (e.g., country, state, or datacenter). Examples of resources include storage, processing, memory, and network bandwidth.
Rapid elasticity	Capabilities can be elastically provisioned and released, in some cases automatically, to scale rapidly outward and inward commensurate with demand. To the consumer, the capabilities available for provisioning often appear to be unlimited and can be appropriated in any quantity at any time.
Measured service	Cloud systems automatically control and optimize resource use by leveraging a metering capability at some level of abstraction appropriate to the type of service (e.g., storage, processing, bandwidth, and active user accounts). Resource usage can be monitored, controlled, and reported, providing transparency for both the provider and consumer of the utilized service.

2.5.1.2 Service Models:

Infrastructure as a Service (IaaS).	The capability provided to the consumer is to provision processing, storage, networks, and other fundamental computing resources where the consumer is able to deploy and run arbitrary software, which can include operating systems and applications. The consumer does not manage or control the underlying cloud infrastructure but has control over operating systems, storage, and deployed applications; and possibly limited control of select networking components (e.g., host firewalls).
Platform as a Service (PaaS).	The capability provided to the consumer is to deploy onto the cloud infrastructure consumer-created or acquired applications created using programming languages, libraries, services, and tools supported by the provider. The consumer does not manage or control the underlying cloud infrastructure including network, servers, operating systems, or storage, but has control over the deployed applications and possibly configuration settings for the application-hosting environment.
Software as a Service (SaaS).	The capability provided to the consumer is to use the provider's applications running on a cloud infrastructure. The applications are accessible from various client devices through either a thin client interface, such as a web browser (e.g., web-based email), or a program interface. The consumer does not manage or control the underlying cloud infrastructure including network, servers, operating systems, storage, or even individual application capabilities, with the possible exception of limited user-specific application configuration settings.

2.5.1.3 Deployment Models

Private cloud.	The cloud infrastructure is provisioned for exclusive use by a single organization comprising multiple consumers (e.g., business units). It may be owned, managed, and operated by the organization, a third party, or some combination of them, and it may exist on or off premises.
Community cloud.	The cloud infrastructure is provisioned for exclusive use by a specific community of consumers from organizations that have shared concerns (e.g., mission, security requirements, policy, and compliance considerations). It may be owned, managed, and operated by one or more of the organizations in the community, a third party, or some combination of them, and it may exist on or off premises.

Public cloud.	The cloud infrastructure is provisioned for open use by the general public. It may be owned, managed, and operated by a business, academic, or government organization, or some combination of them. It exists on the premises of the cloud provider.
Hybrid cloud.	The cloud infrastructure is a composition of two or more distinct cloud infrastructures (private, community, or public) that remain unique entities, but are bound together by standardized or proprietary technology that enables data and application portability (e.g., cloud bursting for load balancing between clouds).

2.6 Cybersecurity terms and definitions

Cybersecurity is the art of protecting networks, devices, and data from unauthorized access or criminal use and the practice of ensuring confidentiality, integrity, and availability of information.[6] Cybersecurity includes application security, information security, network security and operational security.

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 any data, device, or other component (hardware or software) that supports information-related activities. This is the cybersecurity definition of the word *asset,* and that definition is chosen because we use quotations from various cybersecurity documents from outside of media production.

Media 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. As noted, we adopt this term to avoid confusion with the cybersecurity meaning of *asset*.

Content Protection is the protection of the media assets used in the creation of television and motion picture content.

Security is used in this document to mean the application of the discipline of *cybersecurity* to a particular use case such as *content protection*.

A *Security Perimeter* is a cordon around a network infrastructure designed to prevent intrusion and the acts of intruders such as content egress. It is a traditional security mechanism for on-premises infrastructure enforced using firewalls and authorized user access from the outside via a VPN.

A *Vulnerability* is a defect or weakness in a particular system, module, or component that leaves it open to being compromised due to attack, disaster, or other causes.

[6] https://us-cert.cisa.gov/ncas/tips/ST04-001

A *Threat* is any circumstance or event with the potential to cause harm to an information system in the form of destruction, disclosure, adverse modification of data, and/or denial of service.[7]

A *Threat Actor* or *Actor* is an individual or a group posing a threat[8].

An *Incident* is a security event that compromises the integrity, confidentiality, or availability of an information asset.

A *Reportable Incident* is an incident deemed to be significant enough to need to be reported outside of the entity. Reporting includes that required by laws or regulations, and those required contractual by contracts with customers.

Incident Handling is the corrective action to address an issue/incidence in violation of security practices and recommended practices.

A *Breach* is an incident that results in the confirmed disclosure—not just exposure—of data to an unauthorized party.

CVE, short for *Common Vulnerabilities and Exposures*, is a list of publicly disclosed computer security flaws. Each entry in the list is assigned a CVE ID number. See https://cve.mitre.org/.

Entropy, In cyber security, *entropy* is a measure of the randomness or diversity of a data-generating function such as a random number generated used to generate encryption keys. Data with full entropy is completely random and no meaningful patterns can be found. Low entropy data may mean that it is possible to predict other generated values.

2.7 References

2.7.1 Industry

MPA Content Security Program: Content Security Best Practices Common Guidelines. Version 4.08. https://www.motionpictures.org/what-we-do/safeguarding-creativity/additional-resources/#content-protection-best-practices

MovieLabs Common Security Architecture for Production, https://movielabs.com/download/8270/

Cloud Security Alliance's Security Guidance for Critical Areas of Focus in Cloud Computing v4.0. https://cloudsecurityalliance.org/download/securityguidance-v4/

Verizon 2021 Data Breach Investigations Report (DBIR). https://verizon.com/dbir/

Cyber Threat Modeling: Survey, Assessment, and Representative Framework, Homeland Security Systems Engineering and Development Institute. https://www.mitre.org/sites/default/files/publications/pr_18-1174-ngci-cyber-threat-modeling.pdf

[7] The term "threat' is sometimes used to mean "threat actor". Our definition is that used in several NIST publications.
[8] NIST SP 800-150

Security Leaders Handbook, HackerOne. https://www.hackerone.com

Open Source Security Testing Methodology Manual (OSSTMM), https://www.isecom.org/research.html

Penetration Testing Execution Standard (PTES), http://www.pentest-standard.org/

Information System Security Assessment Framework (ISSAF), https://sourceforge.net/projects/isstf/files/issaf%20document/

OWASP Testing Guide, https://owasp.org/www-project-web-security-testing-guide/

B.A.S.E. - A Security Assessment Methodology, SANS, https://www.sans.org/white-papers/1587/

COBIT: An ISACA Framework, https://www.isaca.org/resources/cobit

HTTP Strict Transport Security (HSTS), Internet Engineering Task Force (IETF) RFC 6797. https://datatracker.ietf.org/doc/html/rfc6797

2.7.2 Government

NIST Cybersecurity Framework https://www.nist.gov/cyberframework

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

NIST Cloud Computing Reference Architecture, Special Publication 500-292, https://tsapps.nist.gov/publication/get_pdf.cfm?pub_id=909505

NIST Definition of Cloud Computing, Special Publication 800-145. https://nvlpubs.nist.gov/nistpubs/Legacy/SP/nistspecialpublication800-145.pdf

NIST Protecting Controlled Unclassified Information in Nonfederal Systems and Organizations, Special Publication 800-171, https://csrc.nist.gov/publications/detail/sp/800-171/rev-2/final

NIST Cybersecurity Framework Profile for Ransomware Risk Management (Preliminary Draft) NISTIR 8374 https://csrc.nist.gov/publications/detail/nistir/8374/draft

NIST Computer Security Incident Handling Guide, Special Publication 800-61r2, https://csrc.nist.gov/publications/detail/sp/800-61/rev-2/final

NIST Risk Management Framework for Information Systems and Organizations: A System Life Cycle Approach for Security and Privacy Special Publication 800-37 Rev. 2 https://csrc.nist.gov/publications/detail/sp/800-37/rev-2/final

NIST Guide for Cybersecurity Event Recovery, NIST Special Publication 800-184. https://nvlpubs.nist.gov/nistpubs/specialpublications/nist.sp.800-184.pdf

NIST Recommendation for Key Management, NIST Special Publication 800-57. https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-57pt1r5.pdf

Data Integrity: Recovering from Ransomware and Other Destructive Events, NIST Special Publication 1800-11, https://csrc.nist.gov/publications/detail/sp/1800-11/final

NIST Special Publication 800-115, https://www.nist.gov/privacy-framework/nist-sp-800-115

Ransomware Guidance and Resources, Cybersecurity & Infrastructure Security Agency. https://www.cisa.gov/ransomware

Password administration for system owners, UK National Cyber Security Centre, https://www.ncsc.gov.uk/collection/passwords/updating-your-approach

Mitigating malware and ransomware attacks, UK National Cyber Security Centre https://www.ncsc.gov.uk/guidance/mitigating-malware-and-ransomware-attacks

Australian Government Information Security Manual, Guidelines for System Monitoring, June 2021, https://www.cyber.gov.au/acsc/view-all-content/guidance/event-logging-and-auditing

3 The Cybersecurity Landscape

We cannot consider the security of media production in isolation from the wider cybersecurity landscape. Our industry is not exempt because it does not handle financial or healthcare data, and like those industries threats to media production come from professional criminals as well as amateur. In this section we look briefly at the wider issue of information security.

Cybersecurity is complicated and difficult because we deal with increasingly complex systems and threat actors are ever more capable. Even the smallest enterprise cybersecurity requires:

- An end-to-end security framework[9] tuned to the organization's particular situation.
- Information security continuous monitoring (ISCM).[10]
- Proactive incident response planning and simulation.
- A qualified cybersecurity team with the authority to ensure the organization's cybersecurity plan is followed without exception.

ISCM is defined by NIST in Special Publication 800-137[11] as:

Information security continuous monitoring (ISCM) is defined as maintaining ongoing awareness of information security, vulnerabilities, and threats to support organizational risk management decisions.

This is difficult enough when everything is on-premises and under direct control of the vendor who can install a security perimeter to protect the internal network from the outside world.

Your risk assessment will inform you of the extent to which assets should be protected. Since many assets, like software and devices, can provide unexpected avenues for exploitation, the starting point is the presumption that all assets need some level and form of protection.

While the focus of this document is the security of resources used to provide services to productions, there are many occurrences of sensitive information leakage originating from non-production environments such as development, testing and even backup environments. In many cases, robust security features found in production are not extended into these other environments even though they may have significant attack surfaces.

