

STORAGE DEVELOPER CONFERENCE



Fremont, CA  
September 12-15, 2022

*BY Developers FOR Developers*



A **SNIA** Event

# Container Data Backup: Going beyond CSI



**Rakesh Jain**

IBM

STSM and Researcher, SODA TOC Co-Chair



**Sanil Kumar D**

Huawei

Chief Architect, Head DM(India), SODA TOC Co-Chair

# Container, Kubernetes, Cloud, Cloud Native, Data Lake...



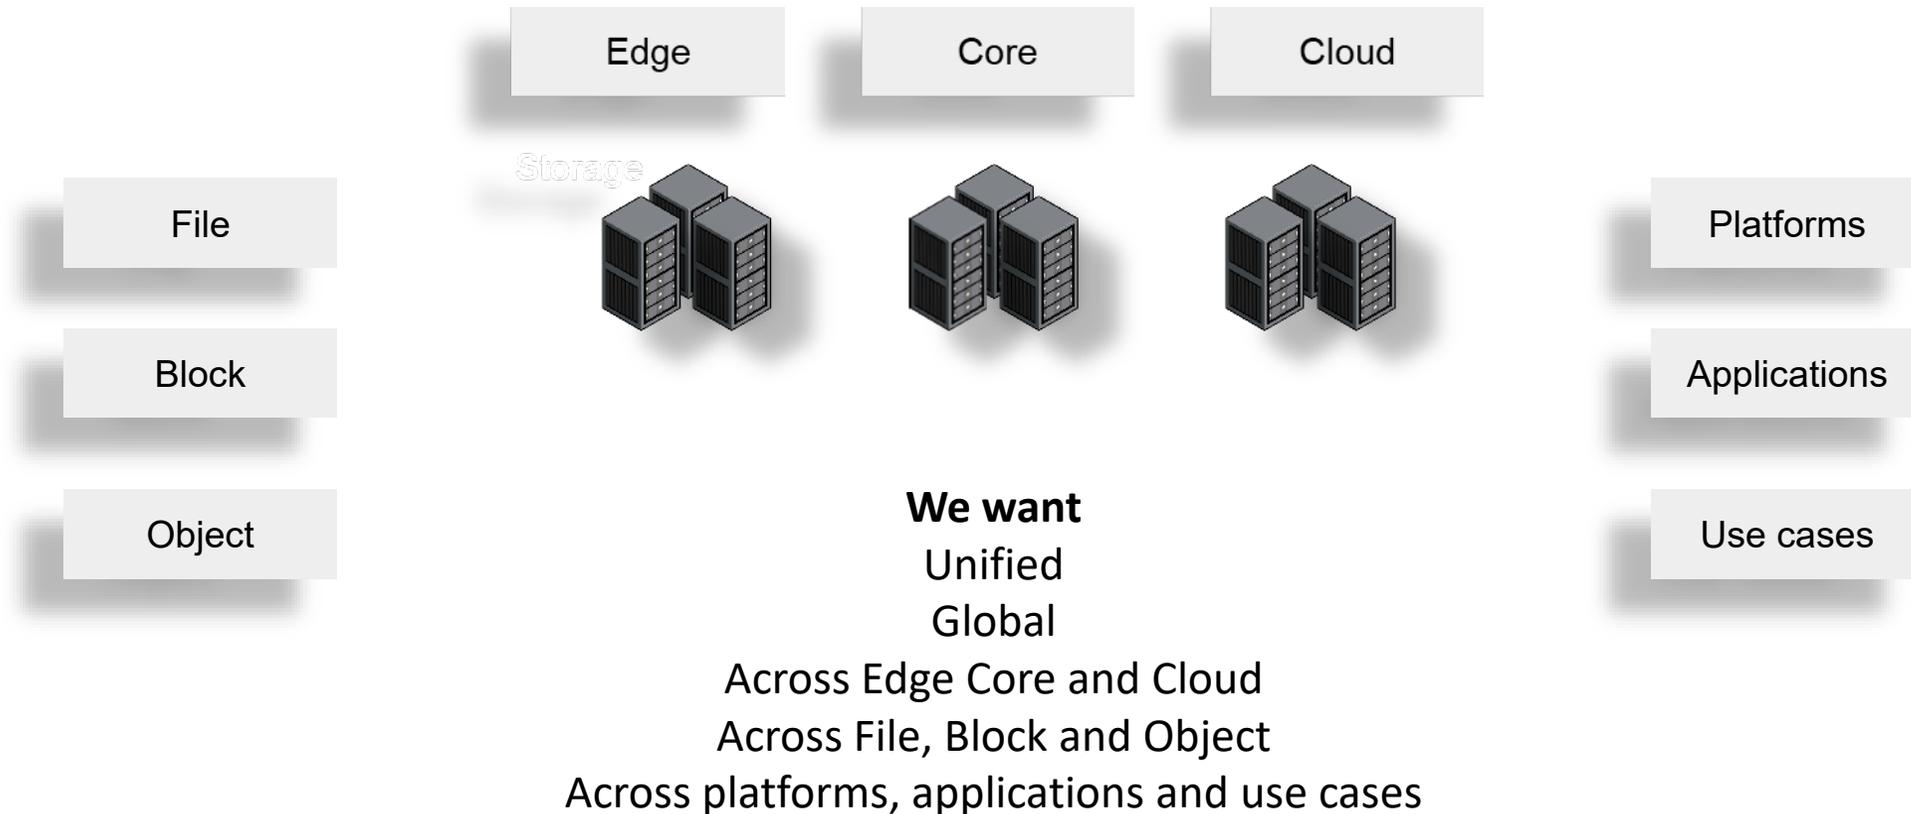
# Setting the context

*Container Application Deployment is growing, hence the container storage deployments!*



**Connected.  
Anywhere**

# Data Management Across



# STaaS to CSTaaS

(Storage as a Service to Container Storage as a Service)

## SODA Data & Storage Survey 2021<sup>#0</sup>

- Platform use pivoting to **container-based environments**
- Over the last 5 years there has been a seemingly **rapid transition from VMs to containers**.
- Kubernetes deployments came top (in fact **top 3 positions - K8S cloud, Hybrid K8S, K8S on prem**)
- Organizations are using **Container Storage for real use cases deployments**

## Other Industry Reports / Surveys

- **Application Container** Market CAGR **26.5%** (2019-25)<sup>#1</sup>
- **CaaS** CAGR **35%** (2021-26)<sup>#2</sup>
- Installed base of **container instances** CAGR **62.1%**(2019-23)<sup>#3</sup>
- **84%** using Container in **Production**<sup>#4</sup>

Container centric storage solutions | Hybrid Cloud Data Management for DR | Momentum towards CSTaaS

# Solutions Trend

- **Logical solutions to provide end to end data management**
- With Storage Boxes + Features
- Outside box features for Data Management, O&M, ...
- Third party storage support
- Container Deployment Support
- Multi-Cloud support
- Application Aware Data Management
- Cross Vendor Product Ecosystems

