Cloud Native Storage Evolution, Ecosystem, & What's New

Michael Cade

Field CTO

Veeam Software

Gunnar Grosch

Principal Developer Advocate

AWS

Agenda

Introduction & Why This Matters

Evolution of Cloud-Native Storage

Cloud Native Storage Building Blocks

Ecosystem

Now, New, Next...

Evolution of Cloud-Native Storage

Pre-Cloud Native – How we got here?

Still very much here today…

Direct Attached Storage

DAS

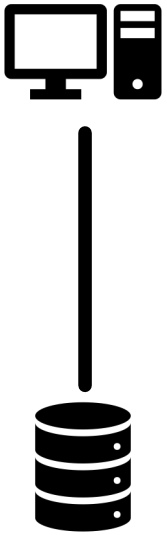

Network Attached Storage

NAS

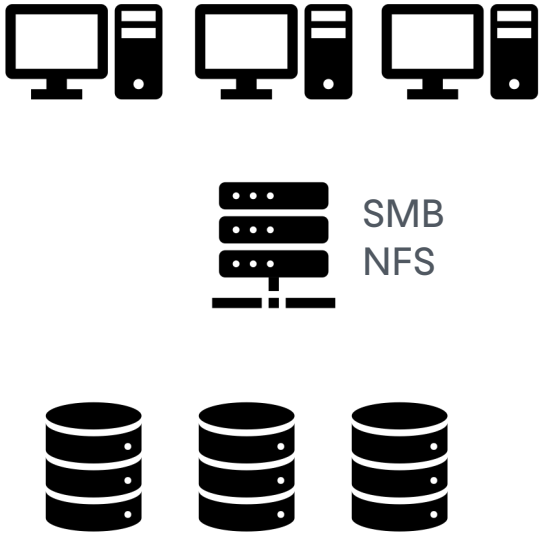

Storage Area Network

SAN

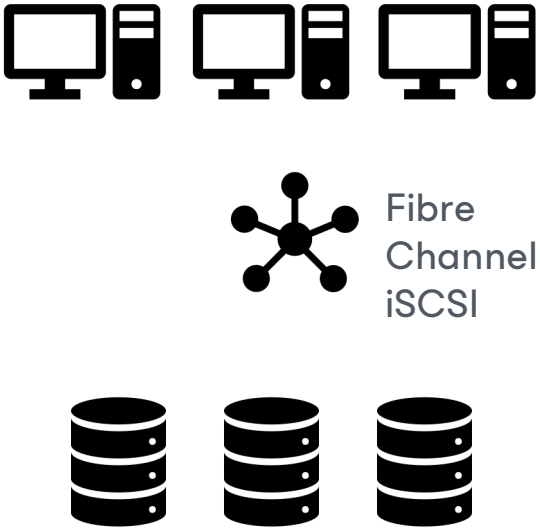

Early Cloud Native Storage

Cloud Native Storage ≠ Just Kubernetes

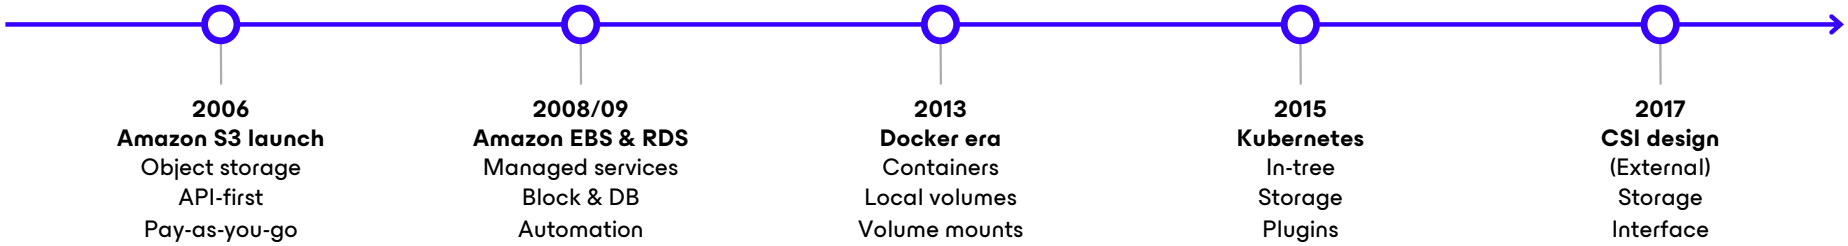

Storage evolution:

- From hardware-centric → API-driven
- From manual → Automated/Managed
- From fixed → Pay-as-you-go
- From limited scale → Infinite scale

Kubernetes as a Catalyst, Not the Centre

One of the platforms that drove cloud-native storage evolution

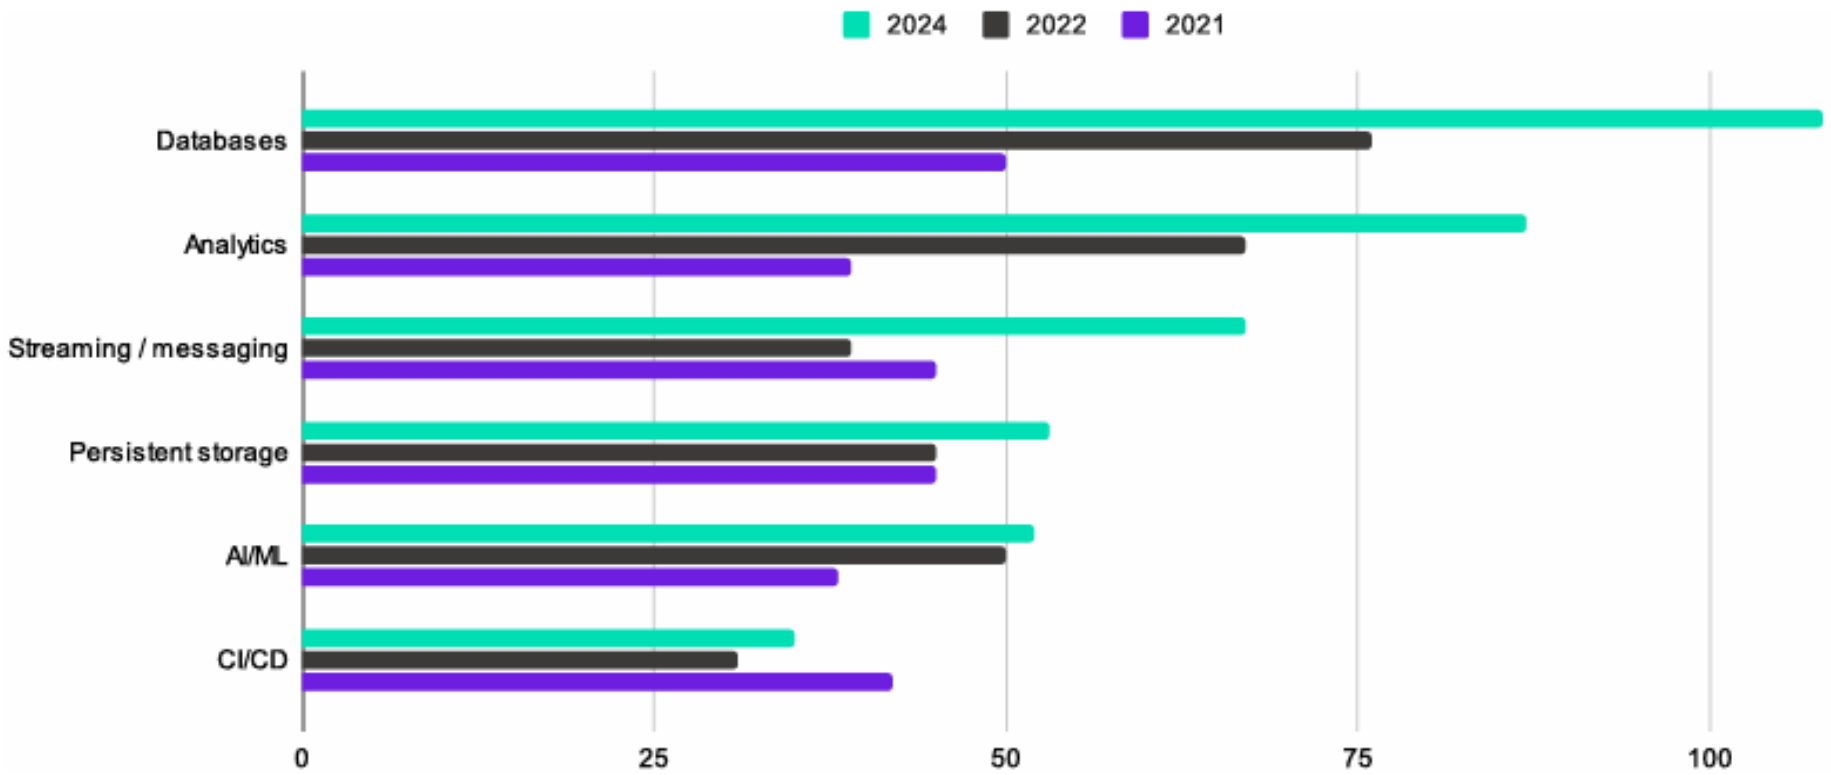

Data on Kubernetes 2024

Maturing Kubernetes

From stateless to stateful

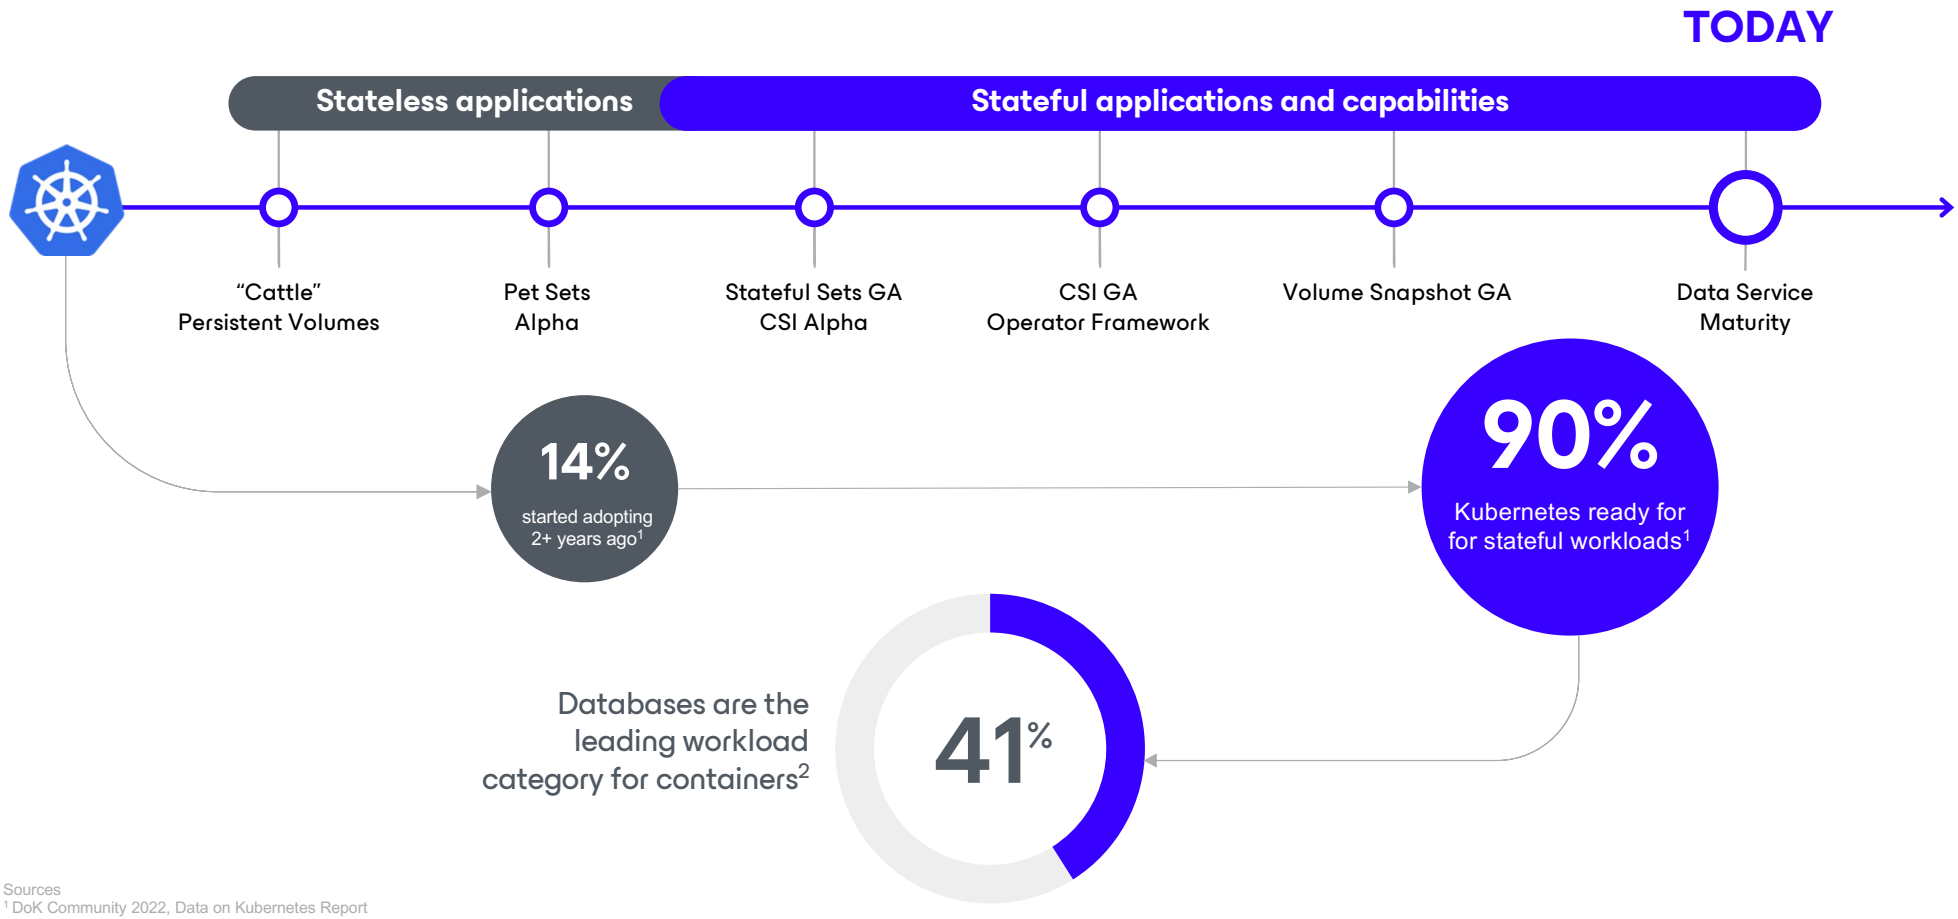

Sources
[1] DoK Community 2022, Data on Kubernetes Report
2 Datadog 2023, 10 Insights on Real-World Container Use

The Evolution of Data On Kubernetes

2021
Early focus on basic stateful workloads, with databases leading adoption

2022
Emergence of analytics and early AI/ML adoption

2024
Full spectrum of data workloads, with AI/ML growing

Our data shows the continuing evolution of workload types on Kubernetes:

Databases
Maintaining their #1 position as the most commonly deployed workload across all three years

Analytics
Moving up to #2 in 2024 from a lower position in previous years

Streaming/messaging
Now #3, showing significant growth

Kubernetes is just another platform

It's Compute, Networking & Storage

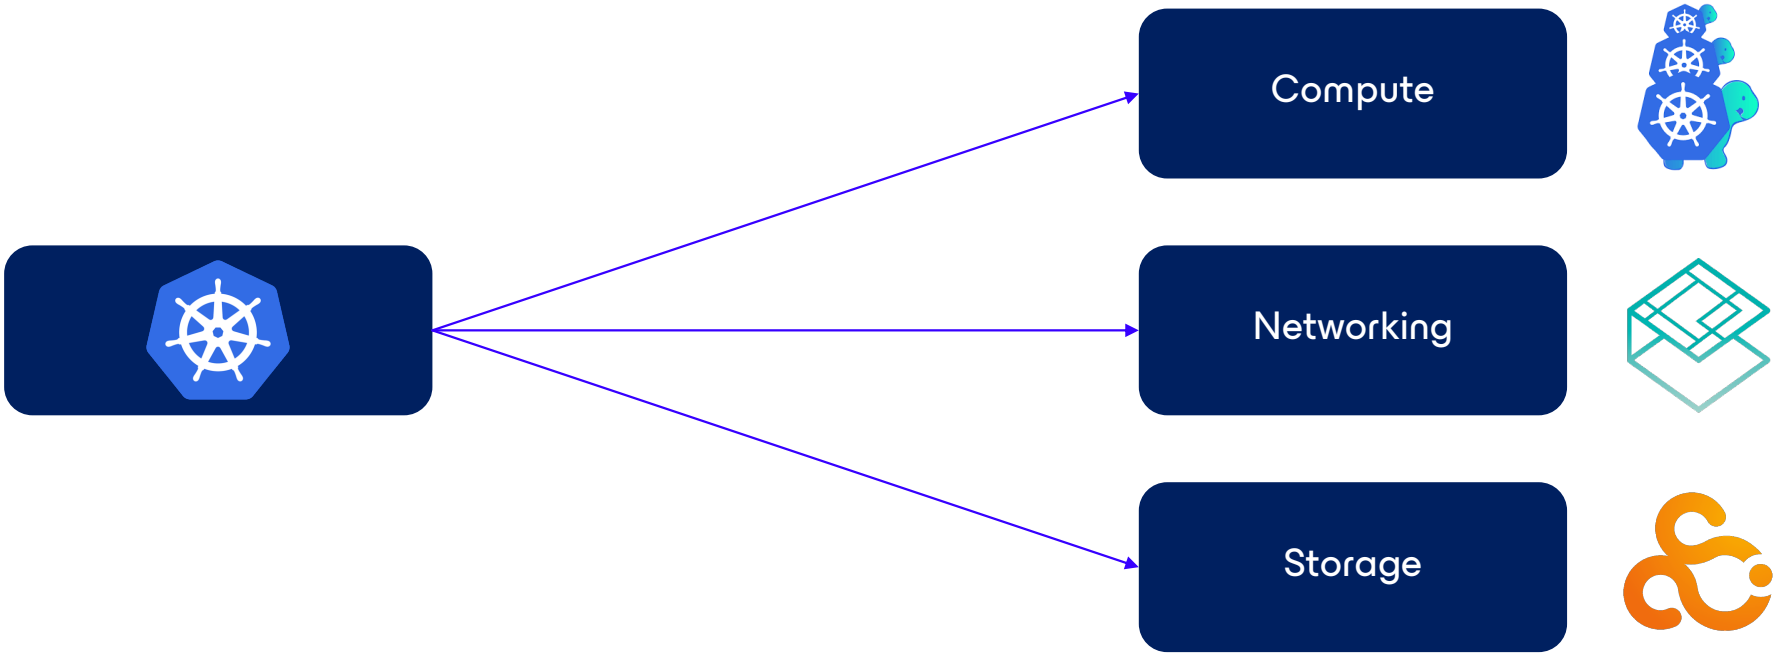

Cloud providers followed and enabled storage for Cloud Native

Abstracting complexity, enabling innovation

1. Infrastructure abstraction
   - From: Physical storage planning
   - To: API call and console clicks
2. Management abstraction
   - From: Storage admin expertise required
   - To: Developer-friendly interfaces
3. Integration abstraction
   - From: Complex storage orchestration
   - To: Managed services integration

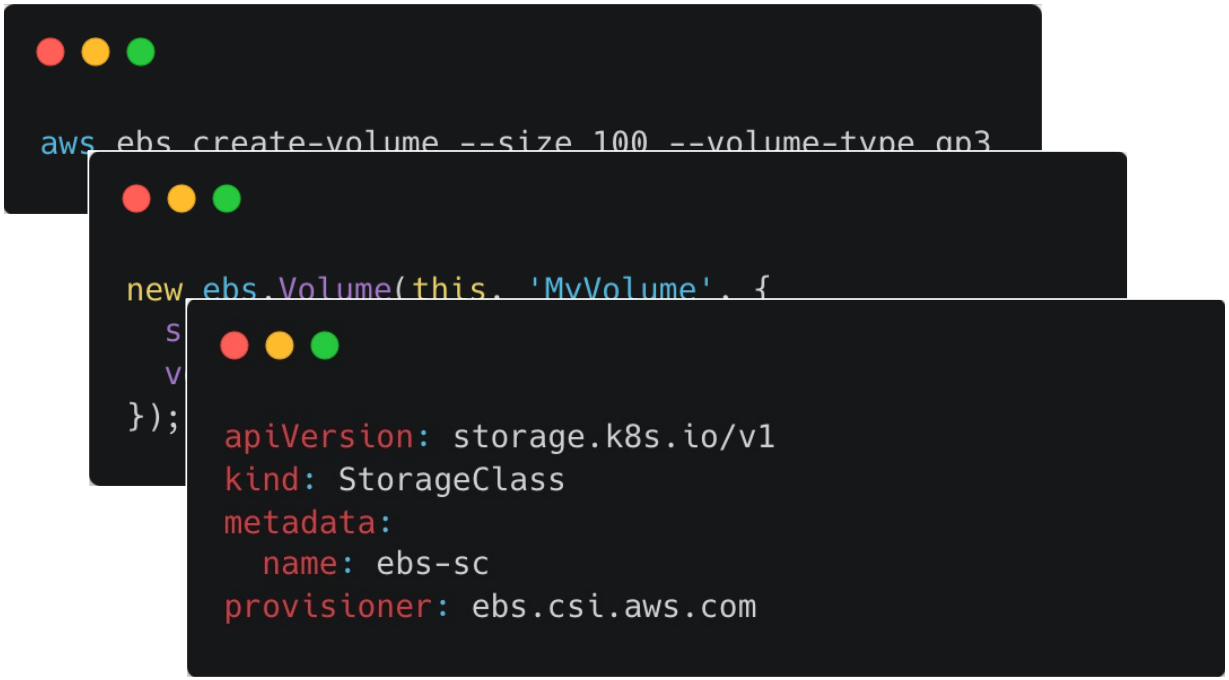

Cloud providers simplified storage through code-first management, unified access, automation, and native integrations.

How many people are running Stateful workloads within the cluster?

Cloud Native Building Blocks

CSI (Container Storage Interface)

Containerised storage plugin deployed using standard Cloud Native primitives

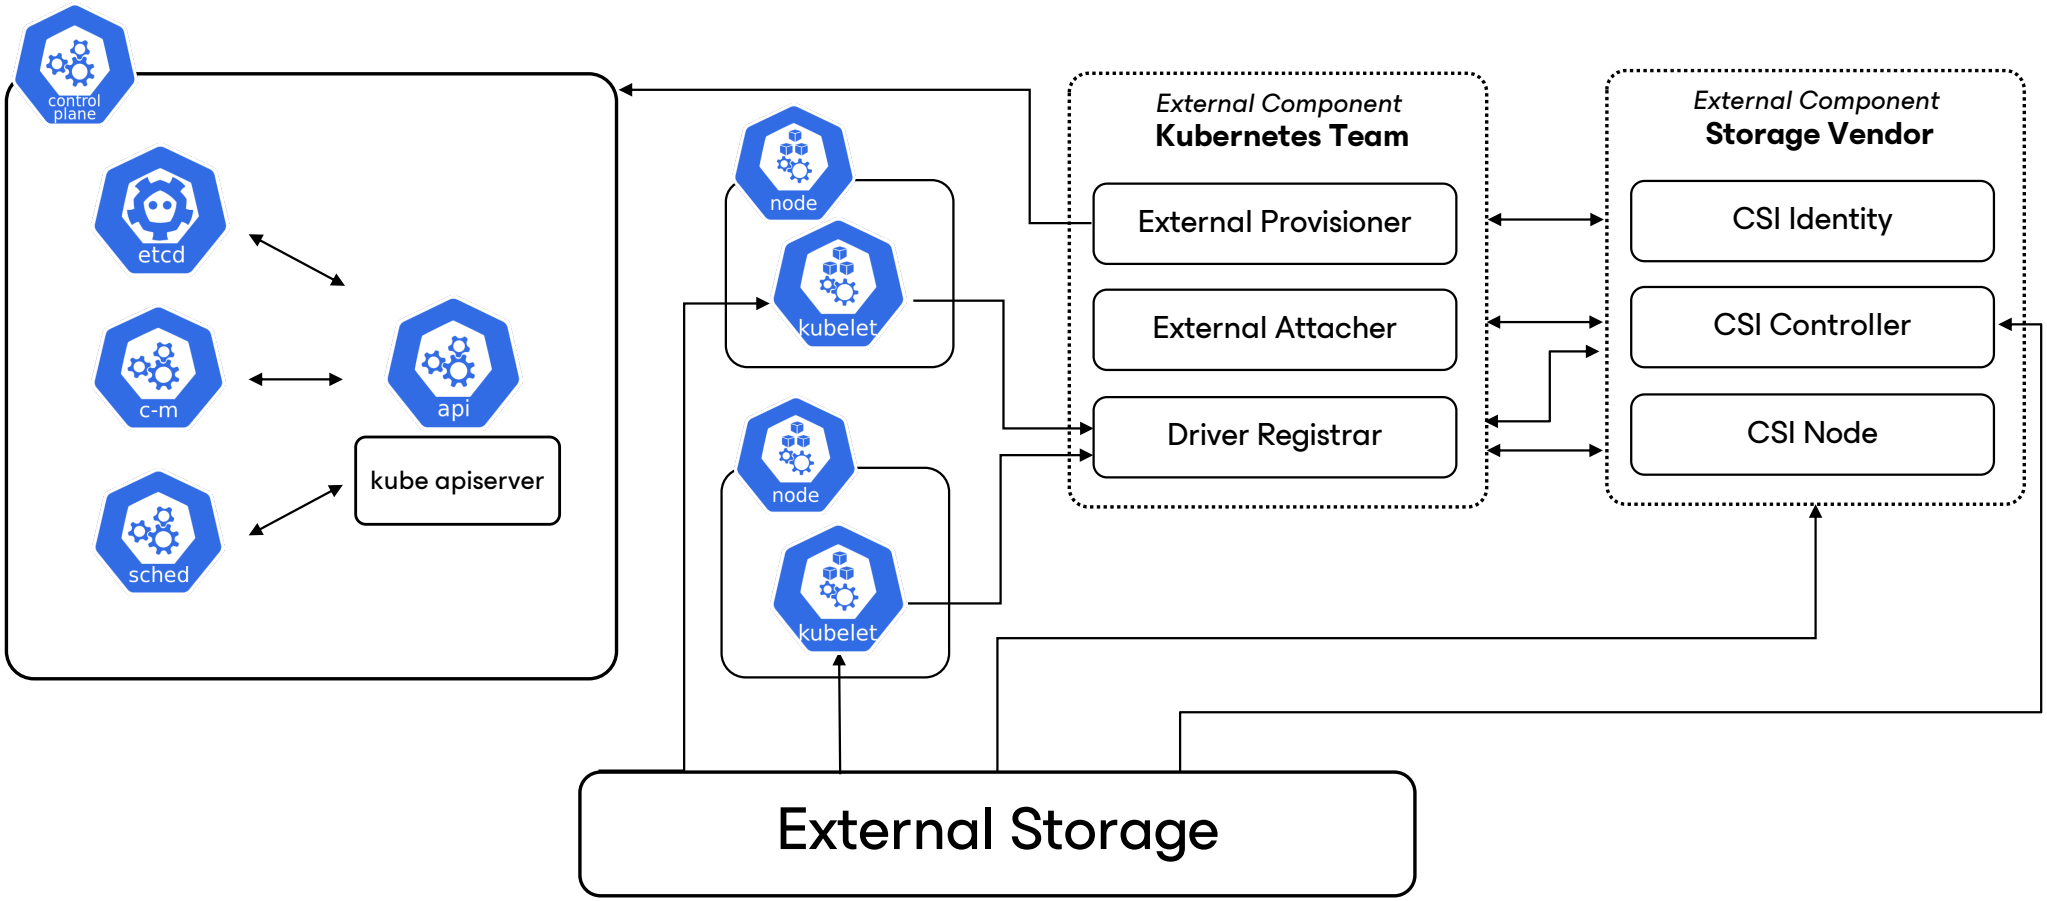

Ecosystem

The Cloud Native Storage Landscape

A diverse vending machine of options

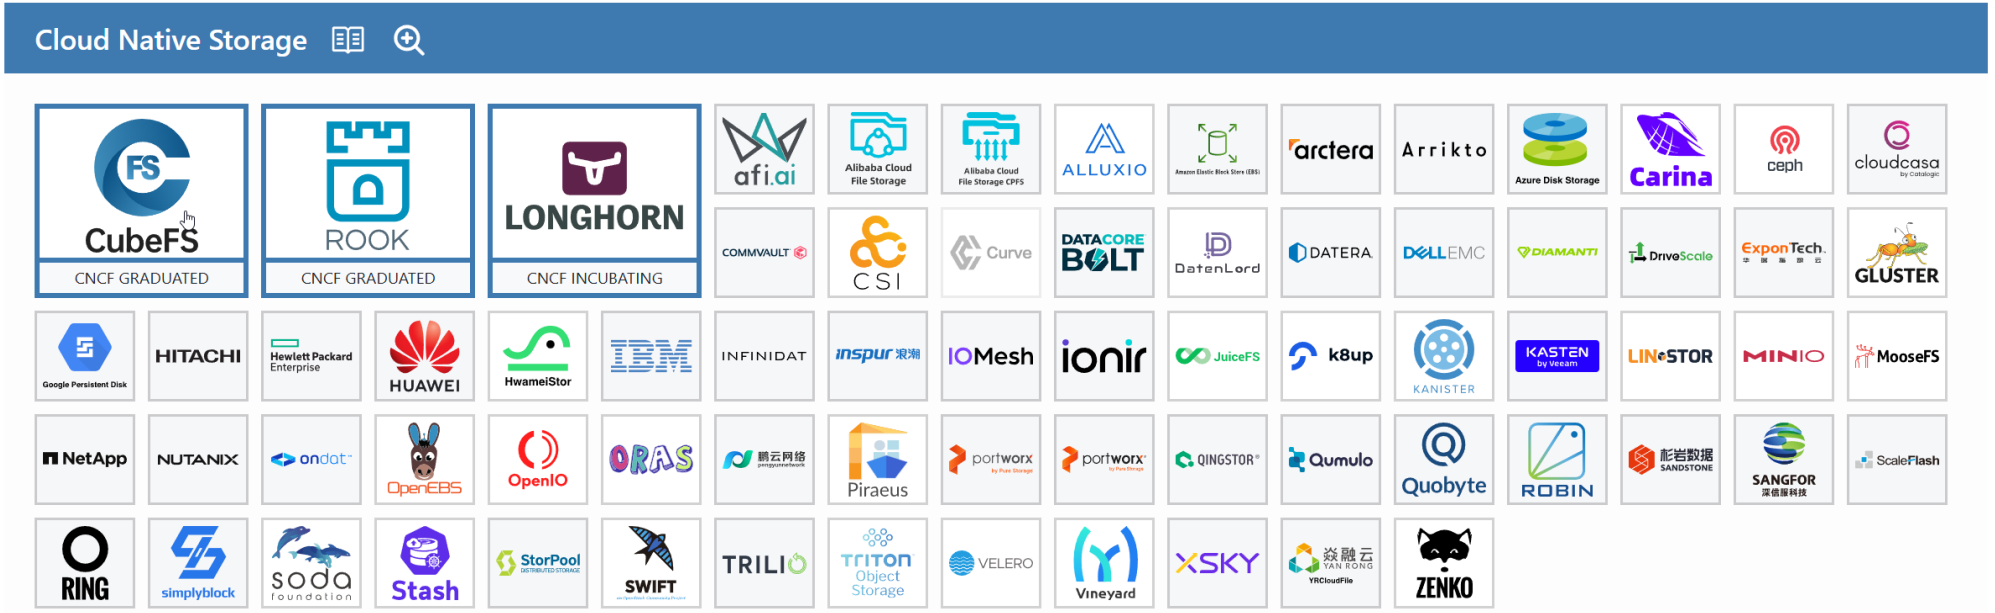

Want to know more about the Data Services ecosystem?

Shameless Plug

13:35 - 14:00

Room 1 O2 Södra
Cloud Native Deployment Models, Operators, Protection & Mobility
Michael Cade, Field CTO, Veeam Software and Gunnar Grosch, Principal Developer Advocate at AWS

Cloud Native Deployment Models, Operators, Protection & Mobility

Let's talk about how we deploy, protect, and move workloads in cloud-native environments. We'll cover deployment models, operators, and what's working (or not) when it comes to keeping things secure and available.

Now, New, Next...

Cloud Native Deployment Patterns

Multiple paths to providing data persistence

Application using data services outside of Kubernetes	Data services in Kubernetes – separate from Application	Application includes data services – all in Kubernetes

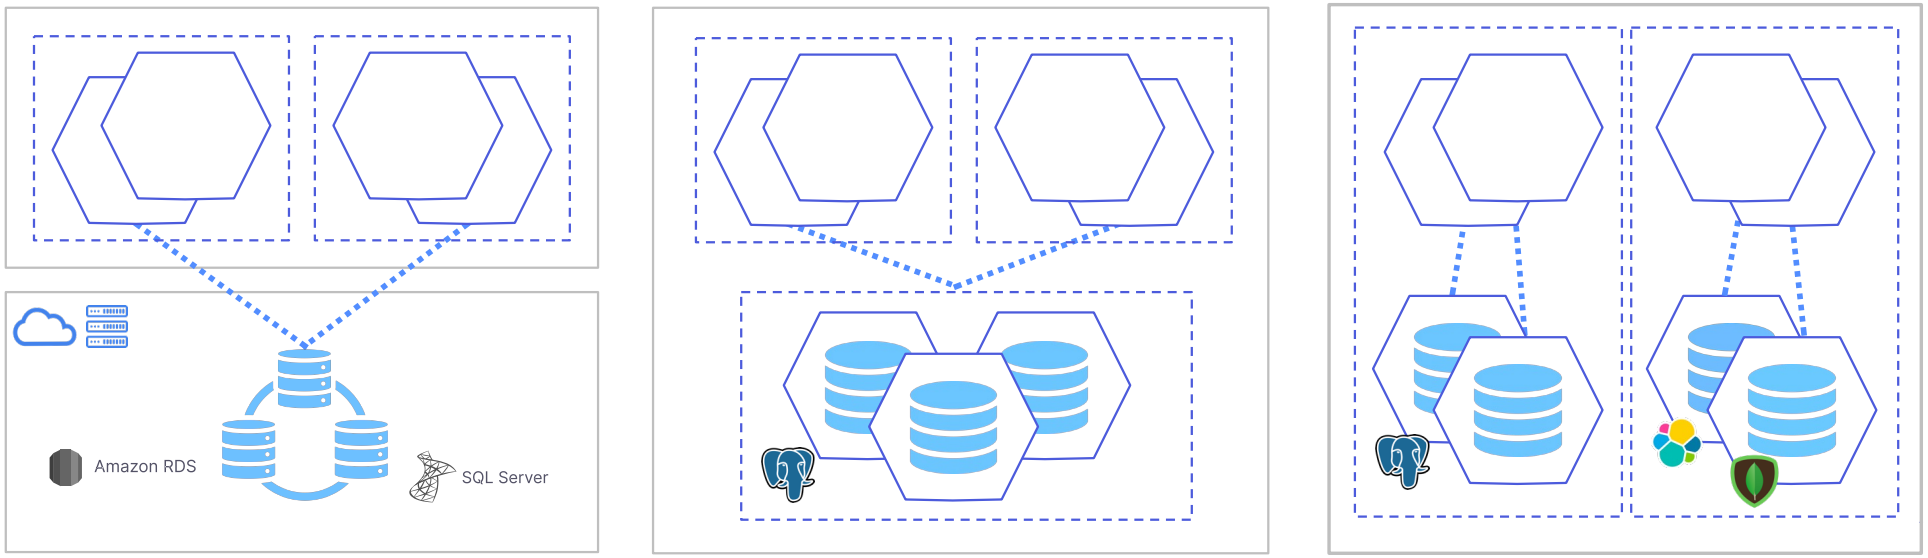

KASTEN K10

Kanister.io

veeam

Recent Innovations

VMs on Kubernetes

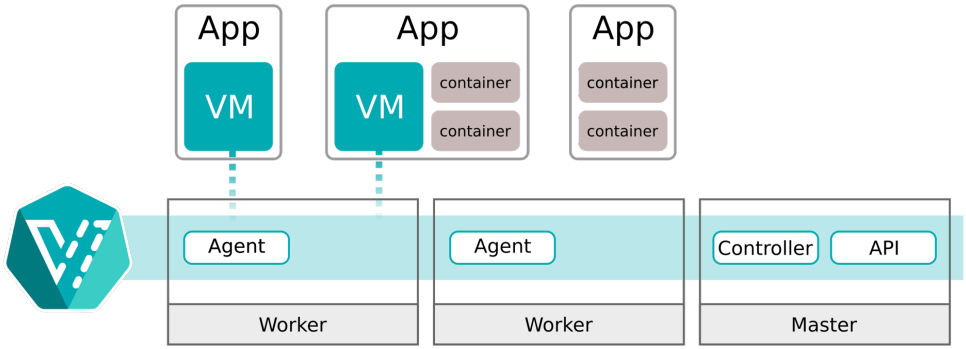

Workflow Convergence

Converging VM management into container management workflows

Using the same tooling (kubectl) for containers and Virtual Machines

Keeping the declarative API for VM management (just like pods, deployments, etc...)

veeam

Recent Innovations

Kubernetes 1.32 introduces **Change Block Tracking (CBT)**

Storage efficiency in cloud-native environments

Reduces costs and improves performance

CBT requires updates to CSI drivers

CBT ensures data is encrypted during transfer

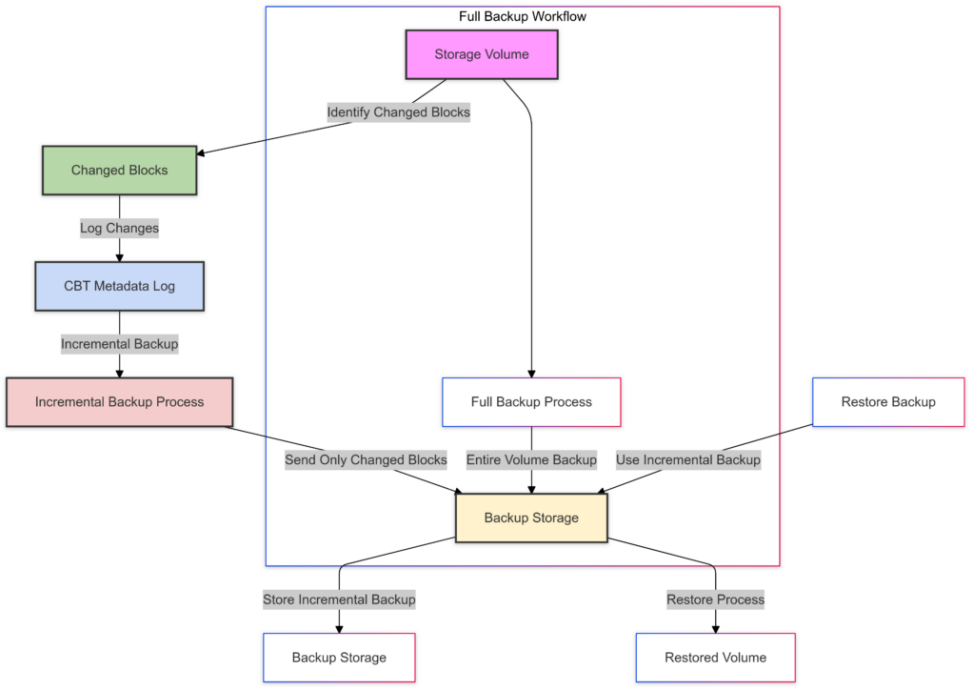

Available as an **alpha feature**

What else came with Kubernetes 1.32 (December 2024)

Automatic Deletion of PVCs for StatefulSets

Dynamic Memory-Backed Volume Sizing

Recovery from Volume Expansion Failures

Enhancements to Dynamic Resource Allocation (DRA)

Volume Group Snapshot

Kubernetes 1.33 (maybe)

Things are not slowing down

KEP-4876 – Mutable CSINode Allocatable Property

KEP-3314 – CSI Changed Block Tracking

Title	Status	Stage
Add storage capacity scoring #4049	Tracked for enhancements freeze	Alpha
Generic data populators #1495	Tracked for enhancements freeze	Stable
Portworx file in-tree to CSI driver migration #2589	Tracked for enhancements freeze	Stable
Always Honor PersistentVolume Reclaim Policy #2644	Tracked for enhancements freeze	Stable
Mutable CSINode Allocatable Property #4876	Tracked for enhancements freeze	Alpha
Remove gitRepo volume driver #5040	Tracked for enhancements freeze	Alpha
Speed up recursive SELinux label change #1710	Tracked for enhancements freeze	Beta
VolumeSource: OCI Artifact and/or Image #4639	Tracked for enhancements freeze	Beta
CSI Differential Snapshot for Block Volumes #3314	Tracked for enhancements freeze	Alpha
Recursive Read-only (RRO) mounts #3857	Tracked for enhancements freeze	Stable

Future Directions

What might we see in the cloud native storage space futures?

Event-driven storage (auto-provisioned, short-lived volumes)

AI-native storage solutions (e.g., AI-driven storage tiering, caching)

Decentralized storage architectures (IPFS, blockchain-backed storage)

Confidential computing (encryption-in-use to prevent unauthorized access)

Storage is becoming more intelligent, ephemeral, and regulation-aware.

Resources

Michael Cade

Field CTO

Veeam Software

Gunnar Grosch

Principal Developer Advocate

AWS

Thank you!

Follow us!

Join the community hub: