Application-Layer Egress Control in Kubernetes

Current Solutions, Future Standards

Joshua Fox

SRECon EMEA 2023, Dublin

Joshua Fox

joshua@doit.com
Senior Cloud Architect

doit

Google Cloud

aws

Experts
2023





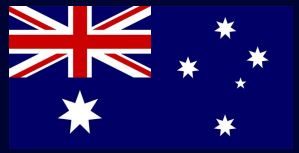
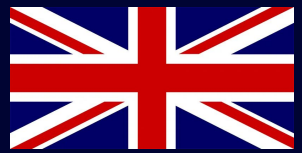

. . .

The challenge
Lock down egress

Threat model

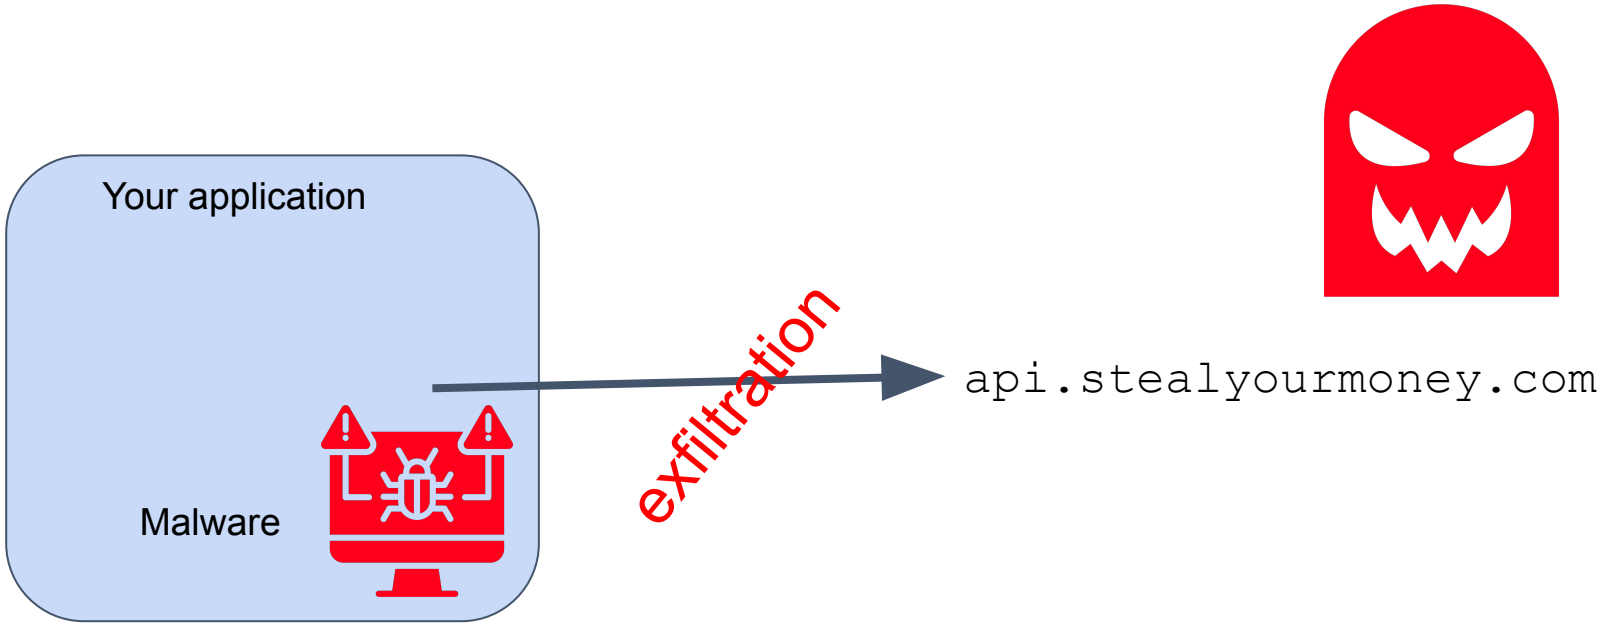

But allow

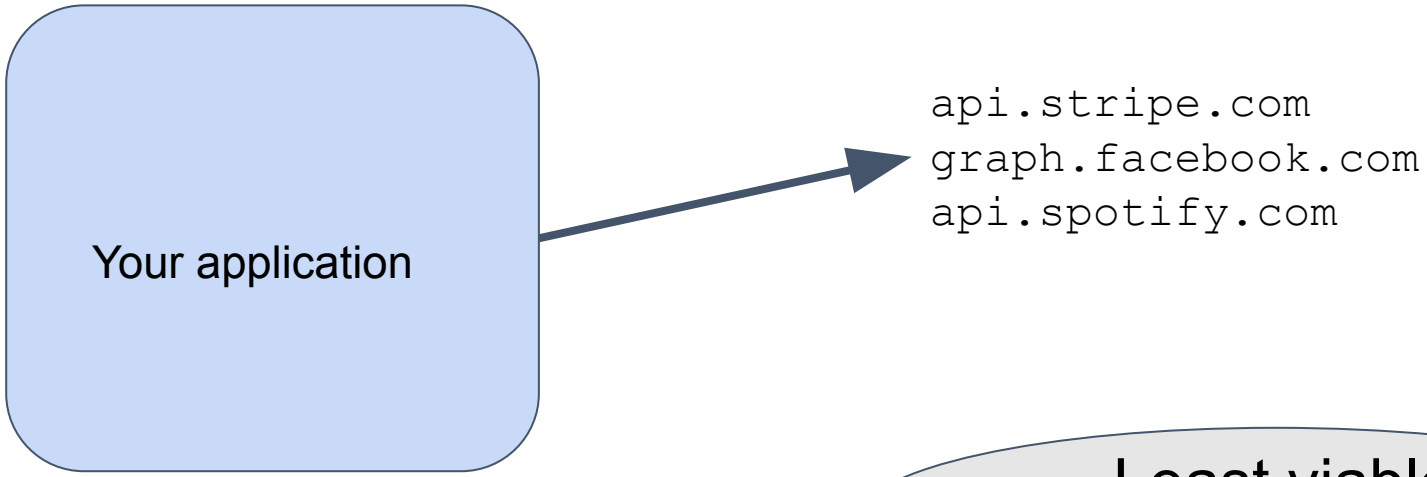

Least viable permissions

Our anonymized example

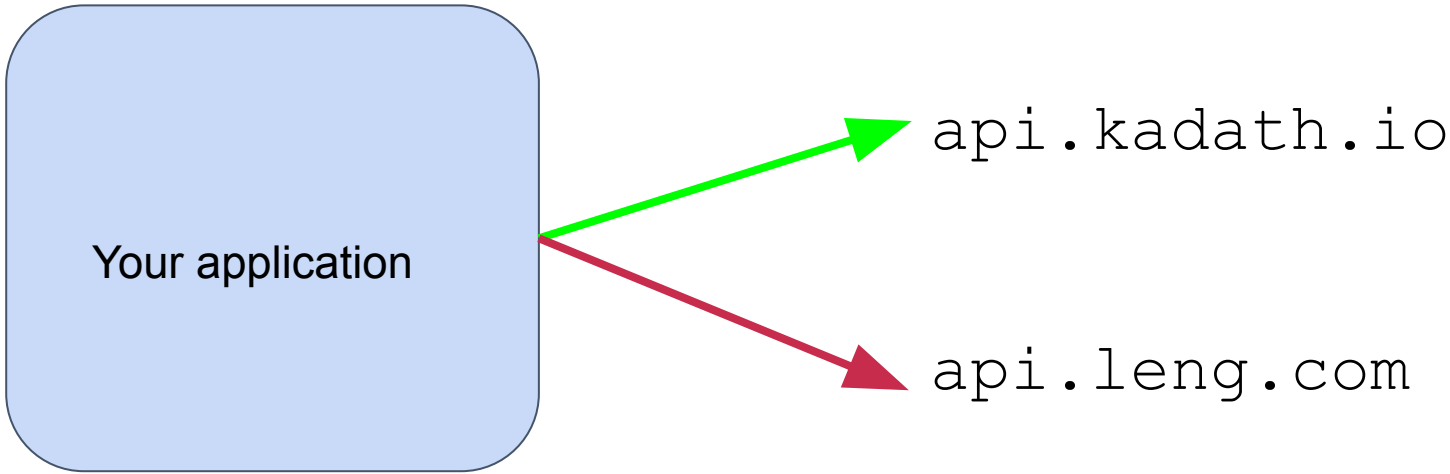

The VPC

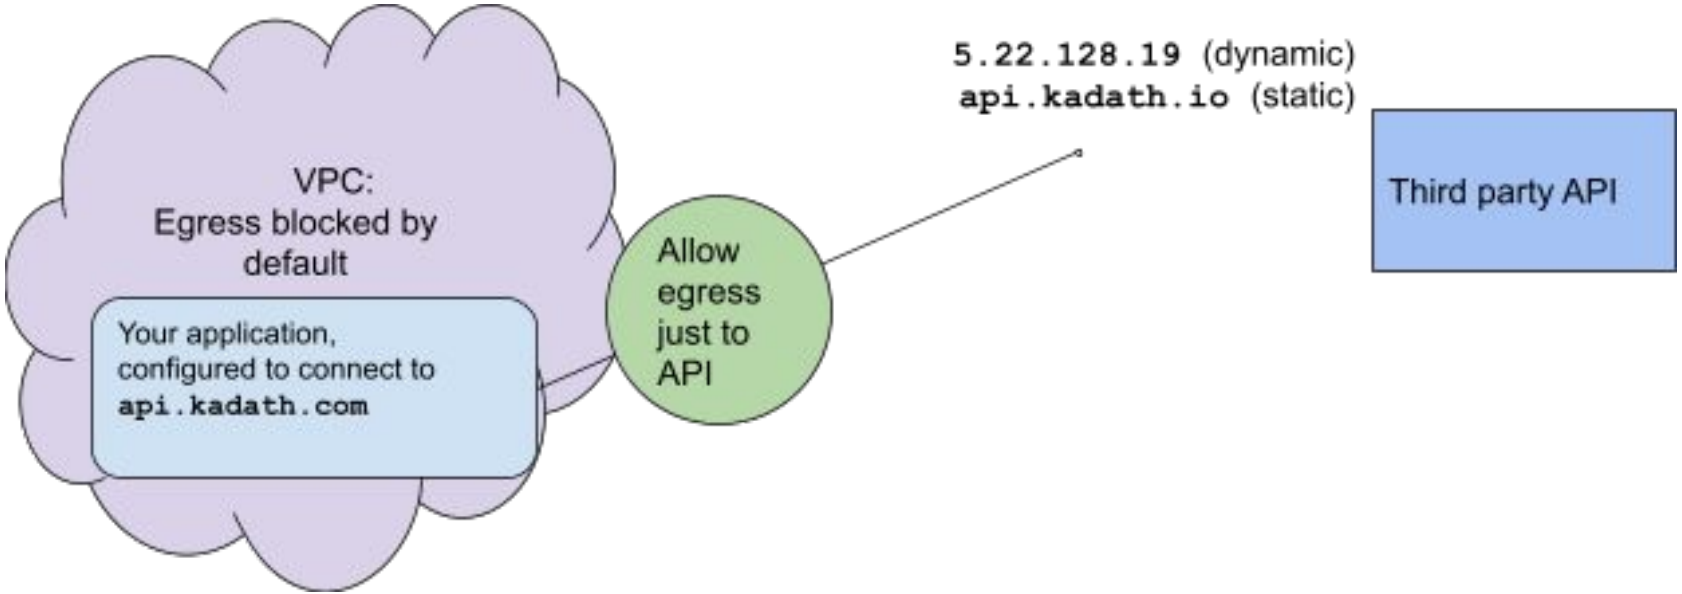
5.22.128.19 (dynamic)
api.kadath.io (static)
Third party API
VPC:
Egress blocked by default
Your application, configured to connect to api.kadath.com
Allow egress just to API

Multiple layers
Each firewall works on a different one

The layers