3.1 Managing security

Clearly, the need for a cybersecurity team may present a challenge to small and, possibly, medium sized vendors and is further confounded by a shortage of talent in cybersecurity but the roles must be assigned to someone.

This document cannot define what is needed by a particular organization because there are too many possible variables.

---

[9] A security framework is a defined approach that intends to make computing free from security risks and privacy threats.
[10] Threat modeling answers questions like "Where am I most vulnerable to attack?" "What are the most relevant threats?" and "What do I need to do to safeguard against these threats?"
[11] And in other NIST publications such as SP 800-150

Here are some basic observations:

- Someone must manage security implications within the organization and that might include strategic, personnel, infrastructure, policy enforcement, emergency planning, security awareness, and other resources.
- The security team[12] and infrastructure management (IT) team should report independently up to the CEO/president because, in some circumstances, securing a system may mean taking it offline which is the antithesis of the goals of the IT team.
- Security configurations and implementation should be reviewed and tested by someone independent of the person who configured the system.[13]
- Security is a highly skilled occupation.[14] If qualified staff cannot be hired for one reason or another, the organization should consider contracting with a company that can manage security appropriately.

In addition to a security team, you need a cloud security operations center. This is a centralized location that deals with security issues throughout your cloud infrastructure. In the simplest case, it is a screen with a real time security monitoring dashboard manned whenever the cloud system is in use.

3.2 Methodology

You need a plan and, generally, the methodology to create the plan would not be much different from the plan you had if you secured on-premises infrastructure except that you will have different tools many of which are wired into the infrastructure.

To get us onto a common understanding as to what is required, we use the NIST Cybersecurity Framework assists us in formulating that plan. We have recommended practices in this section, in part, as a reminder that your plans must be updated.

[12] See Appendix B The Security Team

[13] Schneier's law: "Any person can invent a security system so clever that she or he can't think of how to break it." https://www.schneier.com/blog/archives/2011/04/schneiers_law.html

[14] "And because anyone can design a security system that she or he cannot break, evaluating the security credentials of the designer is an essential aspect of evaluating the system's security." ibid

Figure 3-1 NIST Cybersecurity Framework 1.1

The five components are:

From this framework we have a series of recommended practices.

3.2.1 Identify

- **Identify** Develop an organizational understanding to manage cybersecurity risk to systems, people, assets, data, and capabilities.

This means understanding the business context, the resources that support critical functions, and the related cybersecurity risks enables an organization to focus and prioritize its efforts, consistent with its risk management strategy and business needs.

In media creation this step might include assessing the attractiveness of each type of media asset to an attacker.

For example, and a big disclaimer here we are making a point and not suggesting that this example has any real basis, are cut editorial proxies with sync sound more attractive to an attacker than OCF (original camera files)? Things to consider in this example might be data size (the greater amount of data, the more difficult to exfiltrate it without detection) and the "street" value of the asset (people would watch an early cut with sync sound but would they watch uncut camera files without sync sound?) [15]

3.2.2 Protect

- **Protect** Develop and implement appropriate safeguards to ensure service delivery. The Protect Function supports the ability to limit or contain the impact of a potential cybersecurity event.

Security tools and best practices from the cloud providers are designed to assist you in protecting your use of their services. We will discuss this more later in this document.

[15] Those of you who have been in a color grading suite will recognize the control surface in the photograph in this article about Kim Dotcom. https://www.bbc.com/news/technology-42773038.

3.2.3 Detect

- **Detect** Develop and implement appropriate activities to identify the occurrence of a cybersecurity event.

The Detect Function enables timely discovery of cybersecurity events. The primary tool to do this is monitoring your cloud usage through system logs. Doing this properly will require analysis which we will revisit later. The last thing you want is for the movie you are working on to end up prematurely on the Internet and then to find out that your log analyzer had been trying to tell you about it for six months.

"You can't defend. You can't prevent. The only thing you can do is detect and respond."

Bruce Schneier (https://www.schneier.com/)

3.2.4 Respond

- **Respond** Develop and implement appropriate activities to take action in the event of a detected cybersecurity incident.

Your incident response plan should be proactive and architected to accommodate failure throughout the process. It is a mistake to not be prepared for one part of the plan not working.

Contacts with law enforcement and cybersecurity incident response specialists such as Mandiant (https://www.fireeye.com/mandiant.html) and CrowdStrike (https://www.crowdstrike.com/) should be established before you need them.

The list of actions you must be prepared to take include:

- Mobilize the incident response team.
- Identify the root cause of the breach and close it.
- Shut down all active intrusions.
- Determine the extent of any breach and loss of data.
- Inventory the state of your systems.
- Notify your clients and law enforcement. Both will want detailed information that you may not have yet.

Regular tabletop exercises to practice your incident response plan are being conducted including the use security consultants to conduct "war games." The lessons from these exercises should be used to update/improve the incident response plan.

3.2.5 Recover

- **Recover** Develop and implement appropriate activities to maintain plans for resilience and to restore any capabilities or services that were impaired due to a cybersecurity incident.

For example, in the event of an attack that denies access to data (such as a ransomware attack), what will you do? This should be closely linked to your business continuity plan. That plan should be updated to include recovering cloud resources. Your recovery plan must be able to recover from part of your recovery plan not working.

A security incident recovery plan is not the same as a disaster recovery plan because the nature of a disaster such as an earthquake and security breach are fundamentally different even if both result in the loss of access to data. You need both and they need separate strategies.

We cannot, by definition, predict the unknown unknowns[16] and that means that your incident response plan must allow for things happening that were outside of your risk analysis[17].

The *Building the Playbook* section of NIST SP 800-184 summarizes recommendations described earlier in the document to provide a consolidated list of items that can be included in a playbook. Recovery activities are organized in two phases:

1. The initial tactical recovery phase is achieved largely through execution of the playbook developed as part of the planning efforts for cyber event recovery. This playbook prepares the organization for the recovery actions themselves, building upon activities performed during the protection, detection, and response functions of the enterprise risk management life-cycle process. The actions can be organized into initiation, execution, and termination stages.
2. The second, more strategic, phase focuses on the continuous improvement of the organization risk management process life cycle, as driven by the recovery activities. This second phase looks at how to reduce the organization's attack surface and minimize cyber threats. Actions can be further organized into the planning/execution, metrics, and recovery improvement stages. Lessons learned in exercises and previous recoveries help to identify gaps and to inform the planning and execution of other CSF functions.

(Source: ITL bulletin for February 2017, Guide for Cybersecurity Incident Recovery)

This document is not the playbook you will need to recover from an incident. That is something you will have to assemble based on your cloud usage and workflows.

3.3 Risk Management

NIST framework classification: Identify

It is important to you to understand what risk management is. Without it you may not be addressing the issues you need to address, or you may be spending effort on issues of lesser importance. Here is the 125-word guide to risk management.

The components are:

- Threats

[16] "Reports that say that something hasn't happened are always interesting to me, because as we know, there are known knowns; there are things we know we know. We also know there are known unknowns; that is to say we know there are some things we do not know. But there are also unknown unknowns—the ones we don't know we don't know. And if one looks throughout the history of our country and other free countries, it is the latter category that tends to be the difficult ones." – Donald Rumsfeld, Pentagon news briefing, February 2002. Rumsfeld has acknowledged he was not the first to use the term.

[17] We do not mean everything that is outside of your risk analysis.

    - The actor (e.g., outsider, insider, partner, hazard) that exploits the vulnerability causing the risk to materialize; also, an indication of a potential undesirable event)
- Vulnerabilities
    - Weakness that is exploited and causes the risk to materialize
- Likelihood
    - Probability that the event will occur, resulting in the risk materializing
- Consequence/Impact
    - Consequence or effect of the risk materializing, positive or negative (accounts for asset sensitivity and criticality)

There are several well-known risk assessment methods:

- COBIT
- NIST 800-37, RMF
- SANS Base
- Risk Vector

And of course, guessing. Informed[18] guessing may be good enough provided you consider each of the four components listed above.

Whatever method is used, it requires knowledge of your cloud operations and an expert on the organization.

The important aspect is that the risk assessment can help make judicious use of the security budget by directing efforts in the best direction.

If you want to read more about how to measure and manage risk, see *Measuring and Managing Information Risk: A FAIR Approach* (https://www.fairinstitute.org/fair-book).

3.3.1 Threat actors

The question is, who are the threat actors? Verizon's 2021 Data Breach Investigations Report (DBIR) identifies the top threat actor varieties in 2,277 breaches, showing that approximately 80% of breaches are by organized crime but, leaving aside a small percentage to nation state actors, most of the rest are employees and unaffiliated actors.

Whether or not organized crime is such a significant threat in our industry, we do know that theft of pre-release content and ransomware are both sources of income for criminals.

A small vendor that is victim of a ransomware attack may find the ransom is disproportionate to the vendor's revenue, but the attacker might be looking at the bigger picture – the ransomed data is part of a very expensive production.

We cannot afford to underestimate the threat actors. They are getting much more sophisticated, and they too have an extensive set of SaaS offerings to draw upon. We are not talking about threat actors

[18] Uninformed guessing is likely to be counterproductive.

with the superhuman skills who leap over the 20' chain link fence you just erected, we are talking about threat actors who are very skilled at sweet-talking you into handing over the keys to the gate.

Risk is assessed in terms of possible damage, the likelihood of a breach and the cost of remediation. We do not assess risk in terms of the motivation of the attacker. The question "why would anyone do that?" is best left unanswered.

3.3.2 Threats

Top of the list of threats are unauthorized access to data and business disruption.

- Examples of unauthorized data access include access to media assets, to backend system data including personnel data, and to proprietary software.
- Examples of business disruption include network DDoS (distributed denial of service) attacks and ransomware which is itself a denial-of-service attack.

Well-known attacks have included both elements – exfiltration of data together with the deployment of ransomware either for the attackers' own reasons or for a "double" ransom demand – one ransom to stop the data being published, the other to unlock the data.

Some breaches are an intermediate step. For example, gaining access to an employee's personnel record gives an attacker the employee's job title, their manager and direct reports, their personal email address and mobile phone number all of which may be useful for a social engineering attack.

3.3.3 Threat modelling

Threat modeling provides a systematic approach to aid in finding and addressing security issues early in the design process. Formally, "*threat modeling is a form of risk assessment that models aspects of the attack and defense sides of a particular logical entity, such as a piece of data, an application, a host, a system, or an environment.*"[19]

The formal threat model framework may be more than a small vendor with few cloud-based workflows wishes to conduct, but some sort of threat model is advised.

