KubeCon | CloudNativeCon

North America 2020

Virtual

Extending Service Mesh to the Edge

Stephen Wong

Agenda

- Motivation on Service Mesh to the Edge
- Challenges
- Demo
- Future Developments

Why Service Mesh to the Edge

- What edge are you talking about?
  - OPNFV Clover
    - Addressing Telco/NFV use cases with cloud native technologies
      - Particularly MEC (Multi-access edge computing) for Telco
      - Excess compute on telco edge sites (not sensors)
  - Kubernetes cluster control plane on edge
    - Allows k8s to reschedule and restart pods locally without going to cloud

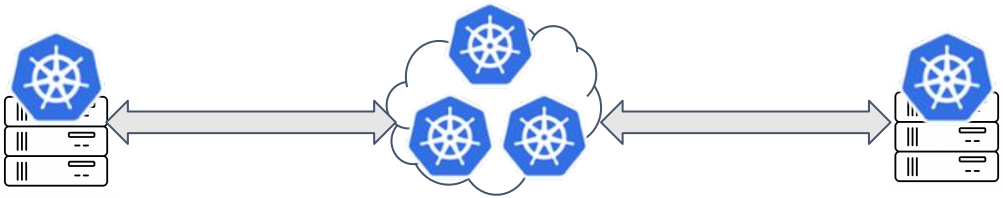

Single Mesh on Cloud and Edge Sites

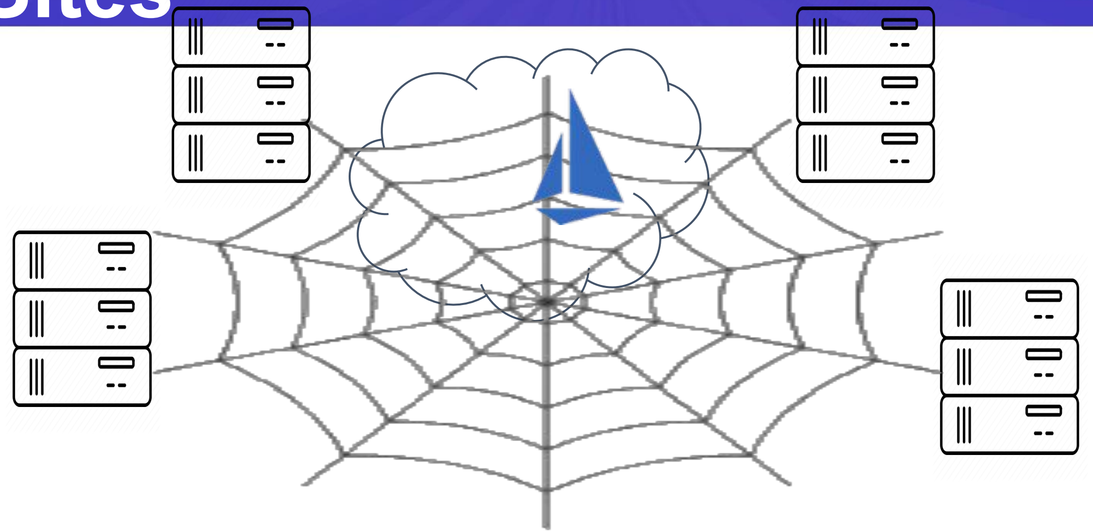

- Multiple k8s clusters, single mesh across all clusters
  - A consistent network policy and telemetry format / gathering framework across cloud and edge sites
  - Thriving ecosystems: ecosystem projects using Istio to run CI/CD, canary releases / testing...etc