Application Layer
FQDN
Layer 7

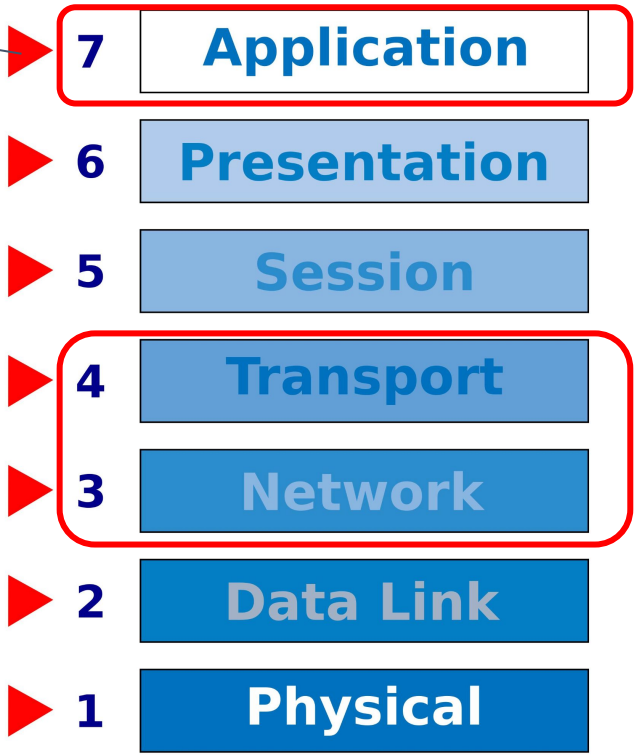

HTTP, DNS

TCP, UDP

IP

Rapidly evolving
My articles and talks

TFiR article re Kubernetes

https://bit.ly/fox-egress

KubeCon | CloudNativeCon

Europe 2023

Re networking

bit.ly/egress-control

Talks


doit
doit
Level 7 Egress Control in Kubernetes
Devopsdays Ukraine, Vilnius
Joshua Fox
joshua@doit.com
Experts
2023


Doit
INTERNATIONAL
שליטה על יציאה מרשת
לפי מתחם
צרות, טעויות, וחוסר פתרונות
יהושע פוקס
אדריכל ענן בכיר
דואיט בינלאומי


Doit
INTERNATIONAL
ווי מען רעגולירט אַרויסגאַנג לויטן ראַיאָן
כטולהון ארט עס נישט
יהושע פּוקס
עלטערער וואָלקן־אַרכיטעקט
כטולהו
עלטערער גאָט וואָס חלומט טויט אין
דיל ים


Firewalls

Firewall

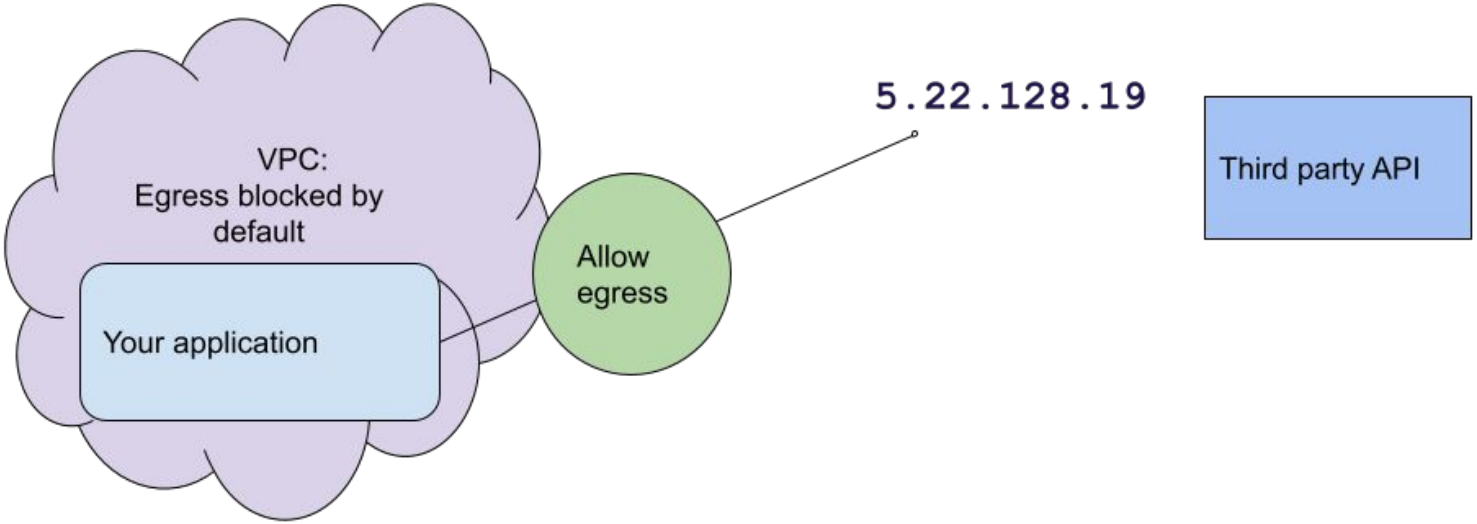
VPC:
Egress blocked by default
Your application
Allow egress
5.22.128.19
Third party API

