



# Pushing Ceph in LATAM market

October 16, 2019

by Álvaro Soto // @alsotoes

CEPH DAY ARGENTINA 2019



# .whoami

---

- Based in Mexico City
- Ceph since Firefly (2015)
- Stacker since Havana (2013)
- Public and private sector
  - Financial
  - Government
  - Health-care
  - E-commerce
  - \*\*\*

## Disclaimer

---

I'm not a storage guy

I'm not a sales guy

I'm an engineer


This talk is about...

---

Lots and lots of comparison.

My personal experience.

# Where are the Ceph users?



# Ceph main features

---



## Philosophy

- **Open Source**
- **Community focused**
- **No single point of failure**

## Design


- **Scalable**
- **Software base**
- **Self managing / healing**

# Ceph main features


---



## Philosophy

- Open Source 
- Community focused
- No single point of failure

## Design

- Scalable
- Software base 
- Self managing / healing

# Open Source

---



≠





# From who do I need help? (Ceph)

---

- Rack & stack
- Sysadmin
- Networking
- Network security
- Monitoring




# Some decisions to make (Ceph)

---

- Should the replicated node be on the same rack or multiple racks to avoid SPOF ?
- Should the OSD traffic stay within the rack or span across rack in a dedicated or shared network ?
- How many nodes failure can be tolerated ?
- If the nodes are separated out across multiple racks network traffic increases and the impact of latency and the number of network switch hops should be considered.
- Ceph will automatically recover by re-replicating data from the failed nodes using secondary copies present on other nodes in cluster . A node failure thus have several effects.
  - Total cluster capacity is reduced by some fractions.
  - Total cluster throughput is reduced by some fractions.
  - The cluster enters a write heavy recovery processes.

# Some decisions to make (Ceph)

- Should the replicated node be on the same host as the primary node? SPOF ?
- Should the OSD traffic stay within the same network or shared network ?
- How many nodes failure can be tolerated?
- If the nodes are separated out across multiple hosts, what are the consequences and the impact of latency and network topology?
- If the nodes are separated out across multiple hosts, what is the number of network switch hops?
- Ceph will automatically recover by using secondary copies present on other nodes in cluster . A node failure
  - Total cluster capacity is reduced by the size of the failed node.
  - Total cluster throughput is reduced.
  - The cluster enters a write heavy recovery processes.



# Any other storage solution

- Defined use cases.
- You will always need support.
- Storage admins.
- Certifications.
- Short communities.
- Mostly no free support.



# Software based

---





# Software based


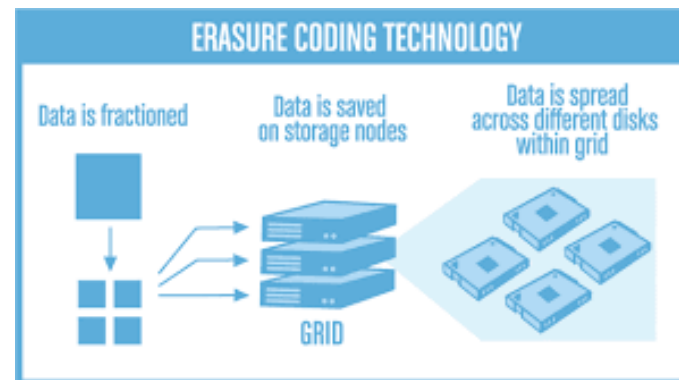



"Well, let's get started now we've got a quorum."

# 3X



Data Age	Usage Frequency
Age < 7 days	20x / day
7 days < Age < 1 month	5x / week
1 mo. < Age < 3 month	5x / month
3 mo. < Age < 3 years	2x / year



# Software based

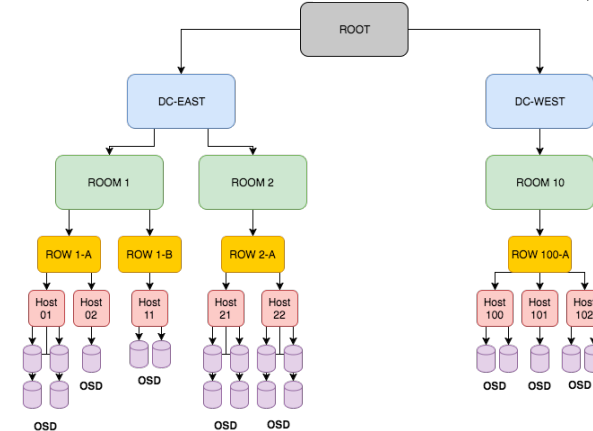
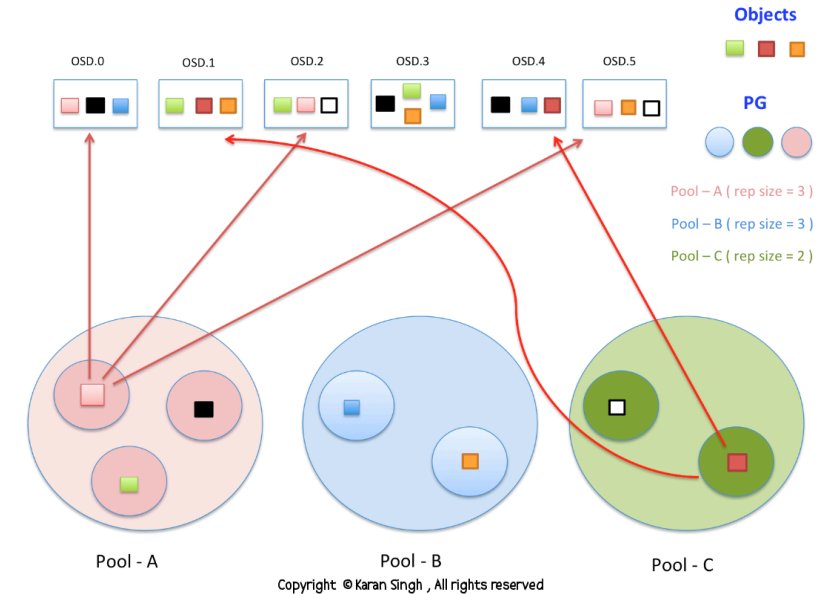
---



Tuning!!! Tuning!!! Tuning!!! Tuning!!! Tuning!!!

# Software based

- Business Requirement
  - Budget ?
  - Do you need Ceph cluster for day to day operation or SPECIAL
- Technical Requirement
  - What applications will be running on your ceph cluster ?
  - What type of data will be stored on your ceph cluster ?
  - Should the ceph cluster be optimized for capacity and performance ?
  - What should be usable storage capacity ?
  - What is expected growth rate ?
  - How many IOPS should the cluster support ?
  - How much throughput should the cluster support
  - How much data replication ( reliability level ) you need ?





# Questions?

---



Thanks!!!

---

**.me**

—@alsotoes



—khyr0n

**#irc**

—<https://headup.ws>



—<https://latam.openstackday.mx/>



ceph