Service Mesh and Related Microservice Technologies in ONAP

Contributors:
Ramki Krishnan (VMware), Srini Addepalli (Intel), Manoj Nair (Net Cracker), Tal Liron (Red Hat), Roger Maitland (Amdocs), Huabing Zhao (ZTE)

Primary Sources:

- ONAP Santa Clara Event: “Towards a Comprehensive ONAP Operations Management Solution, 2017 Santa Clara event,” https://wiki.onap.org/download/attachments/16002054/ONAP-comprehensive-oom-v1.pdf?version=1&modificationDate=1513963503000&api=v2
- Material from various Kubecon Austin 2017 presentations on leading edge open source technologies for platforms & microservices and CNCF Open Source projects

Agenda

“Service Mesh”

- Evolution & Business Benefit
- Architectural Overview & Popular Open Source Projects
- Change Management & Security Examples
- ONAP Problem Statement
- Mapping to ONAP S3P
- ONAP Architectural Recommendations
- Cloud Native ONAP Workloads - ONAP Impact
- ONAP Collaboration Recommendation

Service Mesh - Evolution & Business Benefit

Cloud-native Feature set	Design Pattern	Implementation Examples	Benefit	Challenge
Request-level load balancing, circuit-breaking, instrumentation, hop-by-hop security	Client libraries for load balancing, circuit breaking, telemetry, hop-by-hop security	Stubby at Google, Hysterix at Netflix, Finagle at Twitter etc.	@scale Microservices deployment	**Tight coupling between service owner and platform team**
Same as Above	**"Service Mesh" -** Pluggable sidecar containers & Proxies for Above	Open Source: Istio, Linkerd etc.	Above + **Decoupling service owner from platform team – clear separation of concerns**	Performance impact for certain use cases

Service Mesh - Architectural Overview & Popular CNCF projects

Service Mesh Data Plane - pluggable sidecars for distributed policy enforcement, telemetry agent function etc.:

- Linkerd, NGINX, HAProxy, Envoy, Traefik etc.

Service Mesh Control Plane - centralized policy administration, telemetry collection etc.:

- Istio, Nelson, SmartStack, Conduit etc.

Related Popular CNCF Projects

- **Service Assurance** – Prometheus
- **Distributed Tracing** – Zipkin, Jaeger
- **Visualization** - Grafana

Simpler Architecture

- Minimizes code changes and increases stability
- Most functionality in sidecars (platform containers)
- Application is not aware of sidecars
- Automatic sidecar injection – no deployment modification

Hadoop MapReduce Ecosystem Analogy

- *User needs to implement just a mapper and a reducer*
- *Rest all is provided by Hadoop infrastructure*

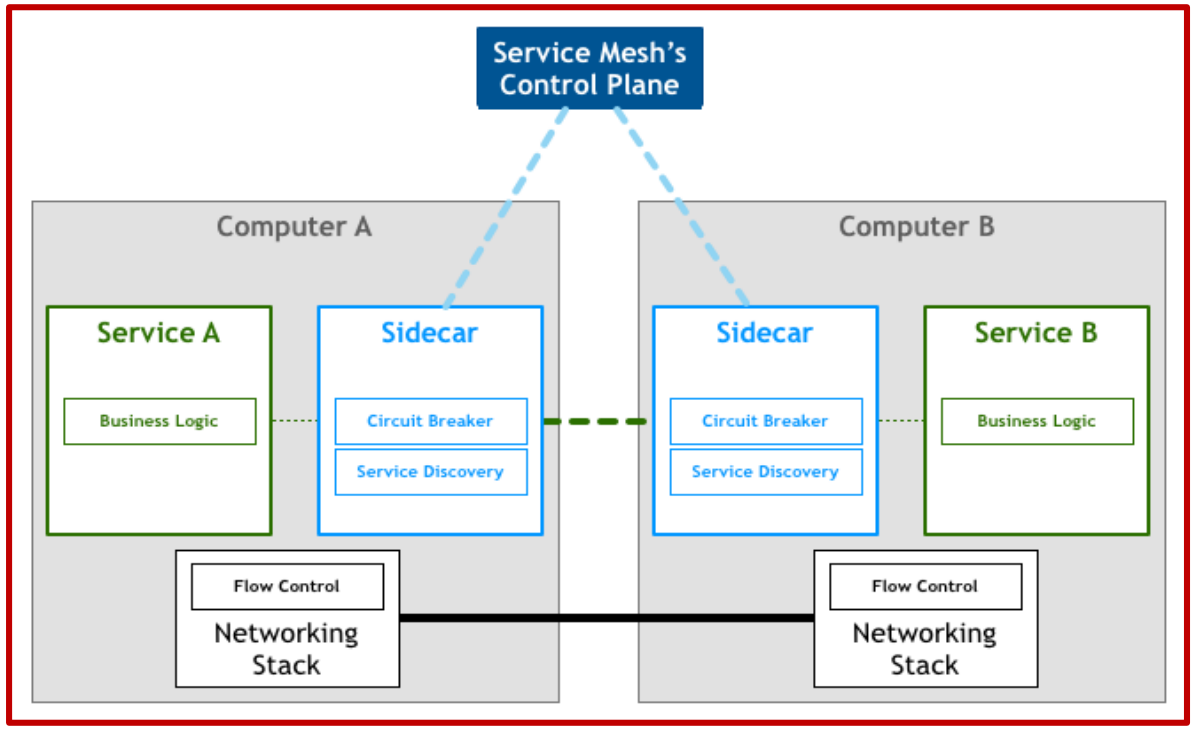

Image Source Philip Calcado, Bouyant Inc.
http://philcalcado.com/2017/08/03/pattern_service_mesh.html

Service Mesh - Live Change Management Example

Automatically maintain multiple versions

- **Istio Pilot/Envoy** service tags provide finer-grained routing (partitioning endpoints between A/B for A/B testing)
- **Live Upgrades can be tested in real production environments without affecting services**

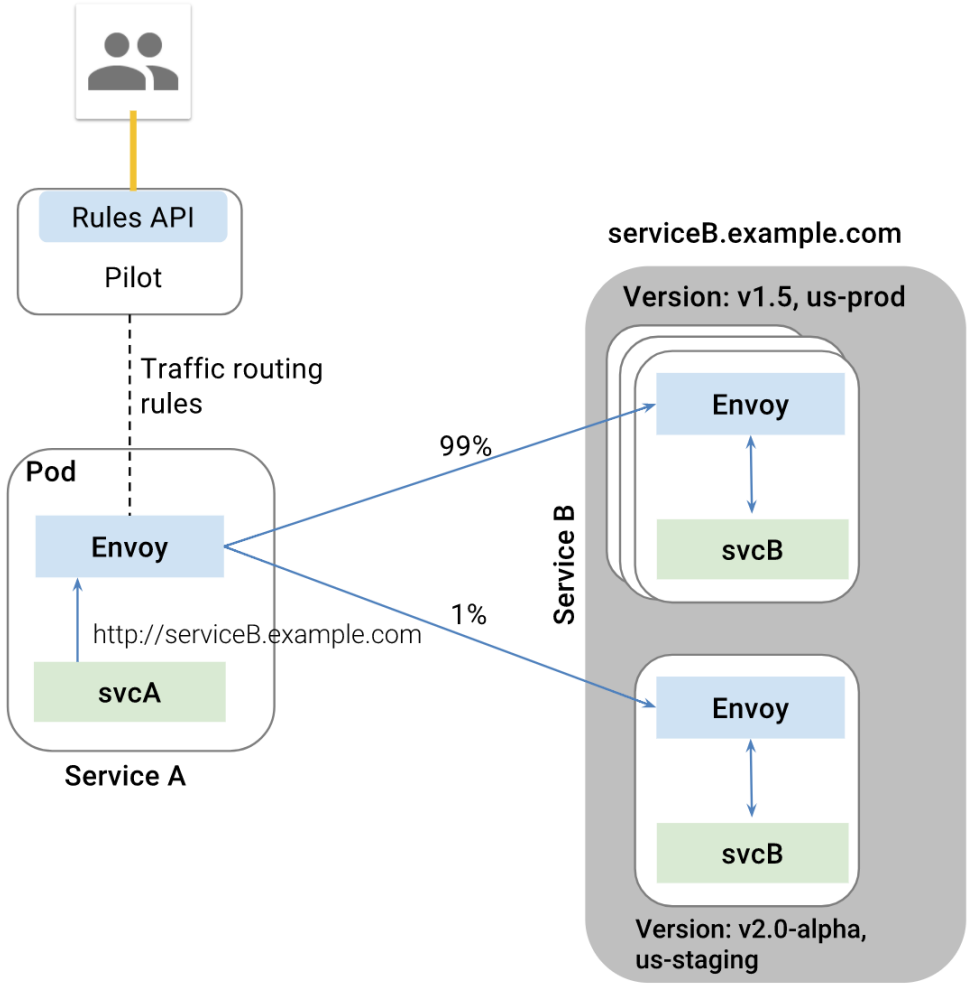

Service Versions

Source: Istio documentation
https://istio.io/docs/concepts/traffic-management/request-routing.html

Service Mesh - Zero-Trust Security Example

Traditional perimeter security via firewalls is not sufficient
- Often, more attacks originate from "inside”
- Definitions of "inside" and "outside" change dynamically

Specialized techniques for single type of security are not sufficient
- EBPF techniques can be used, but not sufficient

Requires mutual TLS between all segments (across each hop)
- Currently, often internal communication is not encrypted

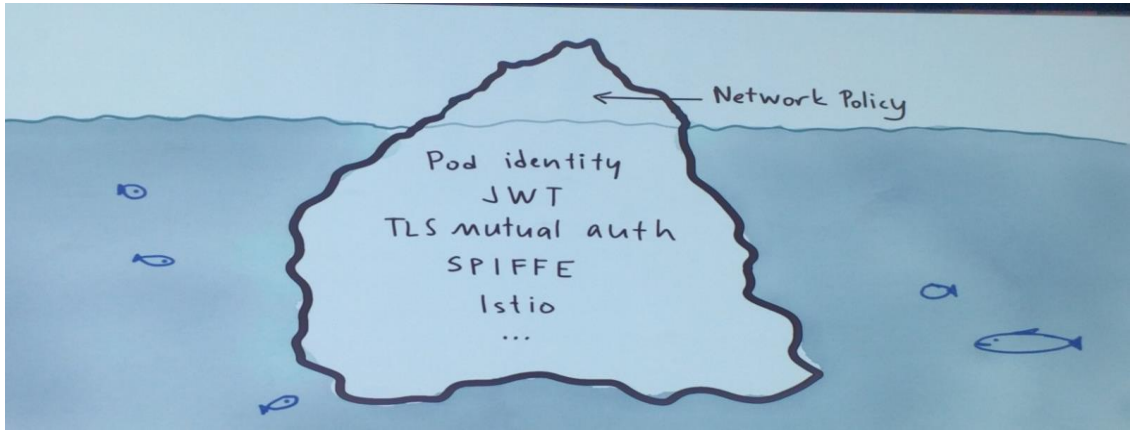

Source: talk by Ahmet Balkan (Google) at Kubecon 2017
via Tweet from Evan Gilman (Scytale; @evan2645)

Policy based orchestration of network security
- Control plane and pluggable sidecar (platform container) proxies

Authentication and secure service-to-service traffic management
- Secret management, regular rotation of credentials, bootstrapping credentials
- Compatibility with variety of external authentication systems

Open Source Standard: Secure Production Infrastructure Framework for Everyone (SPIFFE)
SPIFFE Implementations: SPIRE, Istio Auth

Service Mesh – ONAP Problem Statement

- Challenges in ONAP Micro Service Architecture
  - Need for client libraries for infrastructure services
    - Consumes lot of time by each project; Error prone; Mismatch in capabilities across projects.
    - Few examples:
      - Integration of Authentication and Authorization by every service; Mutual-TLS enablement by every service; Storing Secrets securely by each service.
  - Polyglot challenges
    - ONAP services are written in various languages (Java, Python etc.)
    - Client libraries in various languages and language specific restrictions to use some features uniformly.
  - Visibility of inter-service data for trouble shooting and any kind of analytics & and also tracing the requests across projects.
    - If Mutual-TLS implemented from the service itself, any troubleshooting that is needed during operations is tough and expect code changes in every project for debugging.
  - Service Discovery, Load balancing of requests among services, Circuit breaking, Health Checks and API routing etc. require changes in every ONAP service container.
  - No easy way to address rolling updates and ensuring continuous operations (versioning etc.)
  - Achieving security is dependent on ONAP projects (System call filtering, Storing keys securely in TPM/SGX, Keeping password secure etc.)

- **Solution Direction – Service Mesh**
  - **Eliminate or reduce micro-service infrastructure level tasks in ONAP services**
  - **Ensure that ONAP Services have only the business logic of that service**

Service Mesh – Mapping to ONAP S3P (and more)

Cloud-native Feature set	Service Mesh Data Plane for K8S	Service Mesh Control Plane for K8S	ONAP S3P Mapping
Resiliency (timeouts, circuit breakers, etc.) – (F) in FCAPS	Envoy, Linkerd …	Istio Pilot, Conduit …	Resiliency Level 2 – single site automated recovery
Troubleshooting/Tracing – (F) in FCAPS	Envoy, Linkerd …	Istio integrated with Zipkin/Jaeger, Conduit …	Manageability Level 2 – tracing across components
Fine-grained routing/load-balancing across multiple SW versions – (C) in FCAPS	Envoy, Linkerd …	Istio Pilot, Conduit …	Manageability Level 2 – single component upgrade; **Live upgrade (ONAP S3P+)** Scalability Level 1 – single site horizontal scaling
Visibility/Telemetry – (A) (P) in FCAPS	Envoy, Linkerd …	Istio Mixer with Prometheus/Grafana, Conduit …	Performance Level 2 & 3 Manageability Level 2 – single logging system
Hop-by-hop security – (S) in FCAPS	Envoy, Linkerd …	Istio Auth, Conduit …	Security Level 2 -- internal communication encrypted

Note: Istio's data plane of choice is Envoy, Conduit's data plane of choice is Linkerd

Service Mesh – ONAP Architectural Recommendation (1)

- Service Mesh Component Implementation Recommendation (Metrics: maturity, ease of use, contributing companies)

- Service Mesh Control Plane	- Istio	- Foundational
- Service Mesh Data Plane	- Envoy	- Foundational
- Tracing	- Zipkin/Jaeger addon	- Istio Integration
- Service Assurance	- Prometheus addon	- Istio Integration
- Metric Visualization	- Grafana addon	- Istio/Prometheus Integration
- Service Graph	- Servicegraph addon	- Istio/Prometheus Integration

- Related Component Implementation Recommendation

- Logging Collector	- Fluentd daemon	- Fluentd K8S Integration

Service Mesh – ONAP Architectural Recommendation (2)

- **Casablanca** - Service Mesh & Related Technologies for at least some ONAP Components on K8S
  - ONAP Project Impact & Value Proposition
    - MSB
      - Istio integration (automatic sidecar injection etc.) discussion in progress
    - AAF
      - Leverage Envoy and Istio Auth as a holistic alternative to the current implementation of secure communication between microservices
    - OOM
      - Leverage Istio etc. for reusable DBaaS components, e.g. MariaDB Galera
    - Logging
      - Leverage Fluentd K8S integration
    - Integration
      - Leverage Istio etc. Live Upgrade capability for key ONAP components

- **Casablanca+** - Service Mesh & Related Technologies for all ONAP Components on K8S
  - Expand on Casablanca work

Service Mesh & Related Technologies for Cloud Native ONAP Workloads

ONAP Impact

Service Mesh – ONAP Workloads Architectural Recommendation

- **Cloud Native** ONAP Workload Examples
  - Virtual Network Functions and Edge Applications (ONAP Edge Analytic, Optimization & Context Processing etc. -- https://wiki.onap.org/display/DW/Edge+Scoping)
  - Cloud Native vIMS Example -- https://github.com/Intel-Corp/clearwater-kubernertes

- **Casablanca+** - Service Mesh for ONAP Containerized Workloads (VNFs etc.)
  - ONAP Project Impact & Value Proposition
    - Multi-Cloud
      - Leverage Cloud Provider capabilities for service mesh and related components (Prometheus etc.)
      - Data Model for Infra/App Metrics/Alerts for ONAP Component Service Assurance
    - DCAE
      - Potential to simplify architecture and improve manageability & service velocity
        - Simplified Collection at Source
          - VNF: No VES Agent (Avoid recompiling VNF)
          - Edge Application: No App Libraries (Avoid recompiling Edge Application)
        - Collector Offload
          - Prometheus from Cloud Providers with Multi-Cloud Integration can provide VES Collector functionality
      - Deliver value through Analytics Microservices

Service Mesh & Related Technologies
ONAP Collaboration

Service Mesh – ONAP Collaboration Recommendation

- Open Source (LF Networking etc.) Collaboration Recommendation
  - OPNFV Clover Project (https://wiki.opnfv.org/display/PROJ/Clover)
    - Clover Project focus
      - Service Mesh for containerized workloads (VNFs etc.)
    - Collaboration Areas
      - Drive commonality on Service Mesh Architecture & Component Recommendation
      - Leverage automation tools for Service Mesh for ONAP Containerized Workloads (VNFs etc.) on K8S

Service Mesh in ONAP
Deep Dive

Running Istio on a ONAP K8S Cluster

Just 6 Steps for running Istio -- https://istio.io/docs/setup/kubernetes/quick-start.html

NAMESPACE	NAME	READY	STATUS	RESTARTS	AGE
istio-system	istio-ca-75fb7dc8d5-dbl6l	1/1	Running	0	1d
istio-system	istio-ingress-8546966f58-j2lmp	1/1	Running	0	1d
istio-system	istio-mixer-566f68f5d6-g6w2h	3/3	Running	0	1d
istio-system	istio-pilot-fd8fb6957-mmm4t	2/2	Running	0	1d
kube-system	etcd-ramki-virtual-machine	1/1	Running	0	8d
kube-system	kube-apiserver-ramki-virtual-machine	1/1	Running	0	8d
kube-system	kube-controller-manager-ramki-virtual-machine	1/1	Running	0	8d
kube-system	kube-dns-86f4d74b45-mgzm5	3/3	Running	826	8d
kube-system	kube-flannel-ds-s72r7	1/1	Running	0	8d
kube-system	kube-proxy-wrzkn	1/1	Running	0	8d
kube-system	kube-scheduler-ramki-virtual-machine	1/1	Running	0	8d
kube-system	tiller-deploy-664858687b-hnbsn	1/1	Running	0	8d
onap	dev-kube2msb-584b9d9c75-jdjf7	1/1	Running	6	8d
onap	dev-msb-consul-797447d5f9-lzw9k	1/1	Running	0	8d
onap	dev-msb-discovery-776497ffd9-zn8cx	1/1	Running	0	8d
onap	dev-msb-eag-66dd455d4b-2pk5f	1/1	Running	0	8d
onap	dev-msb-iag-588956774f-87xln	1/1	Running	0	8d