A risk matrix can be helpful, and it would look something like this:

		Consequence		
		Minor	Moderate	Severe
Likelihood	Likely	*Catching a cold*		
	Medium			
	Unlikely			*A tsunami*

For example, catching a cold is likely but with minor consequences whereas a tsunami is unlikely but with severe consequences. The matrix can be drawn with more granularity.

[19] NIST Draft Special Publication 800-154, Guide to Data-Centric System Threat Modeling. https://csrc.nist.gov/publications/detail/sp/800-154/draft

Your threat model is a combination of:

1. Your risk assessment which should include a list of what needs protecting and identifying the assets, actors, entry points, components, use cases, and trust levels to be protected and what the threats are.
2. The measures taken (or planned) to mitigate those risks including any controls in place or planned.
3. Reviewing the risk matrix to determine if each threat is adequately mitigated.

Threat modeling is most effective when done at the workflow level, ensuring that all context is available for assessment.

3.4 Security of data connections

NIST framework classification: Protect

Historically, network protocols were defined using the OSI network model[20] which divided the network into 7 layers, each of which performs a different function. Networks are implemented using the four layers defined in the TCP/IP protocol set. In both models, the topmost layer is the application layer. However, when we speak of the application layer or layer 7, we mean everything above OSI layer 4.

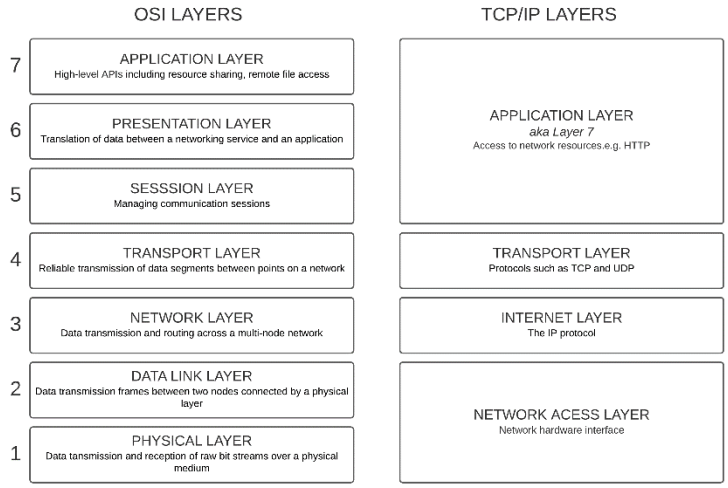

Figure 3-2 OSI and TCP/IP network models

The lower layers of the network architecture are vulnerable to unauthorized physical access to hardware, MAC spoofing[21] and man-in-the-middle attacks.[22] Layers 3 and 4 are also vulnerable to

[20] https://en.wikipedia.org/wiki/OSI_model
[21] https://en.wikipedia.org/wiki/MAC_spoofing
[22] https://en.wikipedia.org/wiki/Man-in-the-middle_attack

protocol flooding,[23] sniffing,[24] spoofing, disruption of packet routing, and so on. There are effective defenses for on-premises infrastructure against attacks at these layers some of which, such as blocking unused network protocol ports (as opposed to physical ports where Ethernet cables are connected), can be applied to cloud services. Security perimeters are focused on protecting networks at these layers.

Understanding these layers is important both when trying to secure communications in workflows and to understand how the security measures employed at each level can be effective.

A security perimeter is created using layer 2 (data link) and layer 3 (network layer) rules in firewalls and routers to create a trusted infrastructure. Security perimeters can be breached (and often are) and a lot of effort must go into detecting breaches. When cloud services are being used, layer 2 and layer 3 security has a lesser role to play and our security focus must shift further up the OSI stack particularly to layer 7, the application layer. For a virtual network, the cloud service providers offer firewall products such as Azure Firewall, AWS Network Firewall and Google Cloud Firewall which can provide filtering as well as monitoring.

Layer 7, the application layer, is the realm of attacks that exploit software vulnerabilities (especially zero-day attacks[25]) and use techniques such as cross-site scripting,[26] phishing, backdoors, and session hijacking. It is the most exposed protocol layer. It is the hardest to protect and interactions with and within applications cannot easily be characterized in a way that allow the creation of intrusion detection signatures.

Layer 7 is the obvious point of attack with cloud services because the best way or the only way to access data is through an application. Attack vectors involve the acquisition of credentials, often through phishing emails, account takeovers, and credential stuffing. The goal of the threat actor is to access a system with the credentials of an authorized user and then engage in whatever activity suits their purpose such as data exfiltration and business disruption. Methods of privilege elevation[27] can be used to go from the privileges of a regular user to those of an administrator.

It is important is to secure the application layer connection. See www.okta.com/blog/2013/01/securing-layer-7-the-closest-point-to-the-end-user/.

3.4.1 Application layer security management and the fallacy of VPNs

Earlier we discussed that one of the principles of zero-trust architectures is not trusting anything attached to a network unless it has been authenticated. The danger of VPNs is that they create an illusion of security. There is a tendency to trust devices connecting to a network over a VPN (the

[23] E.g., https://en.wikipedia.org/wiki/UDP_flood_attack
[24] https://en.wikipedia.org/wiki/Sniffing_attack
[25] https://en.wikipedia.org/wiki/Zero-day_(computing)
[26] https://en.wikipedia.org/wiki/Cross-site_scripting
[27] https://www.csoonline.com/article/3564726/privilege-escalation-explained-why-these-flaws-are-so-valuable-to-hackers.html

antithesis of zero-trust) as though they were authenticated devices attached to the local network.[28] Regardless of the vulnerabilities in any particular VPN solution, if credentials are stored by the VPN client, then it is, at best, the device that is authenticated, not the user.

3.4.2 Transport layer security

The transport layer security (TLS) protocol is designed to provide communications security over a computer network. Several versions of the protocol are widely used in applications such as email, instant messaging, and voice over IP, but its use as the security layer in HTTPS is the most visible.

The TLS protocol aims primarily to provide privacy and data integrity between two or more communicating computer applications. It runs in the application layer of the Internet and is itself composed of two layers: the TLS record and the TLS handshake protocols.

TLS is the successor to SSL. All versions of SSL have been deprecated[29] because of security flaws. TLS versions 1.0 and 1.1 were deprecated by the major browser providers in March 2020 because of security vulnerabilities. The currently acceptable versions are TLS 1.2 and TLS 1.3. These remove cryptographically weak algorithms from earlier versions and all backward compatibility with SSL.

By default, the TLS protocol which secures connections between the user and a server, validates the identity of the server to the client using X.509 certificates.[30] The authentication of the client to the server is left to the application layer.

TLS is not the only protocol used for this purpose, for example QUIC[31] which uses UDP instead of TCP. The protocol used may depend on the application.

3.4.3 Mutual authentication

The TLS protocol also offers the ability for the server to request that the client send an X.509 certificate to prove its identity. This is called mutual TLS (mTLS) as both parties are authenticated via certificates with TLS.

mTLS is commonly used for business-to-business (B2B) applications and is useful with Internet of Things (IoT) applications to authenticate devices using digital certificates. mTLS can be used to authenticate users and services before granting access to data and other services. Typically, this is done using a private certificate authority (CA).

---

[28] Trusting devices that connects to a network over a VPN as much as directly attached devices is perfectly reasonable if you are using zero-trust and don't anything on the network until it has been authenticated.

[29] The RFC deprecating SSL 3 can be found at https://datatracker.ietf.org/doc/html/rfc7568.

[30] X.509 certificates

[31] QUIC (pronounced "quick") is a general-purpose transport layer network protocol initially at Google. QUIC is used by more than half of all connections from the Chrome web browser to Google's servers. Microsoft Edge, Firefox, and Safari support it, even if not enabled by default.

3.5 Activity monitoring

NIST framework: Detect.

Logging is the collection of event data generated by devices and services.

> ***Event logging policy***
>
> *By developing an event logging policy, taking into consideration any shared responsibilities between organisations and their service providers, an organisation can improve their chances of detecting malicious behaviour on systems and networks. Such an event logging policy would cover events to be logged, logging facilities to be used, event log retention periods and how event logs will be protected.*
>
> *Australian Government Information Security Manual*
> *Guidelines for System Monitoring, June 2021*

Effective collection of event messages will likely result in a lot of messages coming throughout the cloud services you use. If the services come from multiple providers, multiple log formats may be in use. Common formats include:

- JSON (JavaScript Object Notation) Log Format
- Windows Event Log
- Microsoft IIS (Internet Information Server)
- NCSA Common Log Format

Associated with log events are levels which are used to provide more or less detail about the subject of the logging. This is application specific but, as an example, in a firewall the levels might be:

Level	**Meaning**
Alarm	An event has occurred that was flagged as an alarm event, for example, the detection of a DDoS attack
Traffic Allowed	Log of traffic that is allowed through the firewall.
Traffic Denied	Log of traffic that the firewall rejected.
Event	An event is generated when something notable occurs such as the update of a geolocation database
Debug	All possible log entries are created for a particular part of the system
Performance.	Events that report the performance of the firewall.

Interpreting logs is not as easy as it might seem. For example, a log entry of Traffic Denied only tells you that there may be an attempt to access the network and that the firewall, for that packet, did its job. However, a successful attempt to gain access to the network is going to generate log entries at the level Traffic Allowed.

A typical Common Log Format event might look like this:[32]

127.0.0.1 user-identifier john [20/Jan/2020:21:32:14 -0700] "GET /apache_pb.gif HTTP/1.0" 200 4782

The elements of this message are:

- *127.0.0.1 - refers to the IP address of the client (the remote host) that made the request to the server.*
- *user-identifier is the Ident protocol (also known as Identification Protocol, or Ident) of the client.*
- *john is the userid (user identification) of the person that is requesting the document.*
- *[20/Jan/2020:21:32:14 -0700] - is the date, time, and time zone that logs when the request was attempted. By default, it is in the strftime format of %d/%b/%Y:%H:%M:%S %z.*
- *"GET /apache_pb.gif HTTP/1.0" is the client's request line. GET refers to the method, apache_pb.gif is the resource that was requested, and HTTP/1.0 is the HTTP protocol.*
- *200 is the HTTP status code that was returned to the client after the request. 2xx is a successful response, 3xx is a redirection, 4xx is a client error, and 5xx is a server error.*
- *4782 is the size of the object - measured in bytes - that was returned to the client in question.*

You cannot be expected to read the logs even if they are formatted nicely. It is just too much data, and it is coming (or should be coming!) too fast. Furthermore, as we implied above, detecting unauthorized access to a network, for example, requires a combination of log entries that when take together, point the way. Both of these mean that your only hope is to use automated analysis that presents a dashboard view and alerts you.

Setting the combination of events that will cause an alert is too application specific to discuss here but it carries with it a delicate balance. Turn the sensitivity down and something important might be missed, turn the sensitivity up and the operator will be overwhelmed with messages.

3.6 Audit and assurance assessments

However qualified and skilled the security team members are, it is important to have your security audited by independent assessors. This may be a requirement in business agreements.

An essential component of any security measures taken for on-premises and cloud infrastructure is vulnerability scanning and penetration testing. The UK Cyber Security Center defines penetration testing as "A method for gaining assurance in the security of an IT system by attempting to breach some or all of that system's security, using the same tools and techniques as an adversary might."[33] Penetration testing may be automated or may be, in part, manual. Best case, you already know the vulnerabilities present in your system and robust testing only confirms those vulnerabilities. Of course, that's unlikely to be the outcome and testing will, at the very least, expose subtle vulnerabilities in your security.

---

[32] Example and explanation courtesy of *Graylog's 2020 Must Reads series, Log Formats – A (mostly) complete guide.* https://www.graylog.org/post/log-formats-a-complete-guide. *The web page offers descriptions for major log formats.*

[33] https://www.ncsc.gov.uk/guidance/penetration-testing

The lifetime of a testing report is limited by a change to the tested system or the disclosure of a new vulnerability in any of the software or hardware in the system. To put that another way, testing needs to be repeated. Like painting the Golden Gate Bridge, it's a job that is never finished.

These extracts from the UK Cyber Security Center advice on testing are informative.

> *Penetration testers can be used to perform a wide range of testing. The following list is illustrative, not comprehensive.*
>
> *1. Test basis*
>
> *Tests can be carried out by testers armed with varying amounts of information about your system:*
>
> - Whitebox testing – *Full information about the target is shared with the testers. This type of testing confirms the efficacy of internal vulnerability assessment and management controls by identifying the existence of known software vulnerabilities and common misconfigurations in an organisation's systems.*
> - Blackbox testing – *No information is shared with the testers about the internals of the target. This type of testing is performed from an external perspective and is aimed at identifying ways to access an organisation's internal IT assets. This more accurately models the risk faced from attackers that are unknown or unaffiliated to the target organisation. However, the lack of information can also result in vulnerabilities remaining undiscovered in the time allocated for testing.*
>
> *2. Test type*
>
> *Each of the tests described below can be run as either a blackbox or whitebox operation:*
>
> - Vulnerability identification in bespoke or niche software – *Most commonly used in web applications. This type of testing must give feedback to developers on coding practices which avoid introducing the categories of vulnerability identified.*
> - Scenario driven testing aimed at identifying vulnerabilities —– *The penetration testers explore a particular scenario to discover whether it leads to a vulnerability in your defences. Scenarios include: Lost laptop, unauthorised device connected to internal network, and compromised DMZ host, but there are many others possible. You should consider, based on previous incidents, which scenarios are most relevant to your organisation.*
> - Scenario driven testing of detection and response capability —– *In this version of scenario driven testing, the aim is to also gauge the detection and response capabilities your organisation has in place. This will help you understand their efficacy and coverage in the particular scenario.*

The pen testing report should preferably be delivered in a security industry recognized format such as the Common Vulnerability Scoring System.[34] The SANS Institute also provides a template here https://www.sans.org/blog/tips-for-creating-a-strong-cybersecurity-assessment-report/.

Your report will either give you a clean bill of health, or it will require action. The former is unlikely. Unless you are prepared to fix everything, the next step is risk assessment. The pen test report should help you determine the impact of an attacker exploiting a vulnerability would be and how likely it is to occur (the risk matrix mentioned earlier). This is input to your risk assessment and management process where controls (preventive, detective and corrective) are put in place to mitigate the risks.

Whitebox or blackbox testing should be accompanied by application security testing.

Application scanning Web Apps can be done using dynamic application security testing (DAST) and static application security testing (SAST).

- DAST tools communicate with a web application through the web front-end in order to identify potential security vulnerabilities in the web application and architectural weaknesses. It performs a black-box test.
- SAST is used to secure software by reviewing the source code of the software to identify sources of vulnerabilities.

A discussion on scanning of executable code that runs outside of a cloud environment as well as in a cloud environments is outside the scope of this document and is in the category of things you should be already doing.

[34] https://www.first.org/about/mission

4 Platform agnostic cloud tools

Aside from the tools provided by the cloud providers there are many good tools that are not tied to a specific cloud platform. Security must be designed into systems and these tools are as much about a single point of configuration and management as security.

The selection of tools in this section illustrates some of what is available. It is not an endorsement of any of these products nor is the reader to read into this section any implication about the superiority of any of them with respect to their competitors.

4.1 Terraform by HashiCorp

The notion of infrastructure as code (IaC) is the process of managing and provisioning computer data centers and cloud infrastructure through machine-readable definition files. Terraform uses a declarative[35] approach which code pushed to the servers being configured.

https://www.terraform.io/

4.2 Aqua by Aqua Security

Aqua is cloud native security platform built to secure cloud native applications.

It is advertised as securing the build, e.g., DevOps, securing the infrastructure such as Kubernetes containers, and securing the workload such as microservices.

https://www.aquasec.com/

4.3 Prisma Cloud by Palo Alto Networks

Prisma Cloud is an agent framework designed to secure hosts, containers and Kubernetes, on-demand containers and serverless functions through a single dashboard. It works across public and private clouds.

https://www.paloaltonetworks.com/prisma/cloud

4.4 Scout Suite

Scout Suite is an open-source tool for auditing cloud security enabling security posture assessment of your cloud environment. It was designed by security consultants/auditors.

https://github.com/nccgroup/ScoutSuite

NCC Scout is a cloud account monitoring platform which is part of the offerings of the NCC Group.
https://www.nccgroup.com/

[35] In computer science, declarative programming is a programming paradigm that expresses the logic of a computation without describing its control flow. It describes what a program must accomplish rather that describing how to accomplish it.

5 Shared responsibility models

Once some part of the system is not hosted on-premises or owned infrastructure, the responsibilities for the security are divided between those providing the services and those using the services. It is critical that the demarcation line is understood by the cloud service provider's customer.

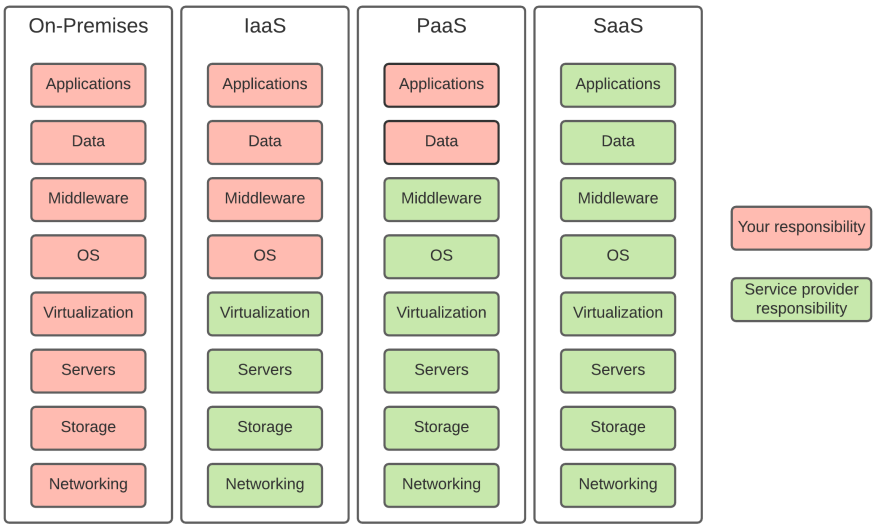

Figure 5-1 Cloud security shared responsibility

What we see from this figure is how the security responsibility is divided up for each of the service types we have identified: IaaS, PaaS, and SaaS.

Cloud providers publish more detailed shared responsibility models which delineate which part they are responsible for securing and which part their customers are responsible for securing.

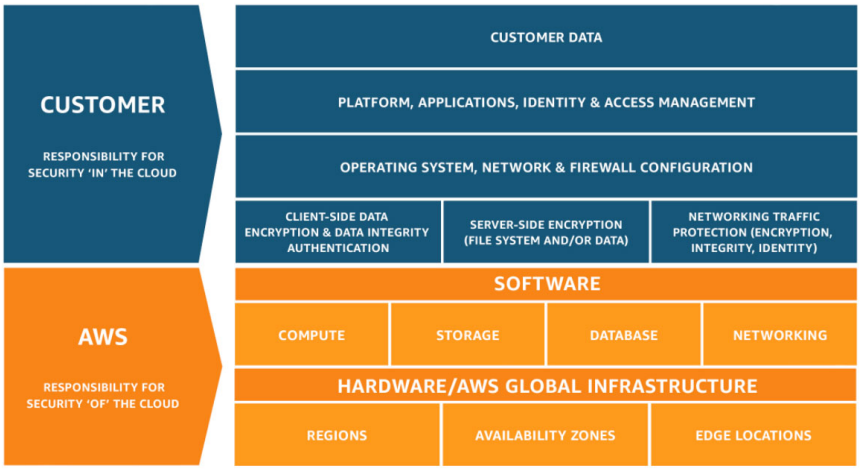

Figure 5-2 The AWS shared responsibility model (Source: AWS)

This division is perfectly understandable, but it does mean that anyone saying “our systems are secure because they are running on <insert cloud provider here>” has not grasped the shared responsibility model. Undoubtedly, if your data is stored on AWS, GCP or Azure, you don’t need to worry about someone stealing the hard drives, but as you can see from Figure 5-2, there is still much to worry about.

Viewing this as a risk matrix can help clarify what this means.

