Kubernetes and Fibre Channel: A Compelling Case

Live Webcast
January 20, 2022
10:00 AM PT/1:00 PM ET

FIBRE CHANNEL INDUSTRY ASSOCIATION

About the Presenters

Mark Jones
Director Technical Marketing, Broadcom
FCIA Board of Directors

Scott Sinclair
Practice Director
Senior Analyst
Enterprise Strategy Group

Dr. Chip Copper
Senior R&D Engineer
Broadcom

Nishant Lodha
Storage Technology Evangelist
Marvell

Matt LeVan
Technical Product Manager
IBM Storage

About the FCIA

The Fibre Channel Industry Association (FCIA) is a mutual benefit, non-profit, international organization of manufacturers, system integrators, developers, vendors, and industry professionals, and end users.

FCIA
FIBRE CHANNEL INDUSTRY ASSOCIATION

About the Fibre Channel Industry Association (FCIA)

25+ Years
Promoting Fibre Channel Technology

Industry Leading
Member Companies

150+ Million FC Ports
Shipped Since 2001

Agenda

- Cloud market trends
- How organizations are deploying cloud
- Introduction to containers and persistent storage
- The role of Fibre Channel in Kubernetes
- Evolution of container storage interfaces and CSI
- Fibre Channel fabric requirements and considerations

2021 Data Infrastructure Trends Survey

State of Data Storage, Cloud, and the Impact of Containers

Scott Sinclair, Practice Director

Data Fuels and Defines Modern Business

59% Identify Data as their Business

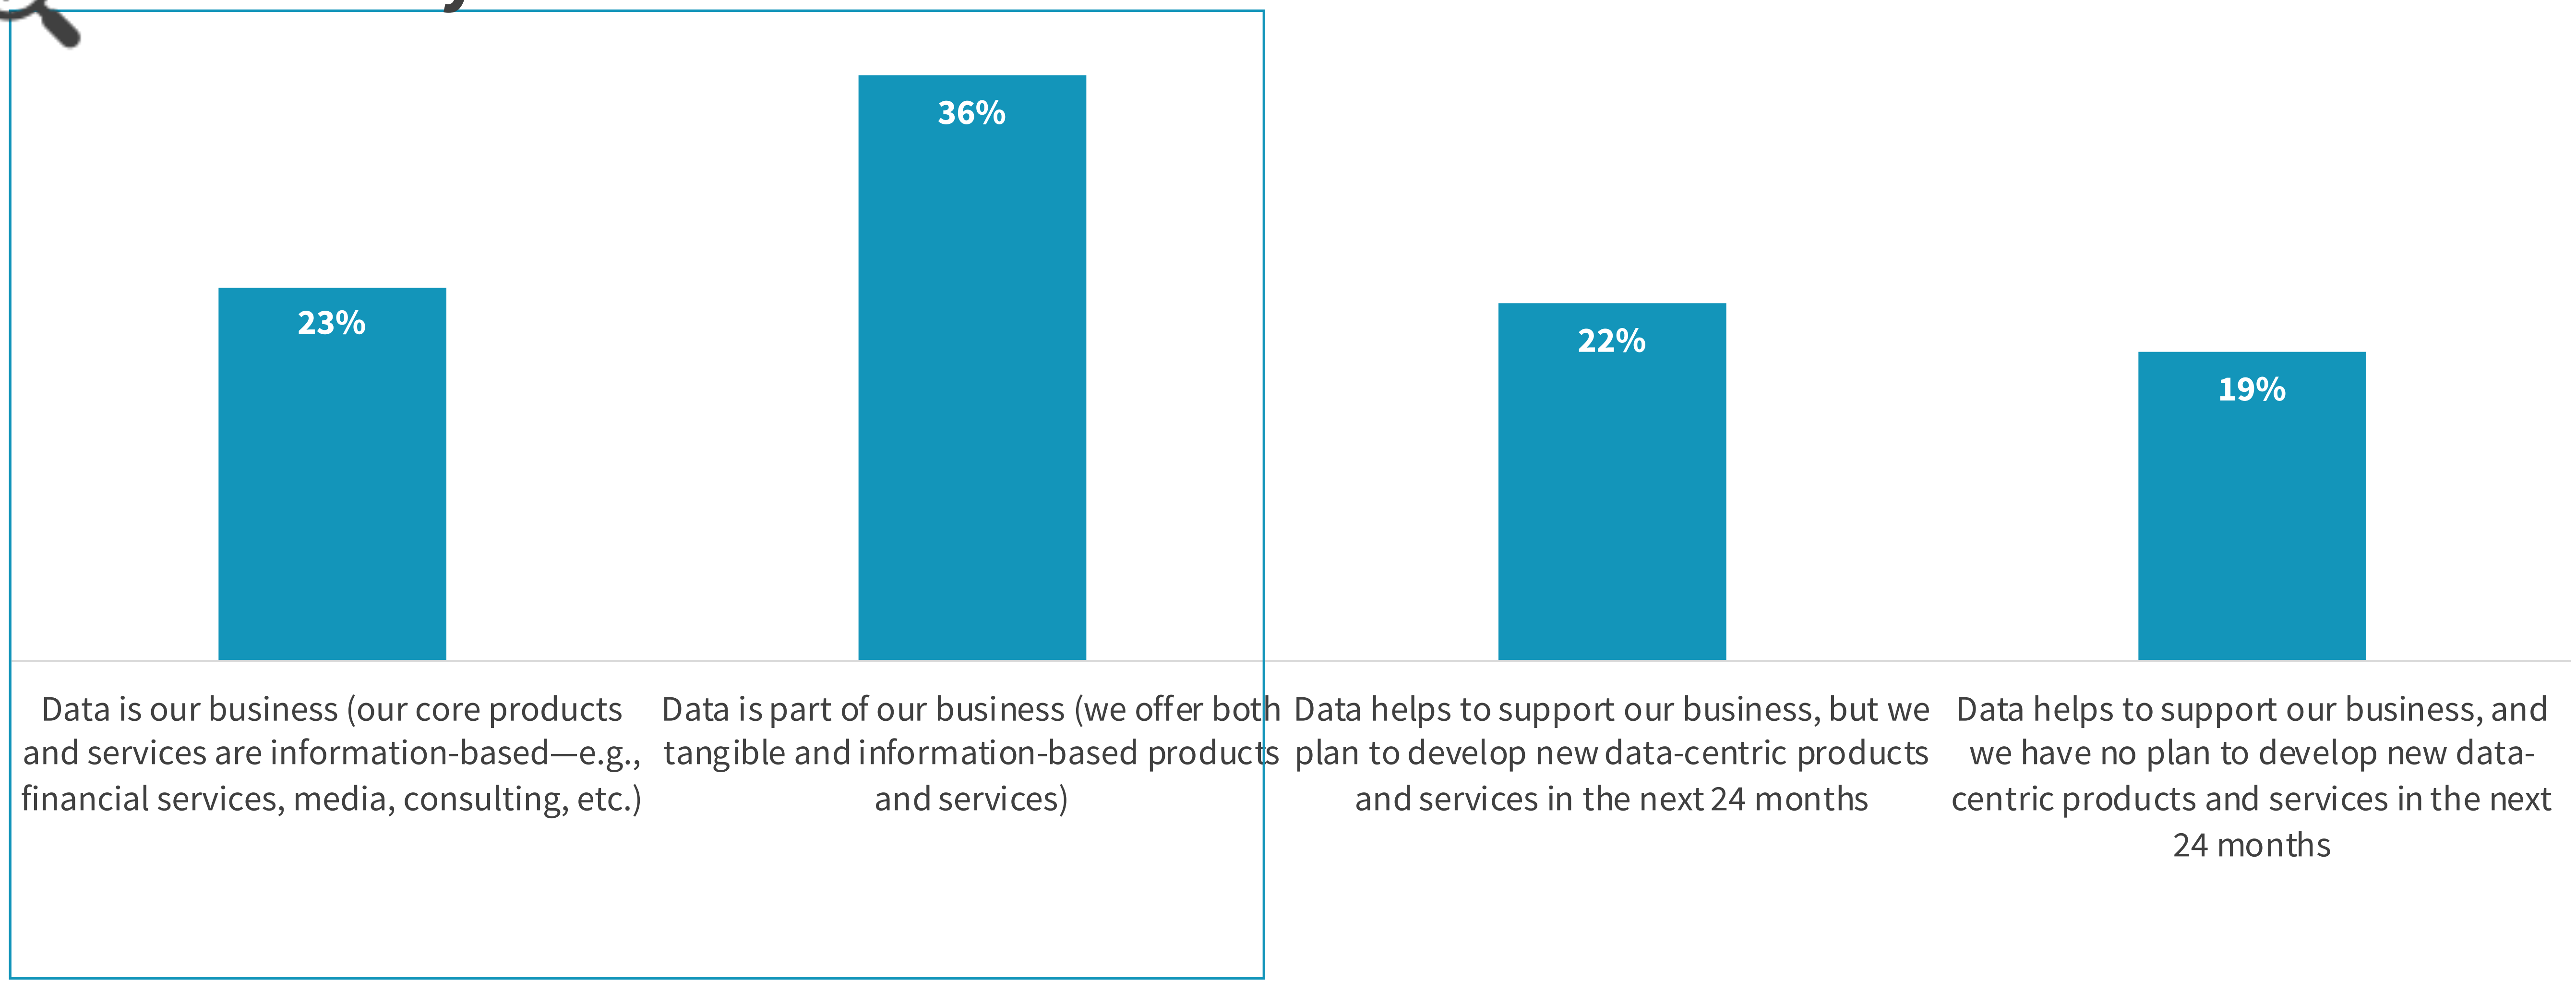

Question text: Which of the following statements best describes your organization's perspective on data? (Percent of respondents, N=359)

Sizing the Data Storage Landscape

48% report **1 PB or more** of active data storage across their entire environment. (Mean of 6.7PB)

44% of data stored **on-premises** in their data center(s), on average.

35% annual growth for **on-premises capacity** on average. (double in less than 3 years)

39% annual growth for **public cloud capacity** on average. (double in little over 2 years)

The Complexity of IT Slows Operations & Fuels Need to Simplify

67% agree we are under pressure to accelerate IT infrastructure provisioning/deployment to support developers/LoB teams.

64% agree that the complexity of our IT infrastructure slows IT operations and digital initiatives.

64% agree that data storage infrastructure requirements and spending are hard to predict for our organization.

63% agree that it is often challenging to properly size workloads for the optimal infrastructure environment.

64% view data center design as strategic and can lead to a competitive advantage

Question text: Please rate your level of agreement with the following statements related to the data storage/management environment at your organization. (Percent of respondents, N=359)

Top Challenges with Persistent Storage for Containers

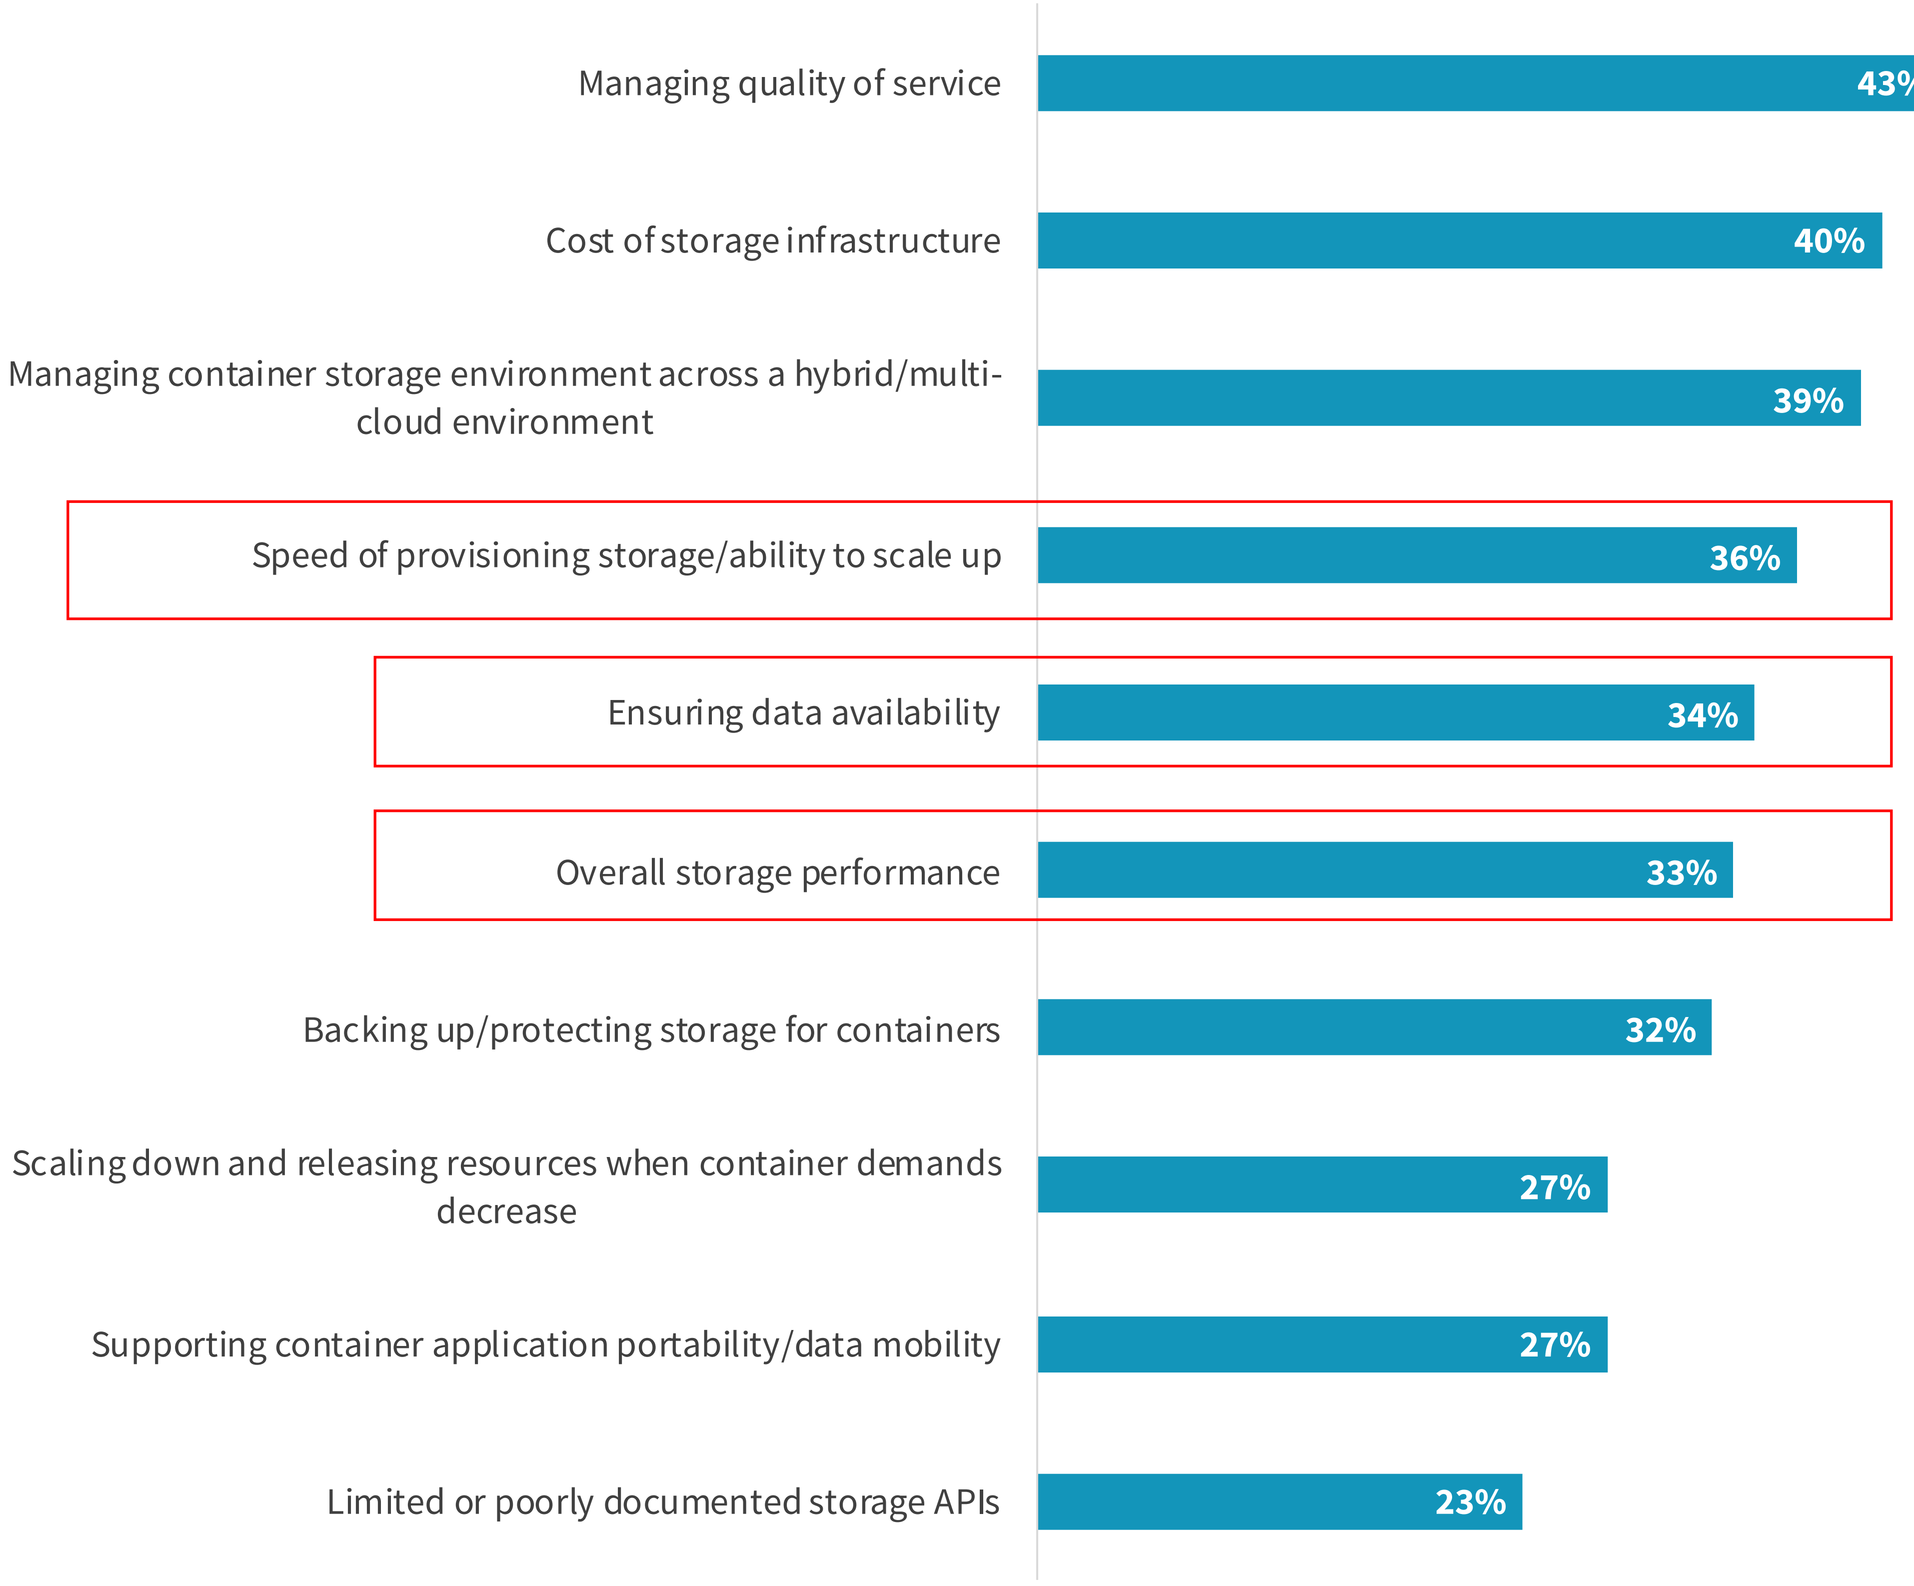

Question text:
In general, what would you say are your organization's biggest challenges with persistent storage for its container-based environment? (Percent of respondents, N=288, multiple responses accepted)

57% of IT Organizations Repatriated Workloads Back from the Public Cloud

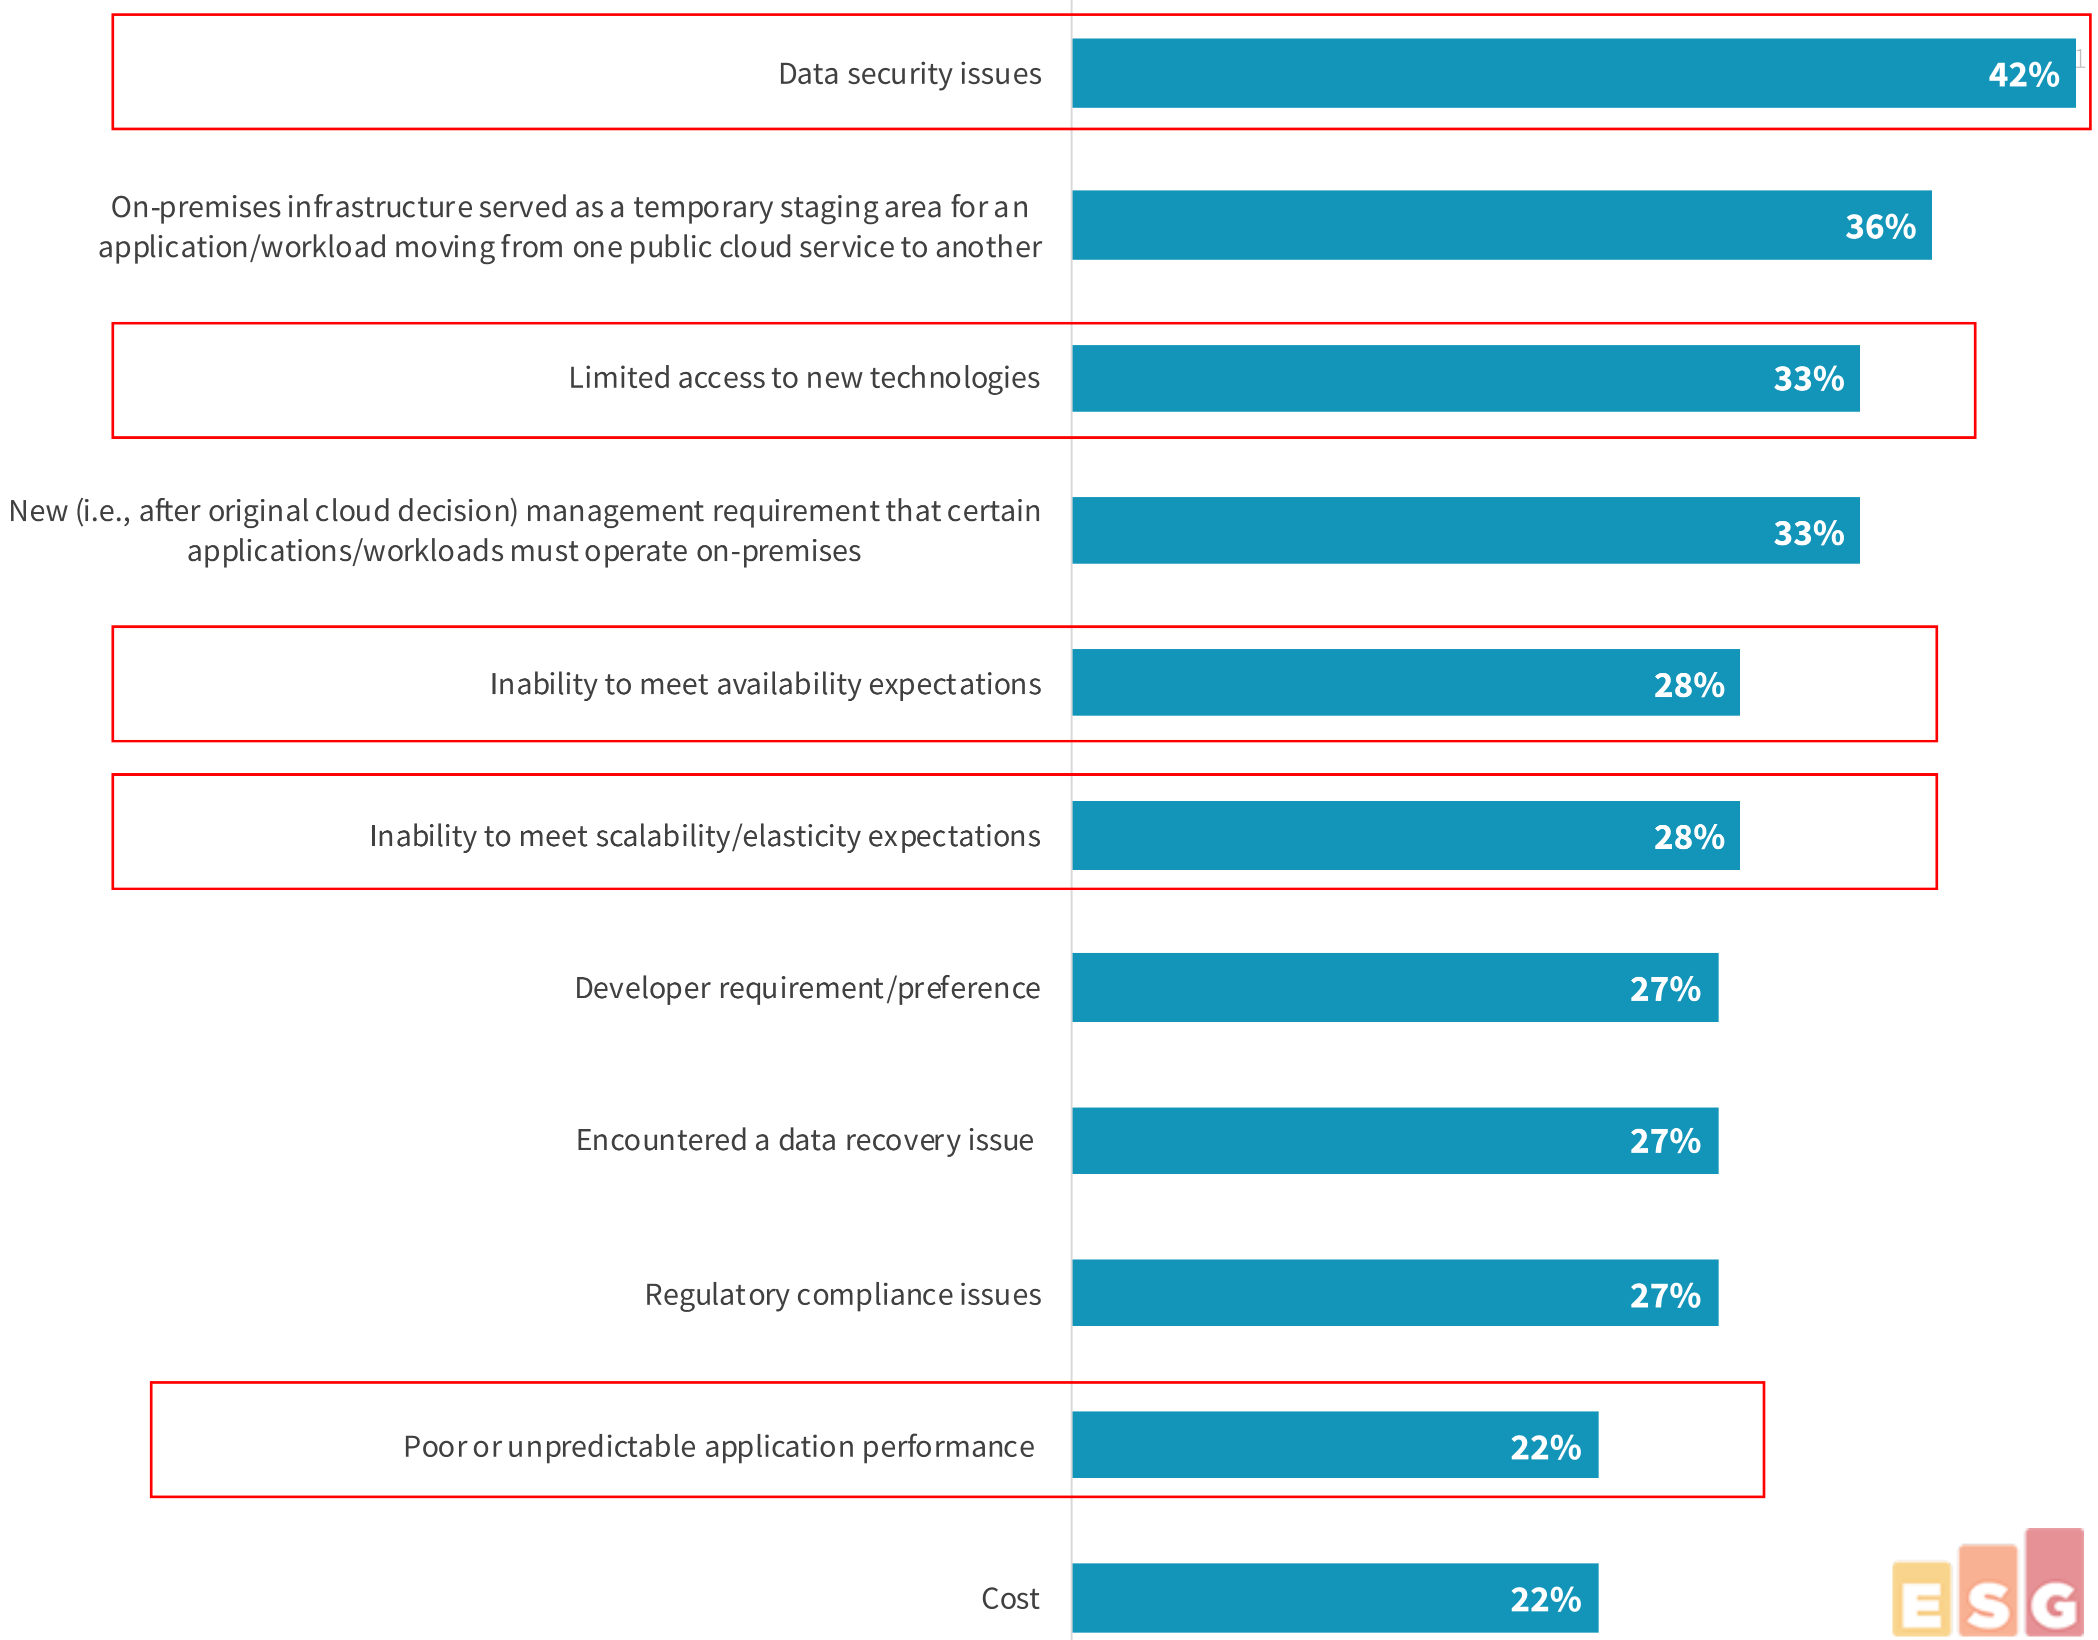

Question text:
What were the reasons behind your organization's decision to move an application(s)/workload(s) back to on-premises infrastructure? (Percent of respondents, N=190, multiple responses accepted)

FC Investment Accelerates, NVMe Adoption Poised to Increase

50% of **FC SAN users** expect to **accelerate their investments over the next 24 months**, with an additional 46% expecting to maintain their current spending levels.

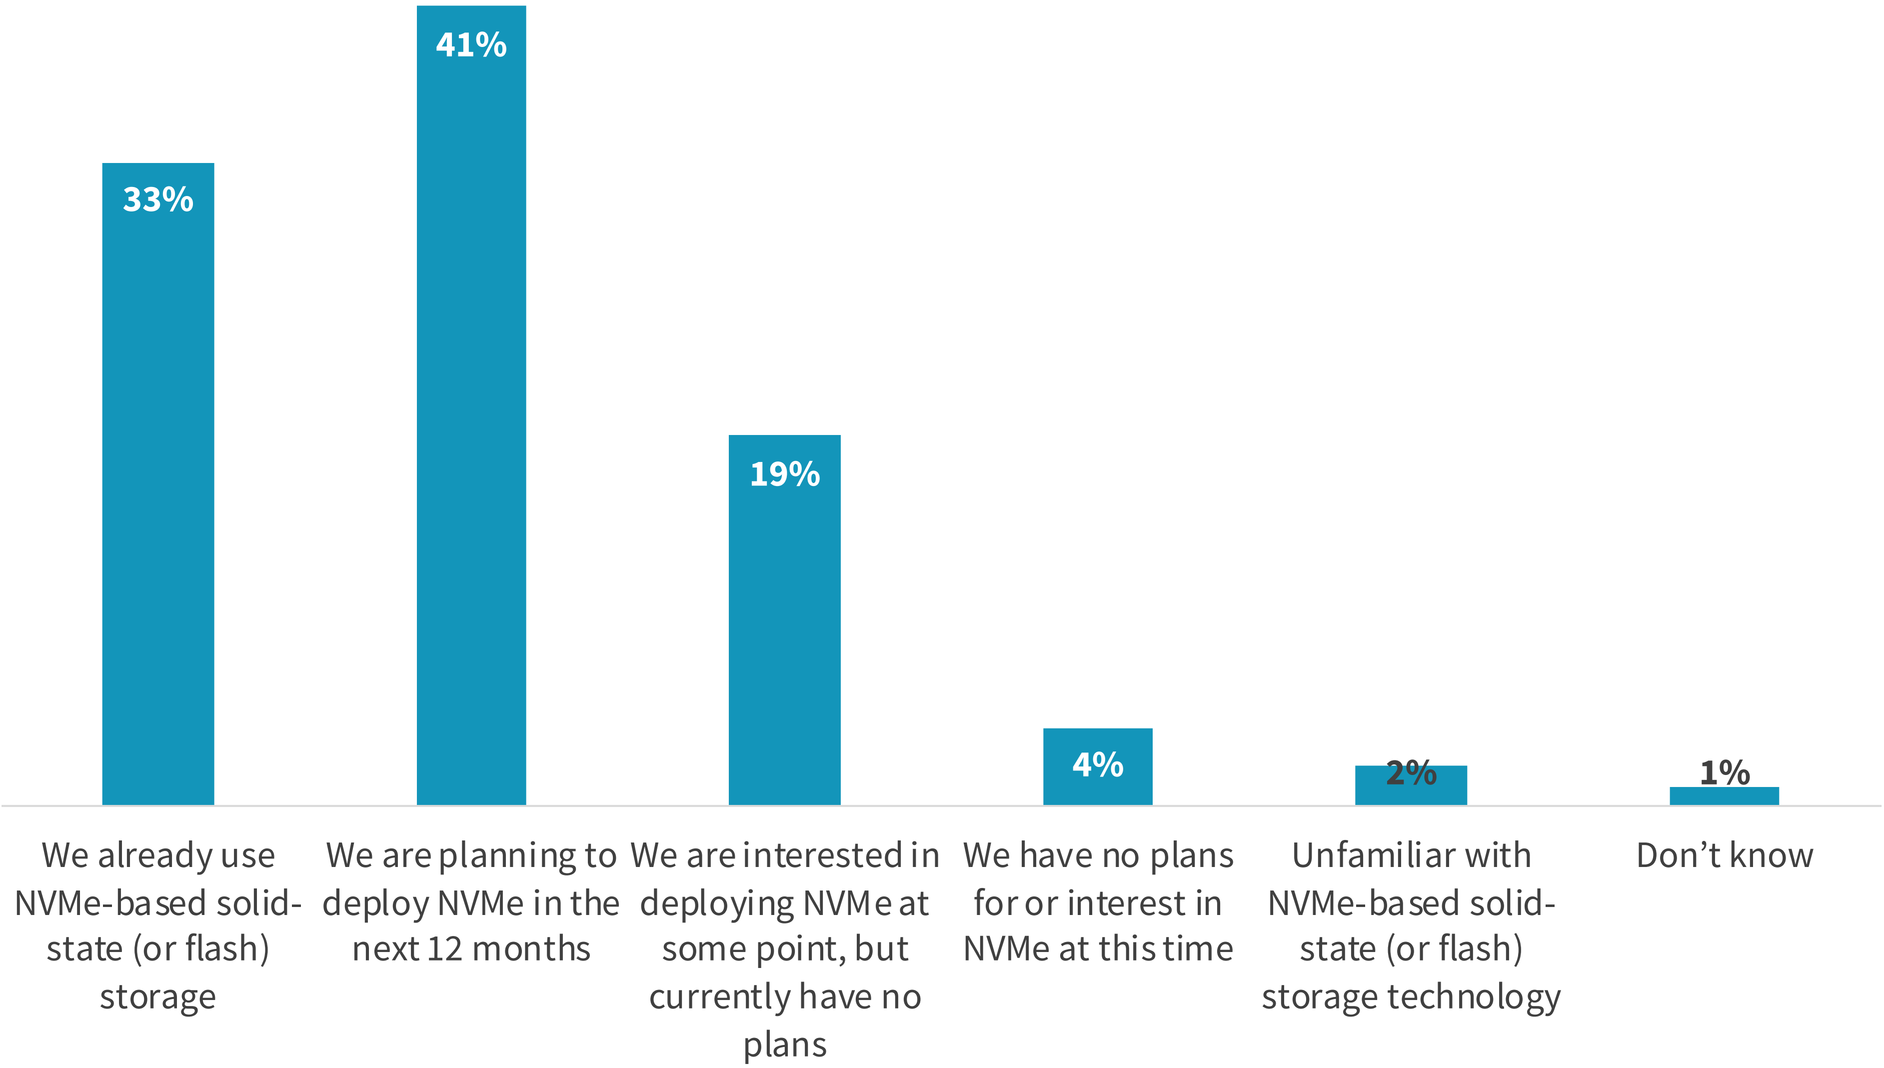

Question text: To the best of your knowledge, has your organization deployed or is it considering deploying on-premises NVMe-based solid-state (or flash) storage technology? (Percent of respondents, N=359)

Realized Benefits of NVMe Flash Storage

Performance for existing and future applications top list of realized benefits.

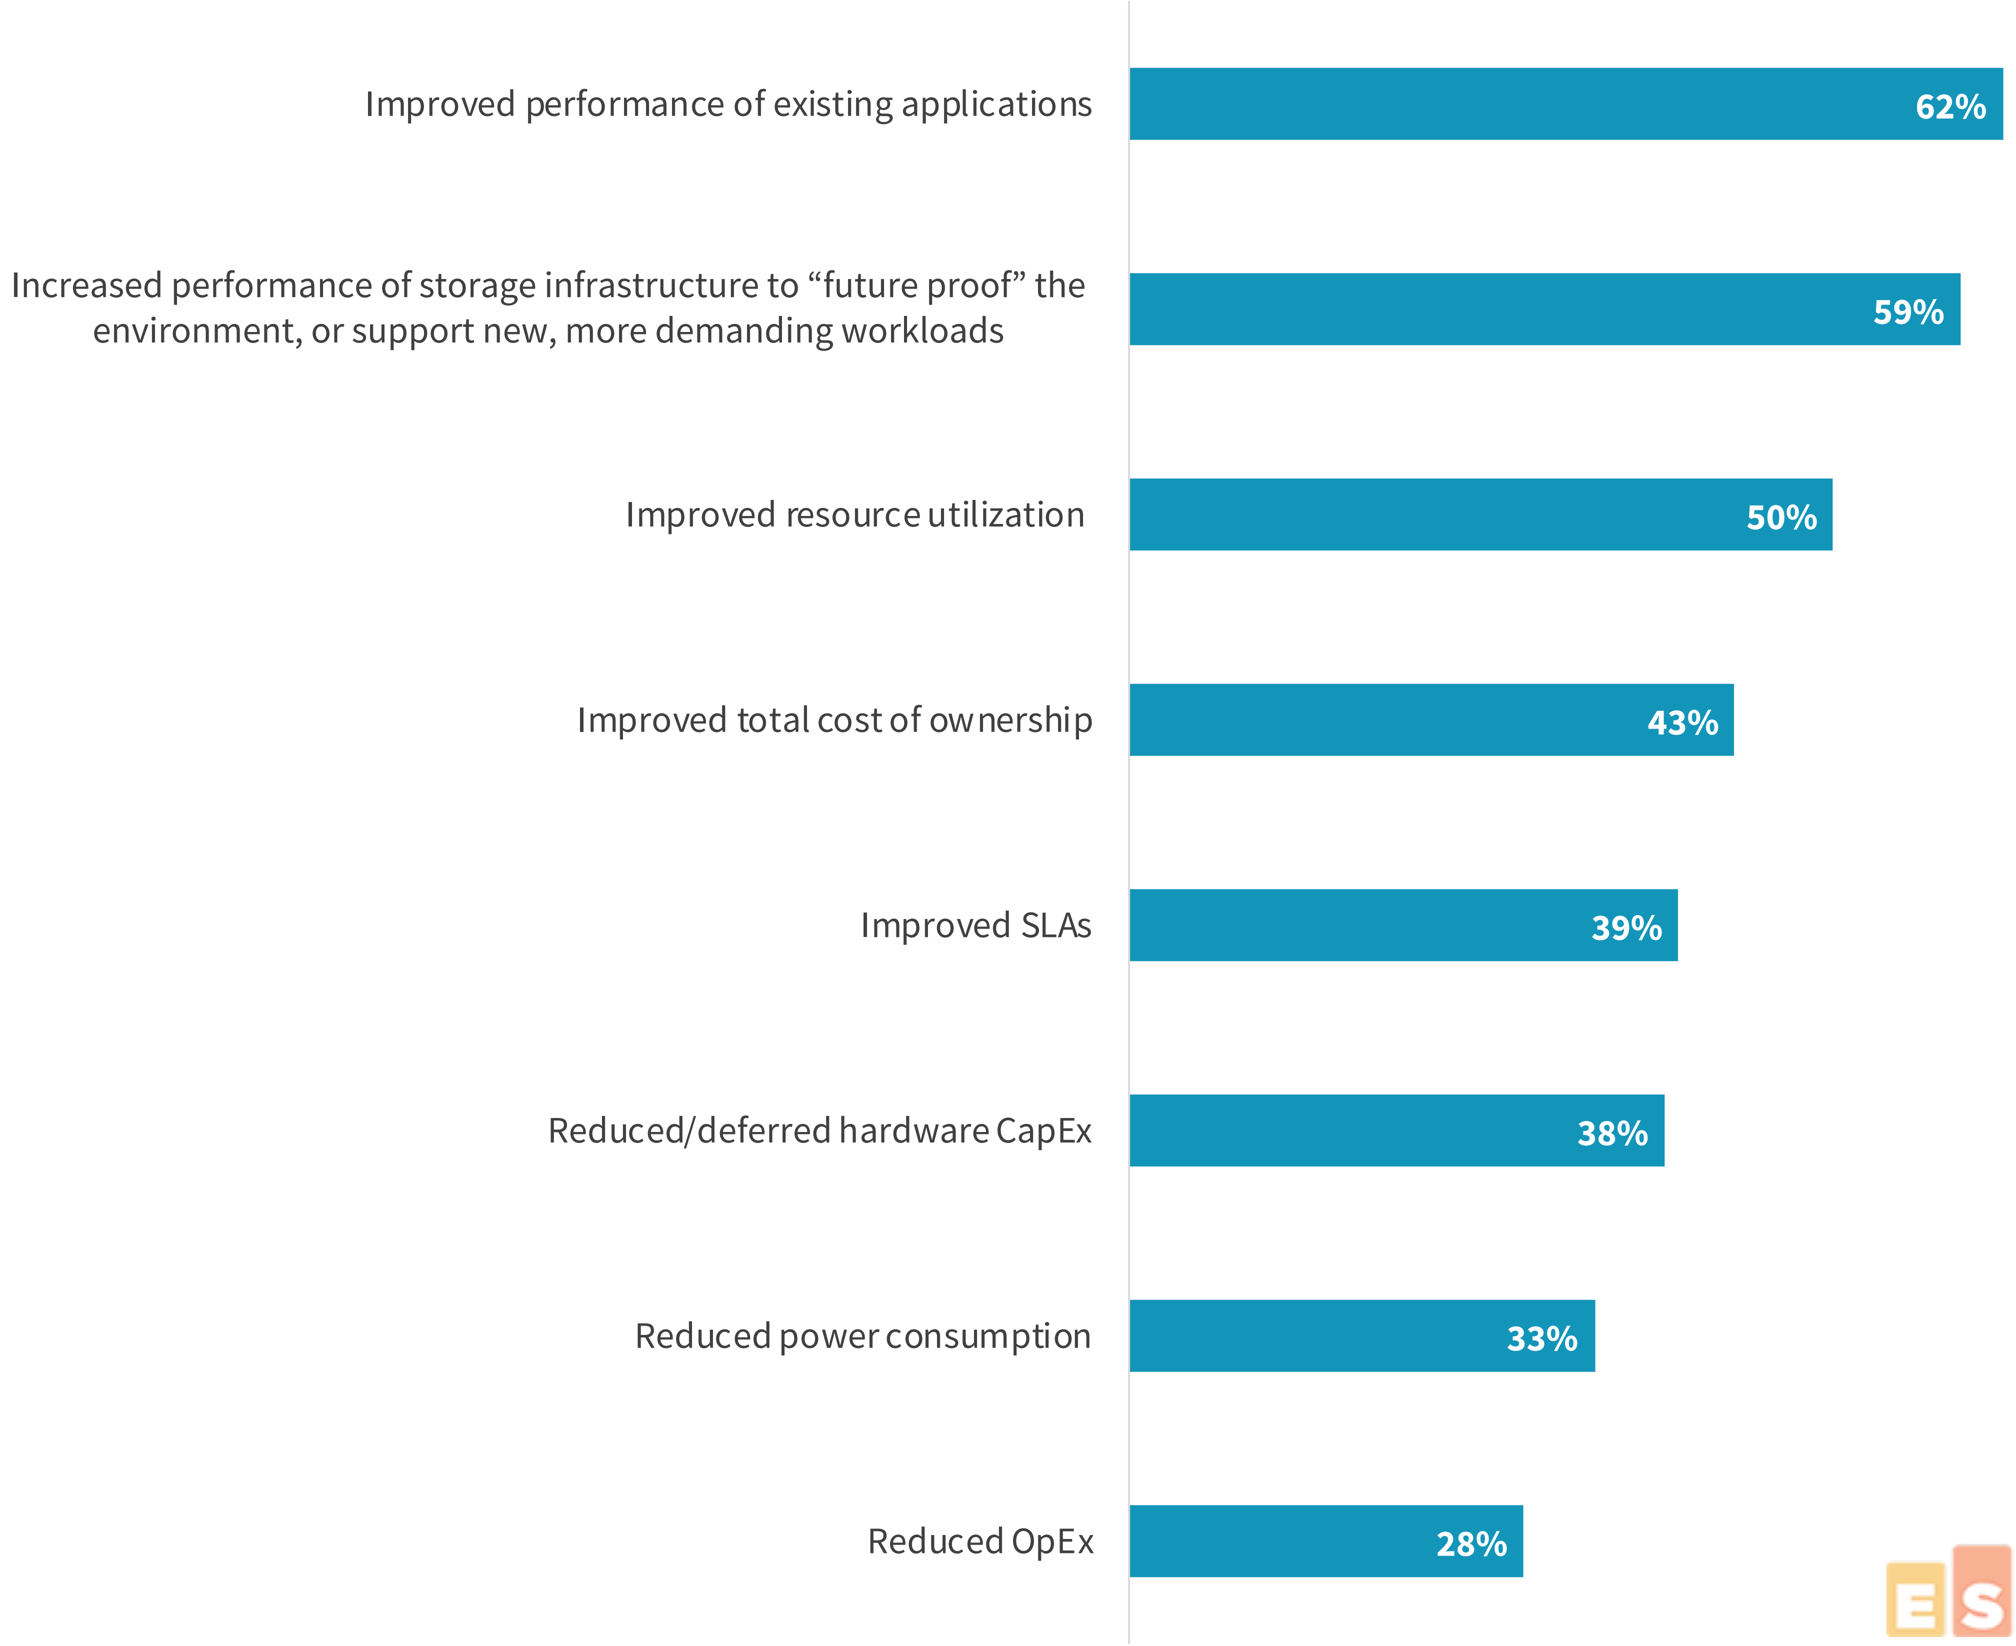

Question text:
Which of the following benefits has your organization realized as the result of deploying on-premises NVMe-based flash storage technology? (Percent of respondents, N=119, multiple responses accepted)

Introduction to Containers and Persistent Storage

Chip Copper

FIBRE CHANNEL INDUSTRY ASSOCIATION

Applications Lifecycles and Data

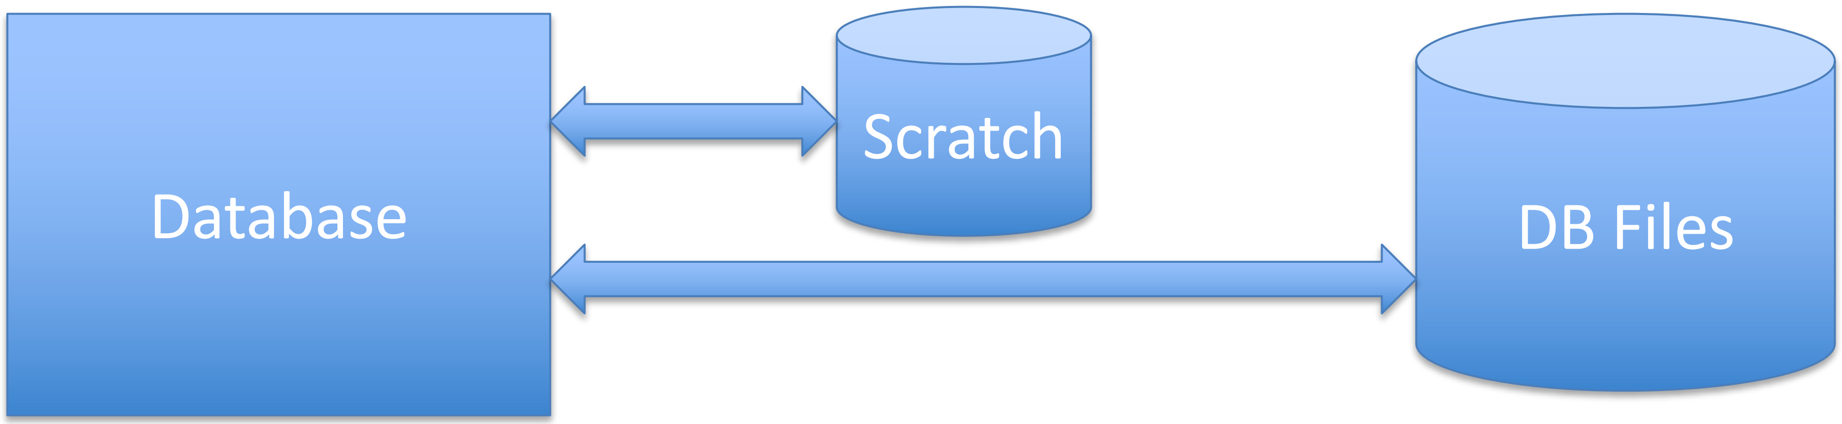

Bare Metal Deployments: Application, Temporary Files, Persistent Storage

Applications Lifecycles and Data

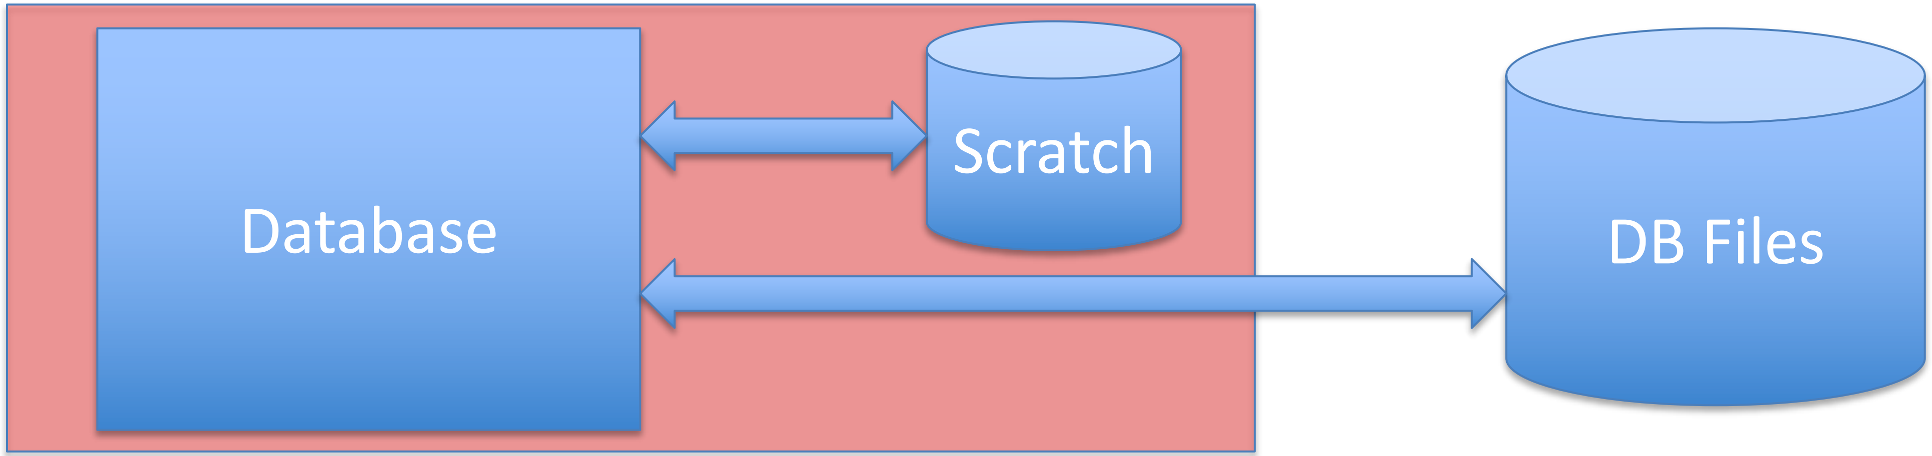

The application and scratch files may go away, but long-term state information stays

Applications Lifecycles and Data

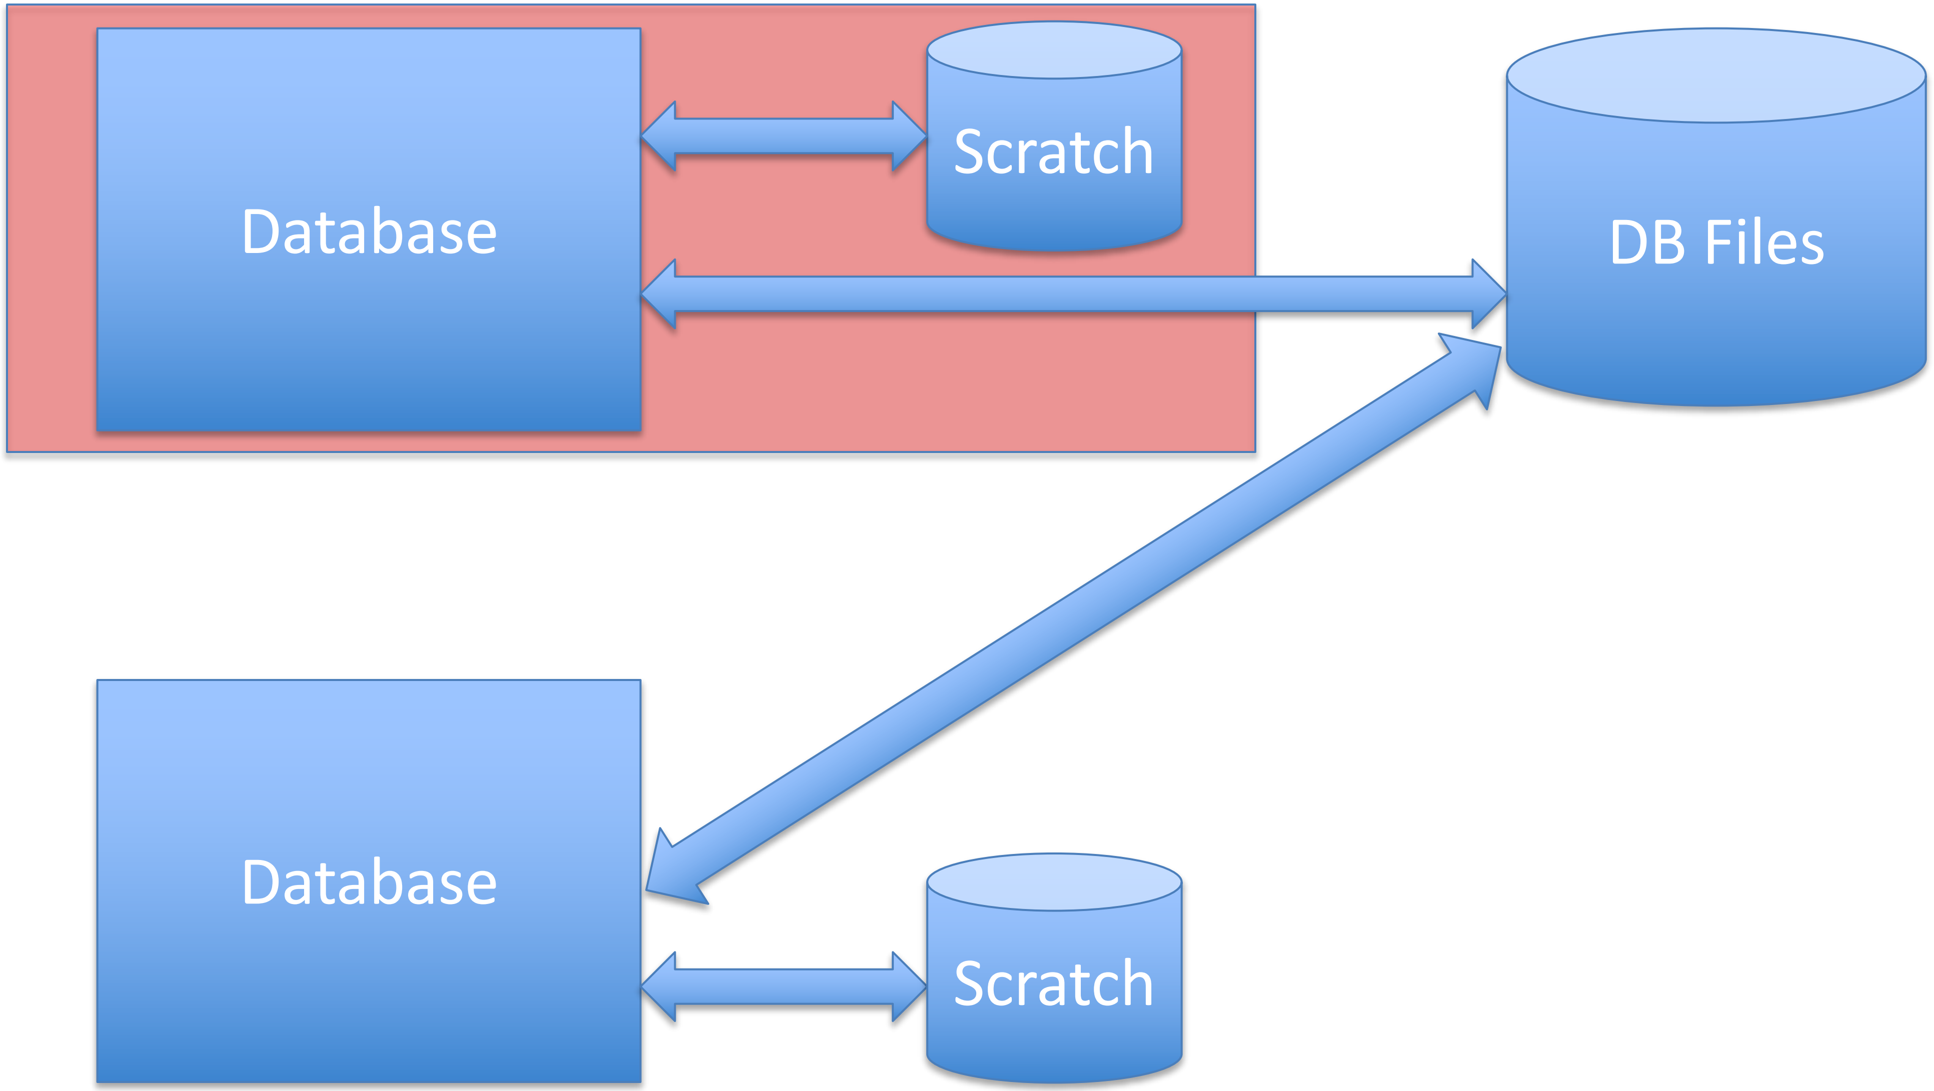

New application instances will attach to the saved state

Containerized Applications

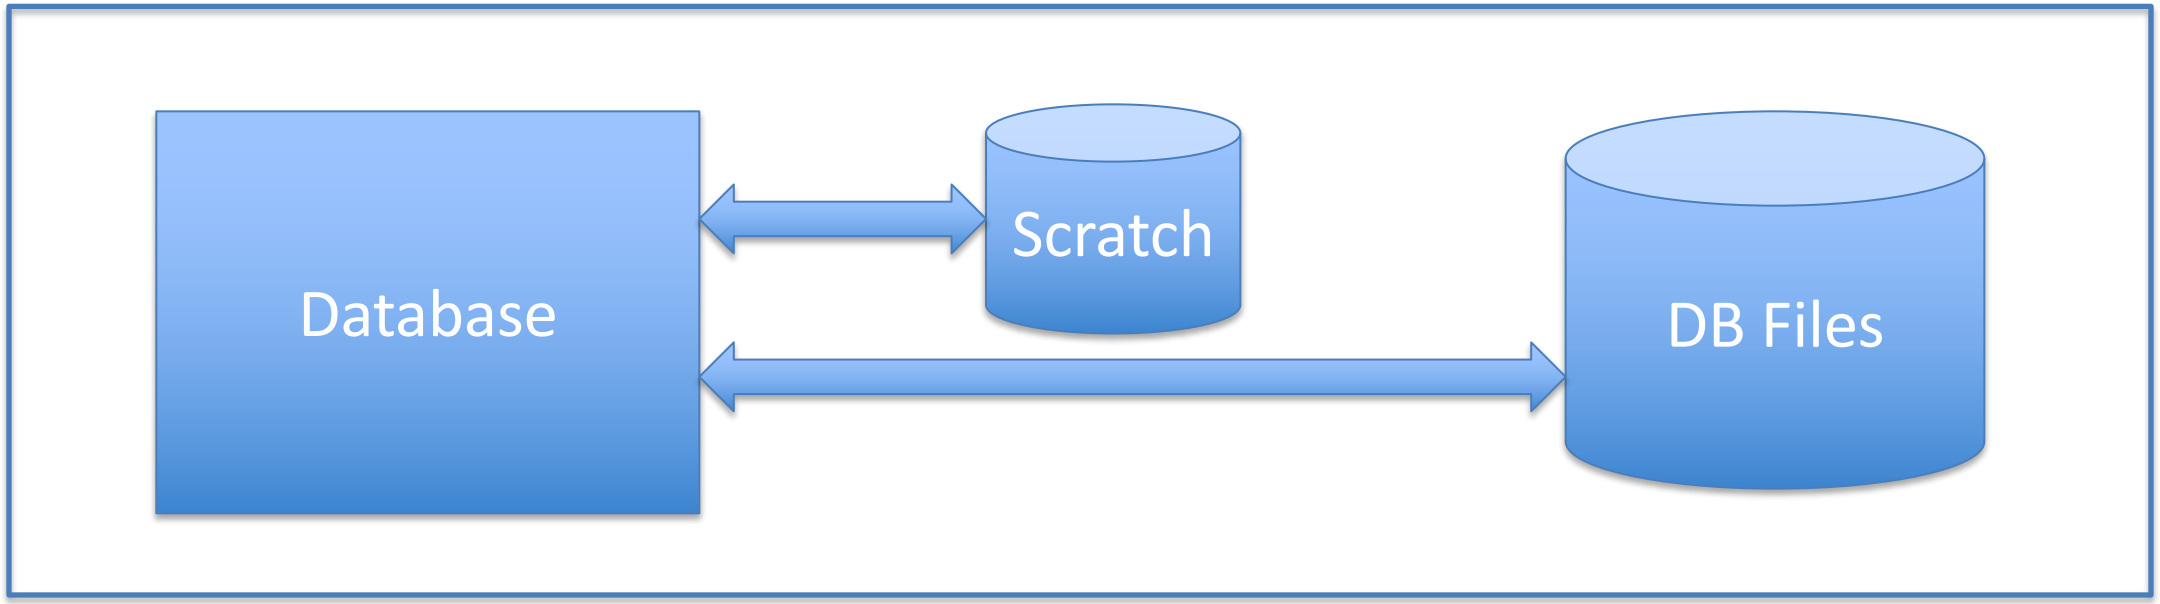

In a containerized environment, volumes can reside inside of container instances

Containerized Applications

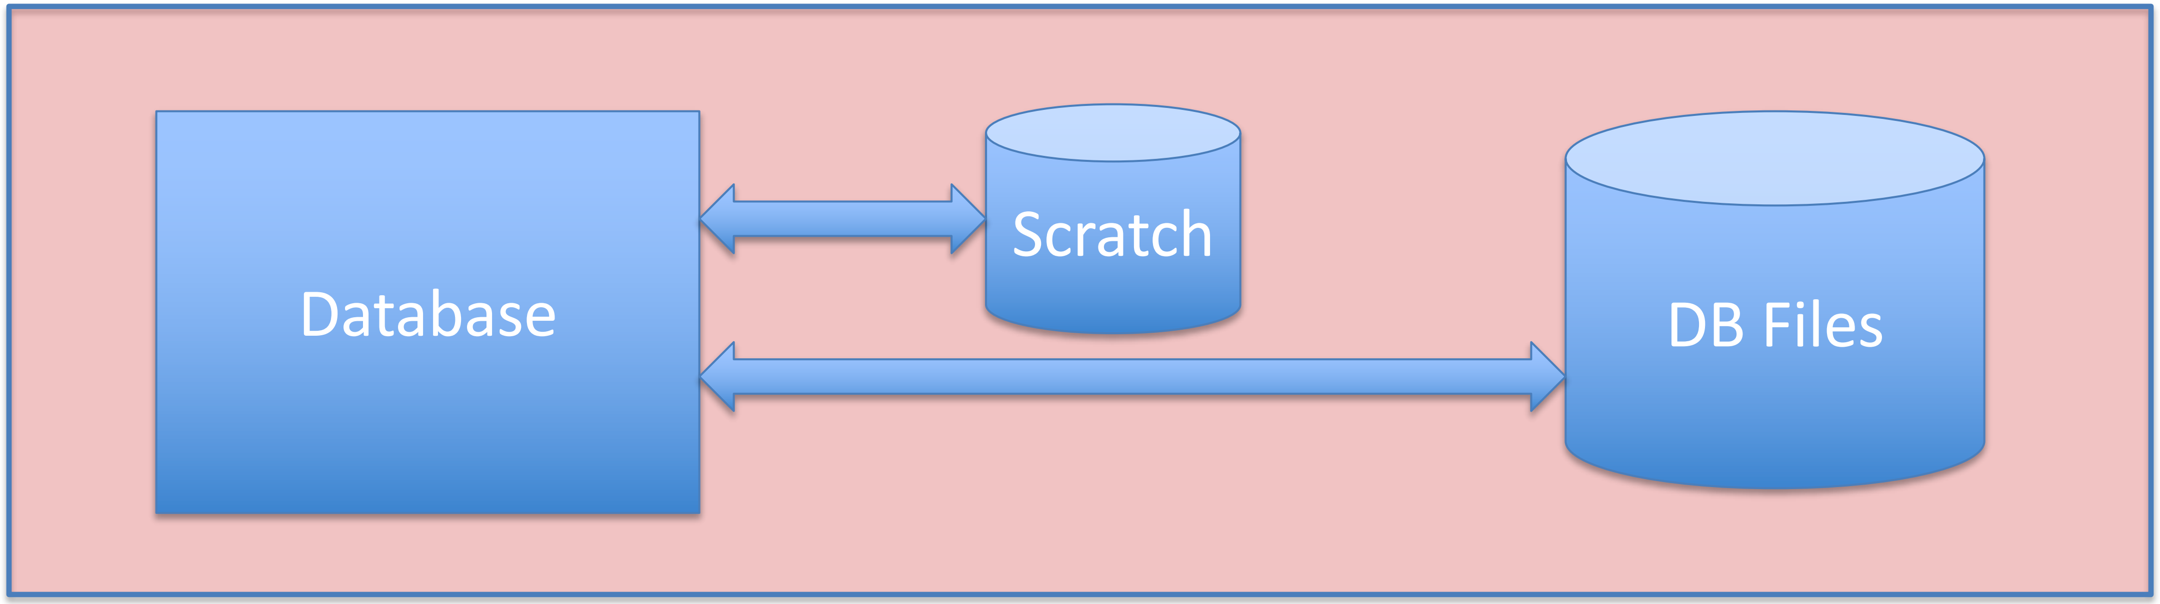

Termination of the container deletes all components inside

Containers

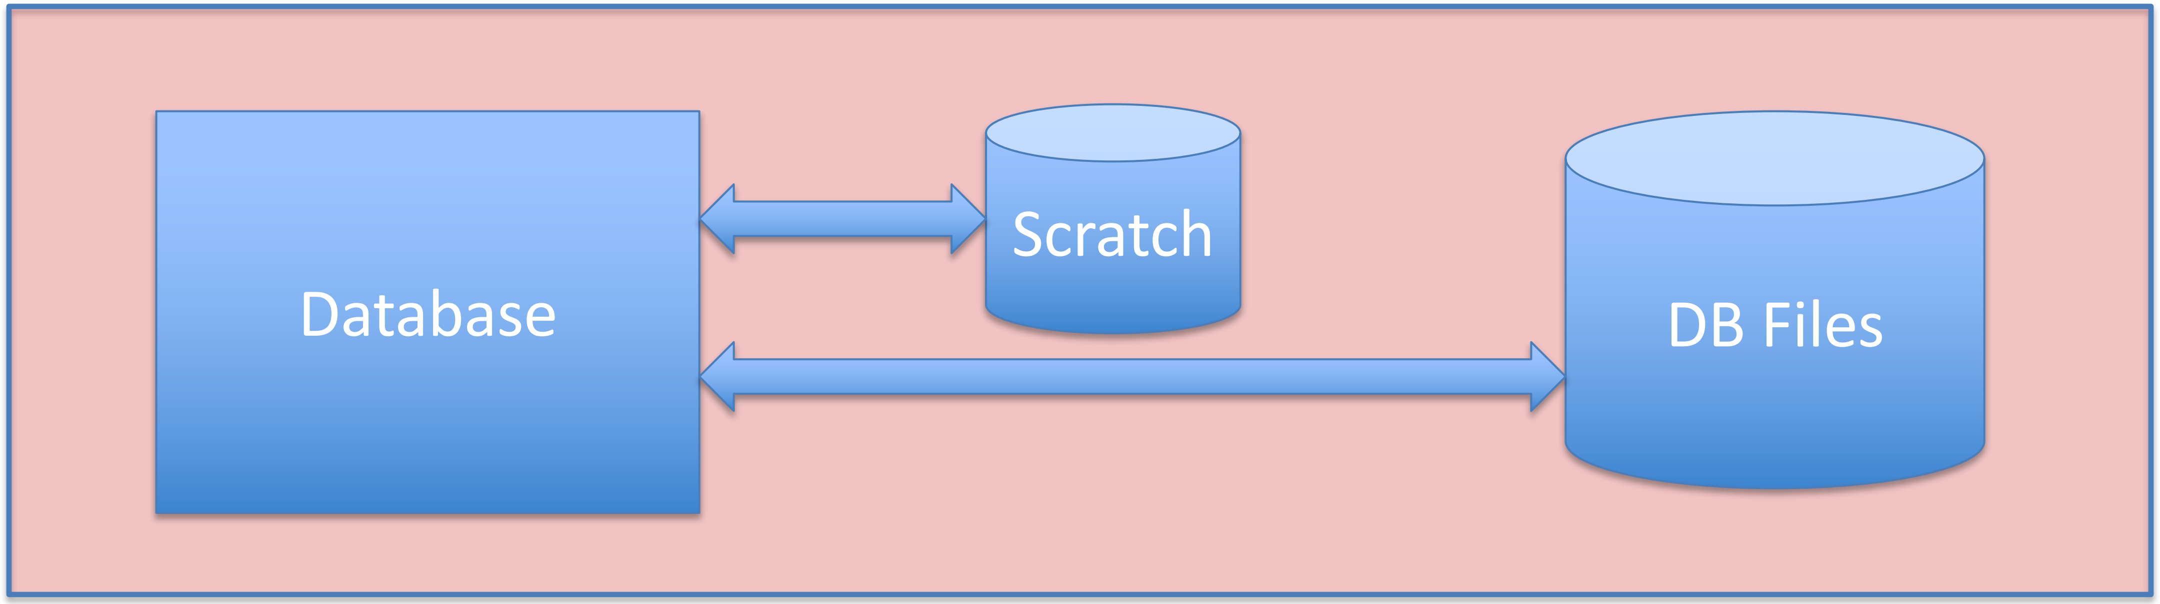

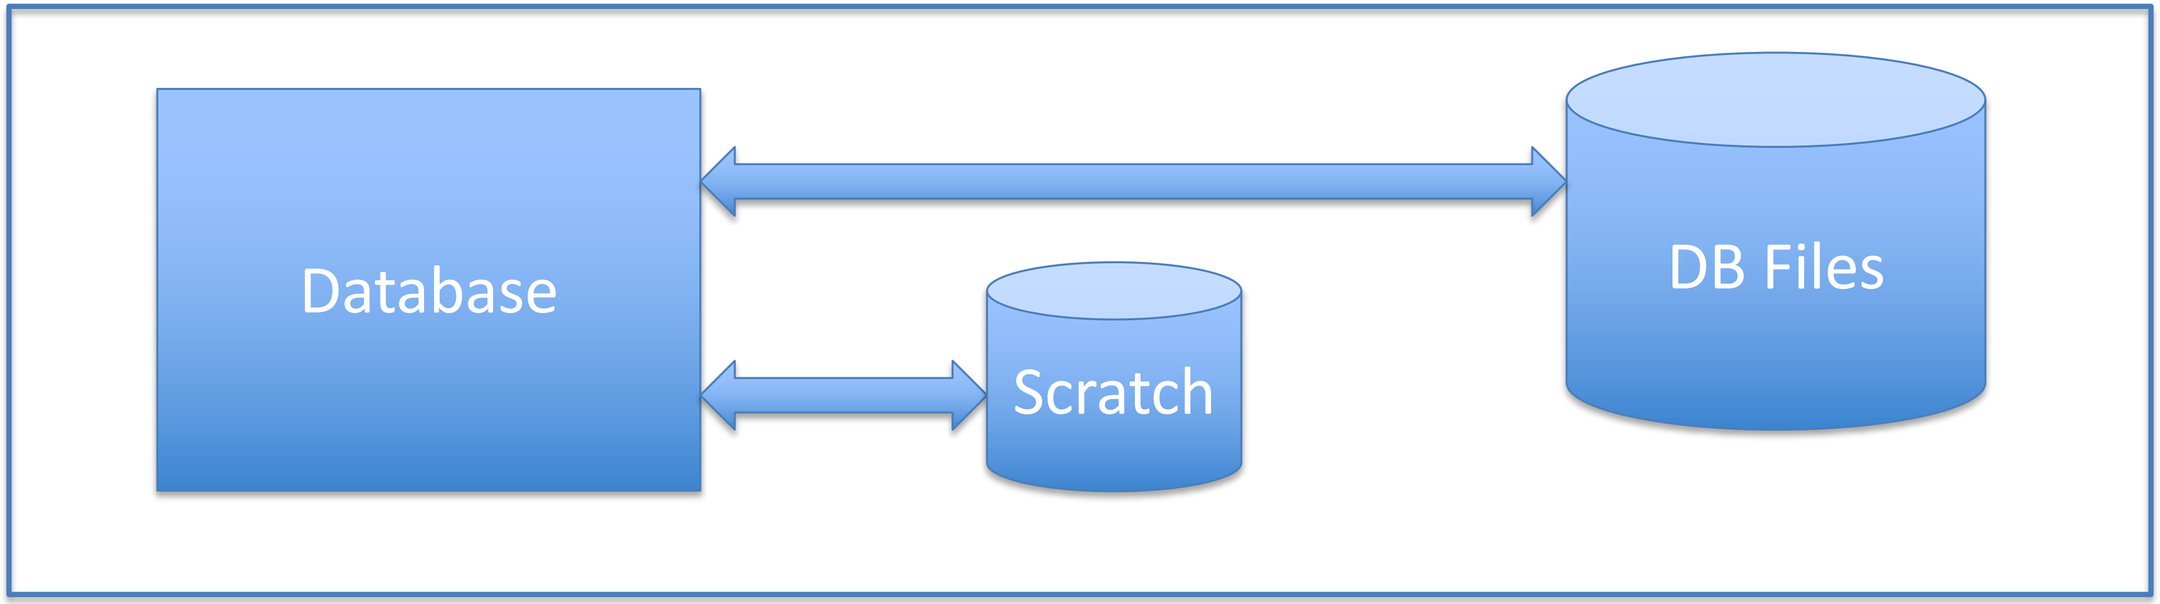

A new container contains new volume instances

Containers

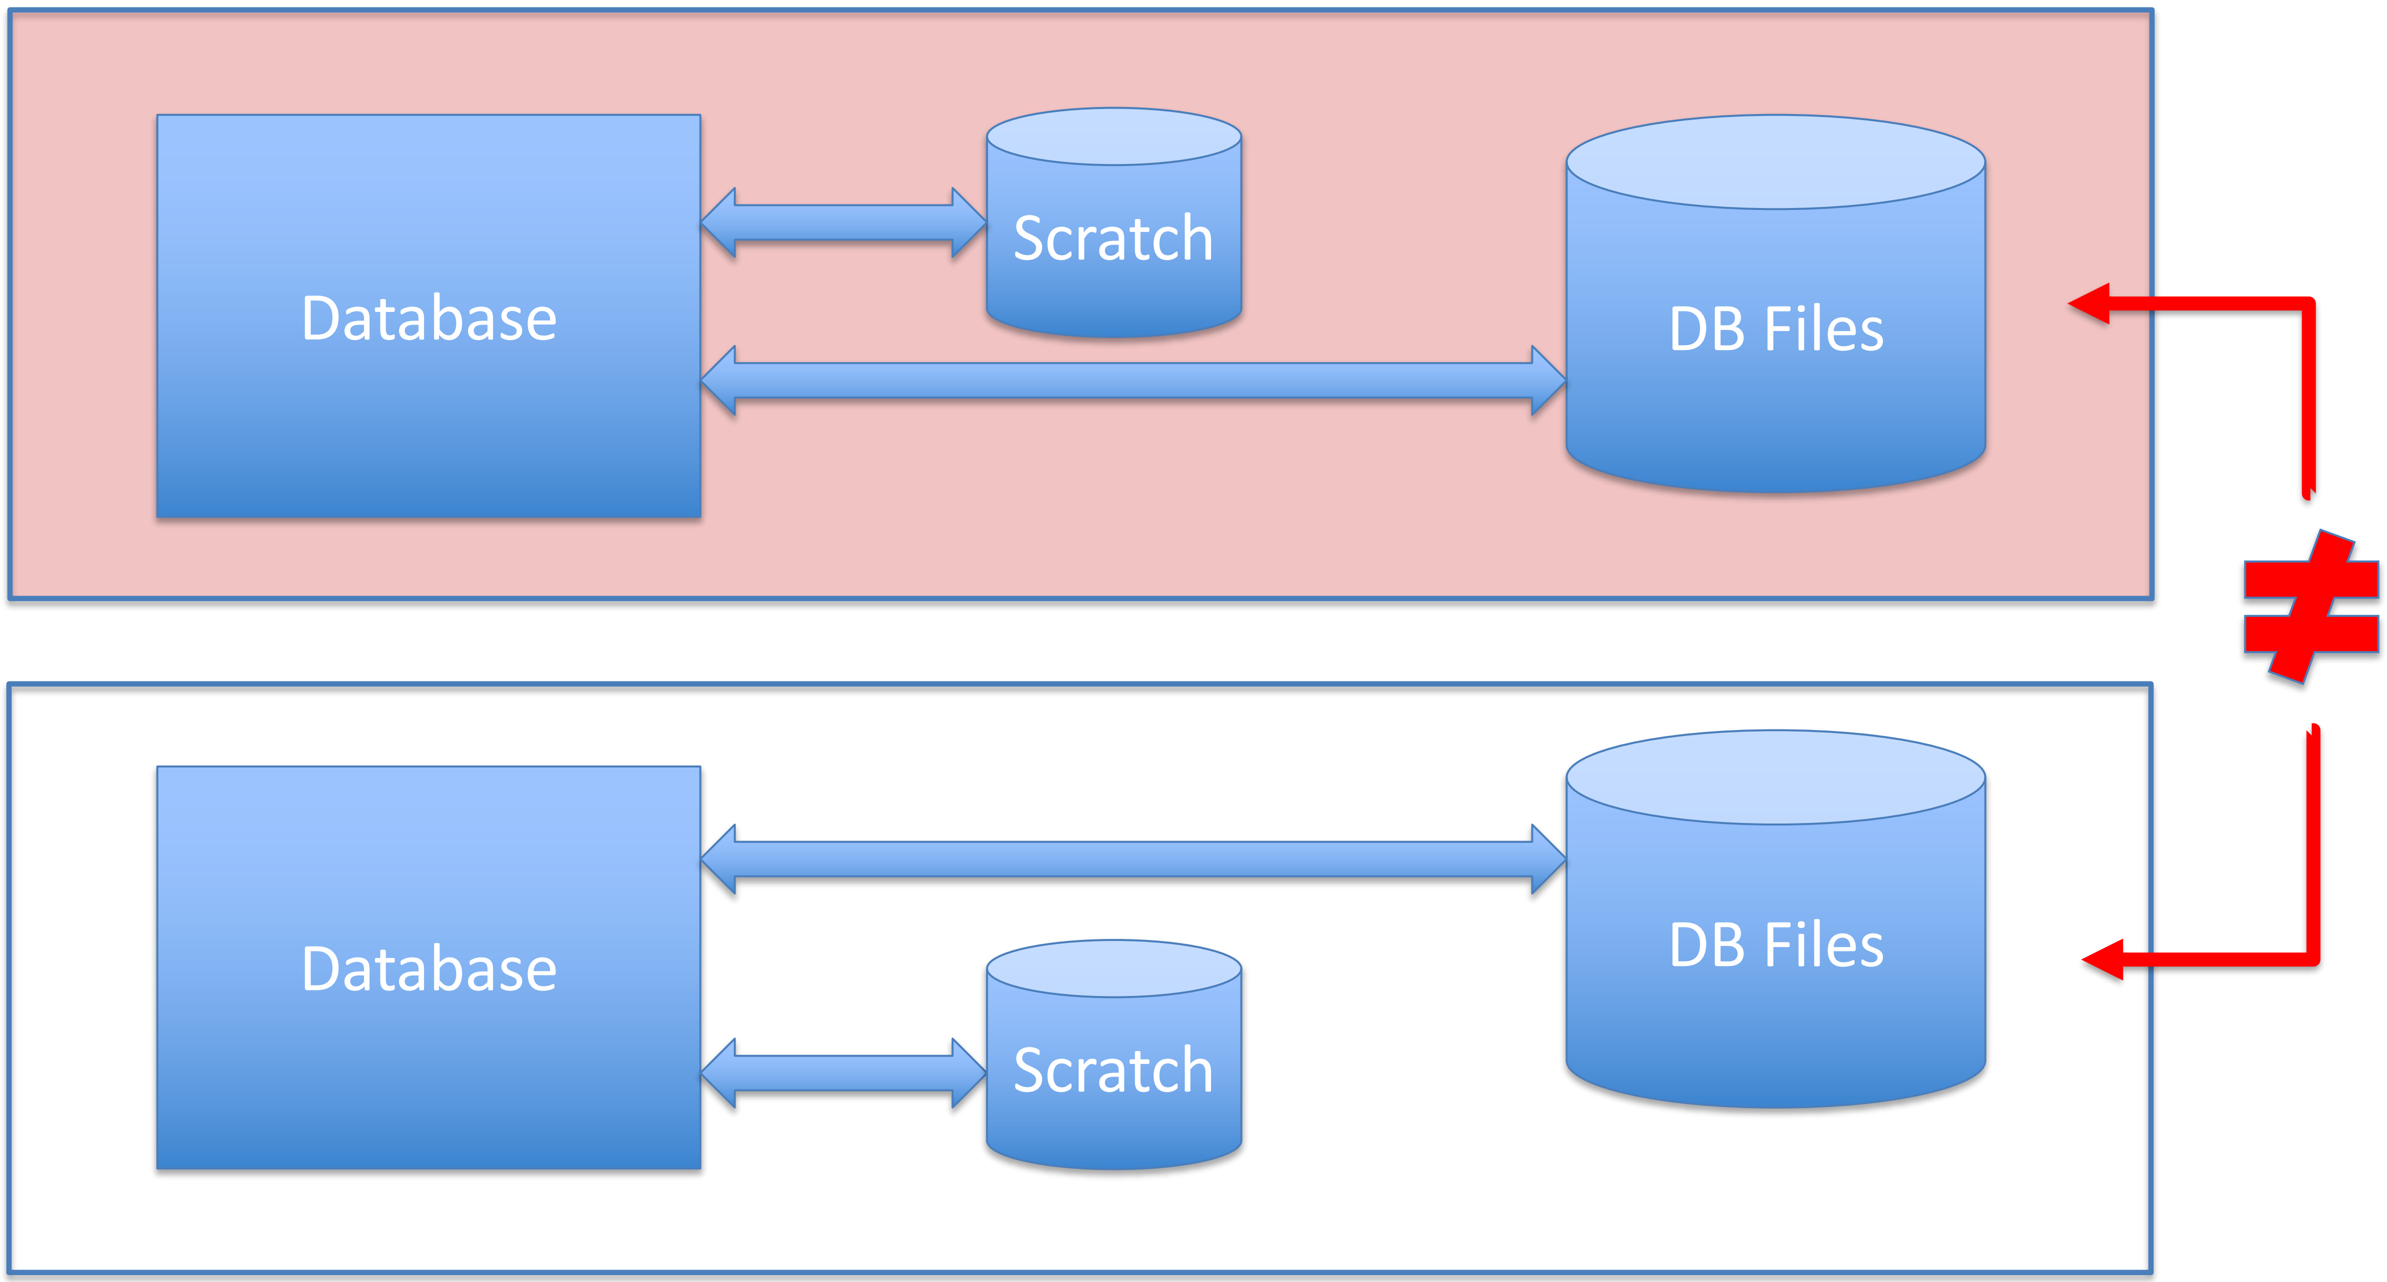

Containers

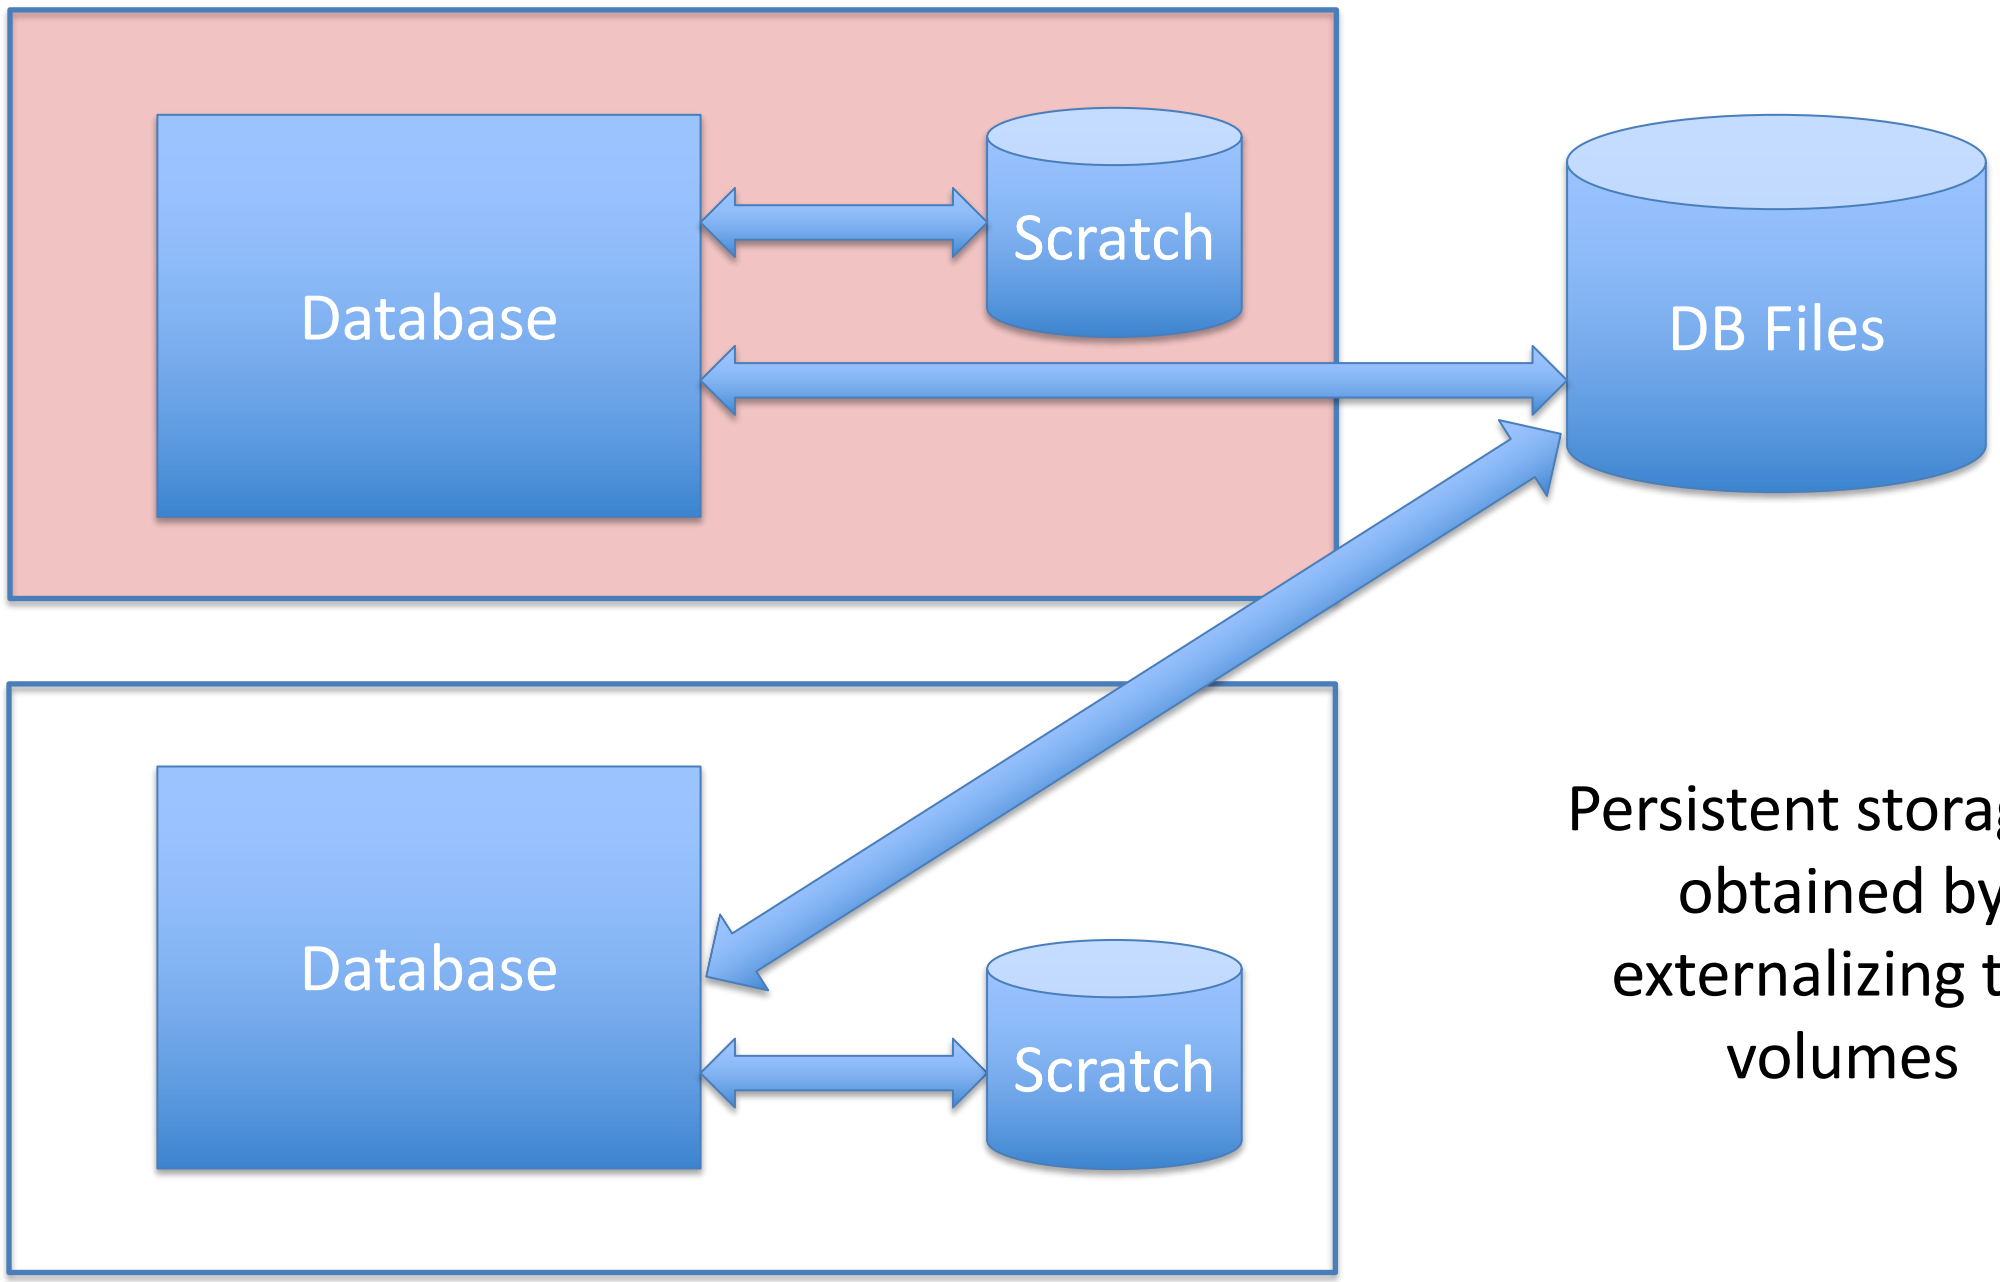

Persistent storage is obtained by externalizing the volumes

Logical Separation

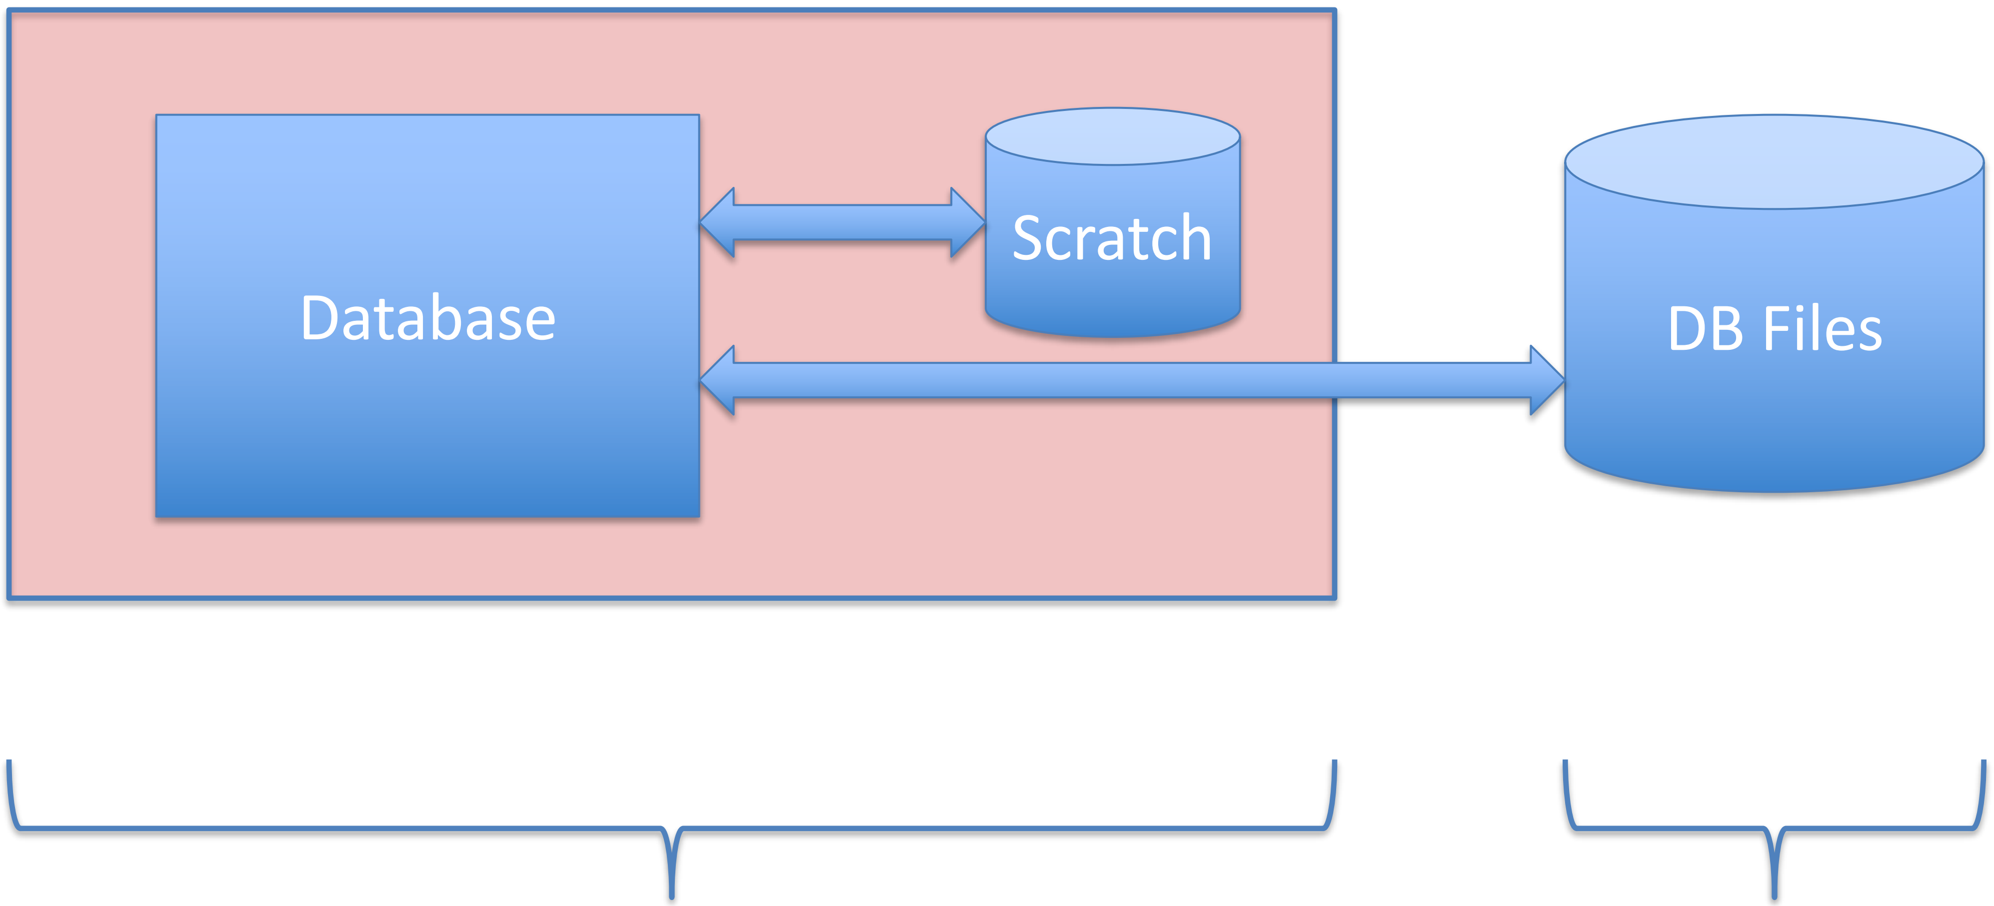

Containers

Volumes

Container and volume lifecycles are managed independently

Logical Separation

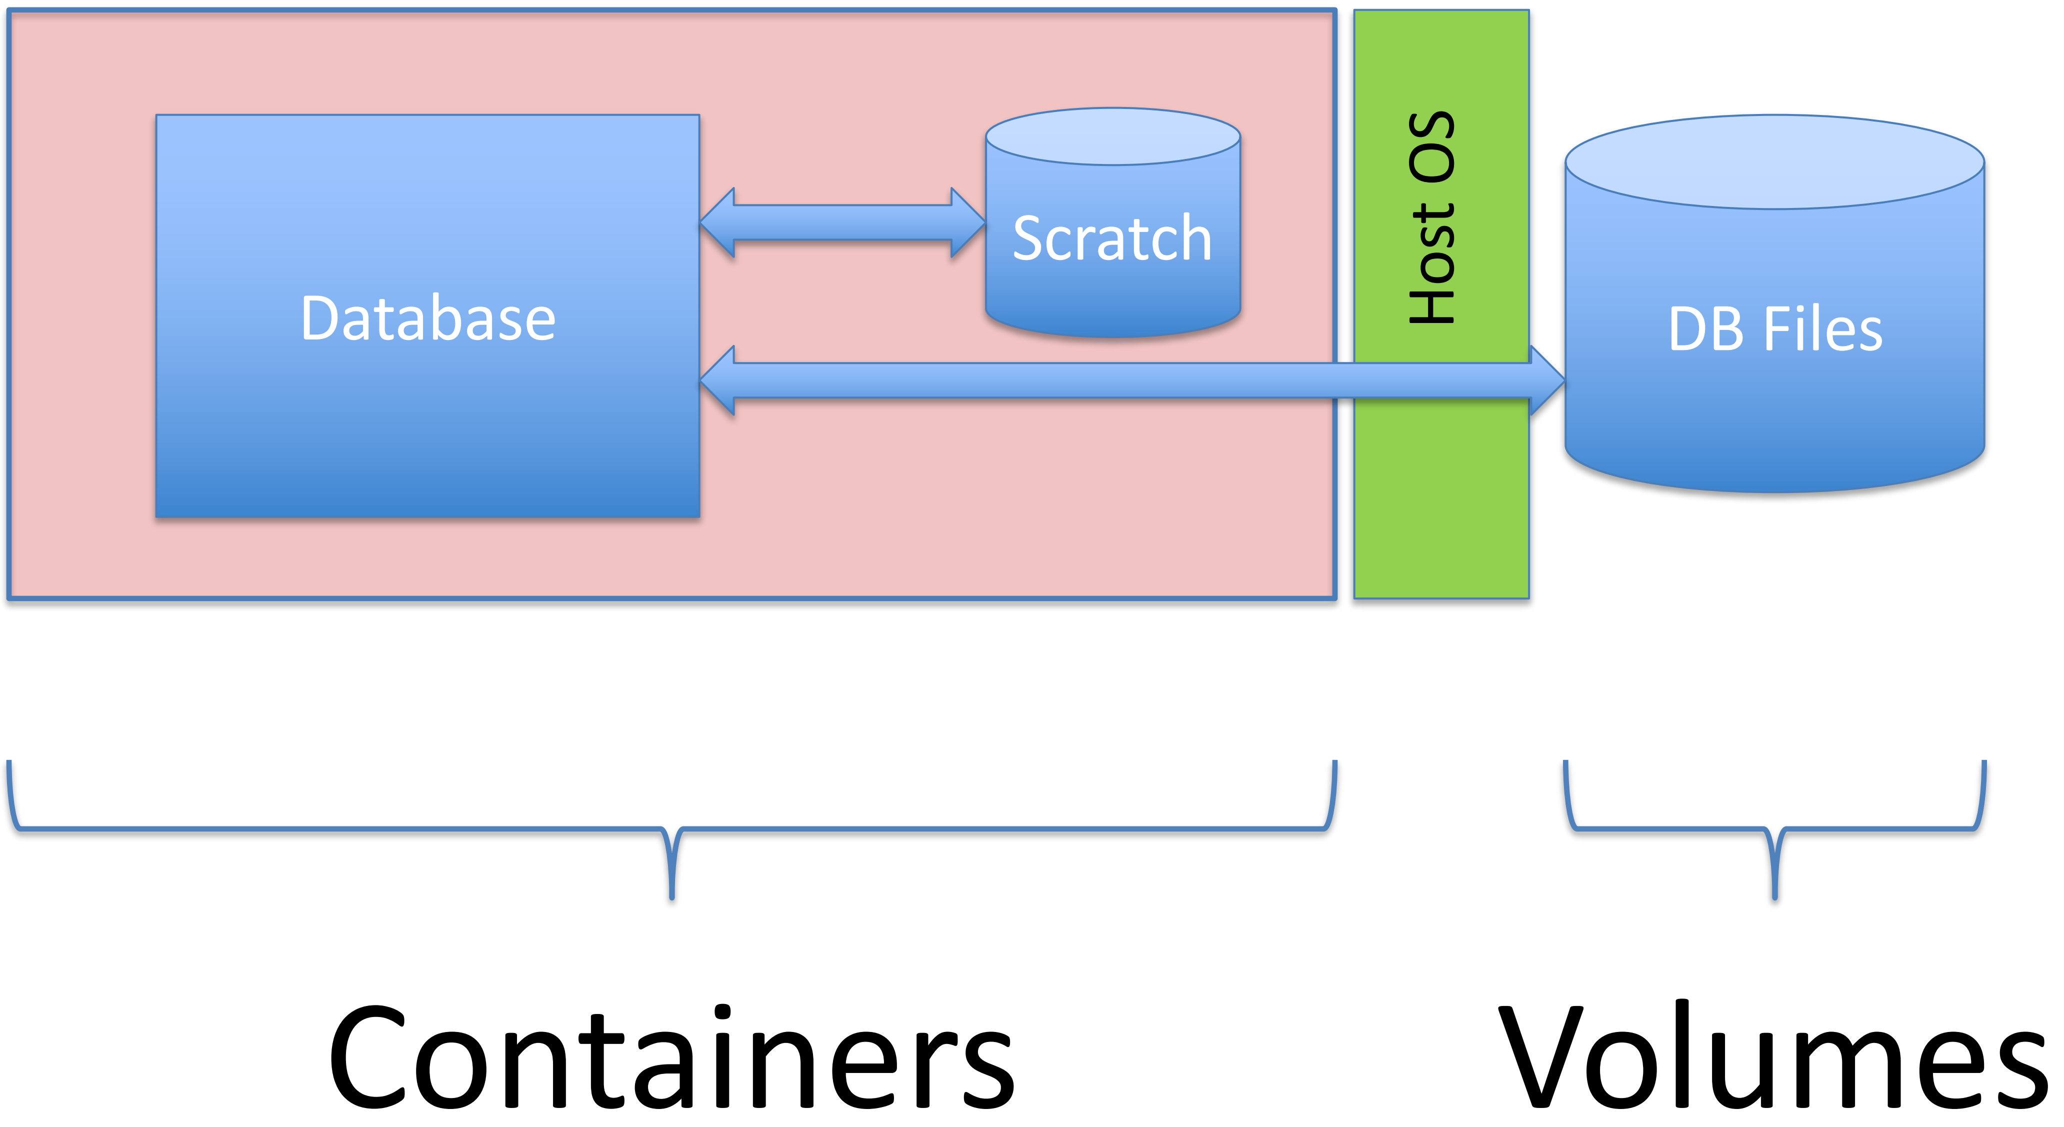

The Host System provides the presentation of volumes to containers as filesystems

Orchestrators Make The Connection

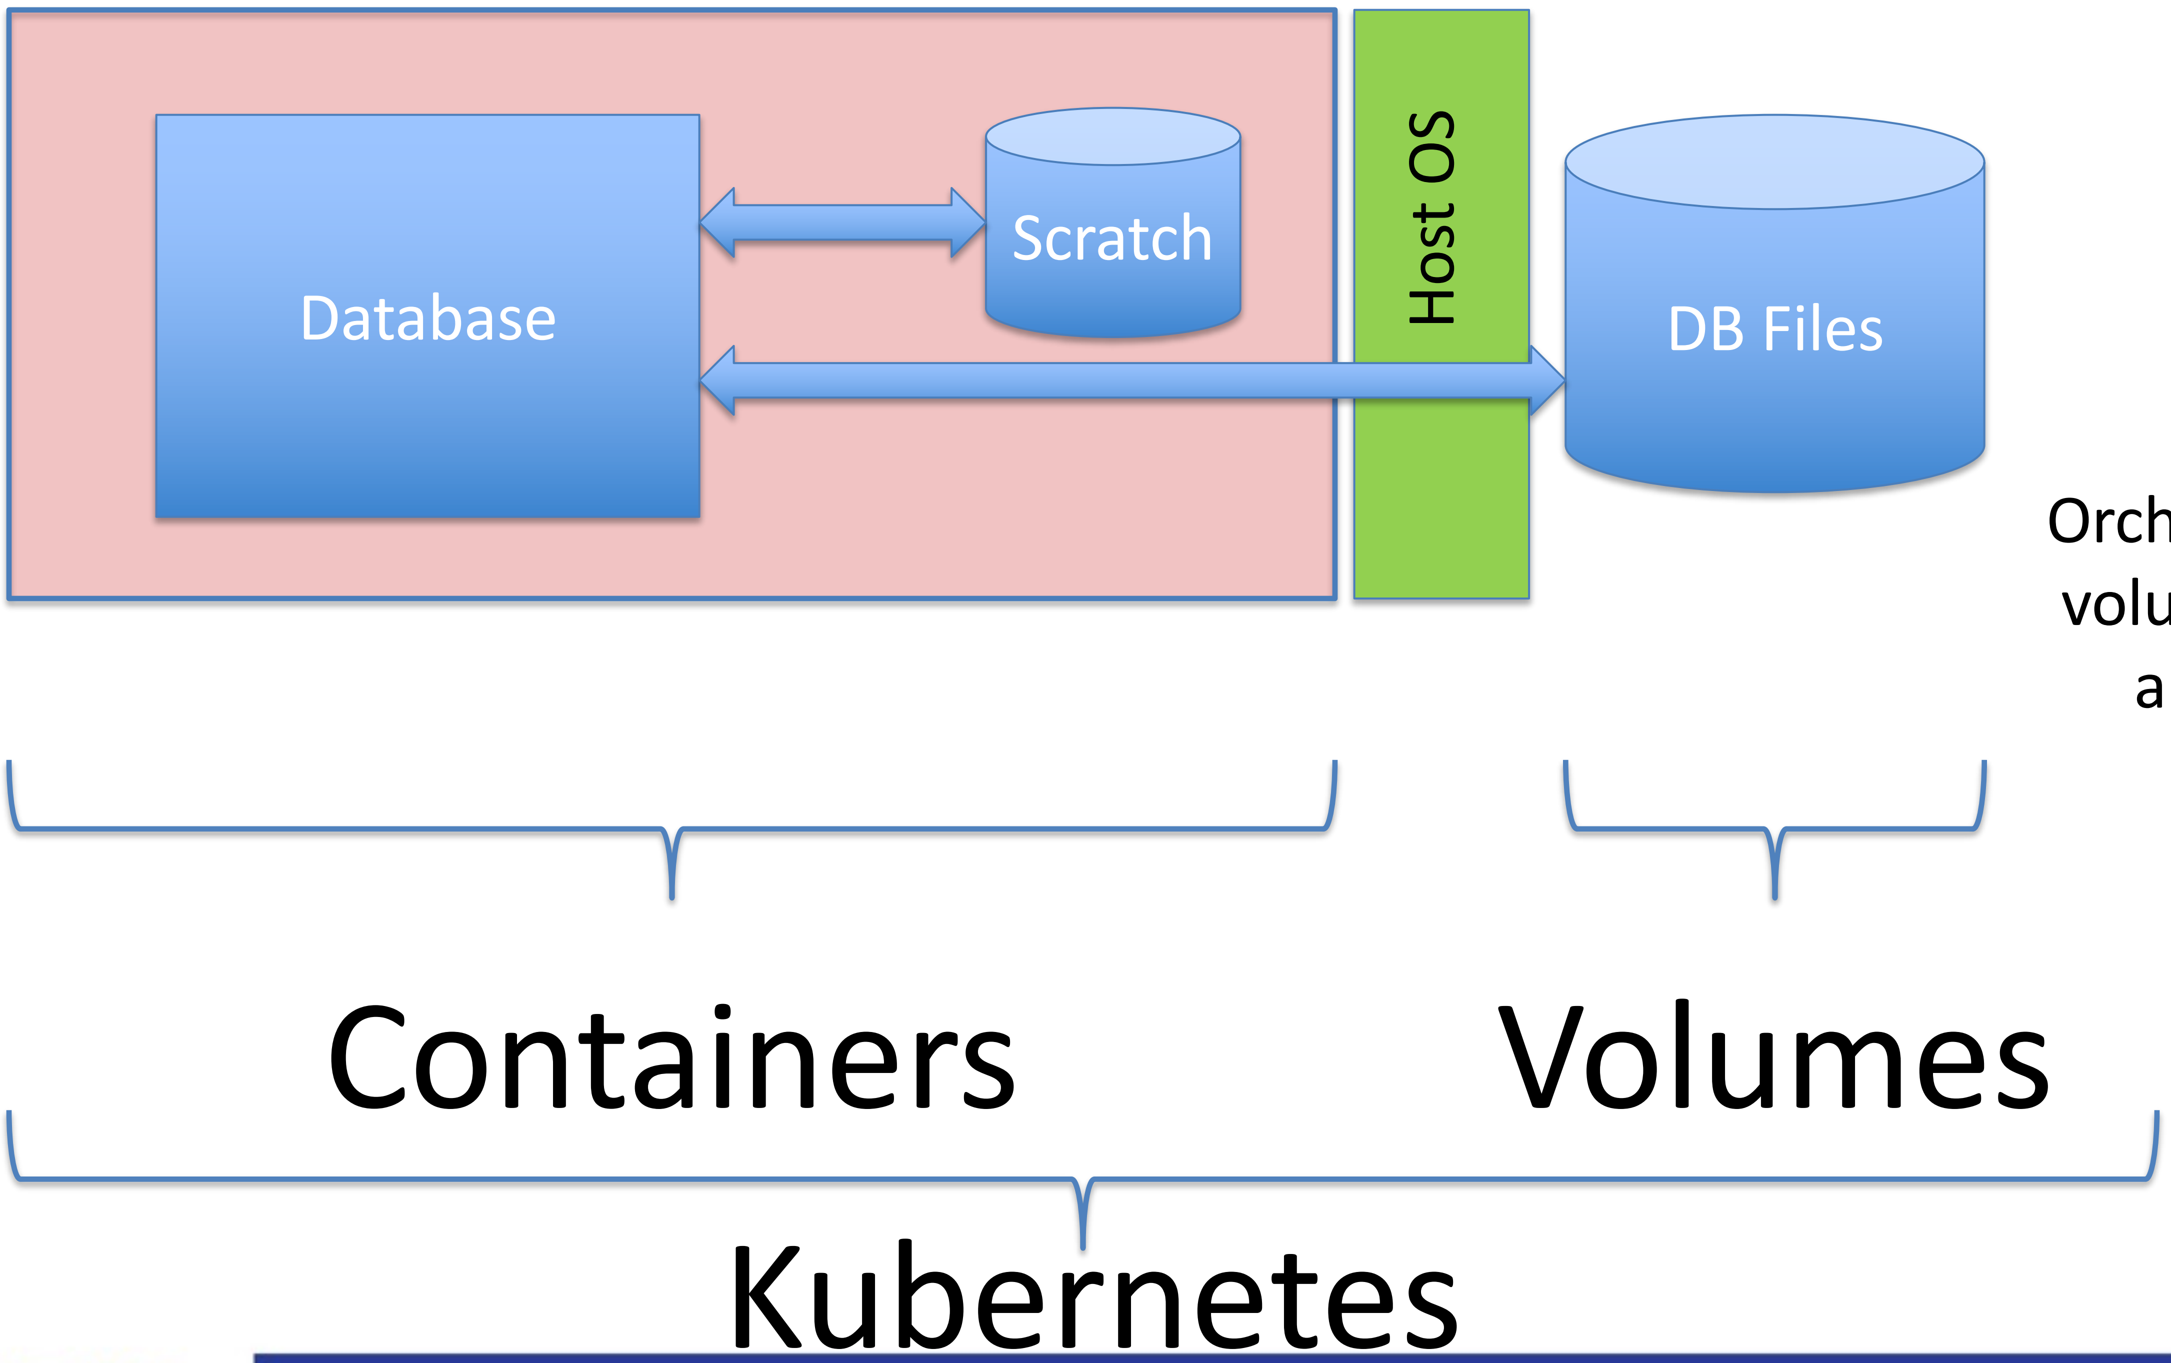

Orchestrators provide volume coordination and association

Kubernetes Volume Lifetimes

- Volume lifetimes may be linked to a container lifetime or not
- Ephemeral volumes
  - Follow the lifecycle of the associated container
  - May be prepopulated with files for staging or predetermined content
  - Are destroyed with the exiting container instance
- Persistent volumes
  - Created and managed independently from containers
  - May be pre-provisioned or created on demand
  - Are made available via the host system's volume pool

Kubernetes Volume Characteristics

- Volumes are presented to containers as filesystems
- Filesystem drivers are supplied by the host operating system
- Container runtimes make the logical volumes available
- Kubernetes orchestrates the management of logical volumes
  - Kubernetes does not supply storage or filesystems
  - Kubernetes depends on the container runtime and OS for them
- Storage system characteristics are reflected in containers
- The user must specify the appropriate type of volume for the app

Volume Types In Kubernetes

- awsElasticBlockStore
- azureDisk
- azureFile
- cephfs
- cinder
- configMap
- **csi**
- downwardAPI
- emptyDir
- fc (fibre channel)
- gcePersistentDisk
- glusterfs
- hostPath
- iscsi
- local
- nfs
- persistentVolumeClaim
- portworxVolume
- projected
- rbd
- secret
- vsphereVolume

Takeaways

- The choice of persistent storage for containers matters
- All of the characteristics of storage alternatives are visible
  - Reliability
  - Reachability
  - Scalability
  - Performance
- Container clients allow users to choose the right storage
- Kubernetes orchestrates but does not implement volumes
- Fibre Channel continues to be differentiated in a containerized application environment

The Role of Fibre Channel in Kubernetes

Nishant Lodha

Persistent Storage Requirements for K8s

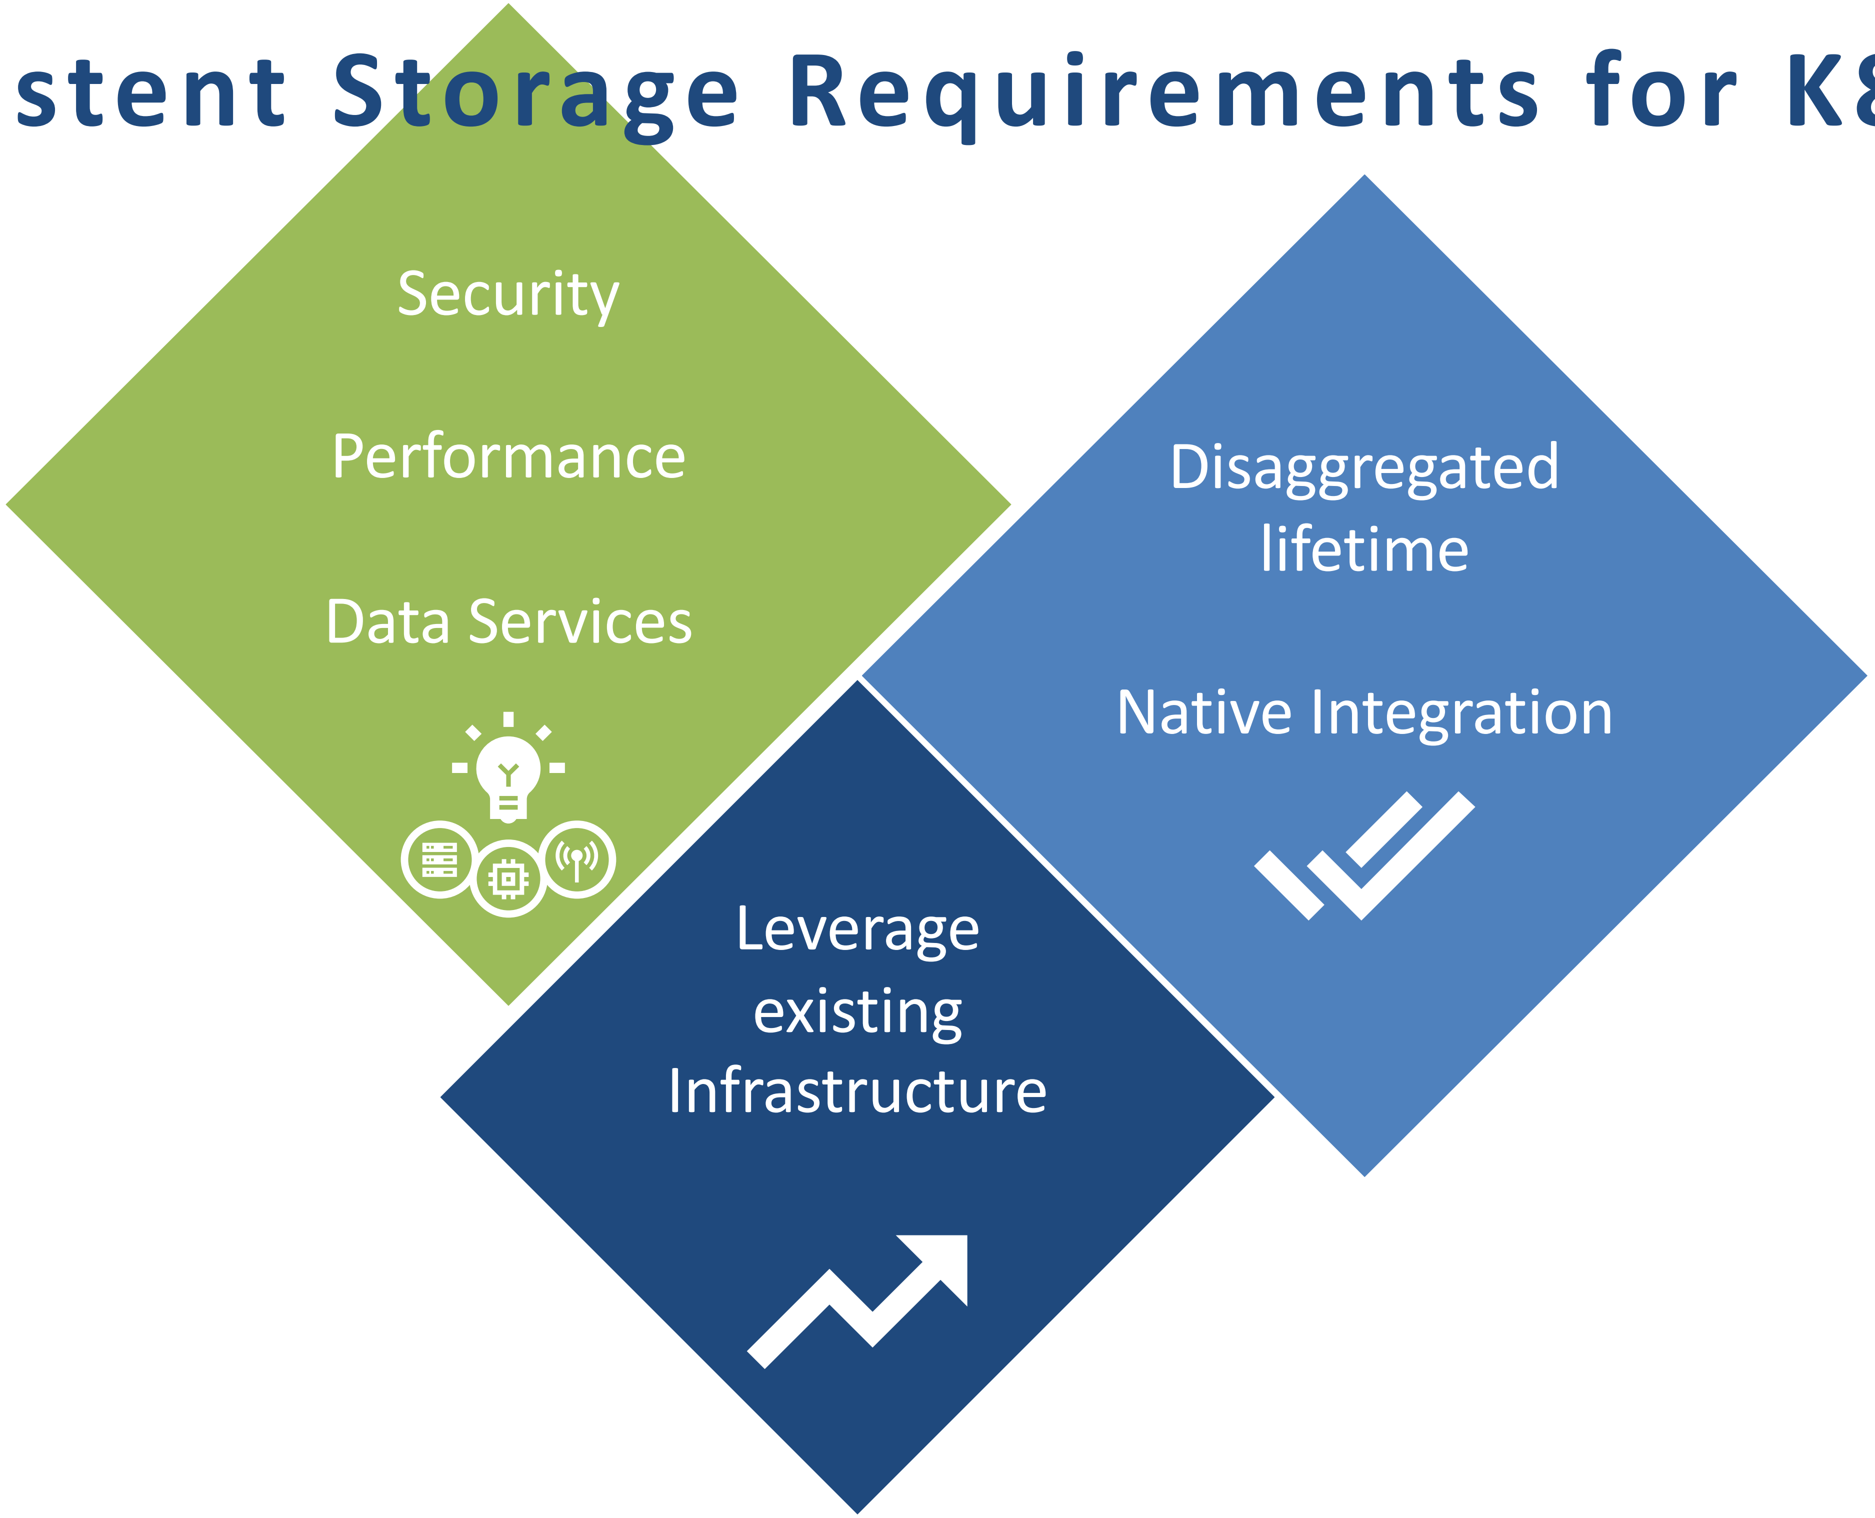

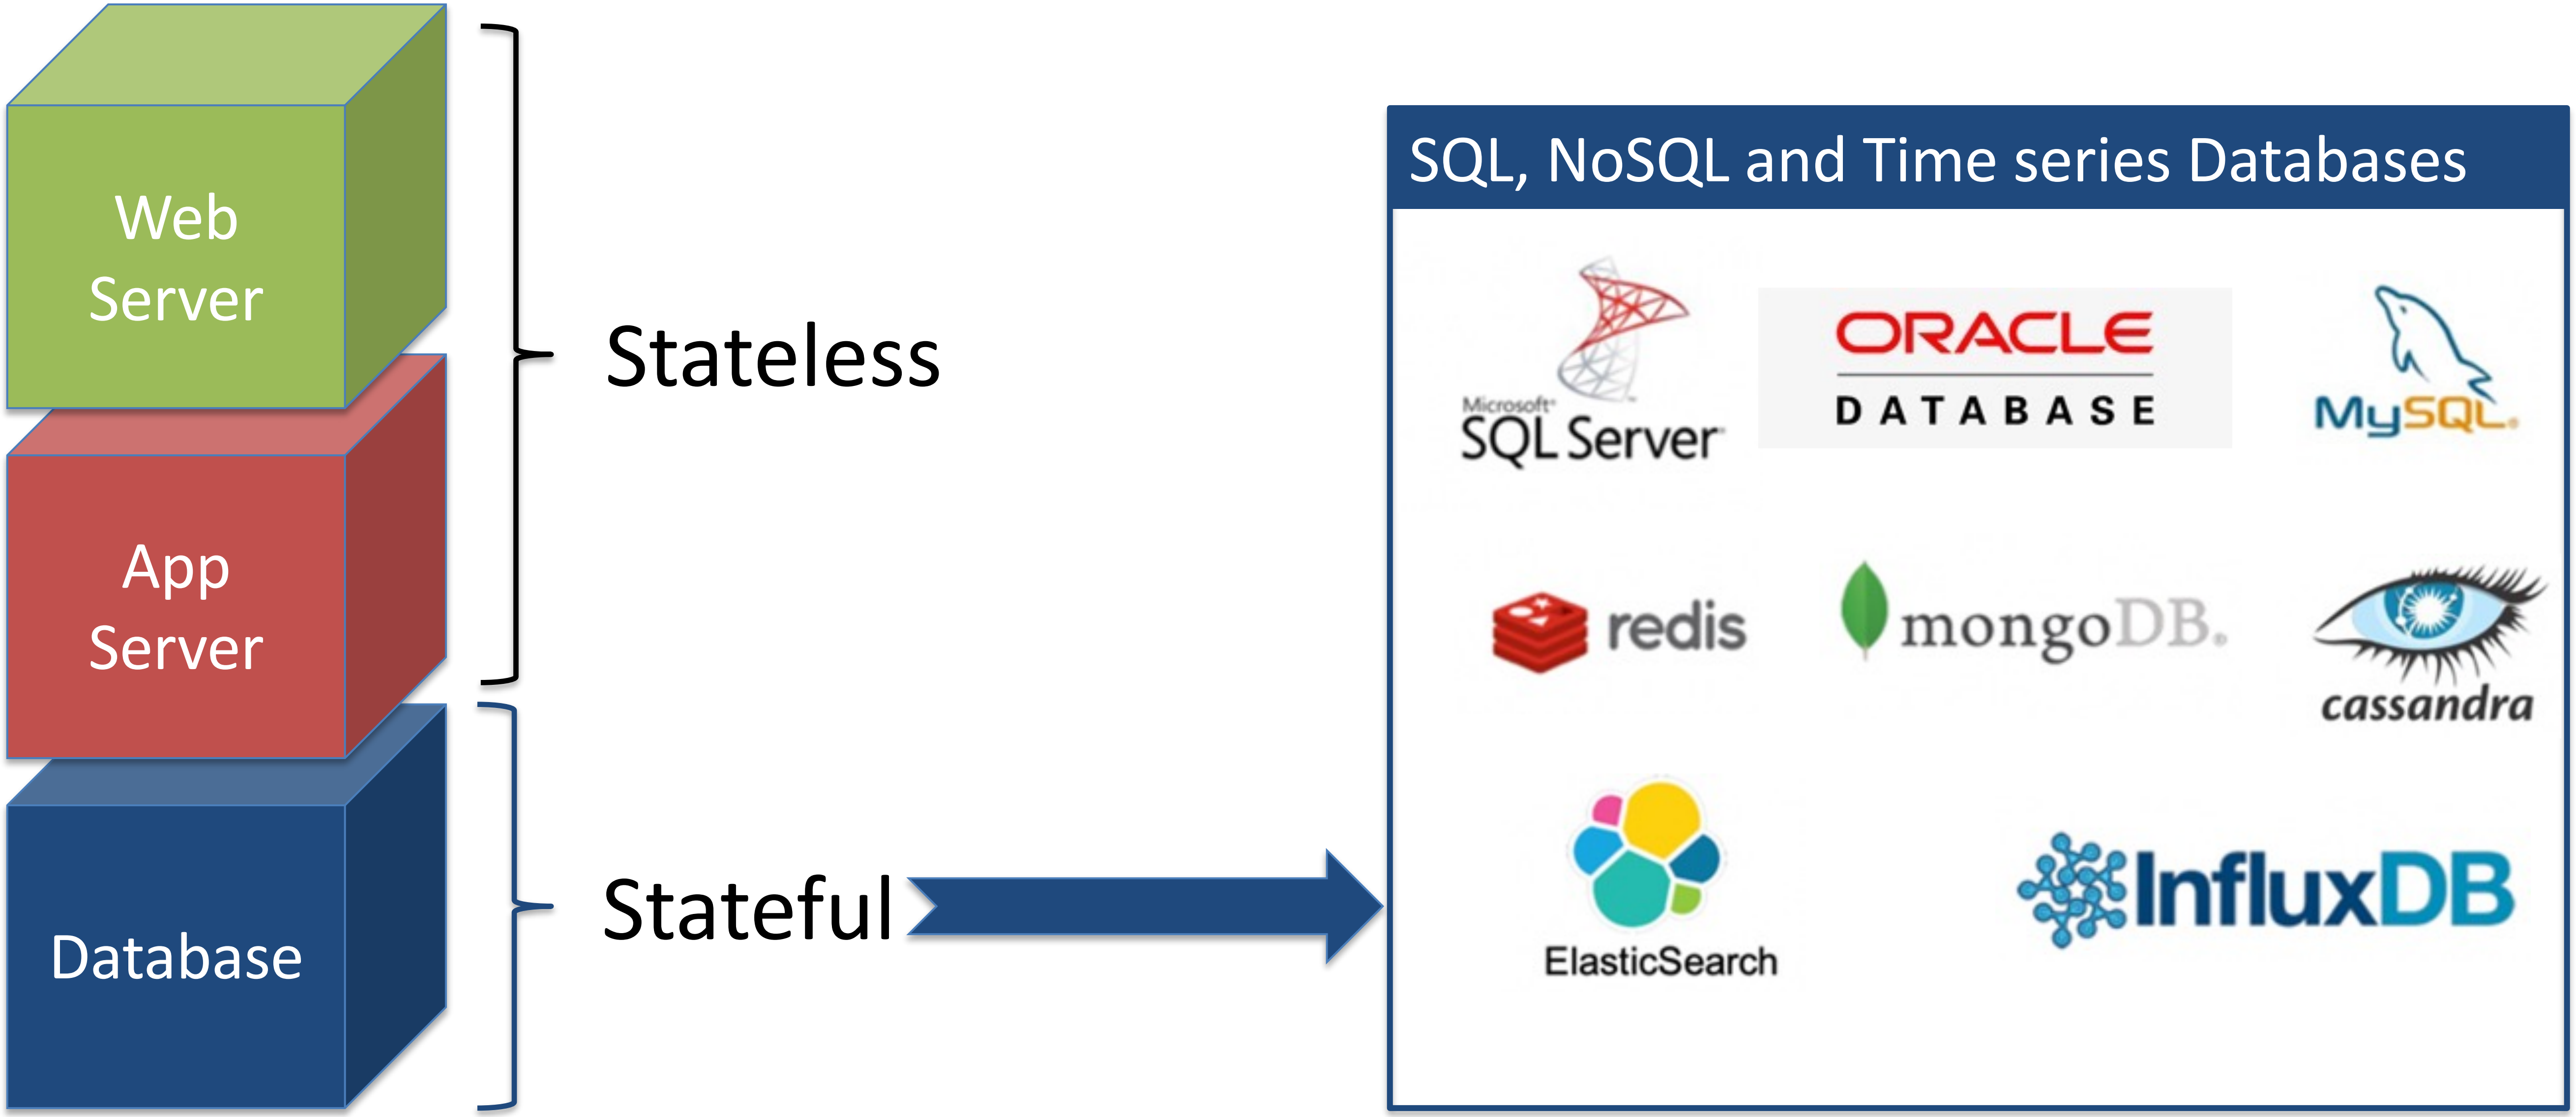

Stateful Applications have special storage requirements

Stateful Applications Require Stronger Guarantee of Storage

Shared Data Sets

Scale Out

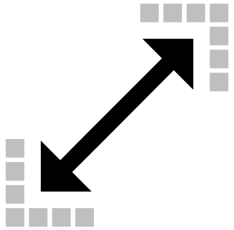

QoS

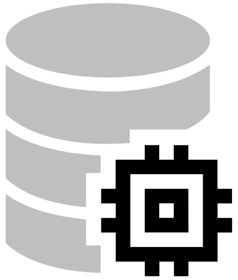

IO Offloads

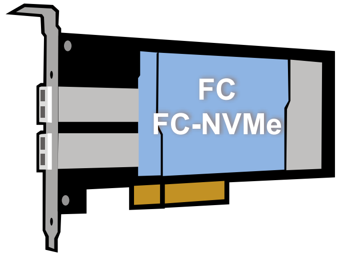

Snapshots

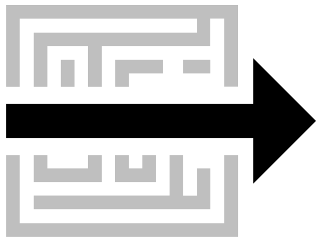

Encryption

Ecosystem

Disaster Recovery

Fibre Channel

FUTURE PROOF

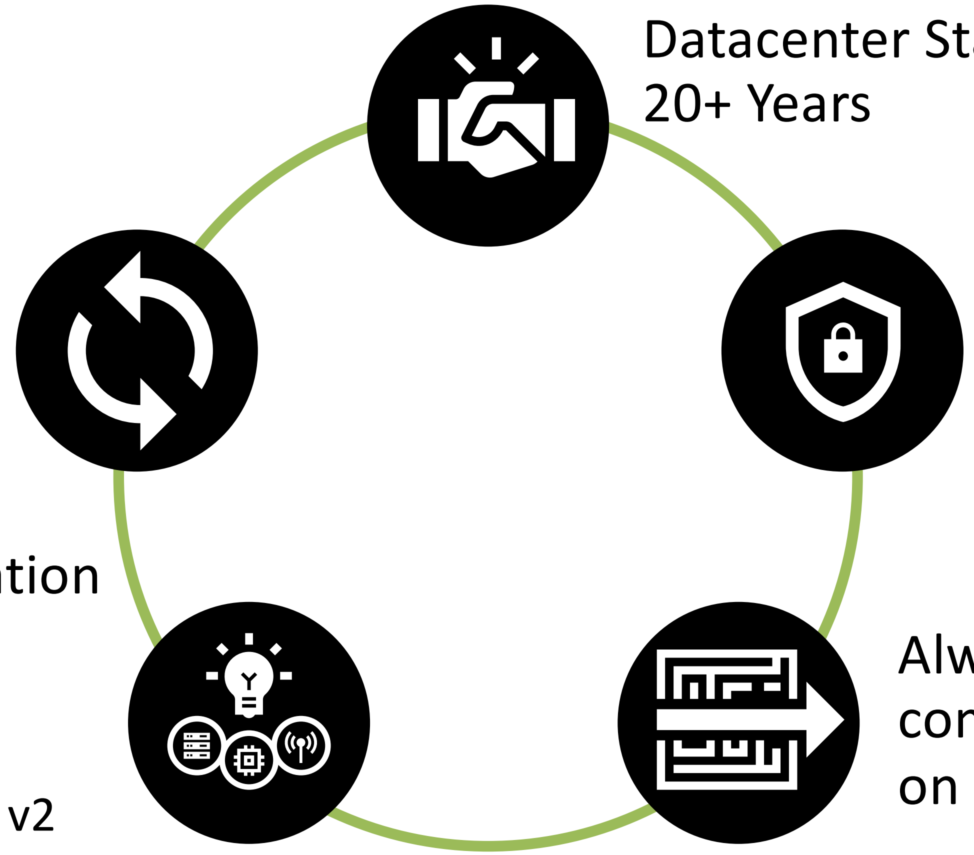

- Leading standard for storage connectivity
- Designed and Trusted for mission and business critical workloads
- Low latency and high performance, with low overhead

FC delivers abstracted, persistent, shared and high-performance storage for K8s

Journey to Cloud Native

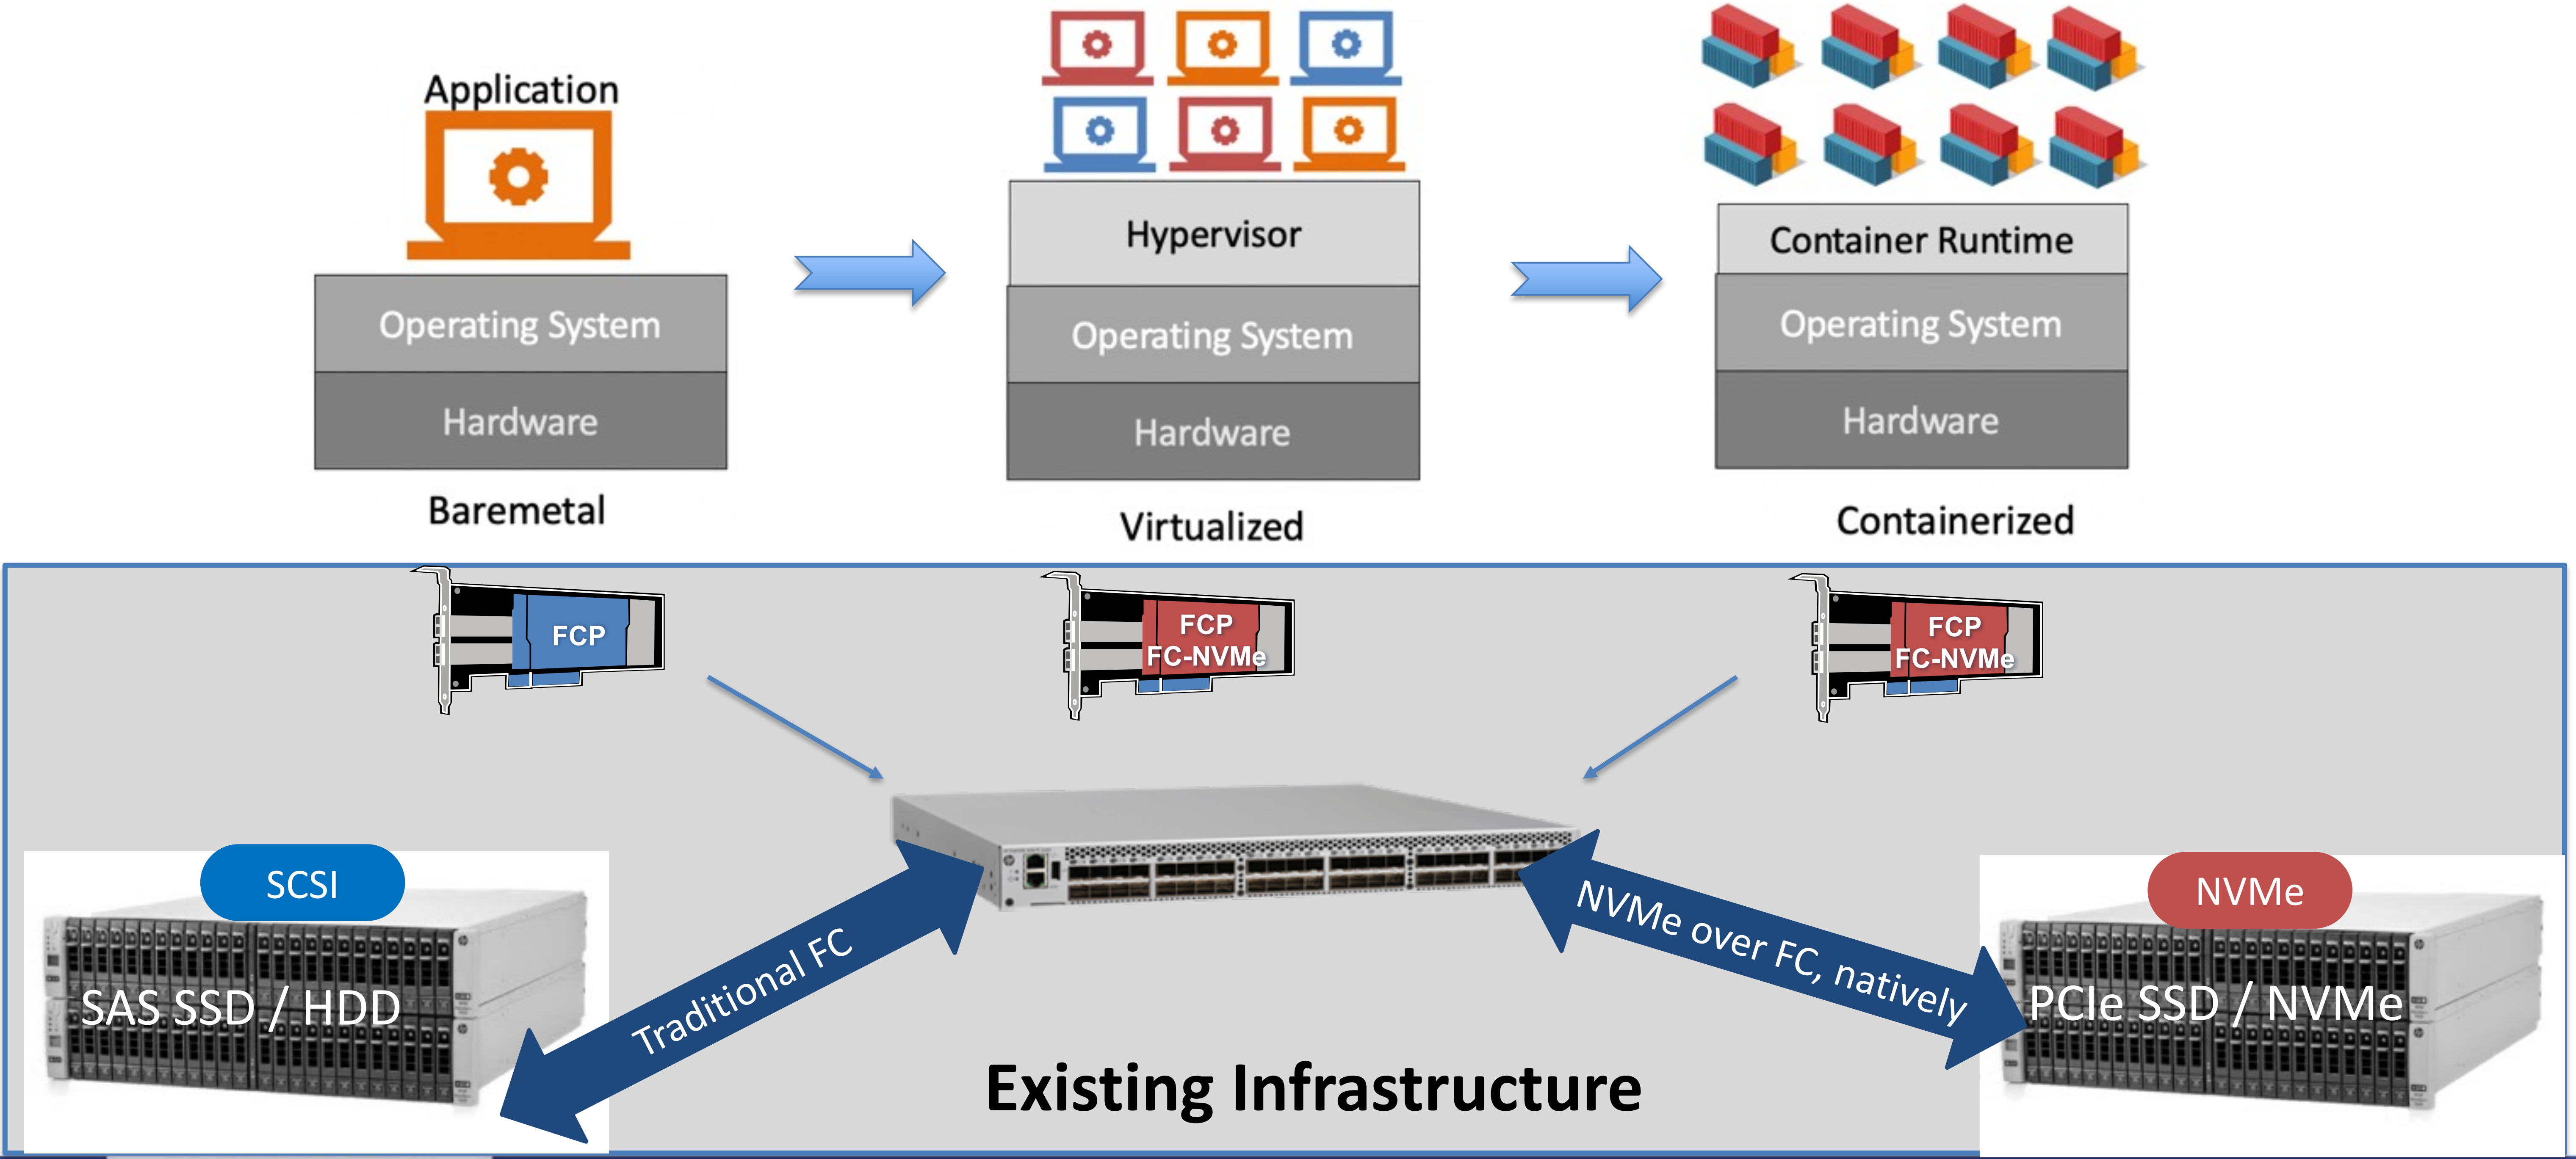

Convergence with FC SANs

- Not all applications can be containerized
- Shared storage and Fibre Channel provides an excellent way to transition and co-exist into cloud native
- Virtual Machines can integrates directly into existing K8s clusters
- Schedule, connect, and consume VM resources as container-native
    - CSI, CNI

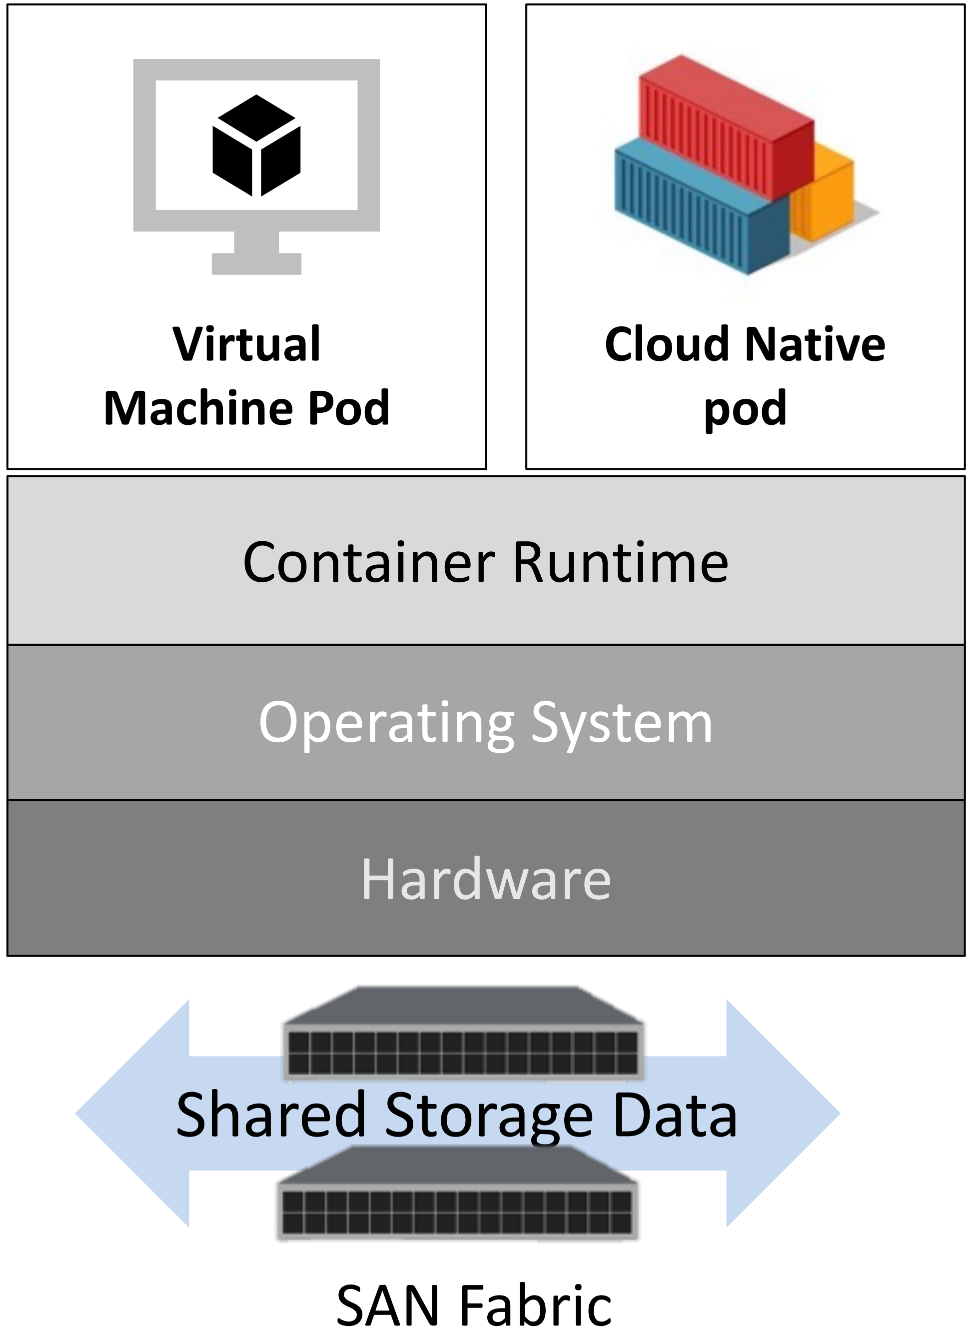

The Evolution of Container Storage Interfaces and CSI

Matt LeVan

3 Types of Kubernetes Storage Drivers

In-Tree Drivers

- Drivers were originally built, linked, compiled, and shipped with core Kubernetes binaries and extend the core Kubernetes API
- Challenging to add support for new volume plugins to Kubernetes
- Third-party storage code caused reliability and security issues in core Kubernetes binaries
- As such, In-Tree Drivers have **not** been accepted since Kubernetes 1.8

FlexVolume Drivers

- Driver uses an execution-based model to interface with other drivers
- FlexVolume drivers are installed on every node in the volume plugin path
- If a client requires that master node(s) have control over attach capabilities, a FlexVolume driver can also be installed on master node(s)

Container Storage Interface (CSI)

- The Container Storage Interface (CSI) is a standard for exposing arbitrary block and file storage storage systems to containerized workloads on Container Orchestration (CO) systems such as Kubernetes
- Using CSI third-party storage, vendors can write and deploy plugins that expose new storage systems in Kubernetes — without having to touch the core Kubernetes code
- CSI has become the **standard** best-practice method for using storage with Kubernetes orchestration

Basic Storage Operations for FlexVolume and CSI

• Create/Delete Volume

- Responsible for the creation and deletion of storage volumes

• Attach/Detach Volume

- Responsible for attaching / mapping volumes to a node (or the whole cluster, depending on driver)
- Likewise, can also detach / unmap the volume from the node
- Typically these operations are called from the Controller Manager
- FlexVolume storage can be handled by master or non-master nodes, depending on driver

• Mount/Unmount Volume

- Responsible for 'mounting' device(s) to a global path, which individual containerized apps (pods) can then bind mount
  - Called only from Kubelet "node agents" that run on each Kubernetes node
  - Some drivers will also create a file system for Block storage devices
- Likewise, can also 'unmount' device(s) from the global path
  - All bind mounts to individual Kubernetes pods must be unmounted, before unmounting from the global path
- Operation is handled by Kubernetes sidecar for CSI drivers

Container Storage Interface (CSI)

- **Standardized storage interface for Container Orchestration (CO) systems**
  - Include: K8s, Mesos, Docker, Cloud Foundry
- **CSI makes installing new volume plugins as easy as deploying a pod**
- **Driven and specified by the CNCF (Cloud Native Computing Foundation)**
  - Since Kubernetes v1.13 (Dec 2018) and onward, CSI has been the recommended specification for Kubernetes storage
- **FlexVolume plugins coexist with CSI plugins**
  - Storage SIG will continue to maintain the FlexVolume API so that existing and future plugins will continue to work

CONTAINER STORAGE INTERFACE

What's Coming for Container Storage Interface (CSI)

- **New volume features will be added <u>only</u> to CSI (not to current FlexVolume)**
- **Future volume features, to be added to CSI and adopted by future Kubernetes versions, include:**
  - Snapshots (Beta in Kubernetes 1.17)
  - Cloning (GA in Kubernetes 1.18)
  - Topology aware volume provisioning (GA in Kubernetes 1.17)
  - Volume expansion (Beta in Kubernetes 1.16)
  - Raw block support (GA in Kubernetes 1.18)
  - Pod Info on Mount (GA in Kubernetes 1.18)
  - Skip Kubernetes Attach and Detach (GA in Kubernetes 1.18)
  - Ephemeral Volumes (Pod Inline Volume Support
    - CSI Ephemeral Inline Volumes (Beta in Kubernetes 1.16)
    - Generic Ephemeral Inline Volumes c
  - Volume Limits (GA in Kubernetes 1.17)
  - Storage Capacity Tracking (Alpha in Kubernetes 1.19)
  - Volume Health Monitoring Feature (Alpha in Kubernetes 1.19)
  - CSI Driver fsGroup Support (GA in Kubernetes 1.23)

CSI in Kubernetes — High Level Architecture

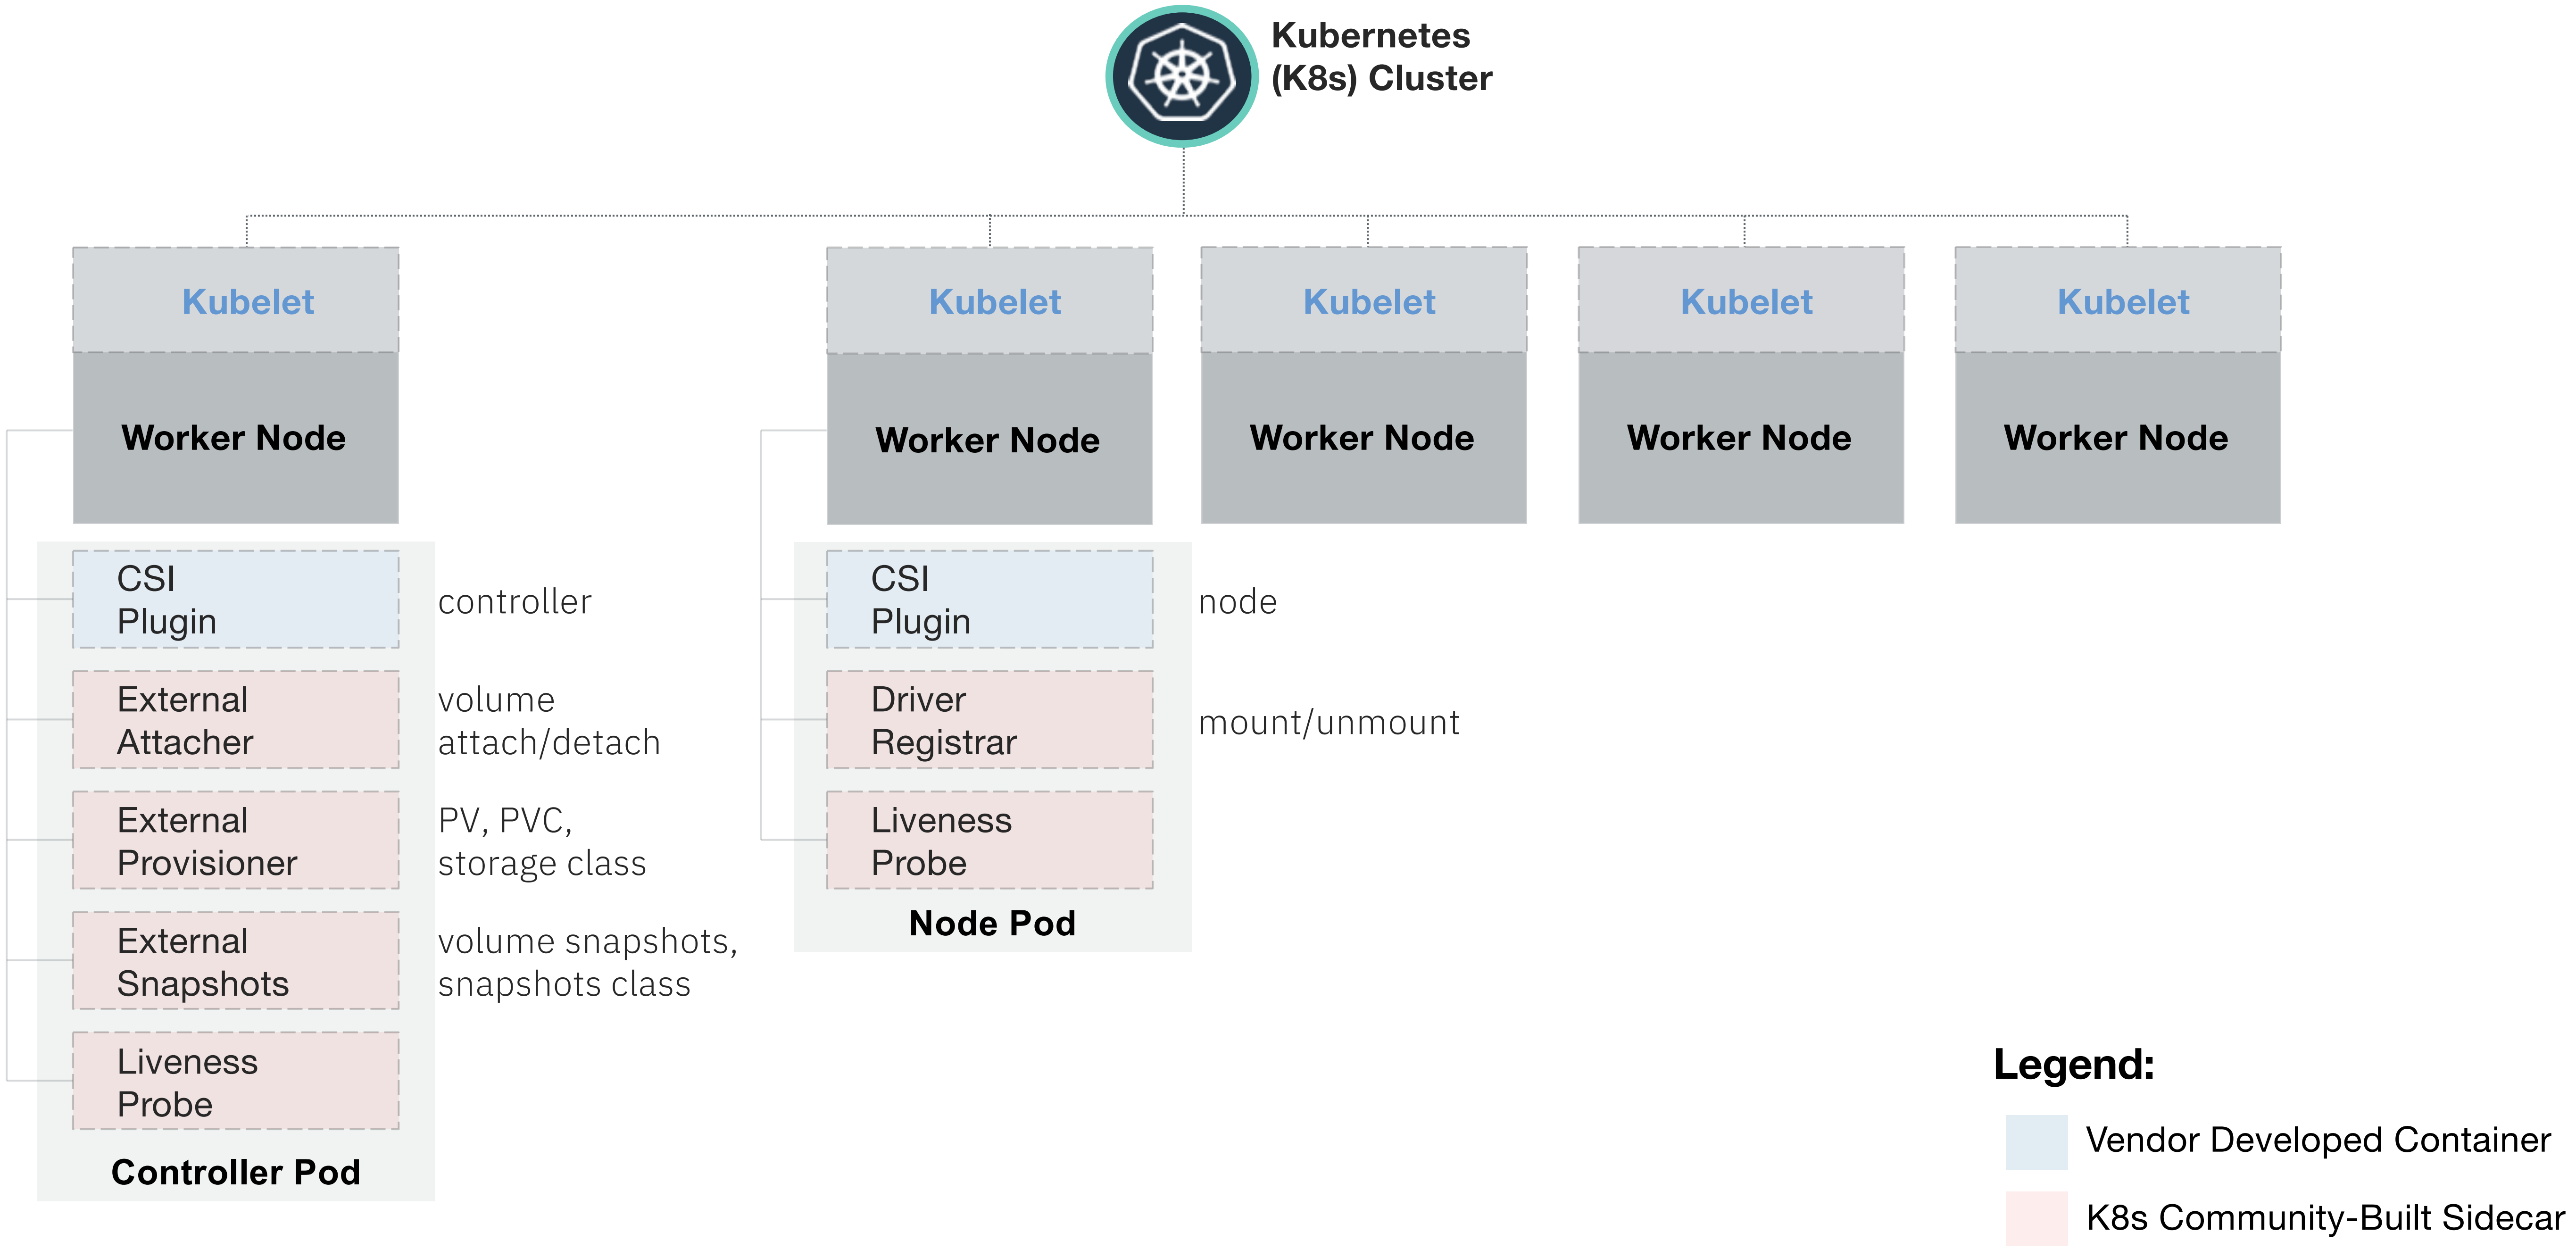

CSI + Kubernetes Integration

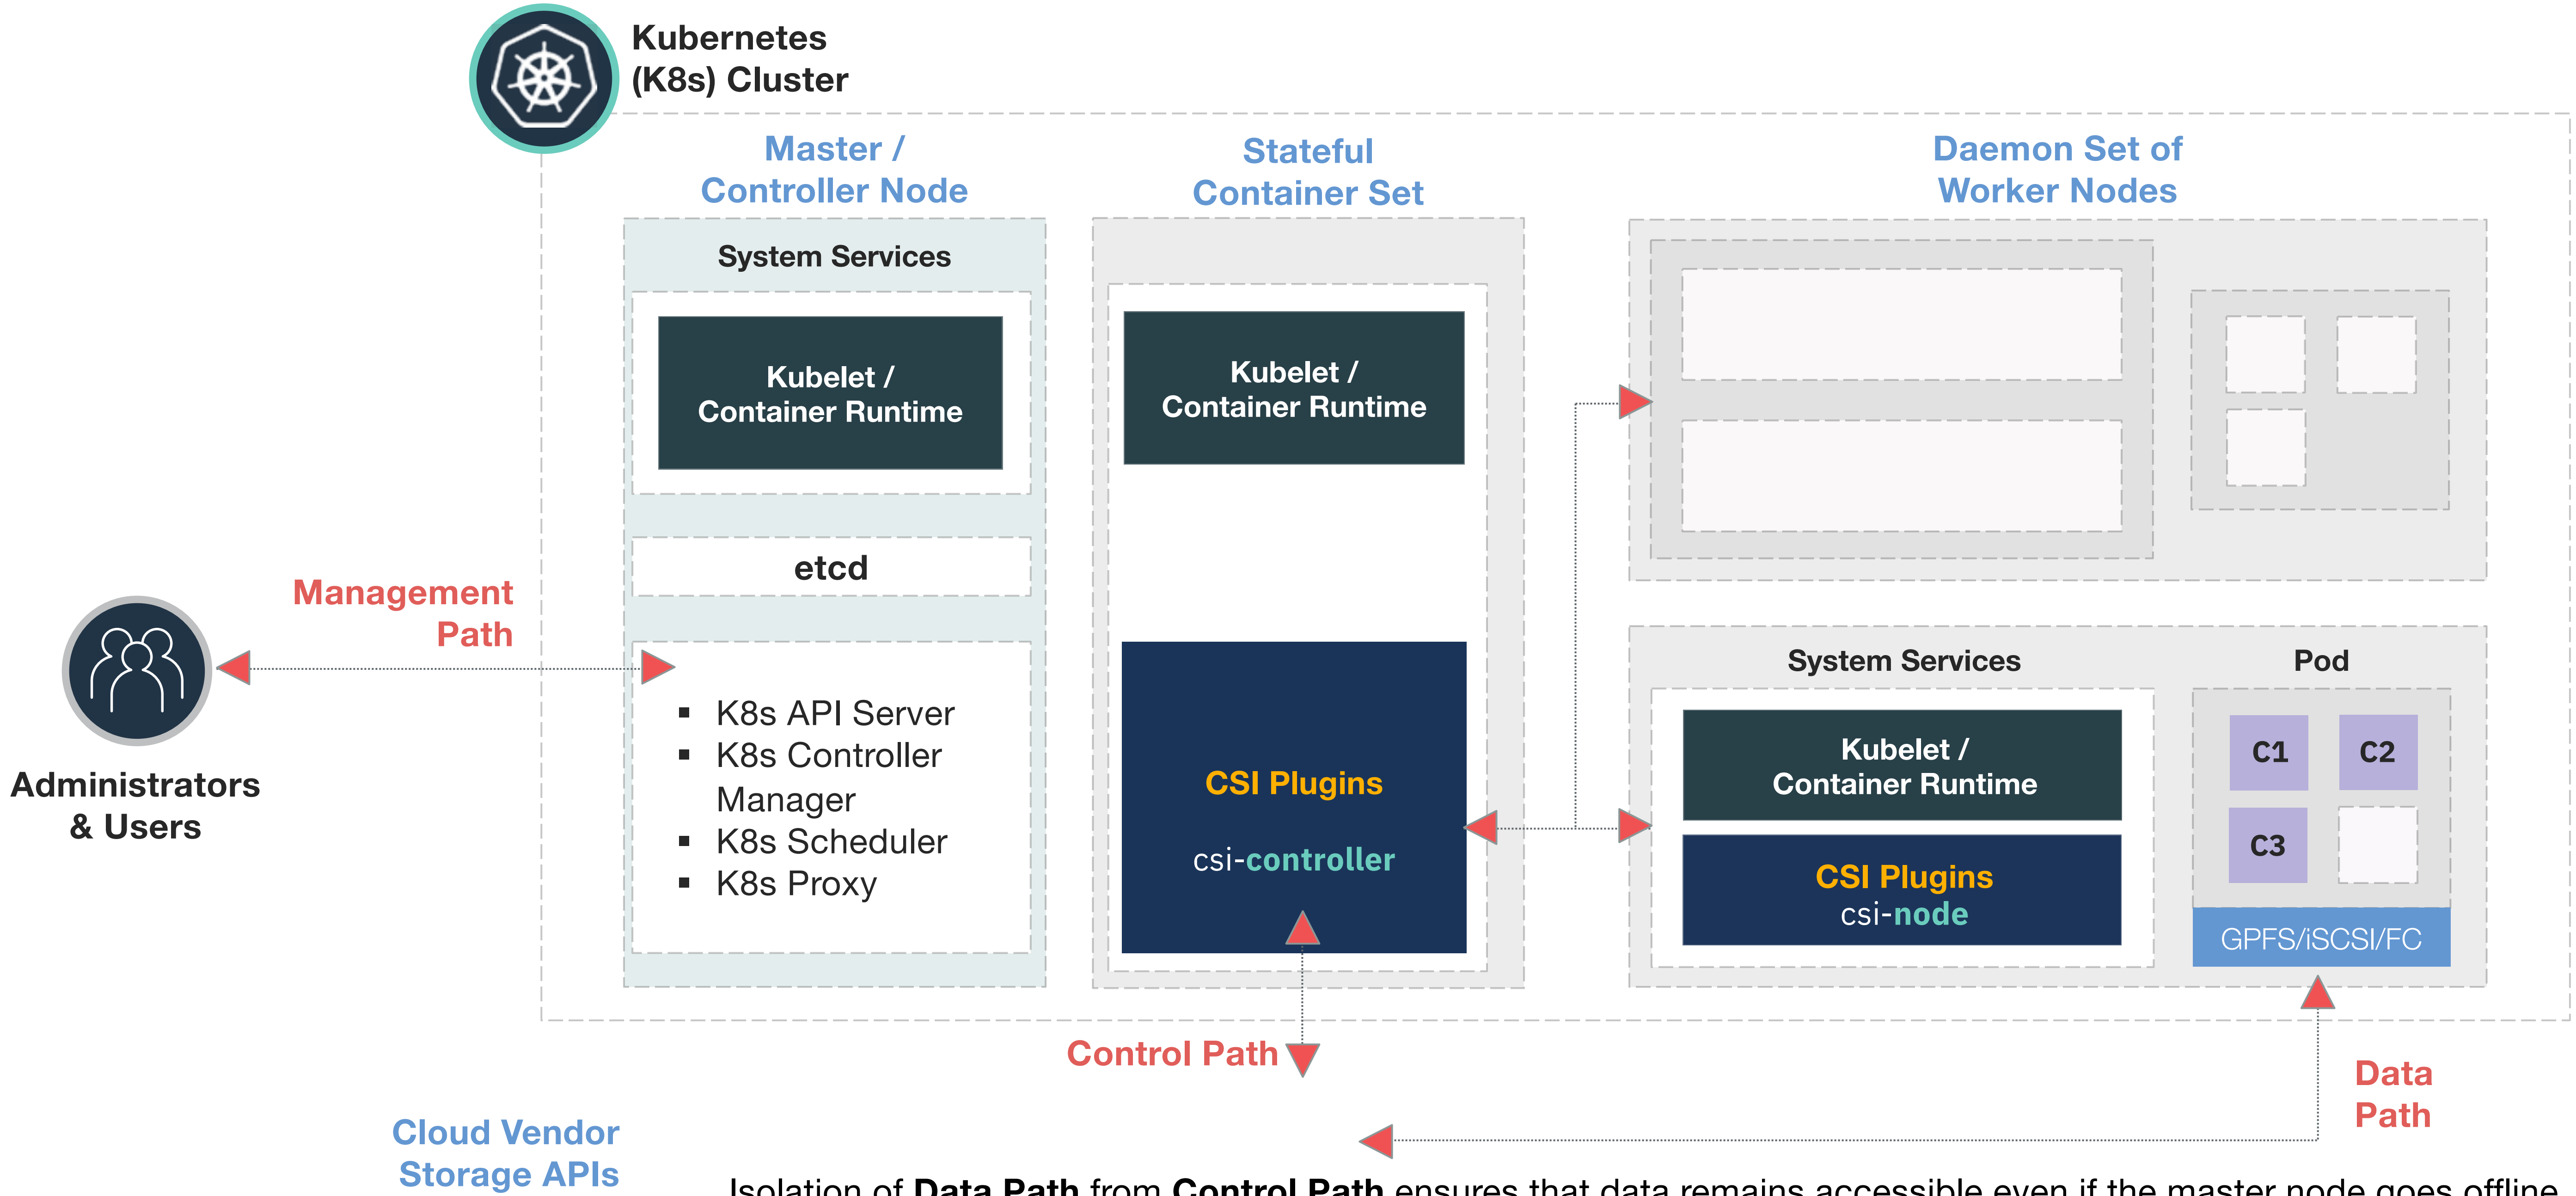

Isolation of **Data Path** from **Control Path** ensures that data remains accessible even if the master node goes offline.

Kubernetes Storage Terminology

Persistent Volume (PV) is a unit of storage that has been provisioned by an administrator or dynamically provisioned via a storage driver or plug-in.

Persistent Volume Claim (PVC) is a request for storage by a user. PVCs consume PV resources.

StorageClass (SC) provides the ability to describe the classes of storage that can be dynamically provisioned. SC will be used to map to quality-of-service, backup or other storage features.

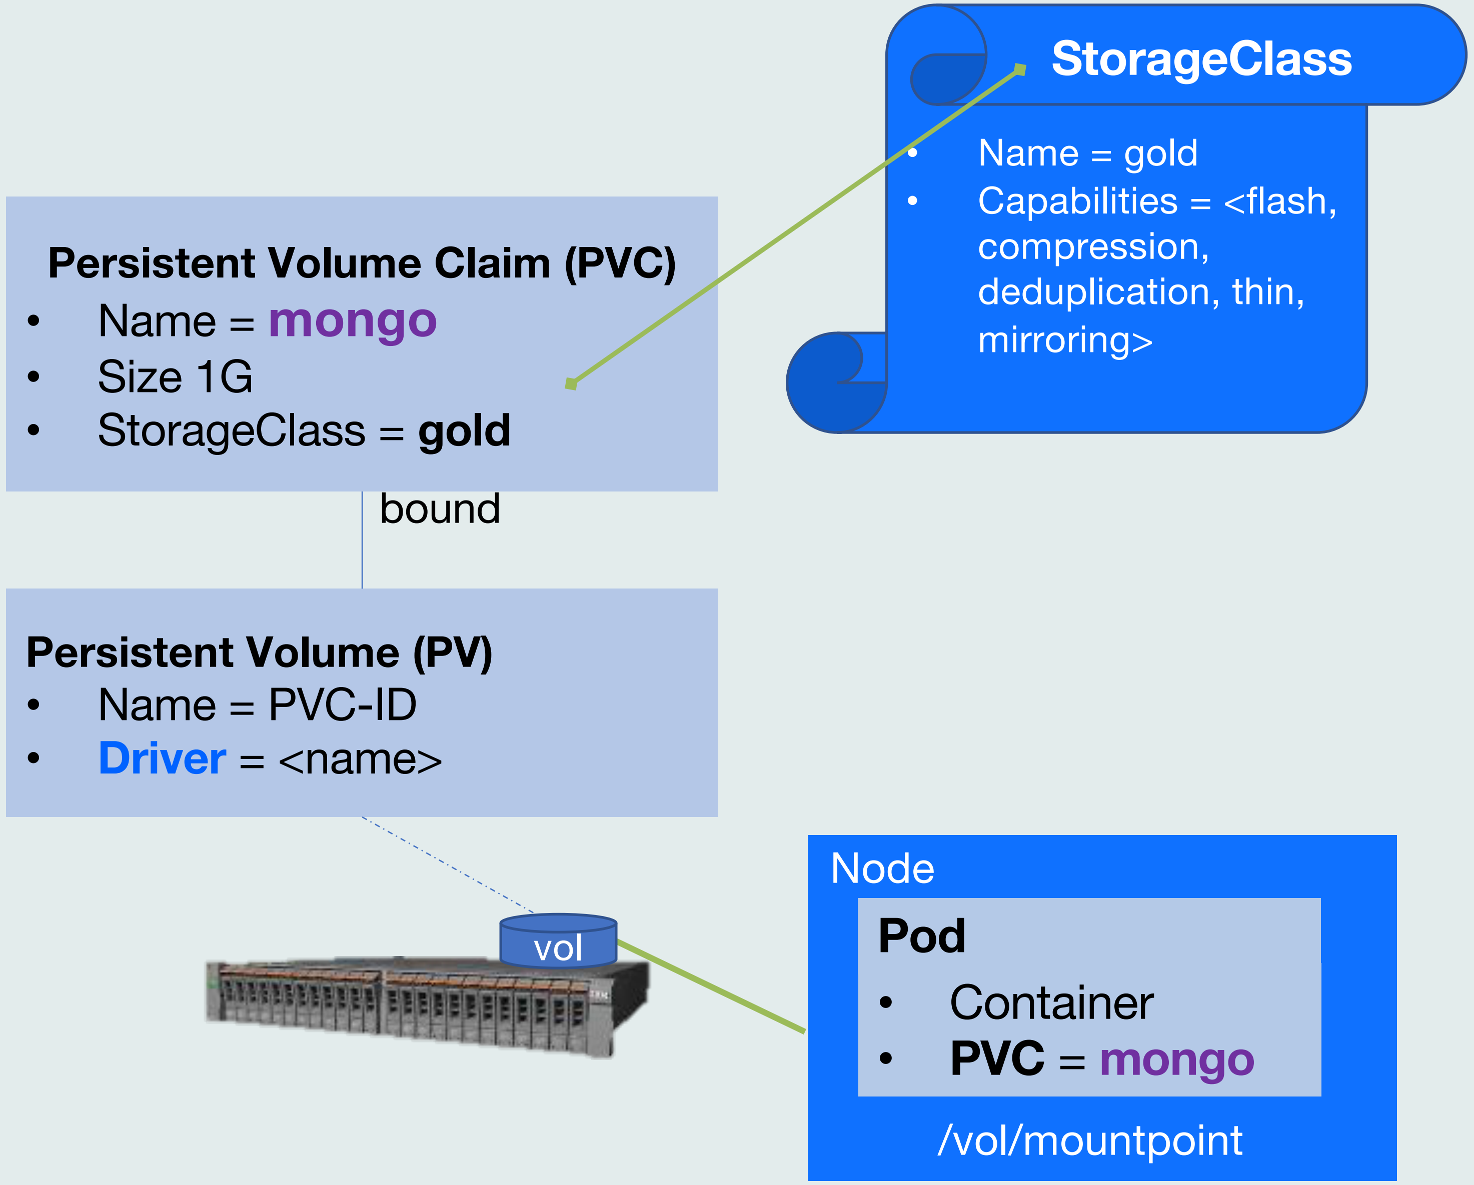

New Definitions of Storage for Containers

Container-ready Storage

- Various classes of storage, including:
  - Storage Area Network (SAN) devices
  - Software Defined Storage (SDS)
  - Network Attached Storage (NAS)
- Leverage your existing infrastructure, processes, management, and monitoring
- Features such as snapshots, clones, and replication can still be used
- Offerings include:
  - SAN
  - NAS

Container-native Storage

- Storage is deployed inside containers and presented to application containers
  - Hyper-converged approach to storage & compute for containerized applications
  - Use internal drives or external storage for consumption
  - Management through a single control plane within K8s
- Depending on protection scheme, can be up to 3 copies of the data stored
  - Can be based on object, file, block, or other types
  - Generally includes snapshots, clones, and/or replication
- Offerings include:
  - Red Hat OpenShift Container Storage (Rook + Ceph + Multi-Cloud Object Gateway)
  - Portworx
  - Robin.io

Container-ready and Container-native Storage

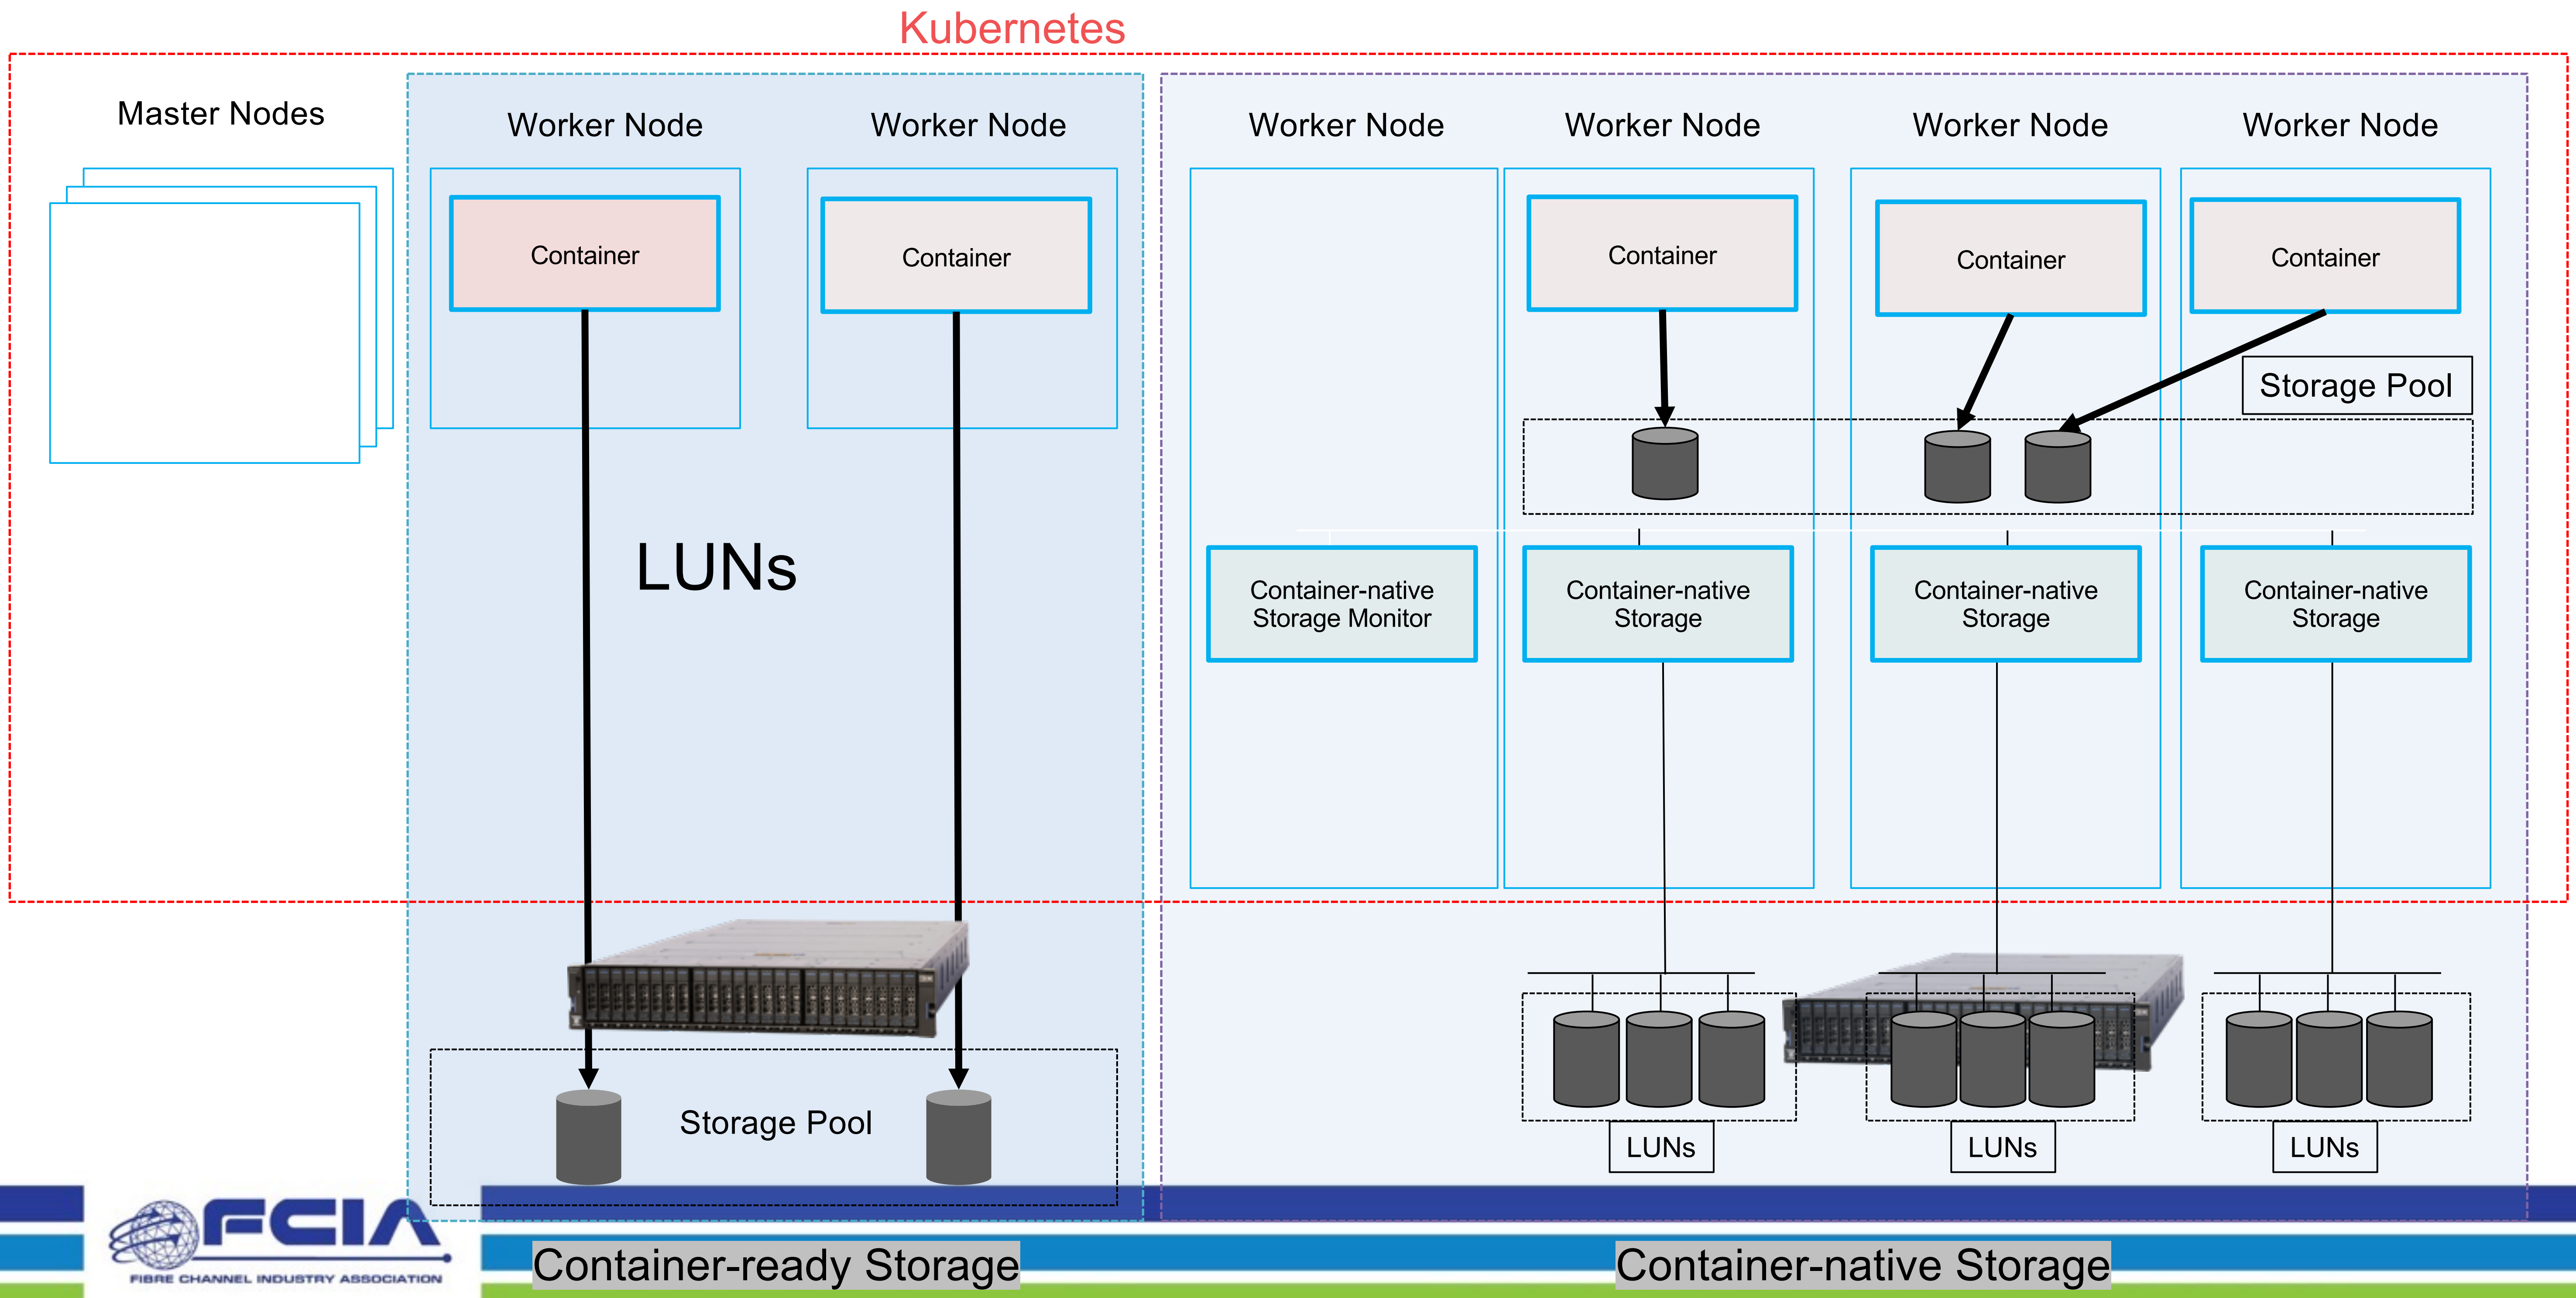

Summary and Q&A

- Organizations are making new investment in storage to return workloads from the cloud for security and performance reasons
- Fibre Channel is the perfect choice for secure, high performance persistent storage for Kubernetes orchestrated containers
- CSI has become the standard best-practice method for using third-party SAN connected storage solutions with Kubernetes orchestration
  - Ask your array vendor for their CSI plugin
- Fibre Channel connectivity is included in Kubernetes distributions today

After this Webcast

- Please rate this event – we value your feedback
- We will post a Q&A blog at http://fibrechannel.org/ with answers to the questions we received today
- Follow us on Twitter @FCIAnews for updates on future FCIA webcasts
- Our library of FCIA on-demand webcasts http://fibrechannel.org/webcasts/ includes:
  - Fibre Channel Fundamentals
  - FC-NVMe
  - Long Distance Fibre Channel
  - Fibre Channel Speedmap
  - FCIP (Extension): Data Protection and Business Continuity
  - Fibre Channel Performance
  - FICON
  - Fibre Channel Cabling
  - 64GFC
  - FC Zoning Basics

Thank You

FIBRE CHANNEL INDUSTRY ASSOCIATION