Secure, Low-Friction Data Access with SPIFFE

Bloomberg

STAC Global Live
May 18, 2021

Phil Vachon
Security Architect, Office of the CTO

TechAtBloomberg.com

Bloomberg by the Numbers

20,000+ employees, worldwide

6,500+ engineers

1,700+ journalists and analysts

2,300+ data specialists

200+ employees working on AI and ML applications

230 billion ticks per day, from every asset class & market you can imagine

96,000 companies, **3.3M** bios of executives, leaders and govt officials

2 million News Articles ingested per day, from **150,000+** vetted sources

100+ Alternative Data sources across multiple sectors

… all to bring the highest quality insights and analysis to decision makers around the world!

Identity for Data Access Needs to Be...

Secure - some data needs to be protected at all costs; all data's integrity needs to be preserved

Usable - should be understandable and usable by mere mortals

Scalable - need to deal with new infrastructure paradigms without falling back on old bad habits

Timely - no complicated rituals to get a new credential to enable access to a data set, once authorized

Adaptable - data is stored in a variety of locations, from public cloud through to relational databases in our data centers

Managing credentials often becomes a risk...

Kerberos Keytabs are tricky to distribute everywhere, and the further you share them, the harder it becomes to control who is accessing a resource...

Shared credentials are a major risk to a firm’s security posture, can be easily stolen and abused...

Client certificates tend to expire in the middle of the night, after everyone forgets how to generate a new one...

Modern Infrastructure is Heterogeneous

Azure

amazon
web services™

openstack®

TRUSTED®
COMPUTING
GROUP

… and has different levels of trustworthiness.

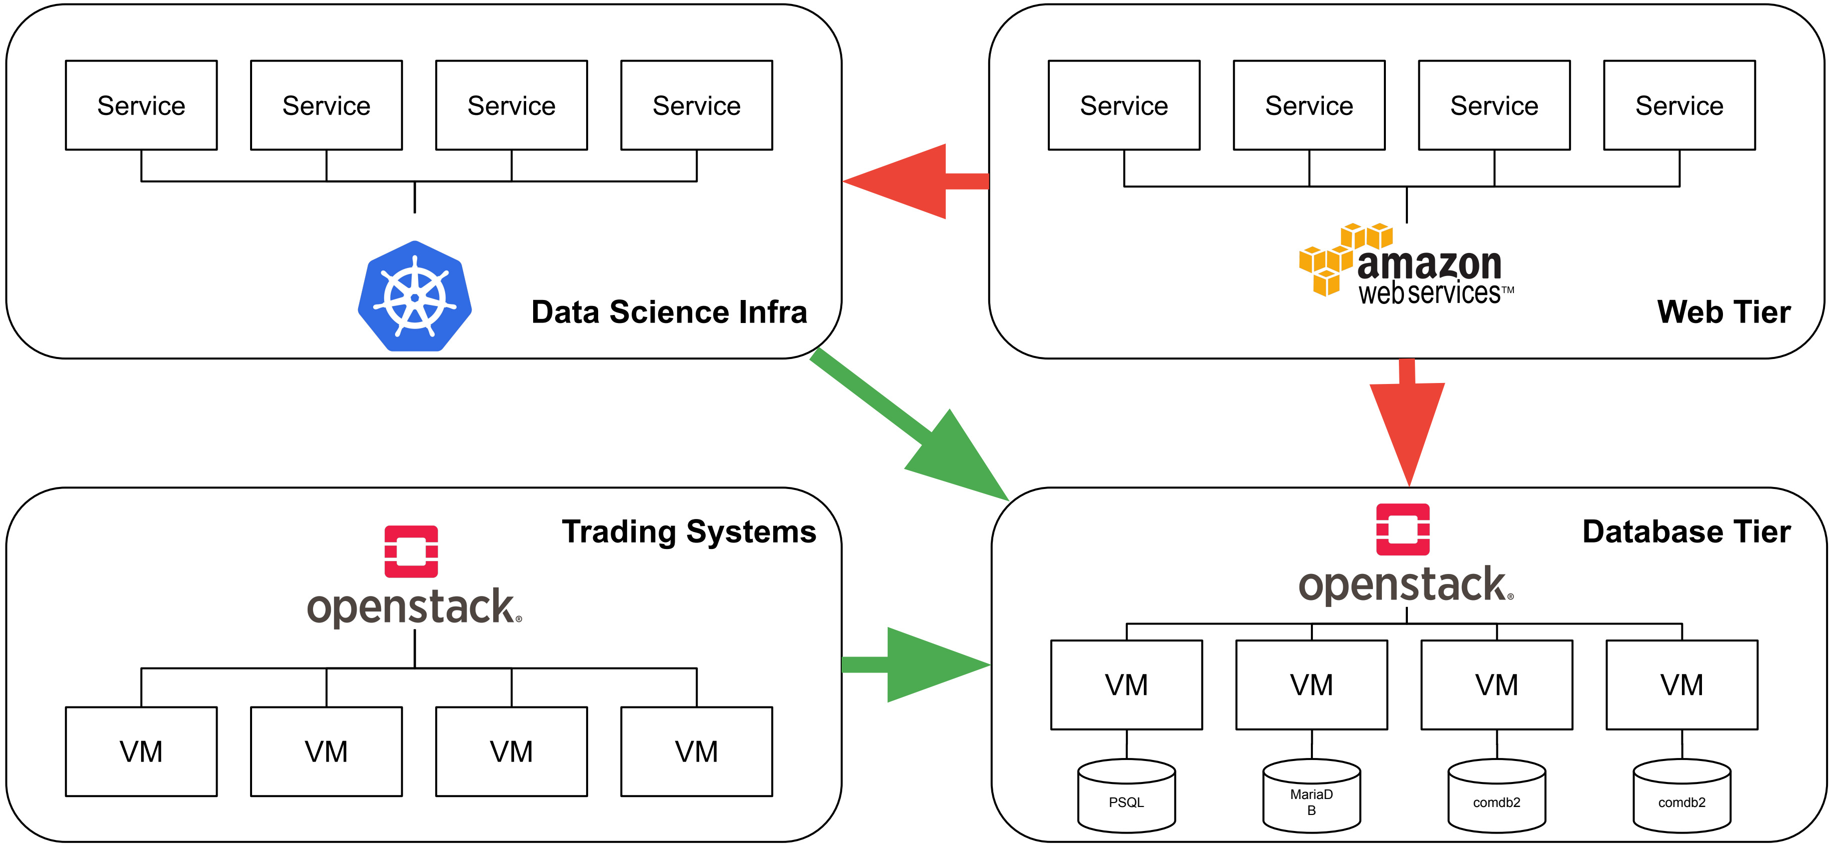

spiffe

Secure **P**roduction **I**dentity **F**ramework
for **E**veryone

What is SPIFFE?

