

ISOVALENT


Introduction to Cilium from Telco and On-premise Perspective

Speaker: Yutaro Hayakawa



Self Introduction

- Yutaro Hayakawa (@YutaroHayakawa)
- Software Engineer at Isovalent
- Dataplane, enterprise networking, etc...
- ex-LINE Verda Network Development Team



Isovalent


- The vendor company provides networking, security and observability solutions for Kubernetes
- The company behind Cilium, Hubble, Tetragon, and more...
- Puts **eBPF** as a company's key technology



ISOVALENT



cilium




HUBBLE



Tetragon

Trivia: Isovalent and SDN



**Thomas
Graf**

Co-founder & CTO

- ex-RedHat, Cisco
- Long-term Linux kernel development leader
- **OVS core contributor**



**Dan
Wendlandt**

Co-founder & CEO

- ex-Nicira, VMWare
- The first PM of Nicira NVP and OVS
- **PM of VMWare NSX**



**Martin
Casado**

Creator of SDN

- Cofounder & CTO of Nicira
- General partner at a16z
- **Serves on the board of Isovalent from a16z**

eBPF

- A way to implement Linux kernel extension
- Safety guarantee by static verification
- Backward compatibility guarantee
 - No “upstream to avoid breakage”
 - No mainline ⇒ distribution gap period
 - Can do anything specific to your business inside the kernel
- **Dramatically speed-ups kernel extension programming/delivery/maintenance cycle**



<https://ebpf.io/>

Originally, created by...



ISOVALENT

Cilium

- CNI, Load Balancer, Firewall, Network Observability, Multi-cluster Networking, Service Mesh, and more...
- Uses eBPF for DPlane implementation
 - Easy to extend, rapid development cycle



cilium

<https://cilium.io>

Cilium from On-prem and Telco Perspective

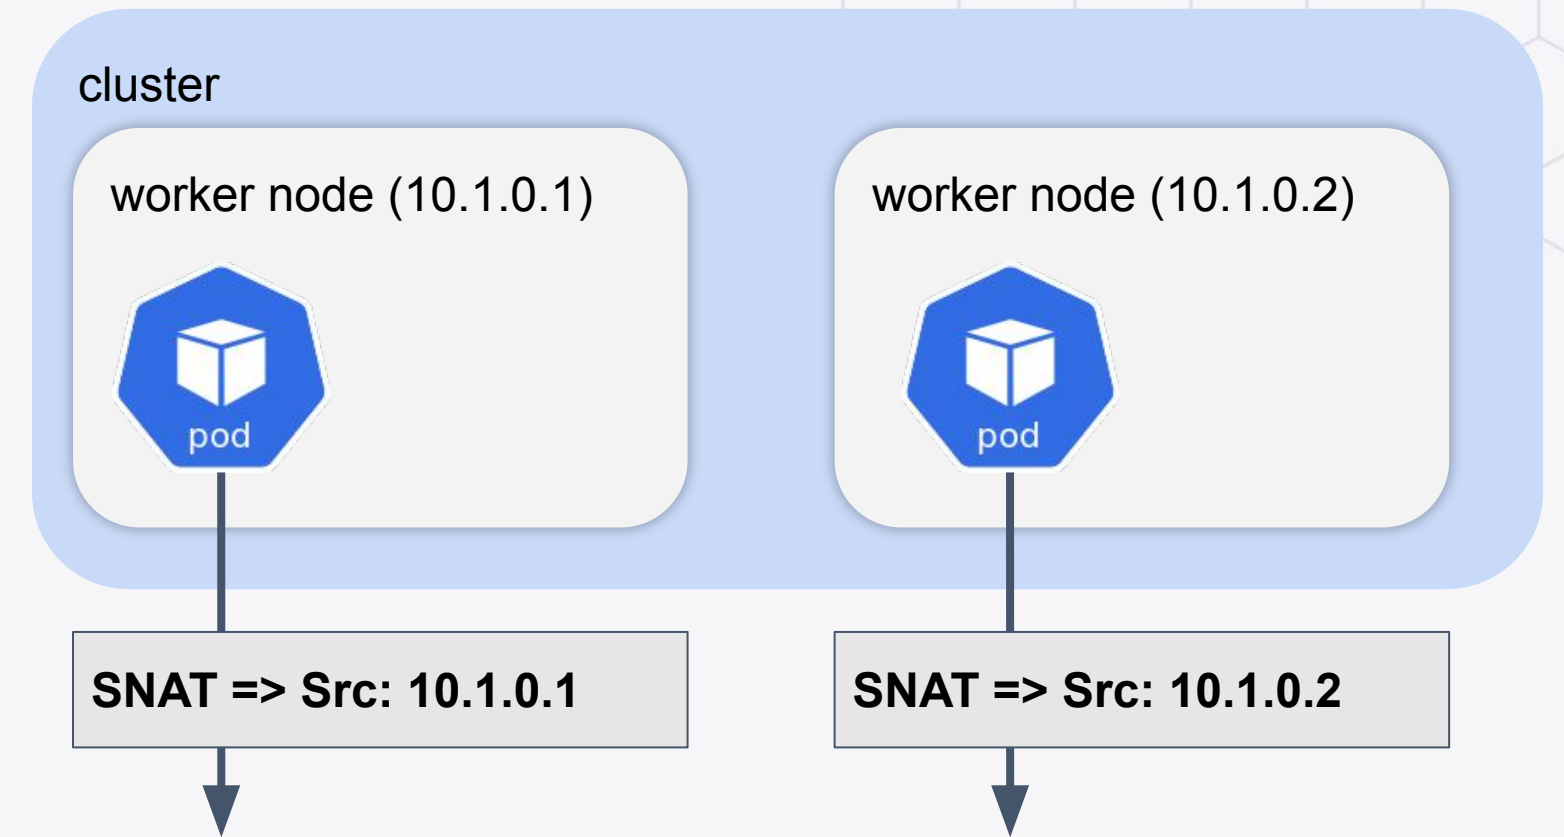


- Traditional Infrastructure ⇒ Cloud native infrastructure transition
 - Static IP vs Dynamic IP
 - Egress Gateway, FQDN-based Network Policy
- Integrating Cilium into DC network
 - How to make k8s network accessible from DC network?
 - BGP Integration
- Integrating Cilium into SRv6 L3 VPN
 - How to make k8s network accessible over VPN?
 - BGP Integration with SRv6 VPN

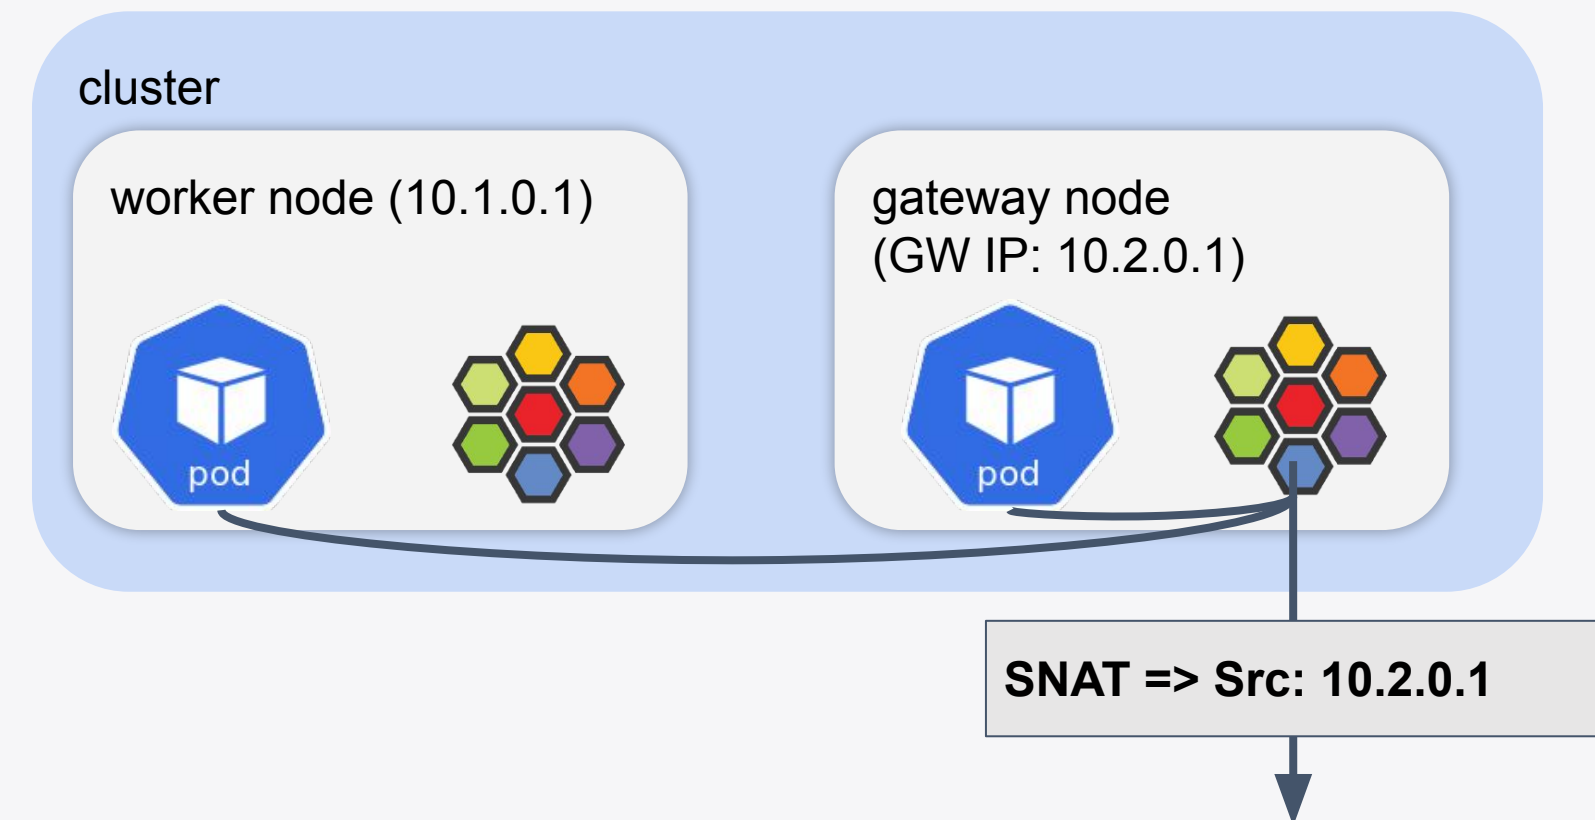
Egress Gateway

- Routes all IPv4 connections originating from pods and destined to specific cluster-external CIDRs through particular nodes
- Use fixed source IP address for egressing the cluster