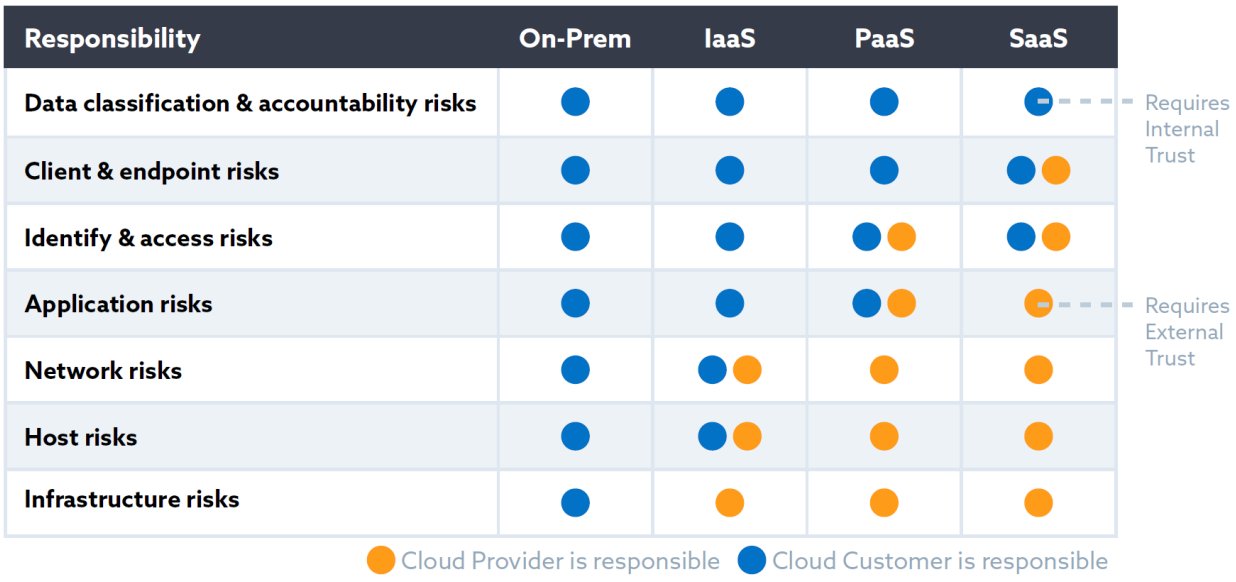

Figure 5-3 Shared responsibility risk matrix. (Source: CSA/Microsoft)

The major cloud providers offer a full suite of tools to support the customer in securing their part of this model, but it requires active engagement by the customer to configure and use those tools.

While these tools are excellent, they do present the cloud customer with a challenge: they are different from the security tools used by the cloud customer for their on-premises infrastructure, and integration can be a challenge. The danger lies in relying on both existing on-premises security tools and the cloud provider’s security tools. That means doubling the effort and increasing risk.

Security of a multi-cloud environment is discussed in Section 9 but in the context of this section, it is worth noting that the shared responsibility models for the cloud providers you are using may not be identical.

6 Global recommended practices for cloud security

In this section, we turn the discussion of section 3 into global recommended practices for cloud security. We will assume that the reader has read section 3 because the principles behind these recommended practices are explained there.

6.1 Managing cloud security

Managing cloud security requires an additional knowledge set from managing on-premises infrastructure.

Recommended Practice 1	*Foundational: Create a security team responsible for securing your system, including your cloud services, and for governance (including risk assessment and incident response).* ***[Baseline]*** *Train security personnel in cloud security practices.* ***[Baseline]*** *Foundational: A security operations center where all system monitoring is centralized. This is where[36] the security dashboards described in the document can be accessed.*

Documentation and processes such as on-boarding/off-boarding, change management, patch management, etc., that you have in place as part of your on-premises security will need to be augmented for the cloud environment.

Recommended Practice 2	*Use tools provided by your cloud provider or third-party tools such as those described in section 4 to configure security and avoid manual configuration.* ***[Baseline]***

6.2 Risk Management

You must understand what you are trying to protect, what you are trying to protect it from, how likely is a particular attack and what it will cost you if an attack is successful. If you don't do that, then you may spend a lot of money and effort that doesn't address the highest risk.

Recommended Practice 3	*Start by conducting a risk assessment and risk analysis of your cloud environment and integrate that into your approach to risk management.* *Foundational: A risk management plan for the current infrastructure that will be connected to the cloud.*

6.3 Asset Inventory

The cybersecurity definition of an *Asset* is any data, device, or other component (hardware or software) that supports information-related activities.[37]

[36] "Where" doesn't necessarily mean a physical location.

[37] We use "media asset" in this document to refer to those things that are called "assets" in media creation.

Recommended Practice 4	*Create an inventory of the cloud services you are using because without that, you do not know what you need to secure. The inventory should contain points of interaction between cloud and on-premises systems.* ***[Baseline]*** *Foundational: An asset inventory of on-premises systems (at minimum, those assets that interact with cloud assets).*

6.4 Threat modelling

Recommended Practice 5	*Update your threat model to include the threats to cloud deployment of workflows and keep up to date with subsequent changes.* *Foundational: A threat model for existing infrastructure.*

6.5 Authentication and authorization

Without effective and robust identity management it is impossible to secure cloud usage (or just about anything else). A pattern in the breaches and incidents described in DBIR is that the user is a part of a high proportion of incidents and breaches. That is not to say the user was the threat actor, but by one means or another, a user's credentials were used, or the user was unwittingly enlisted through social engineering.

Recommended Practice 6	*Everything must be authenticated: users, services, applications, and systems. The participants in an activity (human, organization, and system) must have been authorized to take part.* ***[Baseline]***

Authentication means that users, services, applications, and systems must have identities.

Recommended Practice 7	*Authorization must be given in accordance with the principle of least privilege where the minimum authorization is given to complete the task, and only for the duration the task will take.*

Single Sign On (SSO) is a powerful security tool and one that improves both security and the user experience. It is essential for cloud security because it lowers the complexity of managing authentication in an environment where allowing unauthenticated access is reckless.

Recommended Practice 8	*Use a centrally managed single sign-on (SSO) identity management system that is fully integrated with all the cloud services being used.* ***[Baseline]*** *Foundational: An identity management system. All user authentication activity, including the context of authentication attempts, is being logged and analyzed.*

SSO can use either your cloud provider's identity management system or a 3rd party identity management system integrated with the cloud provider's access controls and, conversely, a cloud provider's identity management system can be used for other resources. The SSO system can also be a federation of identity management systems. See https://en.wikipedia.org/wiki/Federated_identity.

6.6 Layer 7 Security

Recommended Practice 9	*Layer 7 security measures should be used even when the connection is over a VPN. Devices connected through VPNs should not be trusted without further authentication.* *All application layer connections should be protected with TLS 1.2 or TLS 1.3. TLS 1.1, TLS 1.0, or any version of SSL must not be used because of security flaws. Use mTLS to connect multiple entities.* ***[Baseline]*** *The validity of X.509 certificates must be verified with the appropriate trust authority before the asset can be trusted. Certificates may have expired or been revoked by the trust authority.* ***[Baseline]***

A second part of securing at layer 7 is strong session management[38]. It complements authentication and authorization.

Recommended Practice 10	*Strong session management must be used. This includes high entropy session IDs, values that do not expose anything sensitive in the tokens, the use of HSTS[39] and minimum expiration times.*

Depending on how your application is made available to users, a web application firewall (WAF) may help secure it.

A WAF or web application firewall helps protect web applications by filtering and monitoring HTTP traffic between a web application and the Internet. By deploying a WAF in front of a web application, a shield is placed between the web application and the Internet. A cloud access security broker, or CASB, acts as an intermediary between users and cloud service providers. CASBs act a gatekeeper to broker access in real time between your enterprise users and cloud resources they use, wherever users are located and regardless of the device they are using.

A CASB helps because it enables you to better understand your overall cloud posture across cloud services.

Recommended Practice 11	*Employ web application firewalls (WAF) and/or cloud access security brokers (CASB) where appropriate for Internet facing application interfaces.*

The shortcoming of WAFs and CASBs is that they cannot protect against intrusion using legitimate credentials whether obtained by a phishing attack or used inappropriately by an authorized user.

[38] OWASP outlines best practices in its Session Management Cheat sheet. https://cheatsheetseries.owasp.org/cheatsheets/Session_Management_Cheat_Sheet.html

[39] https://en.wikipedia.org/wiki/HTTP_Strict_Transport_Security

6.7 Encryption and key management

To be secure, data needs to be encrypted. There are two approaches, and they are not mutually exclusive.

1. Encrypt the container (be a physical hard drive, a cloud bucket, etc.) where the data is stored or the pipe through which the data is transferred. In both cases, the data is protected from outside attackers but still vulnerable to insiders.
2. Encrypt individual files. This means that the data is secure even if the container or pipe is not. Files can be encrypted on creation and not decrypted until they are consumed.

There are many advantages to the second approach, for example data management probably does not mean granting access to the contents of the file limiting the exposure when keys are securely managed.

However, encrypting data is of little use if the keys are not properly managed. Attackers are going to go after the key management system because trying to decrypt an AES encrypted file without the key is a futile exercise.

Recommended Practice 12	*Data should be encrypted using AES[40]-256[41] when it is stored and when it is being transmitted.* ***[Baseline]*** *Keys must be generated such that they have high entropy and must be rotated frequently.* ***[Baseline]***

[40] Advanced Encryption Standard https://en.wikipedia.org/wiki/Advanced_Encryption_Standard

[41] The number is the key length, the larger the number the more robust the encryption. 256 is generally regarded as the minimum at the time of publication of this document.

Recommended Practice 13	*Follow an established set of recommended practices such as NIST Special Publication 800-57[42] for key management and use a secure key management service such as the key management service provided by your cloud provider. Do not use manual management of keys.* ***[Baseline]*** *Never hardcode or embed keys into software or scripts.* ***[Baseline]***

6.8 Network security

Recommended Practice 14	*Use network segmentation, network security and network protection against external cyber-attacks (e.g., DDoS) for any application or service deployment in cloud.* *Use Network Security Groups to restrict the traffic flow to only required addresses and ports.* ***[Baseline]***

The network can be segmented using tools such as micro-segments on Google Cloud, VNet on Azure and VPC on AWS.

While there is basic DDoS enabled it is important to understand what level of protection is in place (i.e. basic, standard/advanced or none etc.), that it is documented and that it is monitored and metrics are generated.

6.9 Endpoint Security

Recommended Practice 15	*Incorporate endpoint security into all virtual machines running in the cloud.* *Foundational: Centrally managed endpoint security that is incorporated into all workstations, servers and virtual machines running in the on-premises environment.*

6.10 Threat Detection

Your detection system is your launching point to respond rapidly to potential incidents – unless you know about something that may be a sign of a potential incident, you cannot respond to it.

