MODERN APPLICATION DELIVERY WITH

KUBERNETES

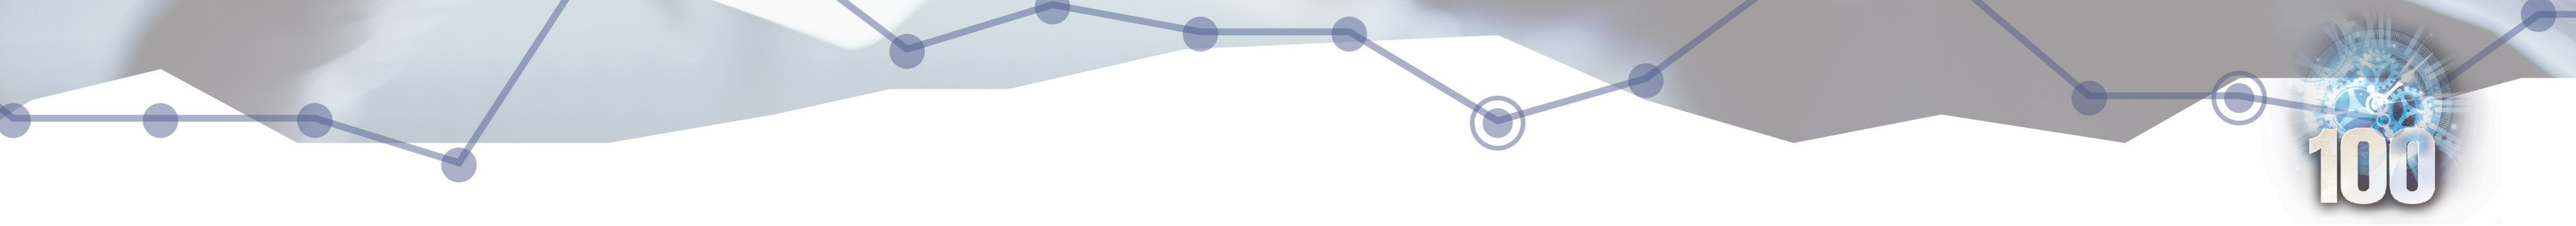

kubernetes

etcd

CoreDNS

CNI

gRPC

Open Policy Agent

Istio

envoy

SECURITY

HELM

APP PACKAGING + DEPLOYMENT

fluentd

elastic

Prometheus

Grafana

MONITORING

OPENTRACING

JAEGER

TRACING

Virtual Kubelet

THE BASE

PLATFORM

WHAT IS KUBERNETES?

- Orchestrates computing, networking and storage infrastructure (and more)
- Portable, extensible open-source platform for managing containerized workloads
- Uses declarative configuration, facilitating automation
- Huge and rapidly growing community
  - Second largest repository on GitHub

WHAT IS KUBERNETES?

- Supports Linux and Windows containers
- Support different container runtimes
  - Docker
  - Containerd
  - and anything that implements the Container Runtime Interface (CRI)
- Many supported network plugins (cloud network, overlay networks, etc.)

WHAT IS KUBERNETES?

- All major cloud providers offer a managed service
  - Google, Microsoft, Amazon, Digital Ocean, IBM, Oracle, and more!
- Platform as a Service (PaaS) offerings
  - RedHat Openshift, VMWare Cloud PKS, Pivotal PKS, and more!
- Self-hosted (on-premise or on IaaS)
  - Kubeadm

WHAT IS KUBERNETES?

- Used by many large organization
  - Google (original creator of Kubernetes)
  - GitHub
  - Reddit
  - Shopify
  - Tinder
  - and more!

ENHANCE YOUR CLUSTER

PLATFORM TOOLS

PROMETHEUS / GRAFANA

- Real-time cluster metrics
  - Broken down by application
- Alerting rules

ELASTICSEARCH

- Centralized application and cluster logging
- Indexed and can be searched

HELM

- Official Kubernetes package manager
- Helm Charts help define, install and upgrade complex applications with ease
- Templating of your applications cluster resources
  - Easily deploy multiple instances of your application
- Provides application lifecycle management
- Many official charts: https://github.com/helm/charts

VELERO (FORMERLY ARK)

- A utility for managing disaster recovery of cluster resources and data volumes
- Takes snapshots of your cluster (and its data)
- Restore across clusters (and soon, cloud providers)


AUTO-GENERATED SERVICE GRAPH (KIALI)

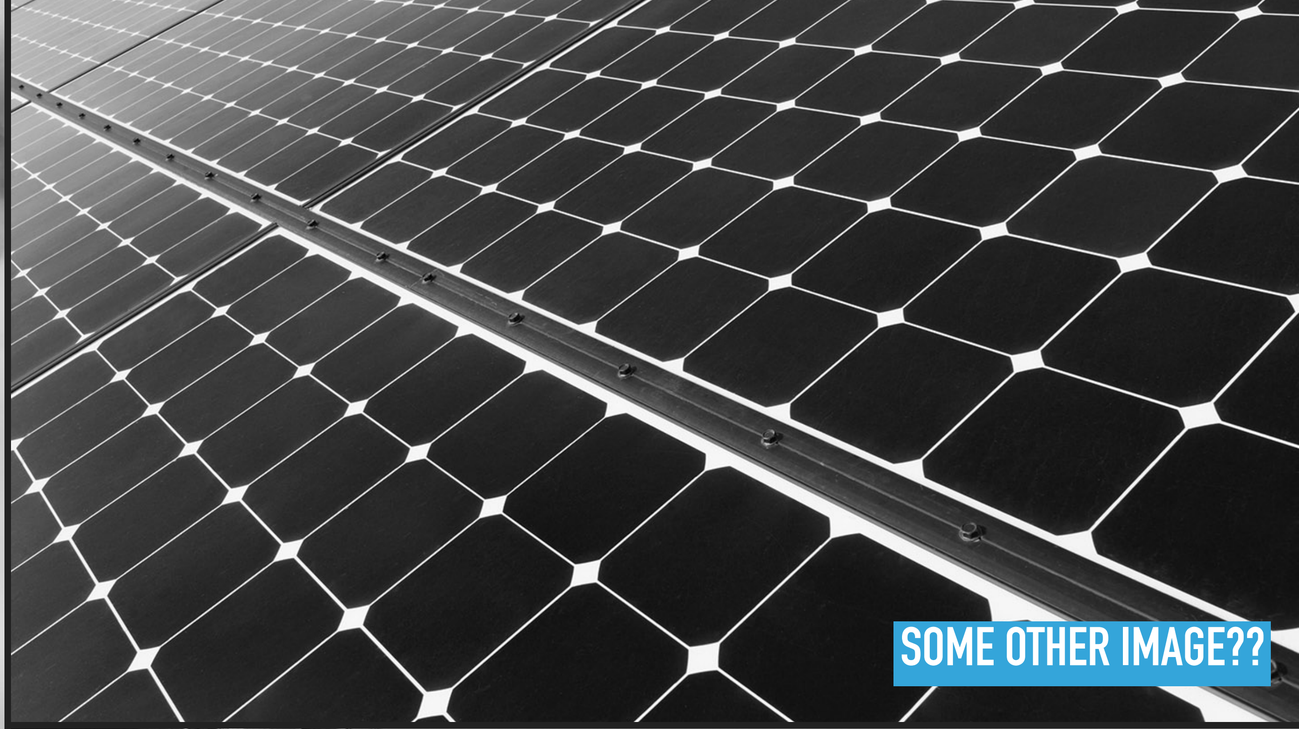
SOME OTHER IMAGE??


TELEMETRY: WHO'S TALKING TO WHO

VIRTUAL KUBELET

- Masquerades as a cluster node but interfaces with external runtimes
  - Azure Container Instances
  - Amazon AWS Fargate
  - and more!
- Hybrid workflow between dedicated machines and serverless infrastructure
  - Cost-benefit: use dedicated machines for base loads, and cloud container runtimes for burst load

OPEN POLICY AGENT (GATEKEEPER)

- Enforce a baseline suite of policies for all cluster resources
- Policy examples:
  - Only allowed images to be pulled from specified sources
  - Deny external load balancers
  - Ensure container limits are specified and reasonable
  - Only allow ingresses in a specific domain and ensure unique ingresses
- Can also add information to new resources:
  - Automatically add necessary taints/tolerations to containers (Windows, special pools, etc.)
  - Copy financing codes from the namespace object to the sub-resources

OPEN POLICY AGENT (GATEKEEPER)

- Policies are written in a language called Rego
- Policies can refer to many data sources, including:
  - The object being created and/or updated
  - Other objects in the cluster
  - External data sources fed into the Open Policy Agent (e.g., users/group membership, etc.)

OPEN POLICY AGENT (GATEKEEPER)

Put some example OPA policies here to show how easy and powerful it is

ISTIO (SERVICE MESH)

- Automatic mutual TLS between services in the cluster
  - Handled by proxy sidecar, which intercepts all network traffic. **No changes to the app.**
- Observability: tracing, monitoring and logging of all network communication
- Enhanced routing rules, including:
  - Timeouts and retries
  - Per-instance routing (e.g, 90% of traffic to v1, 10% to v2)
  - Load balancing between instances
  - Circuit breaking

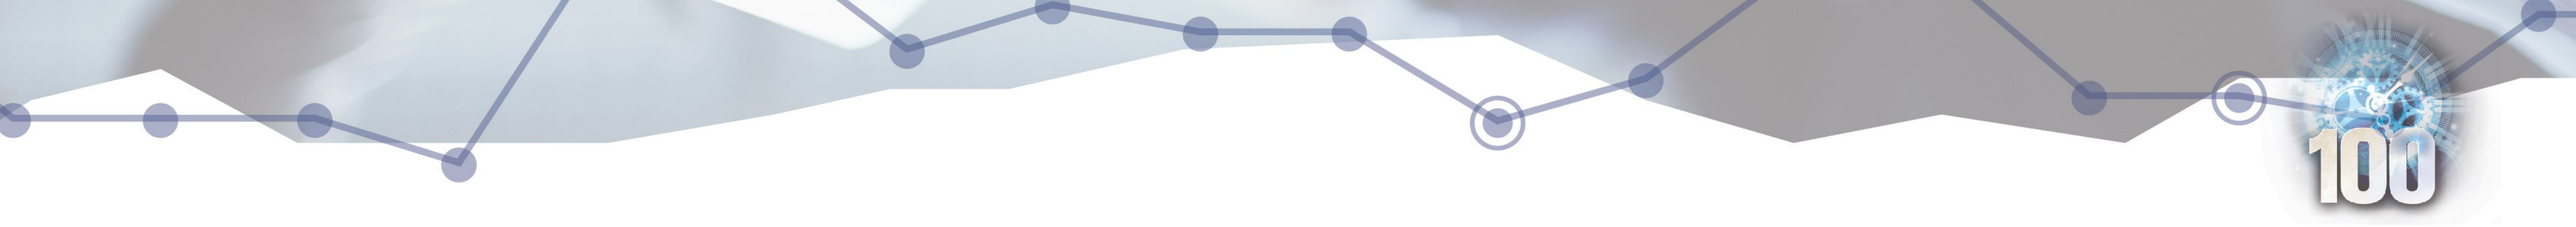

“SERVERLESS”

KNATIVE

Event-driven, sometimes also called serverless or functions as a service, is a computing execution model in which the infrastructure/provider dynamically manages the allocation of machine resources.

Brandan Burns (Kubernetes Founder)

SERVING: HOW YOUR CODE RECEIVES REQUESTS AND SCALES WITH THEM

- Request-driven compute runtime
- Scale-to-zero / scale out per load
- Multiple revisions of the same app
- Route traffic across revisions (powered by Istio)

BUILDING: HOW YOUR CODE IS BUILT AND PACKAGED AS A CONTAINER

- Pluggable model to automatically build containers from source
- Build in-cloud or on-cluster
- Templates available (buildpacks)
  - Don’t need to create a Dockerfile for your app
  - Don’t need domain-specific knowledge

EVENTING: HOW YOUR CODE IS TRIGGERED BY EVENTS AND EXECUTED

- Apps and functions consume and publish to event streams
- Multiple event sources available, examples:
  - Kafka
  - Cloud events
  - Webhooks
  - And more!
- Encourages asynchronous, loosely coupled architecture

WHY KNATIVE?

- Provider agnostic: your code runs inside of your Kubernetes cluster
  - Current solutions like AWS Lambda, Google Cloud Functions, Azure Functions have little interoperability between each other
- KNative, rather than developers, manages cluster resources:
  - Deployments, services, ingresses, routing rules, etc.
- KNative manages application scaling based on load (scale up and down)

WHAT ARE WE DOING?

- Onboarded 10+ developer teams to a unified workflow
  - Hybrid workloads: Linux and Windows
- Use Kubernetes to run our platform services
  - Source code management system, artifact repository, deployment tools, build tools
- We’re currently working with our IT security team to complete our security control profile
  - Isolation of workloads
  - Policies: Role Based Access Control, Network, Pod Security Policies

William Hearn

william.hearn@canada.ca

sylus

william_hearn

Questions?

Zachary Seguin

zachary.seguin@canada.ca

zachomedia

zachomedia