Aerospike on Kubernetes

Gangadharan SA
Sr. Engineering Manager
Adobe

Companies Struggle to Obtain a Complete View of Customer to Drive Personalized Experiences Across Channels

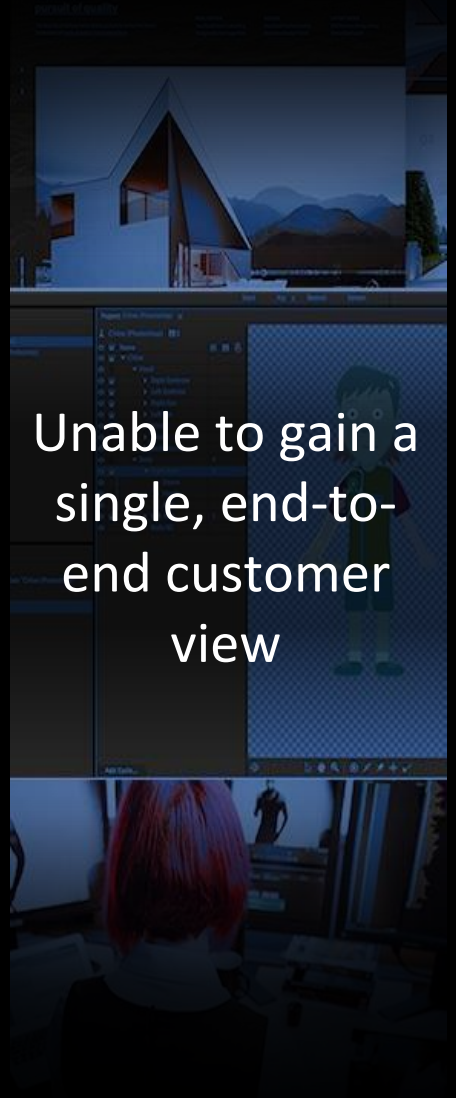

Unable to gain a single, end-to-end customer view

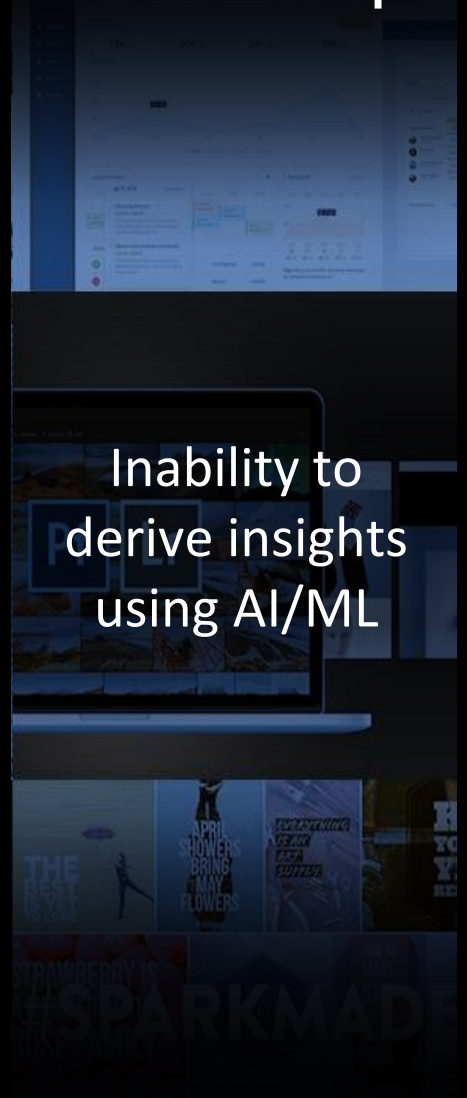

Inability to derive insights using AI/ML

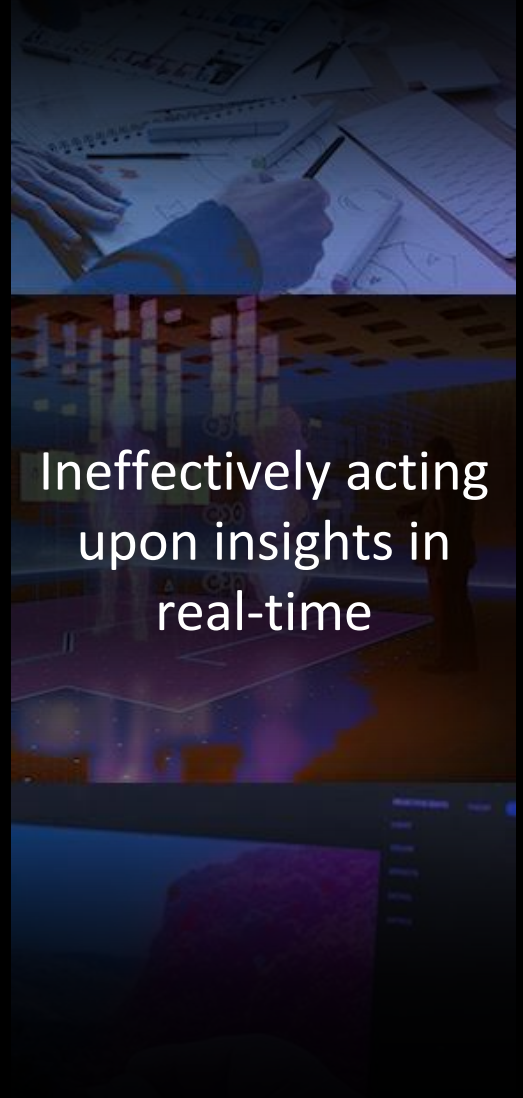

Ineffectively acting upon insights in real-time

Must be able to integrate with existing investments

Adobe Experience Platform

An open and extensible platform that stitches data across the enterprise powering a **real-time customer profile** enhanced through experience intelligence and governance, and activates content to deliver real-time personalized experiences

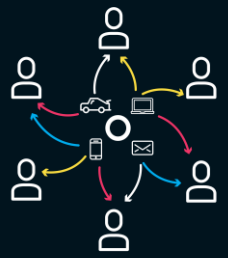

Real-time Customer Profiles

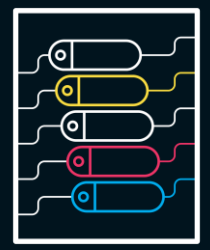

Data Management & Governance

AI & Machine Learning

Application Delivery

Open & Extensible

Motivations to build a K8S operator for Adobe Experience Platform

- **Scale operations to several datacenters & minimize operational costs**
- **Normalize deployment across cloud providers and physical DCs**
- **Different from existing Aerospike operator (as we understand it):**
  - Align with StatefulSets instead of Pods
  - Helm chart packaging for deployment
  - Cover all cloud types and self-owned datacenters
  - AZ fault tolerance
- **Gain experience with K8S Operators!**

Goals

- **Minimize manual management effort**
- **Retain performance qualities**
- **Cloud fault tolerance**
- **Cloud provider neutral**
- **Align stateful with stateless systems**
- **Convert operational knowledge into code**

Introducing Kubernetes

- **Container orchestrator**
- **Open-sourced by Google in 2014**
  - Based on 15 years of in-house experience
- **Highly extensible**
  - By cloud provider
  - By application developer
- **Helm Charts for packaging and deployments**
- **Spinnaker for workflow pipelines**

Kubernetes – Pod: “The Logical Host”

