

# Rethinking the K8s DNS for the Modern Enterprise

KubeCon NA 2019

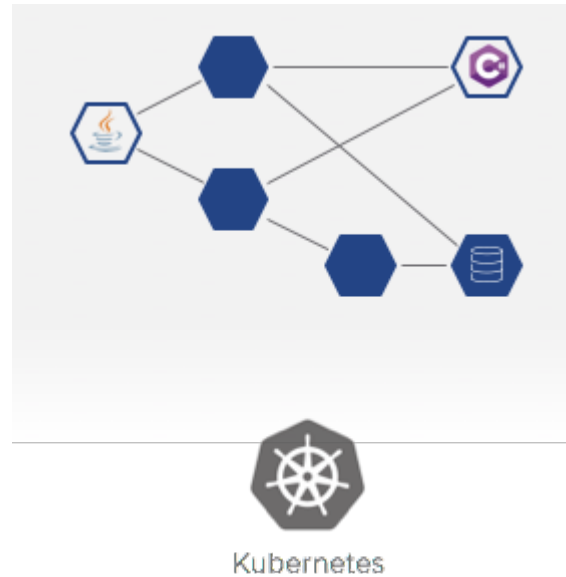
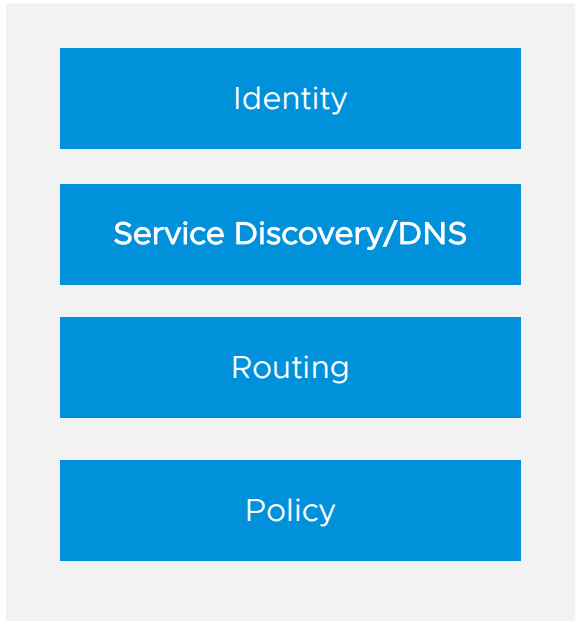
**Deepa Kalani**