Google Firewall

Summary of AWS Firewall types

Security Groups
Network ACL
Network Firewall
Shield/WAF
Route 53 Firewall

Article on AWS Firewall types

https://bit.ly/aws-firewalls

AWS Security Groups

AWS Network ACL

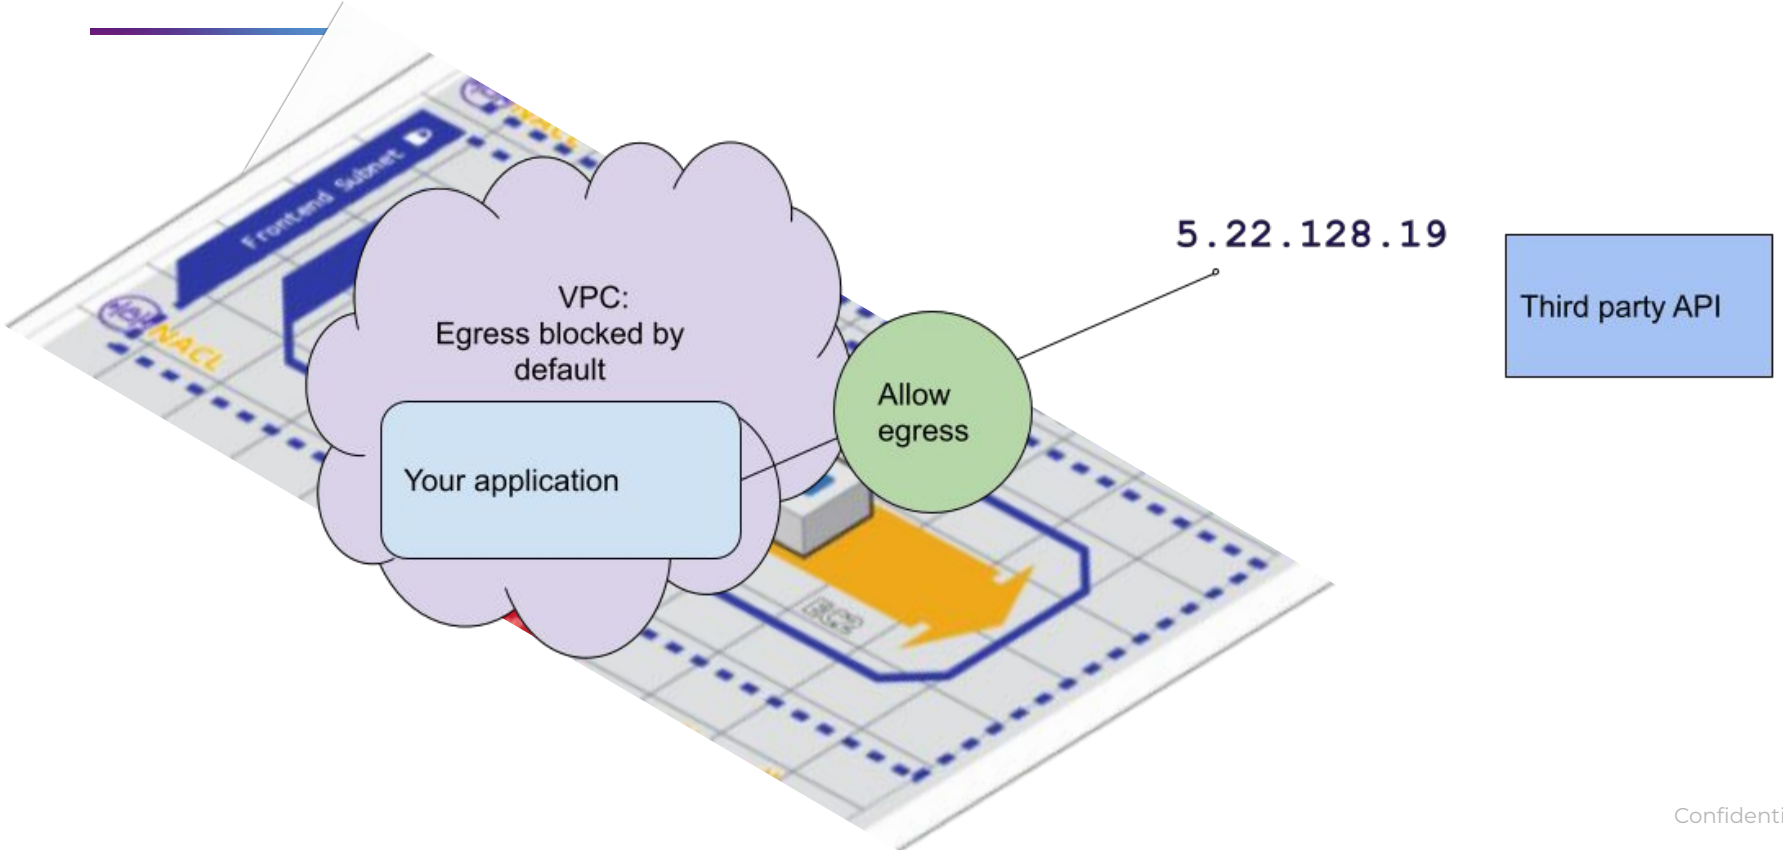

IP address changes

5.22.128.15

api.kadath.io

~~api.leng.com~~

Implementations
Hosted

Implementation

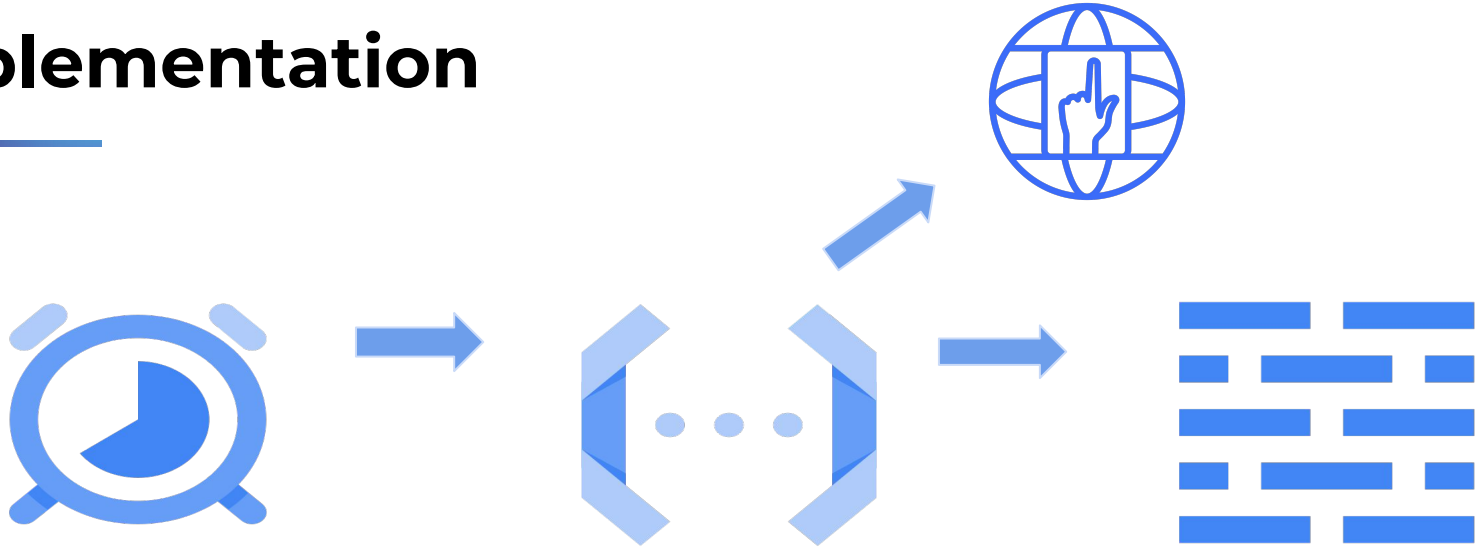

5.22.128.19

5.22.128.21

api.kadath.io

~~api.leng.com~~

Don’t DIY!

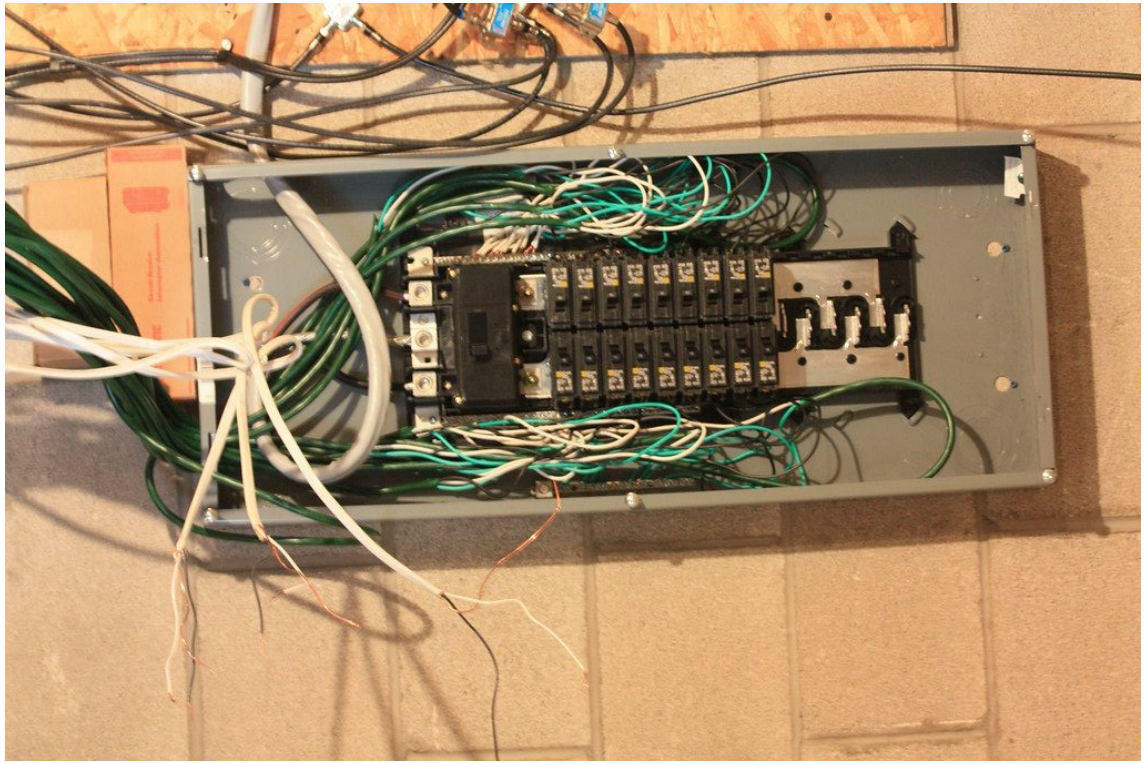

Squid Proxy

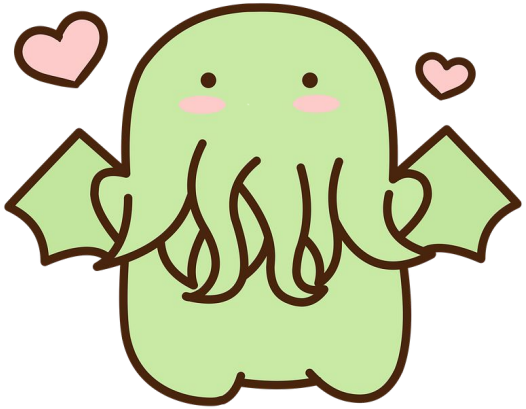

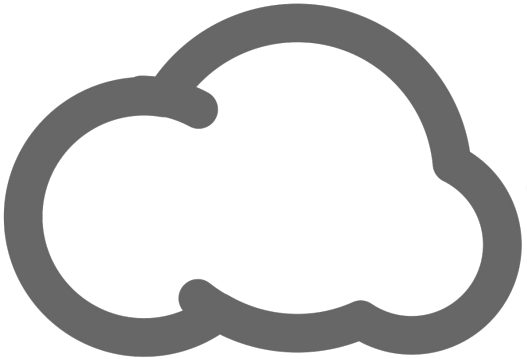

api.kadath.io

5.22.128.19

5.22.128.21

Squid Proxy

```
acl whitelist dstdomain api.kadath.io
http_access allow whitelist

http_access deny all
```