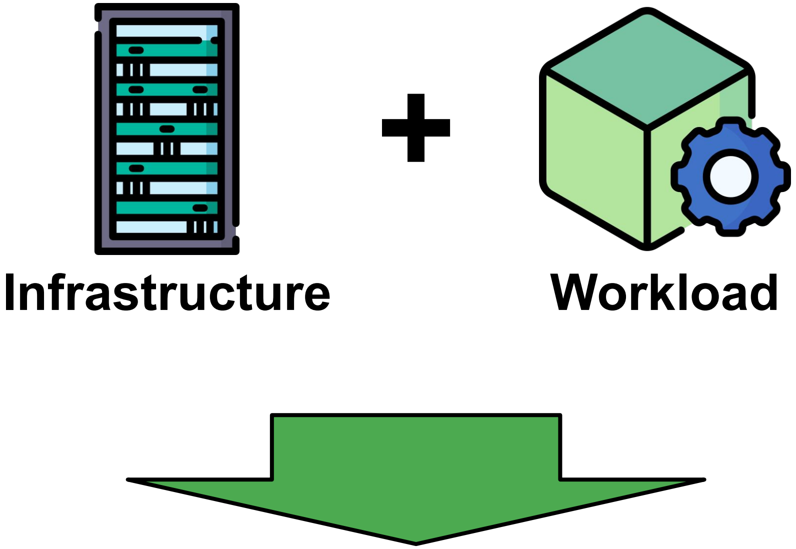

Infrastructure + Workload

spiffe://prod.paas.yoyodyne.org/security/sample_workload

Carrying a SPIFFE Identity

SVID: SPIFFE Verifiable Identity Document

X.509 Certificate

or

JSON Web Token (JWT)

Getting a SPIFFE Identity

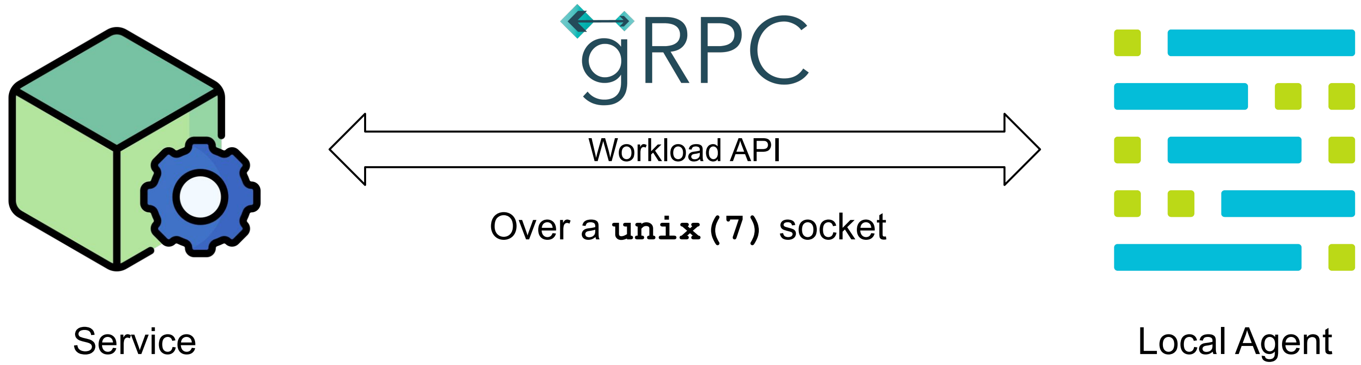

gRPC

Workload API

Over a `unix(7)` socket

Service

Local Agent

Using SVIDs

Authenticate callers to a service, using JWTs in a TLS connection

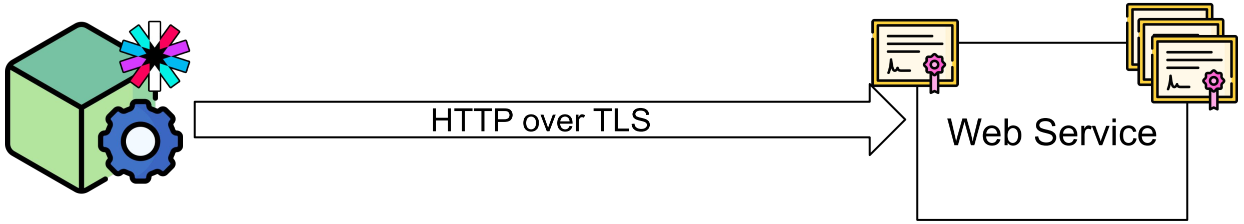

Caller and service mutually authenticate using TLS, with an X.509 SVID

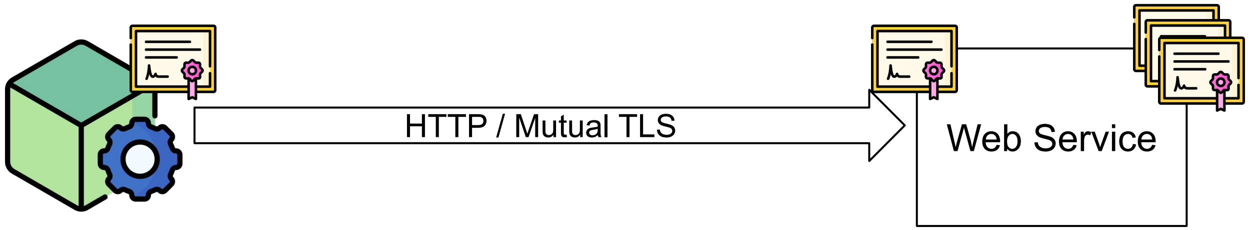

Caller authenticates to database with X.509 SVID, using database's native support for X.509 identities