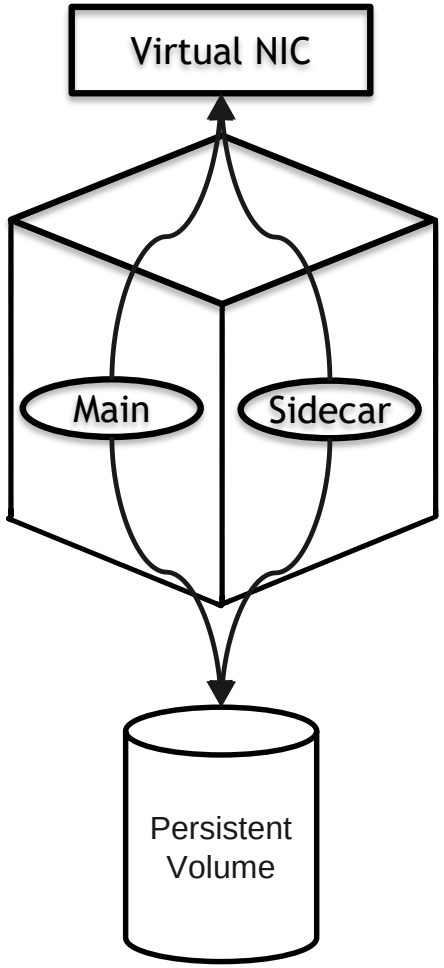

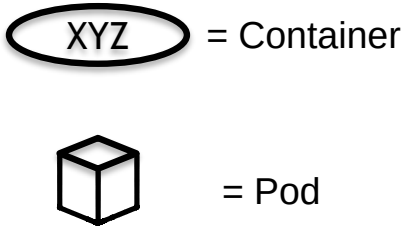

Kubernetes - Architecture

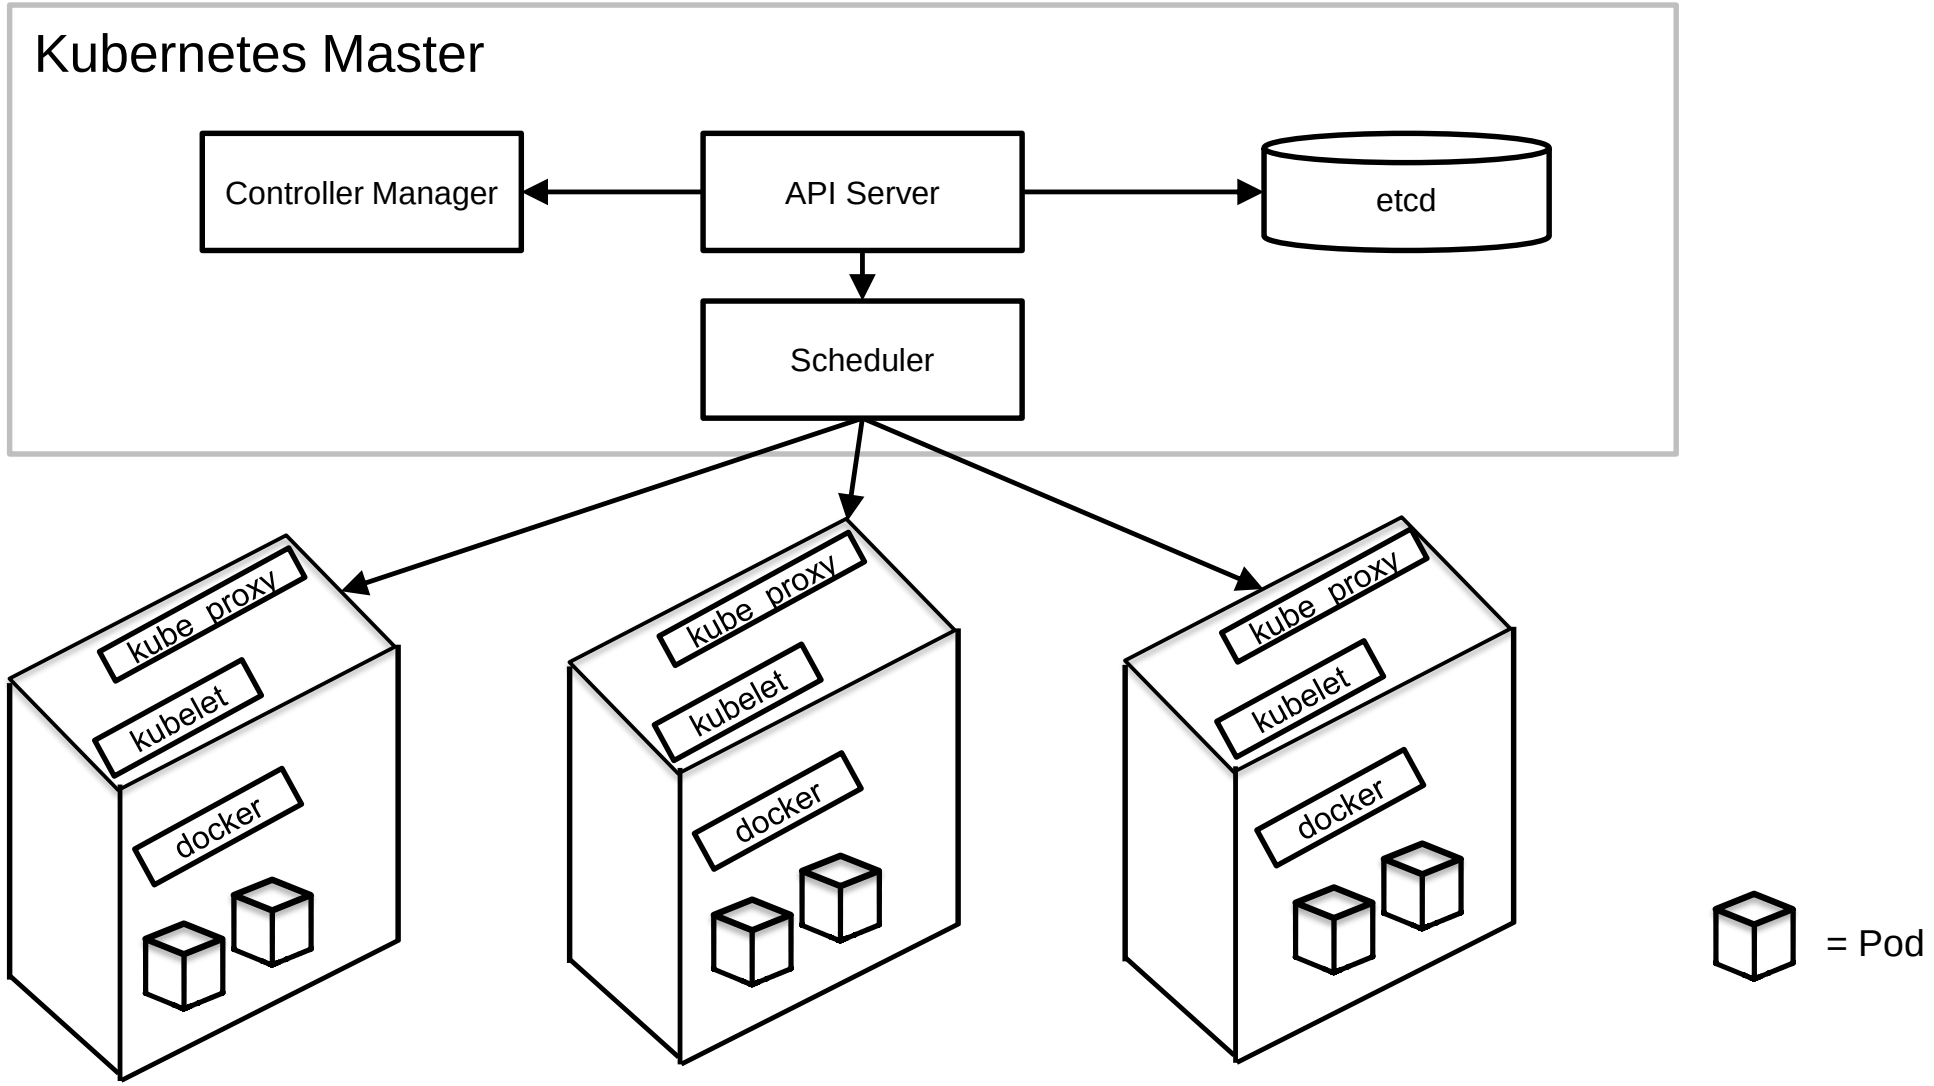

Aerospike State Considerations

- **System Metadata persistence**
- **Local SSD Block device**
- **Shadow device: Persistent NAS block device**

Kubernetes- PVs, PVCs & Storage Classes

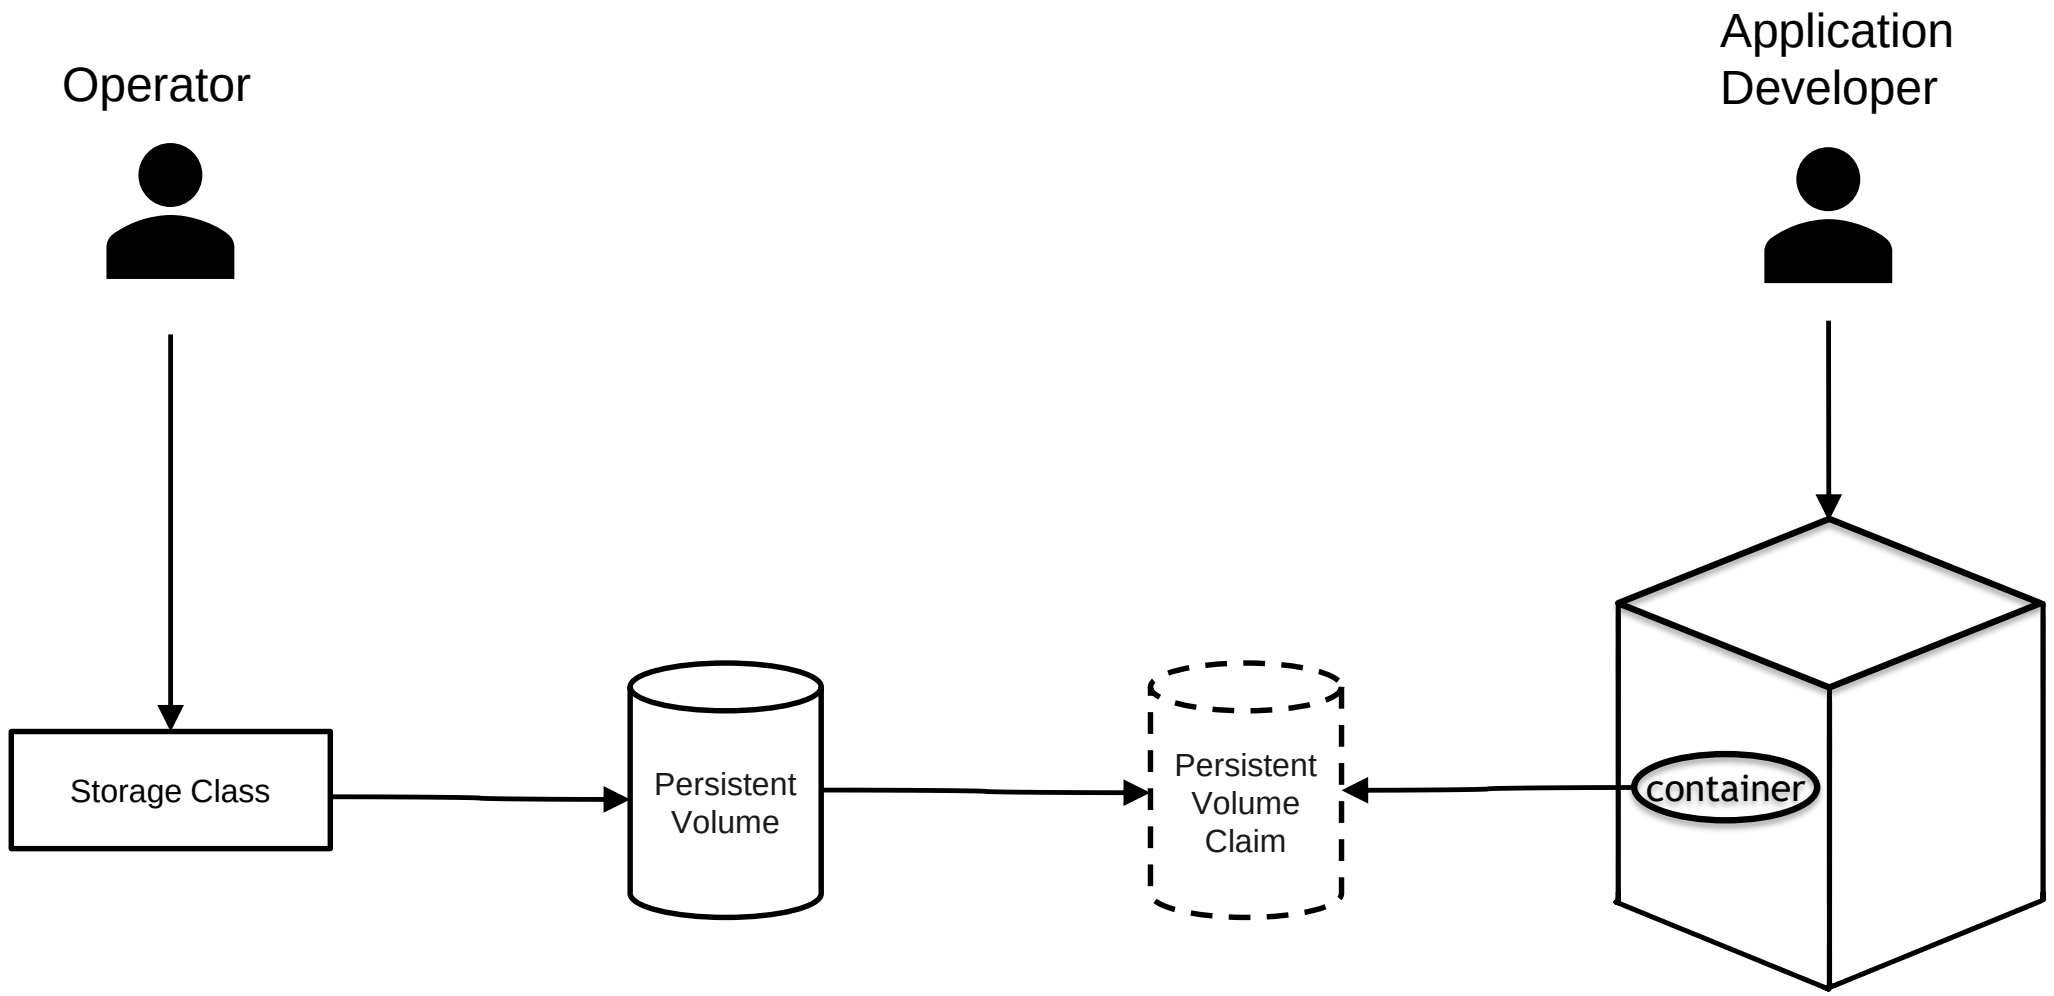

Kubernetes – Aerospike Persistent Volumes

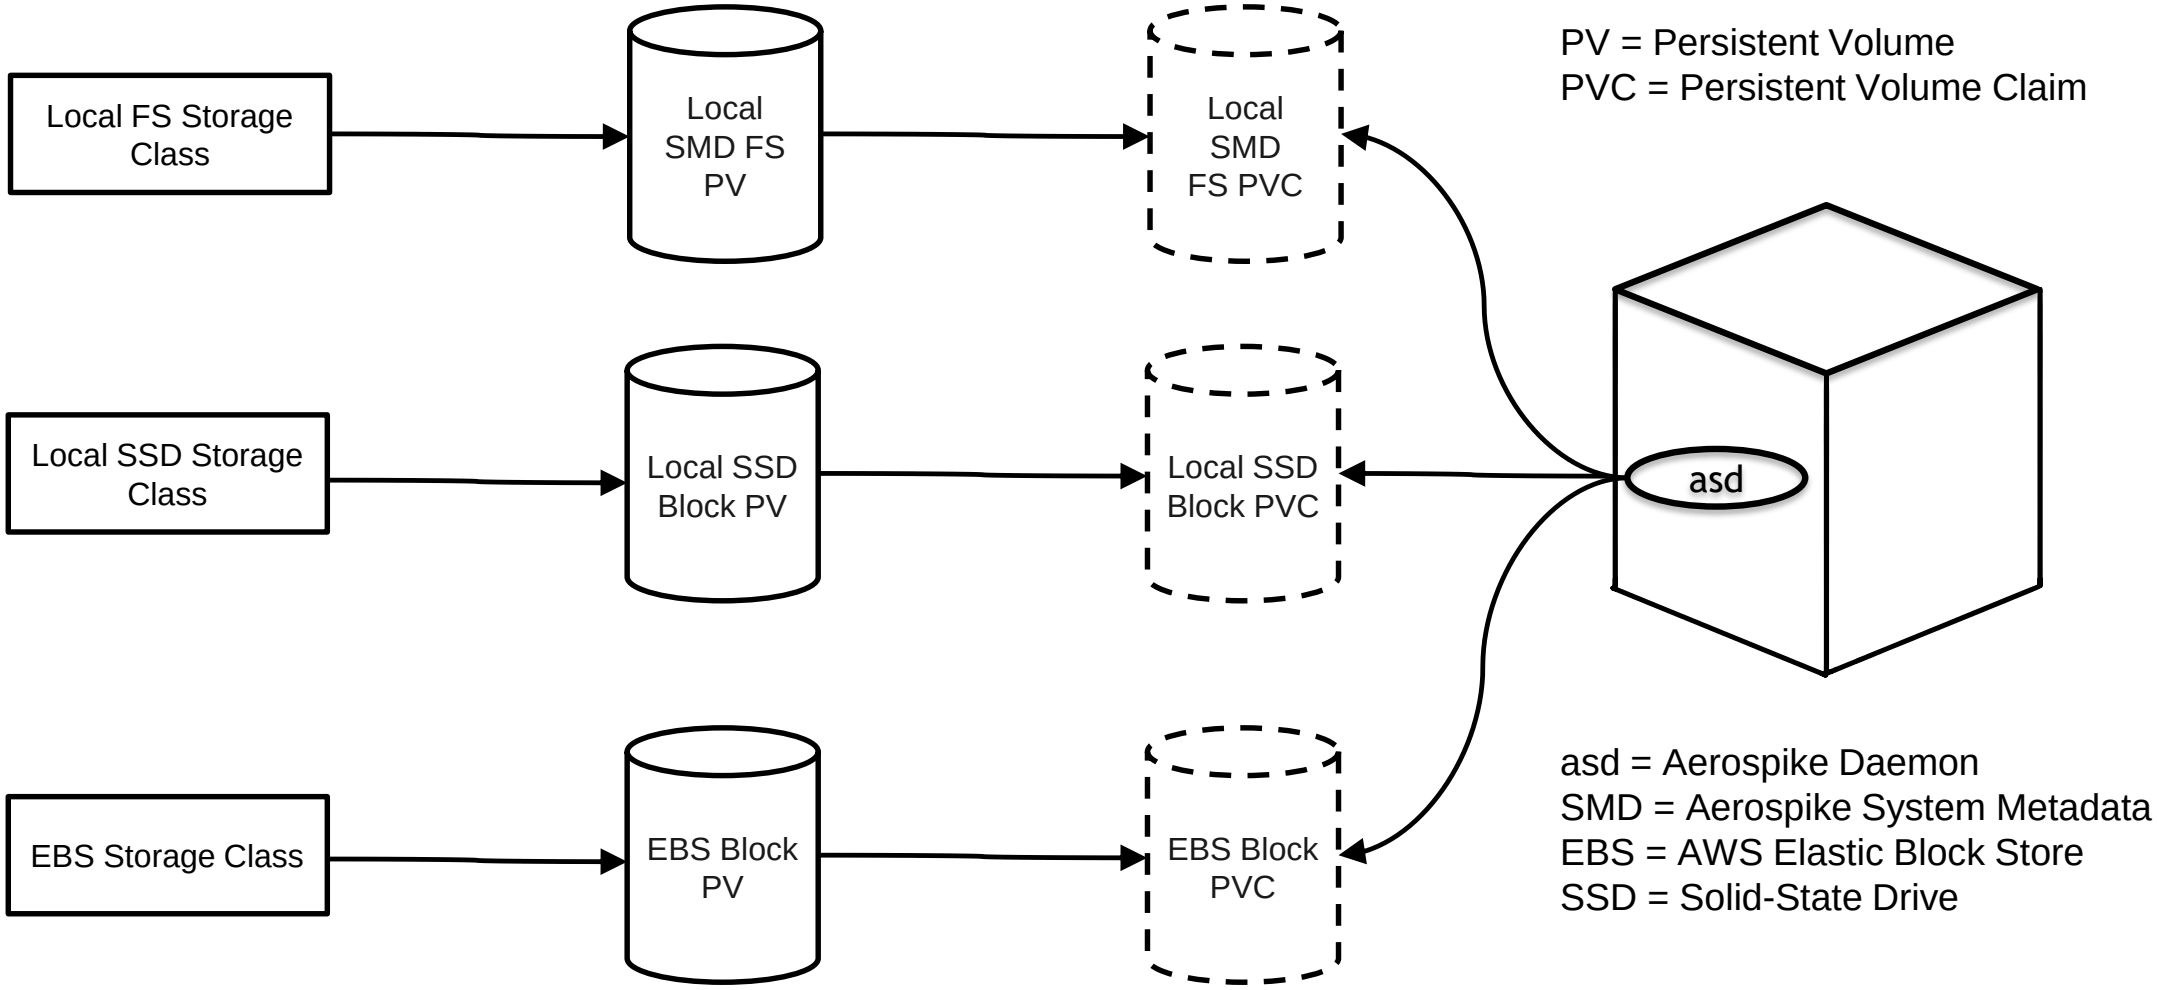

Adobe | AEROSPIKE SUMMIT '19

Aerospike State Considerations → K8S Features

- **System metadata persistence**
  - Local persistent volumes
- **Local SSD Block device**
  - Local persistent volumes
  - Raw block volume support
- **Shadow device: Persistent NAS block device**
  - Cloud Provider Storage classes
  - Raw block volume support

Aerospike Node Considerations

- **Consistent Aerospike node-id**
- **Local SSD Stickiness**
- **Diversify aerospike nodes across failure domains**

Kubernetes - StatefulSets

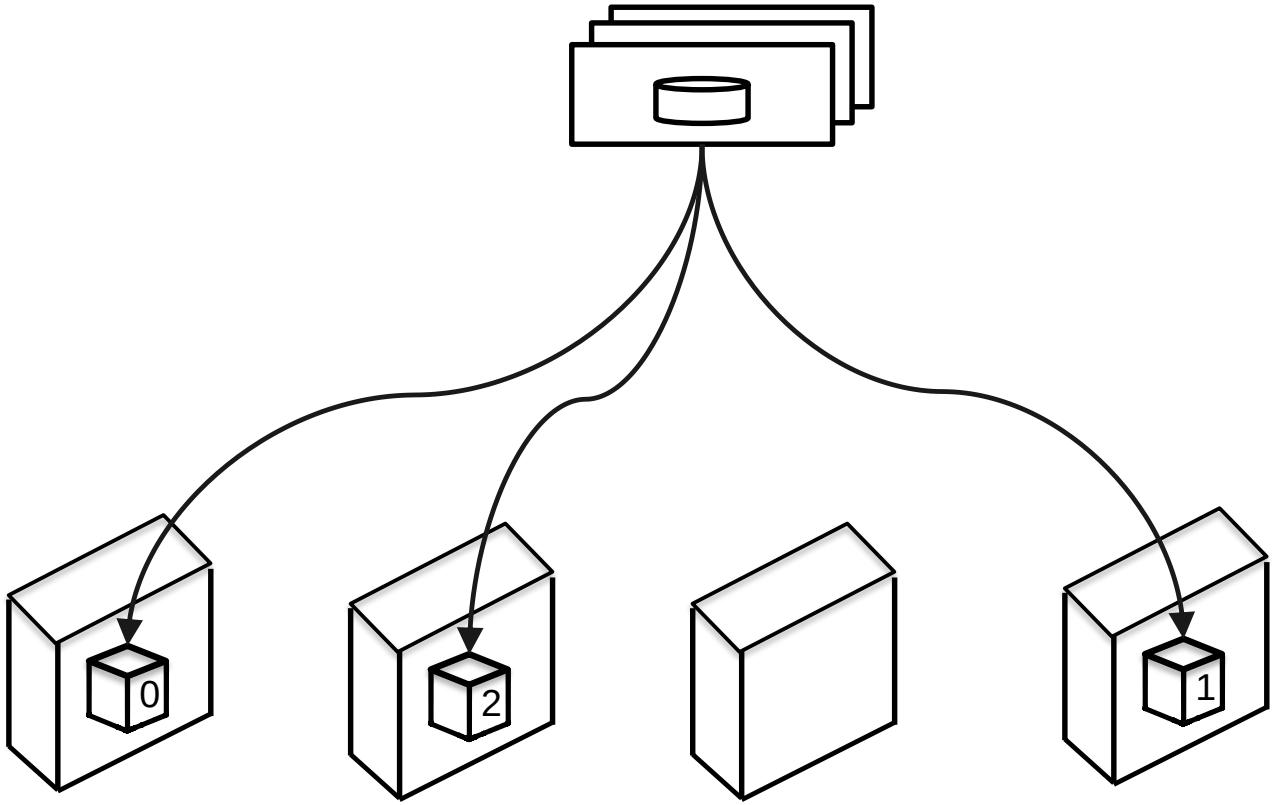

- Sticky pod ids
- Stable pod hostnames
- Deployment & Scaling guarantees
- Keep volumes intact

