STORAGEOS

8 Core Principles of Cloud Native Storage

A paradigm shift in storage design and delivery

Storage must evolve.

To support the convergence of container and cloud adoption.

Developers

Have embraced cloud native and container technology for the ease and flexibility of deploying web scale applications.

Enterprises

Are rapidly adopting cloud and container-based infrastructure as they realize the value of flexible infrastructure, replacing operator-managed VMWare with IaaS, PaaS and pay as you go compute and storage.

Yet storage, needed for all apps, is being pulled in many directions…

…And it cannot keep up.

Enterprises want to:

Increase flexibility, performance and agility with containers and cloud.

Minimize the risks of any transition.

Maximize the usage of their existing resources.

The result? A paradigm shift in storage design and delivery is needed.

And it needs to be

cloud native.

What is cloud native?

- Horizontally scalable
- Built to handle failures / No single point of failure
- Resilient and survivable (self healing)
- Minimal operator overhead – API driven and automatable
- Decoupled from underlying platform

Cloud native storage is needed to support today’s IT.

8

Core Principles of Cloud Native Storage

1

Application Centric

Storage should be presented to and consumed by applications and not by operating systems or hypervisors.

It is no longer desirable to present storage to OS instances, and then, later, have to map applications to OS instances to link to storage (whether on-premises or in a cloud provider, or on VMs or bare metal). Storage needs to be able to follow an application as it scales, grows, and moves between platforms and clouds.

2

Application Platform Agnostic

The storage platform should be able to run anywhere and not have proprietary dependencies that lock an application to a particular platform or a cloud provider.

Storage should be shaped just like an application – capable of being scaled out in a distributed topology just as easily as it could be scaled up based on application requirements. Upgrades and scaling of the storage platform should be implemented as a non-disruptive operation to the running applications.

3

Declarative and Composable

Storage resources should be declared and composed just like all other resources required by applications and services.

This allows storage resources and services to be deployed and provisioned as part of application instantiation through orchestrators.

4

API Driven and Self-Managed

Storage resources and services should be easy to provision, consume, move and manage via an API.

Storage should provide integration with application runtime and orchestrator platforms.

5

Agile

The platform should be able to dynamically react to changes in the environment and be able to:

- Move application data between locations
- Dynamically resize volumes for growth
- Take point in time copies of data for data retention or facilitate rapid recovery of data
- Integrate naturally into dynamic, rapidly changing application environments

6

Natively Secure

Storage services should integrate and provide inline security features such as encryption and RBAC and not depend on secondary products to secure application data.

7

Performant

The storage platform should offer deterministic performance in complex distributed environments and scale efficiently using a minimum of compute resources.

It is no longer desirable to present storage to OS instances, and then, later, have to map applications to OS instances to link to storage (whether on-premises or in a cloud provider, or on VMs or bare metal). Storage needs to be able to follow an application as it scales, grows, and moves between platforms and clouds.

8 Consistently Available

The storage platform should manage data distribution with a predictable, proven data model to ensure high availability, durability and consistency of application data.

Data recovery processes during failure conditions should be application independent and not affect normal application operations.

This is the StorageOS
vision of storage design and delivery for
Cloud Native Storage.

STORAGEOS

a Cloud Native Storage solution

Designed and built from the ground up with a patent pending storage solution.

Allows users to deploy a cloud native storage platform in their environments.

STORAGEOS

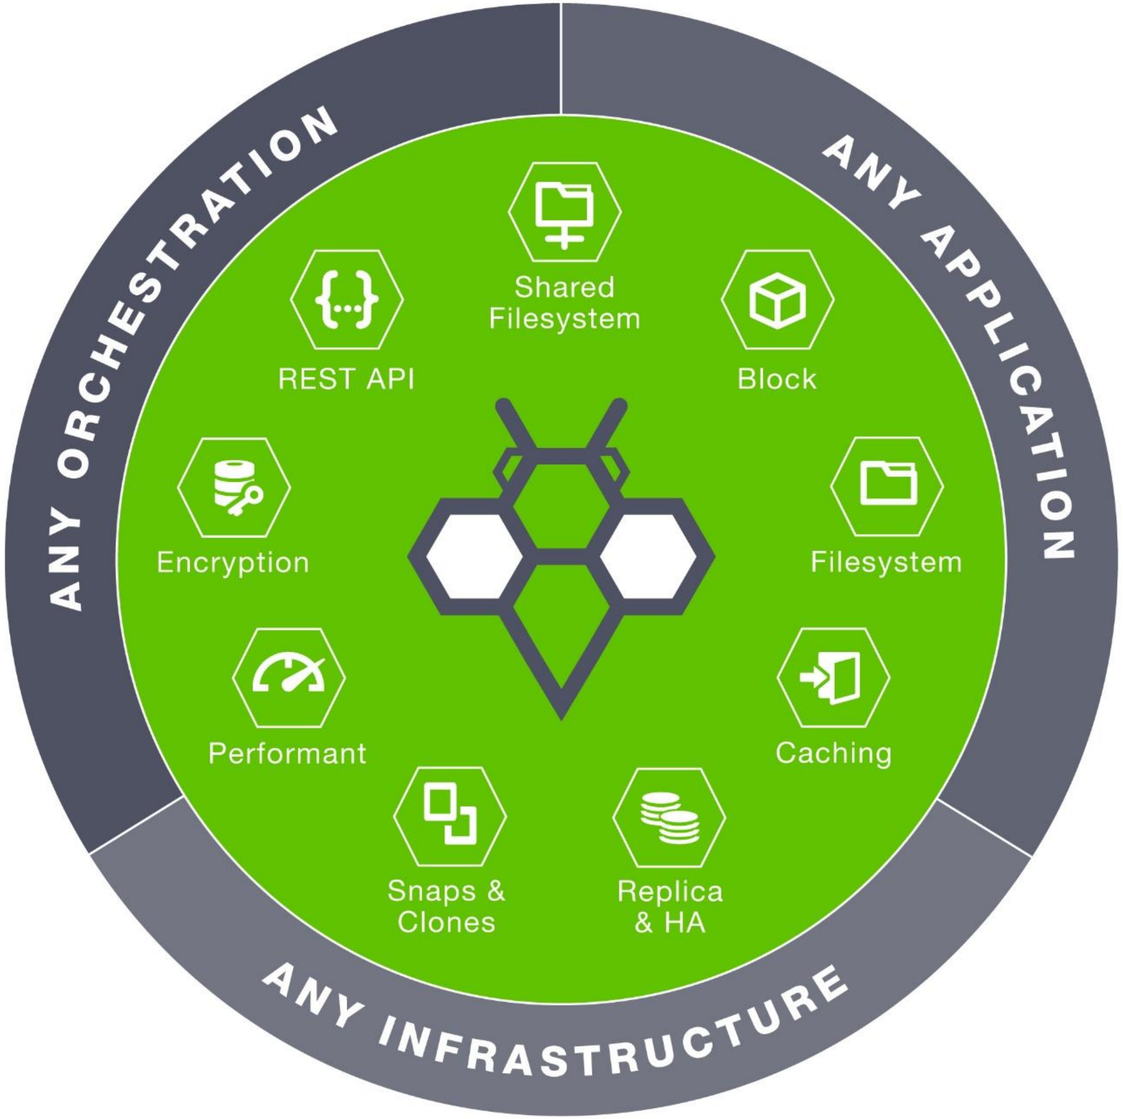

ANY ORCHESTRATION
ANY APPLICATION
ANY INFRASTRUCTURE
Shared Filesystem
REST API
Block
Encryption
Filesystem
Performant
Caching
Snaps & Clones
Replica & HA

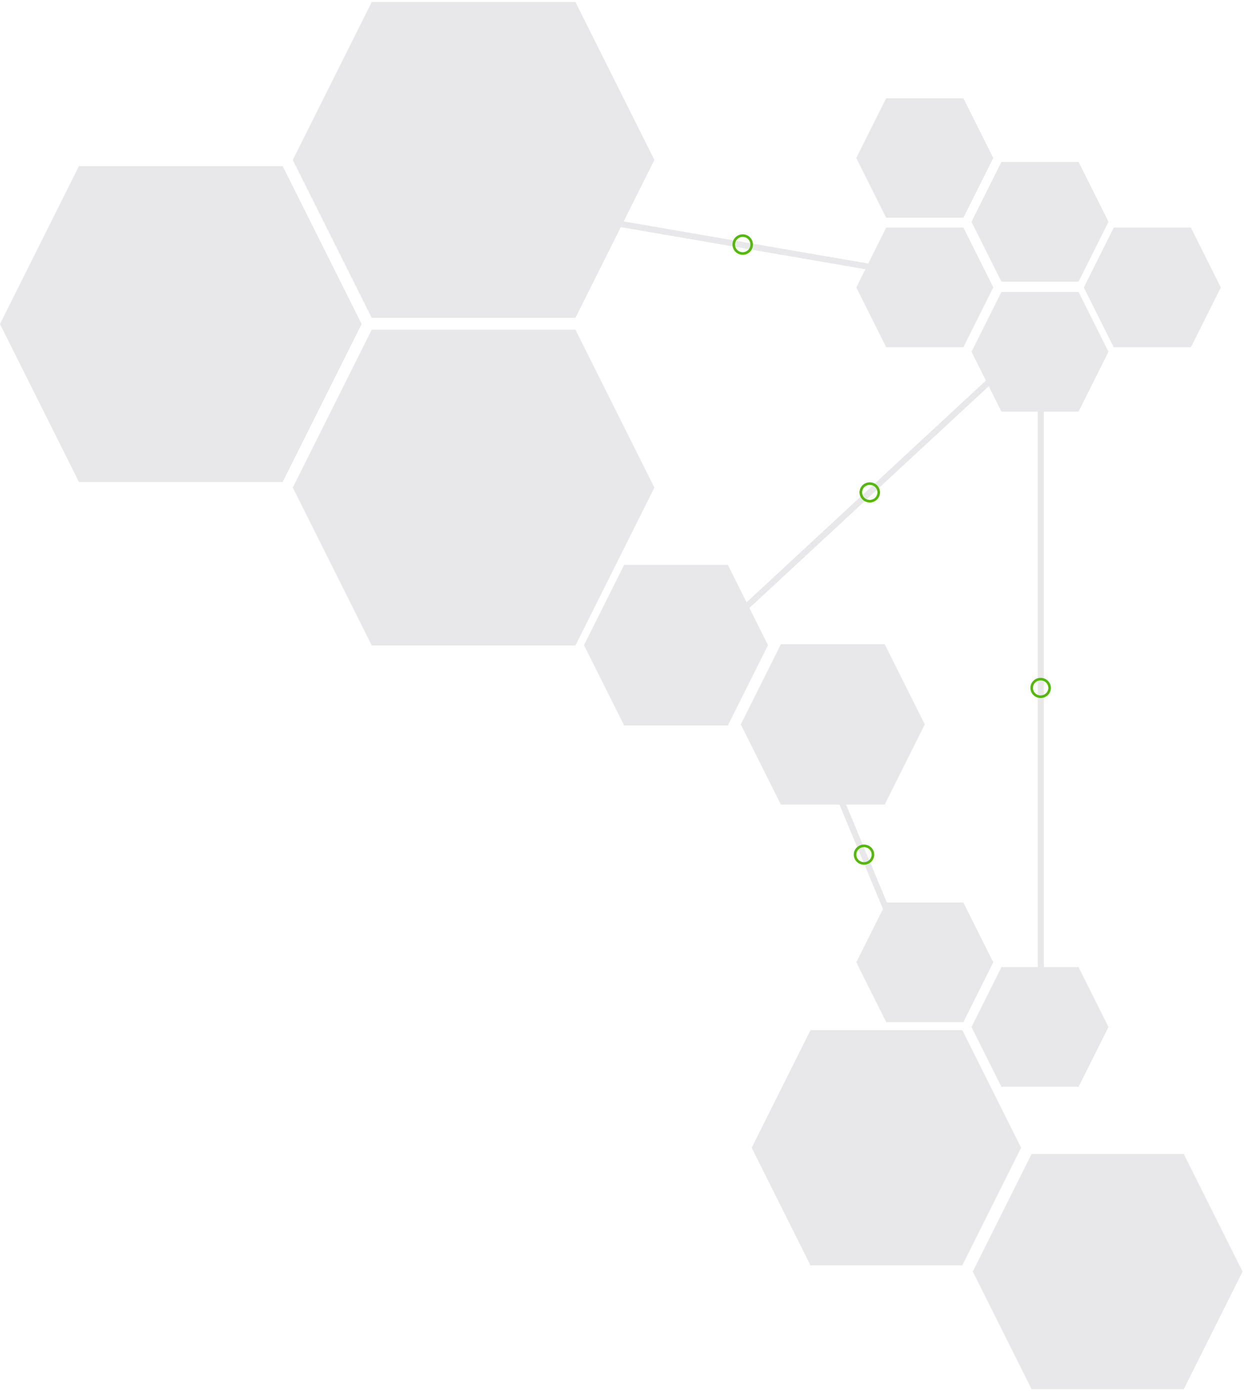

STORAGEOS

Want to try StorageOS?

Get started with StorageOS for free

docker store

Docker Hub

kubernetes

STORAGEOS

STORAGEOS

Get in Touch

slack in