[42] https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-57pt1r5.pdf

Recommended Practice 16	*Your detection system must conduct comprehensive real-time analysis of activity using system logs that not only identifies threats but also the properties of the threat: what is it doing, whether it matters and what the response is.* ***[Baseline]*** *Use a centralized logging system to collect, analyze and report on logs collected from all the cloud services used. Analysis should be done in real time. The logging system analysis and status should be monitored and set to alert the security team.* *Use automated tools, such as those offered by a cloud provider, to inspect system logs for suspicious activity and ensure that alerts are acted upon in a timely manner.*

An example of a centralized logging solution for a cloud service is AWS CloudWatch[43].

If you have any doubt as to what you should be logging, explore the following:

- Logs your client wants collected
- Logs an assessment program wants collected
- Logs necessary to:
    - Detect a potential incident
    - Respond to an incident
    - Recover from an incident

The list below has examples of threat detection offered by the cloud providers.

- Google Cloud Security Command Center
- Amazon GuardDuty
- Azure threat detection

[43] https://docs.aws.amazon.com/solutions/latest/centralized-logging/welcome.html

6.11 Security Testing

Recommended Practice 17	*Establish policies and procedures for testing the security of the use of IaaS and PaaS systems.* ***[Baseline]*** *Foundational: A risk-based plan to mitigate any vulnerabilities discovered.* *Foundational: Frequent testing by the security team and period testing by an independent testing company specializing in security testing.* *Foundational: Use of white box testing to confirm understanding of the security stance and black box testing because that's what your adversaries will be doing.* *Foundational: Use of a standard testing methodology conducted by an accredited third-party testing company. Do not rely entirely on internally conducted tests.*[44]

Some recognized methodologies are:

- Open Source Security Testing Methodology Manual (OSSTMM), https://www.isecom.org/research.html
- Penetration Testing Execution Standard (PTES), http://www.pentest-standard.org/
- NIST Special Publication 800-115, https://www.nist.gov/privacy-framework/nist-sp-800-115
- Information System Security Assessment Framework (ISSAF), https://sourceforge.net/projects/isstf/files/issaf%20document/
- OWASP Testing Guide, https://owasp.org/www-project-web-security-testing-guide/

Recommended Practice 18	*Automated system testing must be extended to cloud infrastructure.* *Foundational: Use of automated systems testing to scan and test for security issues on on-premises infrastructure.*

The issues that can be address include:

- Virtual Machines (VMs) security
- Security of Internet facing interfaces
- Insecure system configuration
- Compiled images for container-based applications
- Unpatched and vulnerable code

6.11.1 Attack simulation

In cybersecurity a penetration test involves ethical hackers contracted for the purpose trying to break into a computer system, with no element of surprise. The blue team (defending team) is aware of the penetration test and is ready to mount a defense. A red team goes a step further, and adds physical

[44] This is consistent with the principle that anyone can create a security system they cannot penetrate.

penetration, social engineering, and an element of surprise. The blue team is given no advance warning of a red team and will treat it as a real intrusion. This is an excellent way to determine the true security posture of an organization.

This is considered to be foundational since it should already be used for on-premises infrastructure.

6.11.2 Security testing cloud infrastructure

When conducting pen testing on systems hosted on cloud provider infrastructure, it is essential to understand what you, as a customer, are permitted to test. For example, this statement from the AWS website captured in August 2021:

> ***AWS Customer Support Policy for Penetration Testing***
>
> *AWS customers are welcome to carry out security assessments or penetration tests against their AWS infrastructure without prior approval for 8 services, listed in the next section under "Permitted Services."*
>
> *Please ensure that these activities are aligned with the policy set out below. Note: Customers are not permitted to conduct any security assessments of AWS infrastructure, or the AWS services themselves. If you discover a security issue within any AWS services in the course of your security assessment, please contact AWS Security immediately.*
>
> *If AWS receives an abuse report for activities related to your security testing, we will forward it to you. When responding, please provide the root cause of the reported activity, and detail what you've done to prevent the reported issue from recurring. Learn more here.*

For the current AWS policy please go to https://aws.amazon.com/security/penetration-testing/.

Azure's Penetration Testing Rules of Engagement can be found here https://www.microsoft.com/en-us/msrc/pentest-rules-of-engagement.

For GCP see https://support.google.com/cloud/answer/6262505?hl=en#zippy=%2Cdo-i-need-to-notify-google-that-i-plan-to-do-a-penetration-test-on-my-project

6.12 Response plan

As we noted earlier, you need a comprehensive and tested incident response plan in place before you experience an incident. When you need it is not the time to find that your incident response has not been updated since 2019 right before you changed cloud provider. If you have a breach or a ransomware incident, there are going to be a significant number of things happening all of which are extremely urgent and an out-of-date or incomplete[45] plan will be a confusion multiplier.

There are many incident response frameworks available there such as NIST 800-61R2, Computer Security Incident Handling Guide. The difference with cloud computing is that these frameworks must

[45] We would include the absence of an incident response plan under the heading "incomplete."

be aligned with the roles and responsibilities of cloud providers and their customers. The cloud can make for a better and more efficient incident response if the tools are leveraged properly.

The three major cloud providers offer incident response guides for their customers.

- AWS Security Incident Response Guide
- Incident Response Overview in Azure's security operations documentation
- Google Cloud's Data incident response process

(Links are correct at the time of writing).

CSA's Cloud Incident Response Framework (CIR Framework) is a guide for cloud provider customers to prepare and manage cloud security incidents.

Your incident response plan must cover the phases in the CIR Framework.

1. Preparation
2. Detection and analysis
3. Containment, eradication, and recovery
4. Postmortem

We will not reproduce the CSA CIR Framework here and strongly recommend you read it.

Recommended Practice 19	*Prepare an incident response plan appropriate for your use of cloud services. The CSA Cloud Incident Response Framework can serve as a guide.* *The plan must include the immediate steps to identify and contain any incident or breach in your cloud infrastructure, who must be contacted, how to initiate a recovery plan.* *A cloud incident response plan should leverage cloud providers' tools, information sources, services, and capabilities.* *Your incident response plan should be proactive and architected to accommodate failure throughout the process. It is a mistake to not be prepared for one part of the plan not working.* *Your CIR plan must be reviewed and updated frequently to accommodate system changes, new and revised workflows, and the changing threat landscape* *Foundational: Contacts with law enforcement and cybersecurity incident response specialists established.* *Foundational: Regular tabletop exercises to practice incident response plan.*

We close with this observation:

> *By failing to prepare, you are prepared to fail*
> Attributed to Benjamin Franklin

6.13 Recovery

Being able to recover rapidly from a breach is critical to your business. Plan now.

Recommended Practice 20	*Create plan to recover from a security incident or breach.* ***[Baseline]*** *Your recovery plan must be reviewed and updated frequently to accommodate system changes, new and revised workflows that use cloud resources, and the changing threat landscape.* *Your recovery plan must include the scenario where none of your data stored in the cloud is accessible, and none of your configurations and unsigned software can be trusted.*

6.14 Keeping up to date on threats

There are multiple sources and feeds of threat information that are relevant to you.

Security responsible vendors, those that understand they must commit fully to securing their products and services, will publish security advisories for their products both on their website and via the CVE list (for example, Teradici's Security Advisories), and will encourage users and researchers to report vulnerabilities either directly or through groups such as HackerOne (for example, see Teradici's Report a vulnerability). HackerOne is an organization of ethical hackers supporting responsible reporting of security vulnerabilities.

Keeping security vulnerabilities confidential is only considered acceptable while no remediation is available, but unresolved vulnerabilities should not be allowed to persist. Not taking every step to remediate a vulnerability as soon possible is not a responsible action. Remediation might be in the form of a patch, configuration changes, or notifying customers so they can make an informed decision as to how to proceed.

Recommended Practice 21	*Identify and track sources of security information for the cloud services and applications you use. Subscribe to and monitor the CVE list*[46]*.* ***[Baseline]***

Be cautious when using products or services from a company that offers no security advisories or has no formal process (such as a bug bounty) for reporting security vulnerabilities by users and researchers.

[46] The CVE list contains an up-to-date list of publicly disclosed cyber security vulnerabilities. It includes both commercial products and open-source software. See https://cve.mitre.org

7 Recommended Practices for IaaS and PaaS

Our quest to secure the cloud services starts with IaaS and PaaS.

Help is at hand as best practices and tools are offered from the major cloud providers and some of their competitors.

We defer to cloud provider best practices for securing services on their infrastructure because cloud providers are continuously offering new services and better tools to configure security. Too many recommended practices abstracted from the realm of deployment are likely to be less useful than referring you to the cloud provider tools.

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. [47]

The basic recommended practice that applies to any cloud service is:

Recommended Practice 22	*Act as though there is no security by default. All security must be configured according to applicable recommended practices from the cloud provider.* *Start with access to everything blocked for everyone. Then enable access user by user based on the least privilege principle.* ***[Baseline]***

The second basic recommended practice is to do with the shared responsibility model:

Recommended Practice 23	*Understand how responsibility is divided between the cloud provider and you as their customer. Identify which parts that you are responsible for.* ***[Baseline]*** *Always follow the cloud provider's recommended practices for security.* ***[Baseline]***

7.1 Starting with an existing cloud security template

Our industry is not the only one facing the challenge of securing cloud resources. Microsoft Azure, for example, documents compliance with nearly one hundred national, regional, and industry-specific regulations for data collection and use.[48]

How does this help? The first is that, without a media production security template to work from, using a comparable standard may provide a partial roadmap. A comparable standard is one that requires at least as high a standard of security in areas that have clear equivalents in media production. There will

[47] The granularity is a policy issue for the organization. For example, it may be that the granularity is larger for developers than users.

[48] https://docs.microsoft.com/en-us/compliance/regulatory/offering-home

be other parts that do not map across and there will be requirements for media production that are not in the standard.

The second way this is helpful is that it offers a way to describe your security to an interested party (e.g., a customer).

However, there is a caveat, and an analogy may help. Your cloud provider gives you a 20-foot-high chain link fence, a padlock, some security cameras, and a check list. It's up to you to decide where you erect the fence, who has keys to the padlock and which tree you attach the security cameras to.

To use Azure as an example again, this does not mean that you can start using Azure's services and your use will be compliant with the hundred national, regional, and industry-specific regulations that Azure is compliant with. Azure's compliance with those standards means that Azure's services can be configured in a way that has been demonstrated to meet a certain regulation or standard. Cloud providers provide documentation and guidance on how to be compliant with a particular certification.

Recommended Practice 24	*If you do not have more specific templates, use cloud services in configurations that meet the comparable parts of the requirements of a certification regime with comparable use and risk factors. Use cloud provider configuration tools and guidance to set up the configuration.*

Some assembly is required.

7.2 Examples of cloud service provider tools

In this section, we look at examples of the solutions offered by cloud providers.

Recommended Practice 25	*When designing your security environment, use IaaS and PaaS provider templates, such as the security design principles in Google Cloud Architecture Framework, Azure Well-Architected Framework, AWS Well Architected Framework, where they are applicable.*

Disclaimer: much of the next three subsections was taken verbatim from cloud providers' websites.

7.2.1 AWS

The AWS Well-Architected Framework helps you understand trade-offs for decisions you make while building workloads on AWS. One of the five pillars of the AWS Well-Architected Framework is the Security Pillar which provides guidance to help you apply recommended practices, current recommendations in the design, delivery, and maintenance of secure AWS workloads.

The design principles of the security pillar apply to whichever cloud provider you use. They are (source: AWS):

- Implement a strong identity foundation: Implement the principle of least privilege and enforce separation of duties with appropriate authorization for each interaction with your AWS resources. Centralize identity management and aim to eliminate reliance on long-term static credentials.

- Enable traceability: Monitor, alert, and audit actions and changes to your environment in real time. Integrate log and metric collection with systems to automatically investigate and take action.
- Apply security at all layers: Apply a defense in depth approach with multiple security controls. Apply to all layers (for example, edge of network, VPC, load balancing, every instance and compute service, operating system, application, and code).
- Automate security recommended practices: Automated software-based security mechanisms improve your ability to securely scale more rapidly and cost-effectively. Create secure architectures, including the implementation of controls that are defined and managed as code in version-controlled templates.
- Protect data in transit and at rest: Classify your data into sensitivity levels and use mechanisms, such as encryption, tokenization, and access control where appropriate.
- Keep people away from data: Use mechanisms and tools to reduce or eliminate the need for direct access or manual processing of data. This reduces the risk of mishandling or modification and human error when handling sensitive data.
- Prepare for security events: Prepare for an incident by having incident management and investigation policy and processes that align to your organizational requirements. Run incident response simulations and use tools with automation to increase your speed for detection, investigation, and recovery.

The AWS Security Pillar resources includes the video Security Best Practices for the AWS Well-Architected Way.

7.2.2 Azure

Microsoft offers the Azure Security Center[49] from which Azure customers can configure and monitor the security of their Azure services. Within the security center, the customer can "score" their security posture and compliance with the security regulations that Azure complies with.

The process is straightforward, from the Azure Security Center the user selects regulatory compliance and is presented with the option to add a compliance regulatory standard. Selecting, for example, we selected NIST SP 800-171 R2, and the Security Policy page looks like this:

[49] https://azure.microsoft.com/en-us/services/security-center/

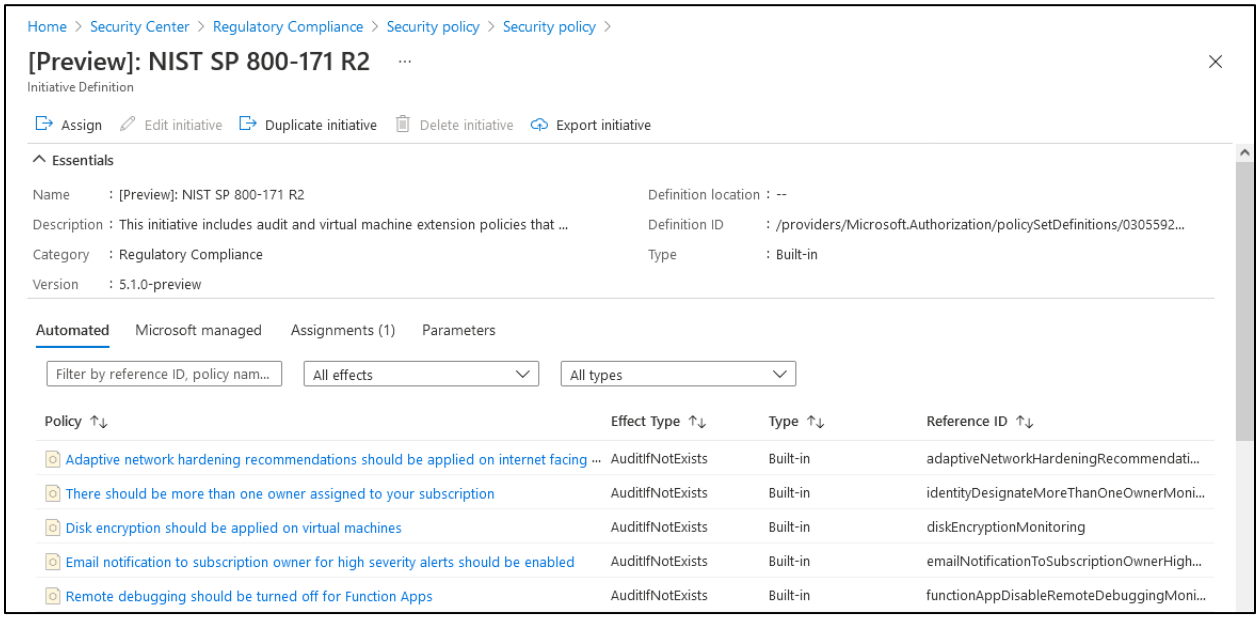

Figure 7-1 Azure Security Center compliance view

The left (blue) column is a list of security policies from NIST SP 800-171 R2. The next column across is the effect type that determines what happens when the policy rule is evaluated.

Azure Security Center enables you to determine where you are and are not in compliance with a particular standard or regulation. It is then up to you to take remedial action or decide that a control does not apply and can be safely left unresolved.

7.2.3 Google Cloud Platform

GCP's CISO Guide to Security Transformation whitepaper outlines steps for a risk-informed, rather than risk-avoidance, approach to security with the cloud. Risk-informed is an important part of any zero-trust architected security solution, but this applies whether you are at zero-trust, have started the journey or still rely on a security perimeter.

In preparing to respond to a threat you need the following information:

- What is the threat?
- What is it doing?
- Does it matter?
- How to respond?

Google Cloud Chronicle detection component is a "data fusion model that stitches events into a unified timeline, a rules engine to handle common events, and a language for describing complex threat behaviors." (Source: Google Cloud).

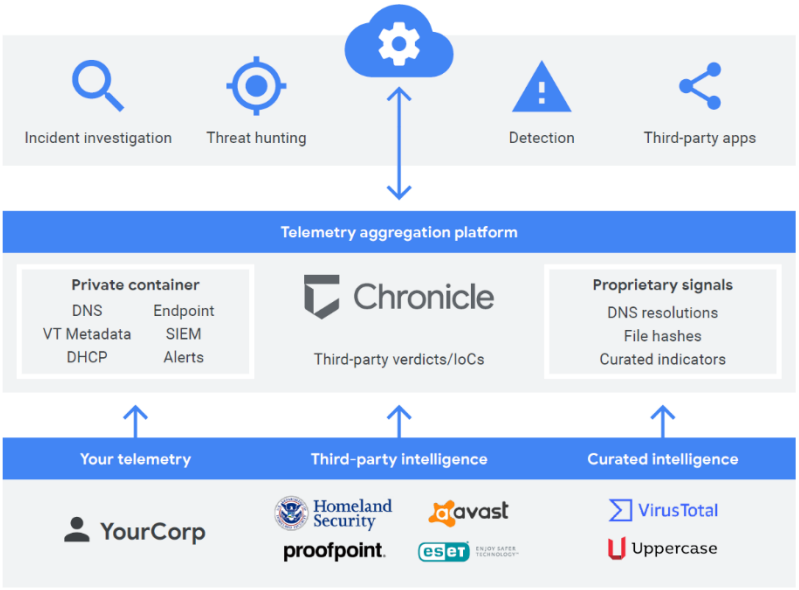

Figure 7-2 Google detection products (Source: Google Cloud)

Two critical features are its high-speed real time operation and its use of a broad spectrum of sources of information of which telemetry is only part, albeit the most critical part.

Feature	Description	Benefits
Continuous enrichment	• Automated IP to host correlation • Automated, continuous, retroactive IoC enrichment	• Faster time to investigate • Greater analyst productivity
Context and insights (threat / IoC, vulnerability, asset, user, file/process)	• Embedded threat intelligence sources (Proofpoint, DHS AIS, OSInt, Avast, ESET) • Customer-provided threat intelligence sources • Asset, vulnerability, and user context • Derived insights	• Faster time to investigate • Greater analyst productivity
Read APIs	• High performance APIs that expose functionality to downstream enterprise and MSSP SOC playbook stages and tools (ticketing, orchestration, dashboarding)	• Automation of SOC playbooks • Integration with MSSP portals • Faster time to remediation
Ingest APIs and Unified Data Model	• High throughput APIs that enable sending data directly to the Chronicle data pipeline without the need for a forwarder	• Faster time to value • Zero deployment footprint
Raw Log Scan	• Access to all unparsed fields • Search any raw security telemetry	• Faster onboarding of all security telemetry
Security & compliance	• Adherence to Google Cloud common controls • SOC 2 and SOC 3 • ISO/IEC 27001 • HIPAA BAA	• Documented, stringent controls to protect your data at every layer • Key attestations and certifications

Figure 7-3 Summary of capabilities and benefits of Google Chronicle (Source: Google Cloud)

8 Recommended Practices for SaaS

There are two classes of SaaS stakeholders that are in scope.

1. Production services vendors that use SaaS services
2. Production services vendors that provide SaaS services

8.1 Using SaaS services

The SaaS component of the shared responsibility model tells us that, as a consumer of SaaS services, your responsibility lies in ensuring that users accessing your SaaS instance/account are properly authenticated, authorized and that the method of access is secured. We refer you to earlier sections:

- Authentication and authorization are discussed in Section 6.5, 0 and 0
- Securing the method of access is discussed in Section 3.4 and in Section 6.6, 0

In addition, you should understand what certifications or standards the SaaS service has or meets. These may provide you with a basis to judge their security.

Your SaaS provider should be prepared to tell you, perhaps under NDA, where their security comes from. For example, if the SaaS provider's applications are running on a cloud provider, do they know[50] what their responsibilities are under the shared responsibility model? How does the SaaS provider address the recommended practices in Section 7 Recommended Practices for IaaS and PaaS?

