



Lessons Learned from Distributed Systems and Applied to Modern Storage Platforms

Avinash Lakshman – CEO & Founder, Hedvig

June 3, 2015

Co-inventor Amazon Dynamo

Goals

- Key value store as foundation
- Solved availability problem.
- Designed to run across data centers
- Designed to run on off the shelf commodity hardware.

What needed improvement

- Design requiring client to resolve version conflicts
- Design didn't take into account operational issues
- Handling data corruption, disk failures etc.

Inventor Apache Cassandra

Goals

- Storage of choice for Facebook Inbox Search.
- Designed to ingest data at insane rates.
- Support range queries via an ordered partition

What needed improvement

- Better than Dynamo from an operational perspective
- Only supported Eventual Consistency

The convergence of data center trends in 2011

- Virtualization was becoming all pervasive
- Enterprises were adopting a true cloud mentality
- Commodity hardware was coming much further along, even then 10 years ago
- Flash and high-capacity HDDs were arriving
- 10/40 GigE networks were becoming ubiquitous

Why do all these trends matter?

Conclusion 1:

Storage had reached a point where it was a distributed systems problem to solve

Conclusion 2:

There had been no fundamental innovation in storage for 10 years

Conclusion 3:

The move to cloud and commoditization signaled enterprise readiness and skill set to adopt a new storage approach

7 Design goals for modern, scalable storage



- Scale-out seamlessly with x86 or ARM
- Consolidate all protocols in one platform
- Support hyper-converged and hyper-scale
- Run agnostic to any hypervisor, container, or OS
- Provide hybrid-aware DR, across any public cloud
- Provision storage with unprecedented flexibility and speed
- Make enterprise-grade storage features fundamental primitives

3

Reasons why a distributed systems is needed for massively scalable storage

Reason 1: Data protection

Data protection in massive virtualized/cloud environments sucks

Value of distributed system

- Self-healing clusters
- Multi-site replication
- No more forklift upgrades or data migration

Example

- Fortune 50 bank
- Hyperscale configuration for VMware environment
- Complex 4-city, 2-region DR failover schema

Reason 2: Data locality

Data locality for big data
(and Hadoop in particular) is a myth

Value of distributed system

- Single data platform to handle multiple distros
- Gain data efficiency across multiple data sets
- Ability to virtualize big data

Example

- Same Fortune 50 bank
- Bare metal Linux environment with HDFS
- Virtualize the various flavors of NoSQL solutions.

Reason 3: Backup

Backup is dead . . .
at least for massive data sets

Value of distributed system

- Eliminates need for separate backup infrastructure
- Hybrid cluster adds cloud tier
- Incremental scale & no more growing maintenance costs

Example

- \$10B mining company
- Backup costs untenable
- Commodity disk-to-cluster better than traditional backup infrastructure

Thank you!

www.hedviginc.com