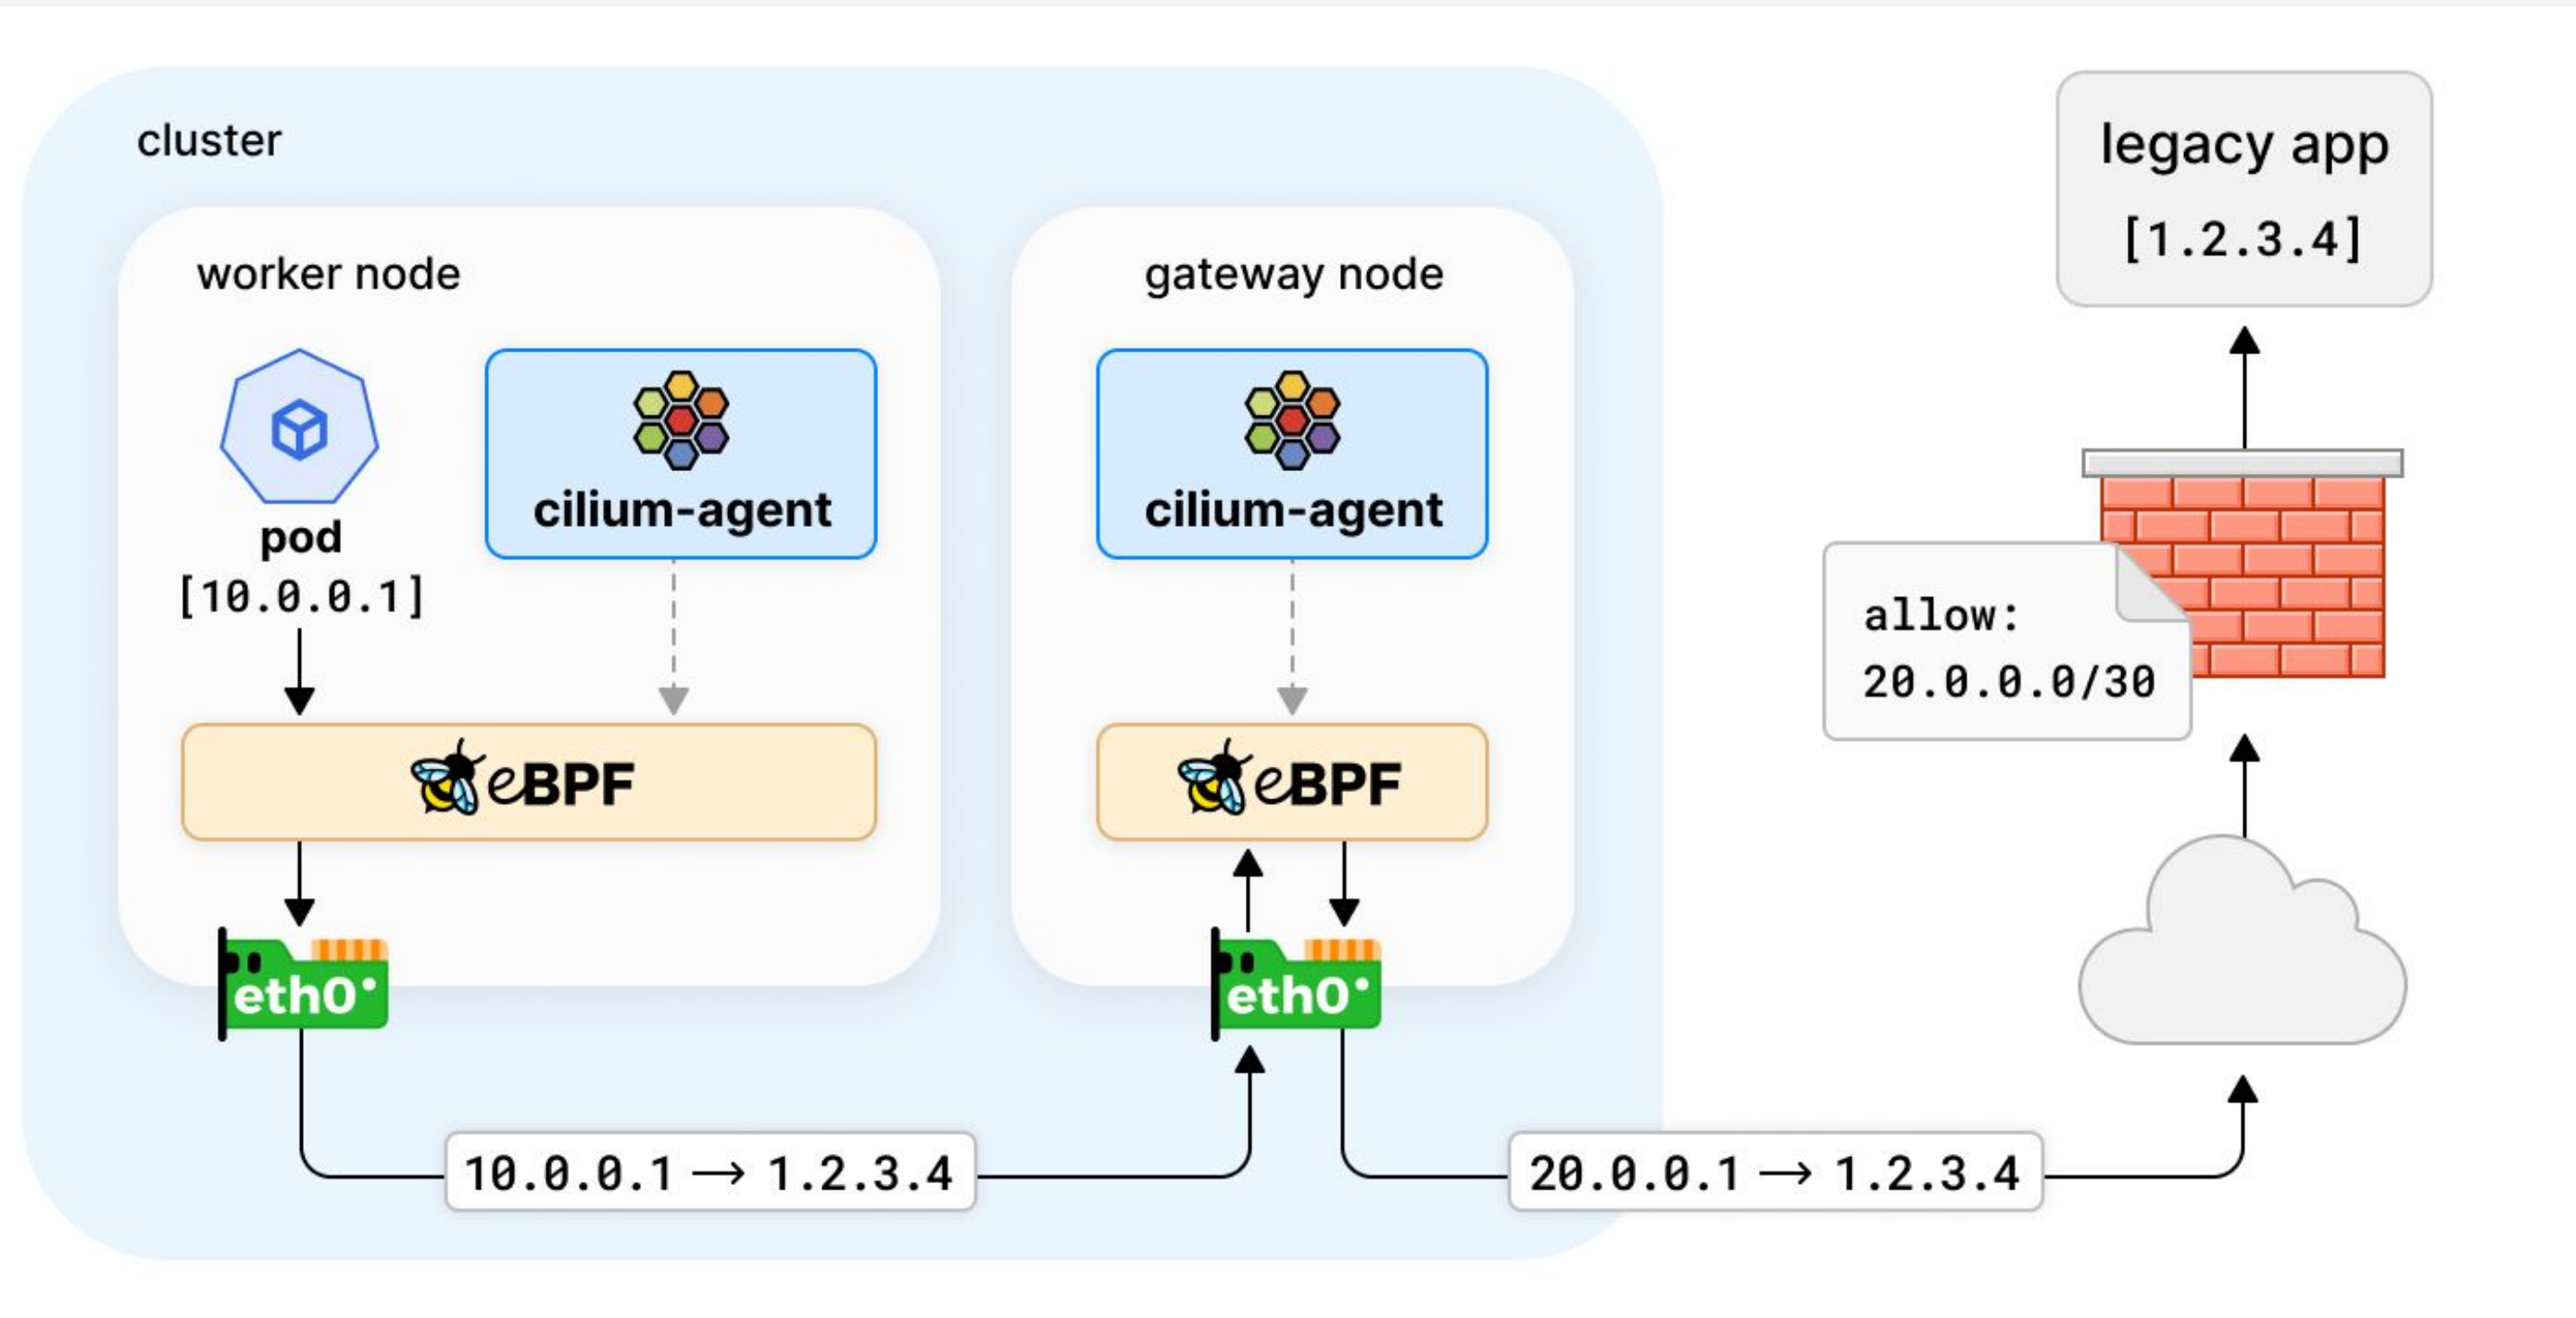
Without Egress Gateway



With Egress Gateway

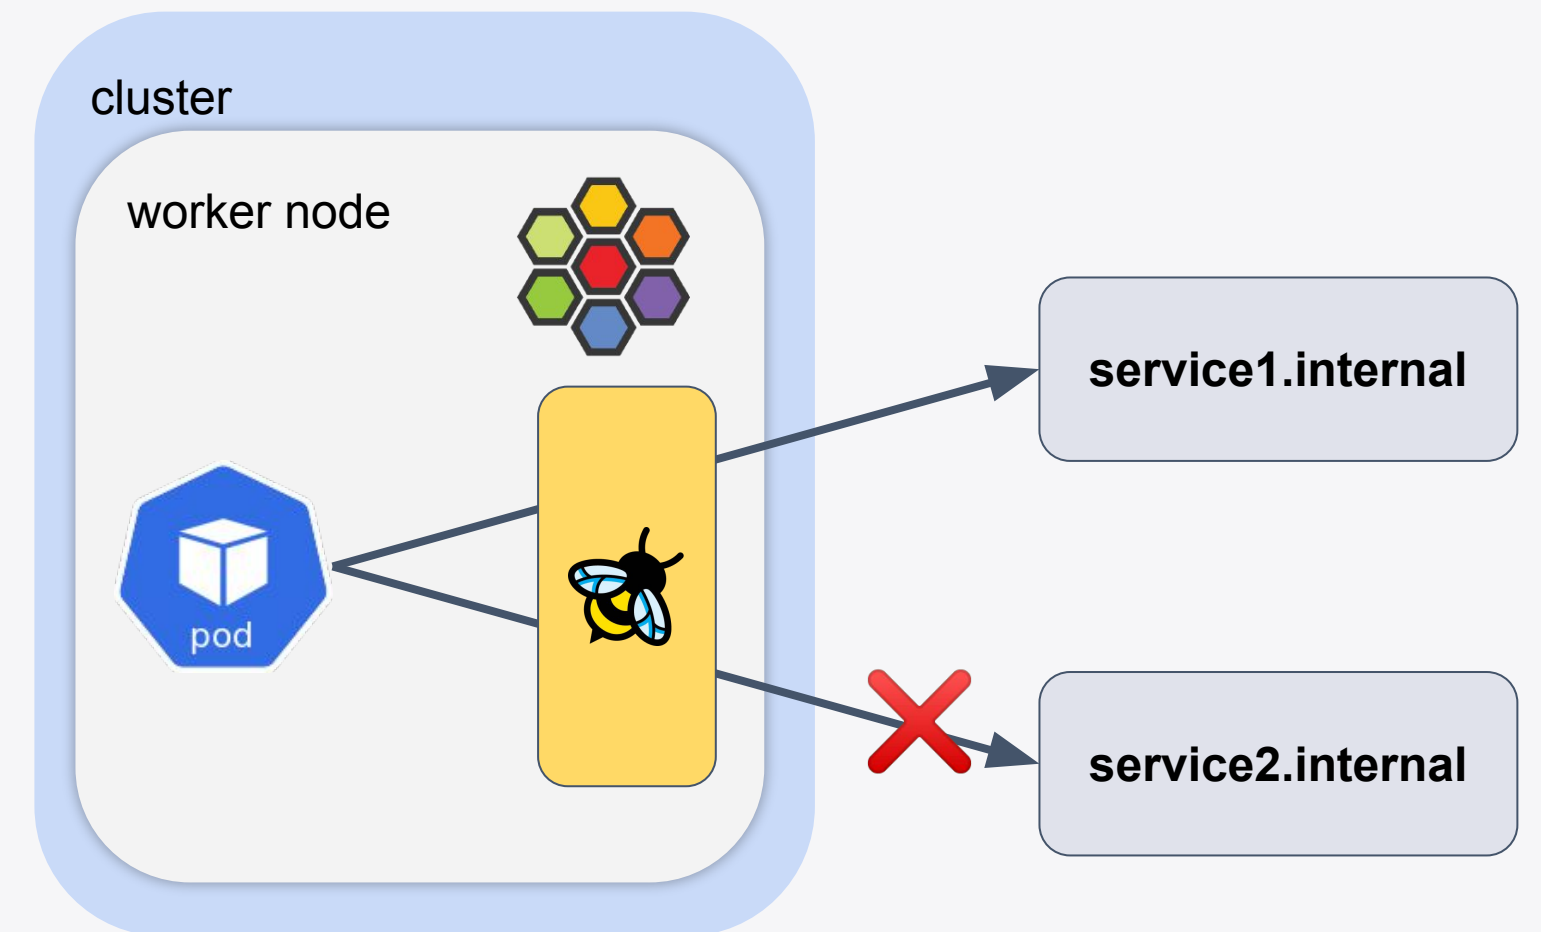


Egress Gateway: Why and How



