SDC 20

BY Developers FOR Developers

Storage Developer Conference
September 22-23, 2020

Track Overview: Cloud Storage

Mark Carlson
Kioxia
Co-chair SNIA Technical Council

Decentralized Platforms Push Edge Networks Closer to the Edge

The connection between edge computing and edge storage

Ben Golub
@golubbe

storj.io
@storjproject

HDD Annual Capacity Shipments

Exabytes Shipped

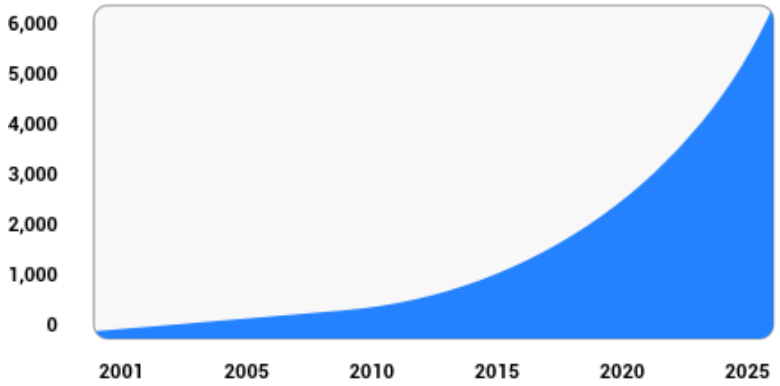

Storage Pricing

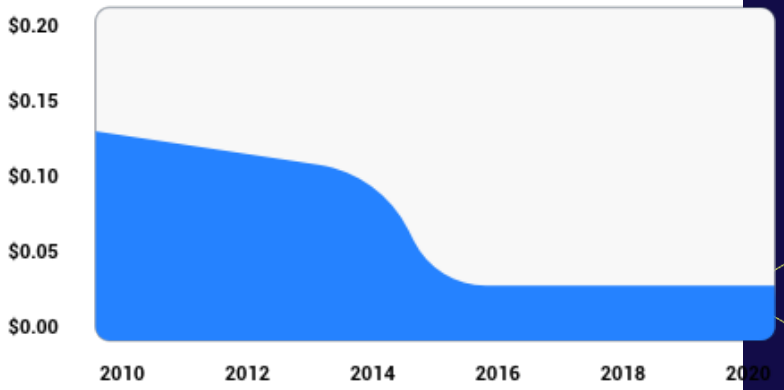

What is a decentralized application?

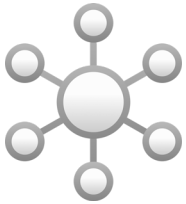

Centralized Systems

- Central Authority
- Single Point of Failure
- Opaque
- Security by People
- Trust Me

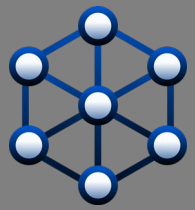

Decentralized Systems

- No Central Authority
- No Single Point of Failure
- Transparent - Open Source
- Security by Math
- “Trustless” (really, trust open code and large community)

SDC 20

BY Developers FOR Developers

Storage Developer Conference
September 22-23, 2020

Decoupling SDS Architectures for Agility

Arun Raghunath (Research Scientist)
Yi Zou (Research Scientist)

Intel Corp.

The Case for Decoupling SDS Architectures

Previously, we have illustrated,

- SDS decoupling via data and control plane separation
  - Ceph, SPDK, and NVMe-oF ecosystem
- Benefits of decoupling when storage is disaggregated
  - Removes extra hop (“datacenter tax”) -> reduce latency & bw cost [SDC2018 talk]
- Decoupled SDS architecture PoC [SDC2019 talk]
  - Integration of next-gen storage and mixed media types [SDC2020: Mortimer talk]
  - Scale-out of multiple heterogenous storage services
  - Separating of failure domains and benefits to recovery

Focus of this talk,

→ ***Agility must be considered in SDS architecture design!***

SDS = Software Defined Storage

A Novel Approach for SDS Agility

Clone OSD top-half

- Need: Place data close to where needed
- Does not contribute to object placement → not part of CRUSH
- Functions as a cache node → does not require durable backend
- Fundamentally different from tiering → not a cache tier pool

Introduce *clonemap* as part of the cluster metadata

- Clones are tracked as part a *clone set* for a given OSD in *clonemap*
- Separate from placement SLA related metadata (e.g., CRUSH)

Clone OSD bottom-half

- Need: I/O bottlenecks from write bursts
- Completely transparent to every entity in cluster except fronting OSD top half
- OSD top-half independently creates/removes clones based on load

Decoupling enables adding SDS clones without data rebalance
Does not cause any object migration across failure domains

SDC 20

BY Developers FOR Developers

Storage Developer Conference
September 22-23, 2020

Architecting Storage Applications for the Public Cloud Economy

Josh Salomon & Orit Wasserman
Red Hat

jsalomon@redhat.com

owasserm@redhat.com

Public Cloud Structure

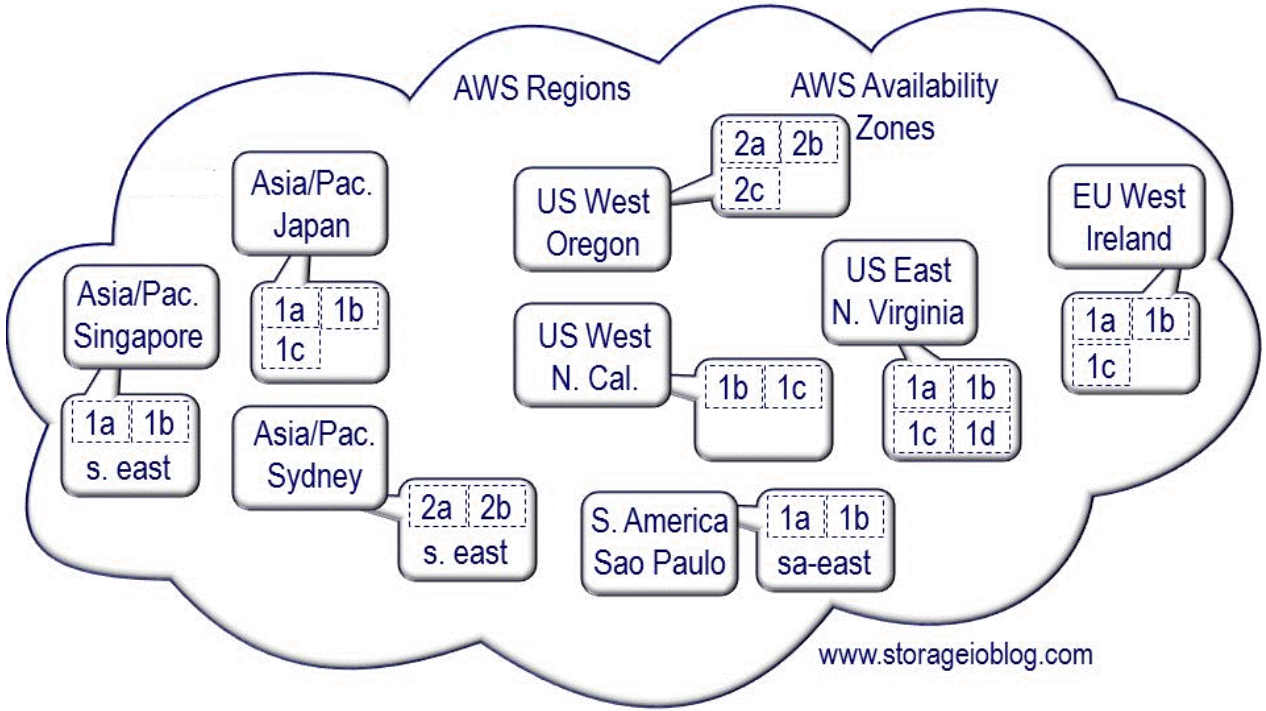

Examples use AWS terminology - https://storageio.com/images/SIO_AWS_Regions.gif

Summary - Key Points for Cloud Systems

- Elasticity
  - While this is the obvious it is also the key for successful cloud implementation.
- Multiple Deployment Options
  - Pricing model can change without notice, system should be flexible enough to adapt quickly
- Think OPEX
  - Need to optimize on more dimensions

Maximize your SDC 2020 Experience

SDC 20

- Participate in our online chat for this track at <TBD>
- Check out the Birds of a Feather (BoF) sessions
- Please be sure you rate each session you watch – you'll see a box under the video
- For additional details see the *Introduction to Virtual SDC* video (https://www.snia.org/SDCintro)
- ***Enjoy the SDC 2020 virtual event!***