

# OPEN POSSIBILITIES.

Empowering Network Operations Teams  
with Pervasive Kubernetes Observability



**OCP**  
GLOBAL  
SUMMIT

NOVEMBER 9-10, 2021

# Empowering Network Operations Teams with Pervasive Kubernetes Observability

Alessandro Barbieri, VP Product Management, Pluribus Networks



OPEN POSSIBILITIES.



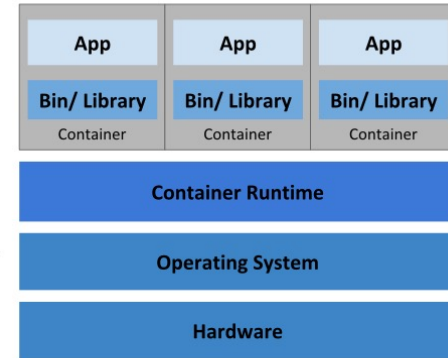
# What Is Kubernetes and What Does It Provide

*“Kubernetes is a portable, extensible, open-source platform for managing containerized workloads and services, that facilitates both declarative configuration and automation.”*

- **Declarative configuration and automation:** A Manifest or configuration file (yaml) describes the *desired state* for deployed containers using Kubernetes, and Kubernetes constantly compares the actual state to the desired state (control loop) and makes the necessary changes to converge the actual state to the desired state.
- **Distribution and load balancing traffic:** Kubernetes load balances and distribute the network traffic so that the deployment is stable and resources are optimized.
- **Self-healing:** Kubernetes restarts containers that fail, replaces containers, kills containers that don't respond, and doesn't advertise them to clients until they are ready to serve.
- **Automatic bin packing:** Based on how much CPU and memory each container needs, Kubernetes can fit containers onto your nodes to make the best use of your resources.
- **Storage orchestration:** Kubernetes allows you to automatically mount a storage system of your choice, such as local storage volumes, public cloud providers, and more.



NETWORKING



Container Deployment

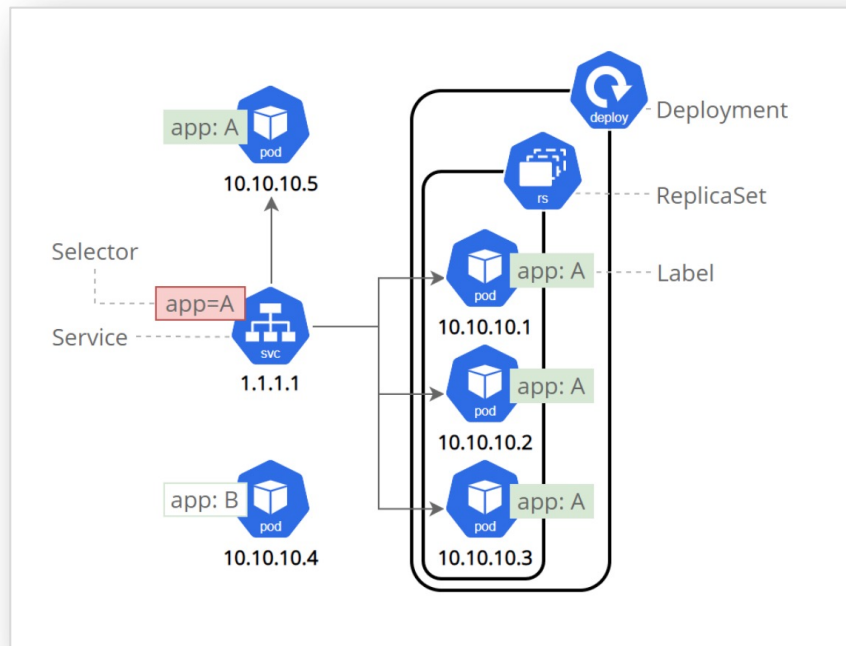


# Kubernetes Components and Basic Concepts for NetOps



NETWORKING

1. **Pods** are the smallest units that can be deployed in Kubernetes. Pods consists of one or more containers that are always scheduled together.
2. Each Pod is given a unique IP address. Pods come and go frequently.
3. **Services** provide a consistent/stable virtual IP address for a group of Pods. Services acts as a **loadbalancers** and are based on Linux IPTable.
4. Pods are organized in ReplicaSets objects (groups of identical Pods), which in turn are managed with Deployments objects



OPEN POSSIBILITIES.