Multiple hosts on an IP address

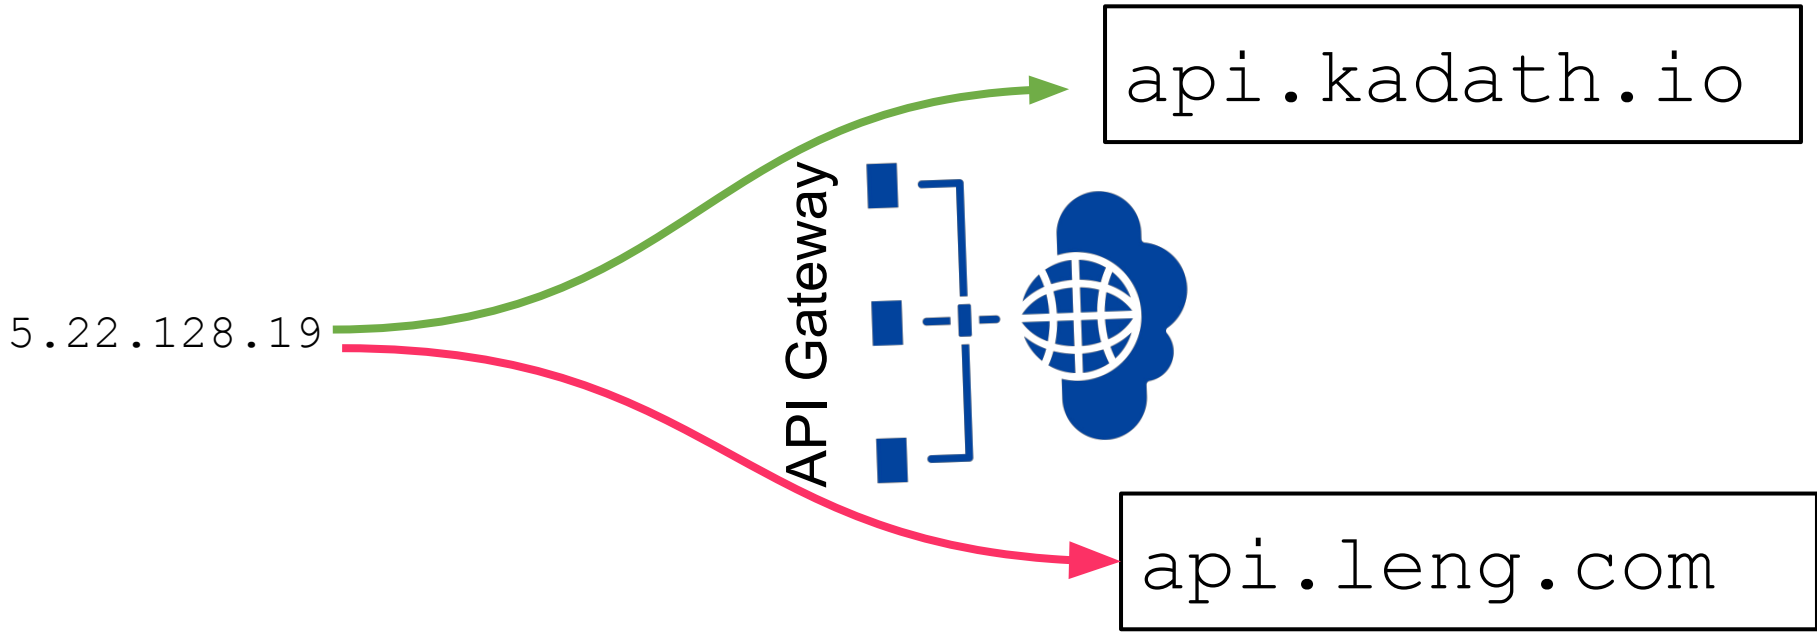

Server Name Indication

HTTPS Flow & SNI

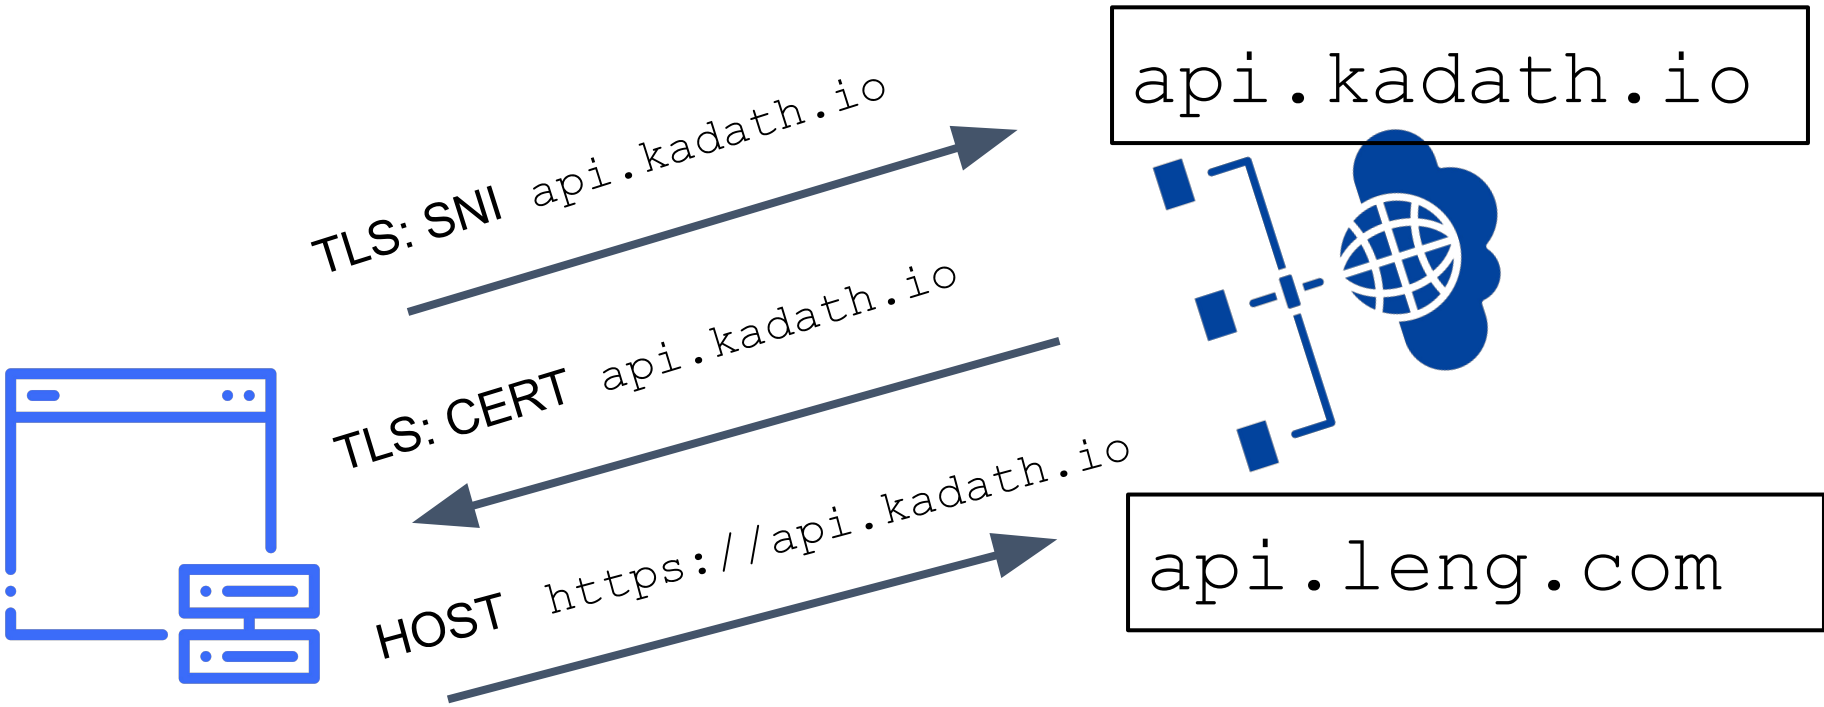

Encrypted SNI

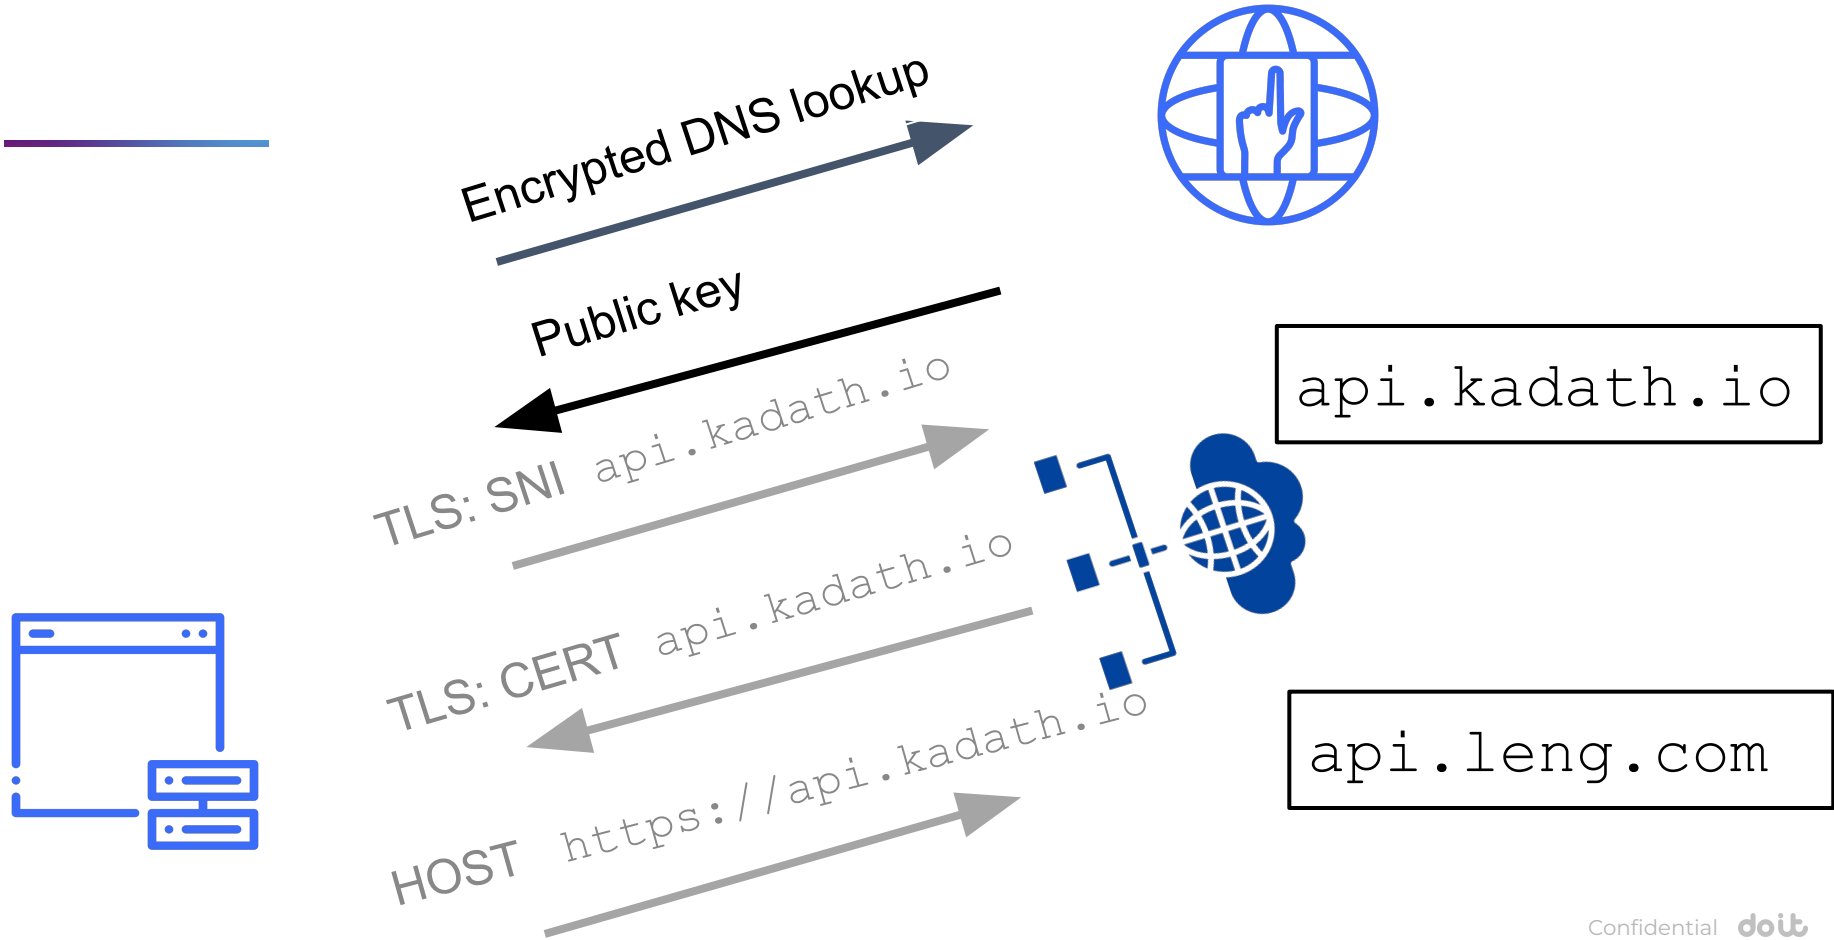

Squid Proxy: SslBump Peel and Splice

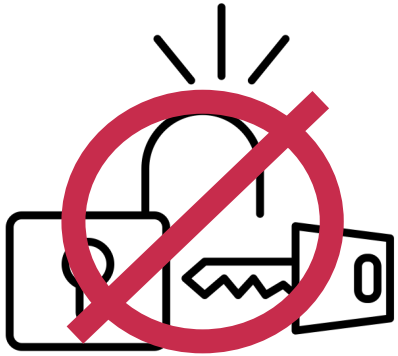

```
acl step1 at_step SslBump1
acl whitelist_ssl ssl::server_name "/etc/squid/whitelist"
acl whitelist dstdomain "/etc/squid/whitelist"
...
ssl_bump peek step1
ssl_bump splice whitelist_ssl
ssl_bump terminate all !whitelist_ssl
```

Commercial self-hosted

DiscrimiNAT

Aviatrix FQDN Egress Filtering on AWS

Don’t use a hosted solution

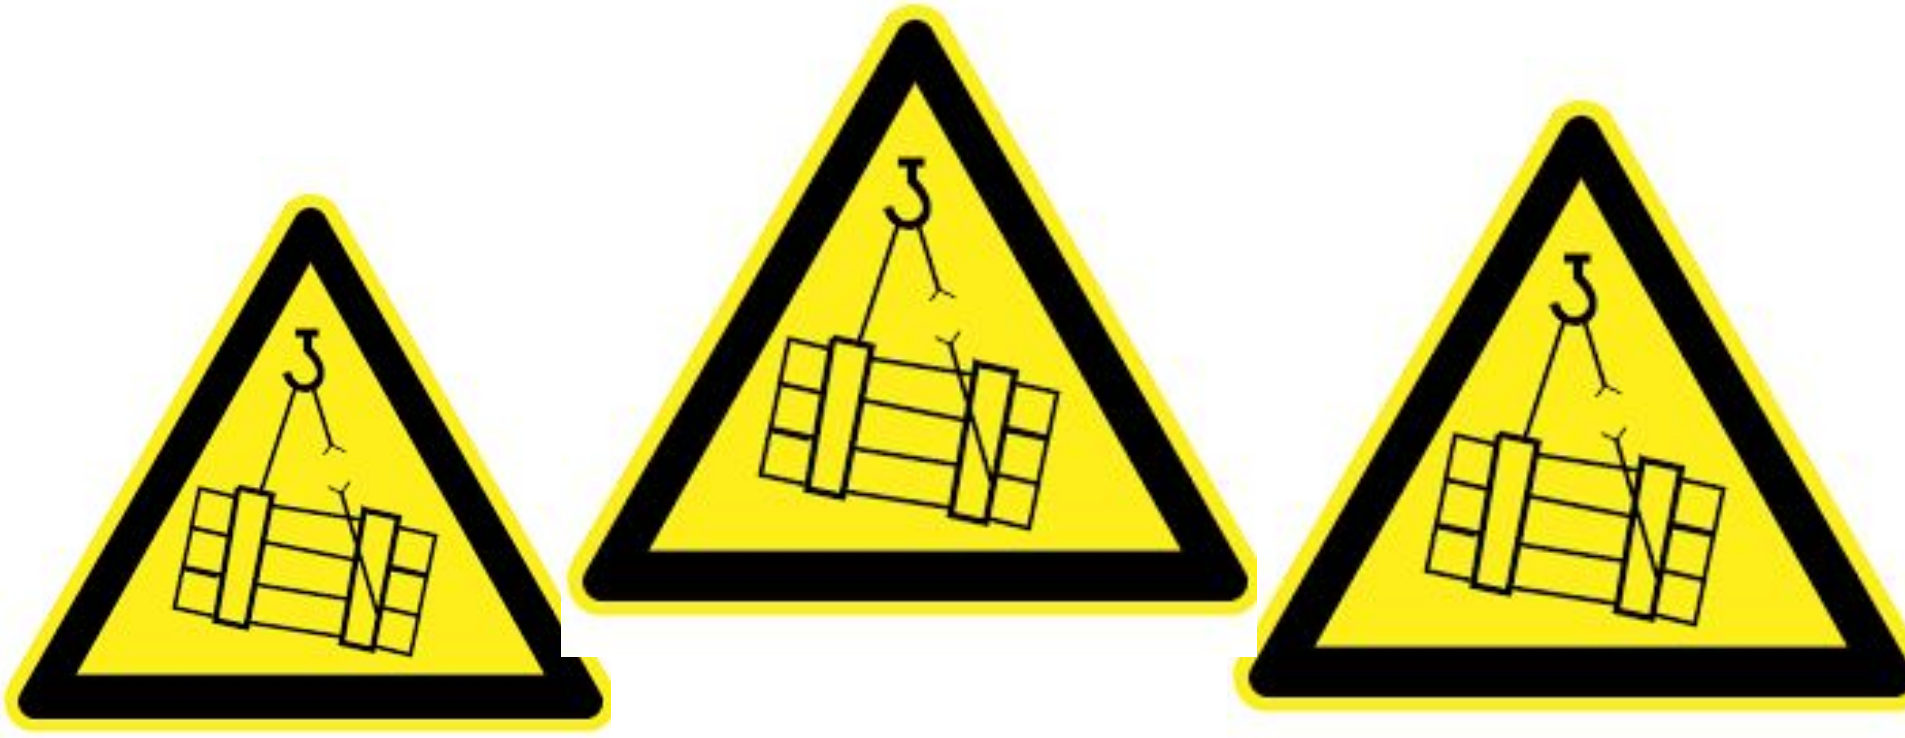

Services,
AWS services

AWS Network Firewall

With SNI!

~~api.leng.com~~

api.kadath.io

AWS Network Firewall

```
{
    "RulesSource": {
        "RulesSourceList": {
            "Targets": [
                "api.kadath.io"
            ],
            "TargetTypes": [
                "HTTP_HOST",  "TLS_SNI"
            ],
            "GeneratedRulesType": "ALLOWLIST"
        }}}
```

HTTPS

Route 53 DNS Firewall

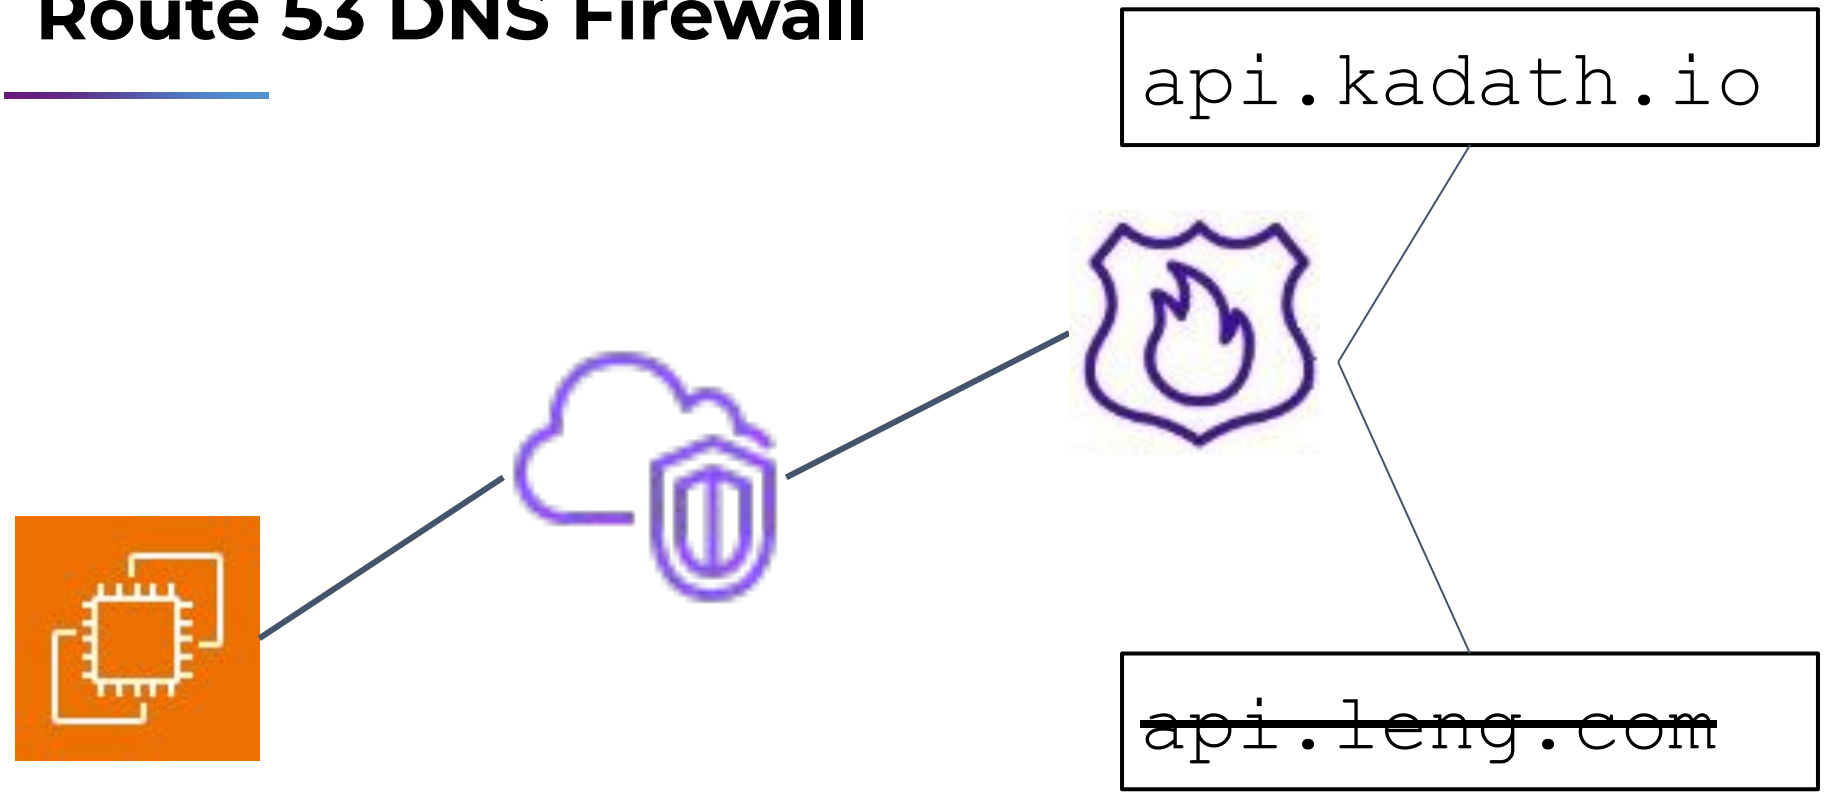

Implementations
Google services

Google Secure Web Proxy

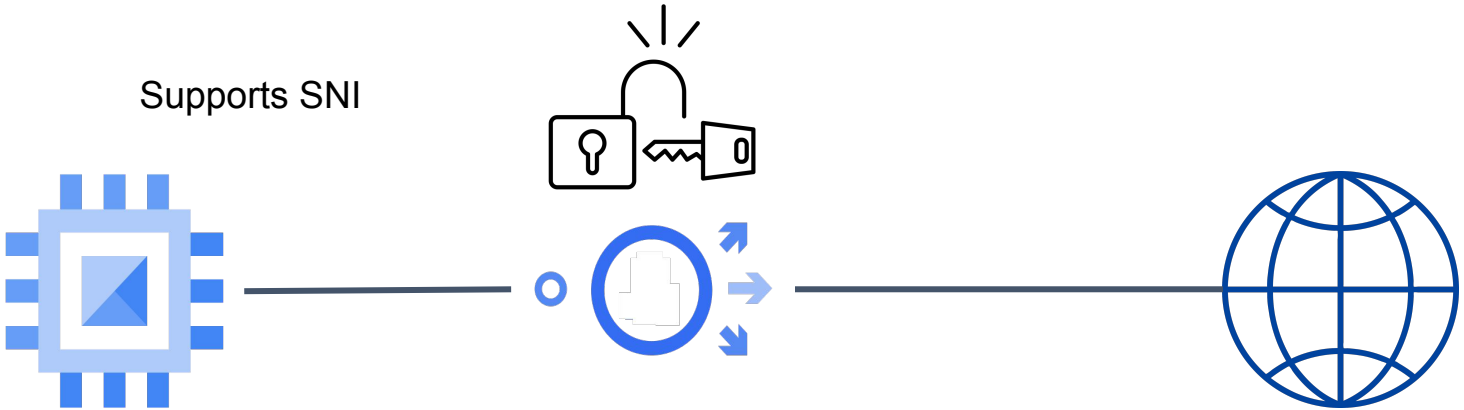

[See the first article by Chimbu Chinnadurai]

Google FQDN Firewall Objects

No SNI

30 sec

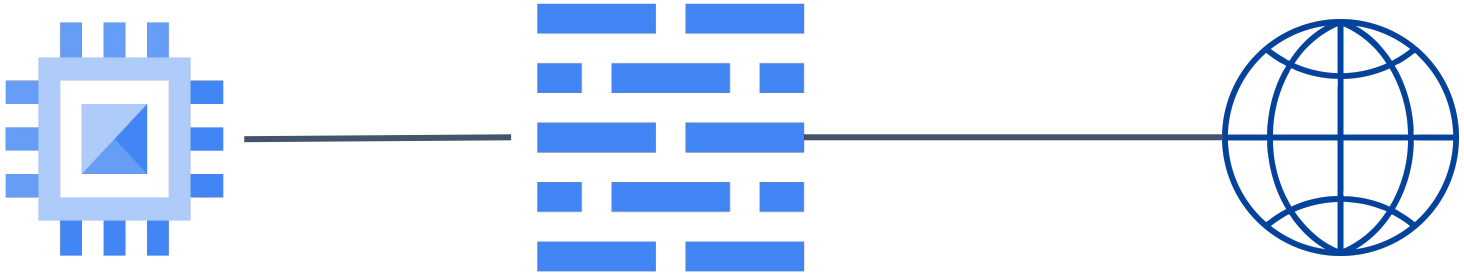

[See the second article by Chimbu Chinnadurai]

Kubernetes
Another world

Cluster and Namespaces

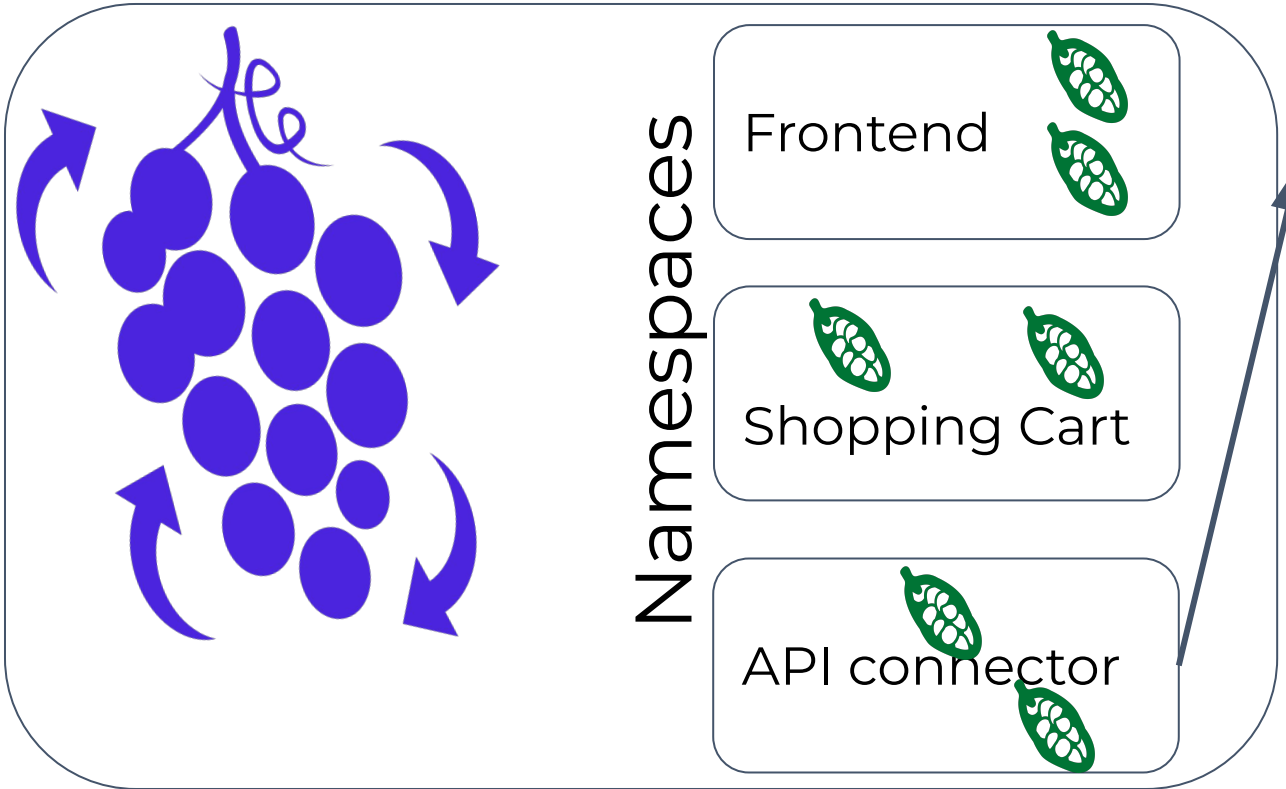

Cluster and Labels

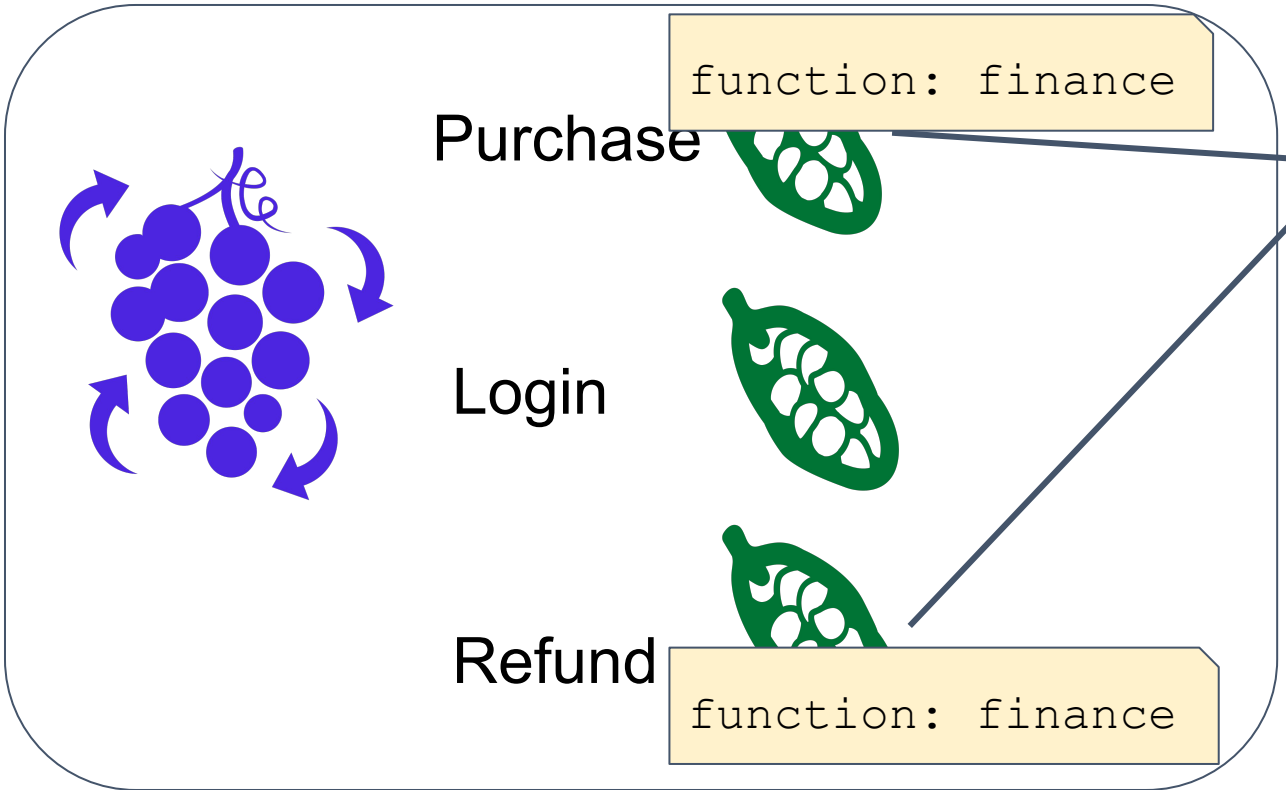

Plain: Kubernetes Network Policy

```
apiVersion: networking.k8s.io/v1
kind: NetworkPolicy
...
spec:
 podSelector:
   matchLabels:
     networking/access-kadath-api: "true"
 egress:
 - to:
   - ipBlock: 5.22.128.133/32
 policyTypes:
 - Egress
```

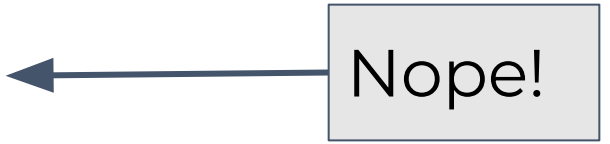

Cilium

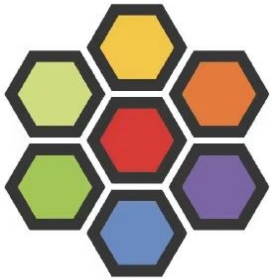

eBPF-based Networking layer

On each host

Kernel-adjacent

Cilium

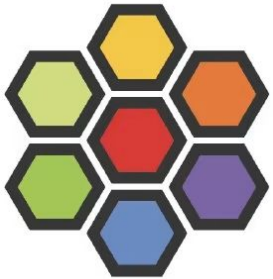

eBPF-based Networking

`Cilium Network Policy`

With SNI!

`Cilium Clusterwide Network Policy`

Cilium Network Policy

```
apiVersion: "cilium.io/v2"
kind: CiliumNetworkPolicy
metadata:
 name: "fqdn"
spec:
 endpointSelector:
   matchLabels:
     function: api-conector
 egress:
 - toFQDNs:
   - matchName: "api.kadath.io"
```

Another clause blocks egress from all other pods.

Cilium

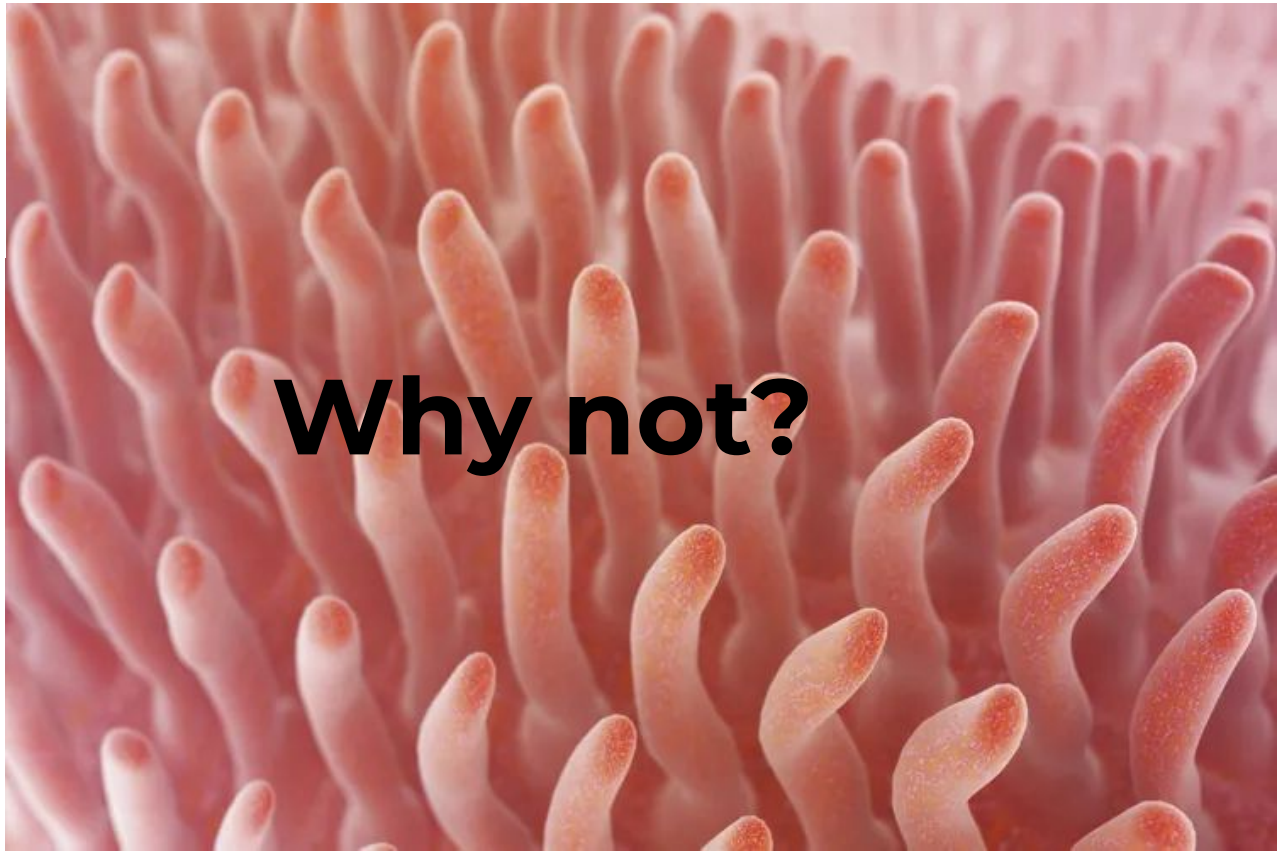

Istio

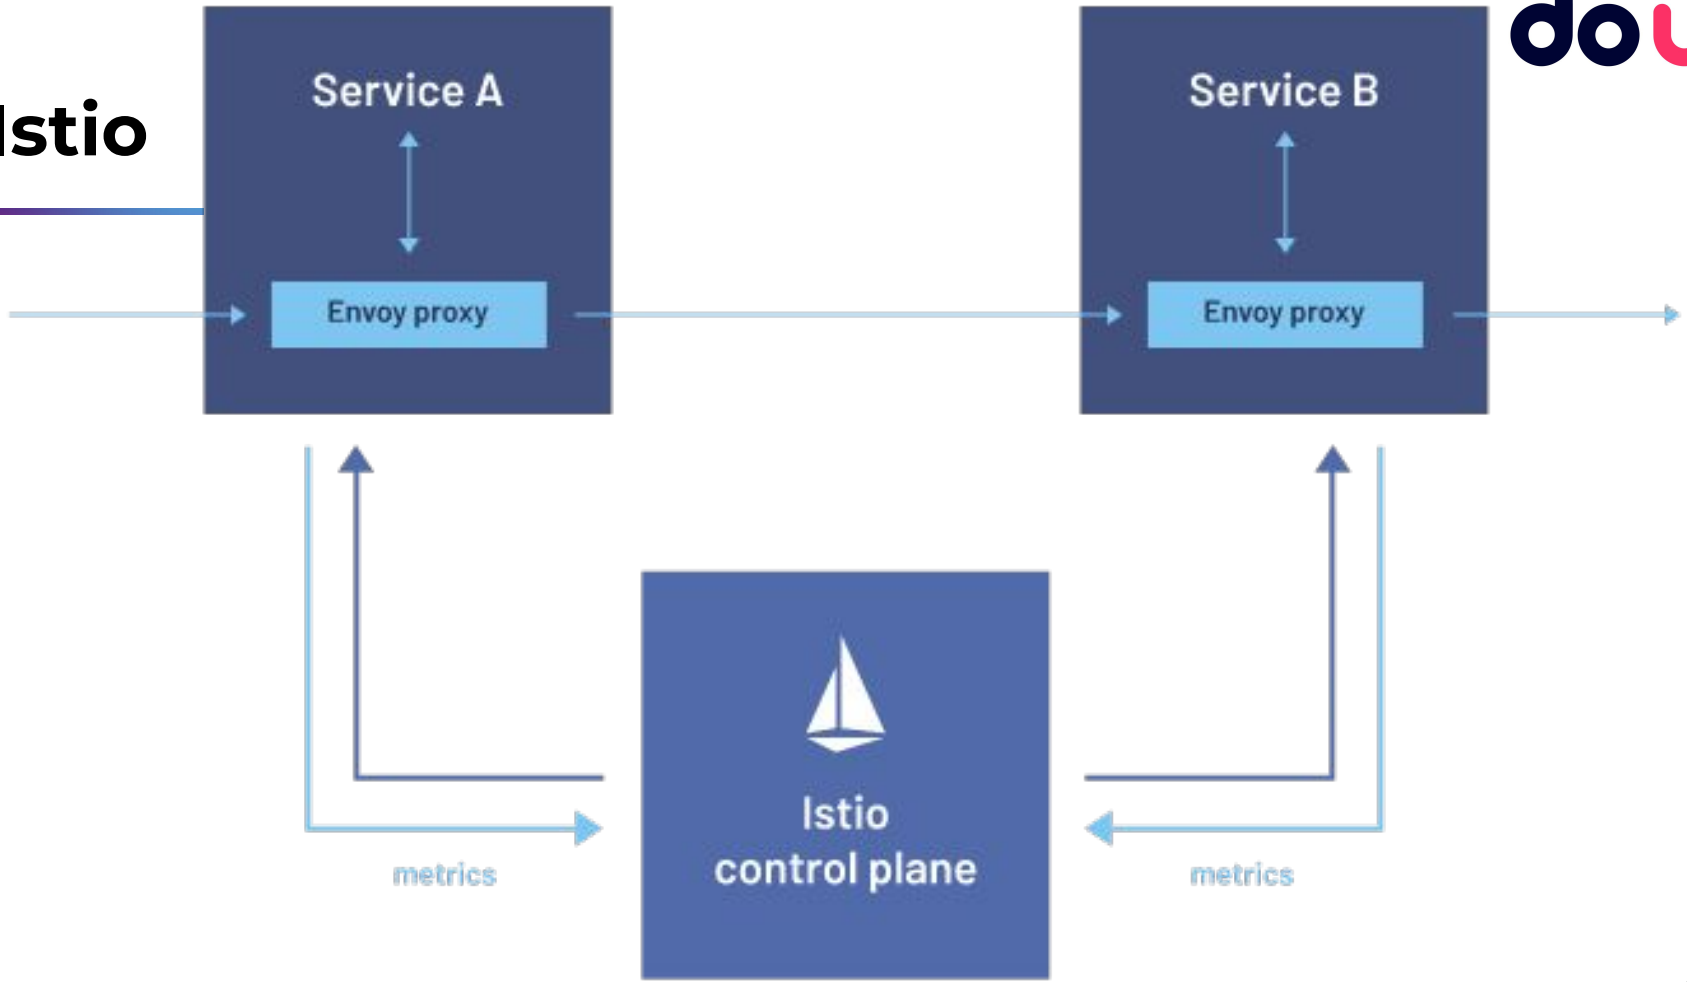
Service A
Envoy proxy
Service B
Envoy proxy
Istio
control plane
metrics
metrics

Istio ServiceEntry

```
apiVersion: networking.istio.io/v1alpha3
kind: ServiceEntry
metadata:
  name: kadath
spec:
  hosts:
  - ".kadath.io"
  - location: MESH_EXTERNAL
  ports:
  - number: 443
    name: https
    protocol: HTTPS
```

With SNI using `sniHosts`

Istio

Why not?

Kubernetes Network Policy

```
apiVersion: networking.k8s.io/v1
kind: NetworkPolicy
...
spec:
 podSelector:
   matchLabels:
     networking/access-kadath-api: "true"
 egress:
 - to:
   - ipBlock: 5.22.128.133/32
 policyTypes:
 - Egress
```

GKE: FDQN K8s Network Policies on DataplaneV2

```
apiVersion: networking.gke.io/v1alpha1
kind: FQDNNetworkPolicy
metadata:
 name: allow-out-fqdnnp
spec:
 podSelector:
   matchLabels:
     run: curl
 egress:
 - matches:
   - name: "api.kadath.io"
   ports:
   - protocol: "TCP" # to allow only HTTPS
     port: 443
```

See the third article by Chimbu Chinnadurai

Committee

- Kubernetes Project
- Special Interest Group Networking
- Extend K8s `NetworkPolicy`

Technique	Advantage	Disadvantage
Squid & other self-hosted	Mature	You manage it, no K8s
Level 7 Firewall-as-a-service	As-a-Service	No K8s
Cilium Network Policies, Istio Service Mesh	Supports K8s, powerful, generally available	Overpowered
GKE K8s Network Policies	Built-in to K8s	Non-standard
Standard K8s Network Policies	Built-in to K8s, standard	Doesn’t exist

See the slides

doit

https://bit.ly/egress-k8s

We’re hiring!

joshua@doit.com

Questions?