FQDN-based Network Policy

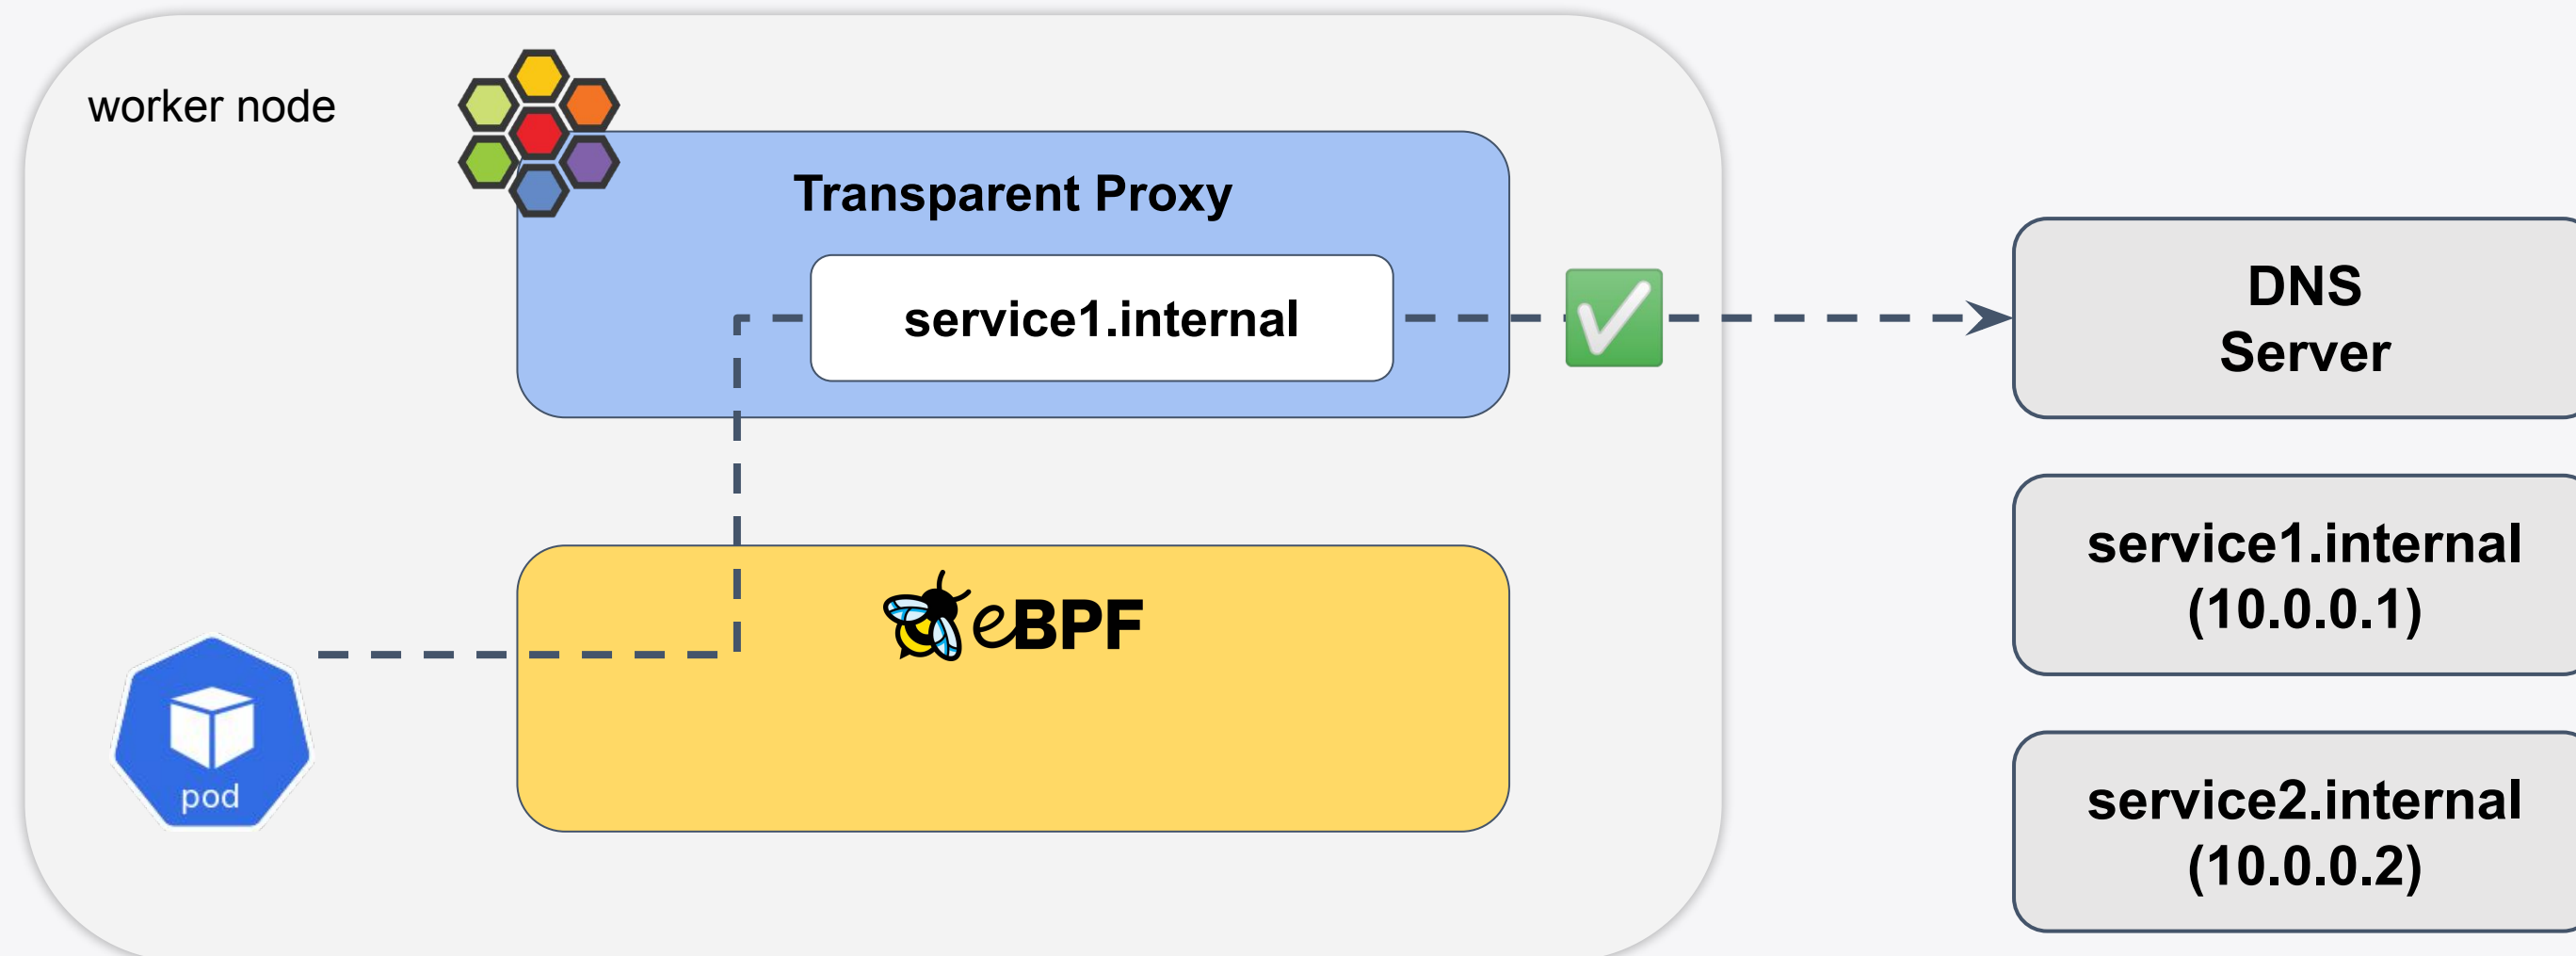
- Filtering egress traffic based on the FQDN (or FQDN pattern) instead of IP addresses



ISOVALENT

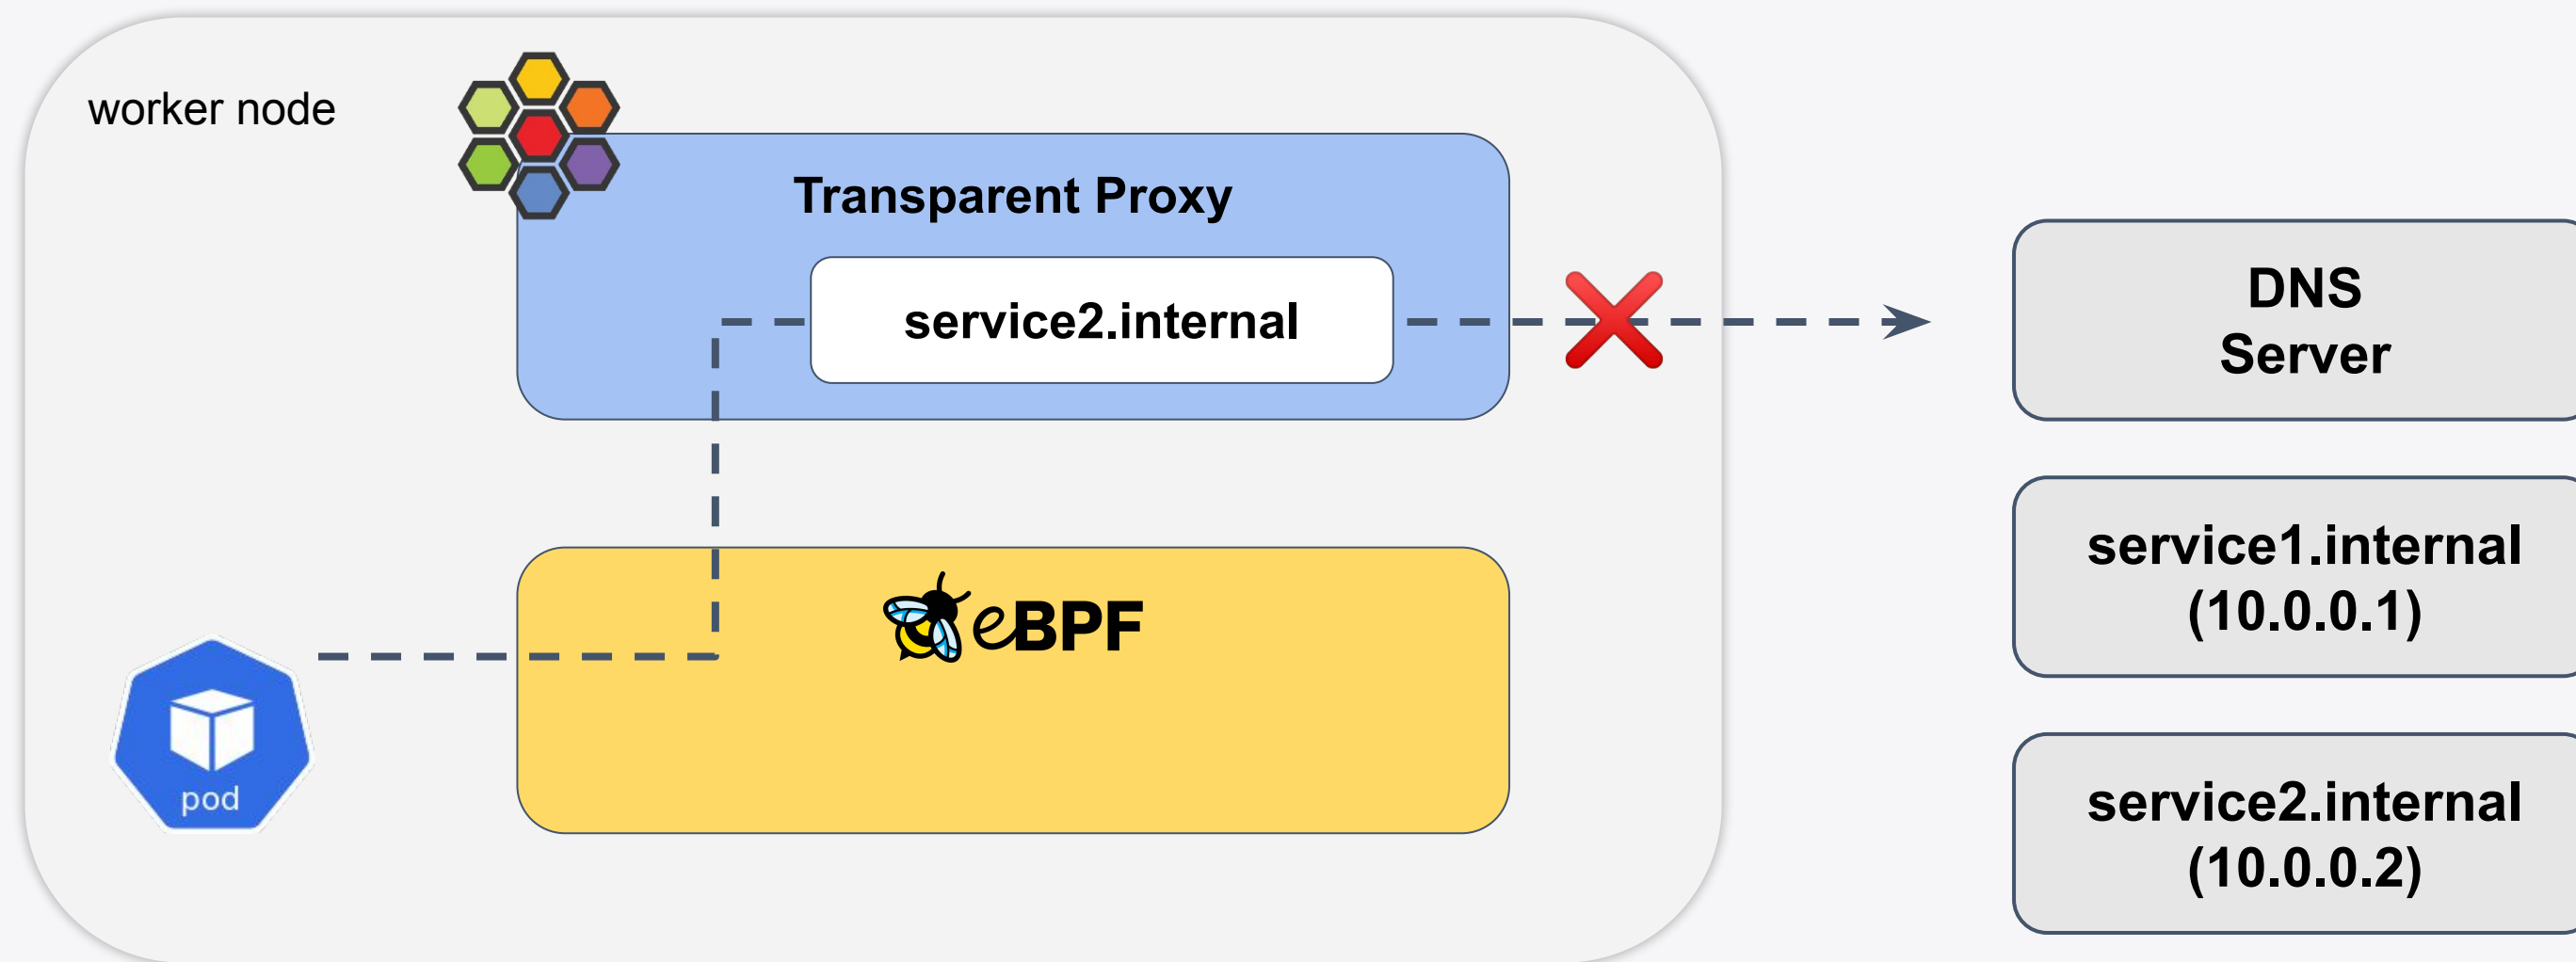
FQDN-based Network Policy Implementation: L7 DNS Blocking