# The NetOps Challenge of Root Causing Application Availability and Performance Issues in a Kubernetes Environment

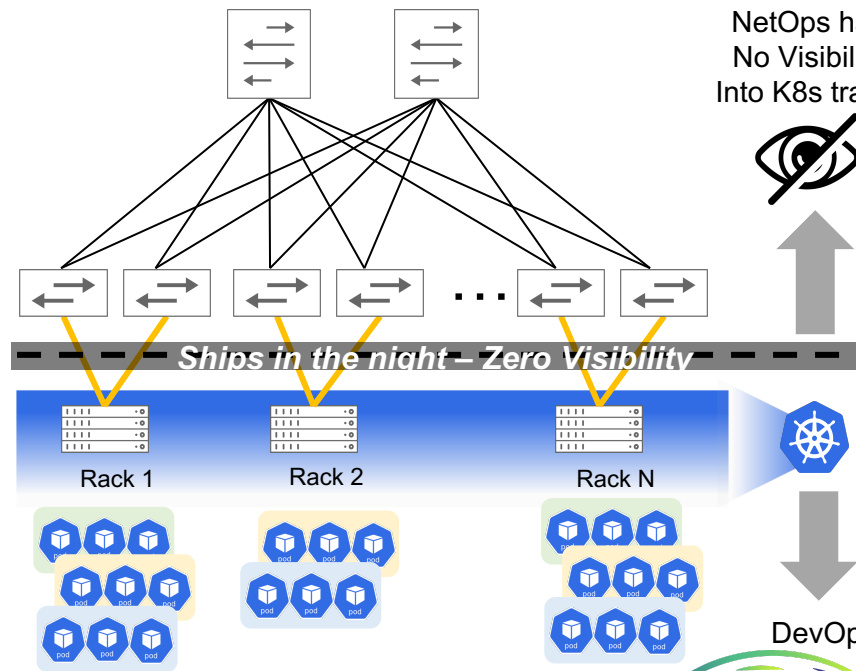


NETWORKING

NetOps has  
No Visibility  
Into K8s traffic



- **Short lifespan:** Containers come and go frequently and may only live for a few minutes. *How can the network admin troubleshoot what is long gone?*
- **Endpoint sprawl:** Containerized apps have 100x more end-points than monolithic applications spread among all the hosts connected to different points in the network. *How can the network admin locate the end-points of a distributed applications when the application has availability or performance problems?*
- **East-west traffic Loadbalancing:** Containerized apps are grouped into *services* communicating through a series of *Virtual IP* constantly loadbalancing traffic among hosts. *How can the network admin trace the path of pod-to-pod traffic flows “bouncing” across multiple cluster nodes?*
- **“Opaque” traffic:** POD-to-POD communication is often encapsulated (e.g. VXLAN or IP-in-IP). *How can the network admin have any visibility into pod-to-pod encapsulated traffic without specialized tools?*