Challenges

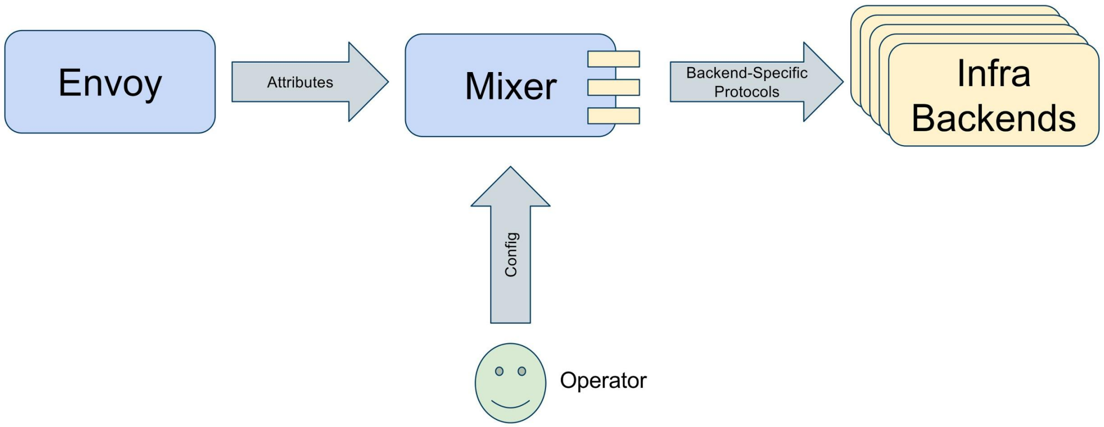

- Mixer
  - Each request needs to be forwarded to Mixer, running on cloud
- Istio community decided to deprecate mixer, opted to instead take advantage of Envoy extensibility to implement custom protocols and complicated policies

Opportunity

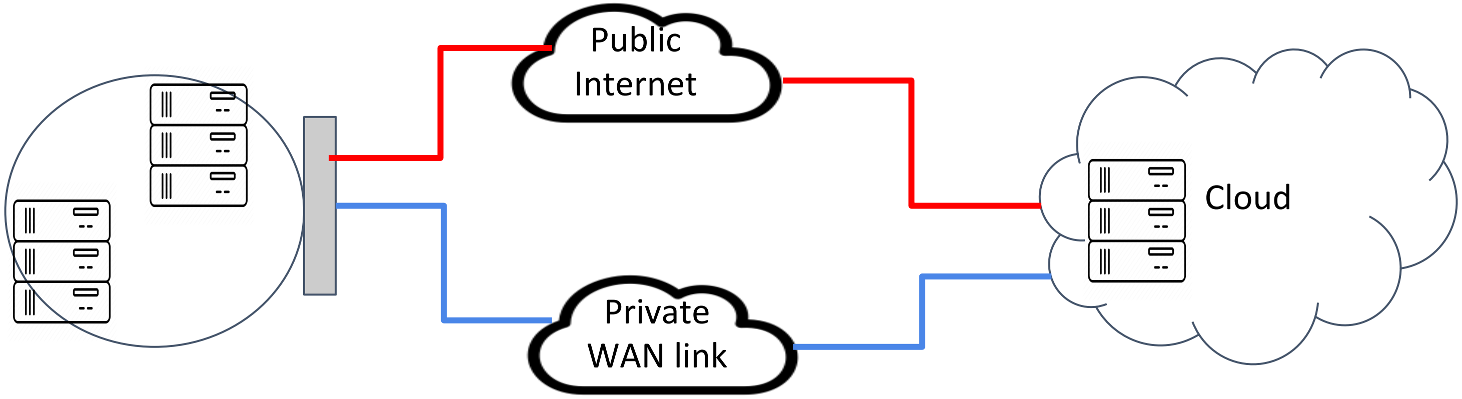

- WAN association
  - With cloud/edge hybrid applications, WAN connectivity becomes part of the communication channel between microservices
  - Service mesh related route rules and policies should influence choice of WAN connectivity

Clovisor

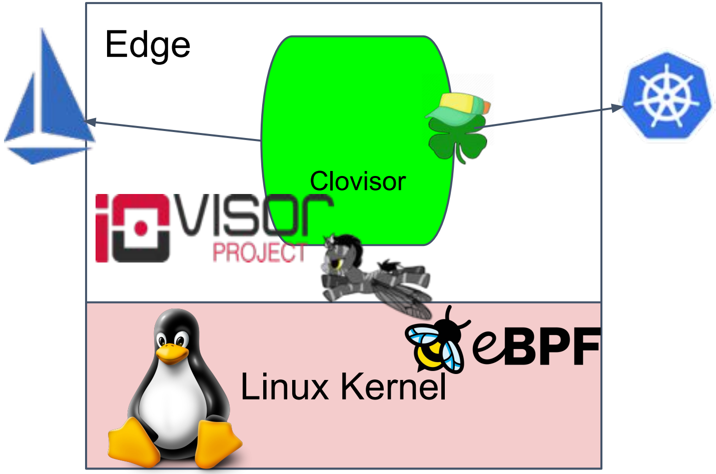

- Clovisor
  - Developed as part of OPNFV Clover, has been spurn out
  - Speaks to both k8s and Istio via go-client
  - Utilizes IOVisor project to compile / load BPF code to kernel
  - Uses BPF to perform both packet tracing and redirection