Staff Engineer 2  
NSX Service Mesh  
VMware  
dkalani@vmware.com


**Venil Noronha**

Member of Technical Staff  
NSX Service Mesh  
VMware  
veniln@vmware.com

# Service Mesh Capabilities




# Names are complicated...



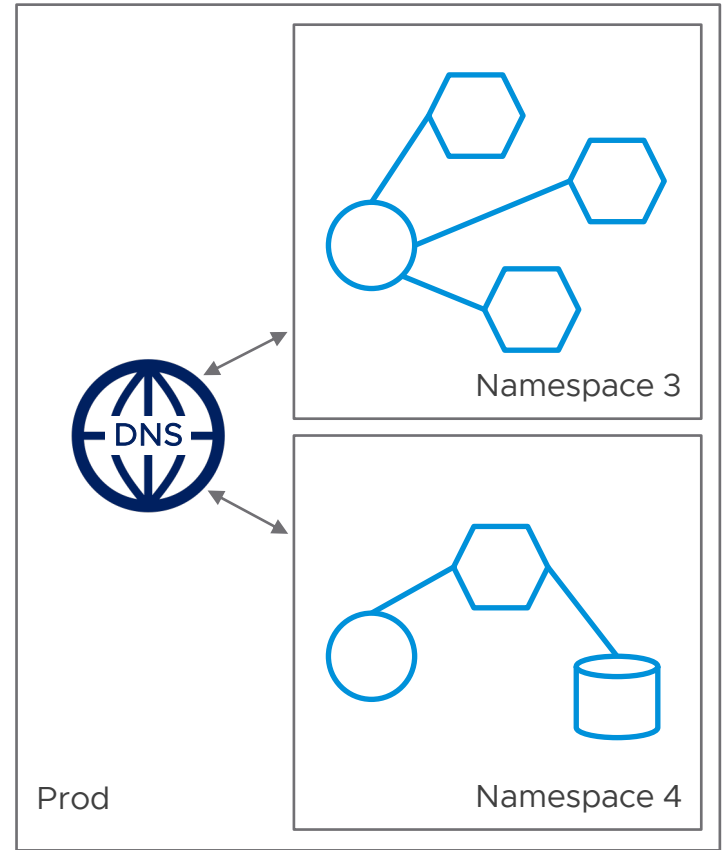
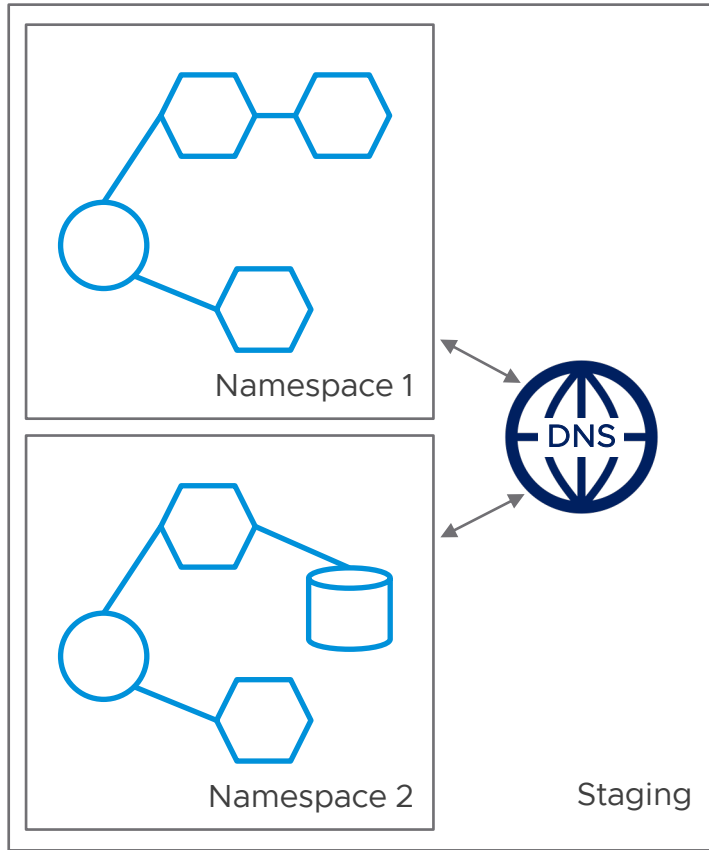
# Application Migration – Simplification through Naming

- Multi-cloud and hybrid-cloud systems
- In a multi-cloud world, applications may be deployed on prem and in the cloud
- Developers should be able to deploy and migrate applications across any cloud provider without changing their native workloads






# DNS Isolation – Enabler for Multi-tenant Clusters




# DNS Observability and Security

- Rich telemetry for DNS queries and responses
- Telemetry per tenant
- Open the door to behavioral analysis based on telemetry data.




# DNS Filtering

- Operators need a way to specify filtering the DNS layer
- DNS policies allow for access control and logging
- Example:
  - *Deny the frontend service from discovering \*.com and log such requests*
  - *TenantA services should not discover tenantb.services*
- Treat DNS just as another entity in the Kubernetes cluster
- Apply L4/L7 policies based on DNS queries/responses



# DNS Evolution

- Some tenants might want to encrypt DNS queries to maintain privacy
- Imperative in a multi-tenant environment
- Upgrade UDP/TCP DNS queries to DoT/DoH

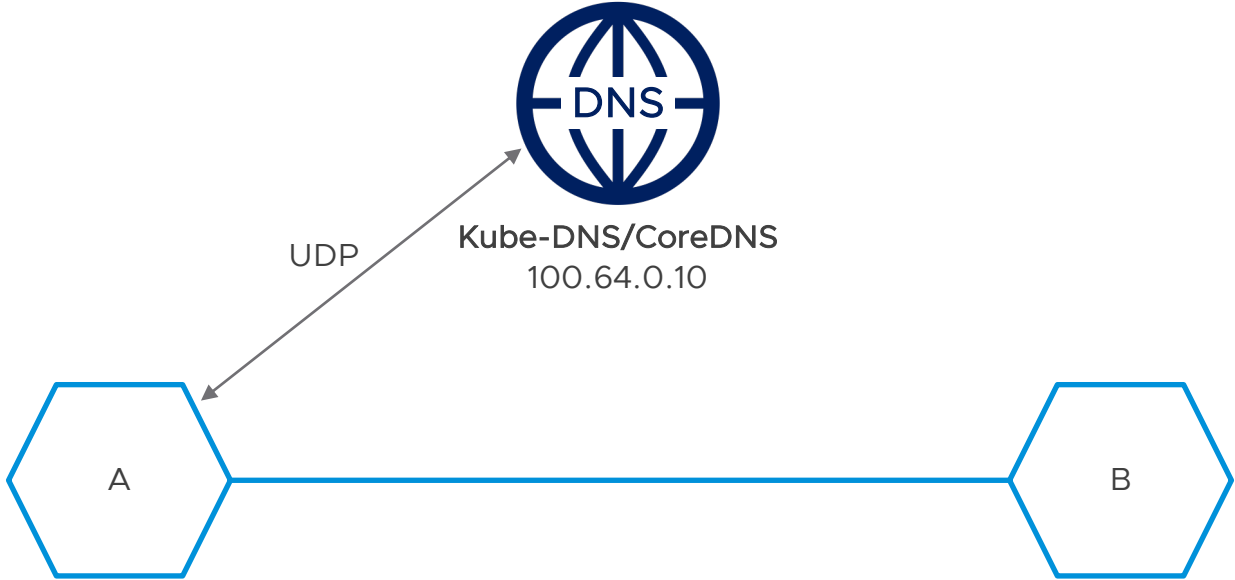


# Current State of Kubernetes DNS

- No tenant isolation for DNS
- No dynamic configuration of DNS
  - Can't configure search domains dynamically
  - Can't configure nameservers dynamically
- Policies cannot be enforced at the DNS layer
- Doesn't provide first-class support for secure DNS
  - DNS-over-TLS (DoT)
  - DNS-over-HTTPS (DoH)