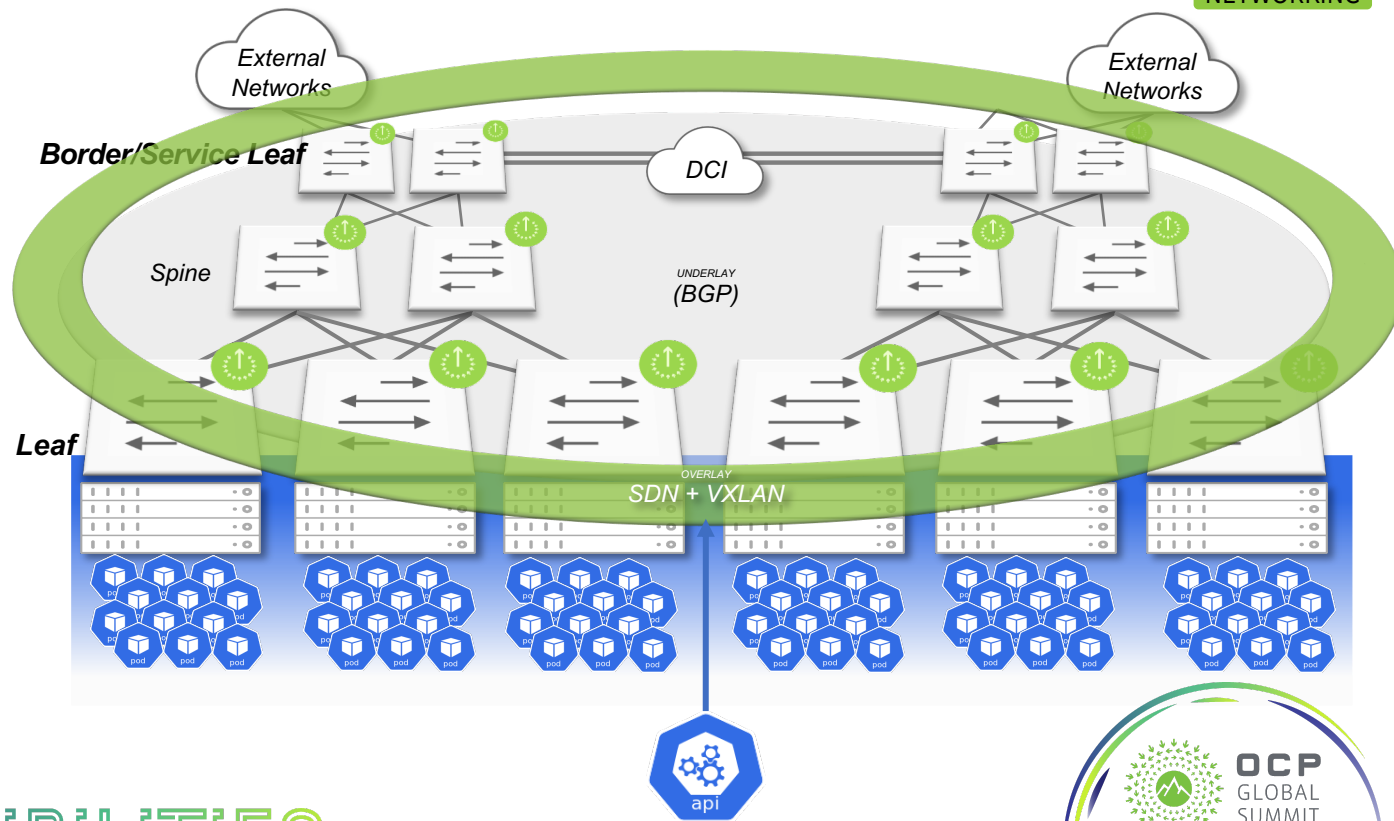


# Components of A “Kubernetes-Aware” Switch Fabric



NETWORKING

- **OCP/ONIE switches** powered by Broadcom merchant silicon
- BGP based fabric with VXLAN overlay
- **SDN control plane to enable end-to-end visibility** into Kubernetes end-points and traffic flows
- SDN fabric ***subscribes*** to K8s **API server** for real time notifications of K8s deployment events and changes