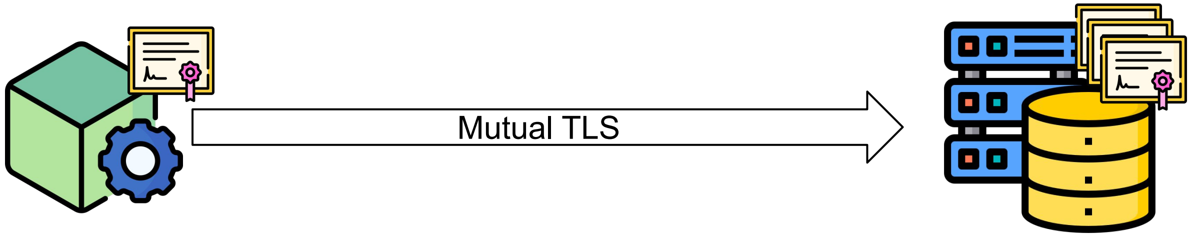

Authenticate to HashiCorp Vault using X.509 SVID

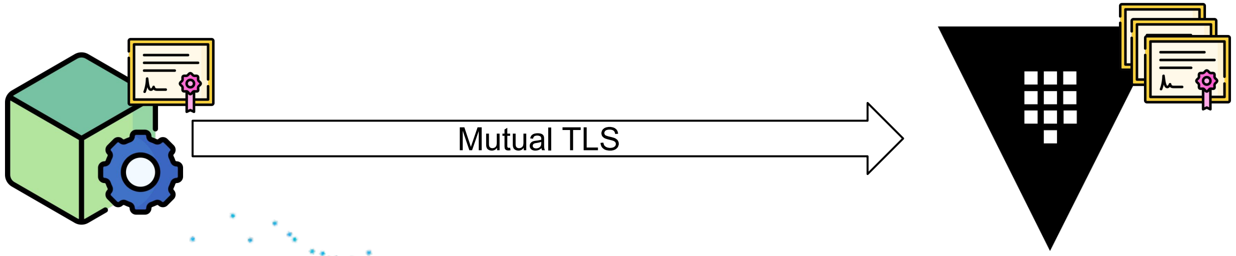

SPIFFE is not...

- …an alternative to WebPKI

- …a means to issue WebPKI certificates (have a look at ACME)

- …an authorization or access control management technology (have a look at Open Policy Agent - SPIFFE identities and OPA fit together nicely)

SPIRE

SPIFFE Runtime Environment

(an implementation of SPIFFE)