- How to allow service1.internal?
- Intercept DNS query from Pod with transparent proxy
- Transparent proxy checks the query and if matches to the allow policy



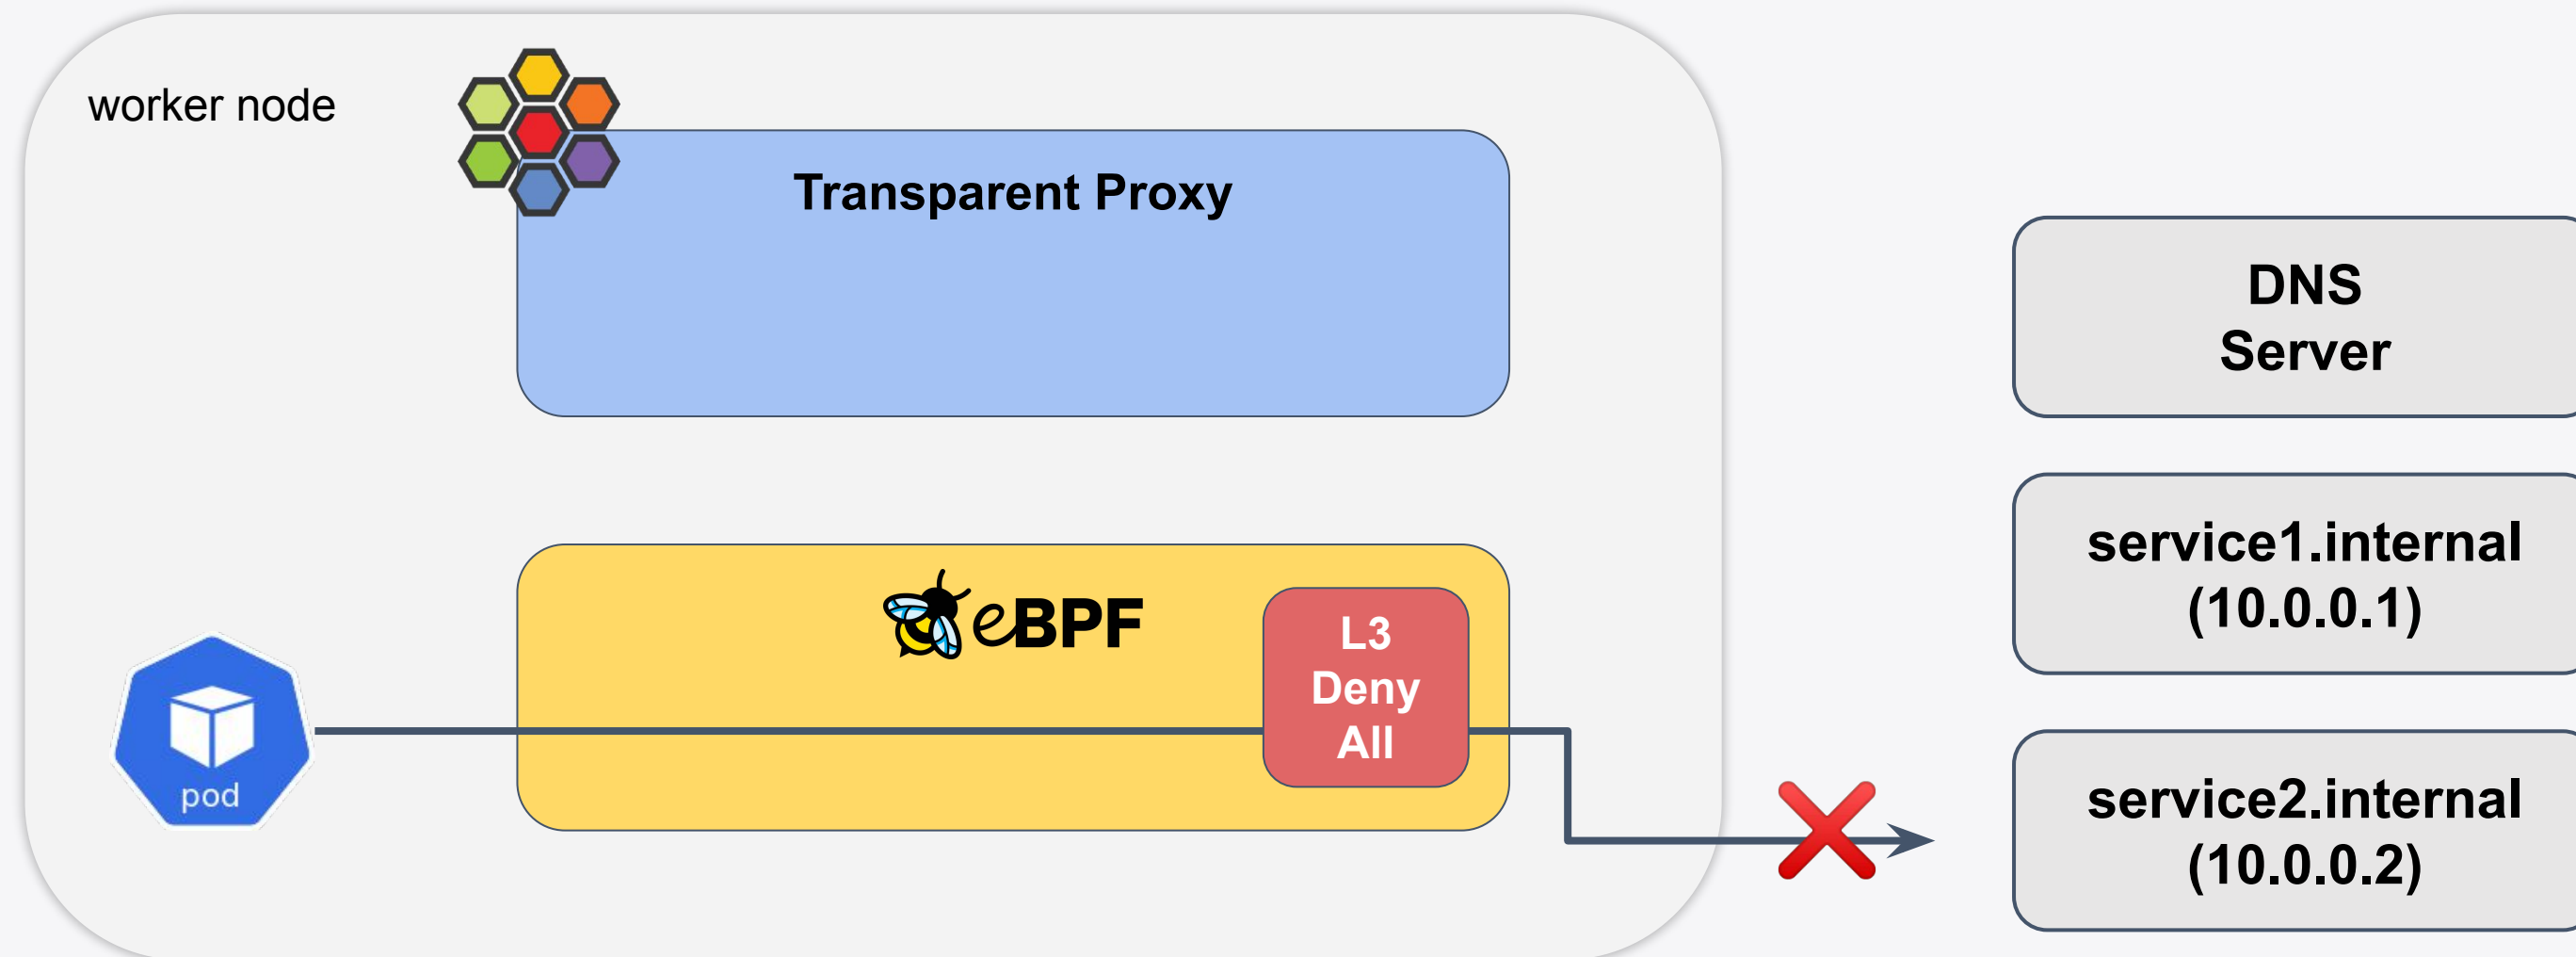
FQDN-based Network Policy Implementation: L7 DNS Blocking

- How to deny service2.internal?
- When the query didn't match to the allow policy, transparent proxy returns error



