Google Cloud

Logs Anonymization

DevOps way to anonymize application logs

December, 2022

Who am I

Leonid Yankulin

Developer Relations Engineer

Cloud architect and technology evangelist focused on Application modernization, cloud migrations and observability solutions.

Engineer at Google Cloud. Prior to Google worked as cloud architect in healthcare (ChangeHealthcare, former McKesson TS) and as a software architect in interactive television (OpenTV)

at in`/minherz`, `/minherz` and `/@minherz`

Non-compliance Costs

- For the majority of organizations, growing compliance obligations are now consuming more than 40% of the security budget.
- 62% of companies surveyed indicate that automating evidence collection reduces overall compliance impact.
- 66% of companies surveyed indicated the move towards continuous monitoring are critical to their compliance program and improved security outcome.

Source: Compliance in the era of digital transformation, Coalfire report

Photo by micheile dot com on Unsplash

Product Data Categories

Business data

- Personal information (e.g. PII, PCI or PHI)
- Financial & business information
- Research and scientific private data

Product data

- Credentials & keys
- Environment info (e.g. Browser info or OS/Runtime metadata)
- Multi-tenant system customer metadata

Audit & Observability

- Logs, traces, application/debug logs
- Observability metadata (e.g. metric labels)

CI/CD Workflows

- Infrastructure & repository credentials
- CI/CD logs & internal configurations

Compliance Controls

Preventive - Intercept

Prevent misconfigurations of resources that violate security, compliance and governance rules from being deployed on the platform. Trust that all resources have the right controls in place with a quick feedback loop integrated in the CI/CD pipeline.

Detective - Audit & Alert

Continuously monitor for violations and be alerted, based on **the same policies** used for preventive controls.

Corrective - Remediate

Quickly remediate any non-compliant resou

Eventually rollback non-compliant changes to the desired and compliant state.

Compliance Enforcement

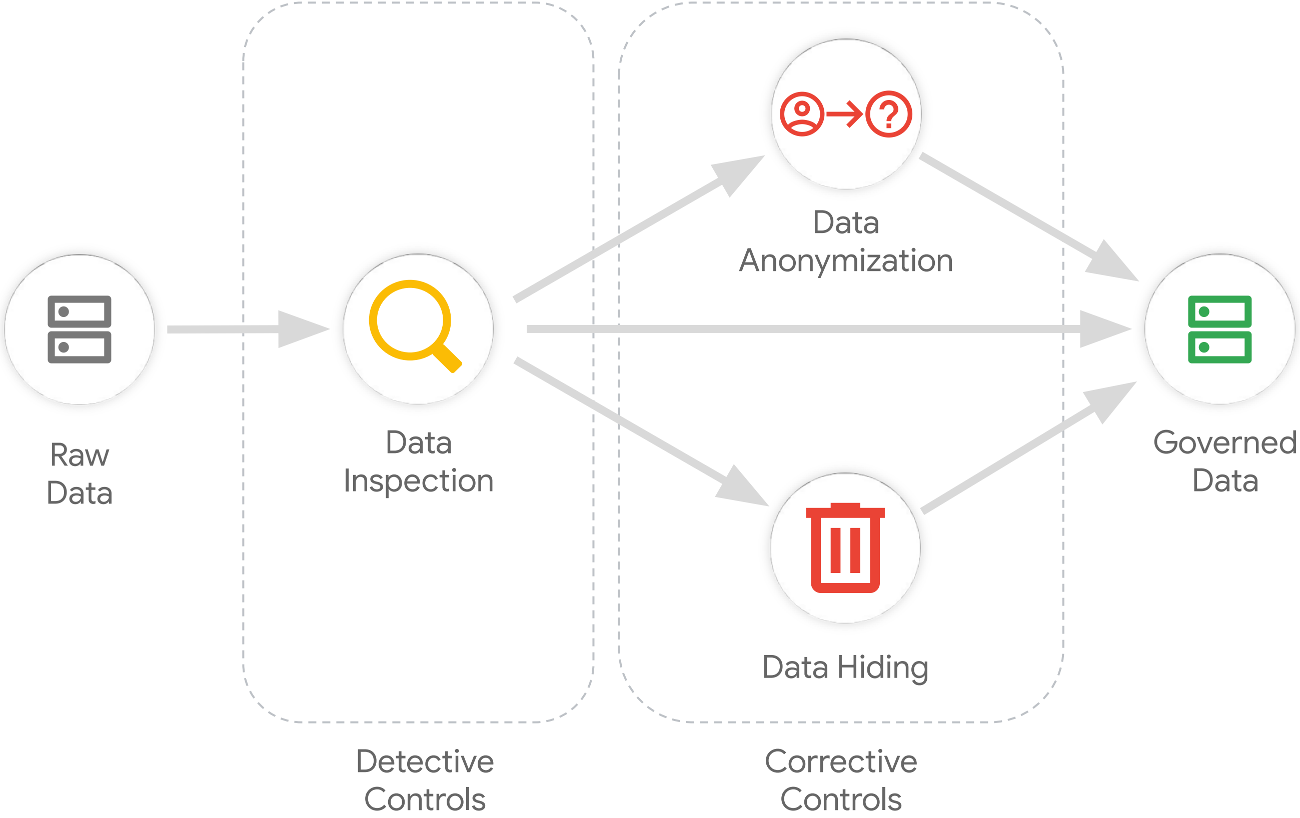

General Architecture: Logs ingestion flow

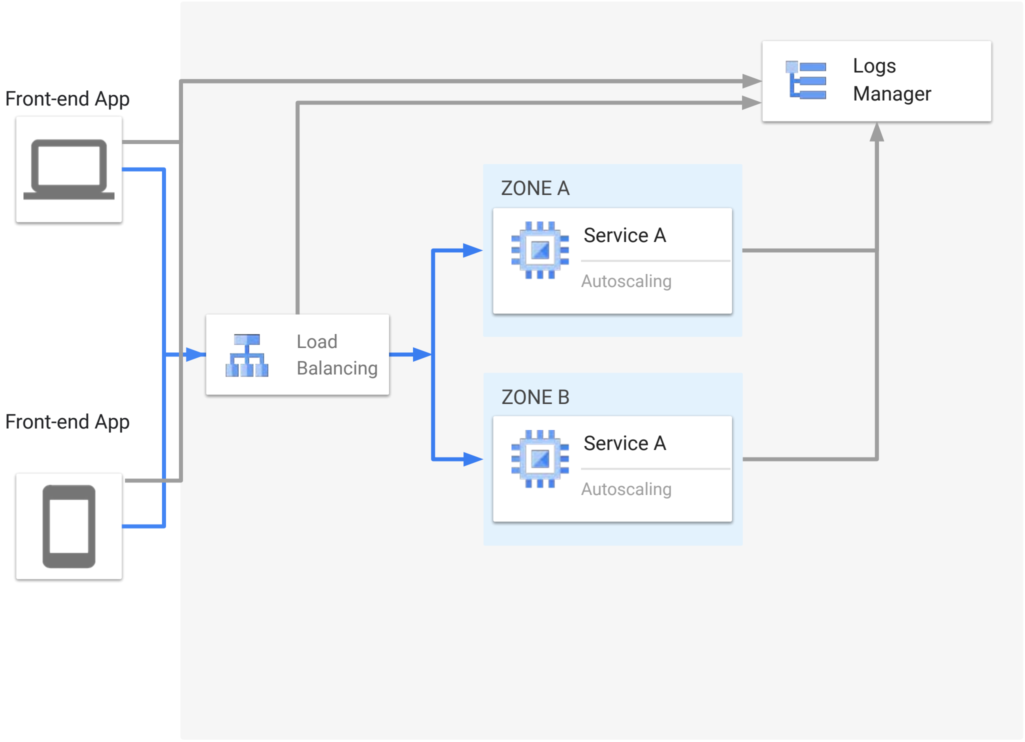

General Architecture: Logs compliance at generation time