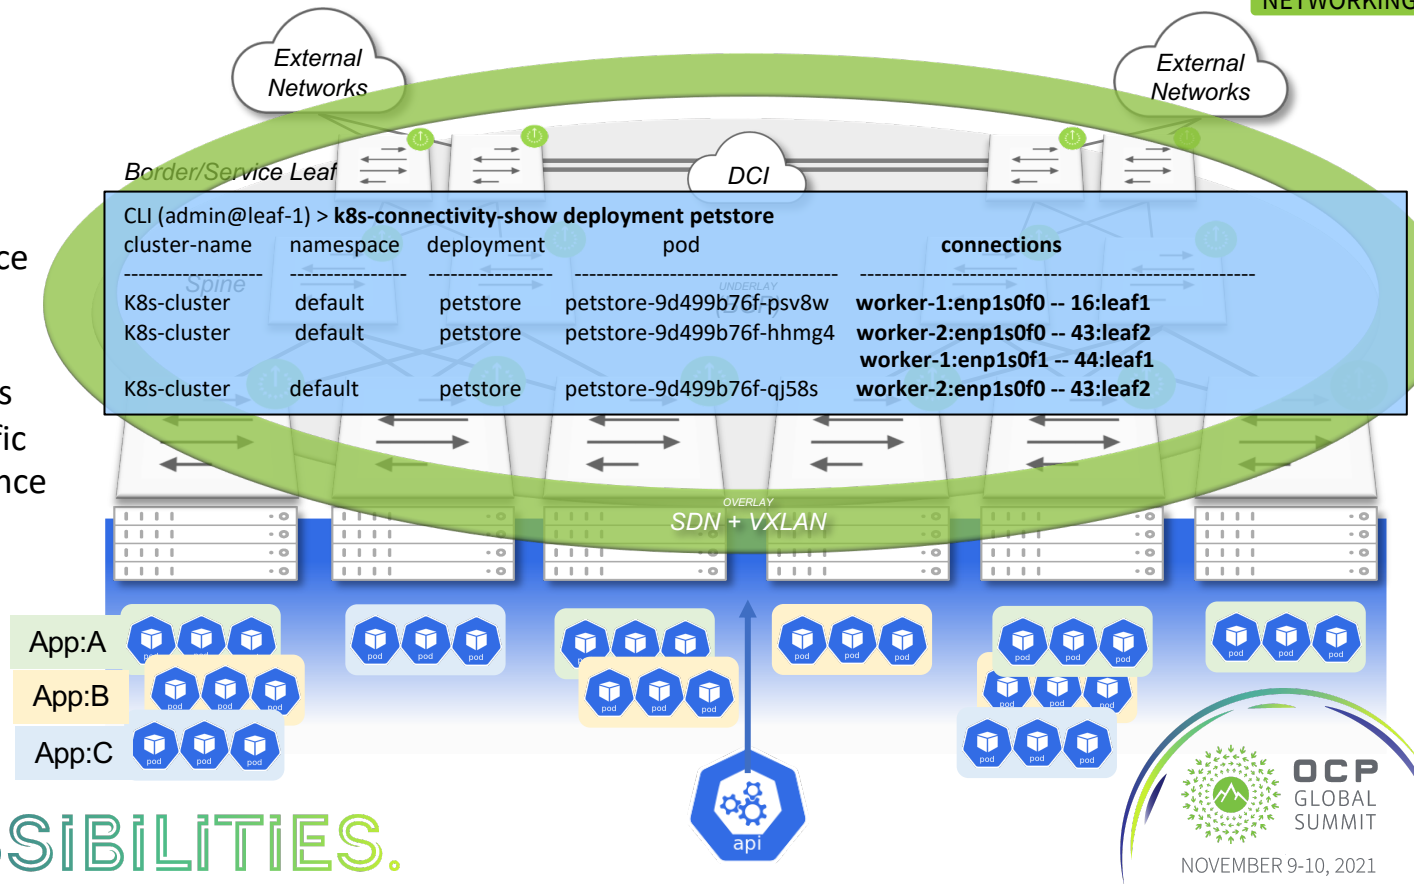
OPEN POSSIBILITIES.

# K8s-Aware Fabric: Locate Application PODs In The Fabric



NETWORKING

- Locate where PODs for specific applications are connected to the Fabric to resolve any application connectivity or performance issues.
- Correlate port traffic spikes and congestion with specific K8s applications performance issues

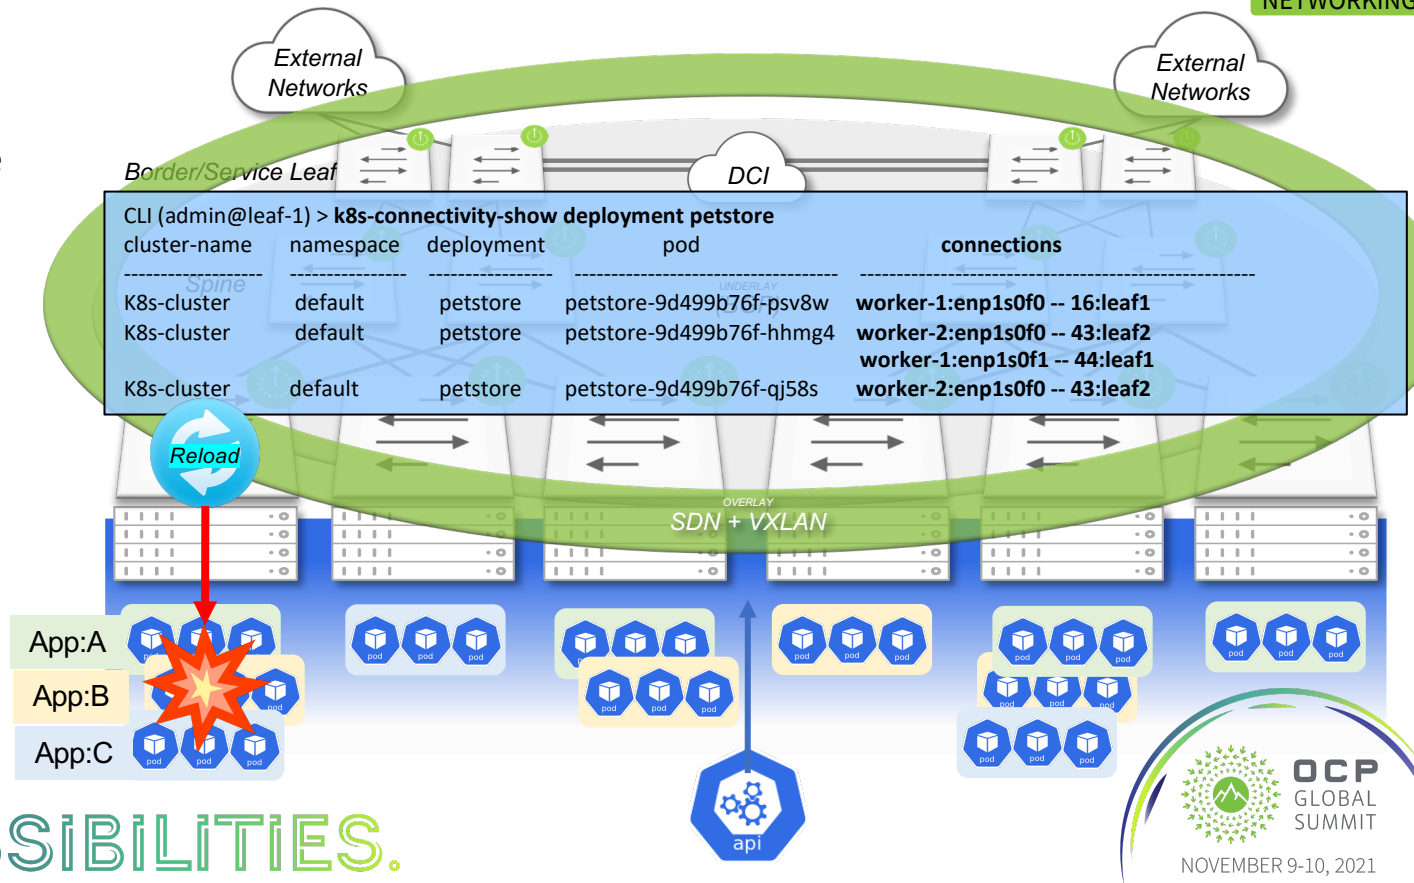


# K8s-Aware Fabric: Assess Impact on Applications of Node Outages



NETWORKING

- Proactively determine the impact on applications of switch outages during service windows (e.g. software upgrades, RMA etc.)



# K8s-Aware Fabric: Track K8s Changes With Time Machine



NETWORKING



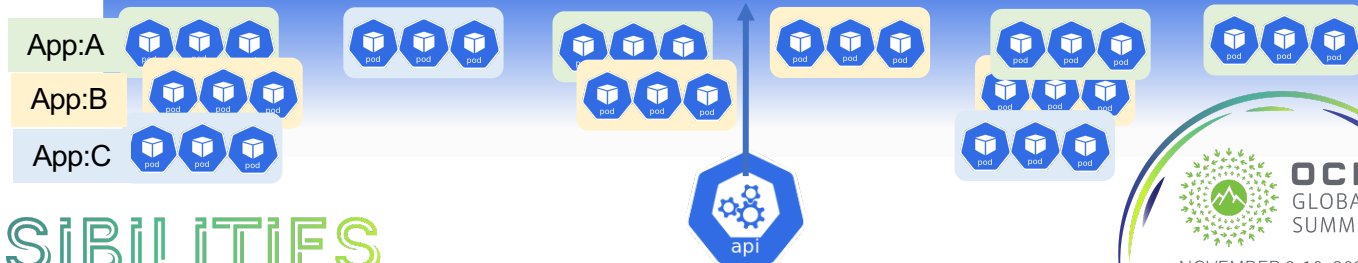
Leverage  
switch  
SSD

CLI (min@leaf-1) > **k8s-pod-show within-last 24h**

cluster-name	name	namespace	ip	state	node	containers	last-changed	action
K3s-TME	hello-world-67b845476-d9fgn	default	10.42.1.60	Pending	worker-1	rancher:1.0.0	09-09,12:13:19	add
K3s-TME	hello-world-67b845476-9pl2f	default	10.42.0.14	Running	worker-1	rancher:1.0.0	09-09,12:13:19	update
K3s-TME	hello-world-67b845476-d9fgn	default	10.42.2.54	Running	worker-2	rancher:1.0.0	09-11,15:40:33	remove
K3s-TME	hello-world-67b845476-9pl2f	default	10.42.2.58	Running	worker-2	rancher:1.0.0	09-11,21:27:32	remove

CLI (network-admin@auriga-colo-1) > **k8s-deployment-show name petstore start-time 9:00:00 end-time 9:30:00**

cluster-name	name	namespace	replicas	ready-replicas	selector	last-changed	action
K3s-TME	petstore	default	4	4	app:petstore	09:11:51	add
K3s-TME	petstore	default	6	4	app:petstore	09:24:35	update
K3s-TME	petstore	default	6	5	app:petstore	09:28:35	update
K3s-TME	petstore	default	6	6	app:petstore	09:29:59	update



OPEN POSSIBILITIES.