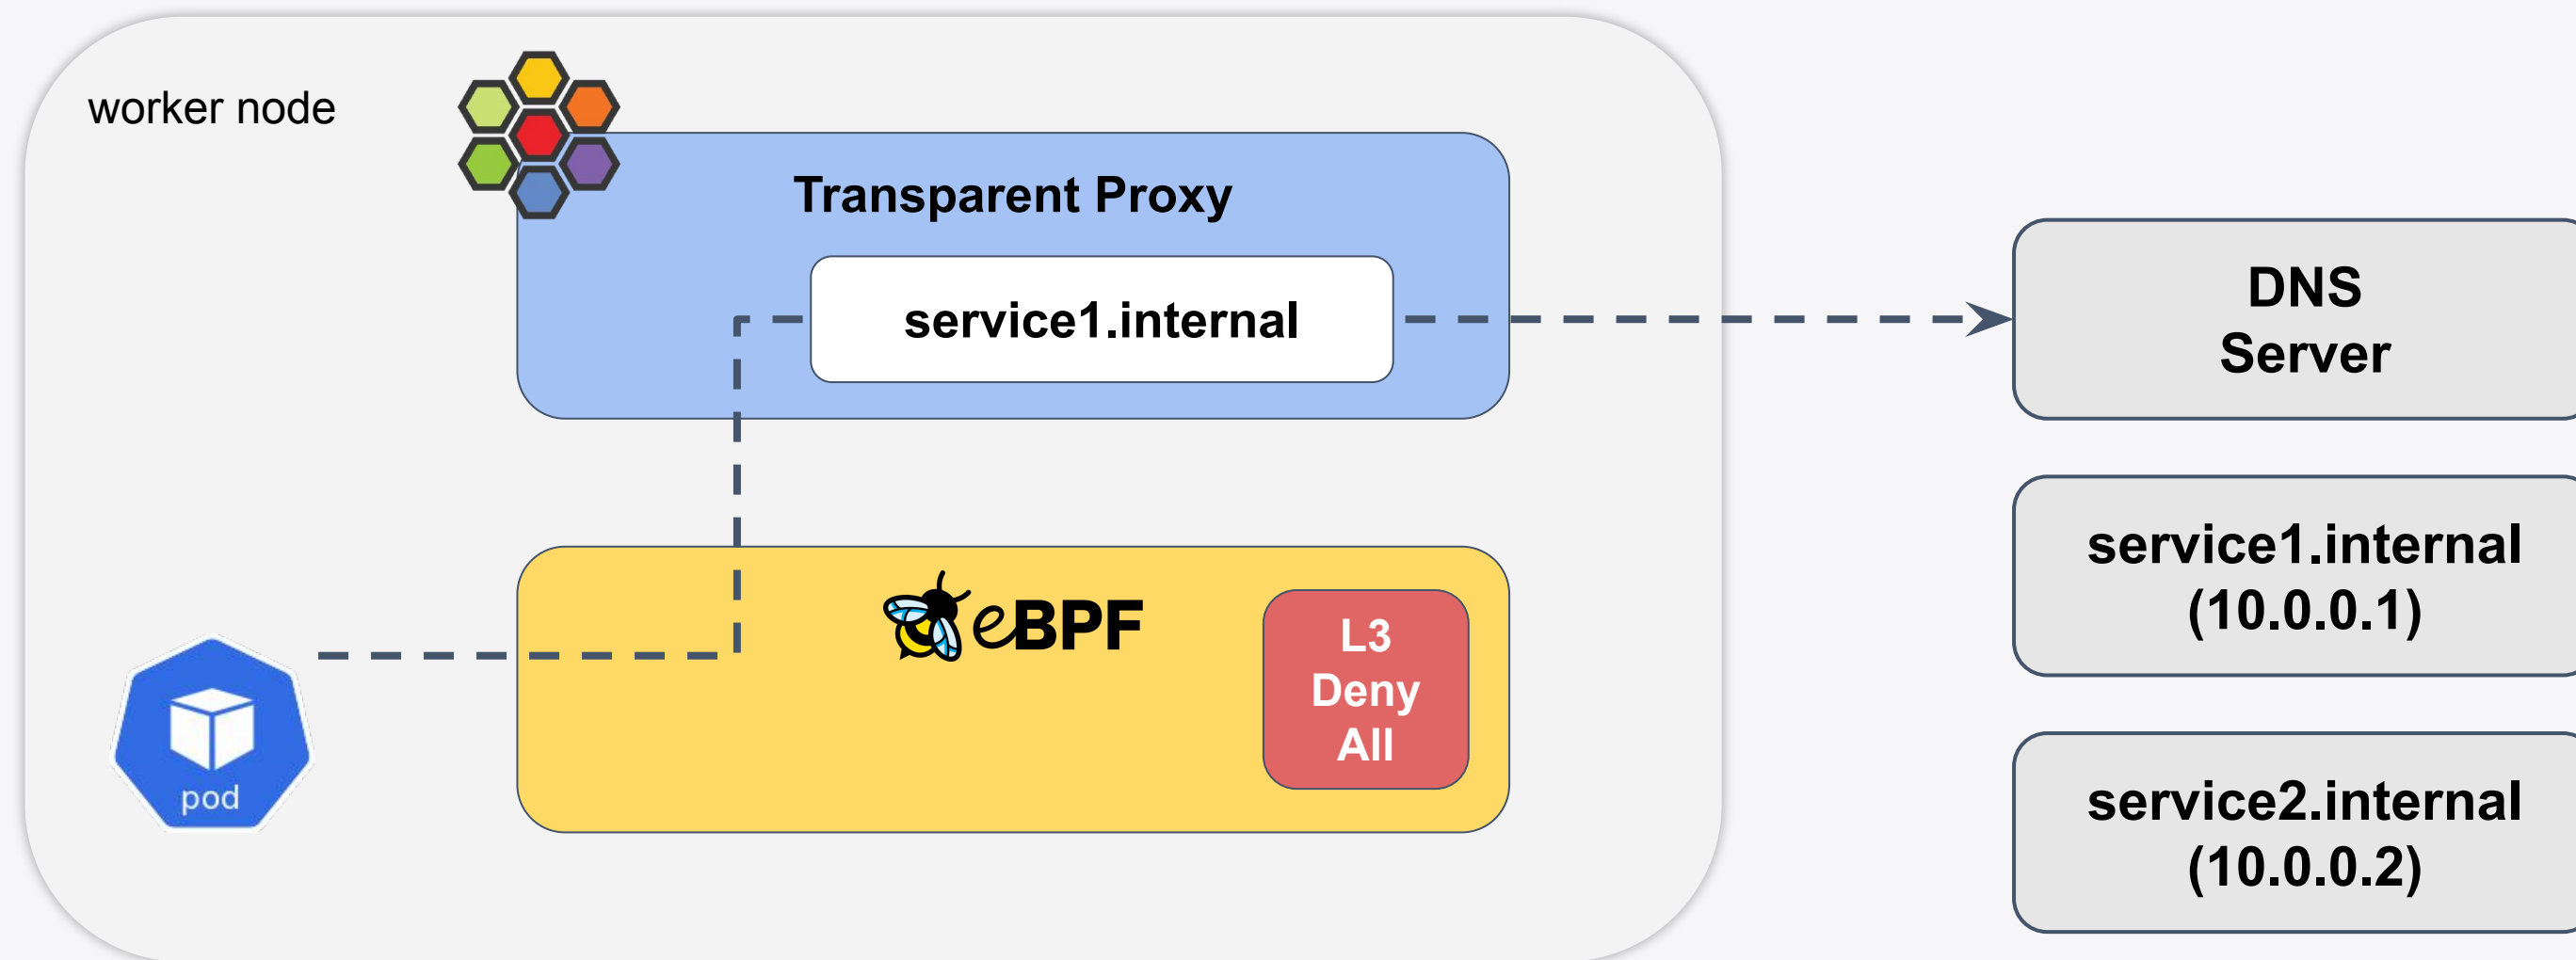
FQDN-based Network Policy: Dynamic L3 Policy Generation

- How to deny service2.internal



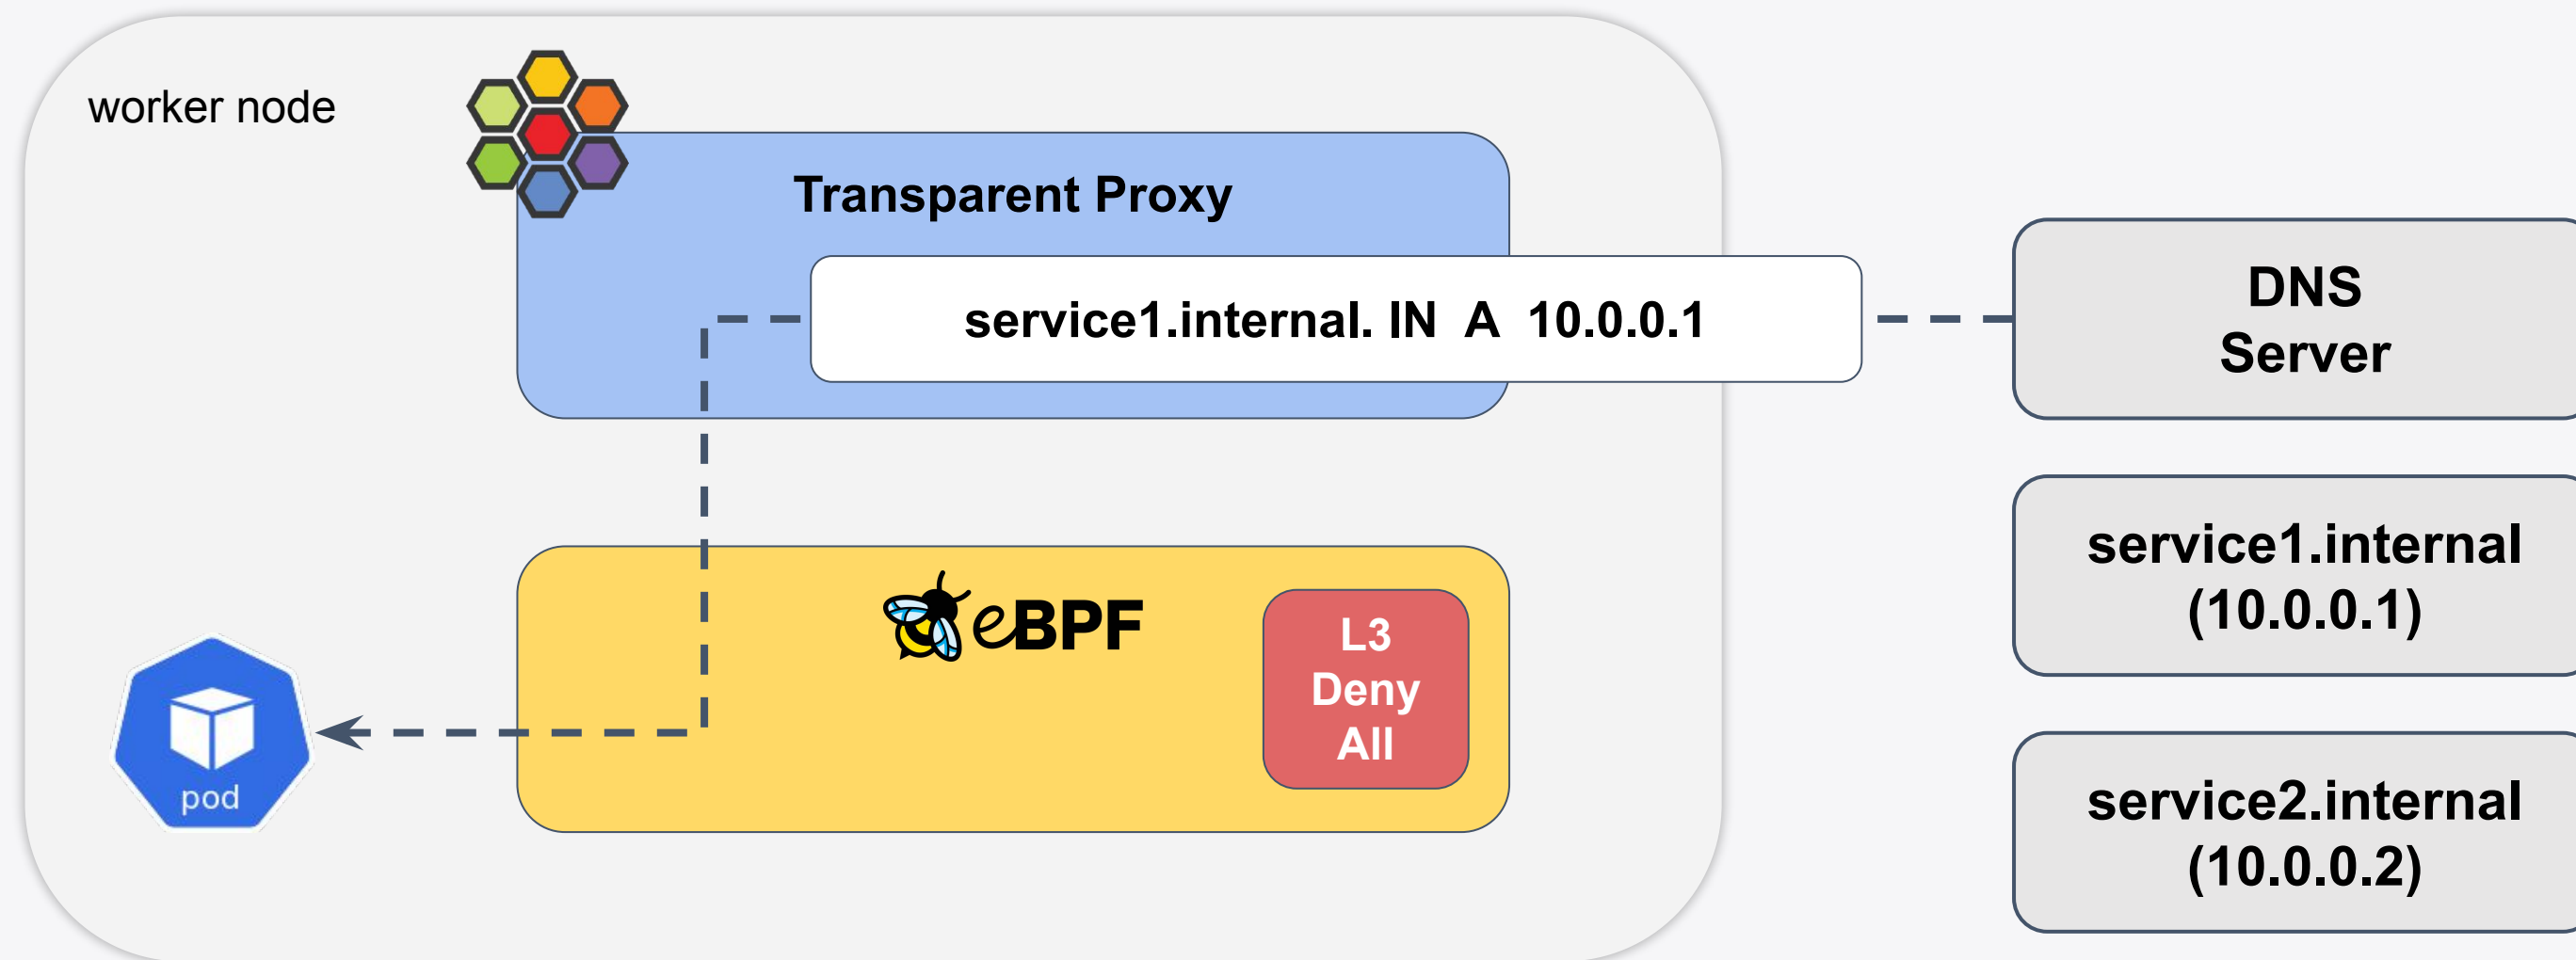
FQDN-based Network Policy: Dynamic L3 Policy Generation

- How to allow service1.internal



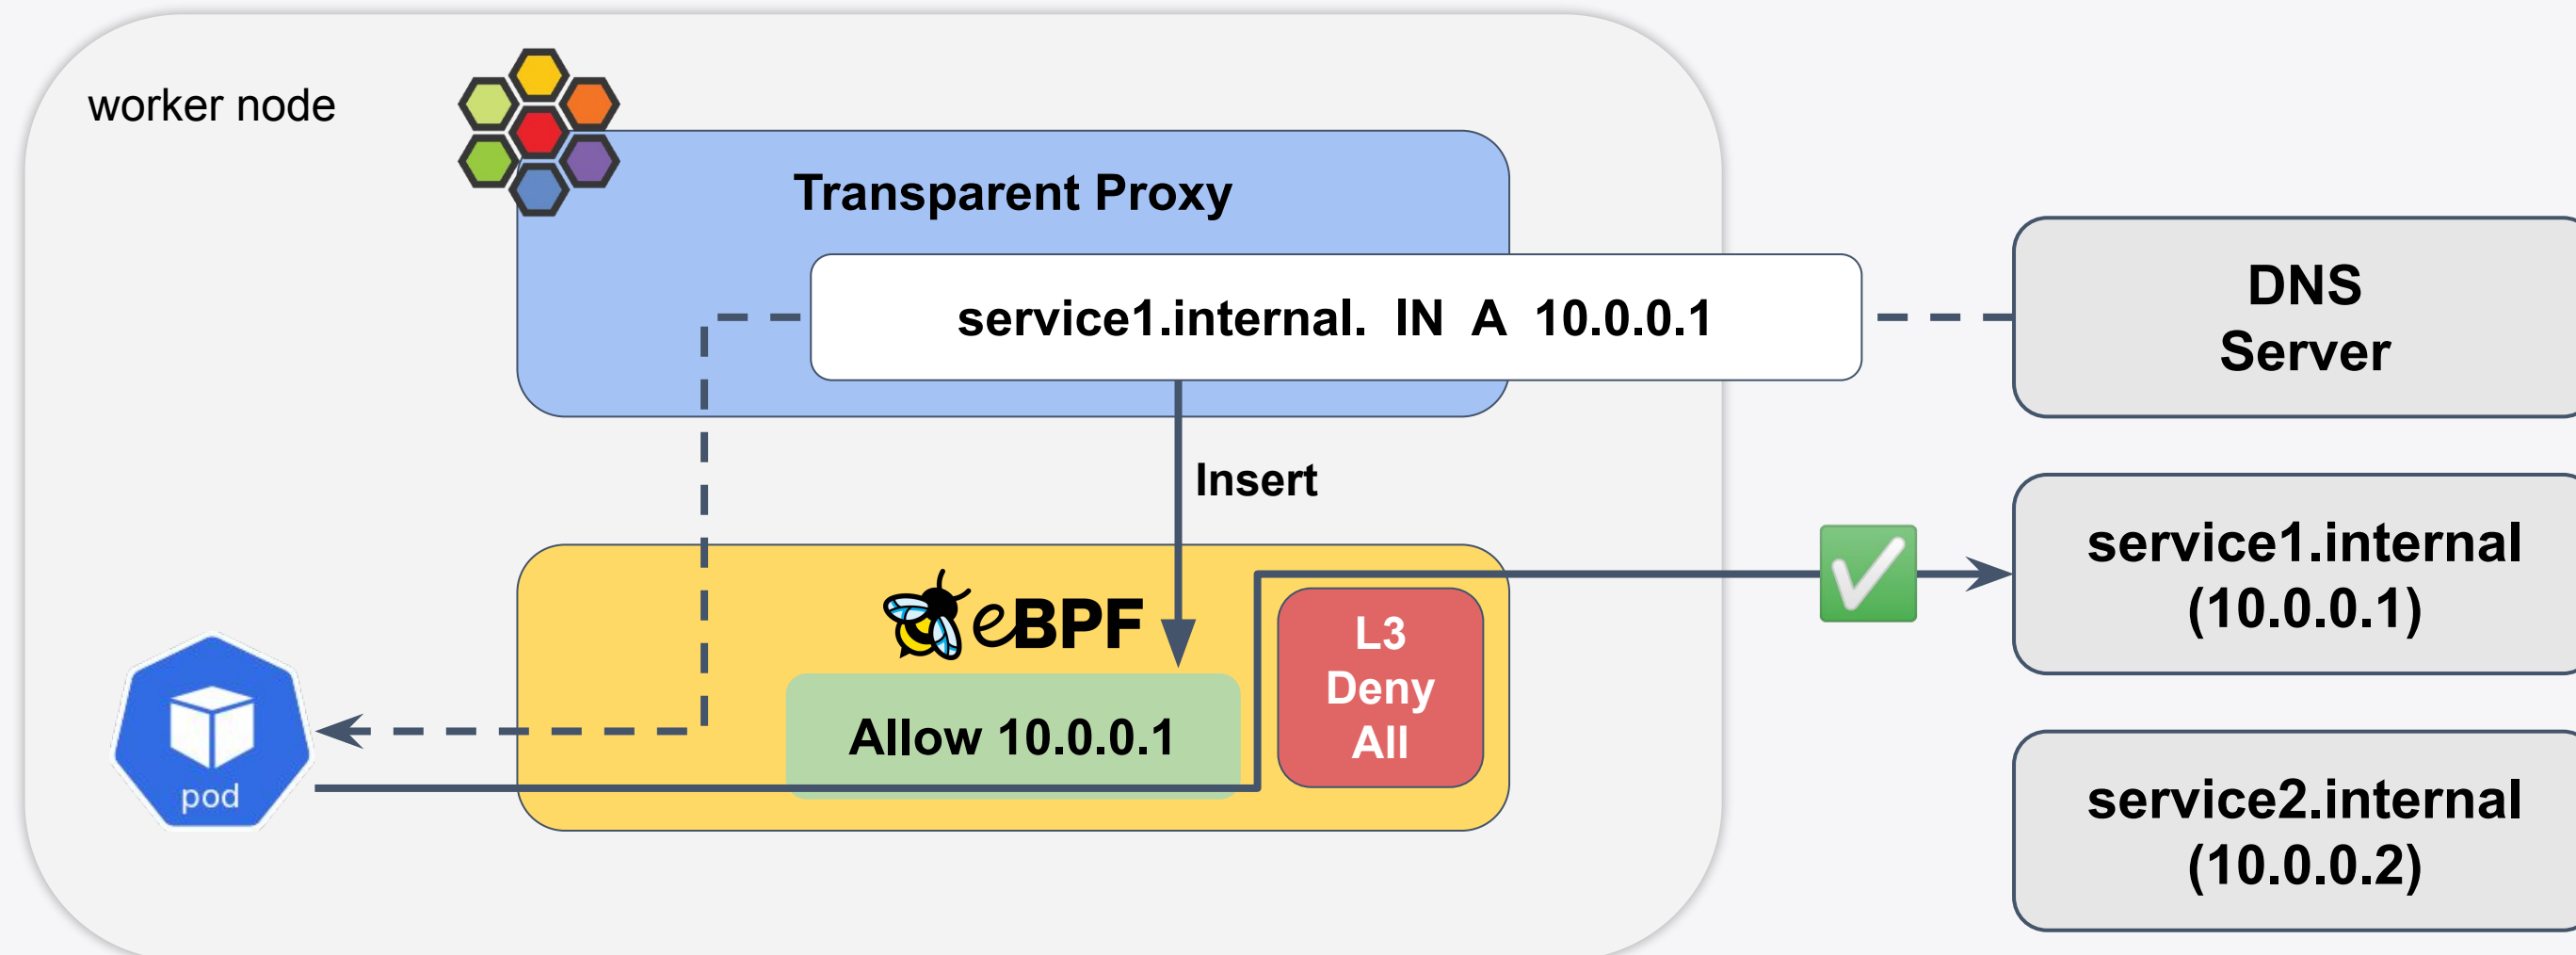
FQDN-based Network Policy: Dynamic L3 Policy Generation

- How to allow service1.internal




FQDN-based Network Policy: Dynamic L3 Policy Generation

- The transparent proxy is TTL-aware



BGP Integration

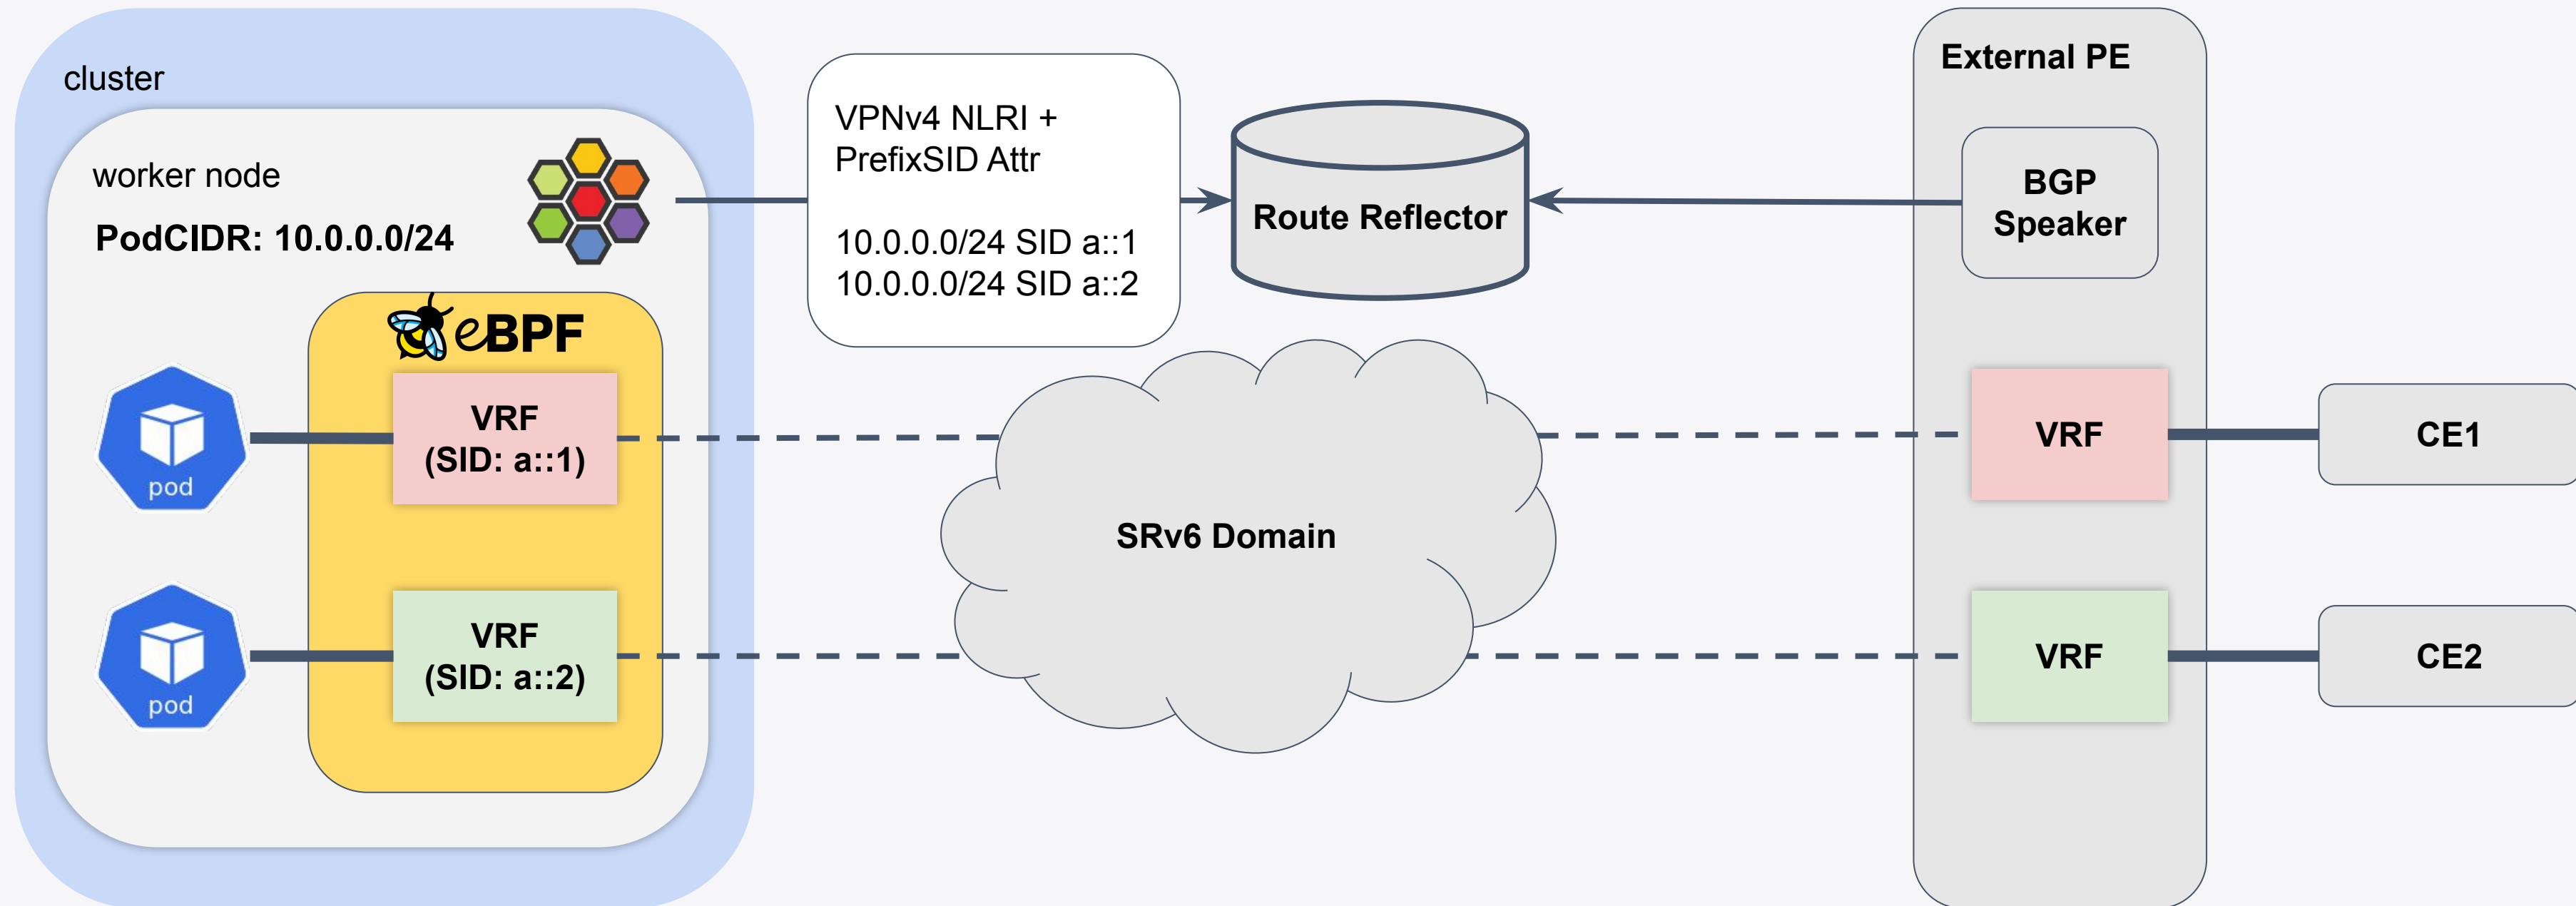
- Cilium speaks BGP to advertise k8s network prefixes (e.g. PodCIDR, Service VIP, Egress Gateway IP, etc...)



SRv6 L3 VPN DPlane + BGP Control Plane

- Cilium speaks BGP to advertise PodCIDR over SRv6 L3 VPN
- Cilium implements eBPF-based SRv6 encap/decap and VRF DPlane
- Telco use case: k8s + Cilium as an SRv6-aware MEC platform

SRv6 L3 VPN DPlane + BGP Control Plane



ISOVALENT

Last Words

- We're use case driven community, so don't hesitate to file issues even if it is niche
- We are hiring



cilium

Join the Cilium community
<https://cilium.io/get-involved/>



ISOVALENT

Join the Isovalent
<https://isovalent.com/careers/>

ISOVALENT

Thank you!

eBPF - Kubernetes Networking Use Cases