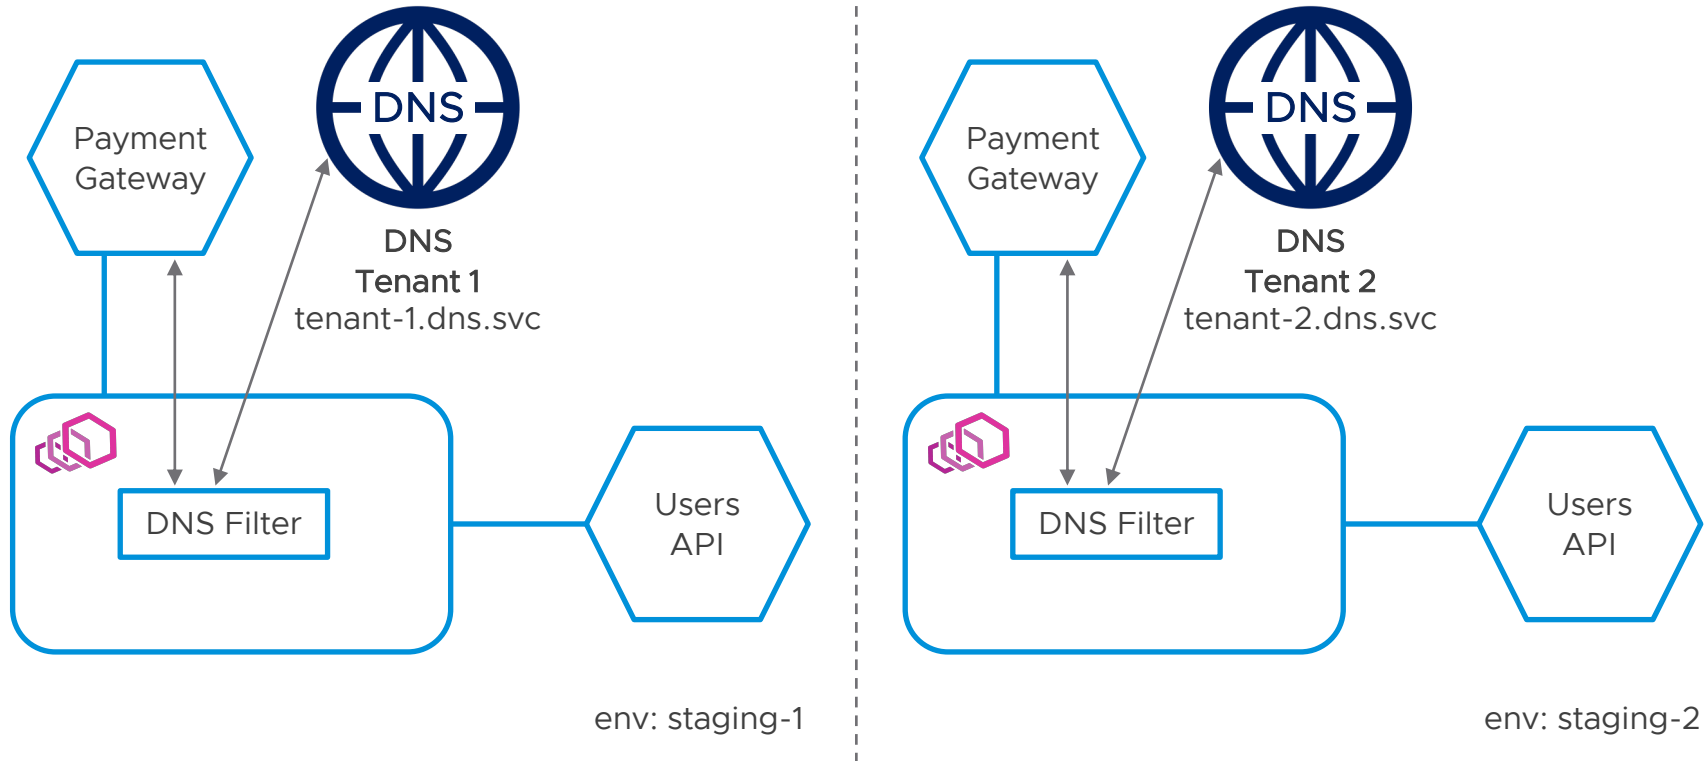
# Plain Old Kubernetes DNS



# DNS Isolation – Enabler for Multi-tenant Clusters

```
1  apiVersion: "networking.example.com/v1alpha1"
2  kind: DnsPolicy
3  metadata:
4    name: dns-policy
5  spec:
6    namespaceSelector:
7      matchLabels:
8        env: staging-1
9    server:
10     # The DNS server address for the tenant.
11     address: tenant-1.dns.svc
```

# DNS Isolation – Enabler for Multi-tenant Clusters

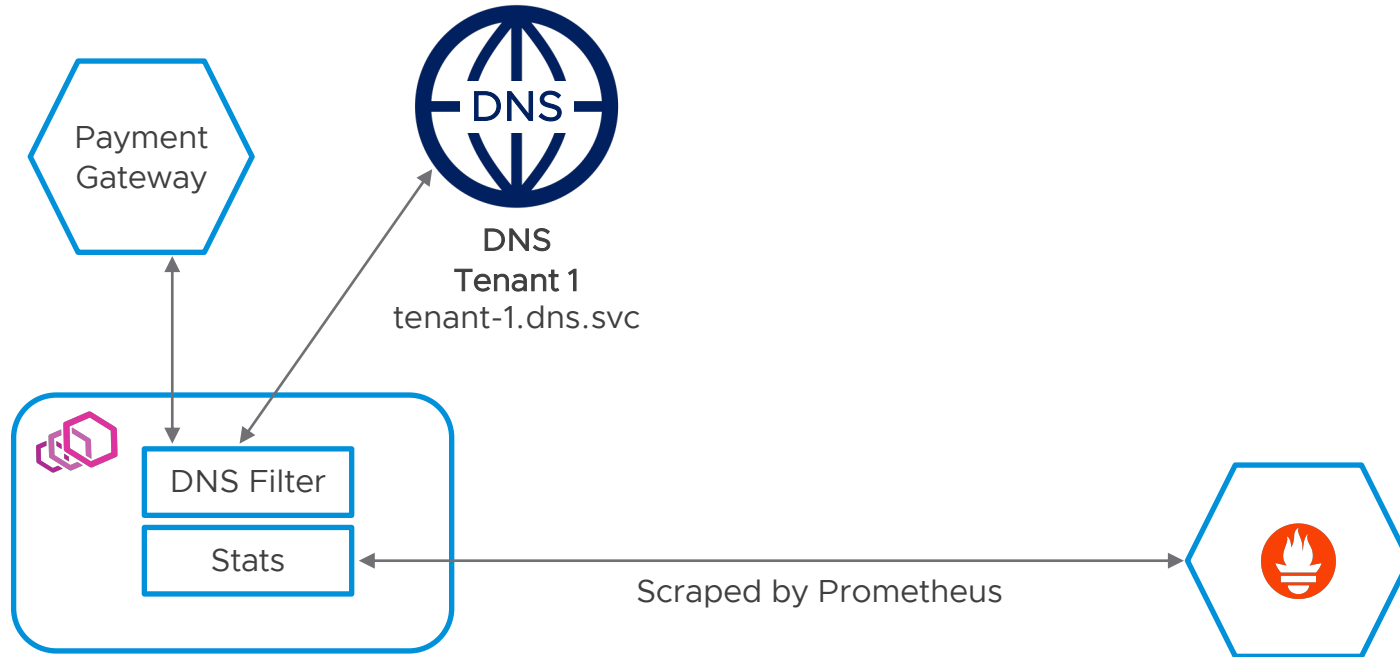




# DNS Observability and Security

```
1  apiVersion: monitoring.coreos.com/v1
2  kind: ServiceMonitor
3  metadata:
4    name: tenant-monitor
5  spec:
6    selector:
7      matchLabels:
8        app: payment-gateway # The service label.
9    namespaceSelector:
10     matchNames:
11       - payments # The service namespace.
12    endpoints:
13     - targetPort: 8000 # The Envoy stats port.
14       path: /stats/prometheus # The Envoy stats endpoint.
```

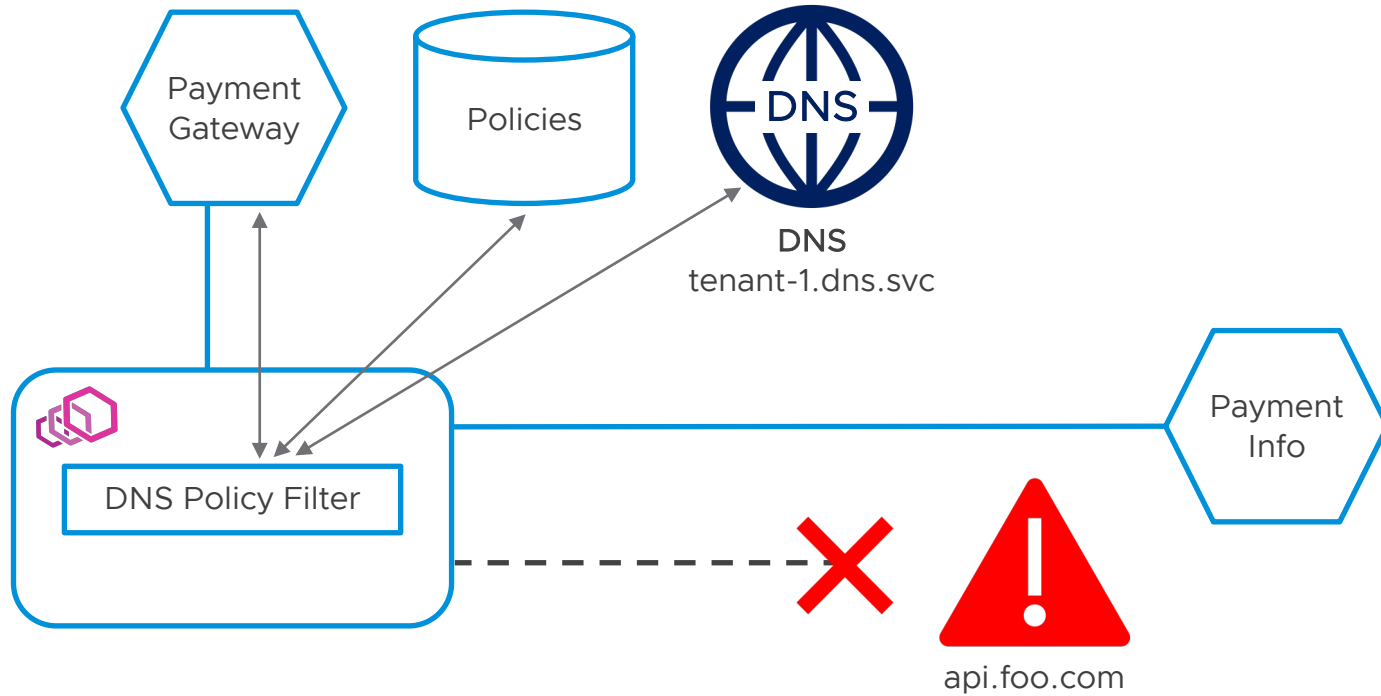
# DNS Observability and Security



# DNS Filtering

```
1  apiVersion: "networking.example.com/v1alpha1"
2  kind: DnsPolicy
3  metadata:
4    name: dns-policy
5  spec:
6    namespaceSelector:
7      matchLabels:
8        workload: payments
9    server:
10     address: tenant-1.dns.svc
11  policy:
12    whitelist:
13     - *.payments.svc.cluster.local # Can list local services.
14    blacklist:
15     - *.foo.com # Can list external domains.
16    defaultAction: WARN # Or DENY or ALLOW.
```

# DNS Filtering



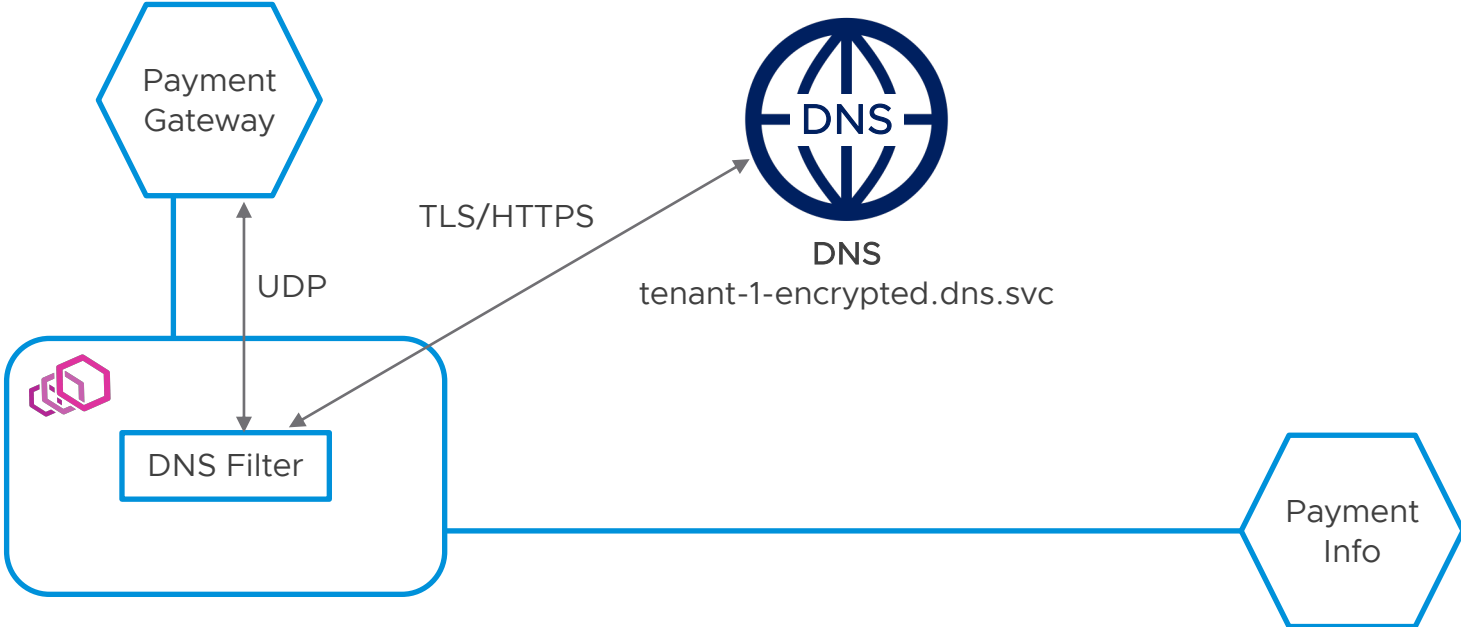
# DNS Filtering

```
5  spec:
6    namespaceSelector:
7      matchLabels:
8        workload: payments
9    selector:
10   matchLabels:
11     app: payment-gateway # Matcher for app level configuration.
12   server:
13     address: tenant-1.dns.svc # The DNS server.
14   policy:
15     defaultAction: DENY
16     server:
17       address: policy-server.acme.com # The DNS policy server.
18       protocol: grpc # The DNS policy server protocol.
```