SPIRE Deployment

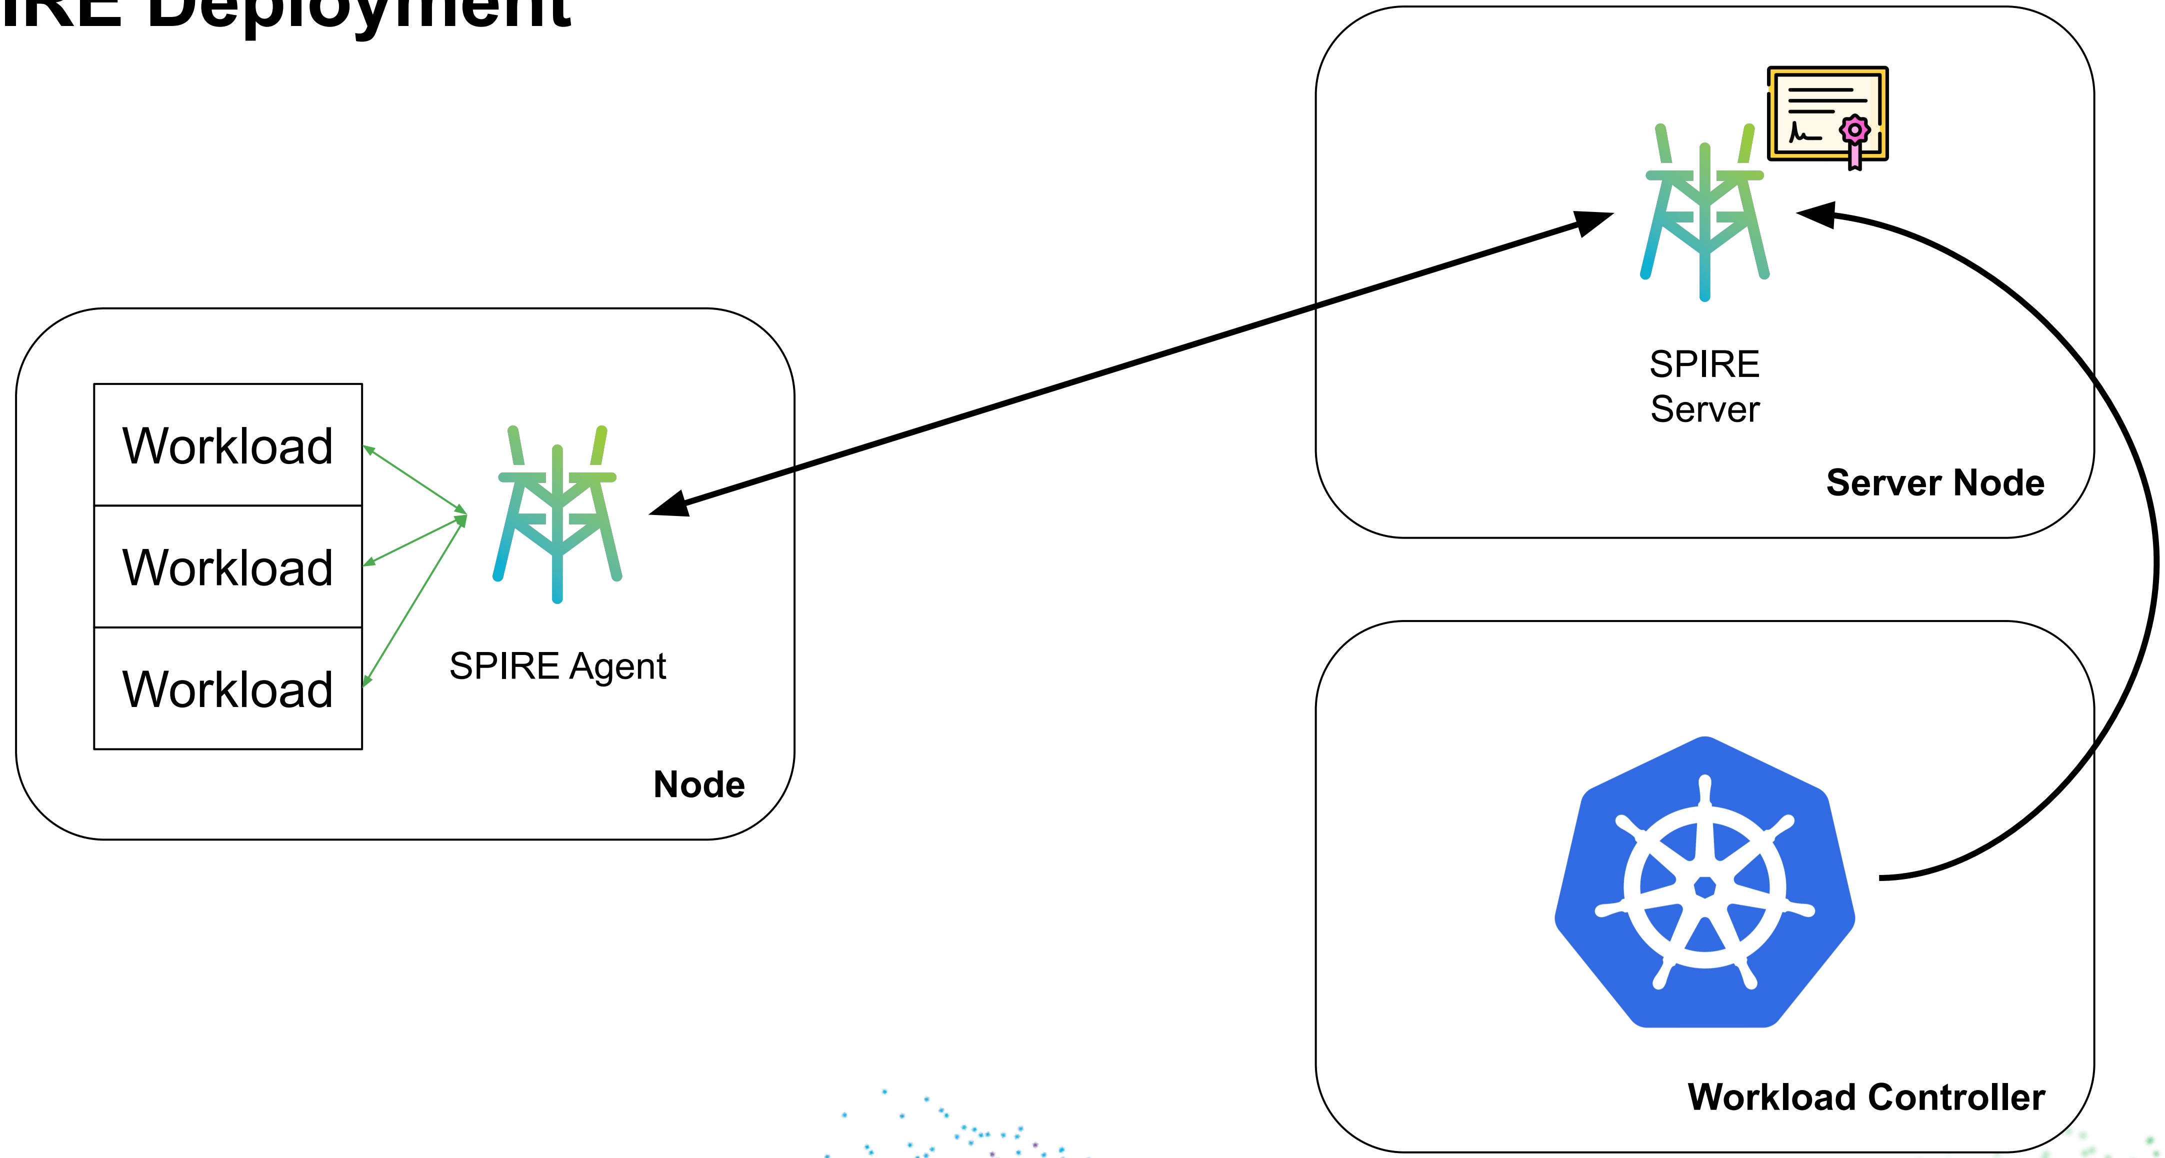

Attestation-Based Identity

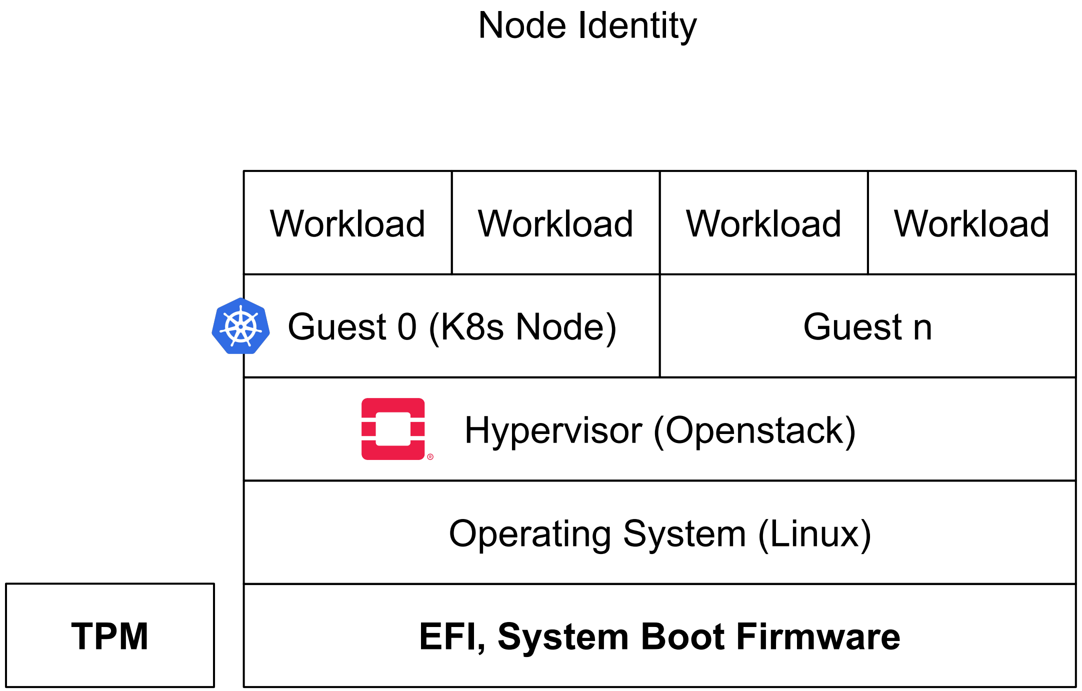

To trust a workload, we want to know:

- Who started this workload/service?
- Who built the kernel and executables running on this virtual machine?
- Who is managing this virtual machine?
- Who built the hypervisor?
- Who built the EFI firmware that booted the machine?

We can answer these questions with:

- Workload orchestrator configuration
- Control plane management
- EFI Secure Boot (boot-time measurement)
- TPM attestation and measurement

Service Identity through Attestation (for K8s)

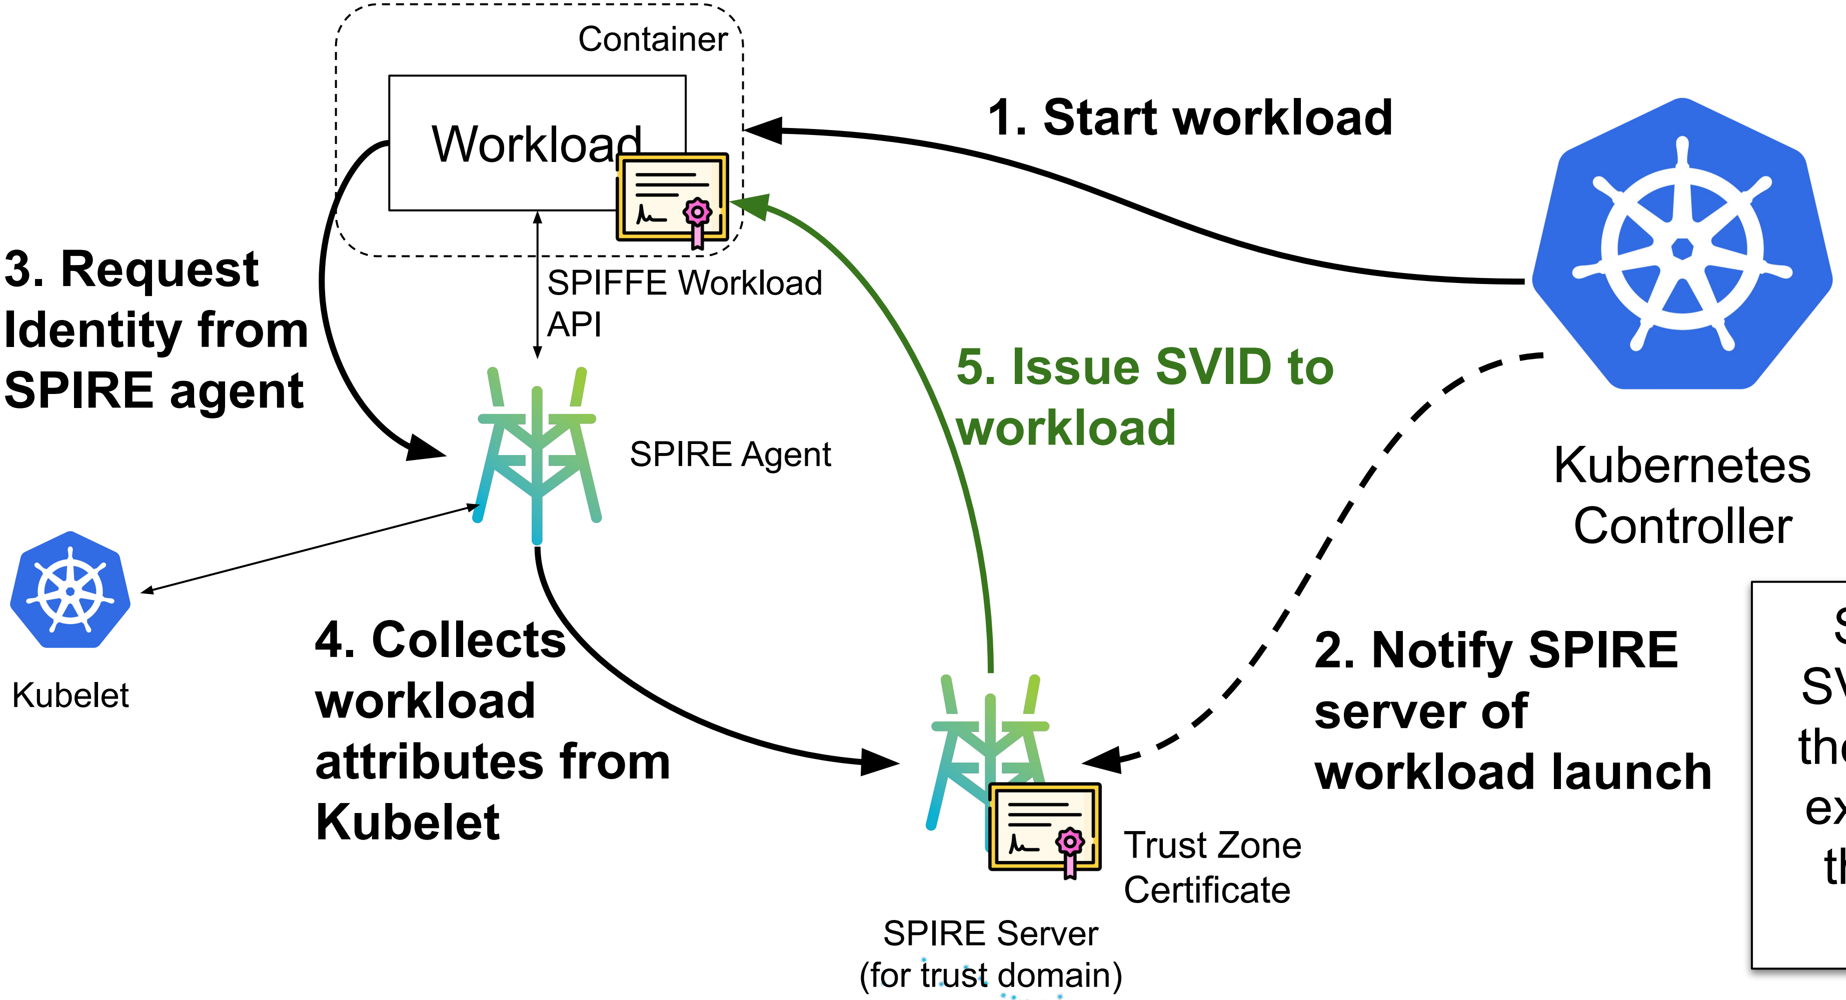

Container
Workload
1. Start workload
3. Request Identity from SPIRE agent
SPIFFE Workload API
SPIRE Agent
5. Issue SVID to workload
Kubernetes Controller
Kubelet
4. Collects workload attributes from Kubelet
2. Notify SPIRE server of workload launch
Trust Zone Certificate
SPIRE Server (for trust domain)

SPIRE issues an SVID ***after*** verifying the workload should exist AND checking the environment's integrity

Different Infrastructure Zones; Federated Identities

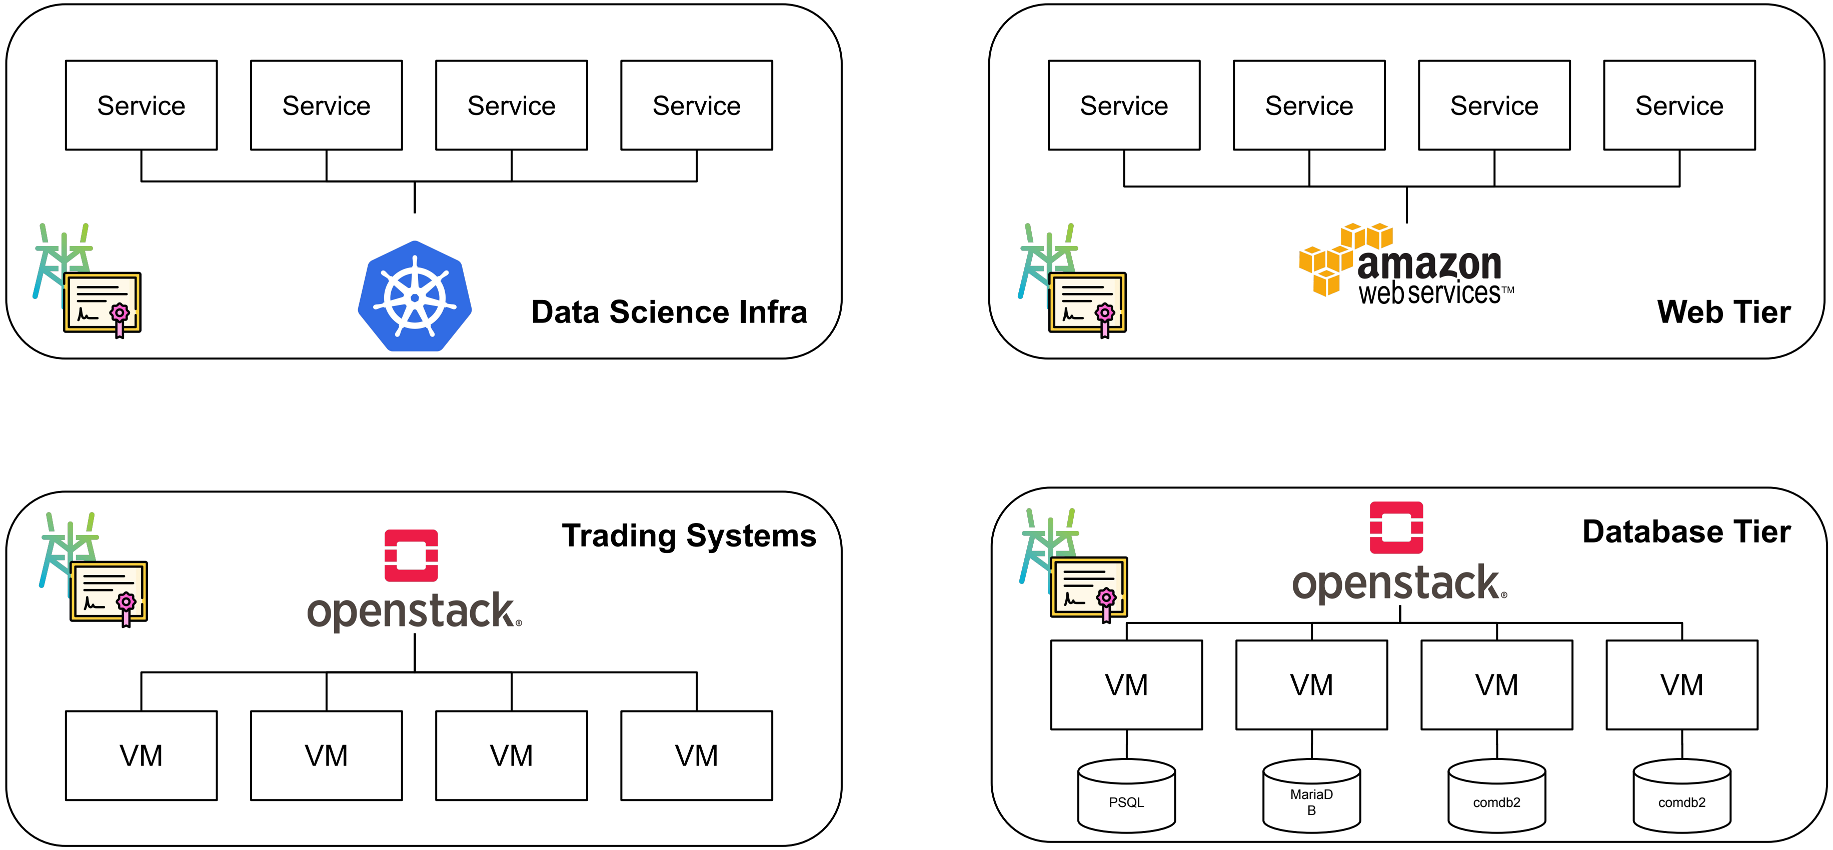

Service
Service
Service
Service
Data Science Infra

Service
Service
Service
Service
amazon web services™
Web Tier

Trading Systems
openstack.
VM
VM
VM
VM

Database Tier
openstack.
VM
VM
VM
VM
PSQL
MariaD B
comdb2
comdb2

Different Infrastructure Zones; Federated Identities

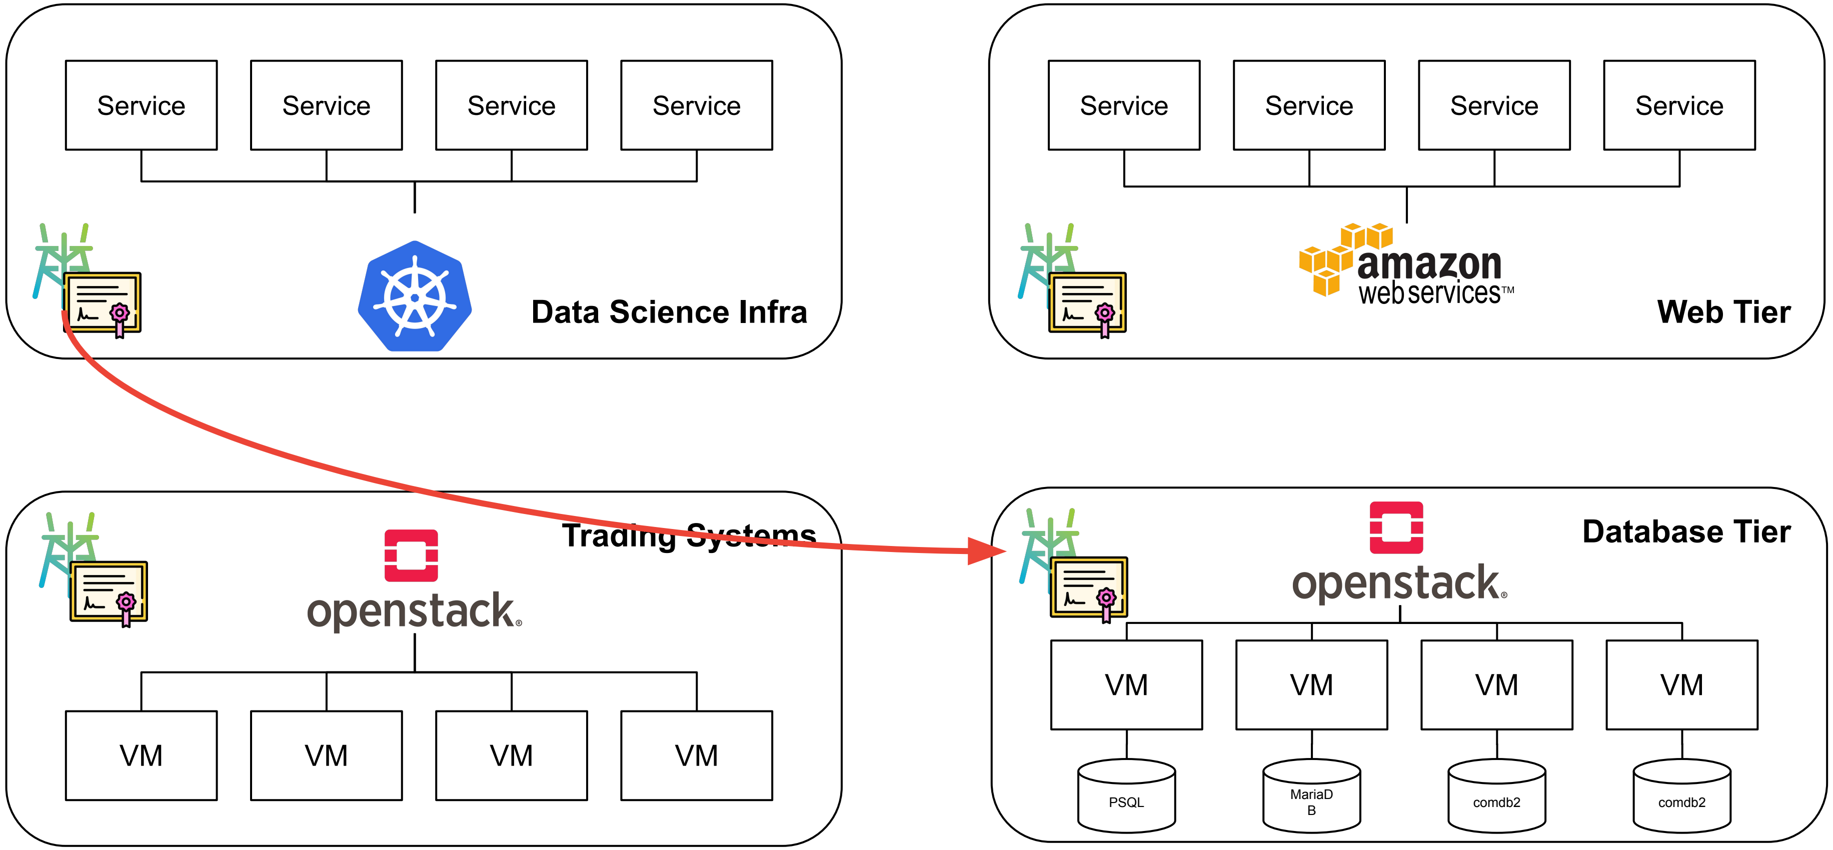

Different Infrastructure Zones; Federated Identities

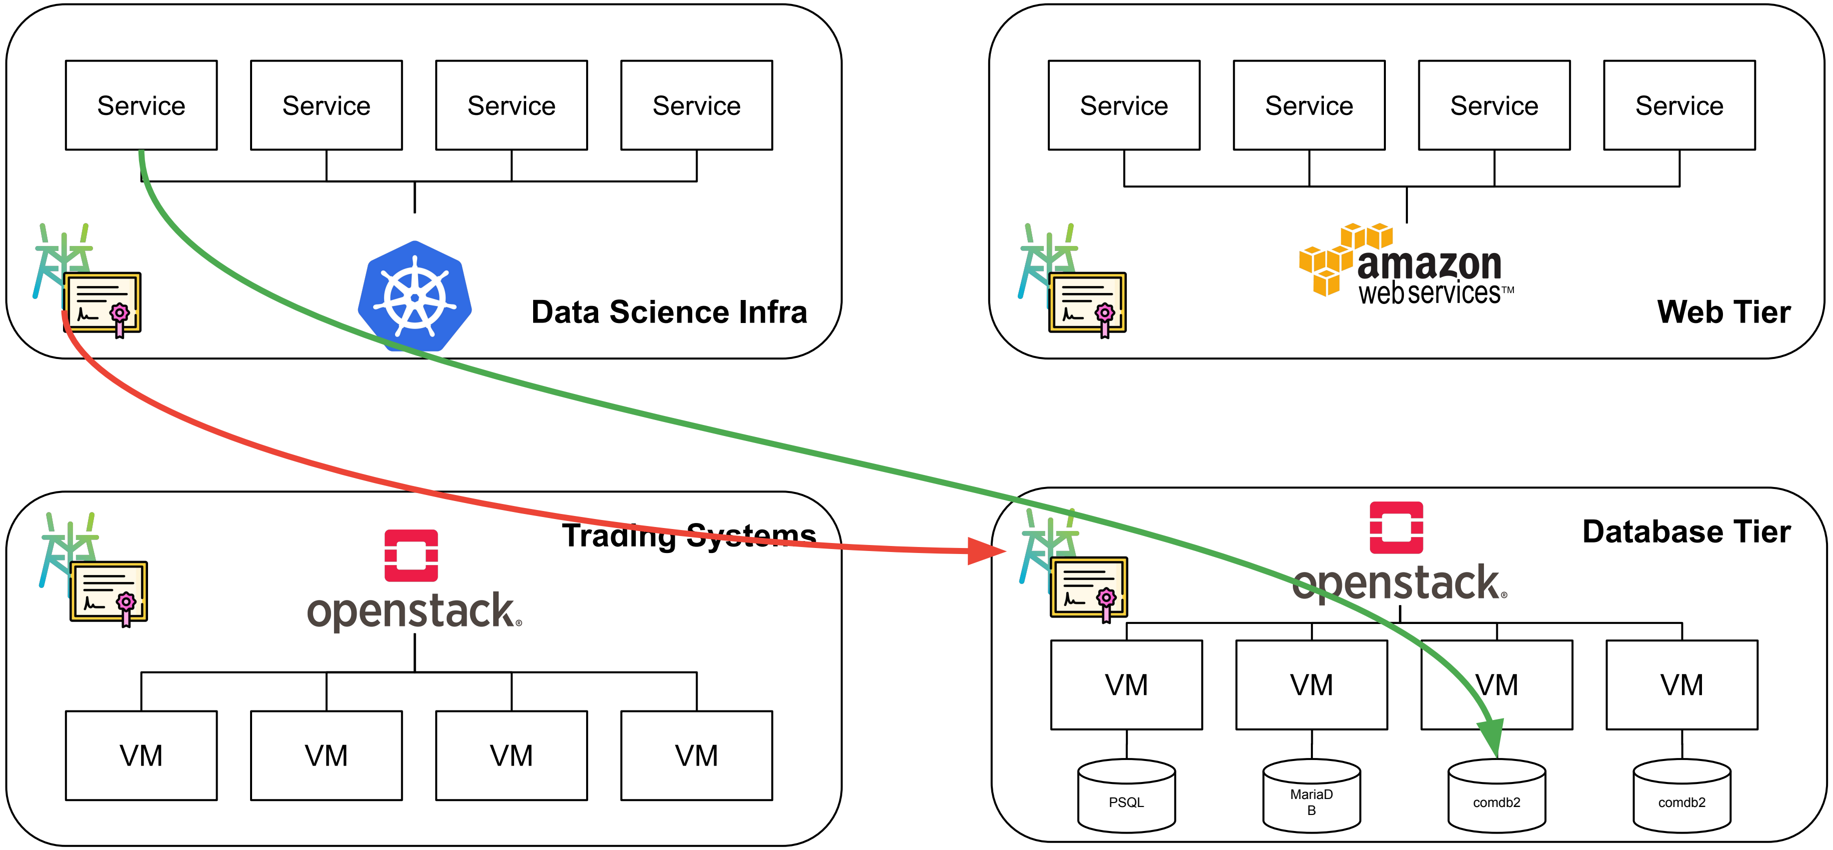

SPIFFE at Bloomberg

Today:

- Attesting and issuing identities for workloads in our PaaS environment

Soon:

- Issuing identities for workloads running in the public cloud
- Managing data center server and service identities

The future:

- Hardware-rooted identities for all services

Thank You!

TechAtBloomberg.com

Bloomberg

Engineering

Questions?

We are hiring for all kinds of Information Security roles!
See https://careers.bloomberg.com to apply, or reach out directly to me.

Goal: Fine-Grained Access Control

Verify identity explicitly

Assume everything is **compromised**

Apply the **principle of least privilege**