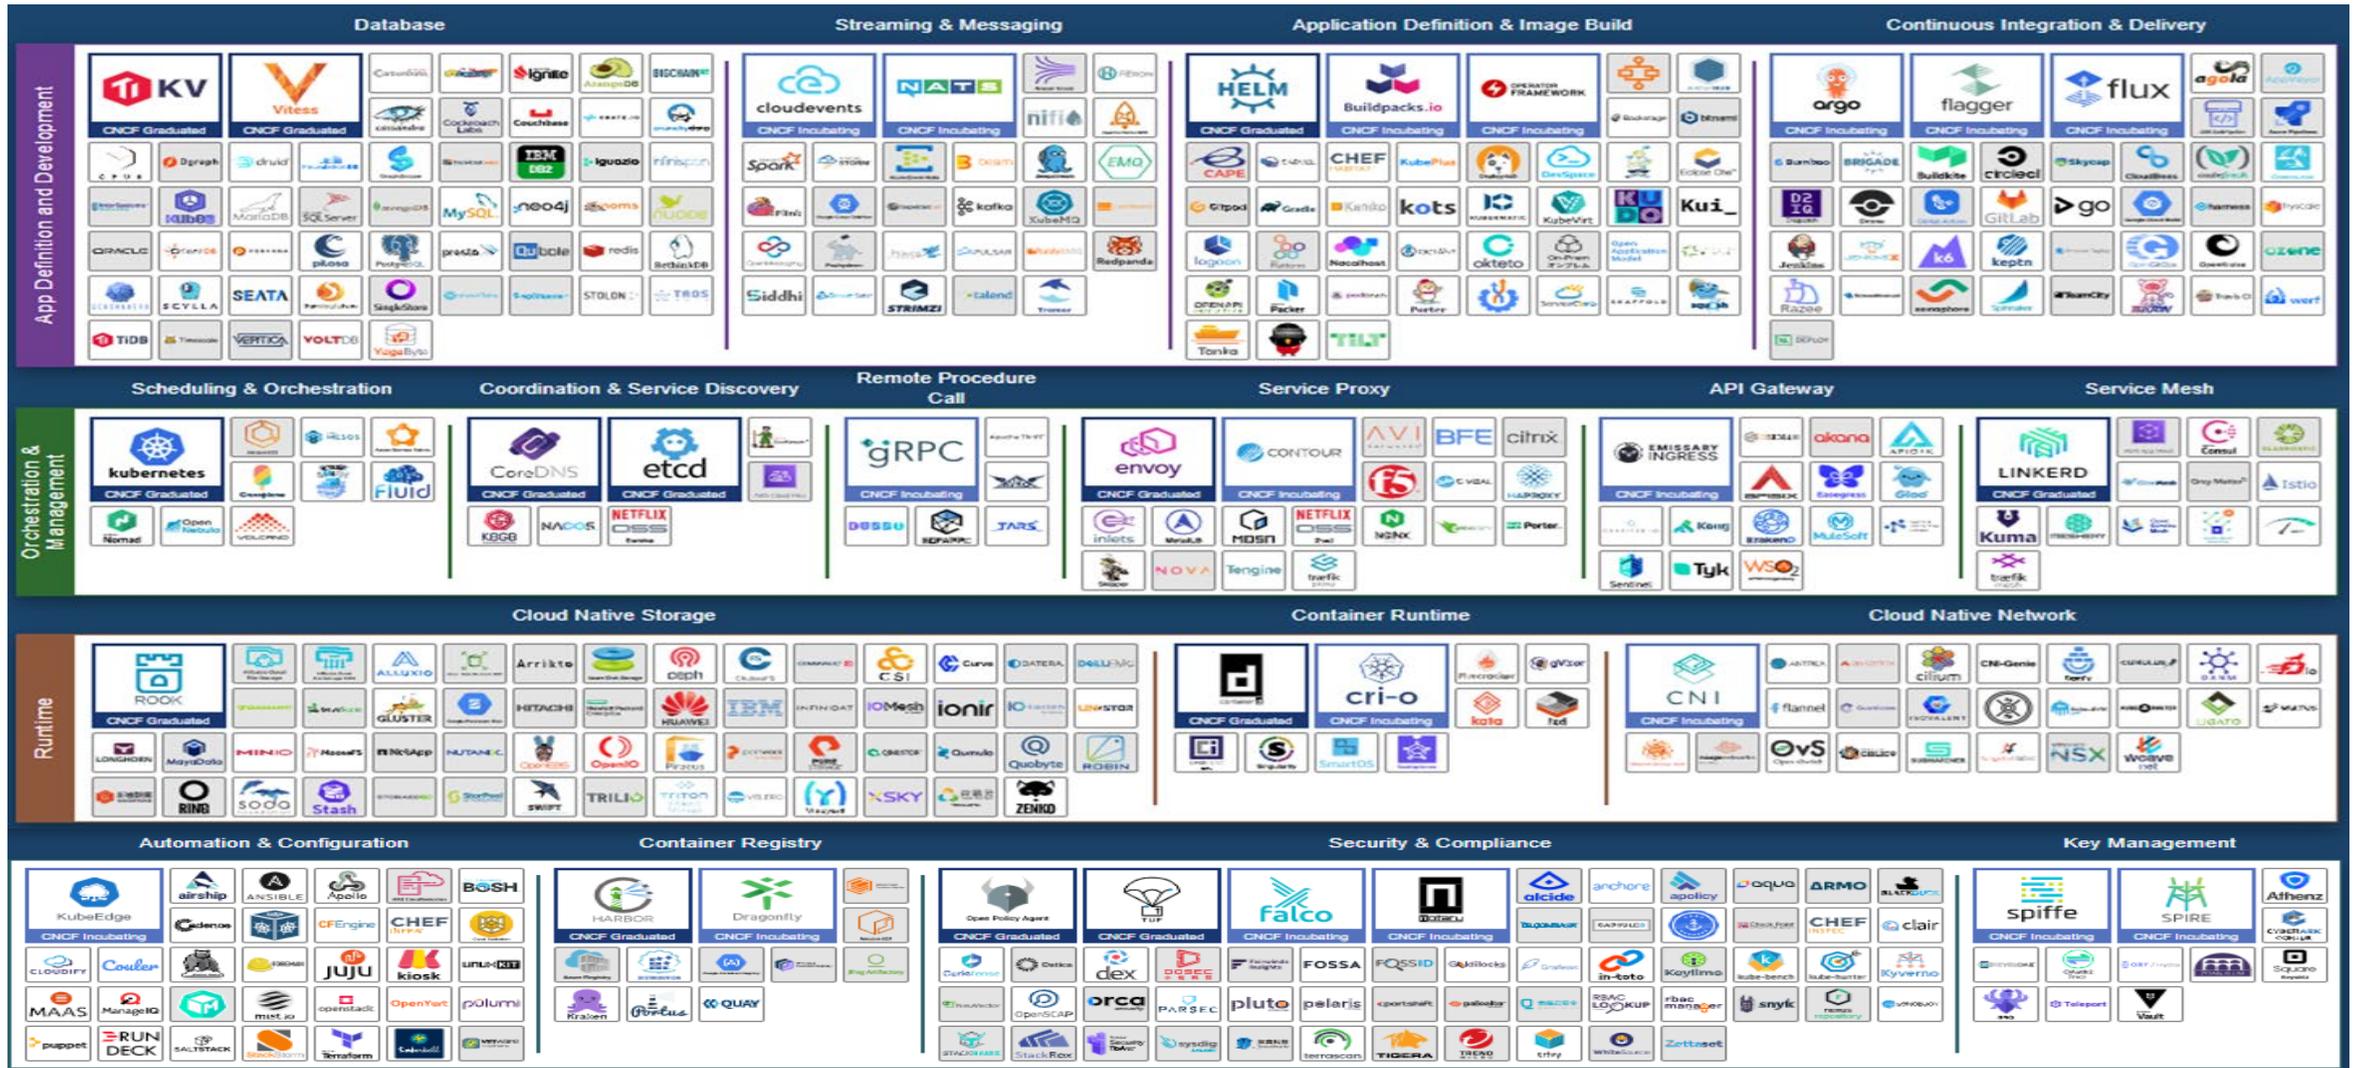
## A few examples:

- NetApp Astra : Application aware data management service built for Kubernetes
- VMware Tanzu & VMware Velero : enables you to build, run and manage modern apps on any cloud, back up and restore your Kubernetes cluster resources and persistent volumes
- Pure Service Orchestrator: Storage as a service for containers; gives developers the agility of public cloud with the reliability and security of on-premises infrastructure
- Product and open source ecosystem across VMware Tanzu, Velero, NetApp / other storage vendor solutions for end to end container data management

*Note: There are similar other products from other vendors and above are just examples for reference.*

**Not limited to Storage Box Features**  
**Best utilize existing open source / own products to build the solution**  
**Application Aware data management & Container Storage push this trend**

# Container Landscape is growing exponentially!



**Hence, Container Data Management and Container Storage are critical!**

# Cloud Native Storage Landscape and Trends

*What is happening and what it indicates...?*

# Revisit the terms - Cloud Native Storage, Container Storage, CSI...

## Container Storage

### Cloud Native Storage(CNS)

Software Defined Storage - that is API driven and customers can auto-provision. A storage solution that is secure, performant and scalable to application demands.

Container centric storage solutions are Cloud Native Storage solutions and more.

Declarative API (GitOps), Auto-healing and 1000x in terms of Volume Churn. Scale up and down.

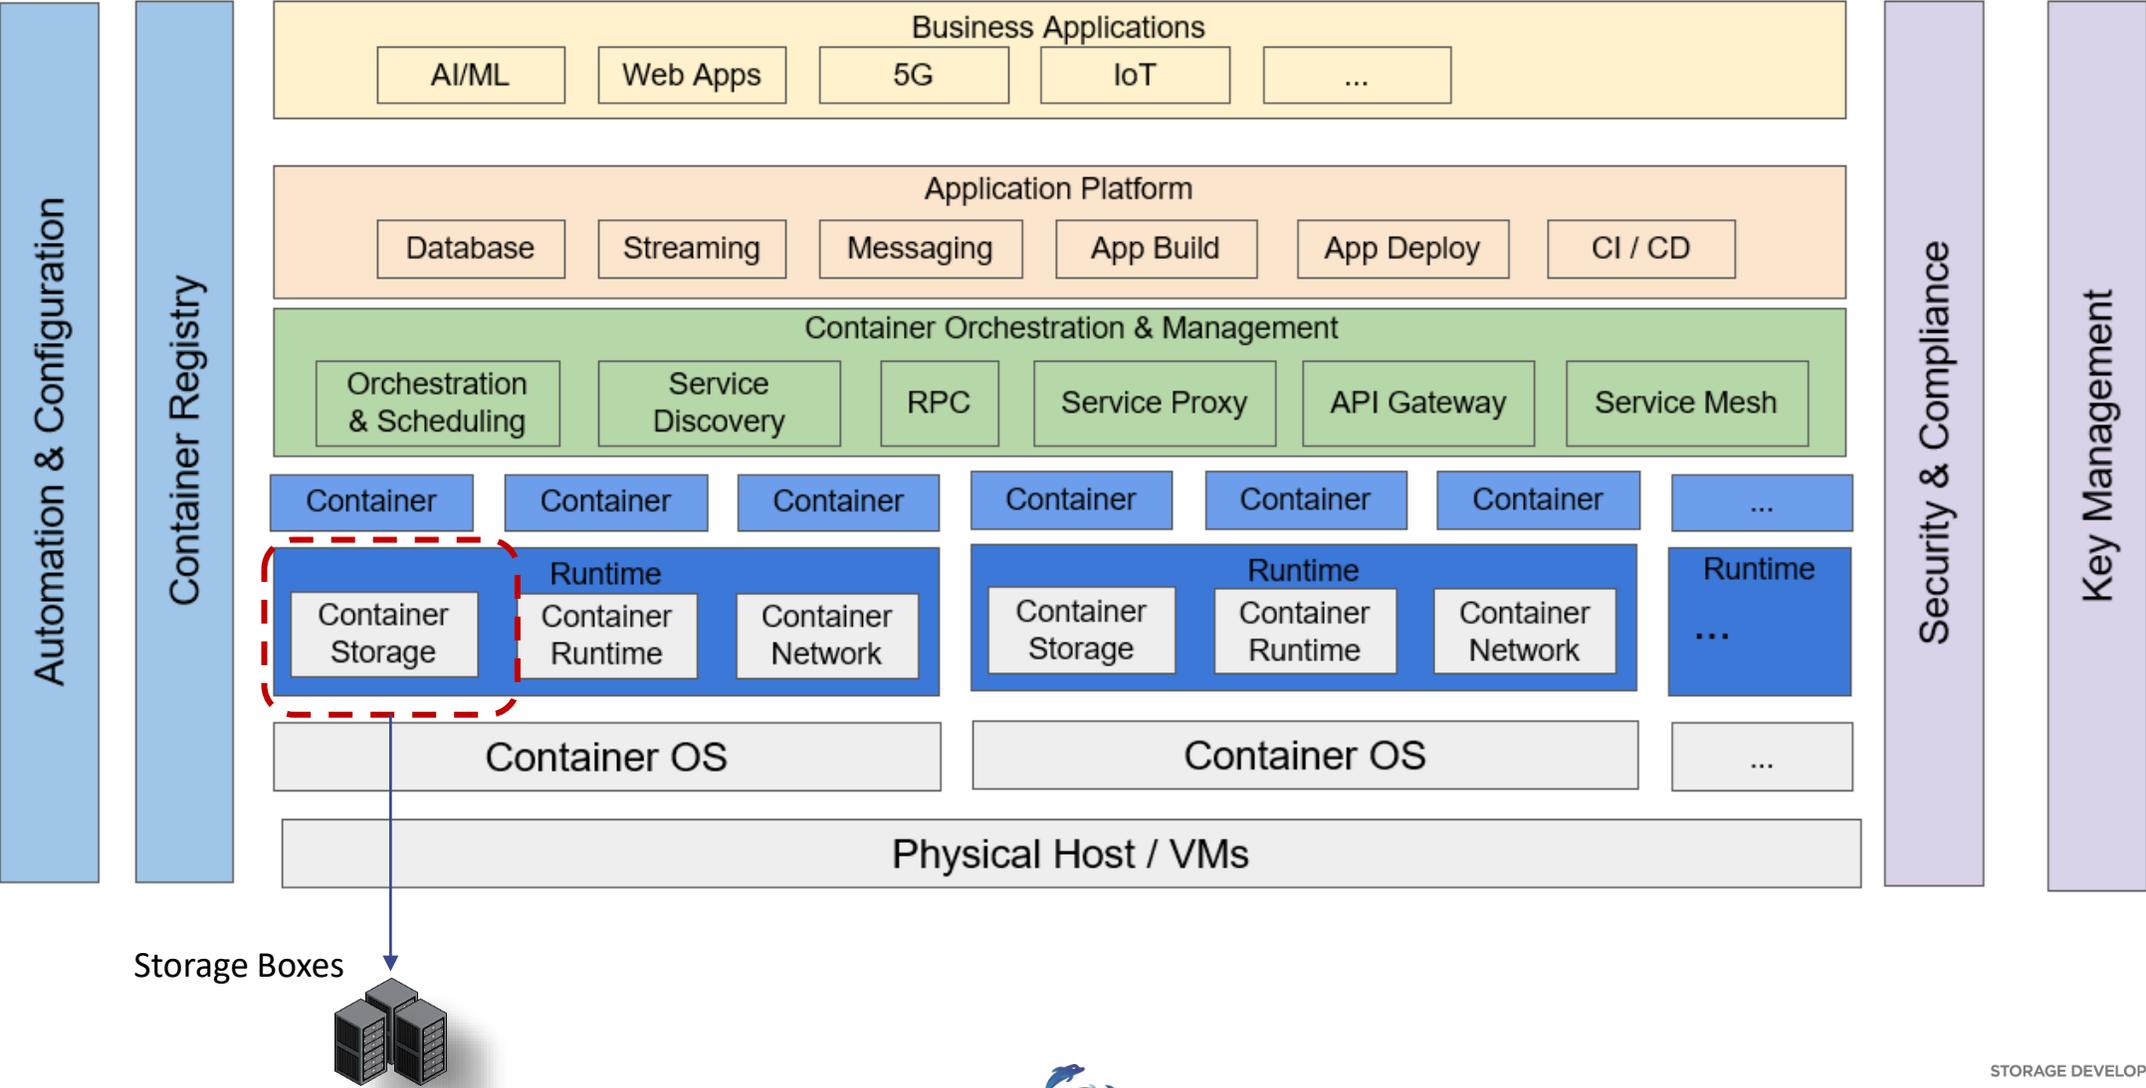
Developer productivity, cost optimization and truly hybrid.

### Container Storage Interface

Container Storage Interface (CSI) is a specification on how Container Orchestrators interact with Storage Solutions to connect to Containers.

Container Storage is CNS, other way need not be true | CSI is an interface specification

# Cloud Native Stack



# Cloud Native Storage Projects

<p>Alibaba Cloud File Storage Alibaba Cloud M.Cap: \$250.7B</p>	<p>Alibaba Cloud File Storage CPFS Alibaba Cloud M.Cap: \$250.7B</p>	<p>ALLUXIO Alluxio Funding: \$73M</p>	<p>Amazon Elastic Block Store (EBS) Amazon Web Services M.Cap: \$1.4T</p>	<p>Arrikto Arrikto Funding: \$15M</p>	<p>Azure Disk Storage Microsoft M.Cap: \$2.1T</p>	<p>Carina Becloud Funding: \$29.8M</p>	<p>ceph Ceph Foundation Funding: \$10.83M</p>	<p>cloudcasa by Catalogic Catalogic Software Funding: \$8M</p>	<p>COMMAVAULT Commvault M.Cap: \$2.6B</p>	<p>CSI Google M.Cap: \$1.6T</p>
<p>CubeFS Cloud Native Computing Foundation (CNCF) Funding: \$3M</p>	<p>Curve NetScout M.Cap: \$39.2B</p>	<p>DataCore Bolt DataCore Software Funding: \$104.7M</p>	<p>DATERA Datera Funding: \$63.9M</p>	<p>DELL EMC Dell EMC Funding: \$63.9M</p>	<p>DIAMANTI Diamanti Funding: \$78M</p>	<p>DriveScale DriveScale Funding: \$26M</p>	<p>GLUSTER Gluster M.Cap: \$118.8B</p>	<p>Google Persistent Disk Google M.Cap: \$1.6T</p>	<p>HITACHI Hitachi M.Cap: \$19.1B</p>	<p>Hewlett Packard Enterprise HPE Storage M.Cap: \$19.1B</p>
<p>HUAWEI Huawei Technologies</p>	<p>IBM Storage IBM M.Cap: \$118.8B</p>	<p>INFINIDAT Infodiat Funding: \$370M</p>	<p>inspur 浪潮 Inspur Storage Inspur Group M.Cap: \$5.5B</p>	<p>IOMesh by SmartX SmartX Funding: \$67.1M</p>	<p>ionir Ionir Funding: \$71M</p>	<p>JuiceFS JuiceData Funding: \$3M</p>	<p>k8up Cloud Native Computing Foundation (CNCF) Funding: \$3M</p>	<p>KASTEN by Veeam Kasten Funding: \$M</p>	<p>LINSTOR LINBIT Funding: \$M</p>	<p>LONGHORN Longhorn Funding: \$3M</p>
<p>MINIO MinIO Funding: \$726.3M</p>	<p>MooseFS Tosca Funding: \$1.29M</p>	<p>NetApp NetApp M.Cap: \$16.2B</p>	<p>NUTANIX Nutanix M.Cap: \$4B</p>	<p>ondat Ondat Funding: \$3M</p>	<p>OpenEBS Cloud Native Computing Foundation (CNCF) Funding: \$3M</p>	<p>OpenIO OpenIO Funding: \$3M</p>	<p>ORAS Cloud Native Computing Foundation (CNCF) Funding: \$3M</p>	<p>鹏云网络 pengyunnetwork Pangyun Network Funding: \$3M</p>	<p>Piraeus Piraeus Datastore Cloud Native Computing Foundation (CNCF) Funding: \$3M</p>	<p>portworx by Pure Storage Portworx Funding: \$55.5M</p>
<p>PURE STORAGE Pure Storage Funding: \$55.5M</p>	<p>QINGSTOR QingStor QingCloud M.Cap: \$253.3M</p>	<p>Qumulo Qumulo Funding: \$347.3M</p>	<p>Quobyte Quobyte Funding: \$84M</p>	<p>ROBIN Robin Systems Robin.io Funding: \$84M</p>	<p>ROOK Cloud Native Computing Foundation (CNCF) Funding: \$3M</p>	<p>SANDSTONE SandStone Funding: \$23M</p>	<p>SANGFOR 深信服科技 Sangfor EOS Sangfor Technologies M.Cap: \$6.4B</p>	<p>ScaleFlash ScaleFlash Funding: \$23M</p>	<p>RING Scality Funding: \$23M</p>	<p>soda foundation Soda Foundation Funding: \$3M</p>
<p>Stash Stash by AppCode AppCode Funding: \$1.29M</p>	<p>StorPool DISTRIBUTED STORAGE StorPool Funding: \$263.3M</p>	<p>SWIFT Open Infrastructure Foundation Funding: \$23M</p>	<p>TRILI Trilio Funding: \$23M</p>	<p>TRITON Object Storage Joyent Funding: \$131M</p>	<p>VELERO Velero VMware M.Cap: \$50.8B</p>	<p>Vineyard Cloud Native Computing Foundation (CNCF) Funding: \$3M</p>	<p>XSKY XSKY Data Technology Funding: \$272.2M</p>	<p>YRCloudFile Yanrong Tech Funding: \$1.3M</p>	<p>ZENKO Scality Funding: \$172M</p>	



# Industry Product Trends : Storage and Platform Vendors

Product Name	Company	Key Use Cases	Supported Infrastructure	Supported Data Types (File/Block/Object)	Deployment Mode	Data Management Technologies Covered	Open Source
Netapp Astra	NetApp	<ol style="list-style-type: none"> <li>Backup/Restore across on prem &amp; cloud (s3 backup supported)</li> <li>Data Migration across clusters</li> <li>Instant clones (working applications)</li> <li>Cloud Bursting</li> </ol>	<ol style="list-style-type: none"> <li>within and across public clouds and on-premises</li> <li>Mainly NetApp on prem and cloud storages</li> <li>k8s multicluster deployments</li> </ol>	File/Block/Object	SaaS	<ol style="list-style-type: none"> <li>DR / Snapshots</li> <li>Data Migration</li> <li>Data Provisioning for k8s (NetApp Trident CSI)</li> <li>Lifecycle Management</li> <li>Cloud Bursting</li> </ol>	No (Trident is Open Source)
PSO	PureStorage	<ol style="list-style-type: none"> <li>Container STAAS (persistent storage service for containerized app)</li> </ol>	<ol style="list-style-type: none"> <li>Within and across public clouds and on-premises</li> <li>K8s multicluster deployments (private namespace)</li> <li>PureStorage FlashArray and FlashBlade</li> </ol>	File and Block	SaaS	<ol style="list-style-type: none"> <li>Data Provisioning (policy based for Containers)</li> <li>DR</li> <li>Data Migration</li> </ol>	No (except CSI SP)
Ezmeral	HPE	<ol style="list-style-type: none"> <li>Easy to deploy applications inside Containers</li> <li>Application Data provisioning and Management</li> <li>Container app data Intelligence</li> </ol>	On-prem and Cloud supported by HPE storages	File and Block	SaaS	<ol style="list-style-type: none"> <li>Data Fabric</li> <li>Container and Container storage Monitoring</li> <li>Container appliation deployment and management</li> </ol>	No
Karavi	Dell EMC	Observability, Data Mobility and Resiliency for stateful applications	On-prem EMC storage	Block	SaaS	<ol style="list-style-type: none"> <li>Data Observability/Monitoring</li> <li>Data Resiliency</li> </ol>	Yes (Karavi is open source)
HSPC	Hitachi	<ol style="list-style-type: none"> <li>Data Protection</li> <li>Snapshot based backup</li> <li>Dynamic storage provisioning</li> </ol>	K8S clusters Hitachi VSP series storage volume	Block	SaaS	<ol style="list-style-type: none"> <li>DR/ Snapshot</li> <li>Data Provisioning for K8S (Hitachi CSI)</li> <li>High availability/DR</li> </ol>	No (except CSI SP)
Velero	VMWare	<ol style="list-style-type: none"> <li>DR</li> <li>Data Protection</li> <li>Data Migration (heterogenous/distributed env)</li> </ol>	<ol style="list-style-type: none"> <li>Within and across cloud and on-premises</li> <li>Any storage provider supporting CSI Plugin</li> <li>multi-cluster deployment</li> <li>FS based snapshot using restic</li> </ol>	Block File Object	SaaS	<ol style="list-style-type: none"> <li>DR/Snapshot</li> <li>Data Migration across distributed env</li> <li>Snapshots (scheduled/on-demand, File and Block)</li> </ol>	Yes
Longhorn	Rancher	<ol style="list-style-type: none"> <li>Snapshot based Data Protection</li> <li>Block storage for K8S</li> </ol>	On-premise and Cloud	Block	SaaS	<ol style="list-style-type: none"> <li>DR/Snapshot</li> </ol>	Yes
Restic	Open Source	<ol style="list-style-type: none"> <li>FS based snapshot</li> <li>uses Object Data Store</li> </ol>	On-Premise and Cloud	File	native	<ol style="list-style-type: none"> <li>DR/FS Snapshot</li> </ol>	Yes
OpenEBS	MayaData	<ol style="list-style-type: none"> <li>Container Attached Storage (CAS)</li> <li>Persistent Storage</li> <li>Stateful</li> <li>SDS</li> </ol>	on-prem and Cloud Plugins for different Storage CSI TP, Cloud Storage	Block File	SaaS	<ol style="list-style-type: none"> <li>Storage (SDS)</li> <li>Availability</li> <li>Persistence</li> </ol>	yes

# Products providing Data Management across Cloud and On prem (1)

## NetApp Astra

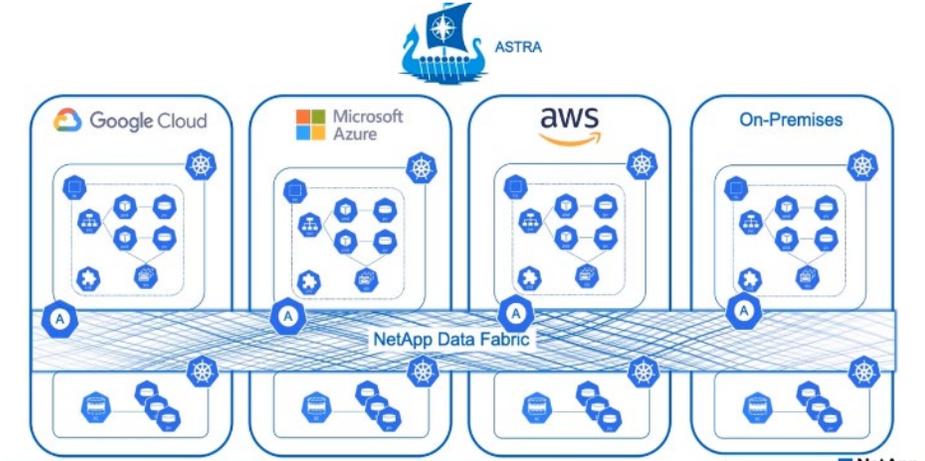
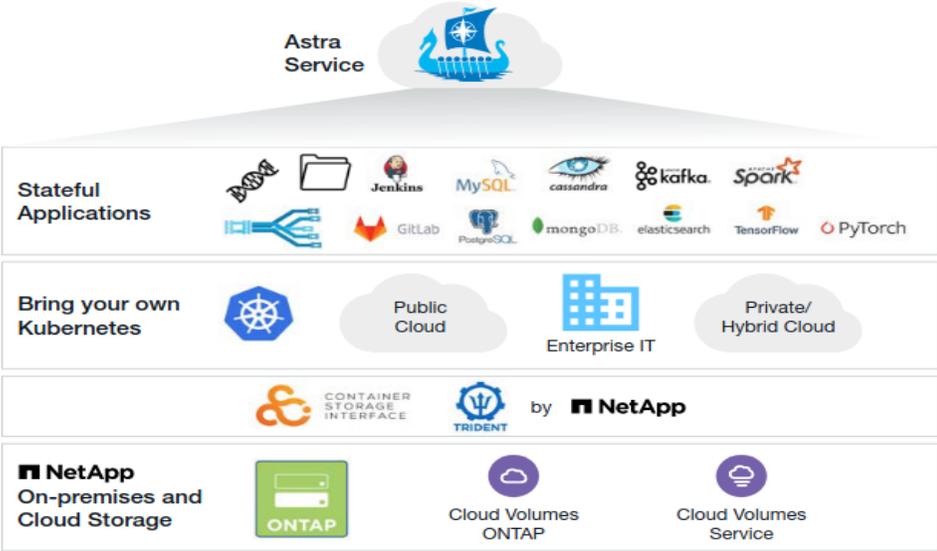
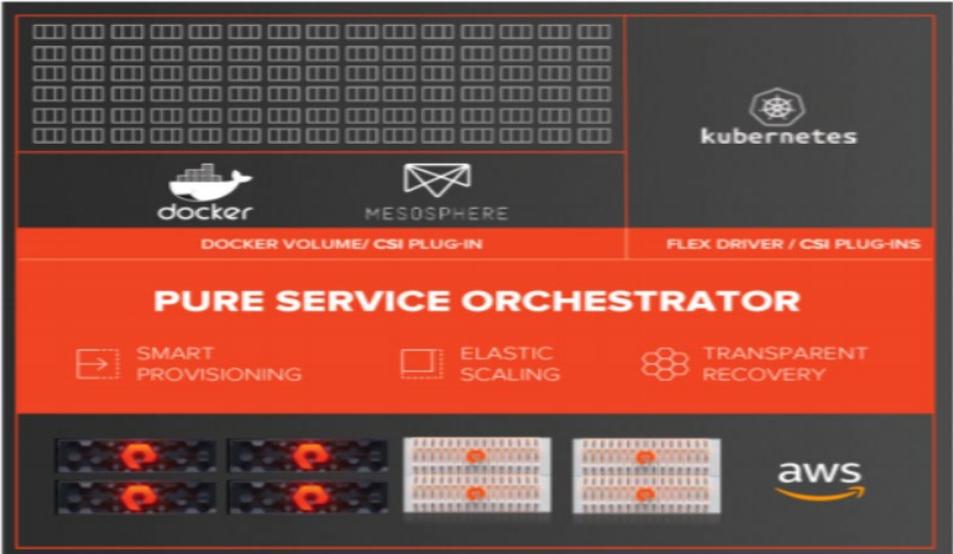
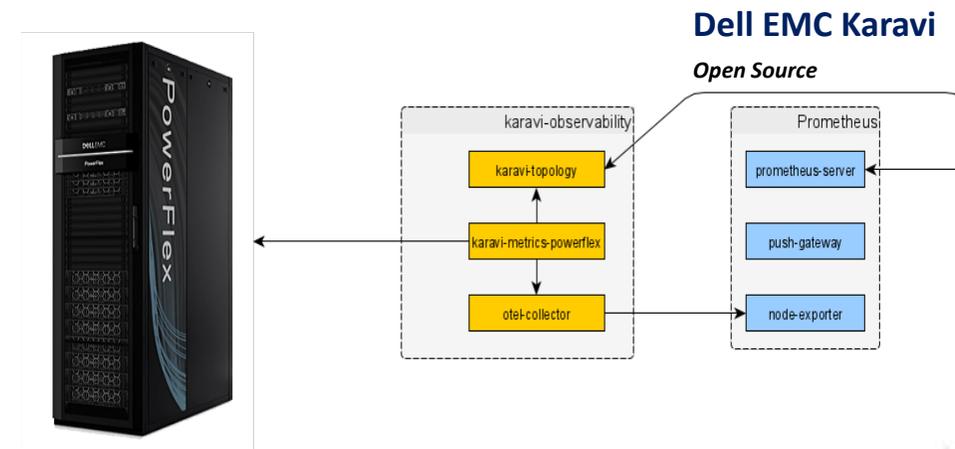
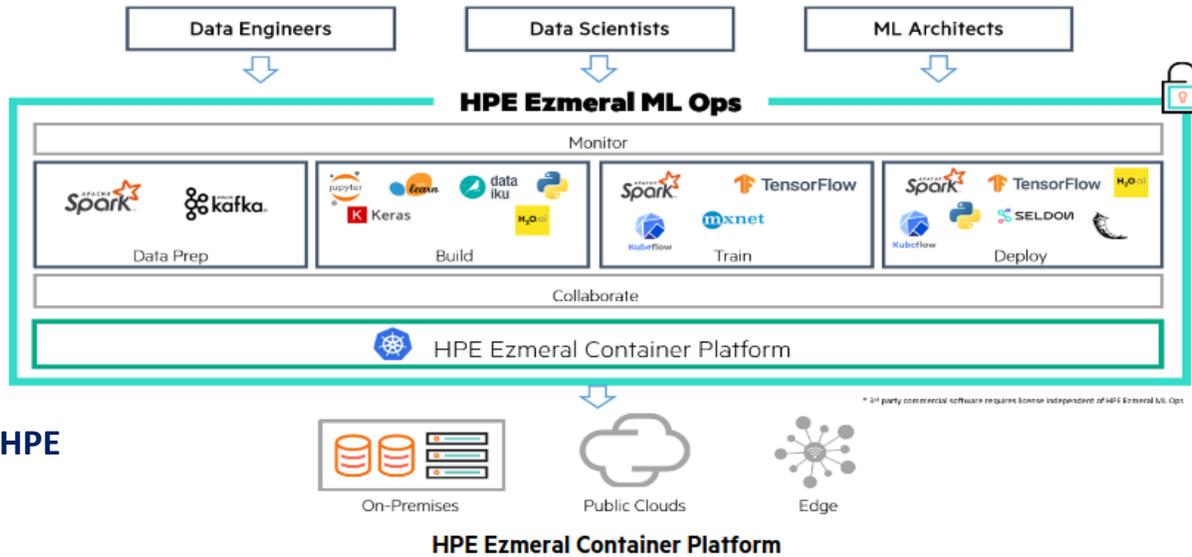


Figure 1: NetApp Astra SaaS for end-to-end application data lifecycle management.

## Pure Service Orchestrator



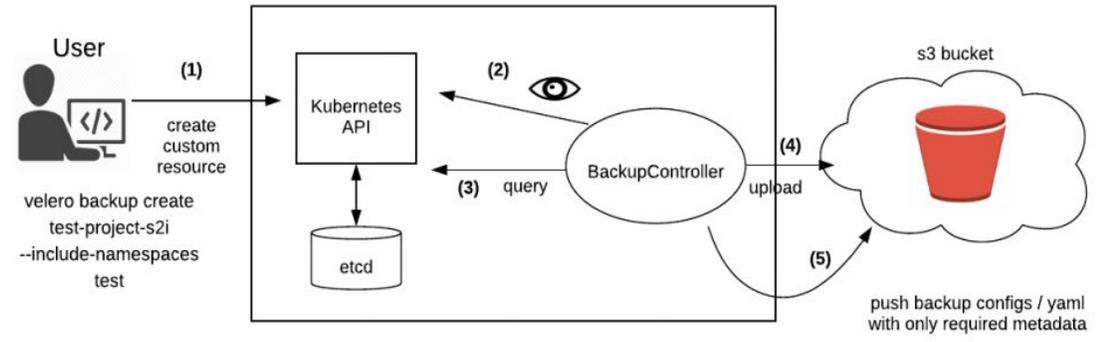
# Products providing Data Management across Cloud and On prem (2)



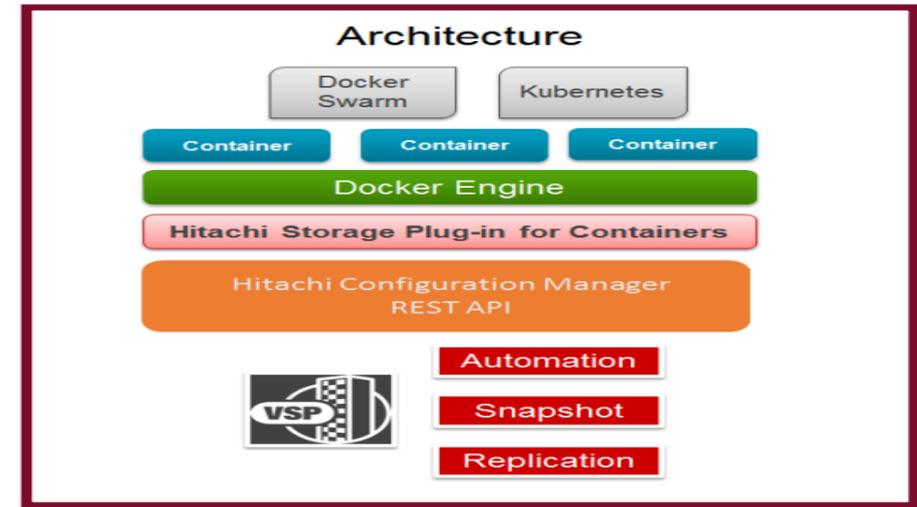
## VMware Velero

Open Source

### Velero Backup Solution Architecture



## Hitachi



**Product Direction is inclined towards hybrid cloud data management, more on cloud native storage support for data protection and mobility**

**Backup-Restore is kind of primary use case under container data protection**

# Data Protection in Kubernetes

# Data Protection is a focus area

## Data Protection Working Group

---

A Working Group dedicated to promoting data protection support in Kubernetes, identifying missing functionality and working together to design features to enable data protection support. Involves collaboration with multiple SIGs such as Apps and Storage. This [work-in-progress doc](#) tracks missing building blocks we have identified and what we are working on to fill the gaps.

The [charter](#) defines the scope and governance of the Data Protection Working Group.

## Stakeholder SIGs

---

- [SIG Apps](#)
- [SIG Storage](#)

## Meetings

---

*Joining the [mailing list](#) for the group will typically add invites for the following meetings to your calendar.*

<https://github.com/kubernetes/community/blob/master/wg-data-protection/README.md>

WIP Document : <https://docs.google.com/document/d/1yHbW0hxHehQzdaL7AWSI81OW4f2OcBoskXTbez92-U/edit#>

# What is available & missing in Kubernetes

## ■ Available:

- Application : Workload APIs: StatefulSet, Deployment, DaemonSet, etc. | Application CRD
- Storage : Basic Building Block: Volume Snapshots

## ■ Missing:

- Volume Backups
- Backup Repositories
- Volume Group, Group Snapshot
- CBT
- ...

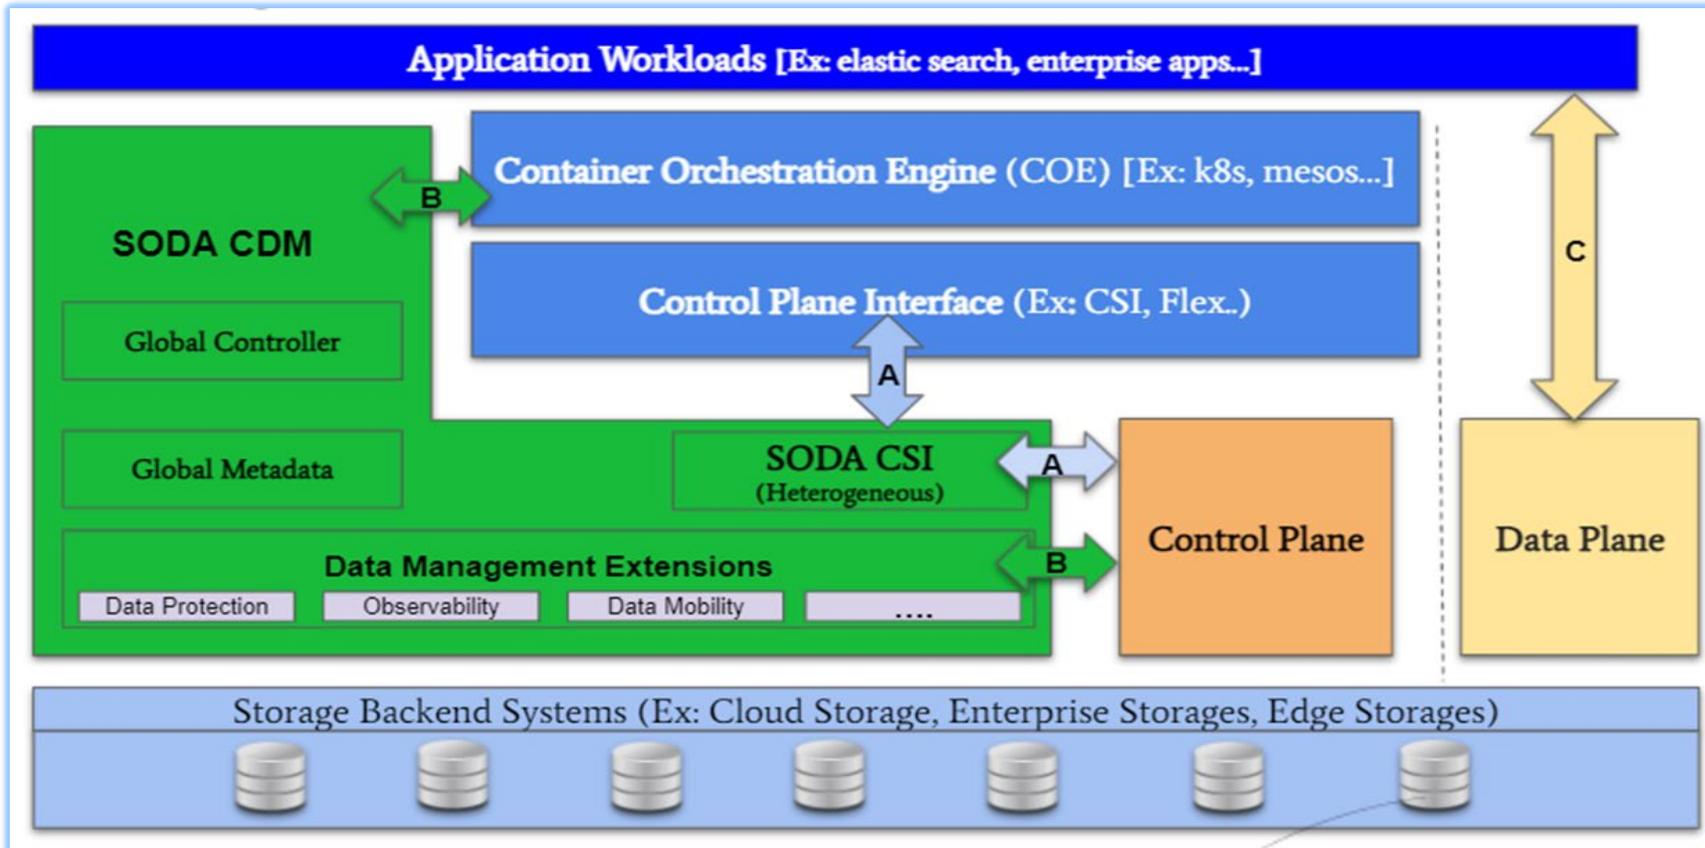
*CSI API spec and implementation are considered for enhancements for above missing features under the Data Protection WG in Kubernetes*

# Multiple projects/products to plug this gap, currently

- Storage vendor specific backup/restore solutions!
- Third Party Commercial products
- Open Source Projects
  - [Velero](#)
  - [Kahu \(SODA Foundation\)](#) – We will focus on this today!

# SODA CDM : Towards Easy Container Data Management

Augment Kubernetes (or COE) capabilities for heterogeneous and hybrid container data management.



- Designed for Container Data Management: Data Protection, Data Observability, Data Mobility and more
- Aligned with CSI
- Unified CSI
- Heterogeneous Ready
- Multi-cluster and Federation Support
- Hybrid Container Data Management (Kubernetes across on prem, cloud and edge)

**A:** Support control plane interface API directly and interact with respective COE

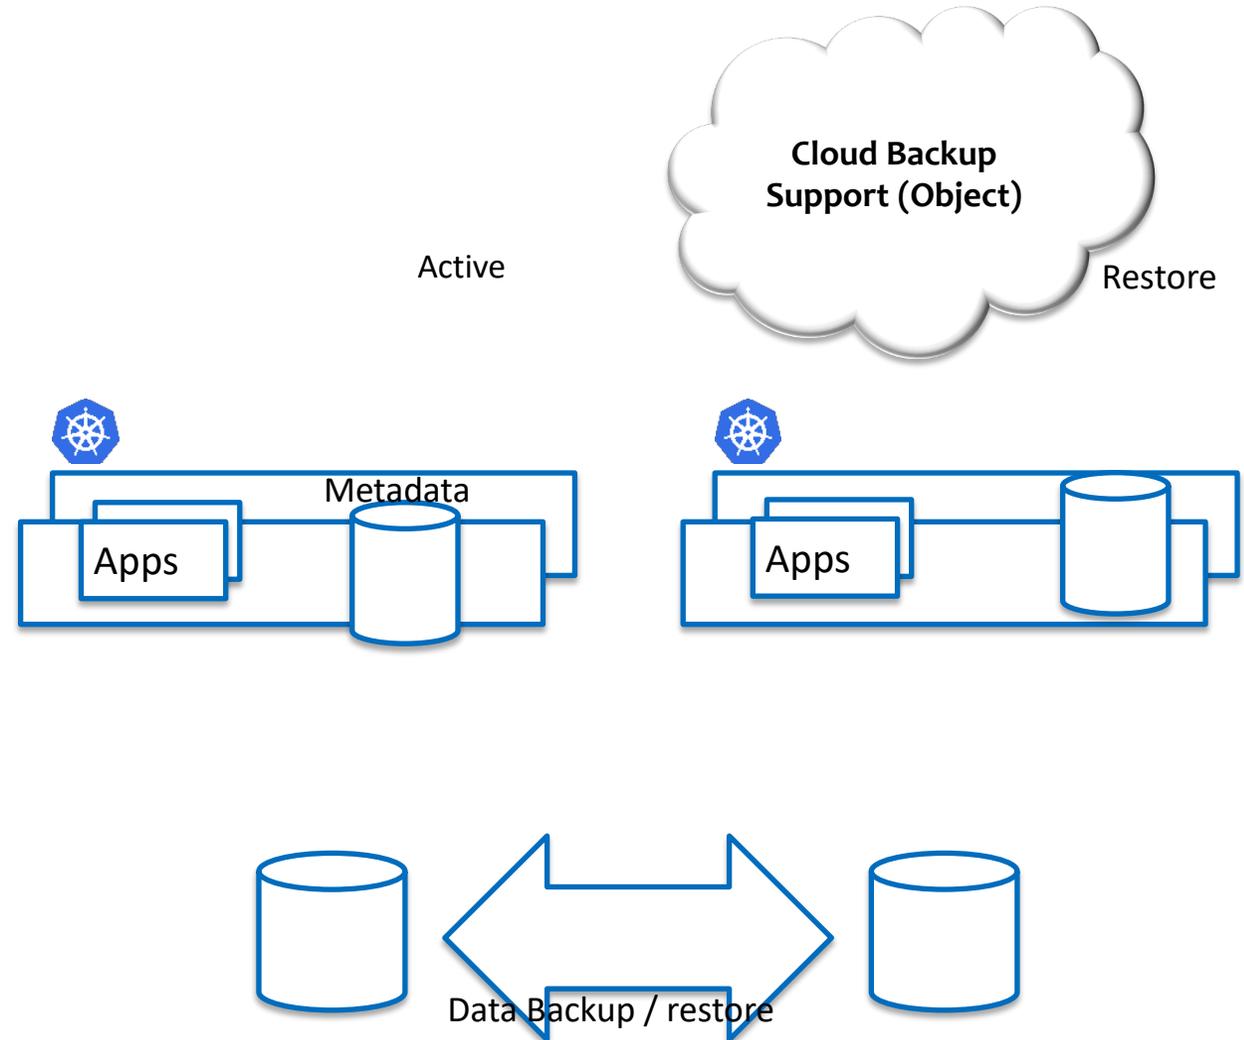
**B:** Interact with Orchestrator through Data Management Framework(SODA)

**C:** Workloads consume storage through data access interfaces.

# Kahu : First project under SODA CDM for Seamless Data Backup & Restore!

Backup and restore of Kubernetes Resource, Config and PV data, across multiple clusters and federation environment

- Backup and Restore for Kubernetes
- Federation / Multi Cluster
- Storage Provider Support
- Data Backup Plugins / Extensions
- Low RPO(Recovery Point Objective) / RTO(Recovery Time Objective)



# Kahu Phase 1 : 2022

## Backup and Restore features

- Metadata backup/Restore of K8S resources at different scope such as cluster, namespace, label selector
- Data backup / Restore of applications in full snapshot way
- Incremental and differential backup
- Custom prehook/post hook
- Backup across storage providers
- CSI based and Non CSI based volume backups.

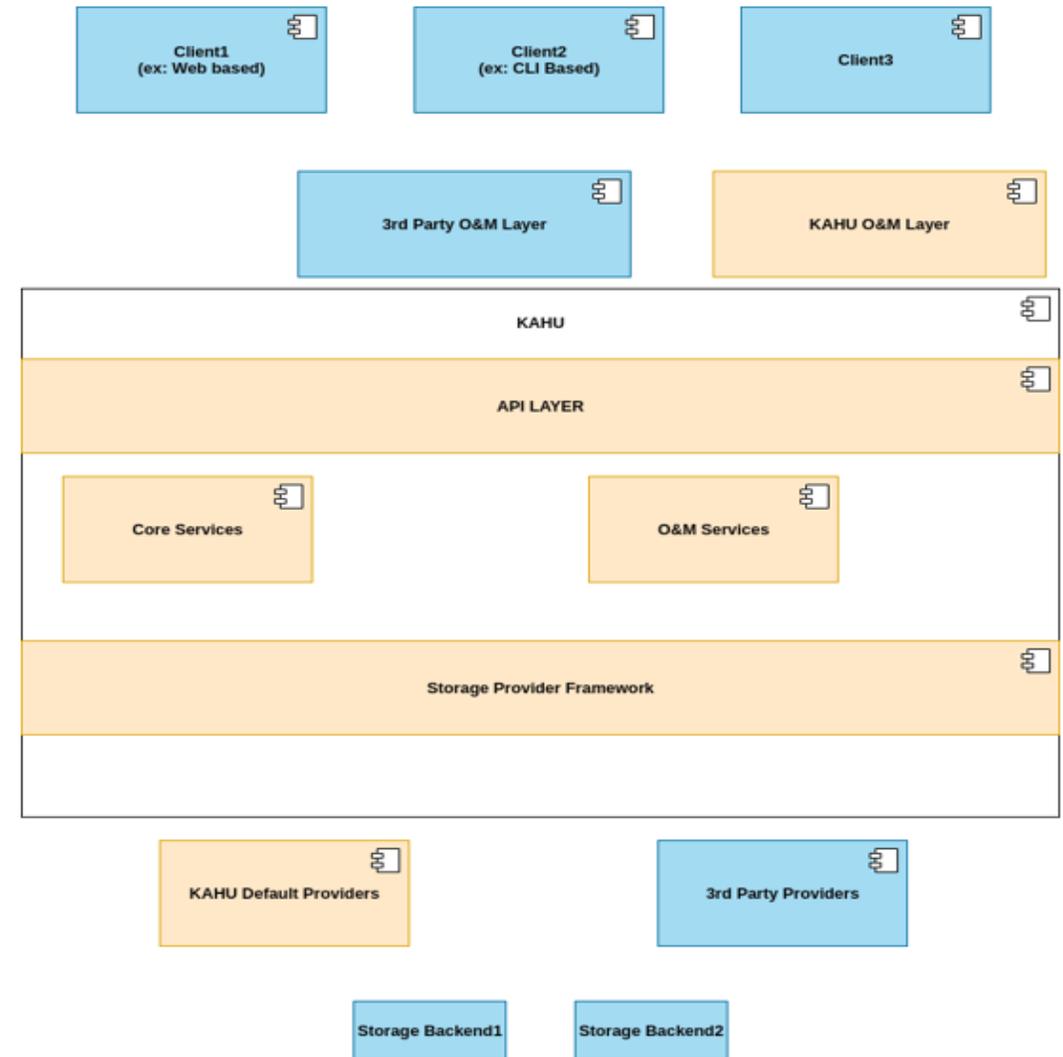
## Storage provider framework support

- Dynamic integration of any storage providers for metadata or volume
- Support coexistence of multiple providers during runtime

## Automation and Orchestration

- Scheduled backup
- Backup based on event driven mechanisms
- Backup based on policies
- Backup service plan – Provider independent service plan

Project initiated : <https://github.com/soda-cdm/kahu>



# First POC of Kahu in SODA Madagascar v1.7.0 Release

Metadata backup PoC is getting tested!

```
root@ubuntu:~# k get pods
NAME                                READY   STATUS    RESTARTS   AGE
kahu-backup-controller-58bd685ddf-wcnq6  1/1     Running   0           9m44s
kahu-nfs-provider-79754df9f5-22cqz      2/2     Running   0           30m

root@ubuntu:~# k get svc
NAME                TYPE          CLUSTER-IP      EXTERNAL-IP   PORT(S)          AGE
kahu-nfs-provider   ClusterIP     10.96.167.143   <none>        8181/TCP         30m
kubernetes          ClusterIP     10.96.0.1       <none>        443/TCP          275d
```

1. Kahu Deployment

```
apiVersion: v1
kind: Pod
metadata:
  name: pod1
spec:
  containers:
  - name: webserver
    image: nginx:latest
    ports:
    - containerPort: 80
```

```
root@ubuntu:~# k get pods
NAME                                READY   STATUS    RESTARTS   AGE
kahu-backup-controller-58bd685ddf-wcnq6  1/1     Running   0           12m
kahu-nfs-provider-79754df9f5-22cqz      2/2     Running   0           32m
pod1                                     1/1     Running   0            8s
```

2. Application pod deployed

```
apiVersion: kahu.io/v1beta1
kind: Backup
metadata:
  name: example-pod-backup
spec:
  includedResources: ["pod"]
```

3. Backup CRD

```
root@ubuntu:~# ls -lart /tmp/nfs
total 16
drwxrwxrwx 2 nobody nogroup 4096 May 25 22:06
drwxrwxrwt 9 root    root    12288 May 25 22:17
root@ubuntu:~# k apply -f backup-cr.yaml
backup.kahu.io/example-pod-backup created
root@ubuntu:~# k get backups.kahu.io
NAME                AGE
example-pod-backup  19s
root@ubuntu:~# ls -lart /tmp/nfs
total 20
-rw-r--r-- 1 nobody nogroup 1132 May 25 22:18 example-pod-backup
drwxrwxrwx 2 nobody nogroup 4096 May 25 22:18
drwxrwxrwt 9 root    root    12288 May 25 22:18
```

4. Backup Created

SODA Madagascar Release : <https://github.com/sodafoundation/soda/tree/master/madagascar>

# Join Us

Right time to collaborate! ( Now is the right time! 😊 )

We welcome developers of all levels of skills and experience to join to build these future technology solutions!



<https://sodafoundation.io/slack>

Channels: #soda-cdm | #soda-lake

# SODA Data and Storage Trends Survey is announced!



<https://www.sodafoundation.io/data-storage-trends-survey-2022/>

# THANK YOU!

<https://www.sodafoundation.io/>

SODA Source Code:

<https://github.com/sodafoundation>

<https://github.com/soda-cdm>

SODA Docs: <https://docs.sodafoundation.io/>

Join SODA Slack:

<https://sodafoundation.io/slack/>

Follow SODA Twitter:

<https://twitter.com/sodafoundation>





# Please take a moment to rate this session.

Your feedback is important to us.