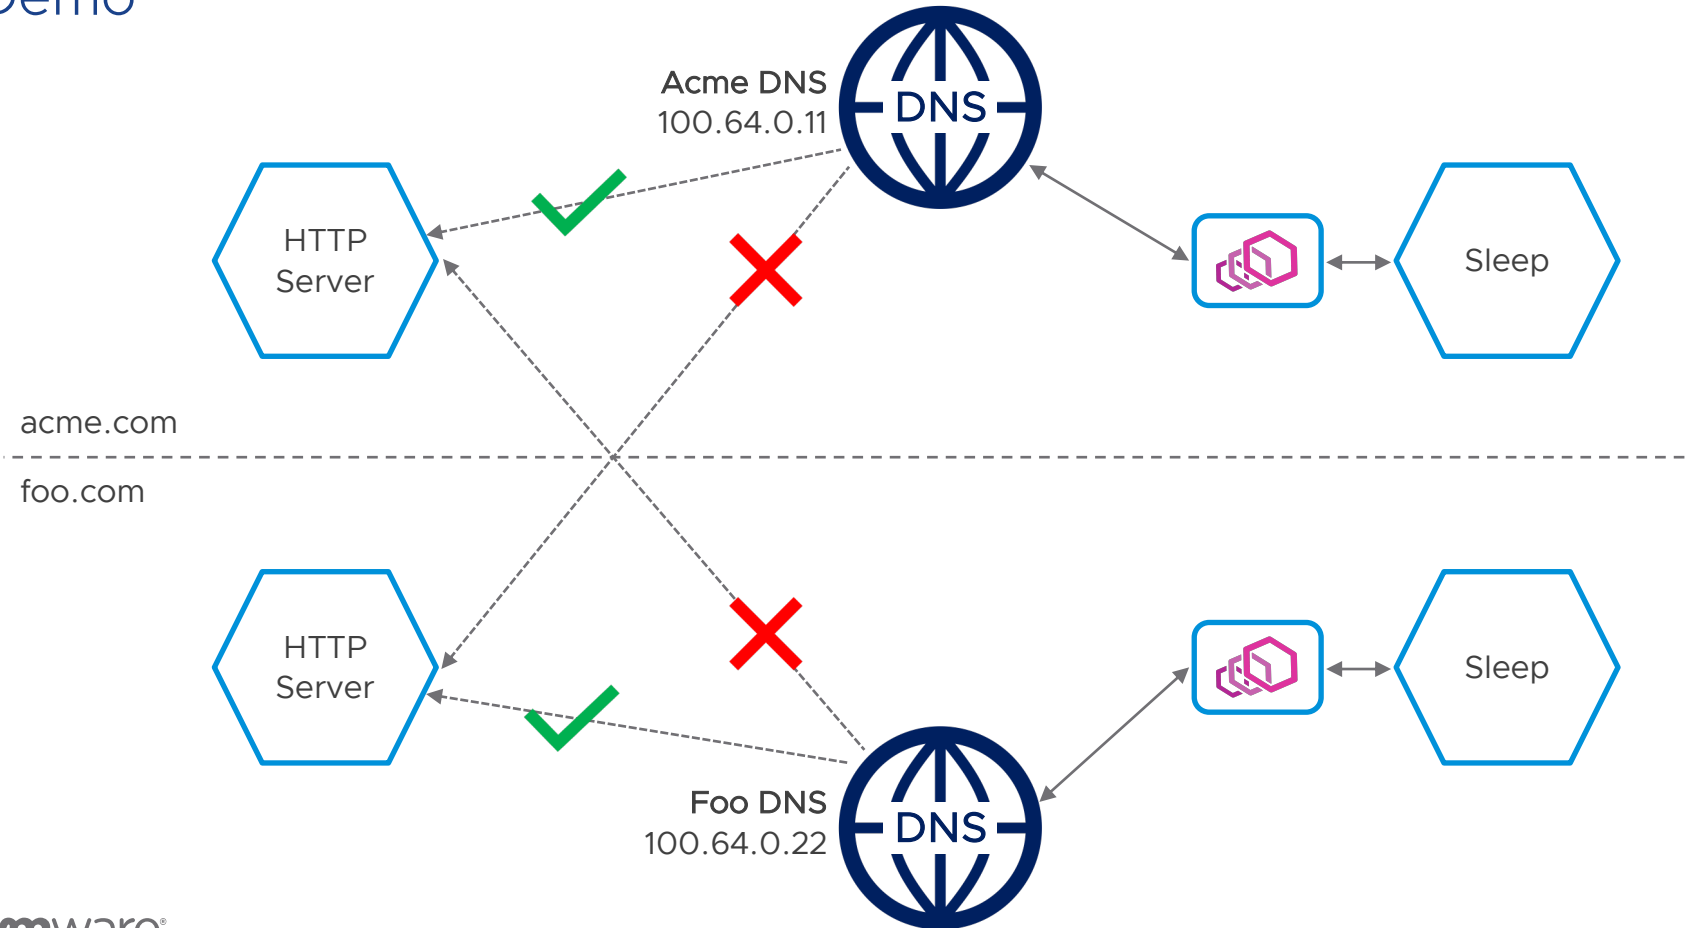
# DNS Evolution

```
1  apiVersion: "networking.example.com/v1alpha1"
2  kind: DnsPolicy
3  metadata:
4    name: dns-policy
5  spec:
6    namespaceSelector:
7      matchLabels:
8        workload: payments
9    server:
10     # A DNS server that supports encryption.
11     address: tenant-1-encrypted.dns.svc
12    protocol:
13     upgrade: true # Upgrade from cleartext to HTTPS or TLS.
14     type: dns-over-https # or dns-over-tls
```