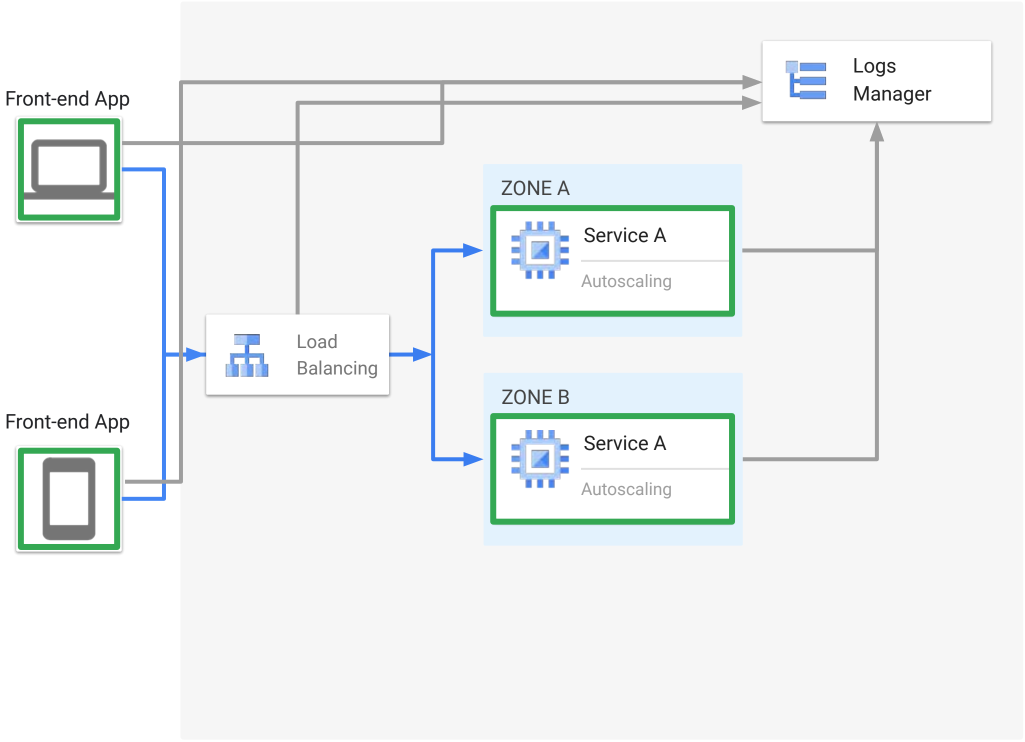

Danielle
Developer

"I implement and maintain compliance controls for logs ingestion according to company's data governing policies"

