

Towards Secure and Leak-Free Workflows Using Microservice Isolation

Loïc Miller, Pascal Mérindol, Antoine Gallais and Cristel Pelsser
May 25, 2021

University of Strasbourg, France

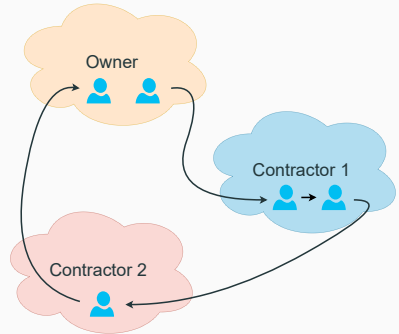


Preventing workflow data exposures with microservices

- There are more and more **data leaks and breaches**.
- They result in important **losses** for businesses.
- Yahoo (2013): 3 billion account details leaked.
- Unencrypted data accessed by an unauthorized third party.
- MikroTik routers hijacked (2018).
- Eavesdropping on > 7,500 routers.

Workflows

- We define a workflow as a *sequence of tasks* processed by a set of actors.
- The instigator of the workflow, the *owner* of the data, interacts with *contractors* to realize a task.
- Actors have *agents*: either an employee or a fully automated service.




How can we enforce a given workflow, which guarantees *data security at rest and in transport*, and *prevents data leaks*?

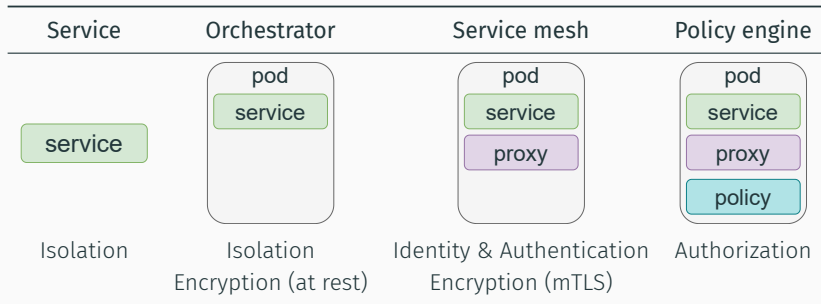
Desired properties

- Data security at rest: stored **encrypted**, access restricted by **isolation** and **policy**.
- Data security in transport: exchanged **encrypted**, with integrity and **authentication** checks.

The data cannot be **leaked** in both cases.

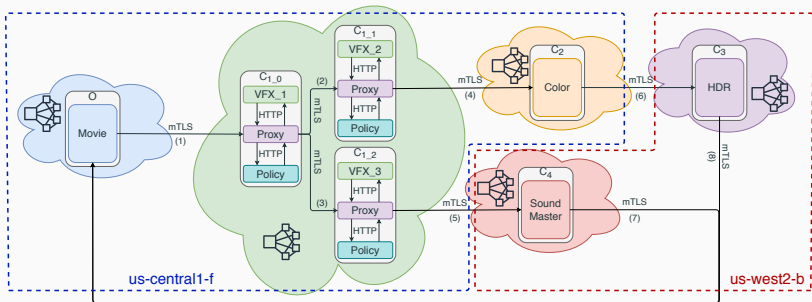


Building block security properties



Encrypted storage, encrypted communications, policy enforcement.

Proof of Concept deployed on Google Cloud Platform



- One Kubernetes cluster per actor (5 in total).
- One n1-standard-v2 per cluster (2 vCPUs, 7.5 GB of memory), except the owner which has two.

How do we estimate the security tradeoff:
Measure two metrics, pod startup time and
request duration.

Effect of OPA on pod startup time

- Independent-samples t-test
- Two deployments: one with OPA and one without.
- 130 observations per pod ($N = 1820$).

Time increased by **2 seconds on average (32.72%)**.


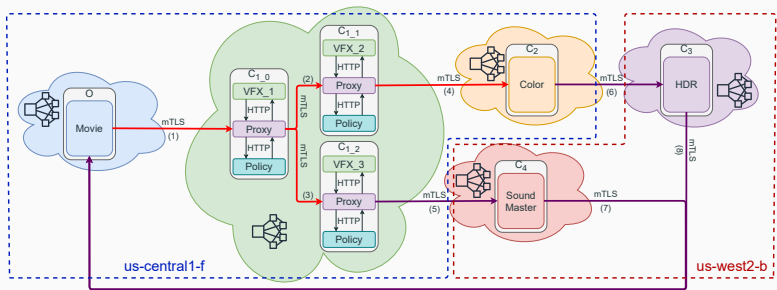


Figure 1: Startup time distribution

- $t(1818) = 43.19, p < 0.001$
- High effect size: $d = 1.985$
- High statistical power: $1 - \beta = 0.999$

Effect of policy size on request duration



We analyze **intra-region** and **inter-region** communications.

One-way between subjects ANOVA.

40 observations per communication per scenario ($N = 1600$).

Policy scenarios: **no opa**, **all allow**, **minimal**, **+100** (+147%), **+1000** (+1470%).


High (low) impact on intra (inter) region request time

Intra-region

- $F(4, 795) = 364.05$,
 $p < 0.001$
- **High** effect size:
 $\eta_p^2 = 0.65$

Inter-region

- $F(4, 795) = 15.23$,
 $p < 0.001$
- **Low** effect size:
 $\eta_p^2 = 0.07$



- Significant difference in request duration between the five scenarios for both types.

Conclusion

- Flexible infrastructure to secure communications in a workflow.
- Proof of concept¹.

Performance analysis

- Startup time using OPA increased by 2 seconds (32.72%).
- Request duration is an important factor in intra-region communications.

¹Code, data and guidance at <https://github.com/loicmiller/secure-workflow>

Thank you!