




# 7 Critical Reasons for Kubernetes-Native Backup





# about us



## Mark Severson

Member of Technical Staff @ Kasten

### Previous Life

Senior Software Engineer,  
Team Lead, and Architect @ Instructure



## Niraj Tolia

Co-founder & CEO @ Kasten

### Industry and Academic Storage Experience

(20+ patents filed, ~3000 citations)

Dell EMC (Cloudboost), Maginatics (Distributed  
File Systems), HP Labs (Next Gen Storage),  
Carnegie Mellon (PhD, Deduplication)





**Kubernetes: Myth vs. Reality**  
**Ready for Stateful Applications?**

# 50%

---

of Top 10  
containers running  
in Kubernetes are  
stateful

# 57%

---

of monitored  
Kubernetes clusters  
are running  
StatefulSets

# 70+

---


Kubernetes storage  
drivers available  
for use today!

2020  
Data


55%

of organizations indicate that **half or more** of their container applications are stateful (451 Research)

Stateful Application Type Deployed on Kubernetes




# Does The Old Data Management Way Still Work?





# what's different? deployment patterns



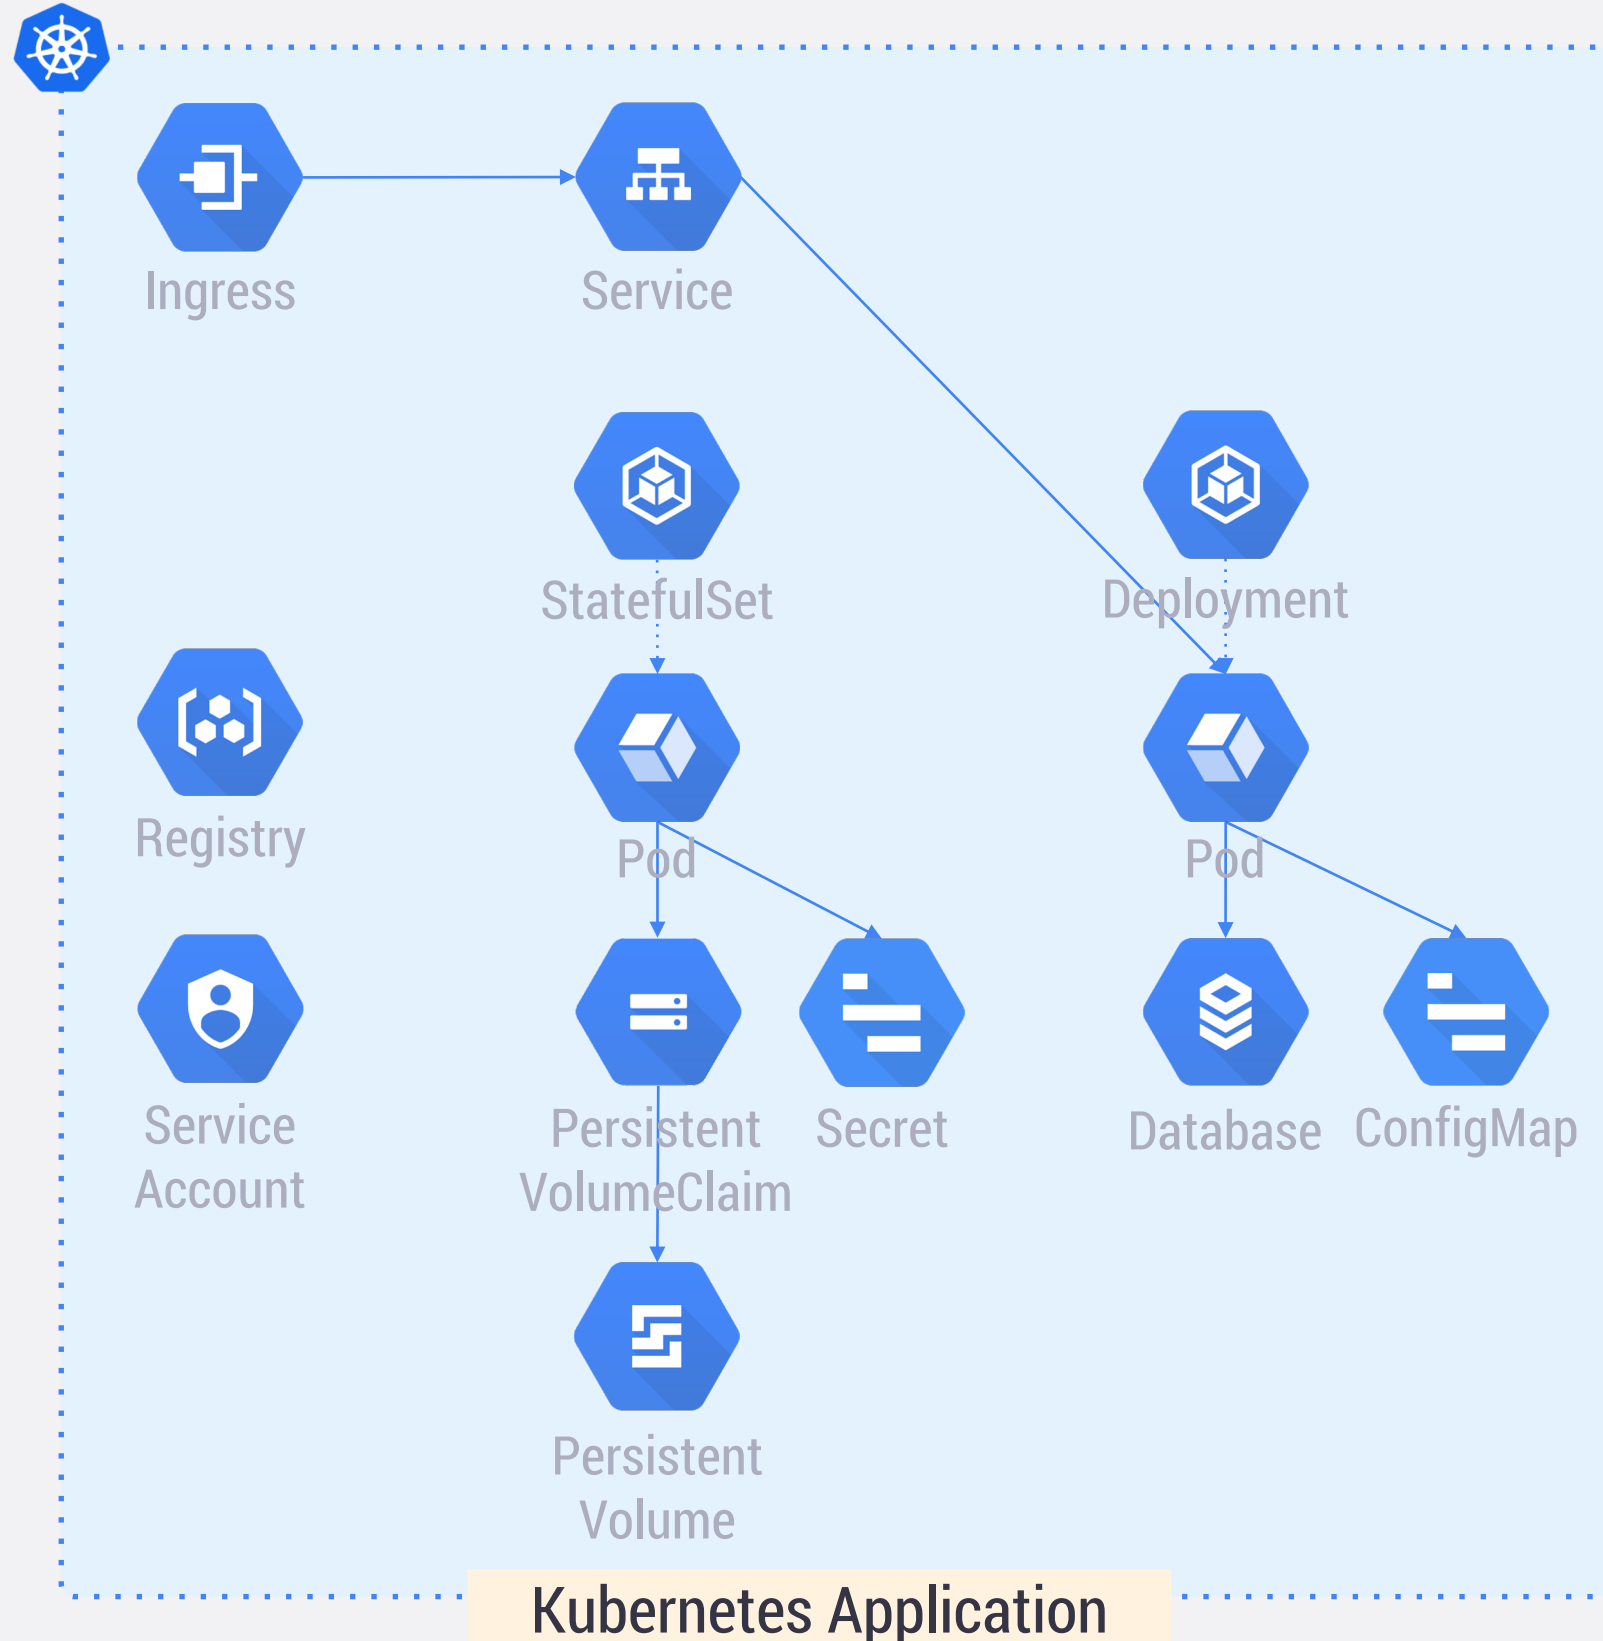
## Kubernetes Deployment Patterns

- No VM  $\leftrightarrow$  Application Mapping,
- Dynamic Rescheduling
- Constant Redeployments
- Multi-Tenant Clusters



# required 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



# what's different? rise of devops




## DevOps and “Shift Left”

- Focus on Applications, Not Infrastructure
- Infrastructure-as-Code
- Self-Service and Dynamic Provisioning
- Greater Accidental Risk



# what's different? scale



## Application Scale

- Explosion in Application Components
- Dynamic Autoscaling (Clusters and Applications)
- Polyglot Persistence
- Multi-Cluster Use Exploding



# kubernetes data management must be hardened for day 2 scale



54 nodes, 216 CPUs, 1.7 TB RAM  
173 Applications/Projects  
Multi-Vendor Storage: 415 Volumes, Multi-TB

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
<b>18,393</b>	<b>Total (average 112/app)</b>




sopra  steria

Top 3 French IT Firm



## DEVOPS RUN

700 dev:2 ops ratio  
Communication challenges



## APPLICATION DIVERSITY



# the old way is infra-centric

## 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





# kubernetes deployment architecture

## a high-level overview



Container Orchestration Platform

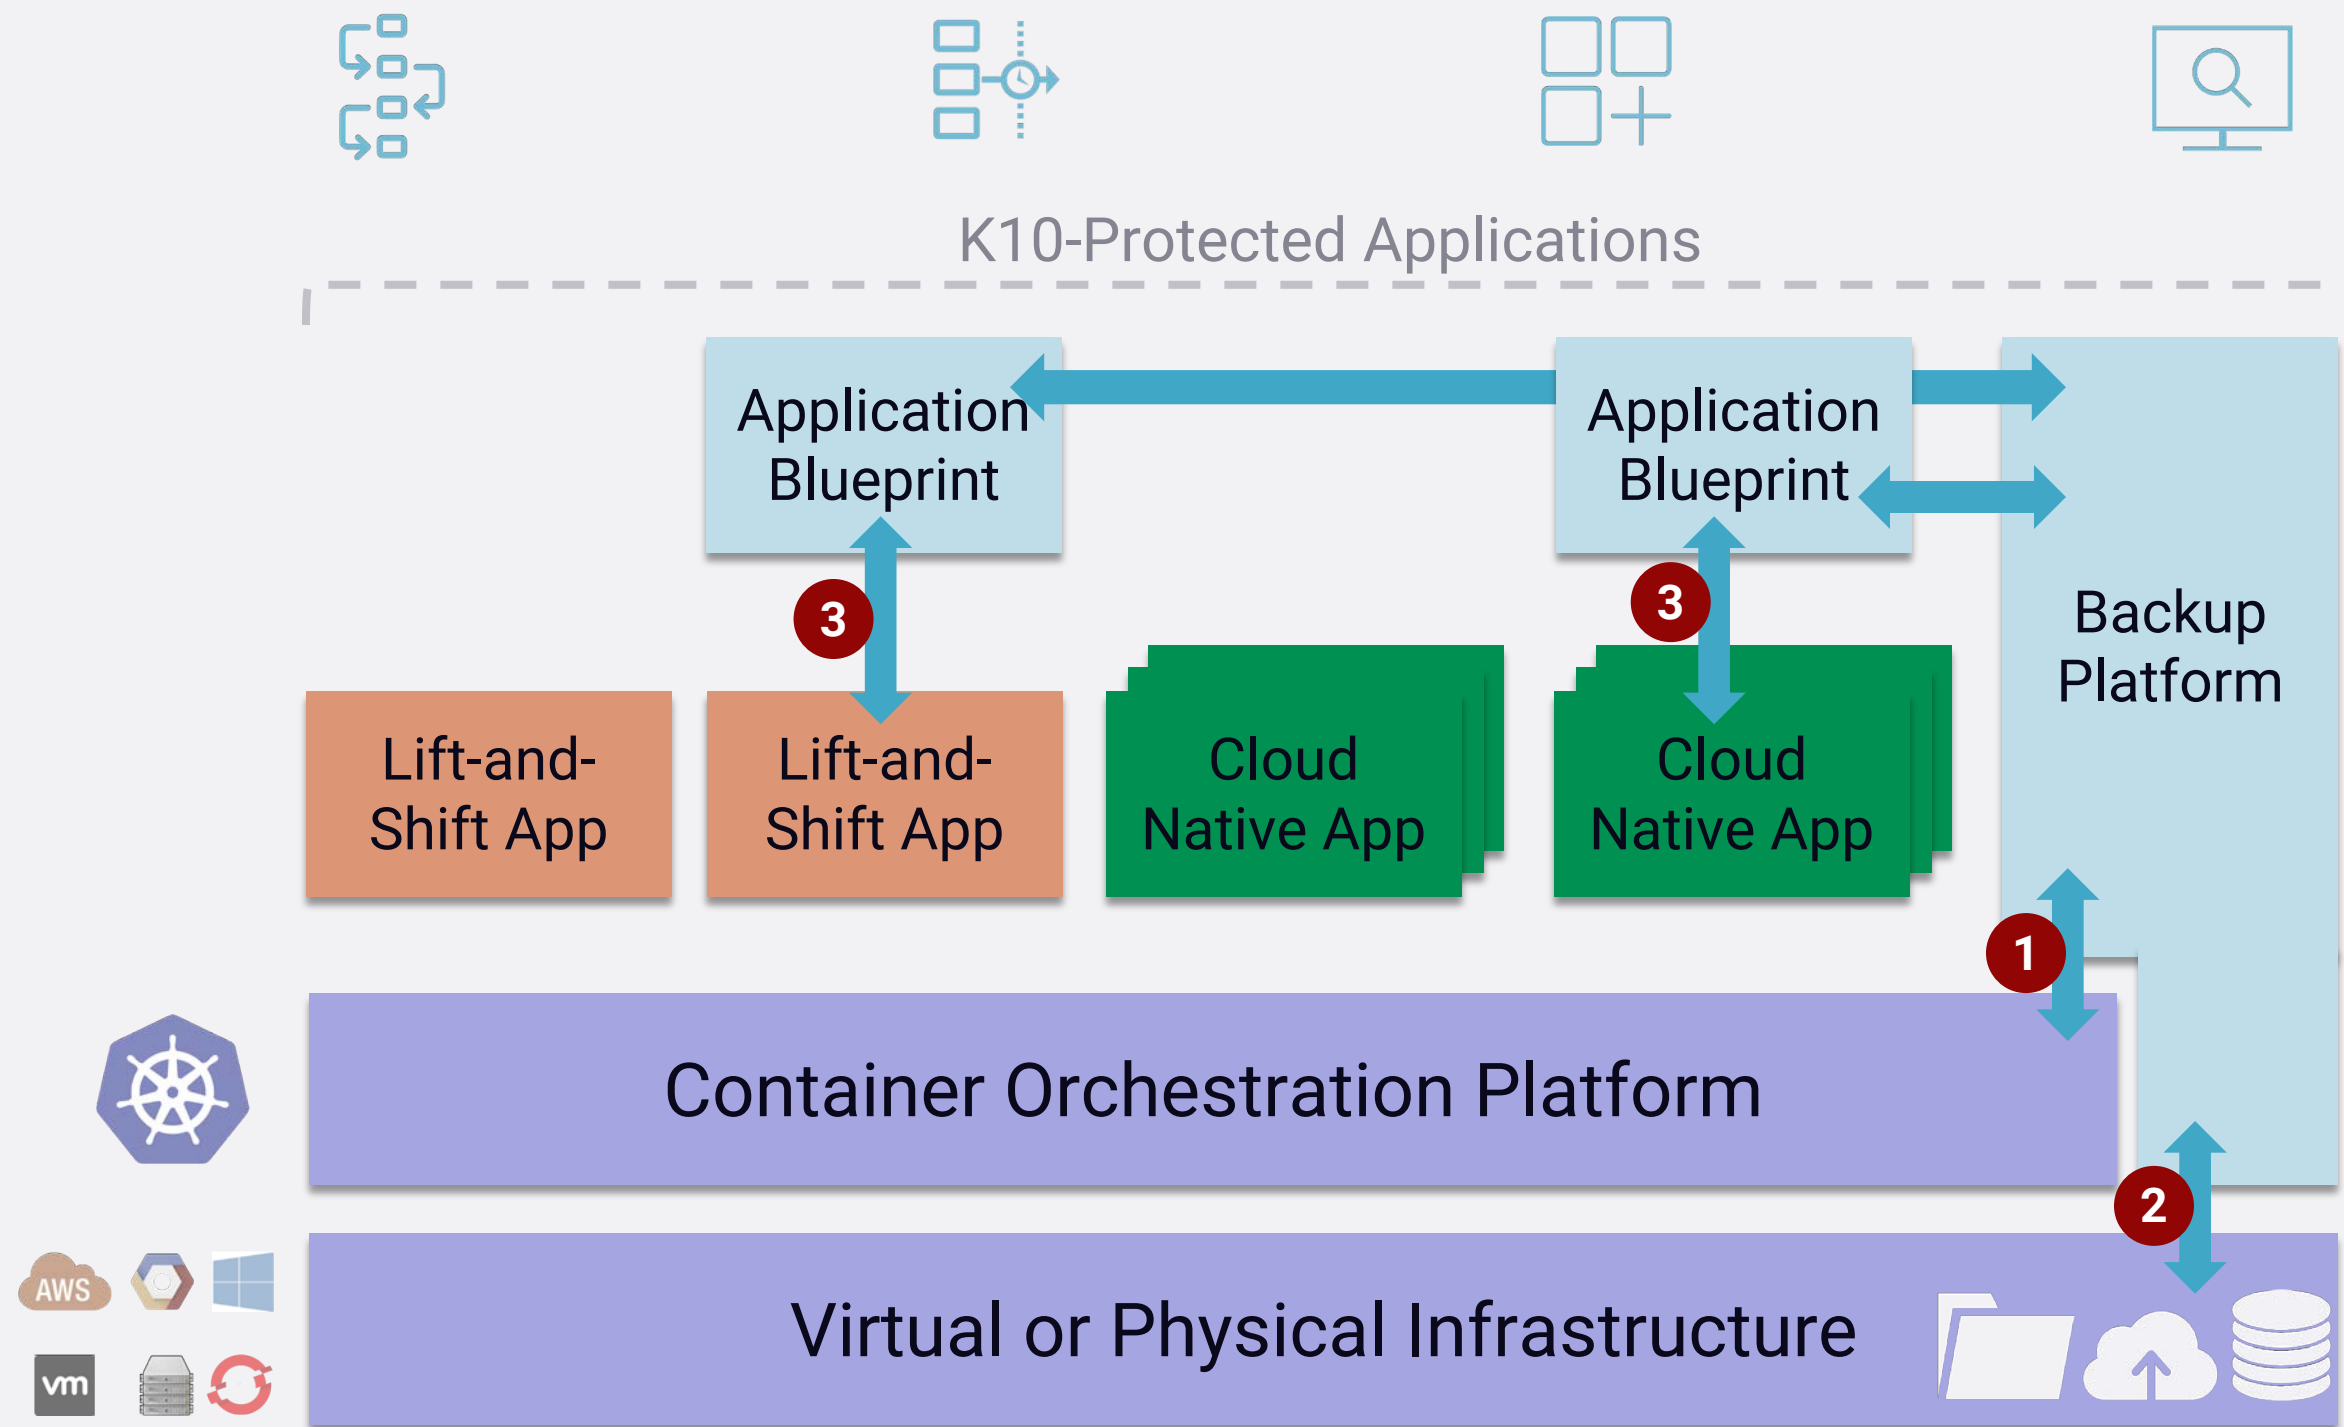


Virtual or Physical Infrastructure



# kubernetes-native backup architecture

## essential touchpoints



### 1 Orchestrator APIs

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

### 2 Infrastructure APIs

**No proprietary storage layer.** Minimal integration with infrastructure specific APIs for the following:

- **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





DEMO




# Other Data Management Concerns to Watch out for





# data management security




## Security and Protection Gaps

- Support for Network Policies
- Authentication (OIDC, Token, etc.) w/ Self-Service
- End-to-End Encryption w/ Customer Managed Keys
- Quick Recovery from Ransomware Attacks





# data management ecosystem integration



## Ecosystem Integration

- Database and NoSQL System Hooks
- Cloud-Native Monitoring and Alerting
- Kubernetes-Native CRD-Based APIs
- Mobility and Freedom of Choice








DEMO

# Summing Up: Things to Look for



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.



# kasten k10

## kubernetes backup and mobility made easy

# K10 PLATFORM

### Use Our Forever-Free Starter Edition!

<https://kasten.io/product>  
[info@kasten.io](mailto:info@kasten.io)