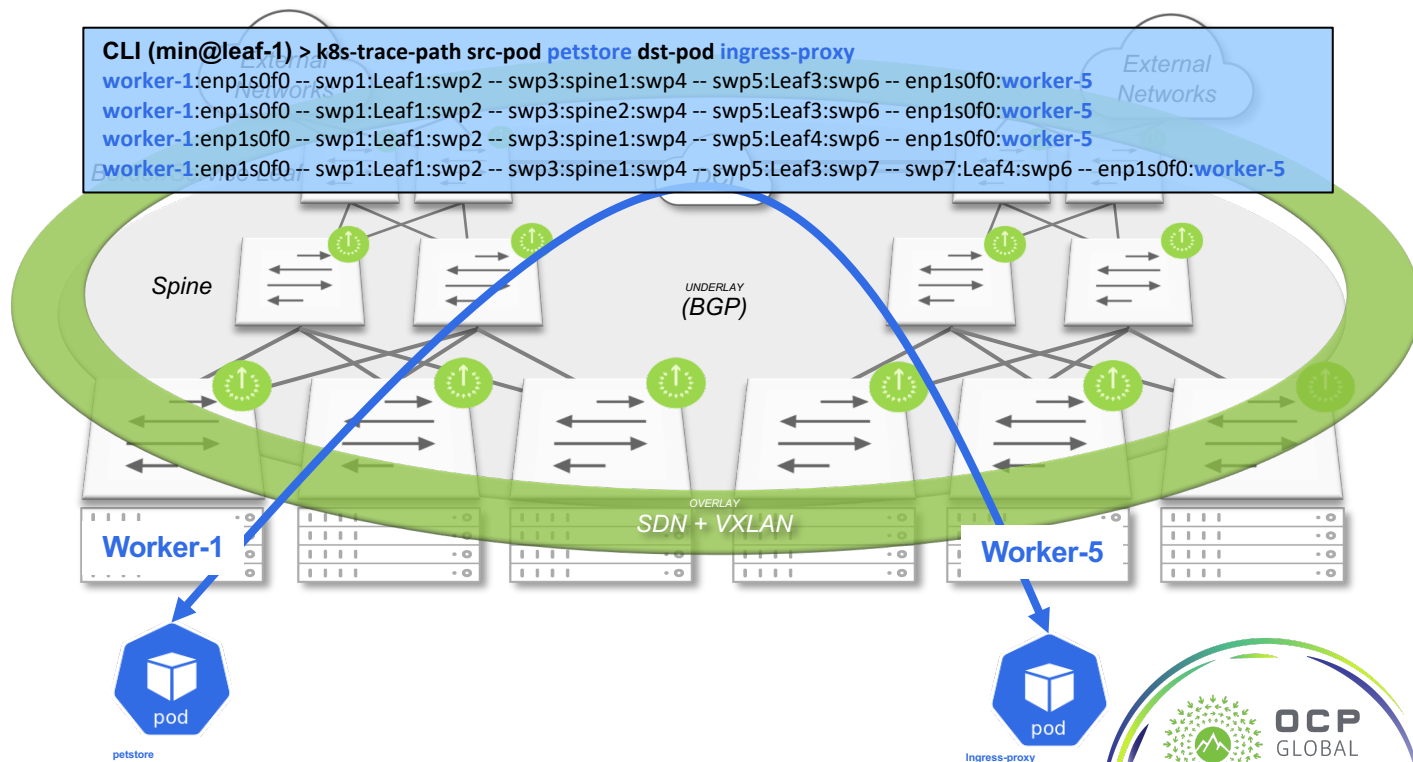


# K8s-Aware Fabric: Pod-to-Pod Path Tracing



NETWORKING

- Identify all the possible paths for pod-to-pod communication across the ECMP fabric
- Locate possible traffic hot-spots and congestion bottlenecks impacting application performance



OPEN POSSIBILITIES.