Demo (description)

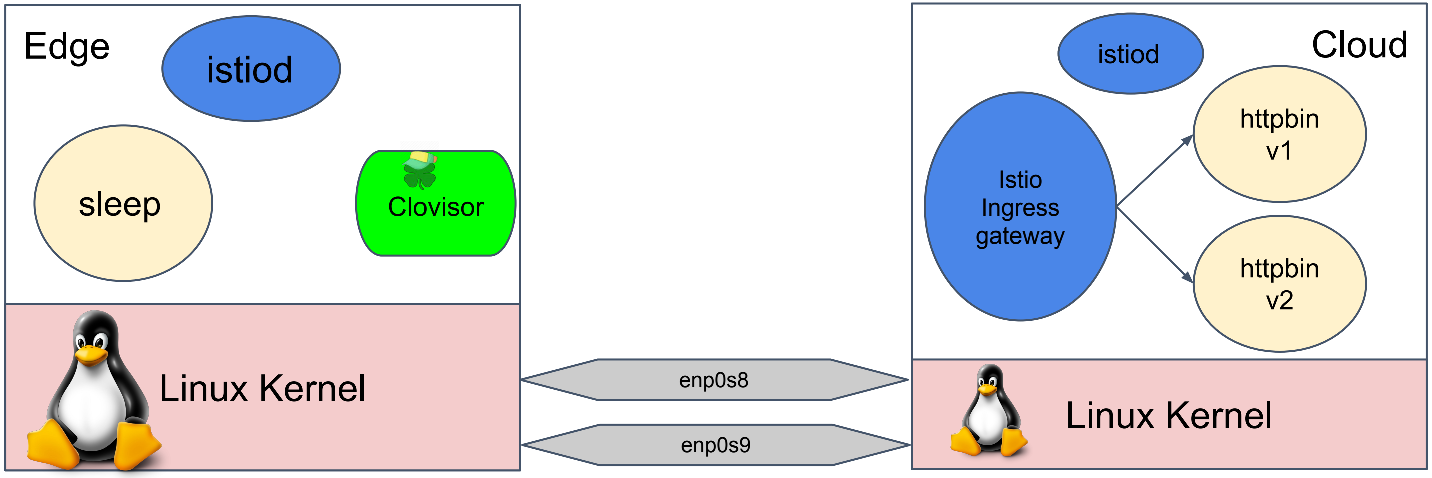

- Setup: Replicated control plane, separate networks
  - Similar to https://istio.io/latest/docs/setup/install/multicluster/gateways/
- Two interfaces simulating dual WAN interfaces from edge node to cloud node
  - Each node is a single Kubernetes cluster

Demo (under the hood)

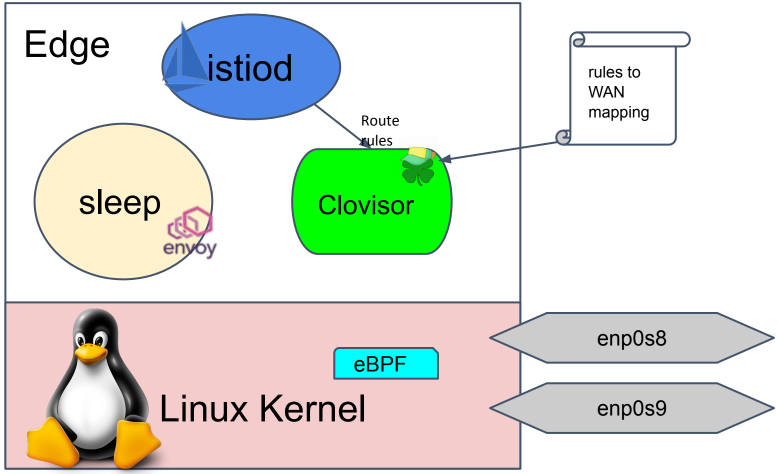

- Rules to WAN mapping loaded to Clovisor
- Clovisor fetches route rules via Istio client-go
  - Clovisor needs to implement the route rule logic
- Istio Envoy Lua filter loaded
  - Lack of service port for sleep

Demo (under the hood)

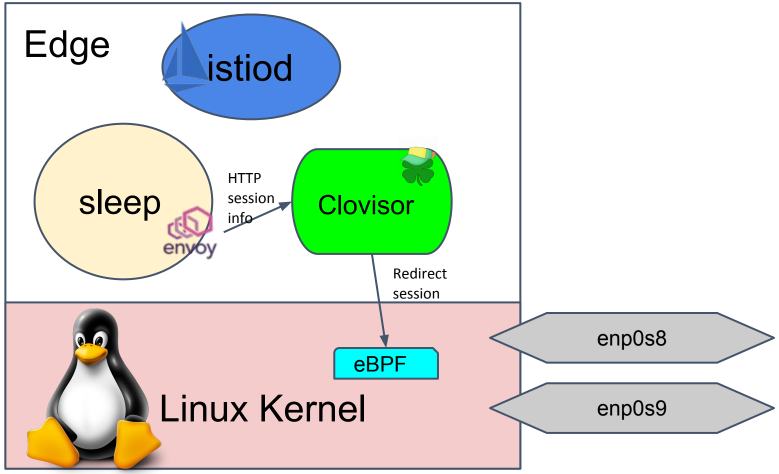

- Request going through the edge side Envoy
  - Route rule does not get applied there --- packet intercepted already
- Envoy Lua filter runs at SIDECAR_INBOUND
  - Updates Clovisor on the session which got matched with the rule classification (user == “boss”)

Demo (under the hood)

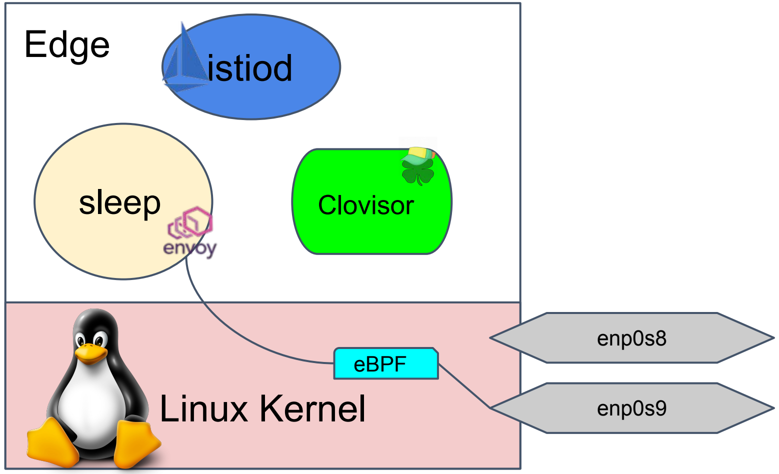

- Clovisor sets redirect rules on egress side of the original outbound WAN interface
    - Packets for user “minion” are set out on second WAN interface
- Traffic now goes through second WAN interface for user “minion”

Future Enhancement

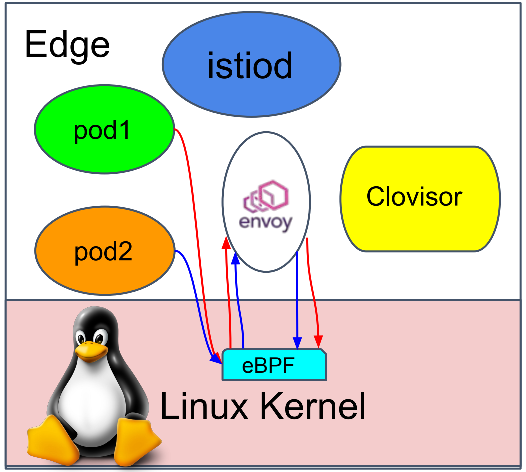

- Single Envoy for multiple pods
  - One per namespace per node
- Great for app that are more CPU bound vs I/O bound
- Limited resources, maximizing CPU power for application containers

Future Enhancement

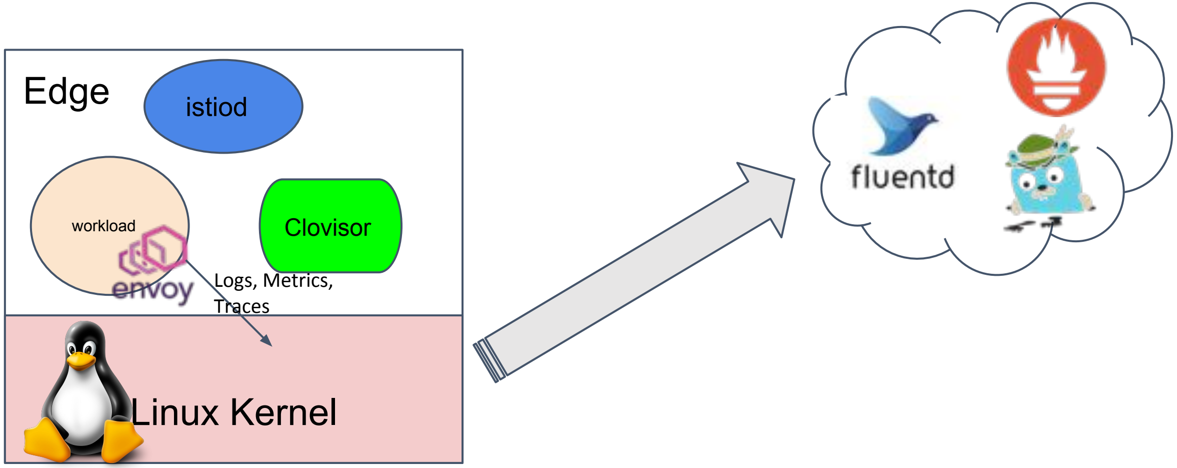

- Control plane elements
  - Logs, metrics (Prometheus), traces (OpenTracing -> Jaeger)
    - how to send, when to send
  - Storage vs WAN utilization

Summary

- Tremendous benefits for running a service mesh across cloud and edge
- WAN association maps applications knowledge to selecting appropriate physical WAN links
- Resource concern on sidecar, and control traffic are two other major areas to address for the infrastructure
- Edge computing is as much, if not more, a networking problem as it is a computing problem.

Summary

- Contact:

  clovisorproject@gmail.com

- Code:

  github.com/clovisor/clovisor

Thank You!!!

KEEP CLOUD NATIVE

KubeCon | CloudNativeCon

North America 2020

Virtual

EVERYWHERE

Summary

Backup Slides

Future Enhancement

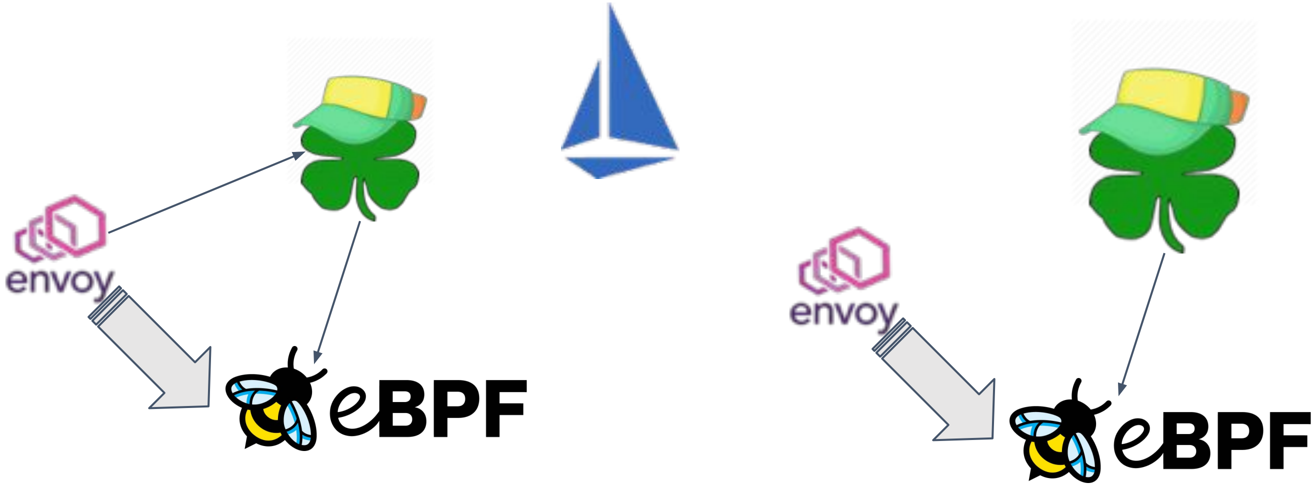

- Tag packet off of Envoy filter
  - Envoy filter (network filter) tags packets (segment that isn't encrypted after service proxy, such as IP ToS byte) to directly map to a WAN link
  - More efficient as it doesn't require communication between envoy filter and Clovisor