Kubernetes – AZ aware custom scheduler

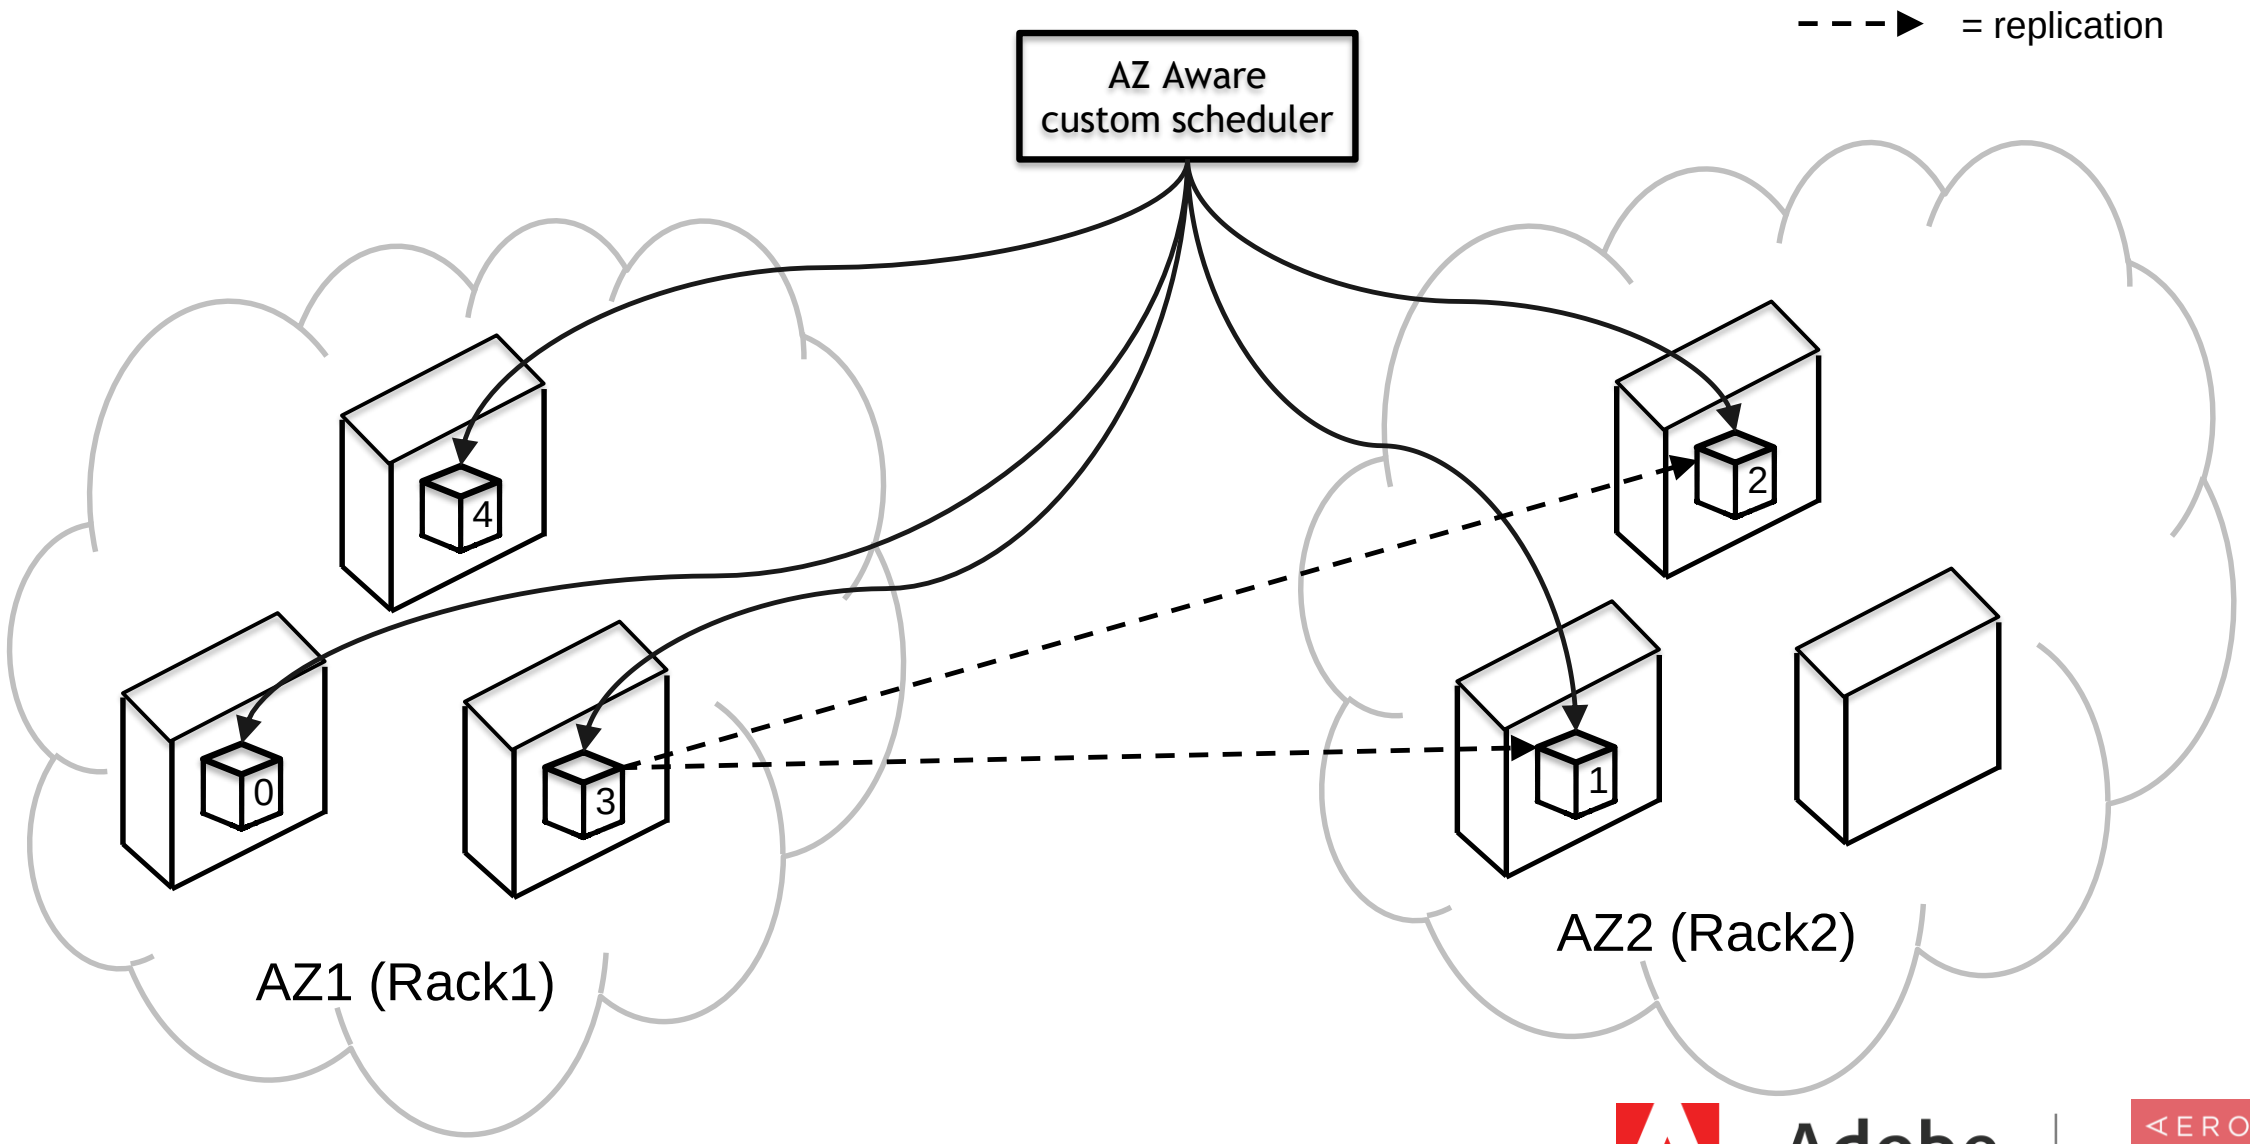

Adobe | AEROSPIKE SUMMIT '19

Aerospike Node Considerations → K8S Features

- **Consistent id**
  - StatefulSets
- **Local SSD Stickiness**
  - StatefulSets
- **Diversify aerospike nodes across failure domains**
  - Custom scheduler: Evenly distribute nodes across availability zones
  - Aerospike Rack ID as a function of availability zone
  - Node anti-affinity

Rolling Updates

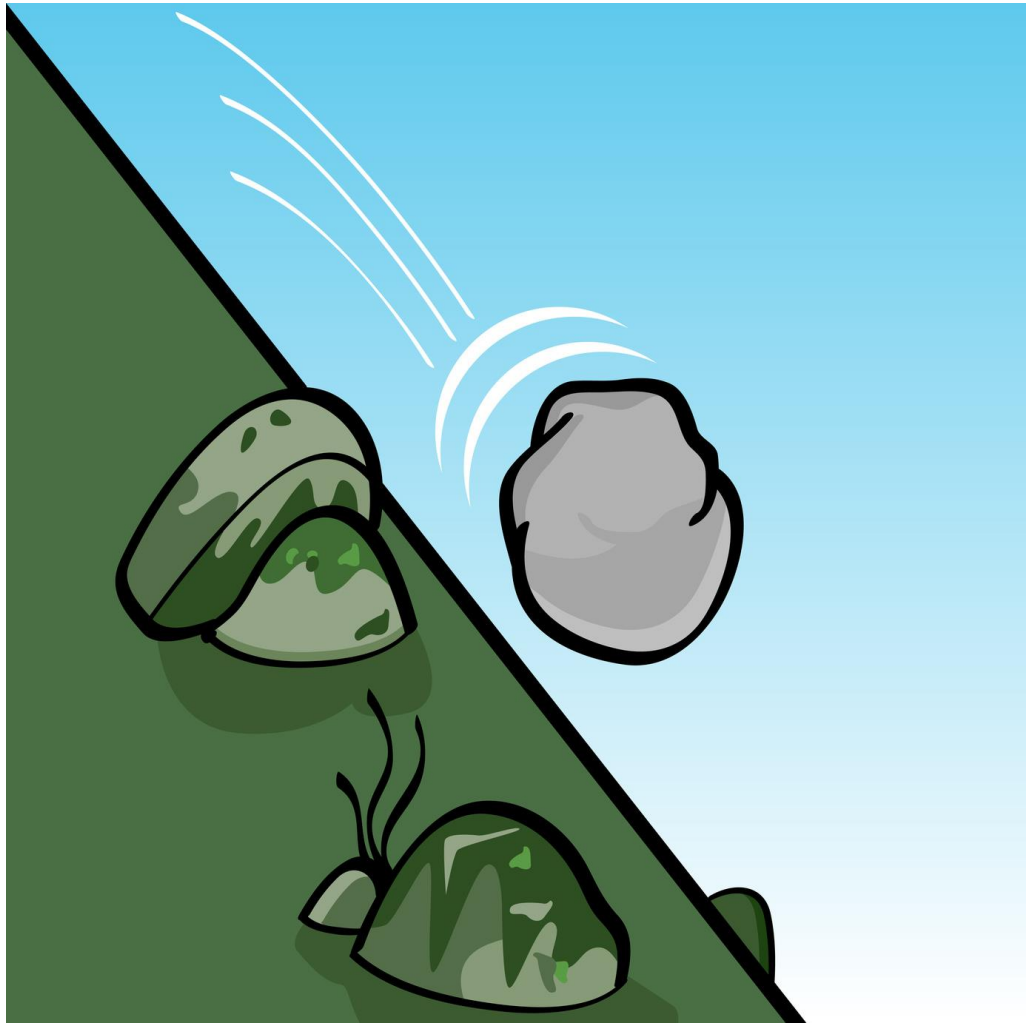

- StatefulSet update strategy
- One pod at a time
- Stop on failure
- Restore failed pod to earlier version

Scaling

- In ordinal order
- Manually delete volumes
- Pacing: Pause for rebalances & migrations

Monitoring

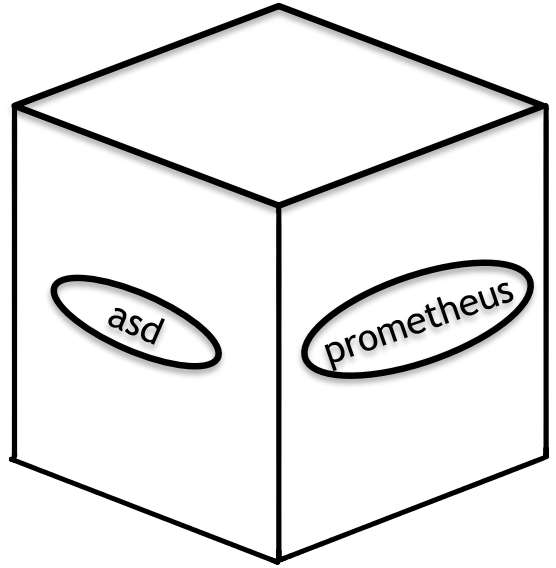

Backup / Restore

- **“Job” - Finite life-timed tasks**
- **Restore as Job**
- **“CronJobs” – Job factory**
- **Backup as CronJob & Job**
- **Prune old backups**
- **Support multiple backends – blob-stores and filesystem**

Taming the Jungle – Operator Pattern

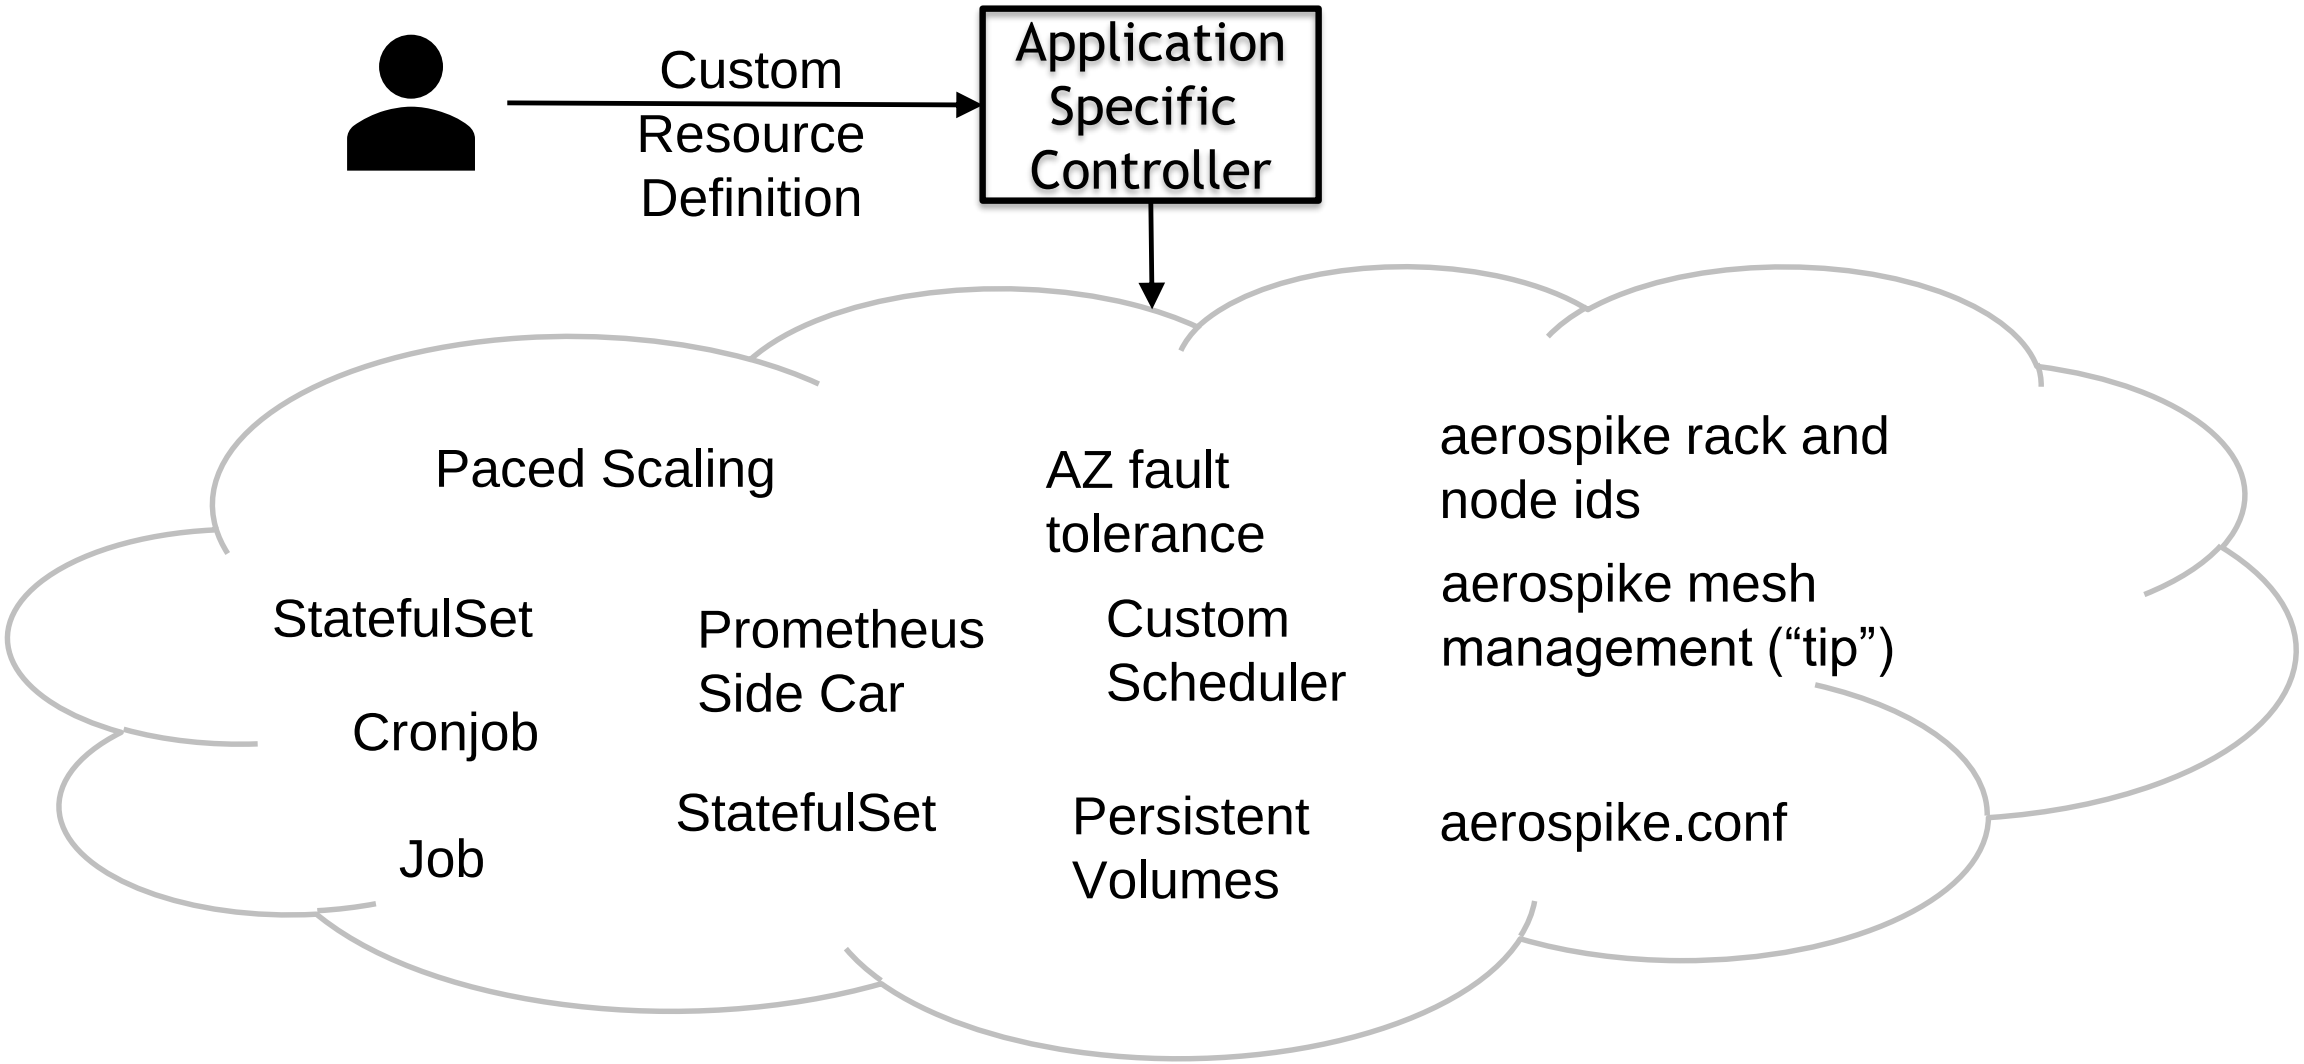

Adobe | AEROSPIKE SUMMIT ‘19

CRD + App. Specific Controller = Operator

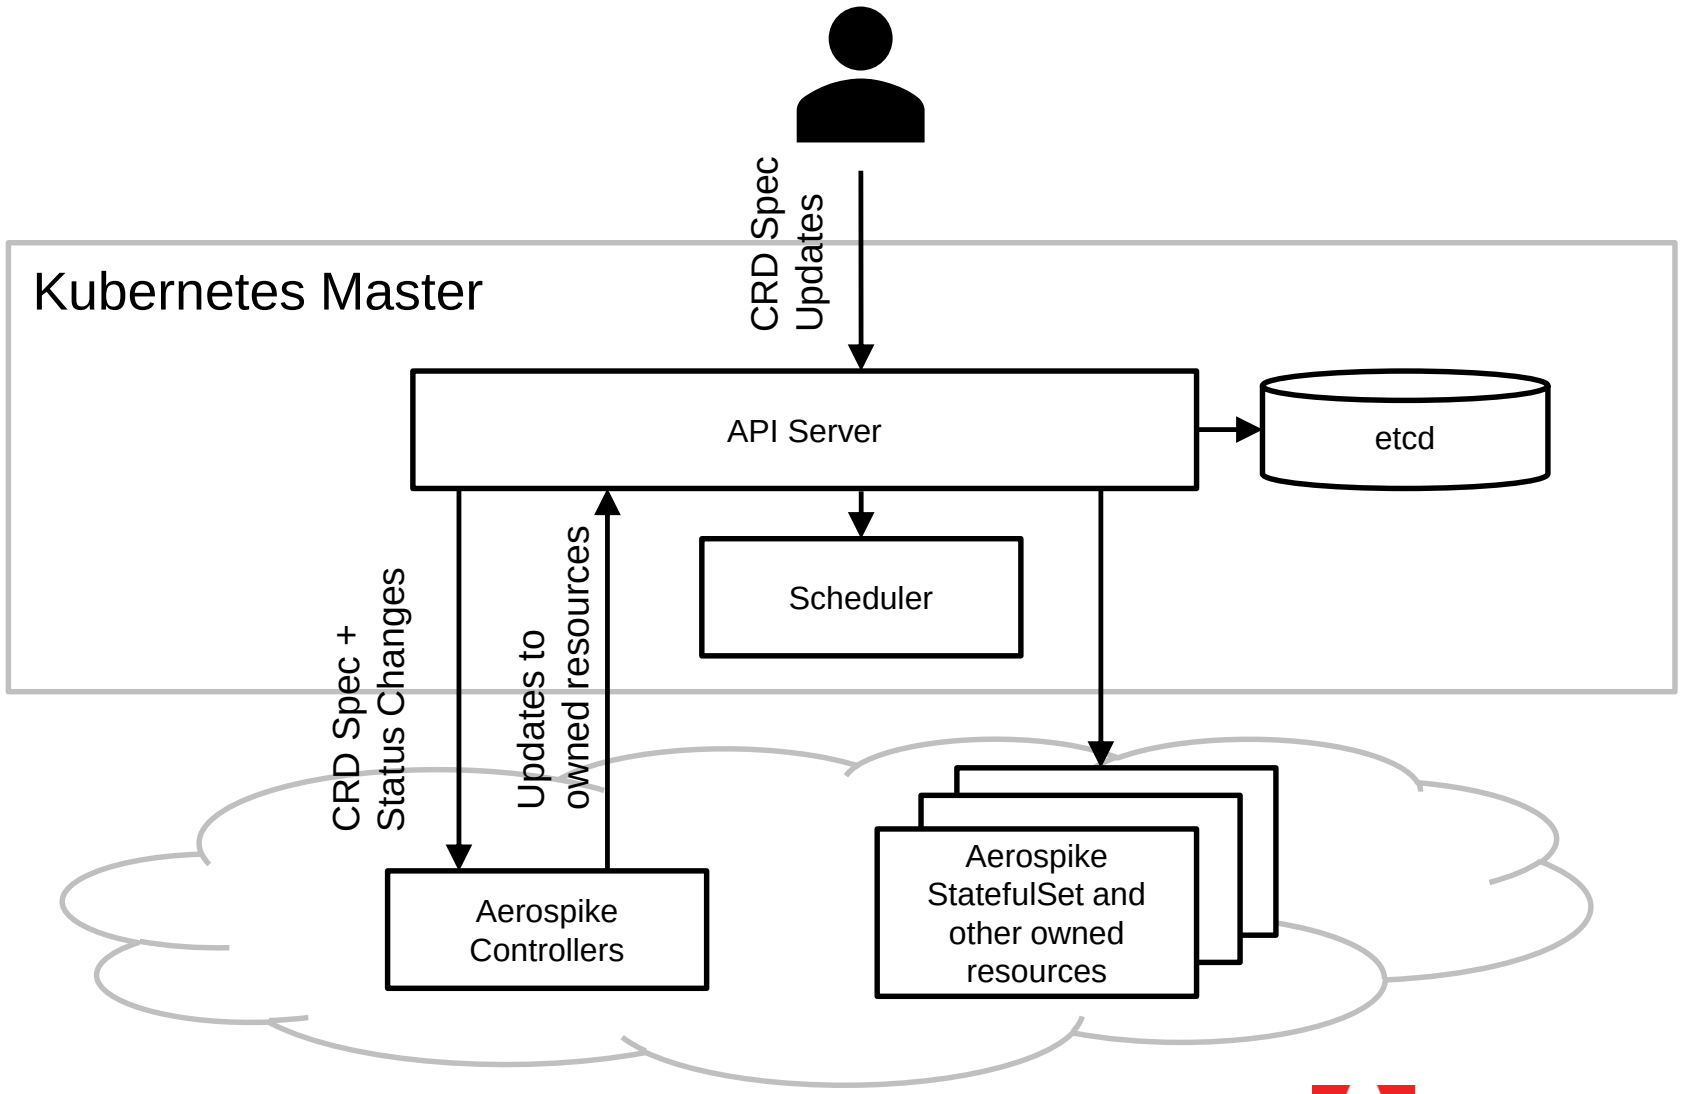

Introducing Helm

- RPM for K8S apps
- Encapsulate dependencies
- Manage updates
- Host on Private or Public servers
- Rollbacks

Spinnaker

- **Workflows and Pipelines**
- **Seamless support for K8S**
- **Pipelines to cut-over to new clusters**
- **Centralize coordination across multiple K8S based edges**

User Flow

- **Install Helm Chart**
  - Install Custom Resource Definitions
  - Deploy Controllers
- **Define Aerospike Cluster Resources**
- **Restore backups as needed**
- **Spinnaker to coordinate cluster cutovers etc.**

Next Steps

- **Planning to open source.**
- **Questions & further details: gangasa@adobe.com**
- **Looking for early testers.**
- **Contribute operational expertise to realize a turnkey Aerospike Operator!**
- **Many more features to think about, like:**
  - Auto-Scaling
  - Cloud Resource Management
  - Benchmark Suites

Thanks!

Adobe

AEROSPIKE SUMMIT '19