General Architecture: Logs ingestion flow using DLP service

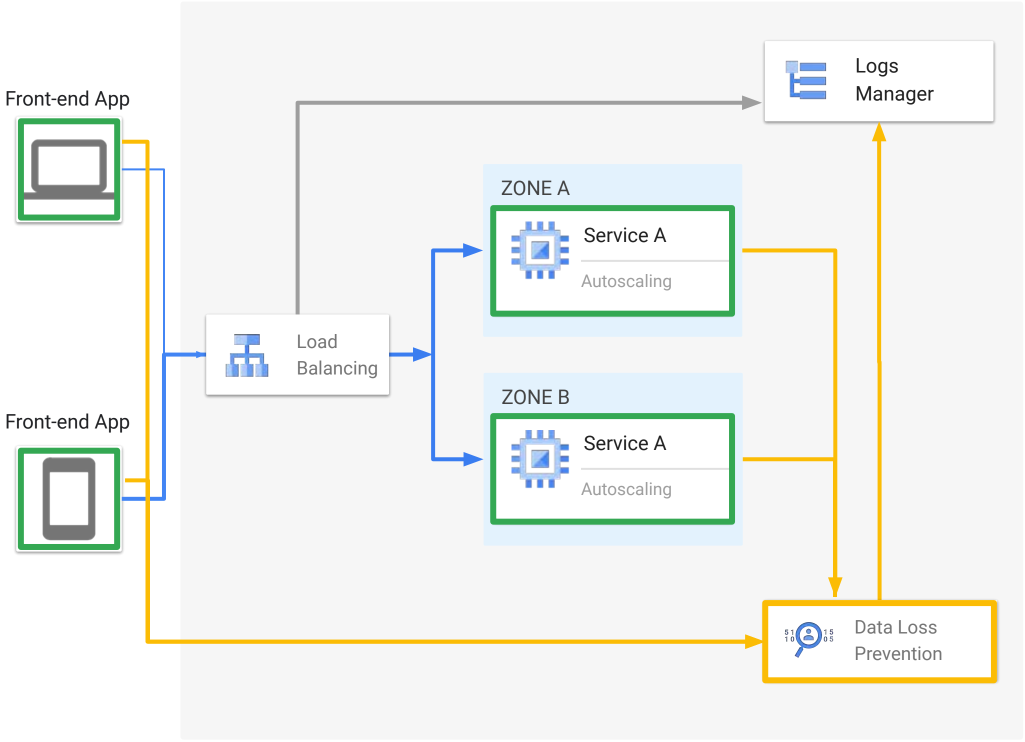

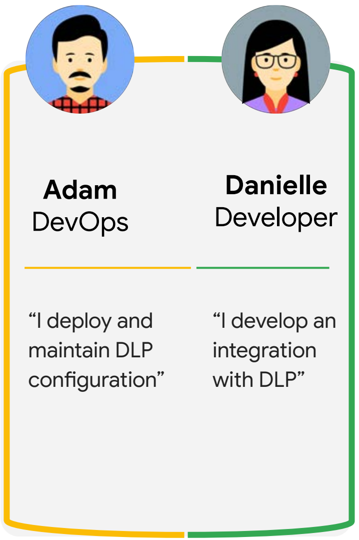

Adam
DevOps

“I deploy and maintain DLP configuration”

Danielle
Developer

“I develop an integration with DLP”

General Architecture: Logs ingestion flow using **managed** DLP service

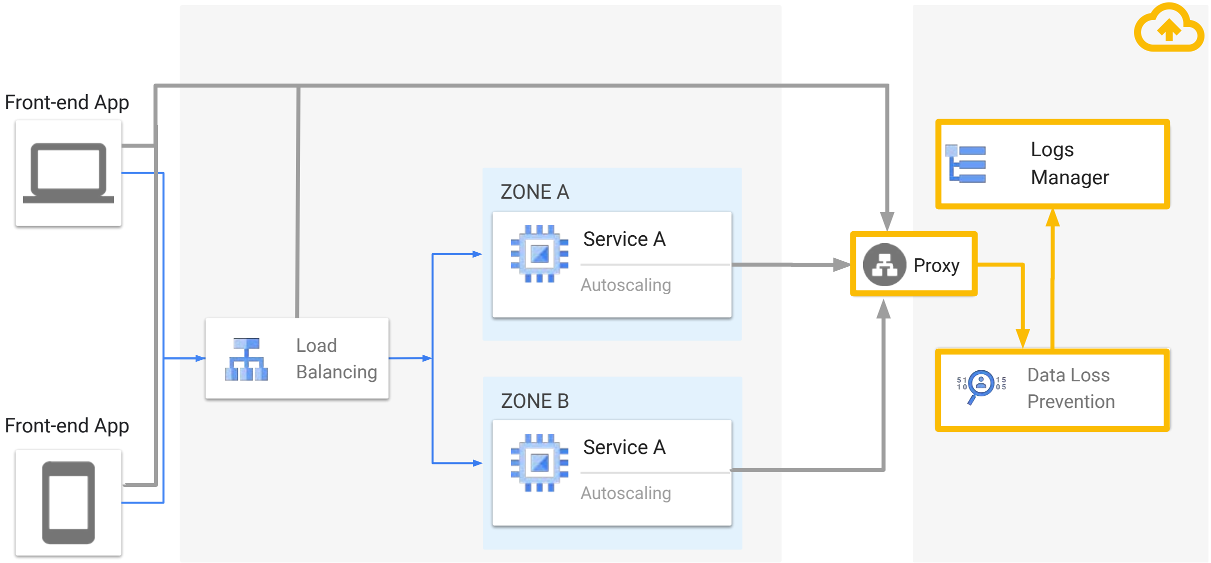

Adam
DevOps

"I deploy and maintain DLP configuration"

Google Cloud hosted PoC

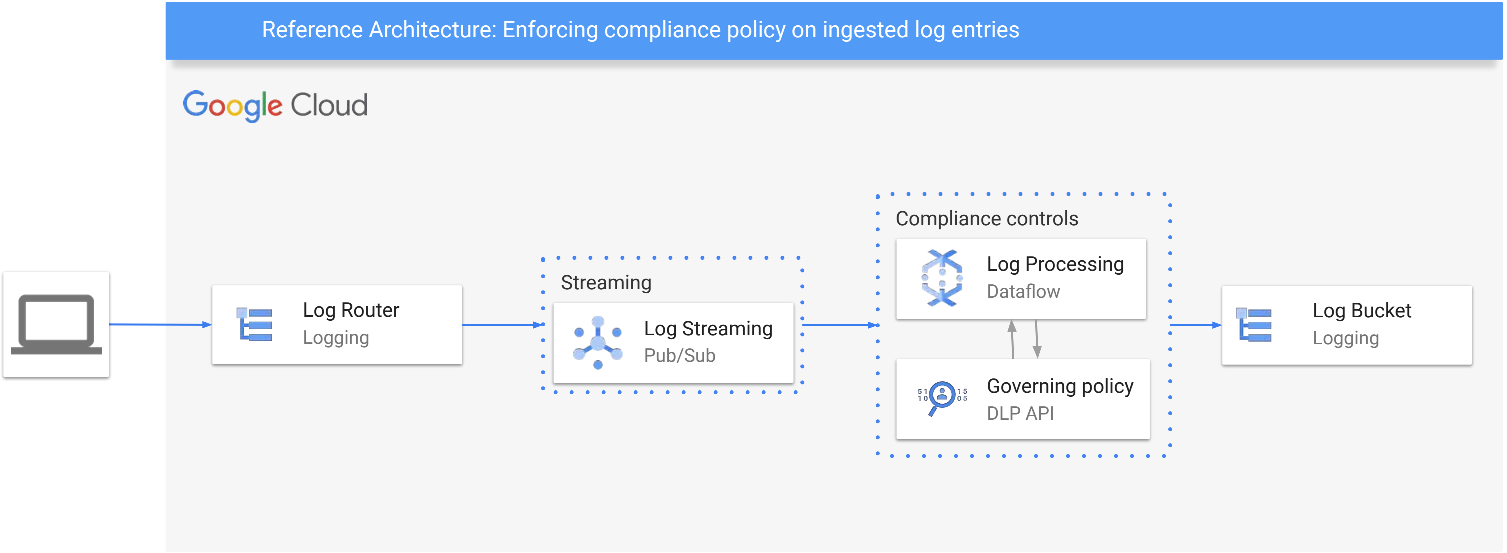

Reference Architecture: Enforcing compliance policy on ingested log entries

ETL Pipeline (in Apache Beam)

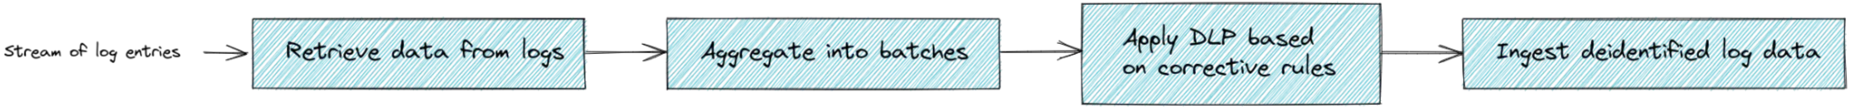

- → Code of the demo on Github
- → Blog post about the solution on Medium
- → DLP documentation

Thank you.

Find `minherz` on Discord's Conf42/devsecops channel: