VEEAM kasten

Kubernetes backup and application mobility

Veeam + Kasten

a strong foundation

multi-workloads

– traditional & containerized

multi-environments

- on-premises and/or cloud

multi-data services

– application aware

multi-storage vendors

– no lock-in

Kasten, a Veeam company for Kubernetes-native

Data Protection and Mobility for Kubernetes

Backup & Recovery

Application Mobility

Disaster Recovery

Multi & Hybrid Cloud

Polyglot Persistence

Multi-Tenancy RBAC

k10 for kubernetes data management

battle hardened for day 2 operations

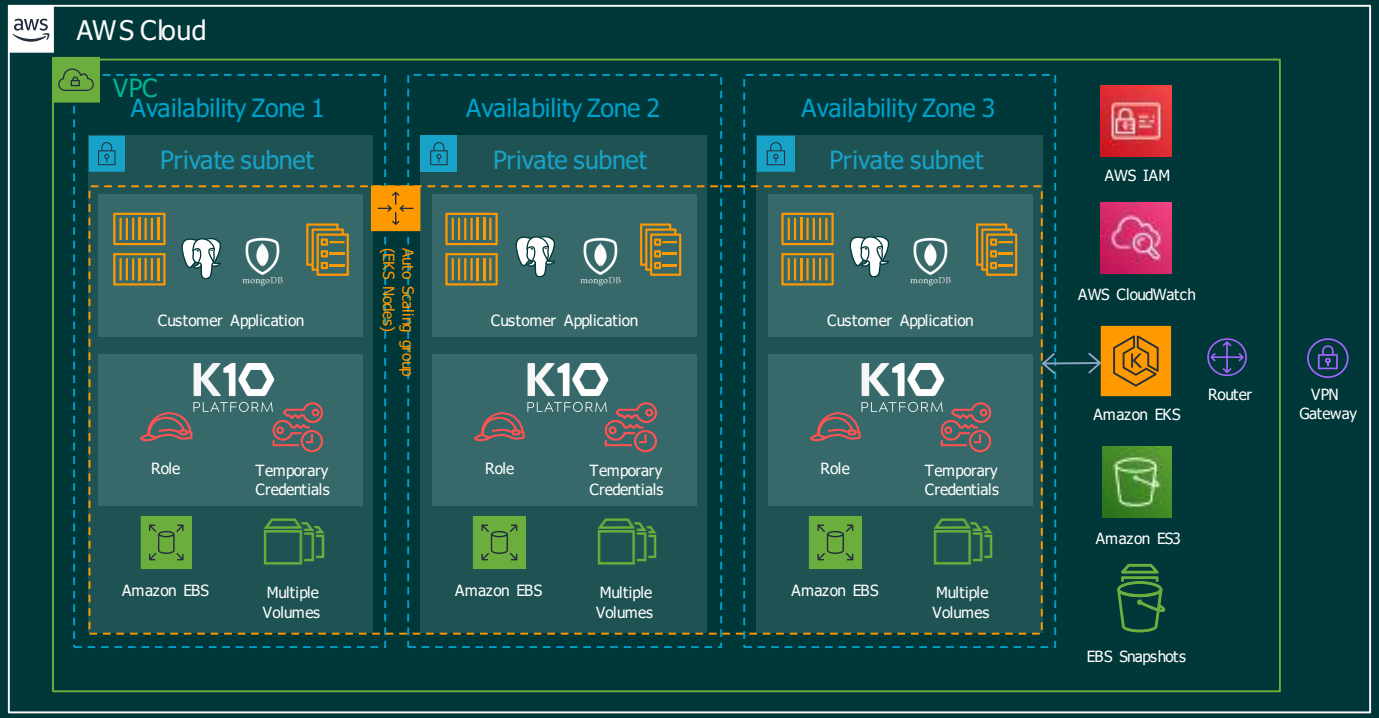

Built for kubernetes

Simplifies platform complexity. Helps scale DevOps teams.

Easy to use

Quick to deploy on-prem and in the cloud. API-First. State-of-the-art UI and UX.

End-to-end security

Enterprise-grade encryption, KMS, IAM, RBAC, Authentication (e.g., OIDC)

Global 10 Customer: Financial Services Use Case: Backup and Disaster Recovery

VEEAM

k10 for kubernetes data management

battle hardened for day 2 scale

kubernetes aws ceph AWS EBS

Number	Component (subset)
2,126	Pods (1,380 workloads)
3,166	Secrets
1,411	Services
3,483	Image Information
768	Service Accounts
915	Configuration
3,484	Role Bindings
5,137	Other Components
18,393	Total (average 112/app)

Sopra Steria: Top 3 French IT Firm
Use Case: Backup and Migration

Devops targeted

Helps scale DevOps teams.
700 dev:2 ops ratio

App mobility

Large migration across clusters (OCP 3.x to 4.x).
Diverse stack (incld. Cobol).

Easy to use

Quick to deploy on-prem and in the cloud. API-First. State-of-the-art UI and UX.

VEEAM

kasten approach: focus on complete application

kubernetes resources and persistent state

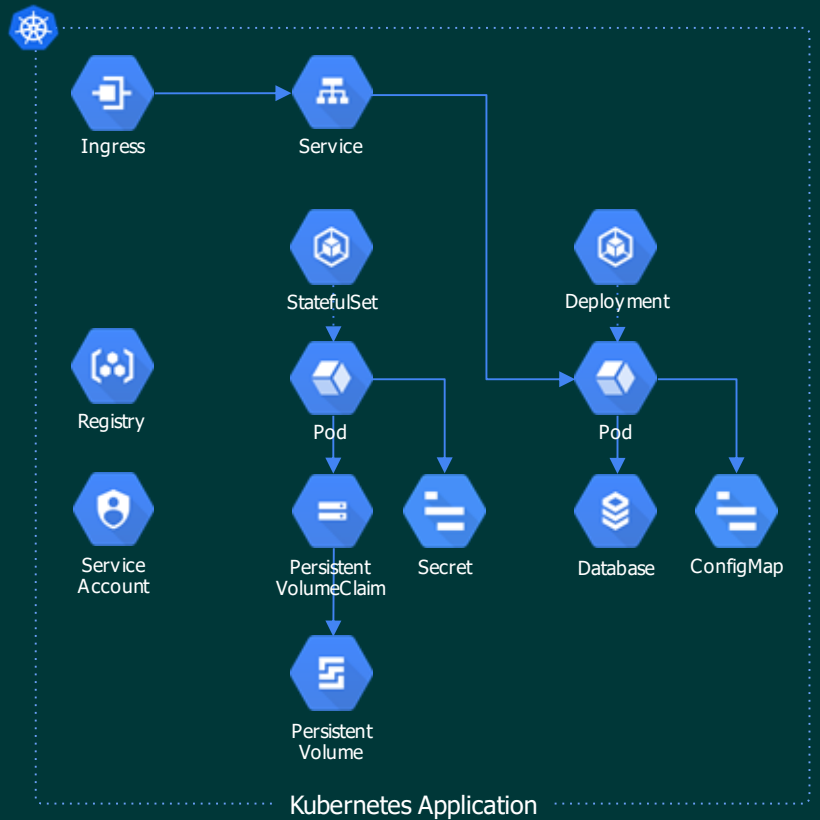

Applications as the operational unit

Automatic and complete application capture
Consistent data and application resources capture
Namespaced objects + non-namespaced dependencies

Abstract underlying infrastructure
Seamless support for storage and data services within and across clusters, regions, and clouds

Perform coordinated operations
Proper sequencing of resource and data operations
Meaningful applications cannot be restored as-is

Unique platform approach:

application-centric data management

Introducing kasten

Software-Only, Easy-to-Use, Secure Data Management
for Cloud-Native Applications

Ops Focused

Simplifies compliance management
Enables policy-based automation
Provides global visibility

Dev Friendly

Simplifies compliance management
Enables policy-based automation
Provides global visibility

Kasten K10

kubernetes backup and mobility made easy

K10 PLATFORM

- Application Discovery
- Policy-driven Automation
- Ease of Use, Simple UX
- End-to-End Security

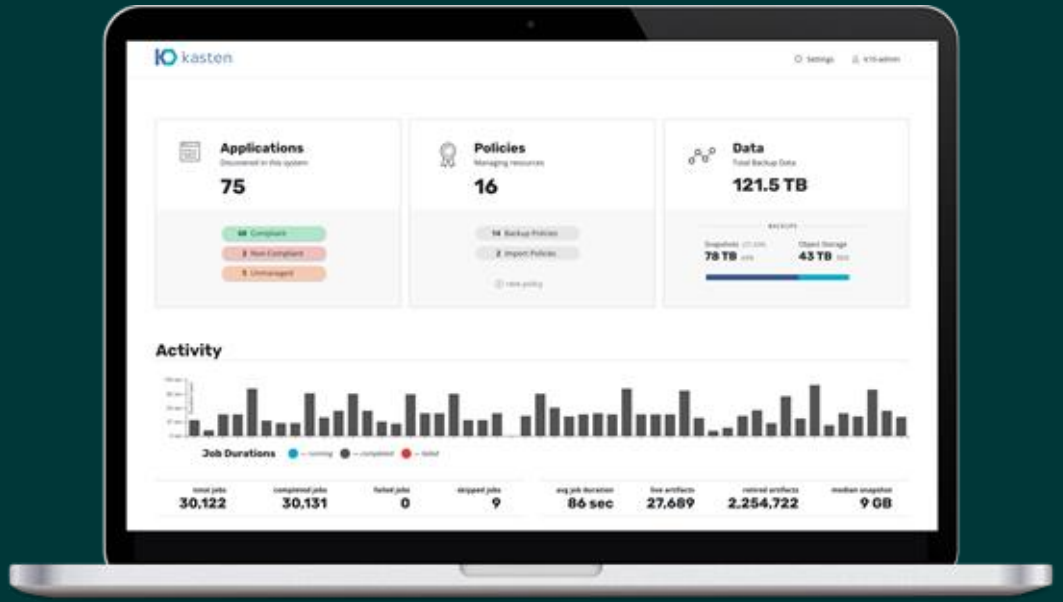

VEEAM

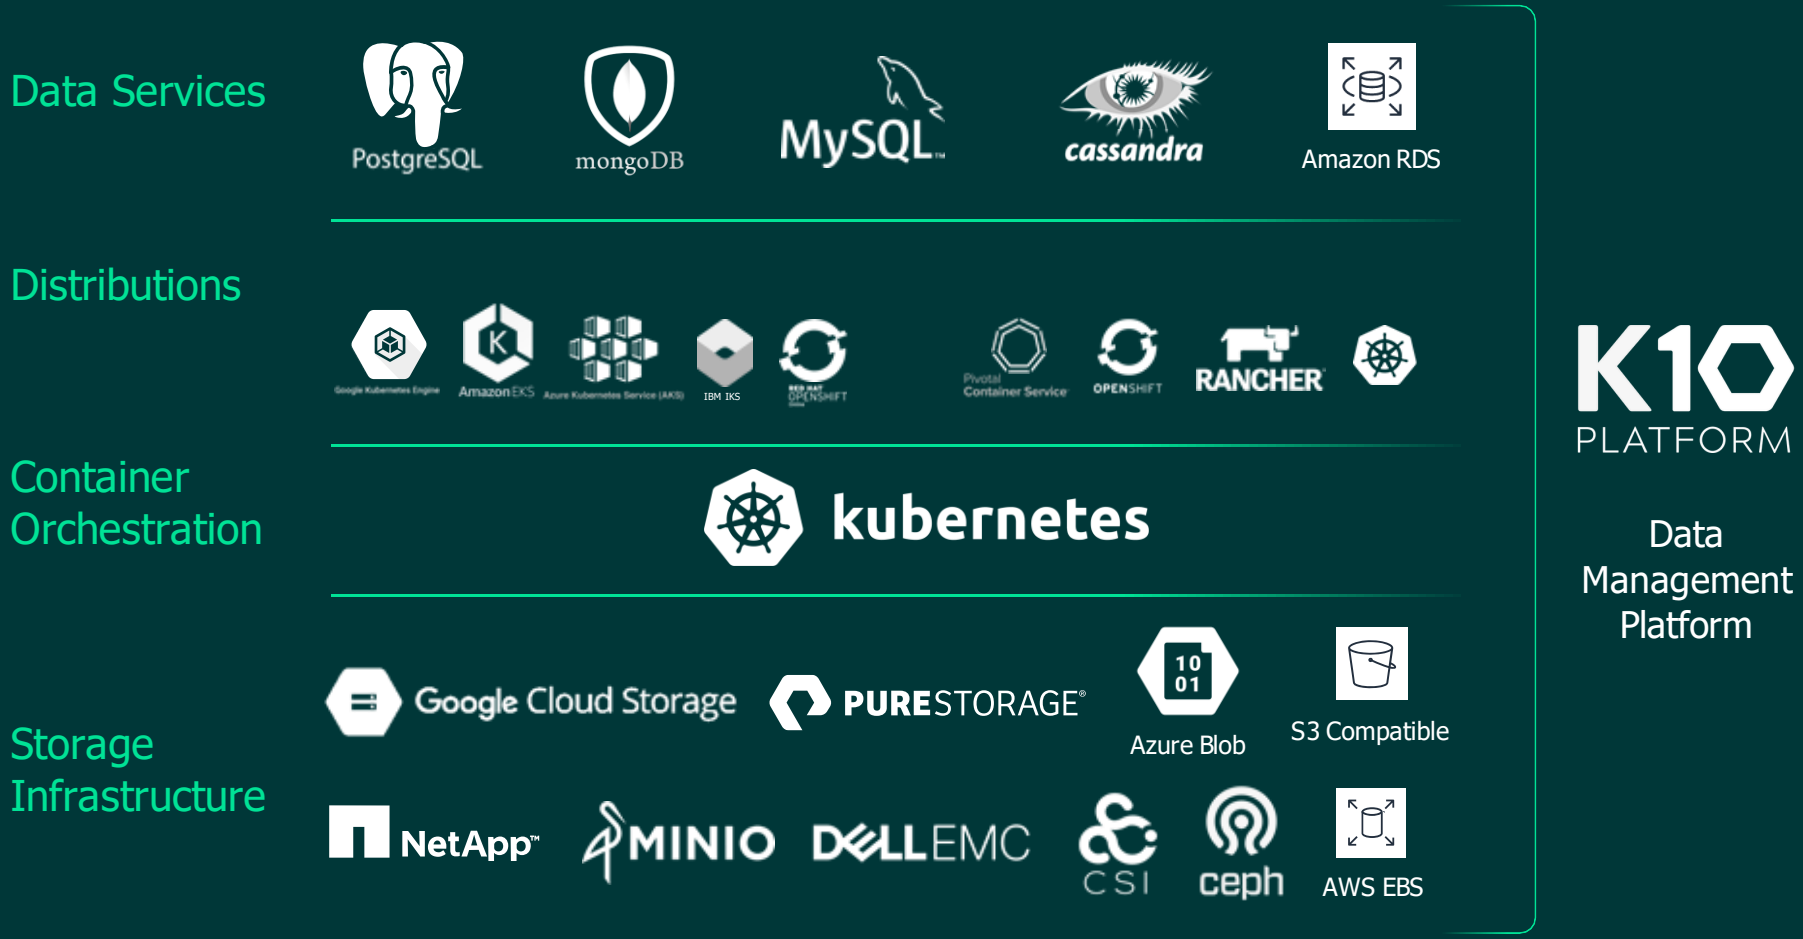

K10 PLATFORM

SELECTED FEATURES

Data Operations

- Volume Snapshots
- Durable Backups
- Change Block Tracking*
- Dedup & Compression
- App-Consistent Backup
- Logical DB Capture
- Managed Data Services
- Log and Replica Capture

Metadata Operations

- Auto App Discovery
- Full Spec Capture
- Spec Transforms
- Global Resource Capture
- Include/Exclude Filters
- Infrastructure Portability
- Global Catalog
- Query API

Backup, DR, Mobility

- Policy-based Operations
- Manual Actions
- GFS Retention
- Independent Schedules
- Application Cloning
- End-to-End Encryption
- Application Hooks
- Blueprint Extensibility

Ops Support

- Enterprise Dashboard
- API-first Design
- Logging Integration
- Monitoring
- Alerting
- Authentication
- RBAC/Self Service
- Air Gap Support
- DR and HA

VEEAM

Solution details

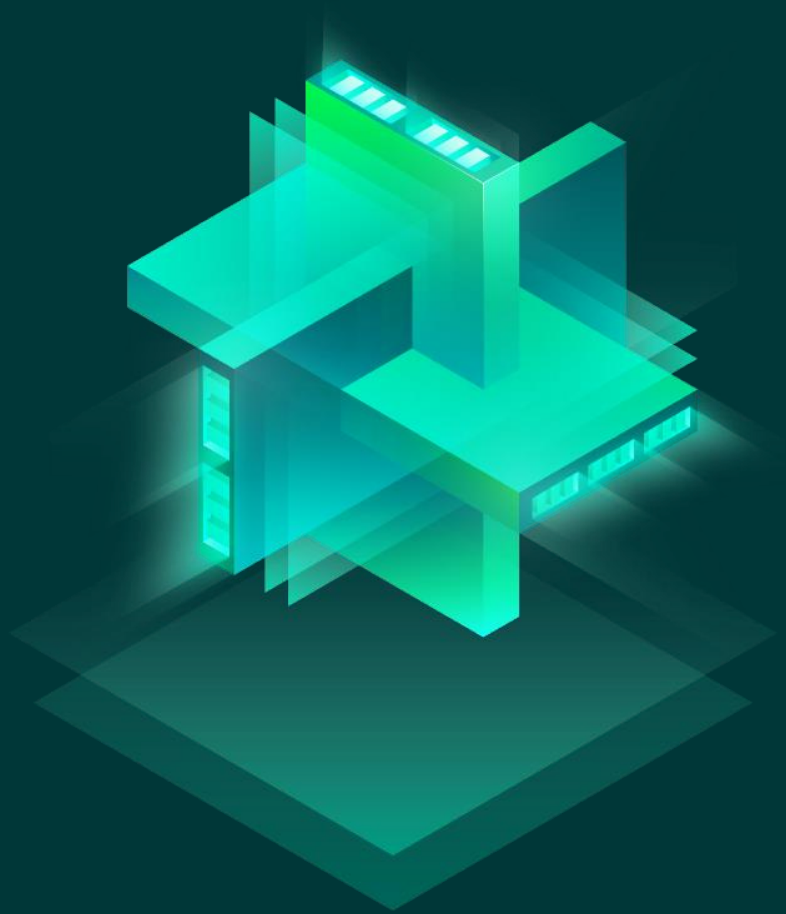

K10 architecture

a high-level overview

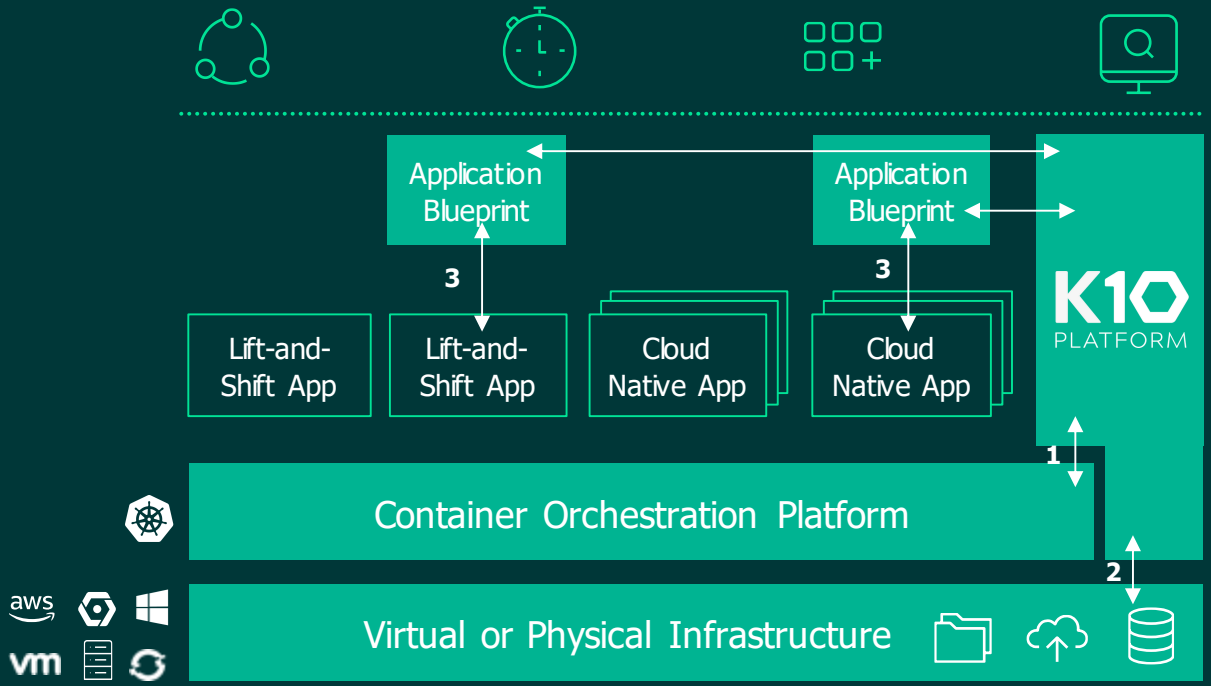

1 Orchestrator APIs

Uses Kubernetes API to discover applications and underlying components and perform lifecycle operations

2 Infrastructure APIs

No proprietary storage layer. Integration with infrastructure specific APIs for:

Block storage provider - Snapshot functionality, snapshot and block copy

Object/file provider - S3-compatible object store or other file storage like NFS for artifacts

3 Application Framework

Optional agentless application-centric hooks can be invoked by easy-to-use blueprints

multi-layer data capture

powerful extensibility, easy to implement

PostgreSQL

MySQL

mongoDB

cassandra

KANISTER

Extensible blueprints

Amazon RDS

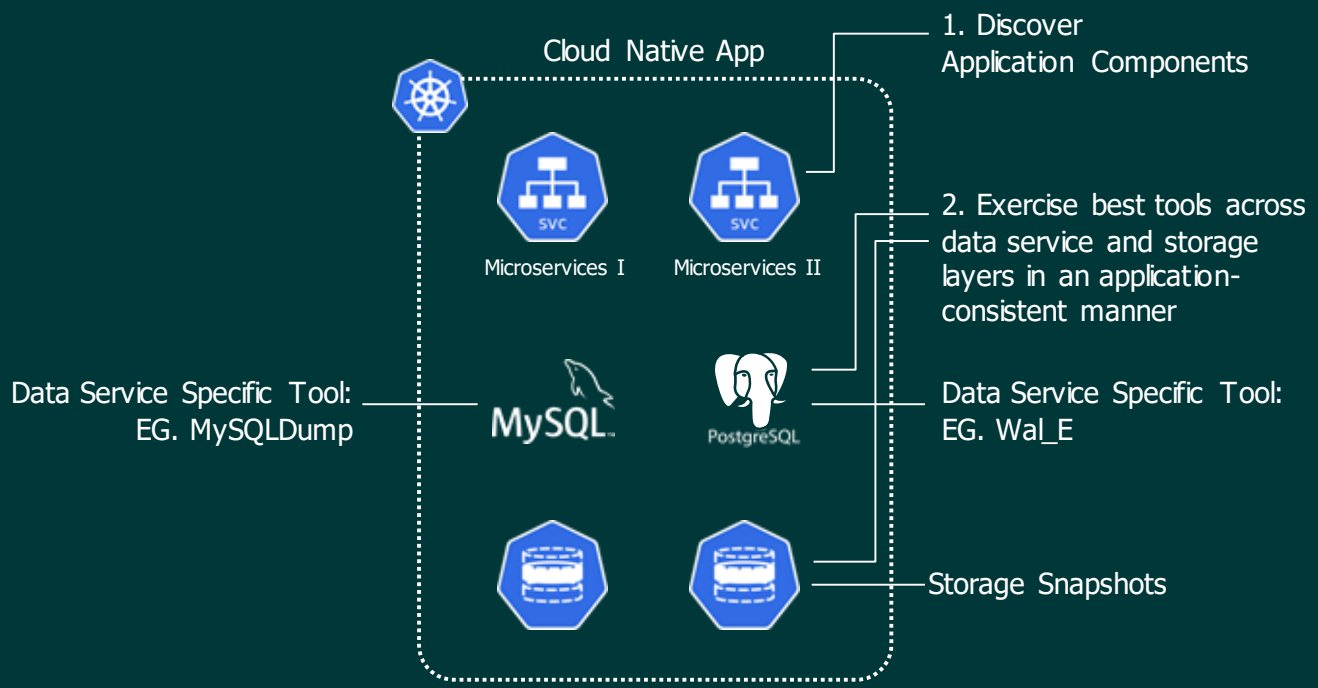

K10 consistency spectrum

range of available options

Crash consistent

Storage snapshots

"App" consistent

Freeze data service
Storage snapshot
Unfreeze data service

DB consistent

Logical dumps via data service-specific tool (e.g., pg_dump)

System consistent

Full app capture
Combination of tools across data and storage layers

K10 workflow

walkthroughs

app capture

Shows application-specific dynamic policy creation with compliance, scheduling, visibility, and auto-discovery

app restore

Illustrates how we generate restore points and restore entire application stacks by repaving infrastructure

app mobility

Demonstrates cloning different application stacks across namespaces, clusters, and clouds

K10 workflow

application capture

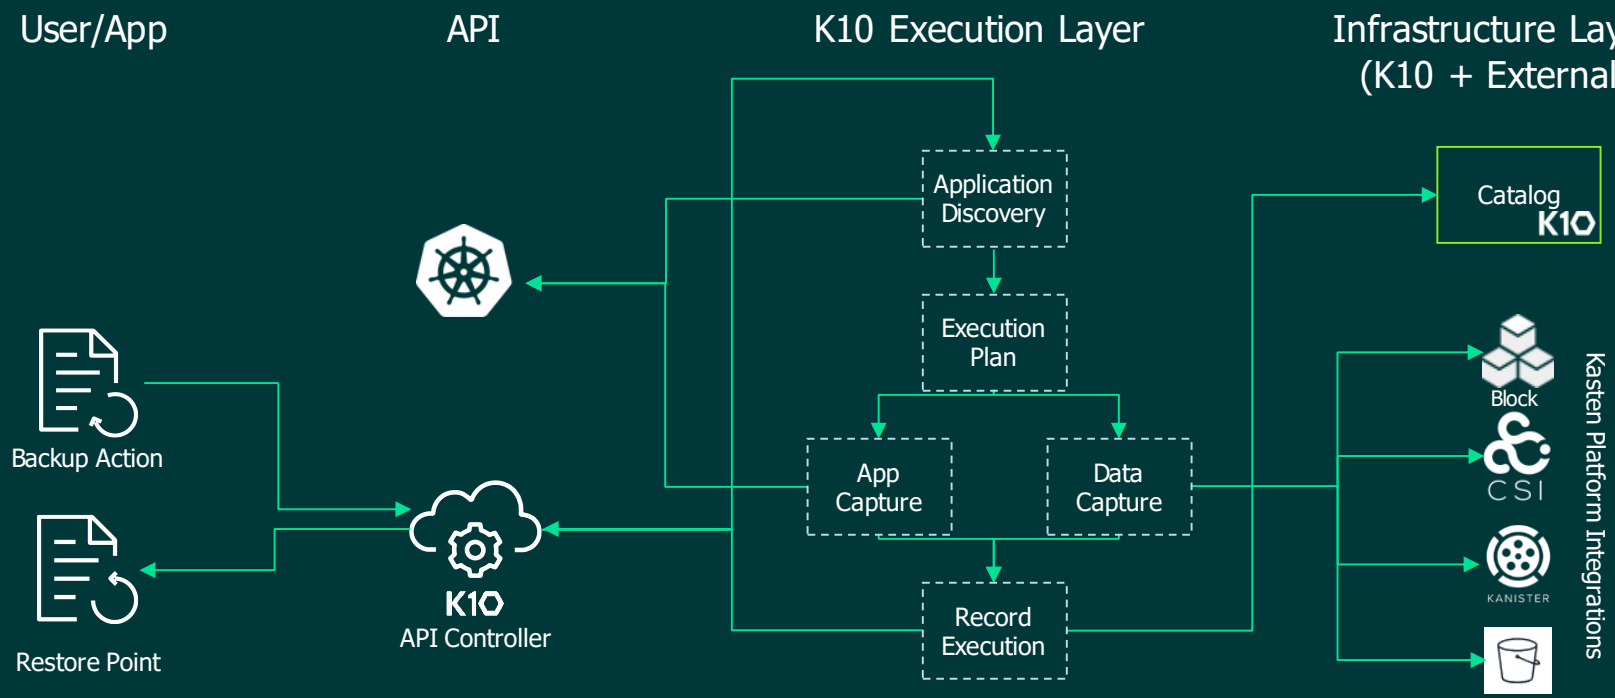

VEEAM

K10 workflow

application restore

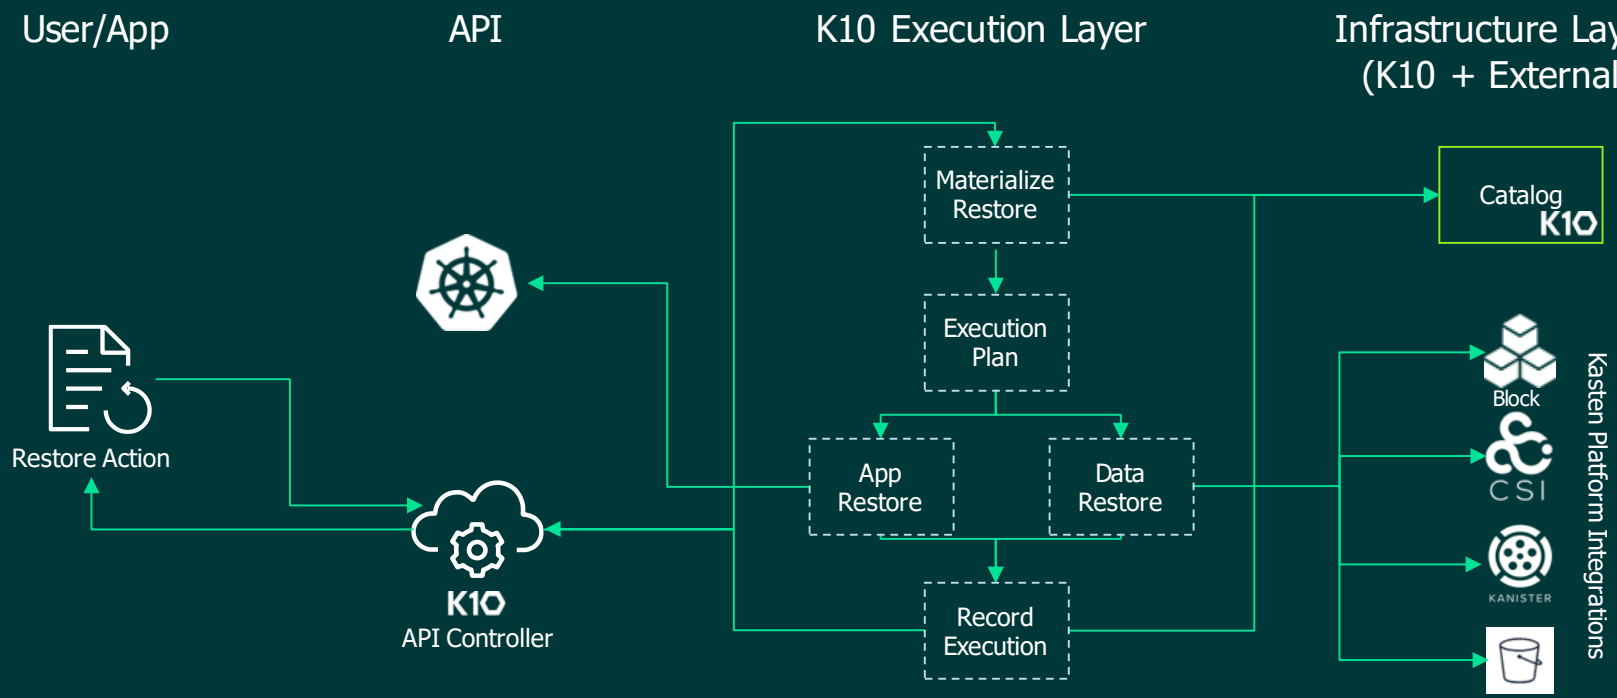

VEEAM

seamless application transformation

application portability requirements

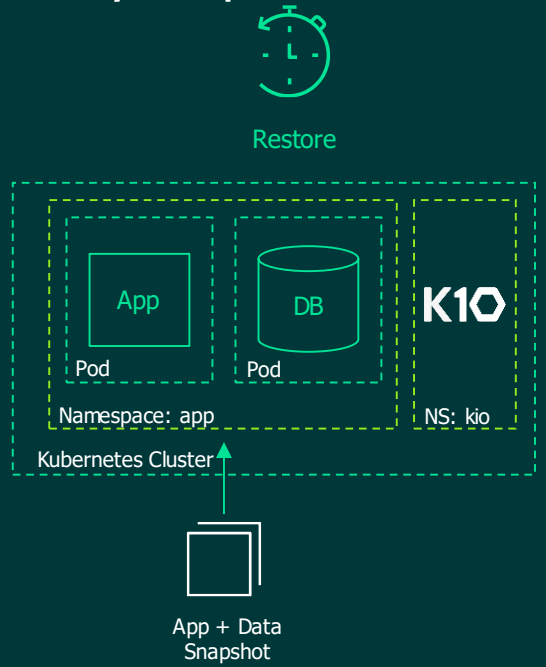

Across:

Storage Systems

Kubernetes Versions

seamless application transformation

application portability requirements

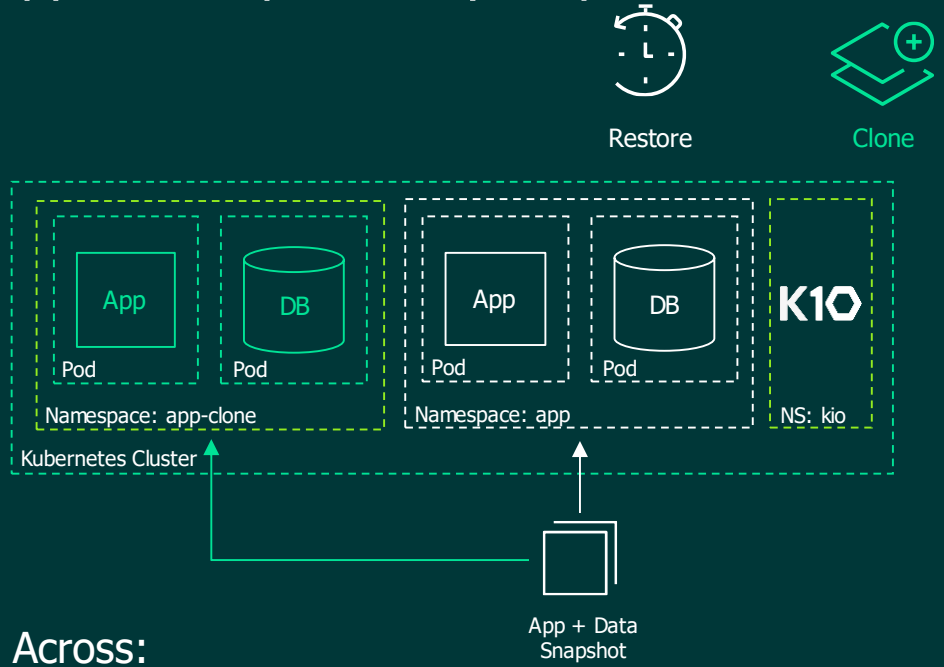

Across:

Storage Systems

Kubernetes Versions

Namespaces

Availability Zones

seamless application transformation

application portability requirements

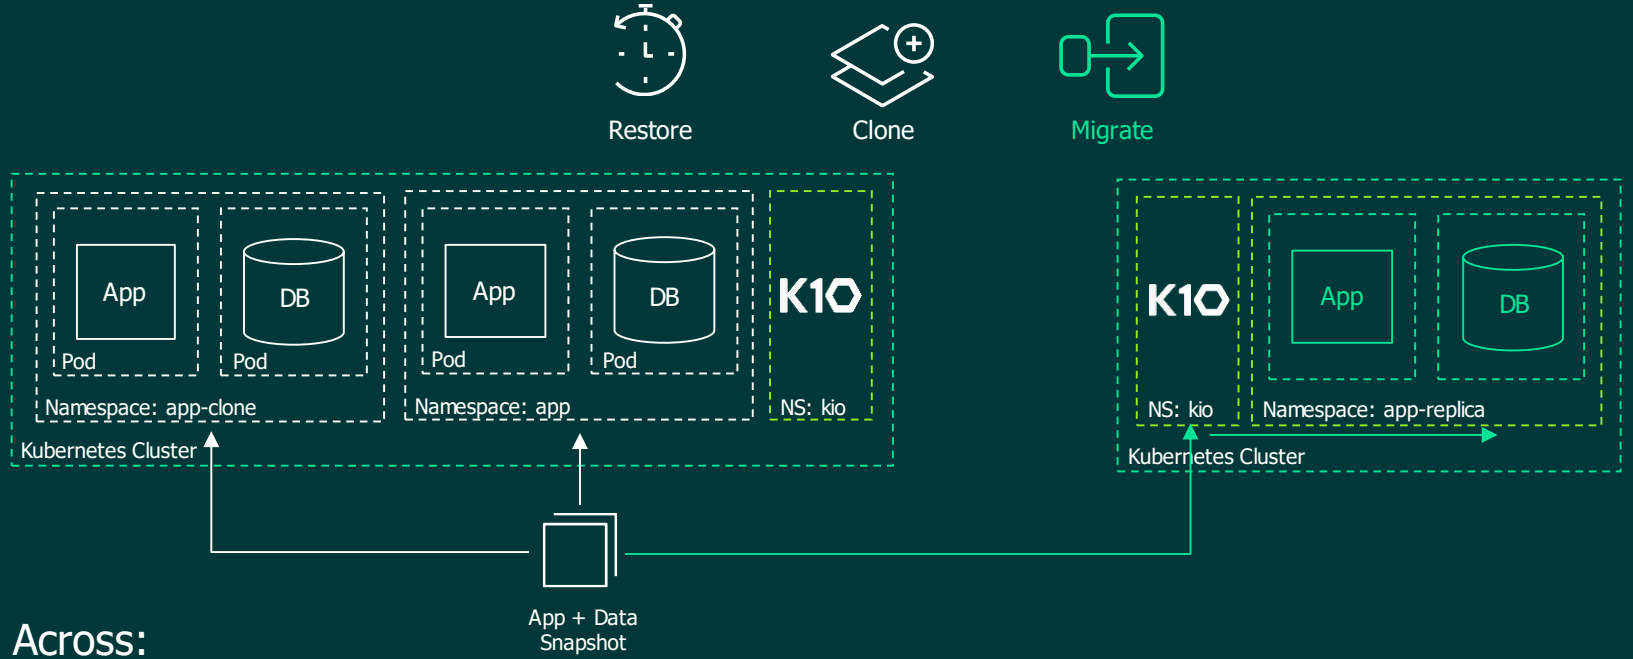

Across:

- Storage Systems
- Kubernetes Versions
- Namespaces
- Availability Zones
- Clusters and Regions
- Cloud Providers
- Kubernetes Distributions
- Prod to Dev Clone

VEEAM

K10: Kubernetes Backup and Mobility Made Easy

Backup & Recovery

Application Mobility

Disaster Recovery

Multi & Hybrid Cloud

Polyglot Persistence

Multi-Tenancy RBAC

Built for Kubernetes

Purpose-built for Kubernetes using cloud-native architectural principles.

Ease of Use

State-of-the-art management interface; cloud-native API, easy install, extensible.

End-to-End Security

Support for RBAC, OIDC, Token Auth, IAM, and industry-standard encryption

Rich Ecosystem

Extensive support across the entire application stack. Select the best tools or infrastructure.

VEEAM

kasten k10 and veeam

span all enterprise data protection requirements

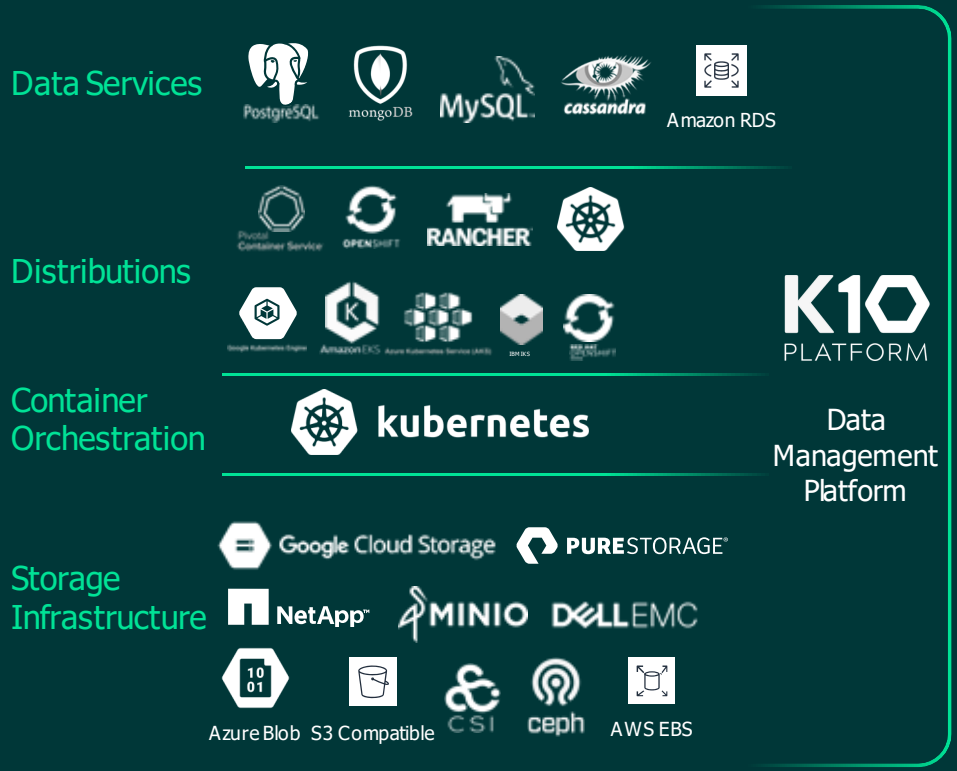

Cloud	SaaS	Virtual	Physical
Microsoft Azure	Office 365	vmware	Windows
VEEAM	OneDrive	Microsoft Hyper-V	Linux
aws IBM Cloud		NUTANIX AHV	ORACLE SOLARIS, AIX

Monitoring & Analytics

Orchestration

Backup & Replication

DataLabs

Universal Storage APIs

Object Storage

NetApp, Hewlett Packard Enterprise, cisco, FUJITSU, NUTANIX, PURESTORAGE, EXAGRID, Lenovo, IBM, DELL EMC

Moving forward with EDM for Kubernetes

Information

E-book - 7 Critical Reasons for Kubernetes
Additional Resources – www.kasten.io/resources
Weekly Demo: https://us02web.zoom.us/s/85354560054

Engage

Schedule a Demo with your Account Team

Deploy

Test Pilot Kasten for Free

VEEAM

kasten by Veeam

Appendix

Multi-workloads
Traditional & containerized

Support workloads running on bare metal, virtual machines and cloud-native platforms – at scale

Multi-cloud
OnPrem and/or in-Cloud

Support multiple clouds, hypervisors, operating systems and Kubernetes distributions

Veeam + Kasten = the whole story

Multi-data services
Application aware

Support software as a service (SaaS), managed services, relational database, NoSQL systems and more

Multi-vendor
No lock-in

Agnostic with special integrations including VMware, Dell EMC, HPE, NetApp, Pure, Amazon, Google, Microsoft

Vms vs. Kubernetes

fundamental platform differences

VMs vs. Kubernetes

Strong impedance mismatch between solutions built for VMs vs. Cloud-Native Platforms

Infra and App Changes

Dynamic autoscaling

Frequent rescheduling

No IP/DNS stability and lack of external visibility

Constant application changes and "repaving"

State and services explosion

User Changes

Application-oriented platforms

Developers owning full stack & infra-as-code

Ops role change focusing more on self-service

Requirement for cloud-native APIs + integration

See https://blog.kasten.io/posts/why-vm-based-data-management-doesnt-work/ for more info

Infra-centric data management

scales poorly and leaves data exposed

Use existing VM-level data protection solutions

- ✓ Data-store snapshots
- ✗ Limited recovery options
- ✗ Weak consistency
- ✗ Complex restore procedure

Let me put together a "quick" script

- ✓ Tailored to application
- ✗ More complex than expected
- ✗ Often tied to infrastructure
- ✗ Difficult to maintain

My storage overlay does backups & migration

- ✗ No fault isolation
- ✗ Lowest common denominator
- ✗ 2X management complexity
- ✗ Performance cost for overlays

K10 ecosystem (selected list)

provides mobility and freedom of choice

Platform Integrations

aws	EKS, EBS, EFS, S3, IAM, RDS
Microsoft Azure	AKS, Azure Stack, Managed Disk, Blob
Google Cloud	GKE, Anthos GPD, GCS
OPENSHIFT	OpenShift 3.12/4.x OCS 4.x (Ceph) OpenShift Extensions
VMware Tanzu	vSphere 6,7u3, 7.x Tanzu/Project Pacific PKS, CNS, FCD
NetApp	Netapp Trident, Cloud: CVO and CSV
PURESTORAGE	Pure Storage: Flash Array and Flash Blade vis PSO
CSI	Linbit, DigitalOcean, etc. v1.0+ compatibility
ceph	Rook, OCS, RBD, CSI
RANCHER	Rancher Kubernetes Engine
	Upstream and certified distros (v1.12+)
MINIO	S3-compatible object storage

Data Services

MySQL
mongoDB
PostgreSQL
elastic
Couchbase
cassandra
FoundationDB
redis

- App-Consistent Backup
- Logical DB Dump
- External Managed Services
- Log and Replica Capture

Community and Kasten driven with support for one or more of the above options

K10 adoption spectrum

zero touch to deep integration

Default operations

Enhanced workflows & application-specific operations

install-only

No developer or application changes required. Easy incremental rollout.

annotations

Minimal deployment changes. Add annotations to spec for tighter integration.

blueprint

Author or extend custom "recipes" for custom application and data workflows.

sidecar

Add sidecar container to application. No core changes but enables additional functionality.

Customer example

north american financial service

IT, Backup, DB admins new to Kubernetes and found it complex

Simplicity

Easy-to-install and use product reduced time-to-market and provided an on-ramp to Kubernetes

InfoSec required scoped roles (RBAC), IAM in AWS, Monitoring

Authentication / Authorization

Native authentication and authorization for APIs and dashboard supported security workflows

Air Gapped Clusters. Data must be encrypted at rest/in flight

End-to-End Encryption

Data and metadata is always encrypted via TLS and AES-256.

Production use case:

1. in-cloud backup

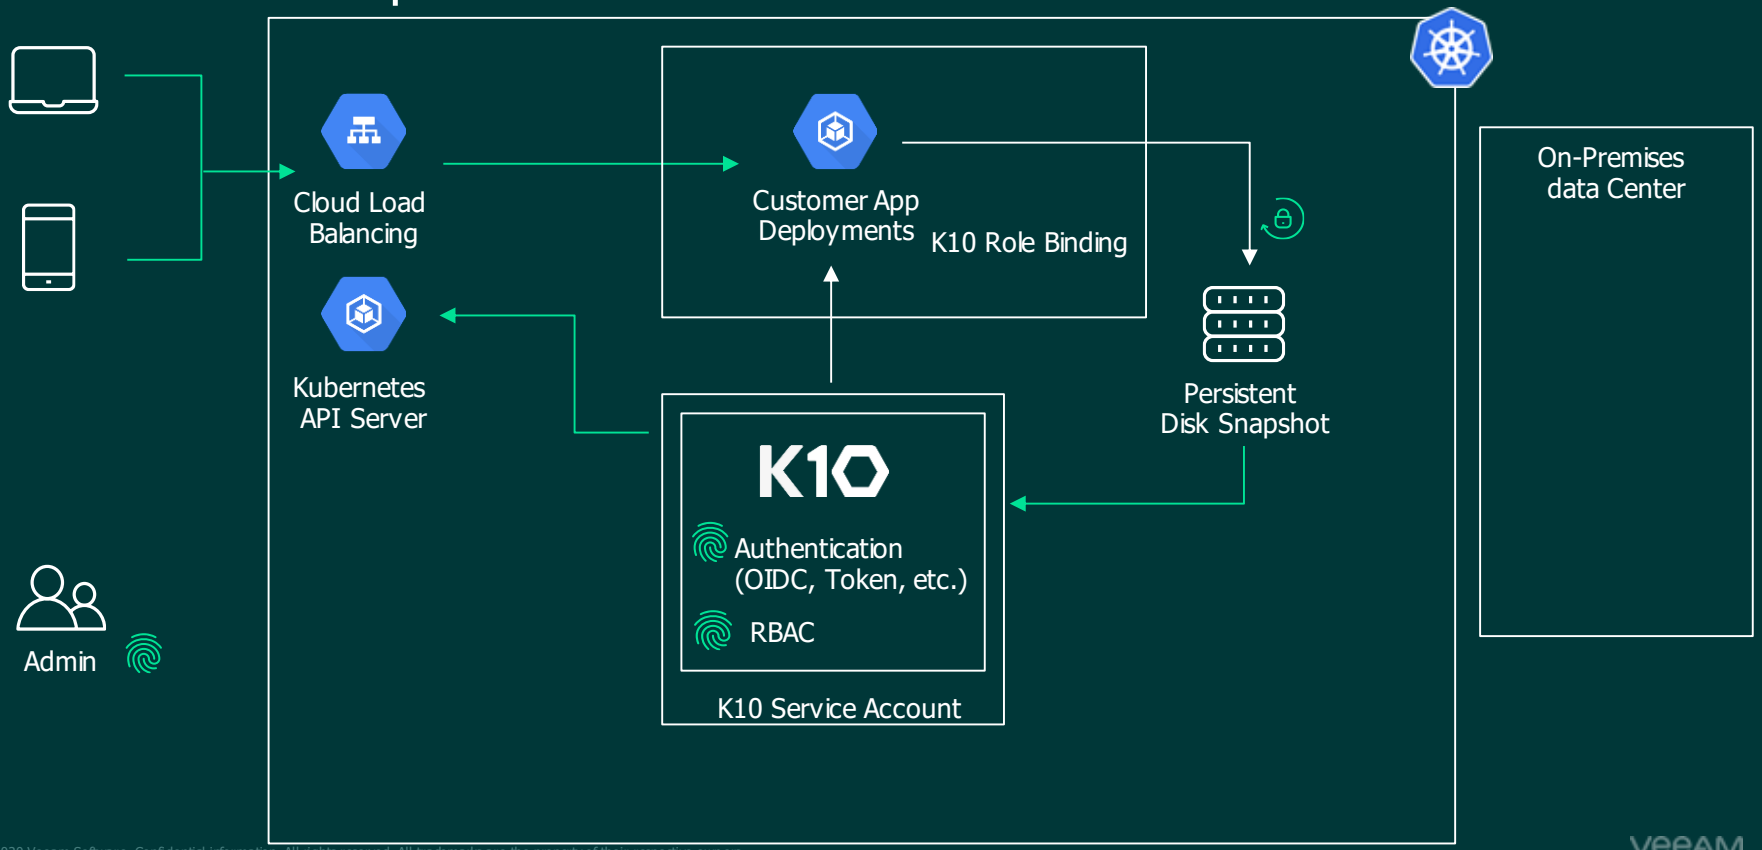

Production use case:

2. hybrid-cloud DR

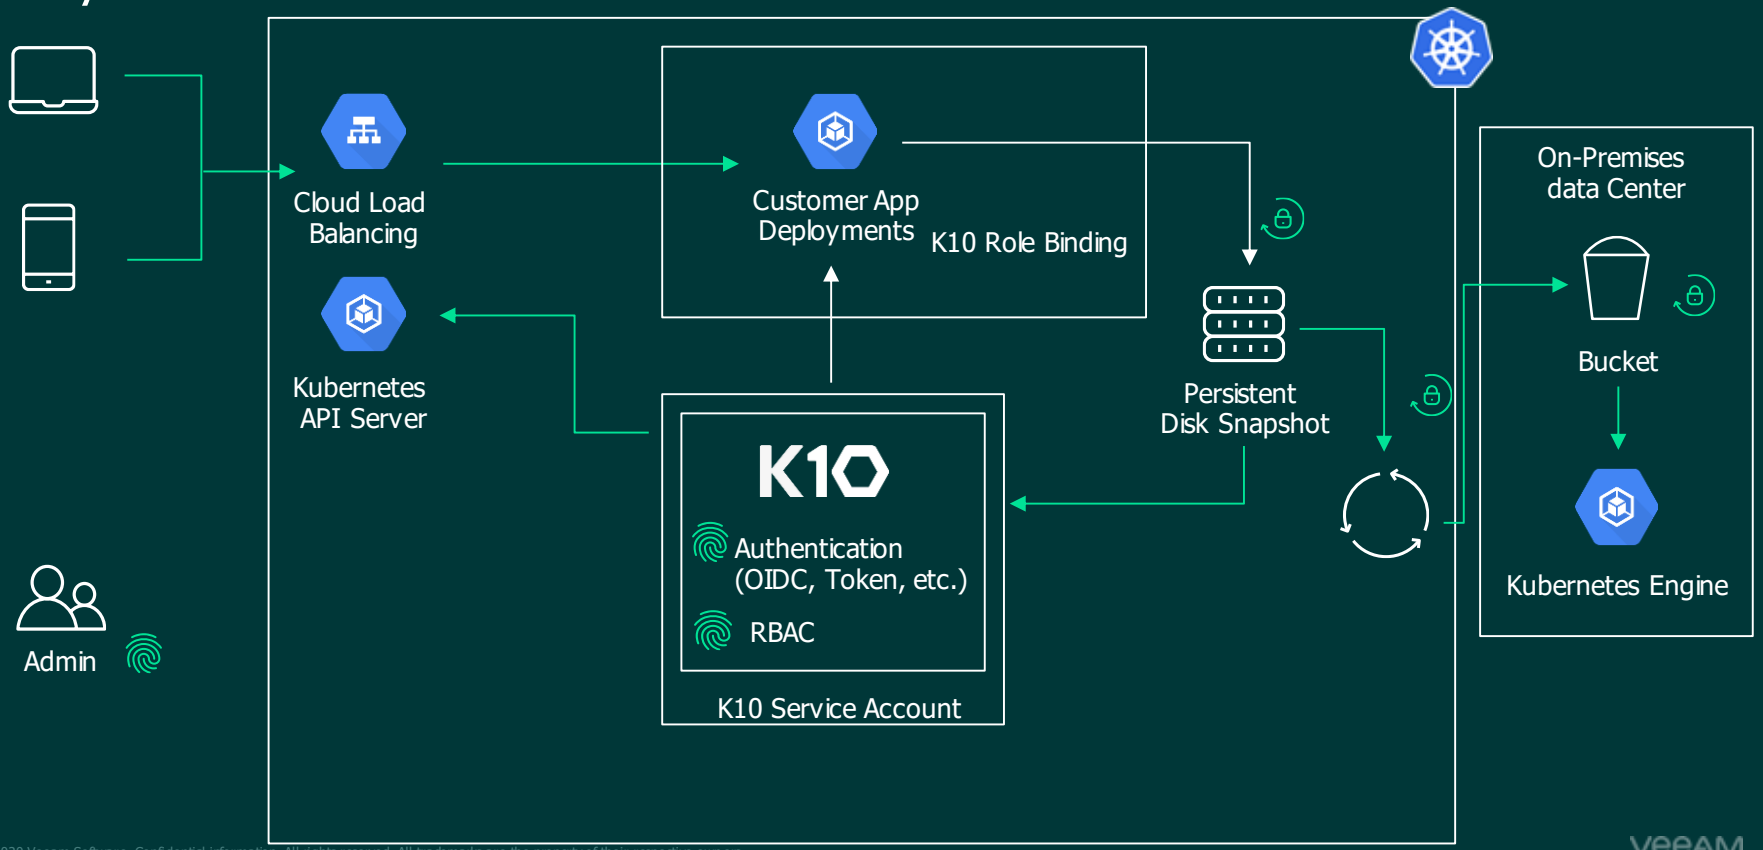

VEEAM

Production use case:

3. ecosystem: monitoring, alerting, logging

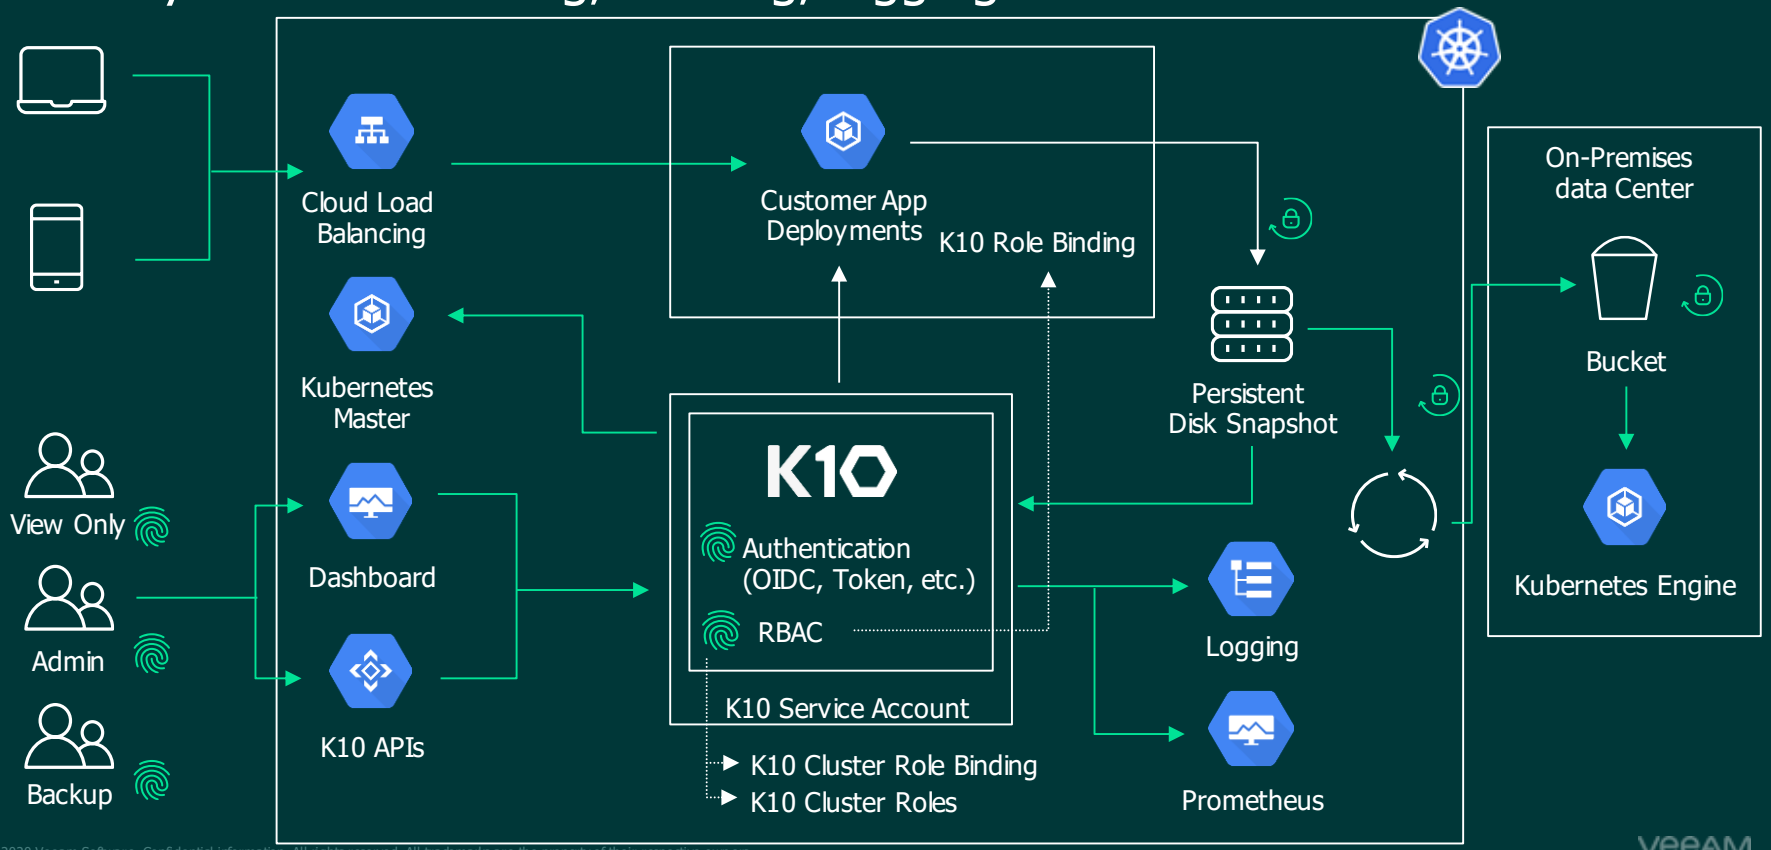

VEEAM