WHITE PAPER

Adobe Operational Security Overview

July 2024

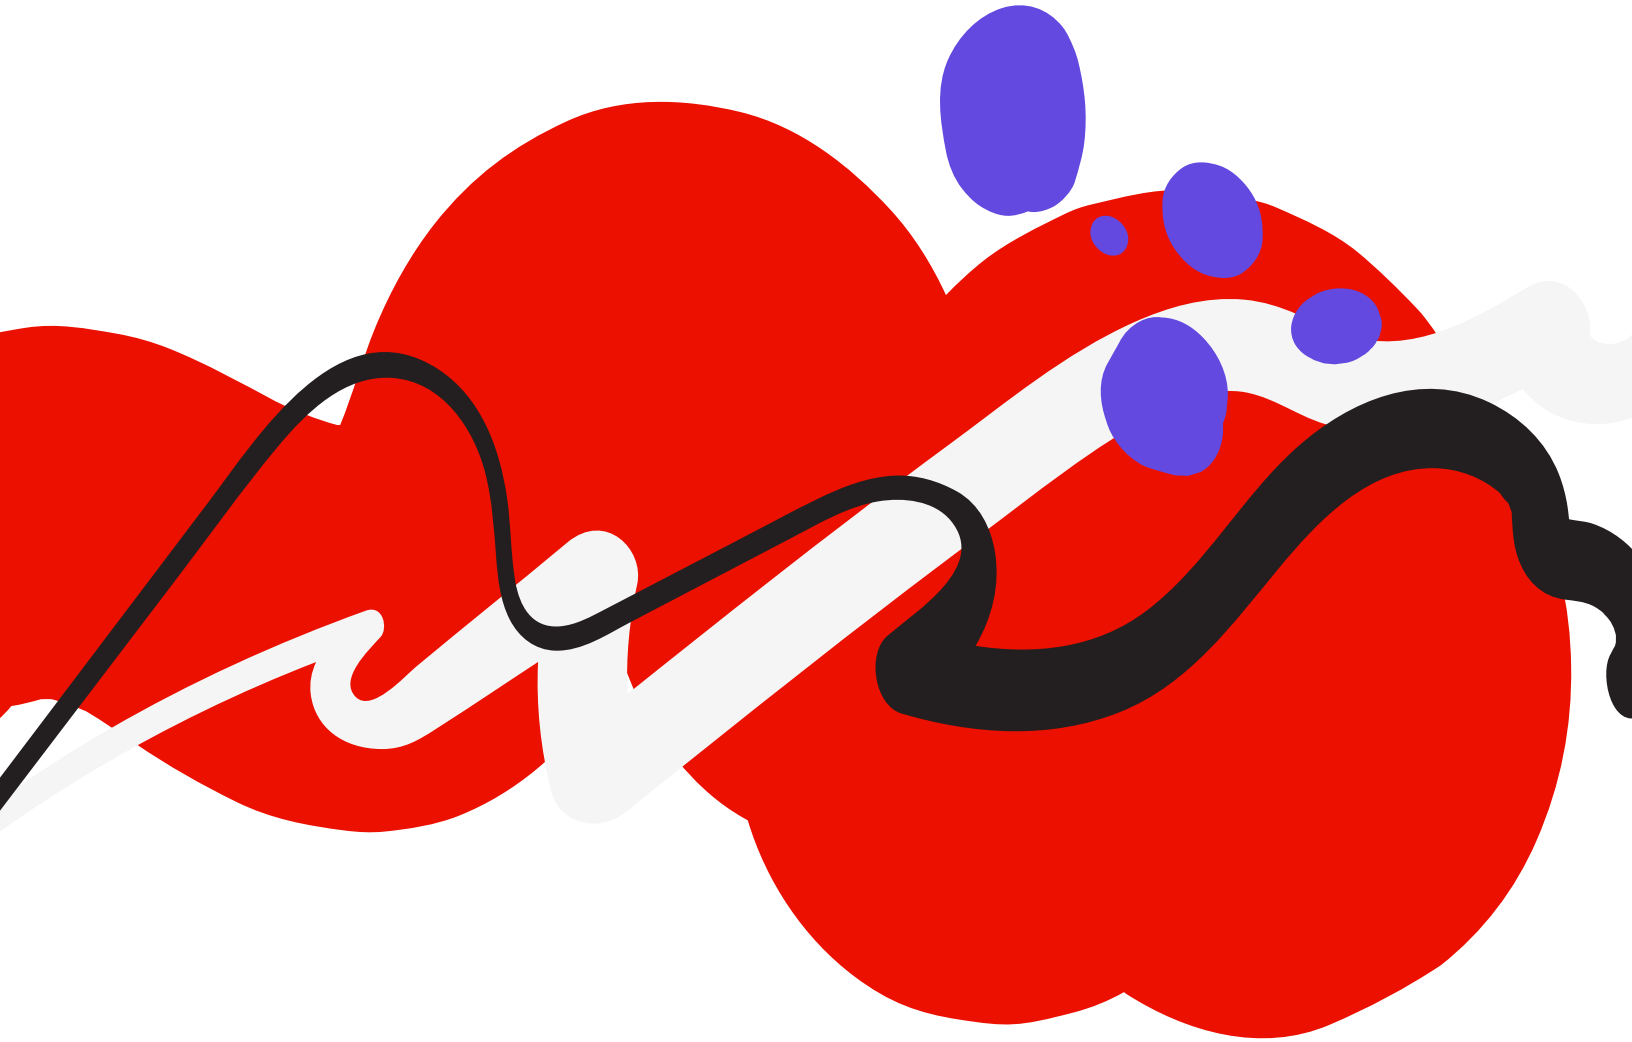

Table of Contents

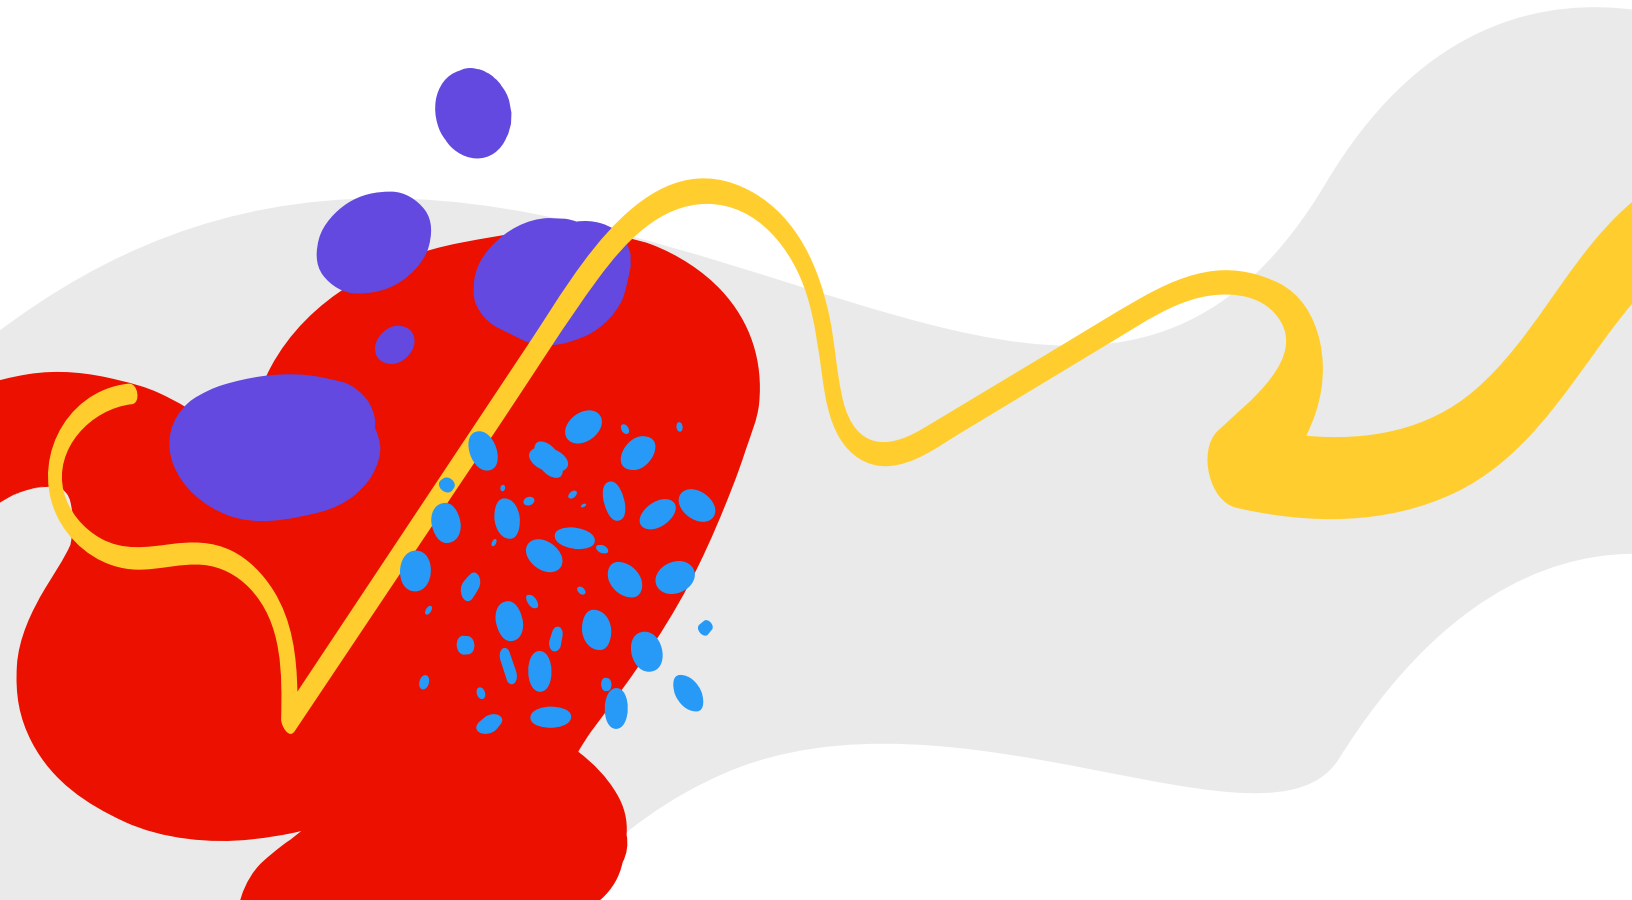

Introduction

With a cloud footprint that includes public and private clouds across different providers, the Adobe multi-cloud strategy requires consistent and repeatable guardrails that are readily available to our product and service teams. To that end, our dedicated operational security team focuses on securing cloud resources at scale and helping ensure the safety and security of customer applications and data within our continually evolving cloud infrastructure operations.

This white paper describes the Adobe secure cloud operations strategy as well as the processes and tools we have developed to help product developers and engineers improve their security posture, reduce risk to Adobe and our customers, and drive Adobe-wide adherence to compliance, privacy, and other governance frameworks.

The Adobe Secure Cloud Operations Strategy

By building security into the core of our cloud processes, Adobe proactively helps prevent potential issues that may occur within the complex security landscape. As our cloud footprint continues to grow and involve multi-cloud environments and emerging technologies — such as containers and orchestrators — standard configurations and policies as well as automation tools help us reduce human error and provide assurance to our customers that multiple layers throughout the infrastructure are protected from potential weaknesses. Scaling security through automation along with regular monitoring of our security posture and quarterly compliance reviews help Adobe detect security drift and other issues before they become critical.

To allow our developers to focus on their areas of expertise and avoid accidental security missteps, we have created standard configurations and security policies that are available to the services we deploy in the cloud. Incorporating security controls in the earliest stages of the development lifecycle helps Adobe not only harden the security posture of our services from design through deployment, but also reduce the discovery of security holes in the later stages of development when remediation is more difficult. Automated enforcement of our security controls and cloud security policies helps improve our overall corporate security posture, while also assuring customers that their security is our utmost priority.

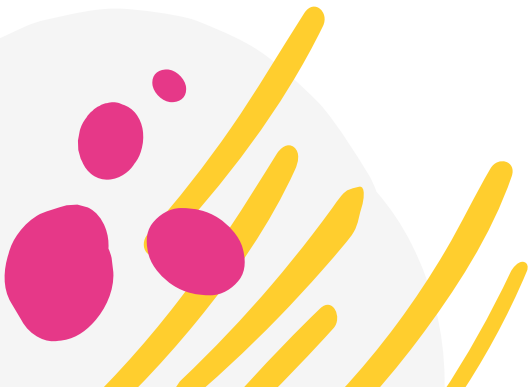

The Adobe Operational Security Stack

Developed by our dedicated operational security team, Adobe Security maintains a consolidated set of tools that helps ensure Adobe products and services are designed with security best practices in mind. Taking the company's multi-cloud security needs into account, Adobe Security's operational security tools are based on two (2) fundamental principles: standardization and prevention. To that end, the toolset includes a standardized set of continuous monitoring and workflow solutions that allows service teams to design their private and public cloud environments with security in mind — from the ground up — and helps proactively prevent security risks.

The stack of operational security tools operates on a variety of cloud resources, provides security at scale, offers standardized capabilities for the entire organization, and helps ensure security visibility into operational environments for Adobe security, audit, and compliance teams. By adopting the same set of tools and processes across our product and service teams, Adobe helps prevent security errors and enable applications to adapt to security solutions without reinventing the wheel.

With a goal of making the secure choice the default choice, Adobe Security's operational security stack includes four (4) distinct layers, each of which includes a broad range of common tools and services that can be leveraged by Adobe product teams, and which provide them with a way to stay on top of fast-changing security best practices.

Monitoring

The monitoring layer includes tools to help ingest log and configuration data from all Adobe cloud environments and regions into a central data warehouse. Once ingested, Adobe security and compliance teams, as well as the Adobe Cyber Defense Center, can analyze this data to help measure security drift and detect security gaps. Gaps may be found through manual review of the data by the security team or through automated security detection tools.

Additionally, our security teams regularly conduct scans on hosts and containers across our cloud environments, from an application as well as a network standpoint, to help detect vulnerabilities. Any vulnerabilities discovered through these scans and penetration tests are assessed, prioritized, and assigned to a remediation plan, if necessary.

The monitoring layer includes the following tools:

- **Endpoint detection and response** — CrowdStrike Falcon, a lightweight, next-generation endpoint detection and response (EDR) agent that is installed on endpoints — laptops, desktops, and servers — within Adobe, protects our data and our systems with real-time continuous monitoring and collection that enables us to identify and respond to threats quickly.

- **IaaS monitoring** — MAVLink, a public cloud data collection tool developed by Adobe, queries Amazon Web Services (AWS) and Microsoft Azure APIs for logging and environment configuration data and then ingests this information into a Splunk data warehouse. By using MAVLink, developers enable Adobe security engineering teams to view the state of the public cloud at a single point in time from a security standpoint. Adobe internal audit and compliance teams can also use this data to determine compliance with many elements of both AWS and Azure security standards.
- **Vulnerability scanning** — With a variety of commercial and in-house developed tools, we periodically scan Adobe data centers as well as our entire cloud footprint, helping pinpoint potential vulnerabilities before they arise.
- **Host scanning** — Using our EDR and vulnerability scanning tools, Adobe conducts the following three (3) activities:
  - Audit – Check host systems against policy files based on the Center for Internet Security (CIS) standard
  - Query – Collect system information to detect intrusions
  - File integrity – Track file changes in key directories
- **Syslog** — Adobe collects system logs and event messages from different machines and stores them in Splunk for monitoring and review.
- **Port scanning** — We periodically scan hundreds of thousands of Adobe IP addresses, which reduces the window of time between initial exposure and remediation. Using the nmap scanning pipeline, teams can quickly detect perimeter port exposure.
- **Container scanning** — Adobe registers and scans container images for known common vulnerabilities and exposures (CVEs) both at build and at runtime.
- **Kubernetes monitoring** — Using Faros, an internal security tool designed to monitor Kubernetes clusters, Adobe security engineers can pull a read-only configuration snapshot of any cluster at a predefined time and run custom assessors for security gaps. We then push findings to Splunk for analysis and ticketing.

Workflow

The workflow layer of the Adobe Security operational security stack helps product developers and engineers deliver the end-to-end security of Adobe products and the company's infrastructure. Enabling our teams to implement security policies efficiently, the tools available in the workflow layer make it easy to perform secure operations, including:

- **Secure host login** — Adobe maintains strict control over cloud virtual machines through enforcement of multi-factor authentication (MFA) policies and least privilege principles. We also log administrative sessions for auditing purposes.
- **Secret storage** — Adobe uses a leading third-party secure vault product to secure, store, and tightly control access to tokens, passwords, certificates, API keys, and other secrets.

- **Central cloud account provisioning** — To streamline management and governance of the Adobe cloud footprint, product teams can create and manage cloud accounts through a central service, which allows the Adobe governance team to more easily manage billing of cloud accounts as well as centrally apply security and operational policies. A single source of truth for cloud account metadata is critical to understanding the size and security posture of our cloud footprint, and centralized cloud account provisioning enables us to understand the correct ownership of accounts and their intended purpose.
- **Hardened operating system images** — By providing centralized hardened images that adhere to the Center for Internet Security (CIS) benchmarks and apply CIS-approved security updates as well as the latest security tools with desired configurations, Adobe makes a secure, out-of-the-box experience available to all our product teams. Our internal security tools scan and harden these images before they are stored in an internally developed application, called Image Factory, and are available for use by our engineering teams. In addition, product teams can use the Image Factory API to integrate the latest machine image directly into their build pipeline.
- **Secure cloud policy enforcement** — While most cloud service providers offer secure default policies, Adobe uses a complementary, internally developed tool to automate policy enforcement and remediation as well as to provide an additional layer of protection against accidental security drift and insecure services deployed in the cloud. The tool uses cloud-native services, such as Azure Policy, AWS Service Control Policies, AWS Config Rules, and Google Cloud Organization Policy, to enforce both policy and resource compliance requirements across all Adobe public cloud accounts. Any non-compliant resource in a public cloud account automatically triggers the appropriate policy action. The tool then logs the action and notifies the affected team or teams, so they can identify what triggered the remediation event.

To protect against the most common misconfigurations, our policies focus on a set of categories that include the most-often-used avenues of compromise by attackers:

- Cloud identity and privilege
- Data privacy and integrity
- Network endpoint exposure

Additionally, our cloud policy operating model dictates that new accounts include all current active policies upon provisioning. When Adobe releases a new policy, our automated enforcement process enforces existing accounts as they become compliant with the new policy. To speed up the path to enforcement after releasing a new policy, we auto-enforce accounts after a defined period (around 30 days) post-release. Non-compliant accounts are automatically ticketed for rectification, including a due date to bring the account into compliance.

Automated enforcement allows our developers to focus on the higher-level work required to become compliant with the policy, while the process ensures the enforcement of the policy itself. Regularly, the automated process checks for newly compliant accounts and automatically enforces the policies.

Infrastructure

Regularly refreshed rich metadata is a key component of the infrastructure layer, which provides the foundation for the monitoring and workflow layers. Using this metadata, the toolset can automatically assign discovered security gaps to the team that owns the offending cloud resource. Other infrastructure tools include:

- **Security Information and Event Management (SIEM)** — Using Splunk to search, monitor, visualize, and analyze the aggregated log data collected in the monitoring layer, the Adobe Cyber Defense Center can conduct deeper analysis on security-related events and incidents.
- **Bug Database** — To help provide a single source for attribution, Adobe logs bugs in Jira using its automated ticketing for accountability and tracking.
- **Identity and Access Management (IAM)** — Adobe uses Microsoft Active Directory in combination with other standard tooling to manage authentication.
- **Cloud Metadata** — Adobe tracks and audits metadata for all public cloud accounts and audits this data on a quarterly basis to help secure the accounts and ensure policy governance. The cloud metadata portal helps product and security teams onboard new cloud accounts according to prescribed workflows. In addition, teams can access the metadata in the data warehouse to filter out noise, eliminate false positives, and prioritize critical threats.
- **Container Inventory** — A rich set of metadata for the container ecosystem at Adobe that helps our product teams gain deeper insights into container orchestration. In addition, teams can use the container metadata to monitor and visualize metrics as well as to gain complete visibility into any Kubernetes environment.

Process

The process layer enables Adobe Security to improve our security posture and implement security best practices on a continuous basis. Data from the three other layers of the Adobe OSS is stored in a centralized data warehouse and ingested by Jira (for security gap mitigation and resolution), dashboards (for management visibility), and other internal partners.

We use Key Performance Indicators (KPIs) to measure how effectively the operational security stack is deployed across the company and to identify outliers. Automated Jira ticketing notifies product teams when their service deviates from a hardened security state and helps our engineering and operations teams meet several control domains in the Adobe Common Controls Framework (CCF), such as configuration management and asset management.

The Adobe Security Operational Security Stack in Action

Leveraging automation, system-level controls, and standardization, each layer of the Adobe Security operational security stack works in concert with the others to help provide security for our deployed cloud resources.

Credentials and access control tooling enable enforcement of identity and access management policies across our managed cloud services, while logging user activities provides deeper insight into potential areas for the improvement of enforcement policies. Developers then create, deploy, and manage secure-by-default cloud services using hardened OS images from our Image Factory infrastructure.

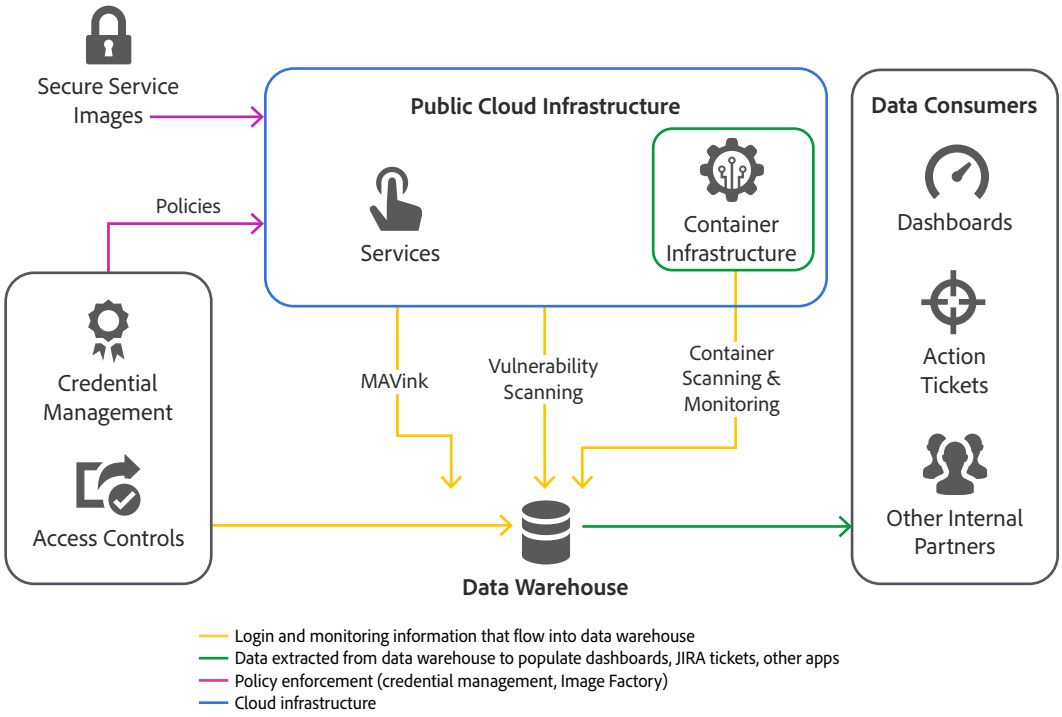

Figure 2: The Adobe OSS Data Flow

After deployment, our monitoring tools continually observe services — whether deployed in the cloud or in our container platform — and send logs and other relevant information to a central data warehouse. Dashboards, Jira tickets, and other applications throughout the company extract data from this central warehouse to populate their user-facing applications.

Conclusion

Adobe's operational security strategy and accompanying set of tools provide our product and service teams with consistent and repeatable guardrails to help ensure Adobe customer offerings are built with security in mind and adhere to our compliance, privacy, and other governance frameworks. Security automation along with continuous monitoring of our security posture through reports, dashboards, and quarterly compliance reviews help Adobe proactively prevent security risks and maintain the end-to-end security of both our products and the company's infrastructure.

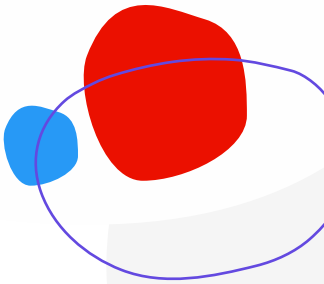

07/24