Recommended Practice 26	*Require SaaS providers to undergo a recognized security assessment, for example a Cloud Controls Matrix[51] assessment or a SOC2[52] assessment.* *Follow the SaaS provider's recommended practices for the secure use of their service so that your use is within the scope of that assessment.*

It's important to realize that your SaaS provider needs to undergo a SOC2 assessment. It is not enough for them to rely on their cloud provider's SOC2 assessment since that assessment only covers the cloud provider's responsibility in the shared responsibility model.

Whether the SaaS provider is running their applications on a cloud service or on their own infrastructure, it is obvious that their infrastructure must be secure but there are so many variations in implementation as to make it impossible to present recommended practices with substance.

Your SaaS provider should publish, at the very least to its customers, security advisories and have a formal process for reporting security vulnerabilities by users and researchers preferably in the form of a bug bounty.[53] They must also notify their customers of any breach.

---

[50] Unfortunately, too many times one hears something like "it's running on [insert name of cloud service provider] so it must be secure" or "we leave our security up to [insert name of cloud service provider] because they know how to secure their cloud."
[51] CCM is published by the Cloud Security Alliance Cloud https://cloudsecurityalliance.org/artifacts/cloud-controls-matrix-v4/
[52] System Operations and Controls published by the American Institute of Certified Public Accountants. SOC2 is one of several types. https://us.aicpa.org/interestareas/frc/assuranceadvisoryservices/sorhome
[53] https://hackerone.com/bug-bounty-programs

Recommended Practice 27	*Ensure there are contractual obligations that require your SaaS provider to notify you immediately of any breach of their system that jeopardizes your data or the processing of your data and there is a defined mechanism for doing so in an expedient manner.* ***[Baseline]***

Recommended Practice 28	*Your incidence response plan must include actions to be taken if the SaaS provider reports a security incident.* *Your recovery plan must take into consideration the nature of a breach of the SaaS provider's security including the possibility that the SaaS provider cannot secure their service in sufficient time.* *Plan for the eventuality that all your data stored in the SaaS service, including any backups on the SaaS service, are lost.*

Your SaaS provider is responsible for the SaaS service up to the interfaces, but you must monitor those interfaces.

Recommended Practice 29	*Implement mechanisms to monitor all SaaS activity to detect suspicious activity. This must include monitoring designed to detect the misuse of credentials.* ***[Baseline]***

8.2 Providing SaaS services

If your SaaS service is running on a public cloud infrastructure, you should also be following the recommended practices described in Section 7 Recommended Practices for IaaS and PaaS

If your SaaS service is running entirely on a private cloud or on-premises infrastructure, securing that falls outside of the scope of this document.

If your SaaS service is running on a hybrid cloud, this document applies to the public cloud portion.

Recommended Practice 30	*Understand what security you are relying on in the cloud platform you are using. Ensure that you will be notified in the event of a security vulnerability in those components.* ***[Baseline]***

The converse of 0 is:

Recommended Practice 31	*In the event you have a security incident be prepared to act rapidly and publish security advisories to your customers at the earliest opportunity.* ***[Baseline]***

In many cases, keeping vulnerabilities secret means that they are known only by you and your attackers: that does not help you customers who need to be able to make informed decisions as to how to act. In many circumstances, it is unwise to assume that a security vulnerability is a secret. Do not be afraid of

publishing a security vulnerability if you can also publish a mitigation (e.g., a patch, a configuration change, etc.)

Note: by 'publish' we do not mean it must be on public facing areas of your website. Publishing through direct communication with your customers is acceptable. You must be transparent with your customers.

Recommended Practice 32	*Have a formal process through which users and researchers can report vulnerabilities*

A bug-bounty[54] program may encourage security researchers to identify security vulnerabilities in your products. It is far better than someone finds a vulnerability and reports it than exploits it.

[54] See https://www.hackerone.com/product/bug-bounty-platform

9 Recommended Practices for Multi-cloud

Multi-cloud in our context is the use of more than one IaaS or PaaS cloud service. To be in scope for this document, at least one of the cloud services is a public cloud. The other cloud service(s) might be public or private. This can be a complex security problem to solve. The challenge is that in our recommended practices we so far have focused on the simplest and most secure approach which is to use the security offerings of the cloud provider.

9.1 Unified security management in multi-cloud

Once you start using a multi-cloud environment, you are likely to encounter a place where separate regimes are managing security on each cloud. Coordinating security management in that situation can be, to say the least, complex but that does not change the applicability of the previous recommended practices.

Recommended Practice 33	*In a multi-cloud situation, use a unified security management solution. At minimum, a multi-cloud environment requires one single sign on (SSO) identity and access management (IAM) across all cloud components.* ***[Baseline]***

The best approach to reducing that complexity is to unify security management by using a single security management solution. To do that, the security management solution must be able to provide the required level of control in each cloud. Options include third-party providers, meaning not one the cloud providers, of security management and using the security management tools offered by one cloud provider to manage security on another.

- In January 2021, Microsoft Azure announced the general availability of a capability to onboard multi-cloud resources to Azure Security Center, including GCP and AWS. This allows you can protect your servers with Azure Defender for Servers based on Azure Arc and includes multi-cloud support in Azure Secure Score.
- GCP's BeyondCorp Enterprise can manage security for applications running on other clouds and SaaS services. See https://services.google.com/fh/files/misc/bce-protected-profiles-whitepaper.pdf.

9.2 Securing a private cloud

Securing a private data center or cloud and any on-premises infrastructure is outside of the scope of this other than were that security is declared to be foundational.

9.3 Hybrid-cloud

To recap our definition of hybrid-cloud:

The cloud infrastructure is a composition of two or more distinct cloud infrastructures (private, community, or public) that remain unique entities, but are bound together by standardized or proprietary

technology that enables data and application portability (e.g., cloud bursting for load balancing between clouds).

Given that your private cloud infrastructure has been secured, the workflows running across the public and private infrastructure should be secured using the recommended practices laid out in the previous sections.

9.4 More than one public cloud

Here we address the issue of using more than one public cloud provider and, potentially, a private cloud.

The issue described in the previous section of using two security systems will probably be compounded when more than one public cloud is used if each cloud is secured using that cloud provider's security suite. It will be very complex to set up and manage, and to develop actionable CIR and recovery plans. Complexity is the enemy of security.

We therefore refer you back to our recommendation to find and use a unified security management solution, 0.

Appendix A The Security Team

In this document we make reference to "the security team." In a large organization it means this:

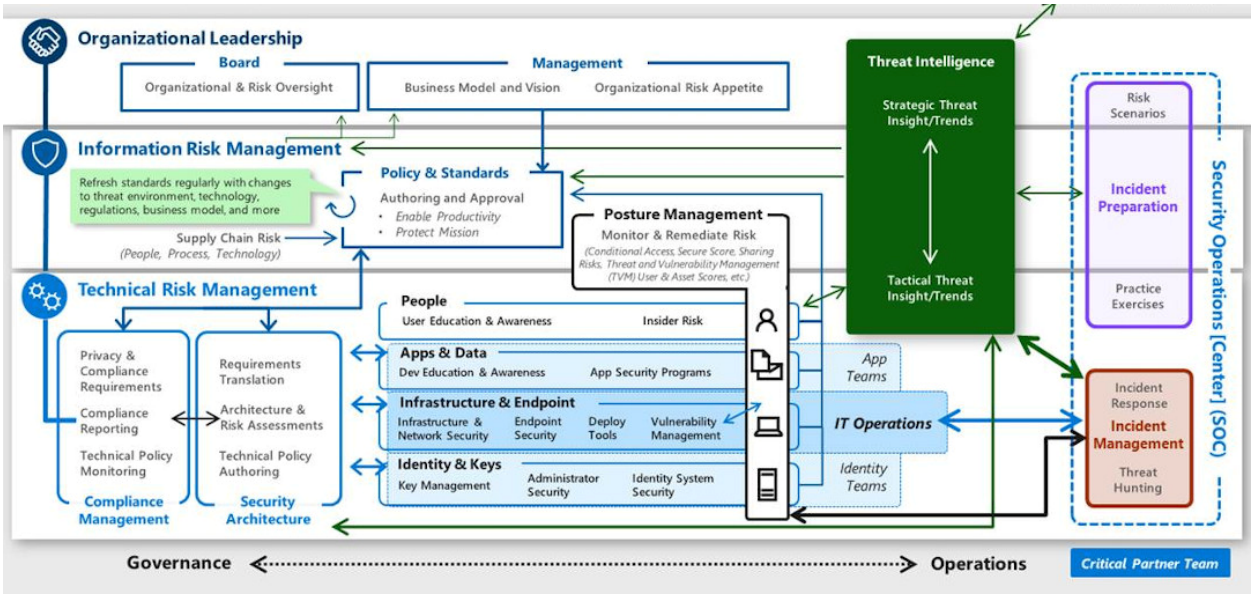

Figure 9-1 "Each function works as part of a whole security team within the organization, which is part of a larger security community defending against the same adversaries." Source Microsoft

In a modest organization, a dedicated security team is out of the question and what is important is that the roles are assigned. The roles are:

Role	Simple Description
Security architecture	You need a security architecture sufficient for your use of cloud services. In the simplest case that is a diagram and a description of the components. Consider it your roadmap.
People security	People security protects the organization from inadvertent human mistakes and malicious insider actions. The extent to which you embark on training programs is going to specific to your organization.
Application security and DevSecOps	If you develop your own applications, you need to develop them for security.
Data security	The main objective for the data security team is to provide security protections and monitoring for sensitive data For many production services providers, this is the goal.
Infrastructure and endpoint security	The complexity of securing cloud infrastructure depends on how extensively you are using cloud services.
Identity and keys	This role provides authentication and authorization of humans, services, devices, and applications.

Role	Simple Description
	This function also plays a significant role in modernizing security by establishing an identity-based perimeter that is a keystone of a zero-trust access control strategy.
Threat intelligence	Security threat intelligence provides context and actionable insights on active attacks and potential threats. This need not be as large a task as it seems, the minimum is monitoring sites where security flaws in the software and the services you use are published.
Security operations center (SOC)	A security operations center (SOC) detects, responds to, and remediates active attacks on enterprise assets. In a small organization answer these questions: who is going to detect there has been a security incident and who is in charge of doing something about it?
Security compliance management	Someone must be responsible for making sure the system has been secured.
Policy and standards	Security policies are necessary so that everyone know what is expected of them. .
Incident preparation	The primary objective for the incident preparation function is to create the “script” for responding to an incident and ensure the script is practiced and refined.
Posture management	Posture management builds on existing functions like vulnerability management and focuses on continuously monitoring and improving the security posture of the organization.
Looking forward	Risk constantly changes: new threats and new ways of working.