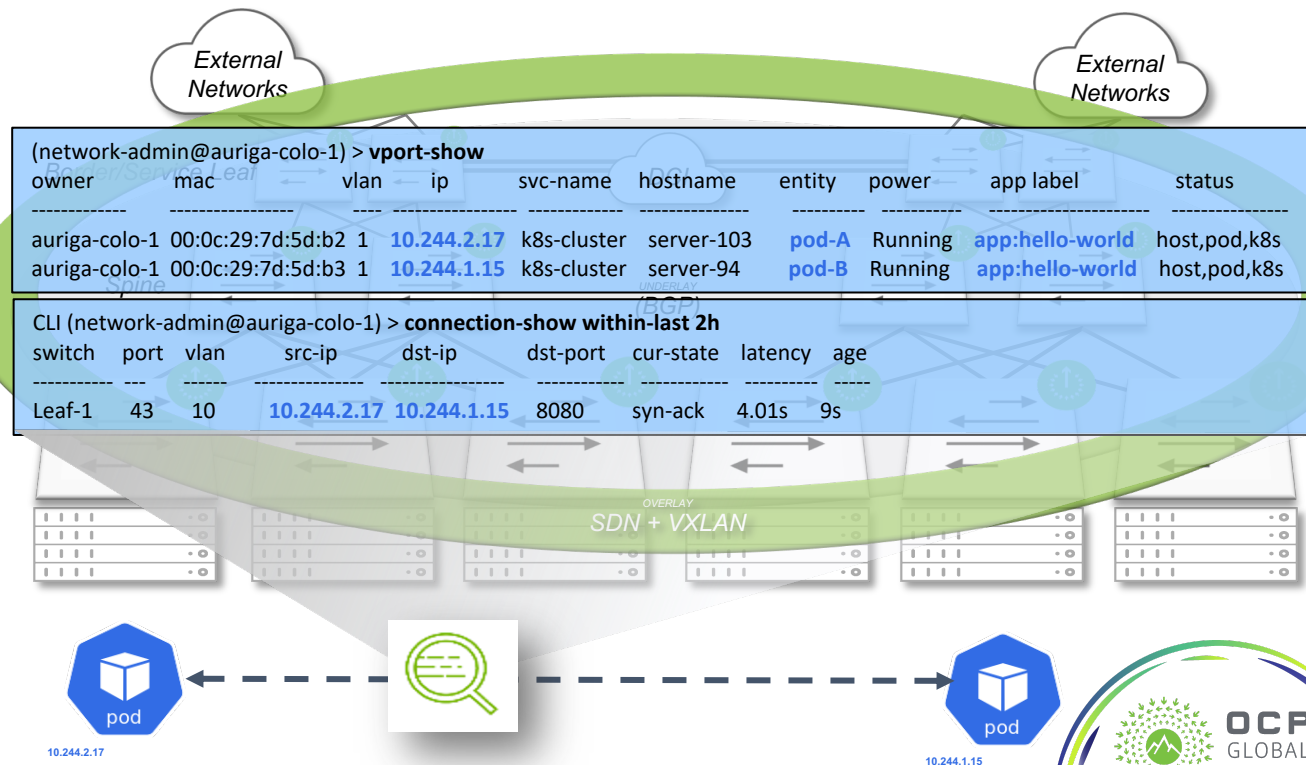


# K8s-Aware Fabric: Pod-to-Pod Application Flow Visibility



NETWORKING

- Troubleshoot application flow with K8s context information
- **No extra probes, taps, host agents and packet brokers:** Just your OCP switches
- Export flow telemetry to external analytics engine



OPEN POSSIBILITIES.

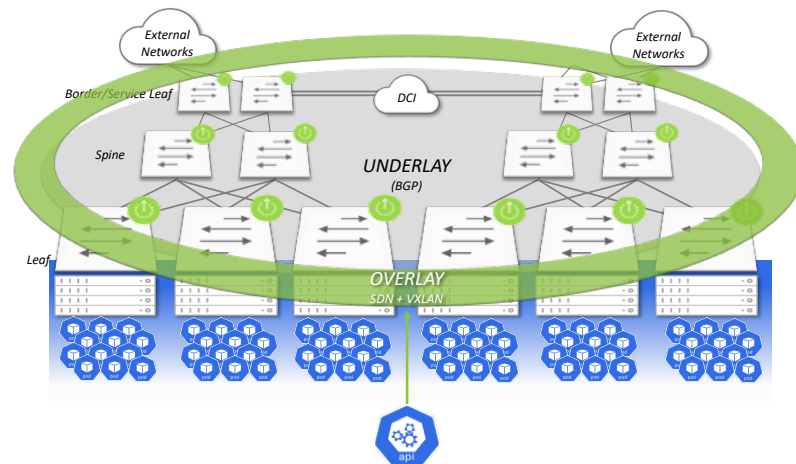


# In Summary: Leveraging The Switch Fabric To Accelerate Root Cause Analysis of Application Performance Issues In a K8s Environment

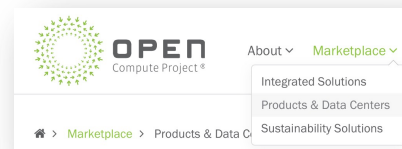


NETWORKING

- OCP switches controlled by an **SDN control plane for end-to-end** visibility across the fabric of end-points and traffic flows.
- SDN control plane integrating a **subscription service to K8s API server** to bring visibility of the K8s cluster into the switch fabric.
- Enables the network admin to rapidly correlate and root cause K8s application performance and availability issues with network congestion, or network outages.
- **No agents required** on the K8s cluster nodes. Works with any K8s distributions.
- Implements Time Machine to eliminate the need to reproduce past events impacting K8s applications.



11 based “OCP Approved”, Broadcom-based switches listed on the OCP Marketplace:



## Software

- ☒ Pluribus (11)
- ☐ ONIE (30)
- ☐ ONL (25)
- ☐ Netris (8)
- ☐ TIP (7)
- ☐ VMWare (7)
- ☐ Red Hat (6)

Show more

OPEN POSSIBILITIES.

# Visit the OCP Marketplace to Learn More

<https://www.opencompute.org/products>

Netvisor ONE OS and the Adaptive Cloud Fabric Software are available on the following OCP Accepted hardware platforms



Pluribus Freedom 9000 Series



**OPEN**  
ACCEPTED™

OPEN POSSIBILITIES.

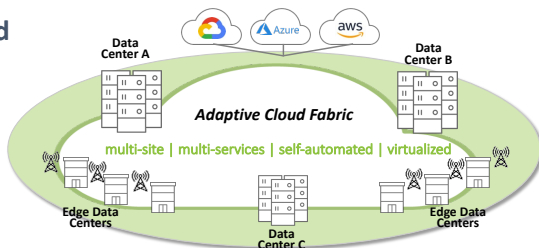


# Pluribus Networks At A Glance

- We build an open NOS for OCP and open networking switches
- The NOS integrates a distributed SDN architecture to radically simplify Private Cloud and multi-site datacenter deployments
- Deployed in 100+ Tier 1 Mobile Operators around the world

Cloud-like, centralized programming model

Simplified, automated, protocol-free overlay services abstraction



Unified Fabric from single site to multiple sites w/o any controllers

SDN overlay interoperable with 3<sup>rd</sup> party BGP EVPN  
Open Standard L2/L3 Underlay

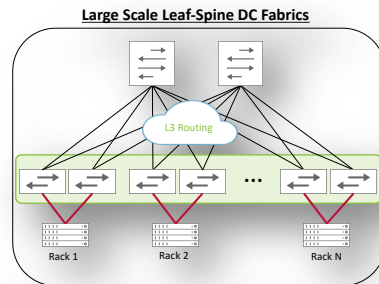
- Partnerships with OCP technology providers like Edgecore and ITRenew



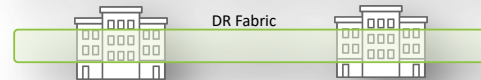
- Strategic Partnerships and Technology Integrations



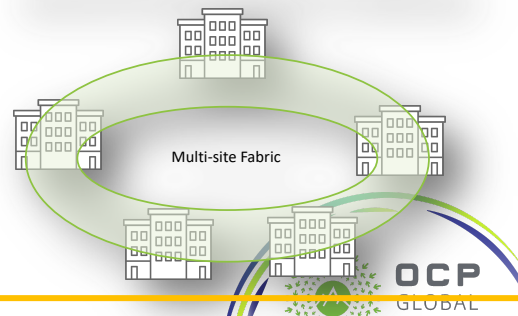
Distributed, Flexible, Adaptable  
To Any Topology



DR / Active-Active DC Topology Managed as a Unified Fabric



Multi-site DC Topology managed as a Unified Fabric



Stop by Booth A26 to learn more about Pluribus and see a demo

Thank you!



NOVEMBER 9-10, 2021