# DNS Evolution



# Demo





```
$ kubectl get pods -n acme
```

NAME	READY	STATUS	RESTARTS	AGE
coredns-77c65cbbc5-bhwzz	1/1	Running	0	5h34m
httpbin-5fc7cf895d-j2gll	1/1	Running	0	5h34m
sleep-5ffdbd896d-hfmb9	2/2	Running	0	5h34m

```
$
```

```
$
```

```
$ kubectl get pods -n foo
```

NAME	READY	STATUS	RESTARTS	AGE
coredns-77c65cbbc5-t4zn4	1/1	Running	0	5h34m
httpbin-5fc7cf895d-x8t2s	1/1	Running	0	5h34m
sleep-5ffdbd896d-6rxc8	2/2	Running	0	5h34m

```
$ █
```

---

[acme:sleep] --> [acme:httpbin]


```
kubectl exec sleep-5ffdbd896d-hfmb9 \  
    -n acme \  
    -c sleep -- nslookup httpbin.acme.com
```

```
nslookup: can't resolve '(null)': Name does not resolve  
Name:      httpbin.acme.com  
Address 1: 100.66.185.144  
Address 2: 100.66.185.144
```

---

[foo:sleep] --> [foo:httpbin]

```
kubectl exec sleep-5ffdbd896d-6rxc8 \  
    -n foo \  
    -c sleep -- nslookup httpbin.foo.com  
nslookup: can't resolve '(null)': Name does not resolve
```



Address 1: 100.66.185.144

Address 2: 100.66.185.144

---

[foo:sleep] --> [foo:httpbin]

```
kubectl exec sleep-5ffdbd896d-6rxc8 \  
    -n foo \  
    -c sleep -- nslookup httpbin.foo.com  
nslookup: can't resolve '(null)': Name does not resolve
```

Name: httpbin.foo.com


Address 1: 100.68.169.65

Address 2: 100.68.169.65

---

[acme:sleep] --> [foo:httpbin]

```
kubectl exec sleep-5ffdbd896d-hfmb9 \  
    -n acme \  
    -c sleep -- nslookup httpbin.foo.com
```



Address 1: 100.68.169.65

Address 2: 100.68.169.65

---

[acme:sleep] --> [foo:httpbin]


```
kubectl exec sleep-5ffdbd896d-hfmb9 \  
    -n acme \  
    -c sleep -- nslookup httpbin.foo.com  
nslookup: can't resolve '(null)': Name does not resolve
```

```
nslookup: can't resolve 'httpbin.foo.com': Try again  
command terminated with exit code 1  
make: [demo] Error 1 (ignored)
```

---

[foo:sleep] --> [acme:httpbin]

```
kubectl exec sleep-5ffdbd896d-6rxc8 \  
    -n foo \  
    -c sleep -- nslookup httpbin.acme.com
```



```
        -n acme \  
        -c sleep -- nslookup httpbin.foo.com  
nslookup: can't resolve '(null)': Name does not resolve
```

```
nslookup: can't resolve 'httpbin.foo.com': Try again  
command terminated with exit code 1  
make: [demo] Error 1 (ignored)
```

---

```
[foo:sleep] --> [acme:httpbin]
```

```
kubectl exec sleep-5ffdbd896d-6rxc8 \  
        -n foo \  
        -c sleep -- nslookup httpbin.acme.com  
nslookup: can't resolve '(null)': Name does not resolve
```

```
nslookup: can't resolve 'httpbin.acme.com': Try again  
command terminated with exit code 1  
make: [demo] Error 1 (ignored)
```

---

```
$ █
```

# Summary

- DNS plays a key role for service discovery and application migration
- Multi-tenancy at the DNS layer is very critical for enterprise systems
- Envoy proxy can solve some interesting challenges with DNS
- Envoy proxy's xDS APIs let us dynamically configure DNS filters
- The DNS filters can also be integrated with third-party systems to provide richer observability, security, and filtering
- Next: Contribute the work to existing open source projects!



# Thank You

 @deepakalani | Deepa Kalani

 @venilnoronha | Venil Noronha

[istio.io](https://istio.io)

[envoyproxy.io](https://envoyproxy.io)

[venilnoronha.io](https://venilnoronha.io)