



Flash Memory Summit

Flash Memory Summit
Flash Storage + Kubernetes + CSI + Databases
Wednesday, May 11, 2022 - Virtual

Flash Storage + Kubernetes + CSI + Databases

Kamlesh Lad

Architect

Cloudcasa.io by Catalogic Software

Who am I ?




Flash Memory Summit

- **Kamlesh Lad**
 - Twitter: @gokam
 - LinkedIn: www.linkedin.com/in/kamlesh-lad-2019
- **Architect**
 - Over 20 years of experience in storage management software, specifically backup/restore and copy data management
- **CloudCasa by Catalogic Software (www.cloudcasa.io)**
 - Backup as a Service (SaaS).
 - Web UI, RBAC, fully hosted, swipe a credit card
 - Kubernetes backup, restore, and migration
 - Kubernetes and Cloud Security Scanning


- Why Postgres ?

- Enterprise class database management system
 - Comparable to expensive commercial databases like Oracle and SQL Server
- Robust feature set
 - First release in 1995 – long history
 - Large open-source development and community, free to use or modify
- Compatible with various platforms using all major languages and middleware.
- ANSI SQL standards conformance, ACID compliance
- Full support for client-server network architecture
- Standby server and high availability
- Object Oriented nature allows additional data types
 - JSON
 - Supports geographic objects



- Why Kubernetes ?

- Automated containerized environments
- Scaling up and down
- Auto provisioning of compute, network, and storage
- Light-weight VM, efficient with resources
- Infrastructure as code
- Portable across on-prem or cloud, allows easy migration of apps and infrastructure
- Multi-cloud
- Strong open source community (CNCF)



- Why Flash Storage Array?

- High performance, smaller footprint, energy and less heat.
- Resilient, no mechanical parts
- Snapshots and cloning
- Shared storage via SAN/NFS/iSCSI
- Performance and shared storage allows more applications to take advantage of flash storage.

Why Postgres + Kubernetes + CSI + Flash Storage

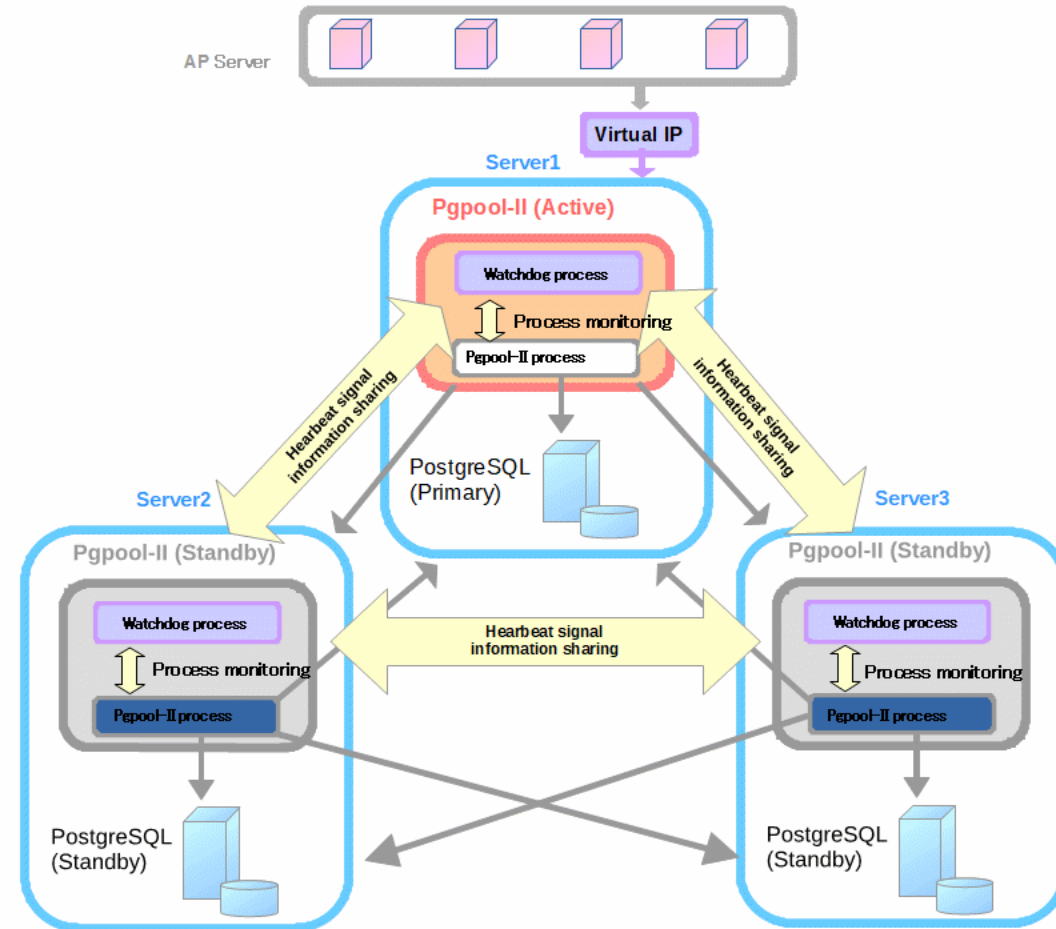
- Why Postgres + Kubernetes + CSI + Flash Storage
 - Postgres scaled and HA environments can get complex
 - Admins must manually provision compute, network, and storage
 - Kubernetes
 - Kubernetes Operators:
 - Kubernetes operators are software extensions to Kubernetes
 - An Operator is an application-specific controller that extends the Kubernetes API to create, configure, and manage instances of complex stateful applications on behalf of a Kubernetes user
 - CSI + Flash Storage
 - Kubernetes Container Storage Interface (CSI) allows storage suppliers to expose their products in containerized applications as persistent storage
 - Expose advanced features using portable standards

Postgres Kubernetes Operator



Flash Memory Summit

- Automatically provisions, modifies, and monitors a PostgreSQL cluster
- Monitors PostgreSQL cluster manifests — nodes, security, fault tolerance, and site topology — and makes necessary adjustments
- Provides automated domain knowledge to correctly scale, upgrade, and reconfigure while protecting against data loss or unavailability
- Application-specific operational knowledge encoded into software that leverages the powerful Kubernetes abstractions to run and manage the application correctly





Which Postgres Operator ?

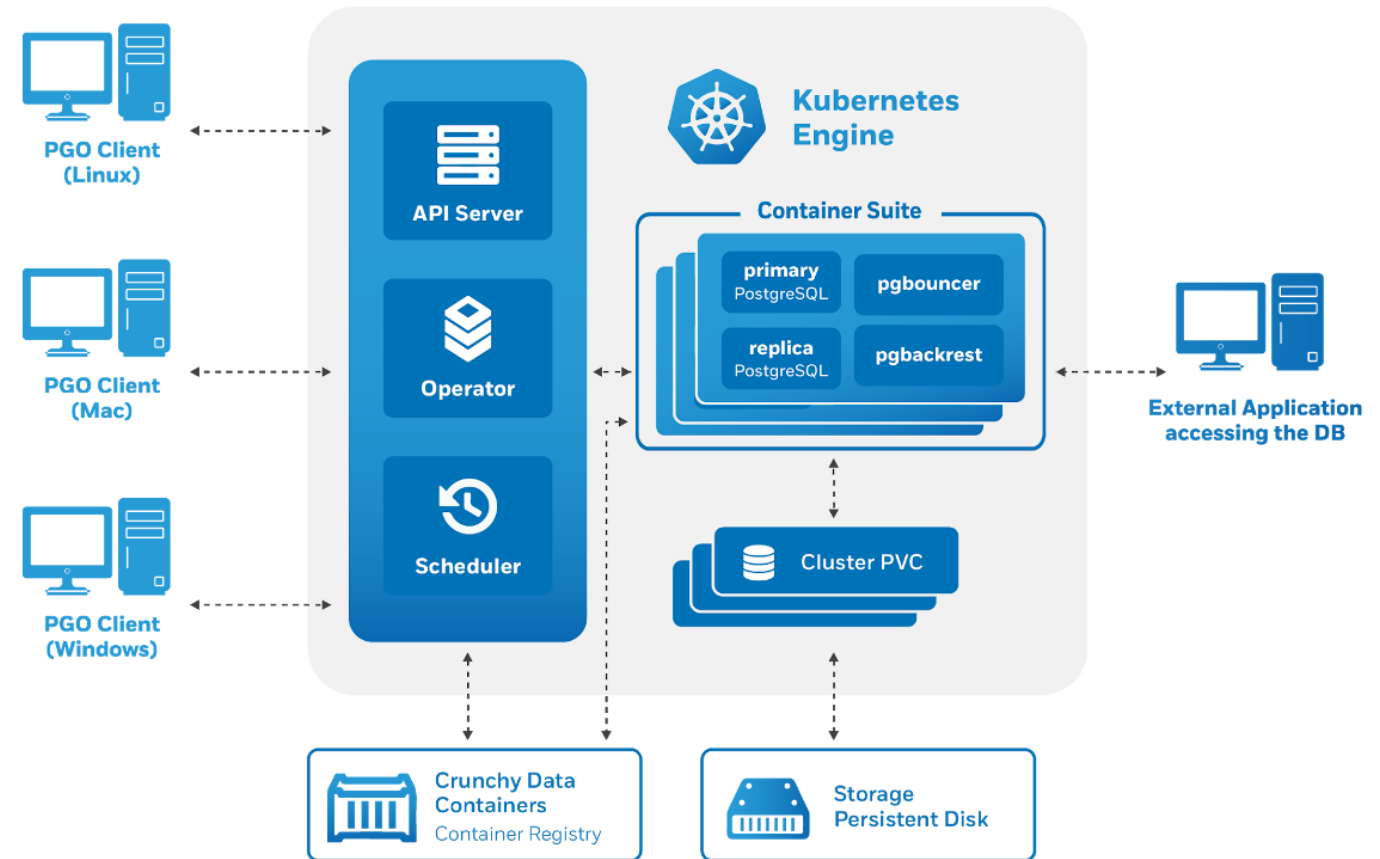
- Several Postgres Operators, each one with it's own advantages.
 - Crunchy Data Postgres Operator
 - Zalando Postgres Operator
 - KubeDB Postgres Operator
 - StackGres



Crunchy Data PostGres Operator

- Crunchy Data Postgres Operator

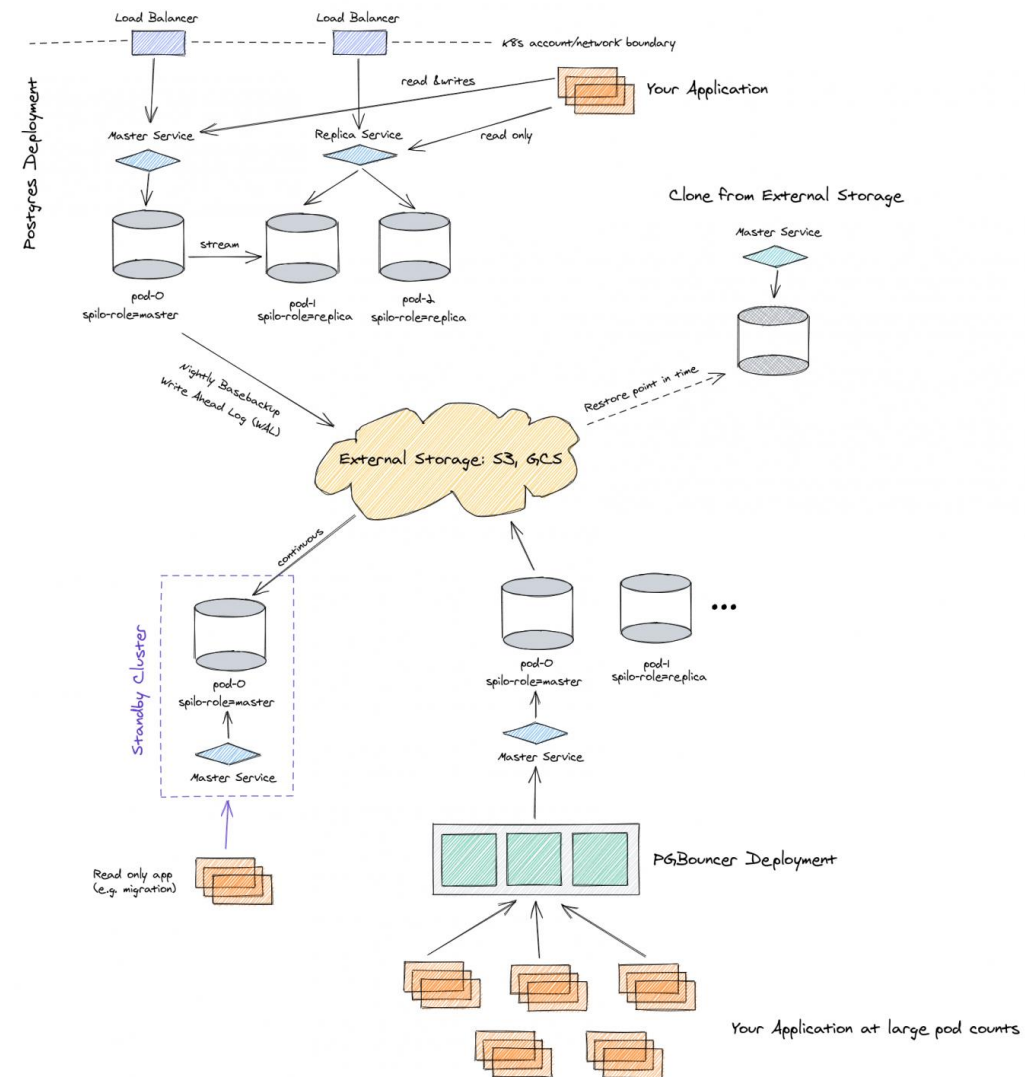
- Allows persistent storage configuration and works primarily with Dynamic StorageClasses, HostPath and the Network File Systems (NFS)
- Enables horizontal scaling for added redundancy and high availability
- Support for Kubernetes Custom Resource Definitions (CRD)
- Cloning allows non-disruptive rolling updates
- Integration with other components (Backup, Audit, WebUI for PostGres, etc)
- Backup to Local Storage, S3, etc
- Supported Platforms:
 - Kubernetes 1.20+, Openshift 4.6+, Rancher, GKE, EKS, AKS, VMware Tanzu



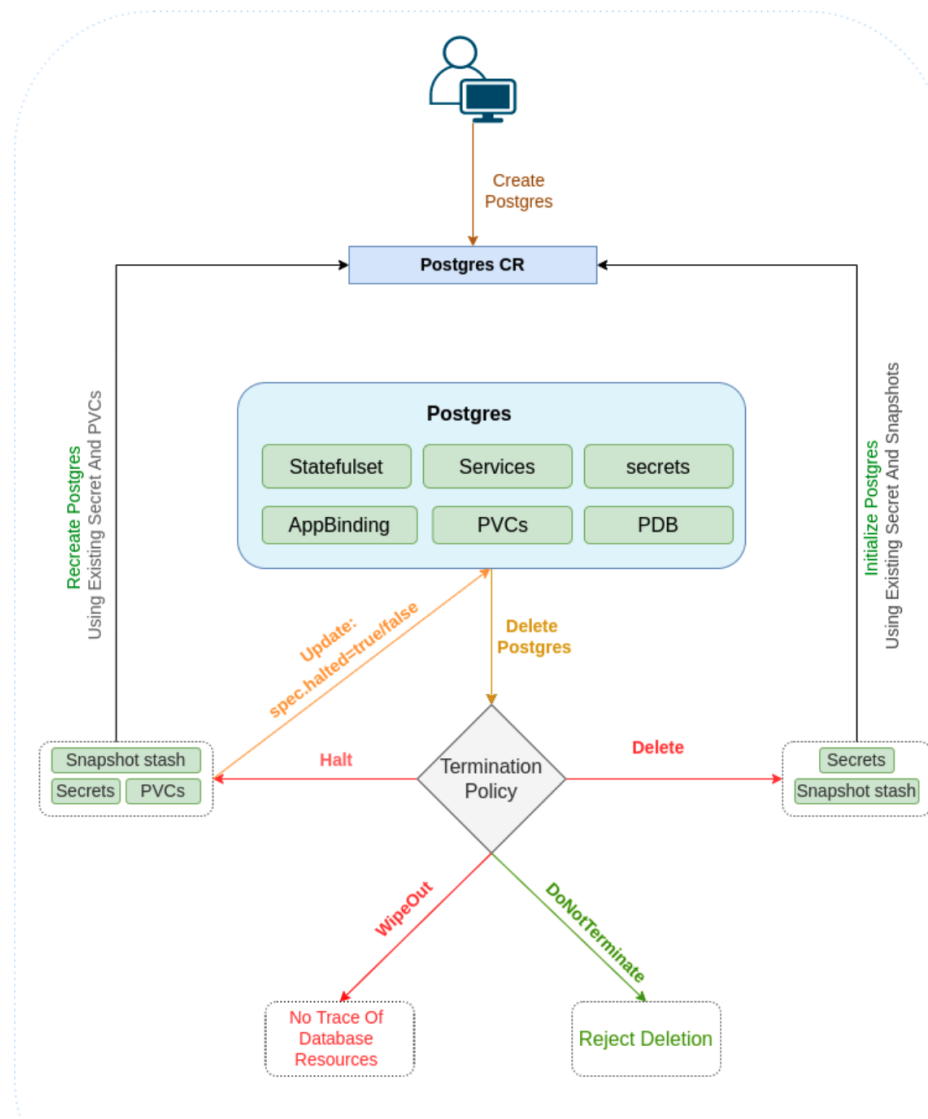
Zalando Postgres Operator

- Integrate into CI/CD pipelines automatically without accessing Kubernetes API
- Managed with Custom Resources
- StatefulSets compatible with sidecar containers
- Simplifies deploying Patroni-powered clusters on Kubernetes — rolling updates, provisioning, and cleaning up Postgres clusters
- Advanced customization for deployments
- Works in conjunction with Patroni which handles the high availability Postgres cluster bootstrapping
- 1500+ stars on GitHub and over 90 contributors

Zalando Postgres Operator : Supported Setups



- One solution for multiple databases, Operator supports PostgreSQL, MySQL, Redis, Memcached, MongoDB, and Elasticsearch
- Lacks customization
- Unique features, ability to restore accidentally deleted databases
- Kubectl plugin, manage databases with well known kubectl
- Well documented, lots of examples
- PostgreSQL operators use PersistentVolumeClaims (PVC) to provision disks for database instances, improving performance
- Additional features are only available in commercial offering such as backups, connection pooling, snapshots, etc)



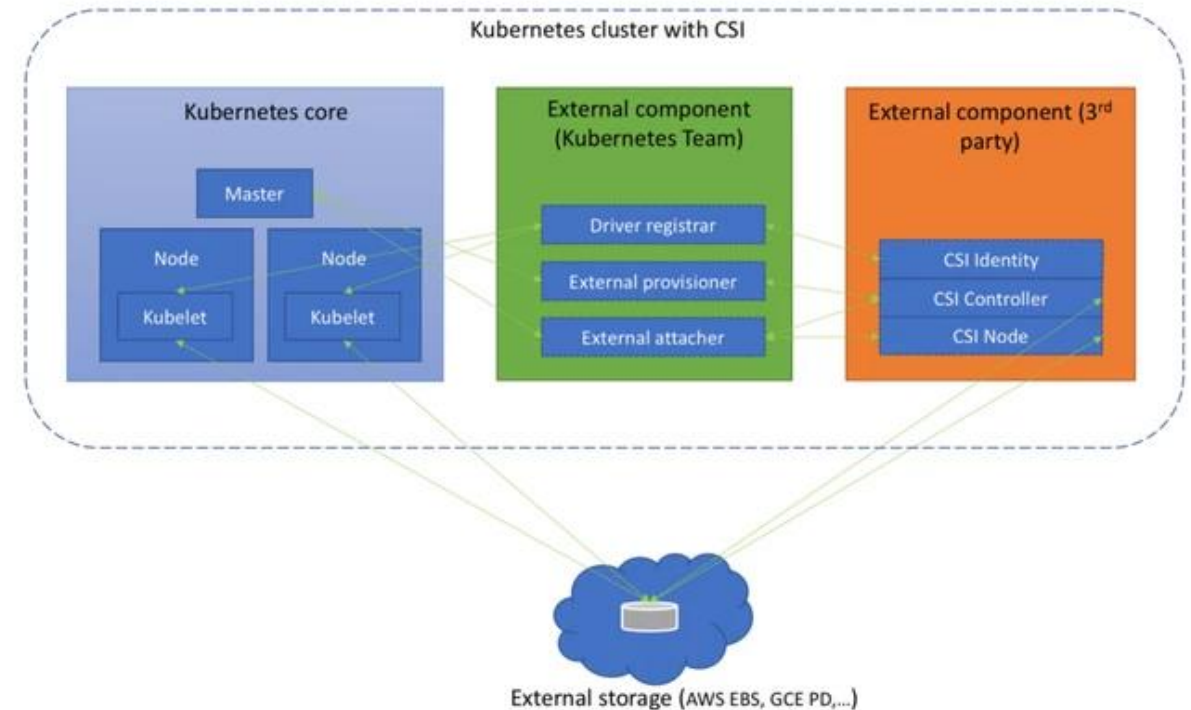
-
- The diagram illustrates the Stackgres architecture, which is designed to be Kubernetes-native. It is divided into several functional layers and components:
- External Interactions:**
 - CRDs (Custom Resource Definitions)** and **REST API USERS** interact with the **Stackgres Controller** via **kubectl**.
 - The **Stackgres Controller** interacts with the **KUBERNETES API**.
 - CLIENTS RW** (Read-Write) and **CLIENTS RO** (Read-Only) connect to the **SERVICE PRIMARY** and **SERVICE REPLICAS** via port **5432**.
 - Control Plane:**
 - The **Stackgres Controller** manages the **REST API** and the **Stackgres UI**.
 - The **Stackgres UI** also interacts with the **KUBERNETES API**.
 - Data Plane:**
 - The **DISTRIBUTED LOG SERVER** (containing **FLUENTD** and **POSTGRES + TIMESCALE**) is connected to a **PV (Persistent Volume)**.
 - The **SERVICE PRIMARY** and **SERVICE REPLICAS** are connected to the **Stackgres Controller** and the **DISTRIBUTED LOG SERVER**.
 - The **SERVICE REPLICAS** are connected to the **Stackgres Controller** and the **DISTRIBUTED LOG SERVER**.
 - The **SERVICE REPLICAS** are connected to the **Stackgres Controller** and the **DISTRIBUTED LOG SERVER**.
 - Pods and Storage:**
 - There are three **PODs**: **POD-0** (role = primary), **POD-1** (role = replica), and **POD-2** (role = replica).
 - Each pod is connected to a **PV (Persistent Volume)**.
 - Dashed lines indicate data replication between the primary pod and the replica pods.

Kubernetes CSI



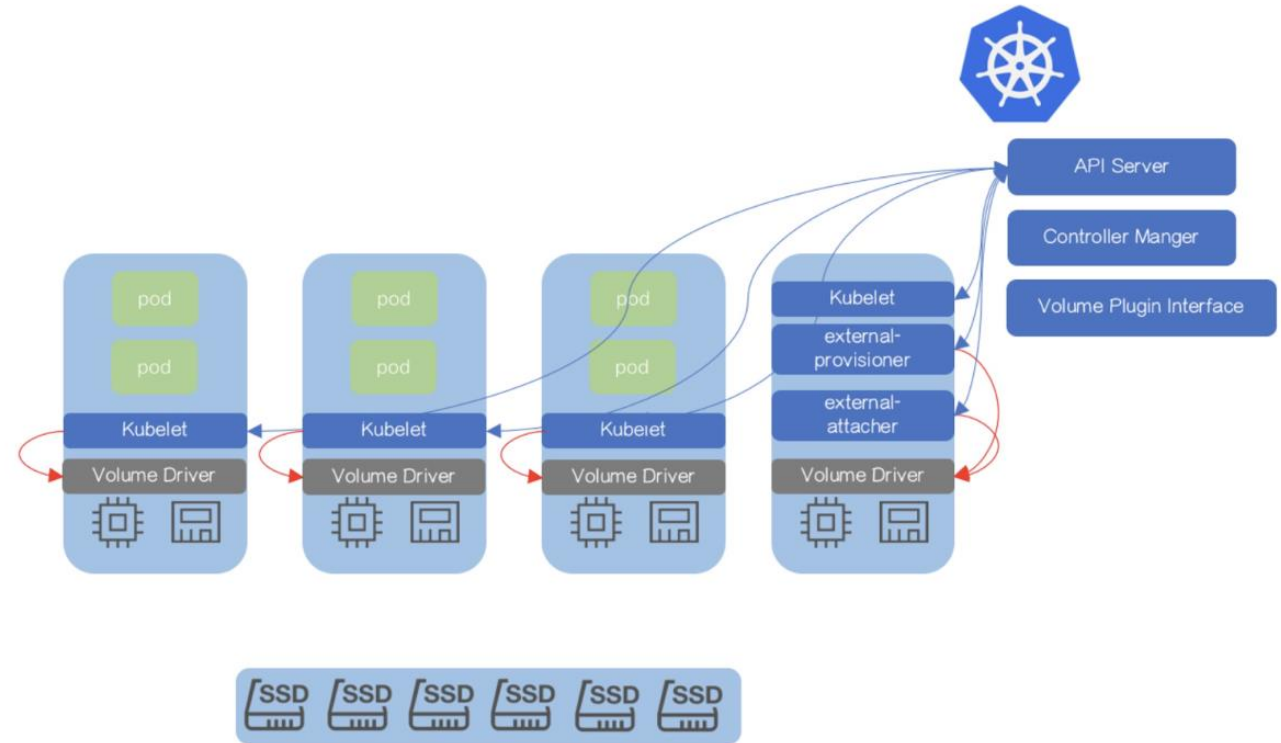
Flash Memory Summit

- CSI is about storage operations within container native environments
- CSI is an API and common model to address storage as an industry standard
- Abstraction layer to manage storage volumes without having to worry about varying storage internals or specifics
- Simple and compatible with broad range of storage technology
 - Block Storage
 - File Storage
 - Raw Devices
 - Local storage
 - Networked Storage



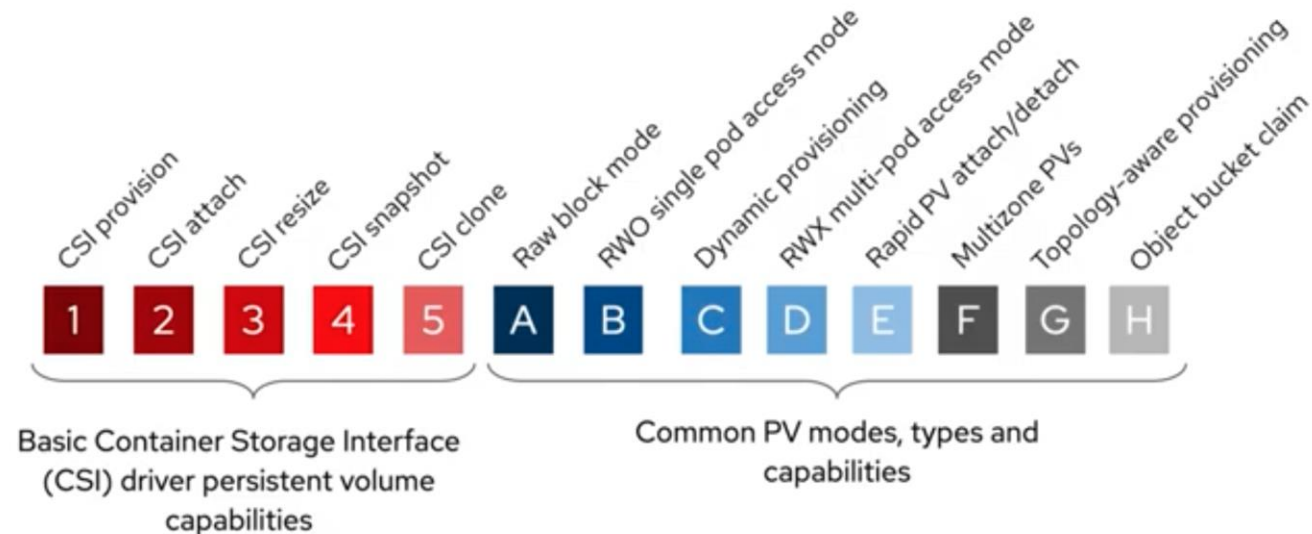
Advantages Kubernetes + Flash Storage + CSI

- Flash Storage
 - Kubernetes Container Storage Interface (CSI) allows storage suppliers to expose their products in containerized applications as persistent storage
 - CSI can expose features in Kubernetes native manner
 - Provisioning/Reclaiming
 - Snapshots
 - Cloning
 - Replication(future)
 - etc



- CSI allows storage vendors to expose their specific features as Persistent Volumes in Kubernetes

Persistent volume capabilities



- Postgres + Kubernetes + CSI + Flash Storage
 - Powerful, enterprise grade database
 - Kubernetes provides scalable portable cloud environment
 - Postgres Operators provide application specific automation for managing large HA environments
 - Kubernetes with CSI and Flash Storage provides performance, instant snaps + cloning

