WHY SOFTWARE DEFINED STORAGE MATTERS

Ross Turk
June 2016

THE DATA EXPLOSION

Web, mobile, social media, cloud

Our digital assets have grown due to web scale services like Facebook, YouTube, and Netflix.

Video on-demand services

Rapid growth of video on-demand has resulted in 50% of households using this service.

Media and entertainment

A staggering amount of content is created during today's optimized production processes.

Medical industry

Medical imaging needs are vast, and regulatory requirements can be demanding.

DATA CHALLENGES

1

Exponential growth in digital content increases pressure on **capacity, scalability, and cost.**

DATA CHALLENGES

1. Exponential growth in digital content increases pressure on **capacity, scalability, and cost.**

2. The need for access to data from anywhere, anytime, on any device requires **unprecedented agility.**

DATA CHALLENGES

1. Exponential growth in digital content increases pressure on **capacity, scalability, and cost.**
2. The need for access to data from anywhere, anytime, on any device requires **unprecedented agility.**
3. Modern services require the **flexibility** to store data on-premises or in the cloud.

DATA CHALLENGES

1. Exponential growth in digital content increases pressure on **capacity, scalability, and cost.**
2. The need for access to data from anywhere, anytime, on any device requires **unprecedented agility.**
3. Modern services require the **flexibility** to store data on-premises or in the cloud.
4. Growing content requires **advanced data protection** that ensures integrity & high availability at very large scale.

THE DATA STORAGE "SHORTFALL"

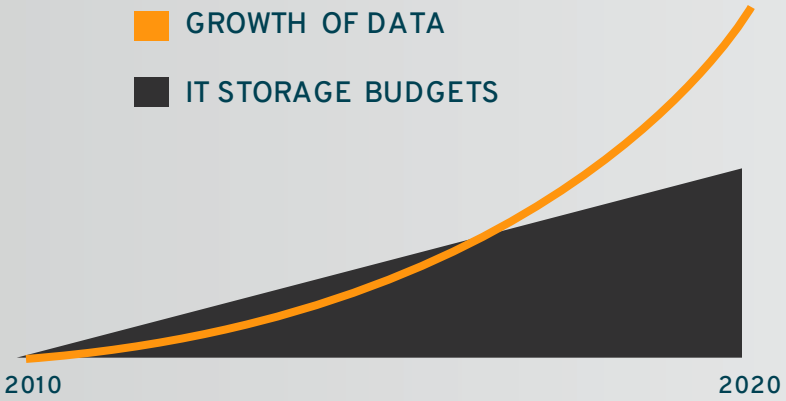

Data stores are growing exponentially, while IT budgets are not

HDDs are becoming more dense, but $/GB decline is slowing

Software and hardware advances are needed to close the gap

WHAT DO WE DO WITH ALL THE DATA TODAY?

PROPRIETARY APPLIANCES

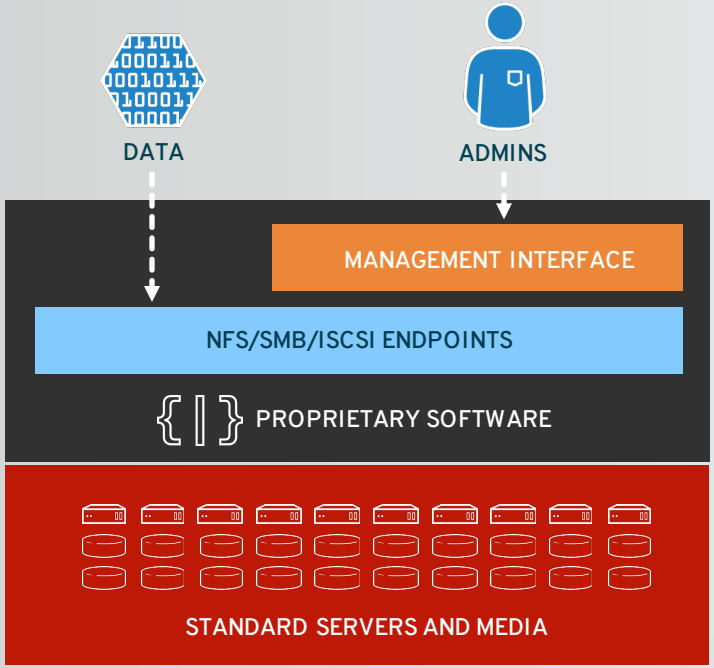

THE TRADITIONAL APPROACH TO STORAGE

- Complexity hidden from end users, along with flexibility
- Vendor lock-in leads to pricing premium
- Price premium over constituent components is difficult to sustain

PUBLIC CLOUD STORAGE

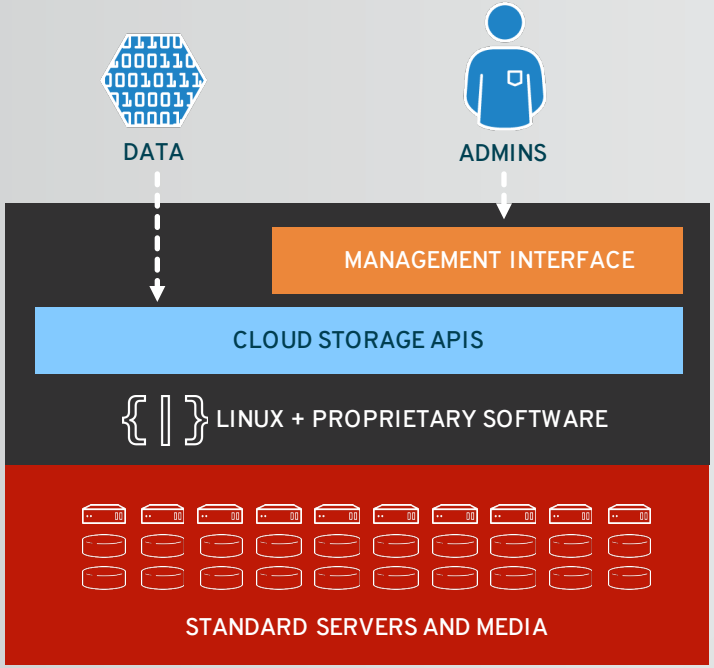

CONVENIENT STORAGE AS A SERVICE

- Complexity still hidden from end users, pay-as-you-go pricing
- Fastest-growing segment of IT storage budgets
- Mostly built with proprietary software (Linux below, “secret sauce” above)

FLEXIBILITY IS EVERYTHING

FLEXIBILITY IS EVERYTHING

ORGANIZATIONS ARE RETHINKING STORAGE

New storage platforms

More efficient use of hardware

Flexible utilization of services

THE DATACENTER IS EVOLVING

Development Model	Application Architecture	Deployment & Packaging	Application Infrastructure	Storage
Waterfall	Monolithic	Bare Metal	Data Center	Scale Up
Agile	N-tier	Virtual Services	Hosted	Scale Out
DevOps	Microservices	Containers	Hybrid Cloud	Software-Defined Storage

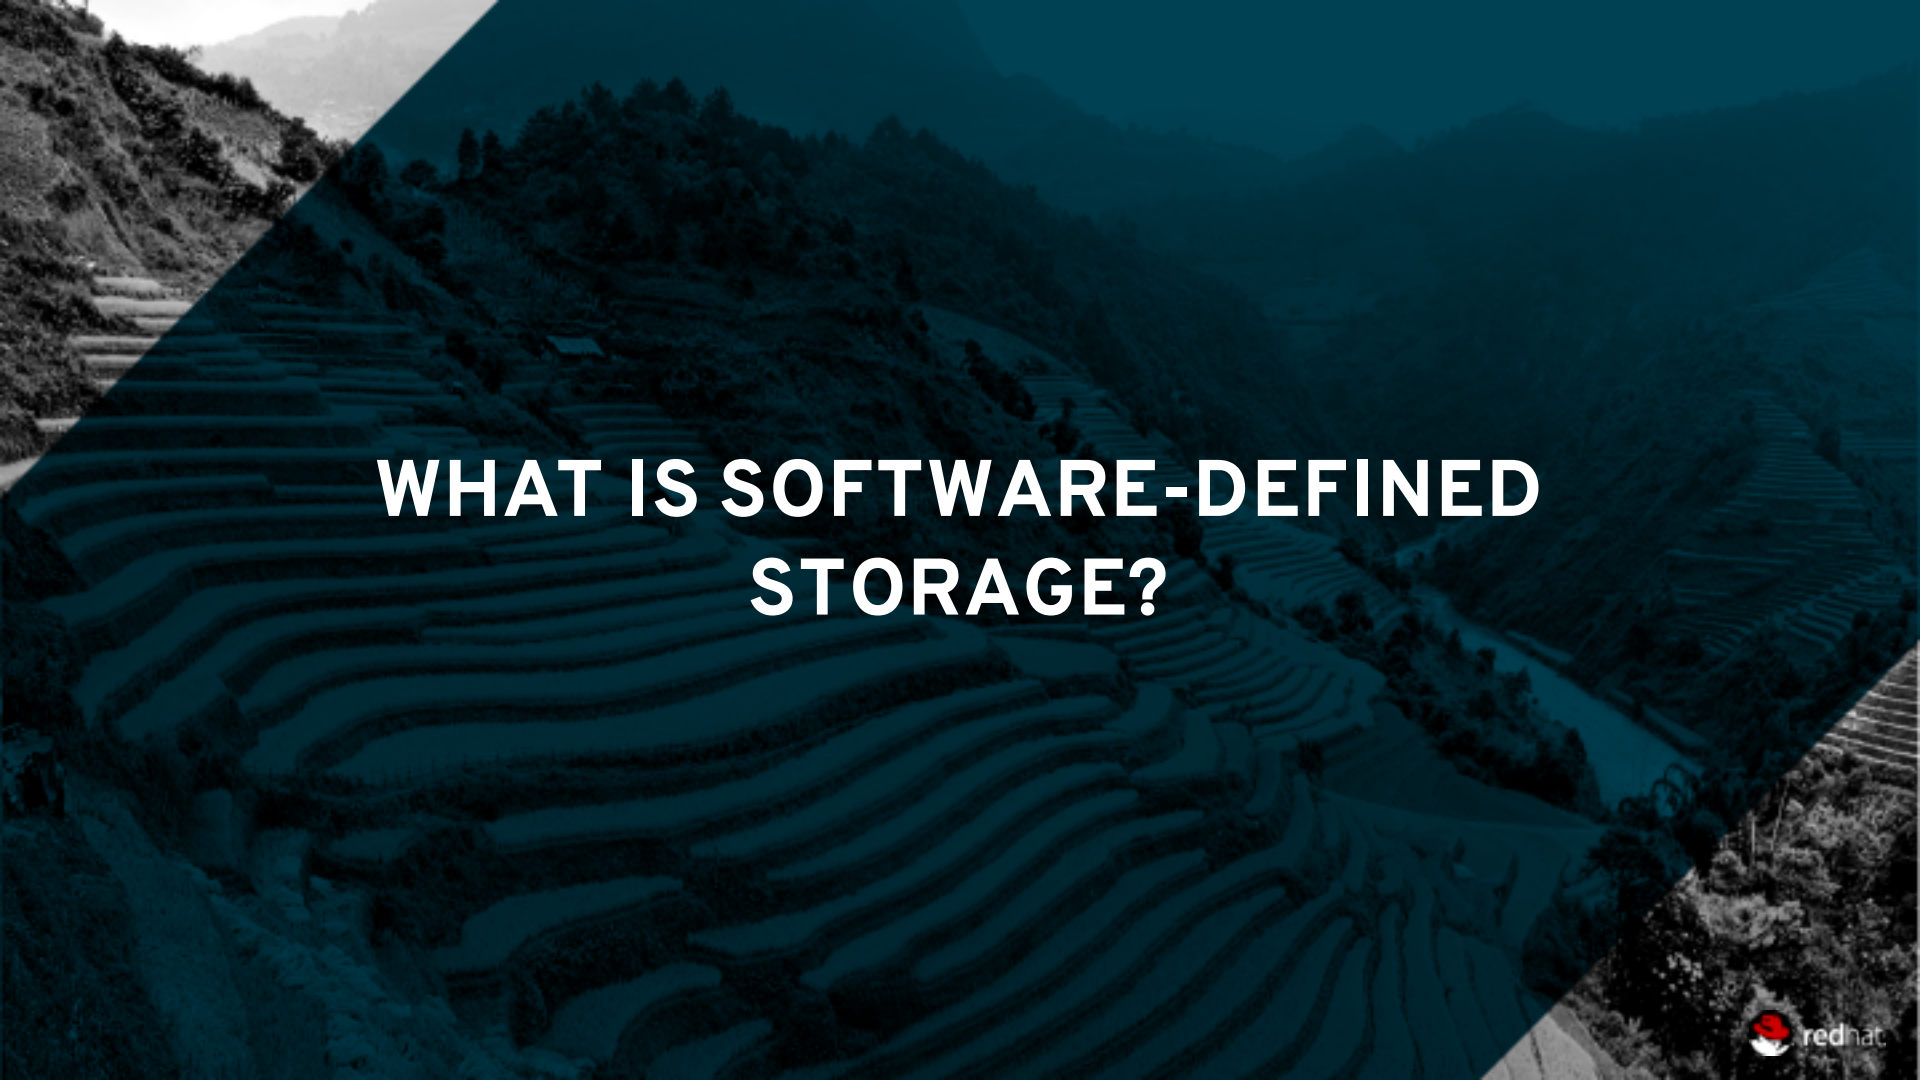
WHAT IS SOFTWARE-DEFINED STORAGE?
redhat

WHAT IS SOFTWARE-DEFINED STORAGE?

Server-Based Storage

Storage Orchestration

WHAT IS SOFTWARE-DEFINED STORAGE?

Server-Based Storage

Storage Orchestration

SERVER-BASED STORAGE

Server-based storage is the use of software and standard hardware to provide services traditionally provided by single-purpose storage systems.

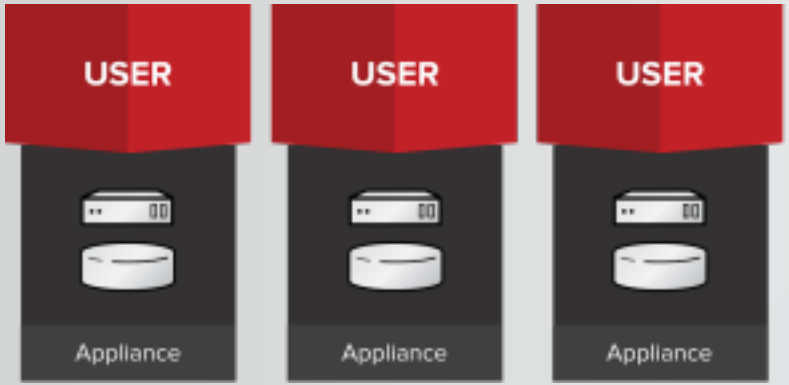

SERVER-BASED STORAGE

Server-based storage is the use of software and standard hardware to provide services traditionally provided by single-purpose storage systems.

VIRTUALIZED STORAGE SCALES BETTER

VIRTUALIZED STORAGE SCALES BETTER

SAN/NAS IS ON THE DECLINE

Changing workloads drive the need for flexible server-based storage.

- Storage in the enterprise has been growing at 40%+ per year.
- Share of storage deployed in servers grew 20%+ between 2010 and 2016.

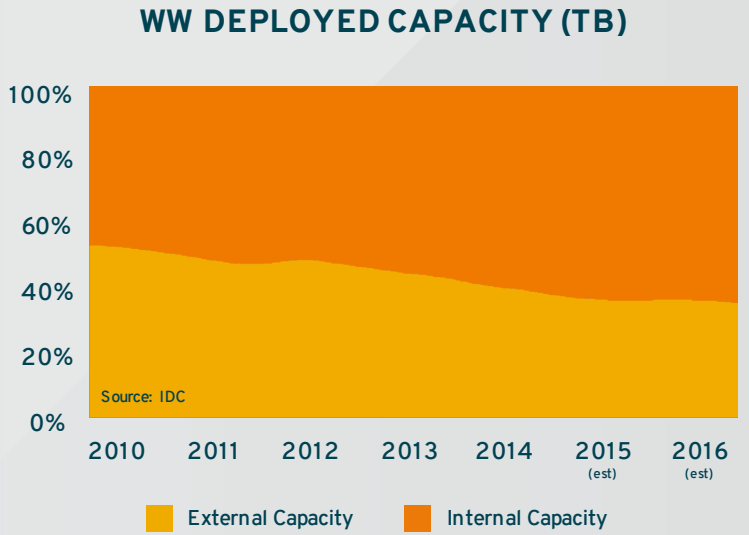

WHAT IS SOFTWARE-DEFINED STORAGE?

Server-Based Storage

Storage Orchestration

STORAGE ORCHESTRATION

Storage orchestration is the ability to provision, grow, shrink, and decommission storage resources on-demand and programmatically.

Web Console
A browser interface designed for managing distributed storage

API
A full API for automation and integration with outside systems

Command Line
A robust, scripable command-line interface for expert operators

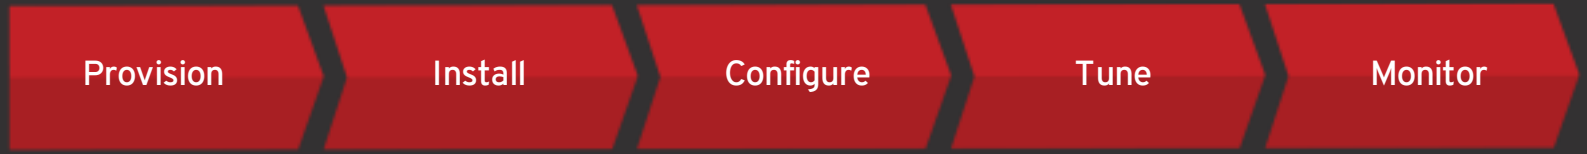

Full lifecycle management for distributed, software-defined data services

A RISING TIDE

Software-defined storage is leading a shift in the infrastructure industry.

- "By 2020, between 70%-80% of unstructured data will be held on lower-cost storage mangaged by SDS."
  Innovation Insight: Separating Hype From Hope for Software-Defined Storage

- "By 2019, 70% of existing storage array products will also be available as software only versions."
  Innovation Insight: Separating Hype From Hope for Software-Defined Storage

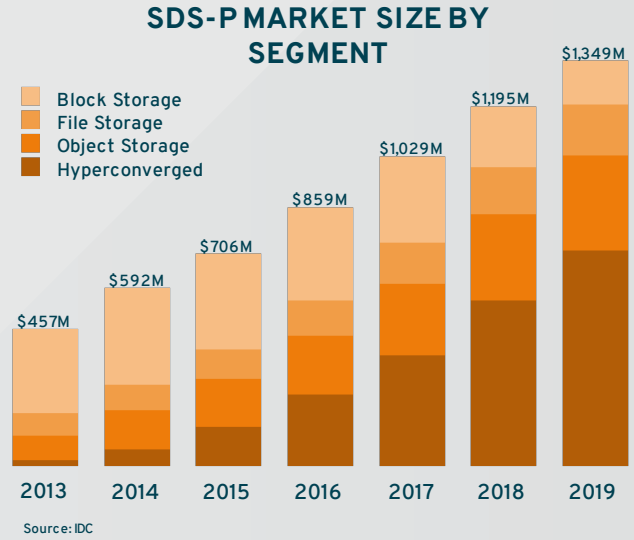


WHY DOES SOFTWARE-
DEFINED STORAGE MATTER?
redhat

FOUR IMPORTANT DIFFERENCES

PROPRIETARY HARDWARE	**Common, off-the-shelf hardware**	Lower cost, standardized supply chain

INDUSTRY STANDARD HARDWARE

Standardization makes storage more convenient

Users can build a cluster using standard hardware from existing vendors that's perfect for their workload.

- Clusters can be performance-optimized, capacity-optimized, or throughput-optimized.
- Need capacity? Add more disks. Too slow? Add more servers.
- Clusters can become larger or smaller with no downtime.

FOUR IMPORTANT DIFFERENCES

PROPRIETARY HARDWARE	Common, off-the-shelf hardware	Lower cost, standardized supply chain
SCALE-UP ARCHITECTURE	Scale-out architecture	Increased operational flexibility

SCALE-OUT PERFORMANCE

Performance should scale as capacity does

Software-defined storage intelligently uses hardware to provide performance at very large scale.

- Traditional appliances perform better when they are empty than they do when they are full of disks.
- Performance in software-defined storage clusters improves as clusters get larger, not the other way around.
- Software-defined storage systems can take advantage of flash technologies from a variety of vendors.

FOUR IMPORTANT DIFFERENCES

PROPRIETARY HARDWARE	Common, off-the-shelf hardware	Lower cost, standardized supply chain
SCALE-UP ARCHITECTURE	Scale-out architecture	Increased operational flexibility
HARDWARE-BASED INTELLIGENCE	Software-based intelligence	More programmability, agility, and control

THE ROBUSTNESS OF SOFTWARE

Software is more flexible than hardware

Software can do things hardware appliances can't. SDS brings the flexibility of software to the enterprise storage world.

- Can be deployed on bare metal, inside containers, inside VMs, or in the public cloud.
- Can deploy on a single server, or thousands, and can be upgraded and reconfigured on the fly.
- Grows and shrinks programmatically to meet changing demands

FOUR IMPORTANT DIFFERENCES

PROPRIETARY HARDWARE	Common, off-the-shelf hardware	Lower cost, standardized supply chain
SCALE-UP ARCHITECTURE	Scale-out architecture	Increased operational flexibility
HARDWARE-BASED INTELLIGENCE	Software-based intelligence	More programmability, agility, and control
CLOSED DEVELOPMENT PROCESS	Open development process	More flexible, well-integrated technology

THE RIGHT TOOL FOR THE JOB

Appliances

are suitable for small-scale, workloads, but they do not scale economically.

Durable, inflexible, convenient, expensive at large scale

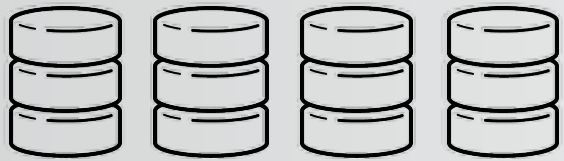

Software-defined storage

has a learning curve, but bring performance and economy at petabyte scale.

Durable, powerful, flexible, economical at large scale

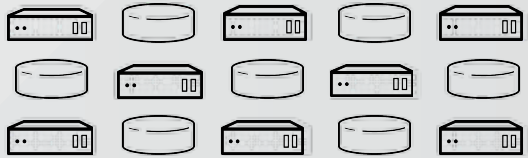

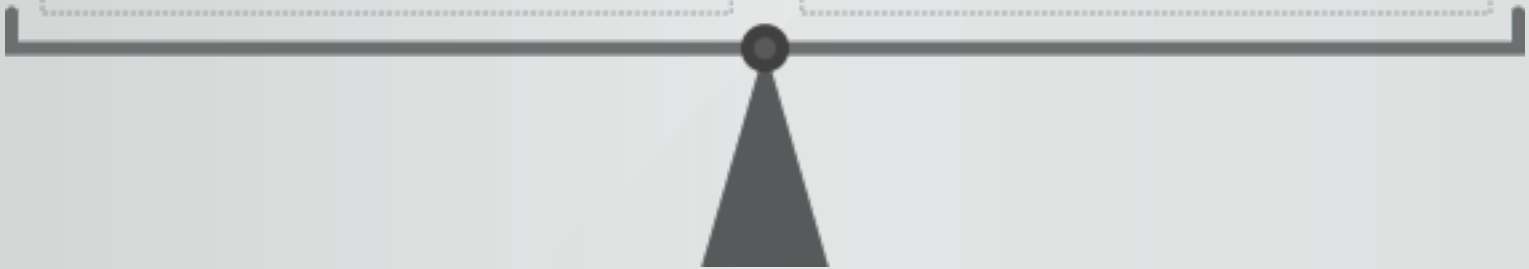


WHAT CAN IT BE USED FOR?

STRENGTH: MASSIVELY DISTRIBUTES DATA

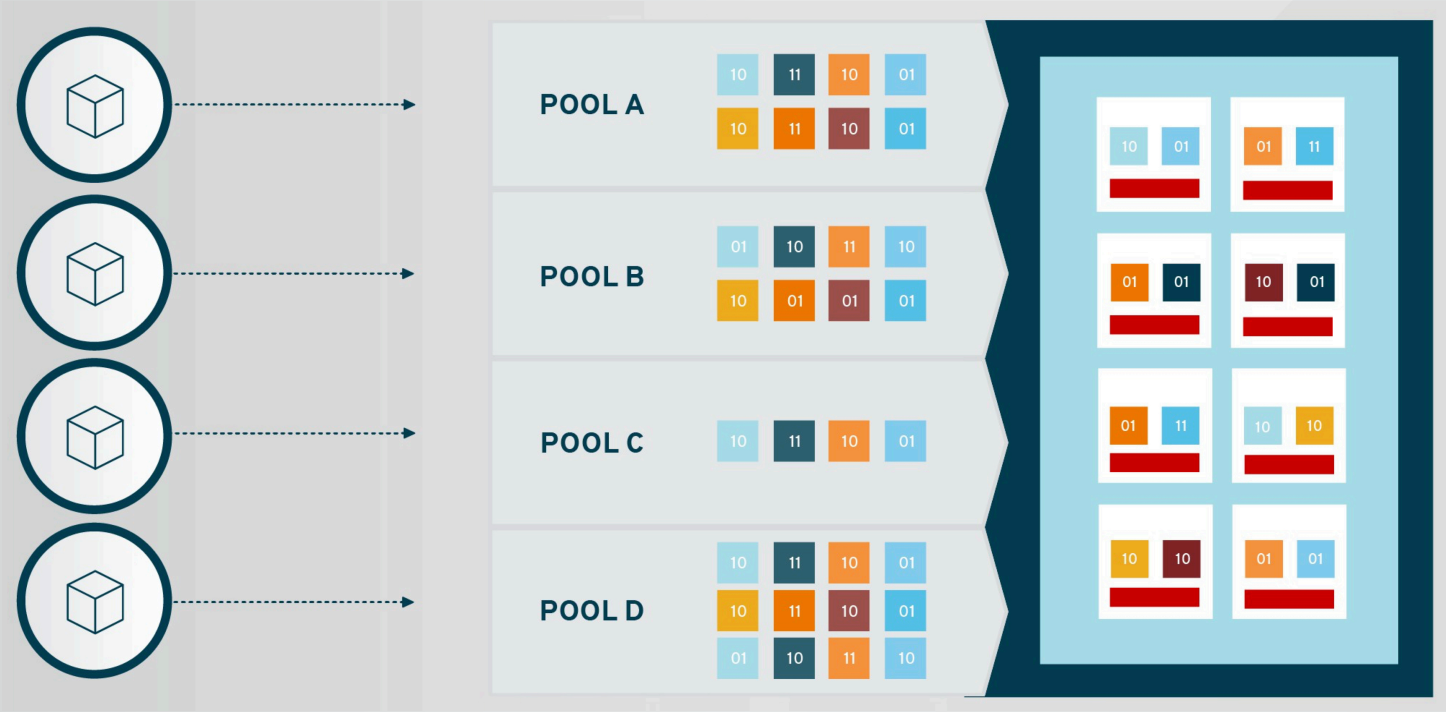

SOLUTION: OPENSTACK

- Allows for instantaneous parallel creation of VMs at massive scale
- Integrates easily and tightly with OpenStack Cinder, Glance, Nova, Keystone, and Manila
- Offers instant backup capabilities
- Provides persistent object, file, and database storage for applications

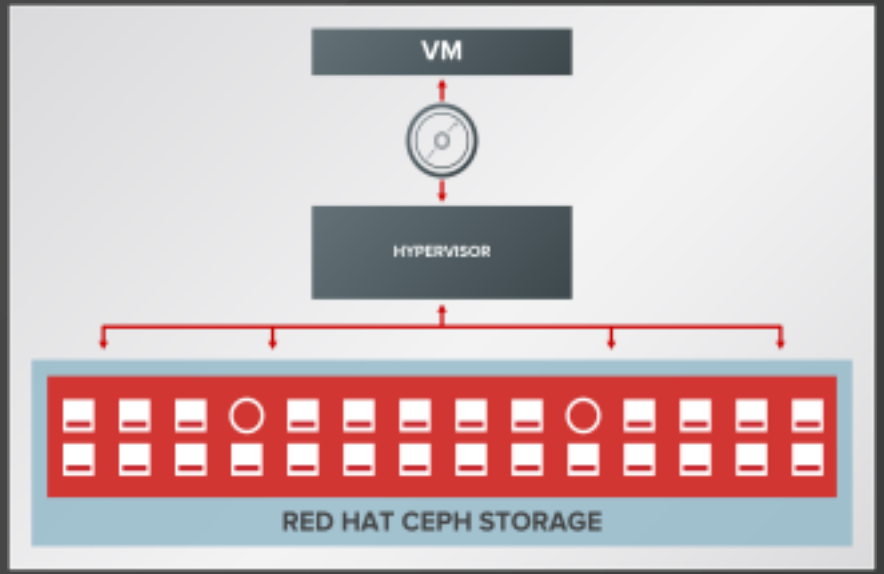

STRENGTH: GROWS TO PETABYTE SCALE

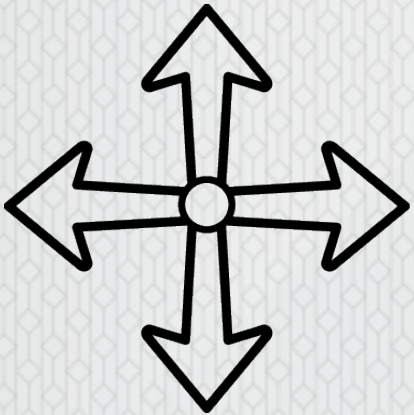

SOLUTION: OBJECT STORAGE

- Stores unstructured data at web scale, using standard hardware
- Works with industry-standard APIs for a wide range of application compatibility
- Spans multiple geographical regions with no single point of failure
- Matches the distributed architecture of software-defined storage

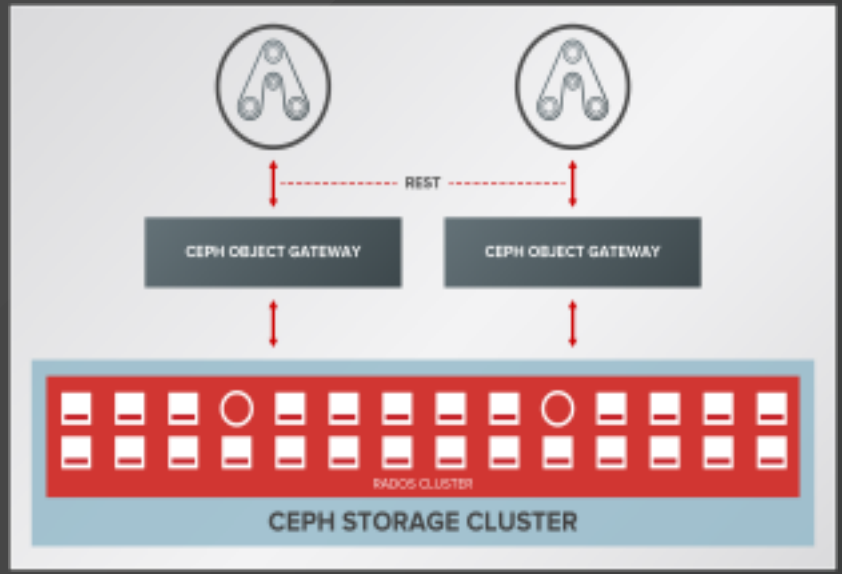

STRENGTH: WORKS IN LOTS OF PLACES

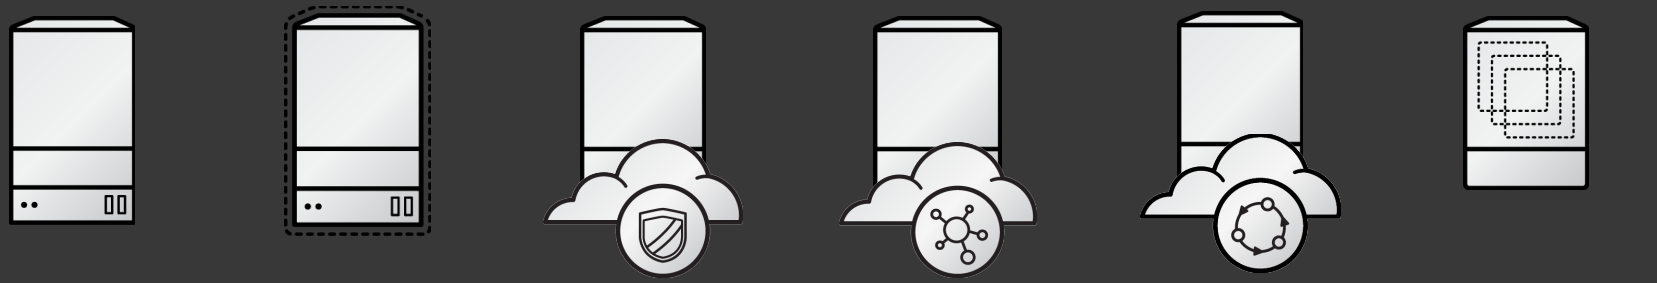

HOSTS

VIRTUAL MACHINES

PRIVATE CLOUD

PUBLIC CLOUD

HYBRID CLOUD

CONTAINERS

SOLUTION: CONTAINERS

- Offers persistent storage to applications running in containers
- Applications and storage can co-exist on the same hardware
- Allows for higher server utilization and lowers operational costs
- Storage generates only 3%–10% overhead on converged servers

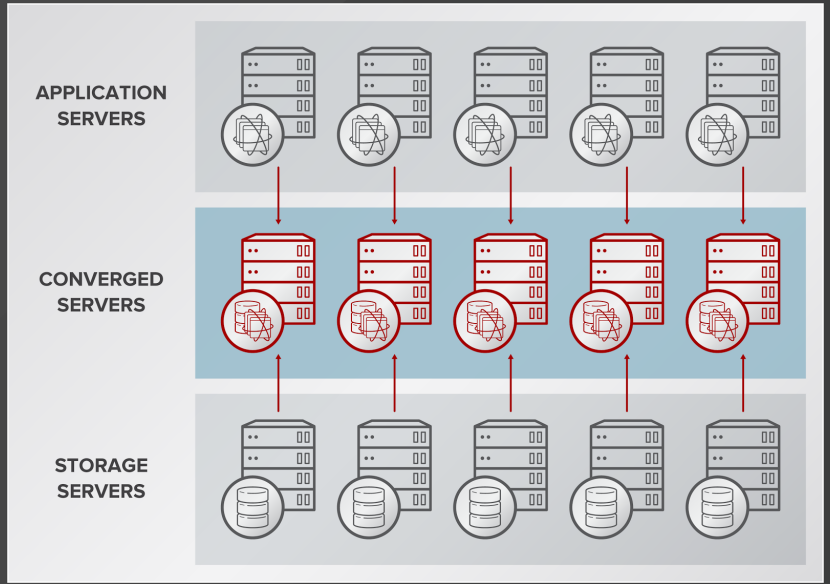


RED HAT®
STORAGE
redhat.

RED HAT STORAGE

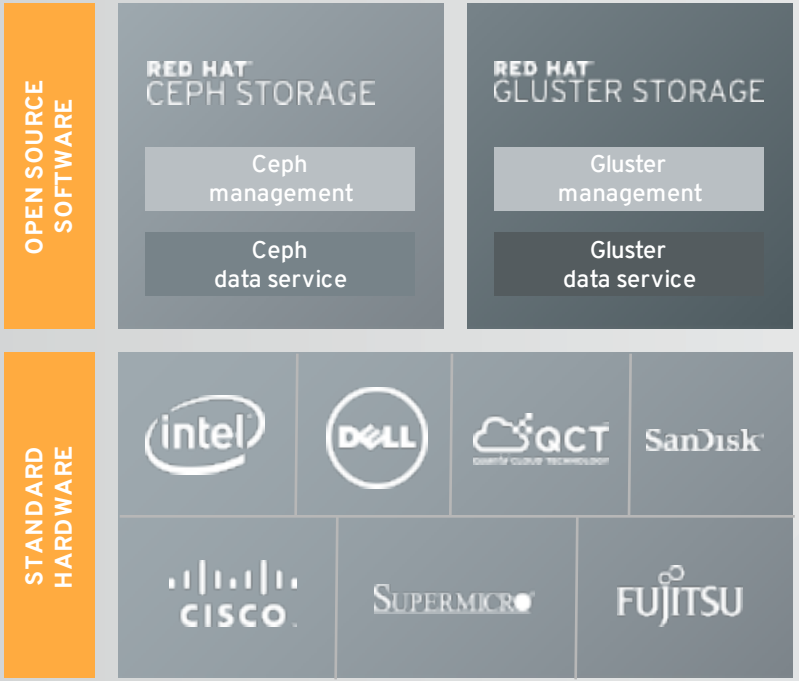

- Share-nothing, scale-out architecture provides durability and adapts to changing demands
- Self-managing and self-healing features reduce operational overhead
- Standards-based interfaces and full APIs ease integration with applications and systems
- Supported by the experts at Red Hat

RED HAT LOVES STORAGE

RED HAT STORAGE

PHYSICAL	VIRTUAL	PRIVATE CLOUD	CONTAINERS	PUBLIC CLOUD
RED HAT CEPH STORAGE RED HAT GLUSTER STORAGE	RED HAT CEPH STORAGE RED HAT GLUSTER STORAGE	RED HAT CEPH STORAGE RED HAT GLUSTER STORAGE	RED HAT CEPH STORAGE RED HAT GLUSTER STORAGE	RED HAT CEPH STORAGE RED HAT GLUSTER STORAGE
RED HAT ENTERPRISE LINUX	RED HAT ENTERPRISE LINUX RED HAT ENTERPRISE VIRTUALIZATION	RED HAT OPENSTACK PLATFORM	OPENSHIFT ENTERPRISE by Red Hat	RED HAT ENTERPRISE LINUX

GROWING INNOVATION COMMUNITIES

ceph

- Contributions from **Intel, SanDisk, CERN,** and Yahoo.
- Presenting Ceph Days in cities around the world and quarterly virtual Ceph Developer Summit events.

97 AUTHORS/MO

2453 COMMITS/MO

260 POSTERS/MO

GLUSTER COMMUNITY

- Over **11M downloads** in the last 12 months
- Increased development velocity, authorship, and discussion has resulted in rapid feature expansion.

33 AUTHORS/MO

97 COMMITS/MO

138 POSTERS/MO

TEST DRIVES

RED HAT®
CEPH STORAGE

bit.ly/cephtestdrive

RED HAT®
GLUSTER STORAGE

bit.ly/glustertestdrive