



**CEPHALOCON APAC 2018**  
THE FUTURE OF STORAGE  
22-23 March 2018 | BEIJING

# **Petabyte Scale Object Storage Service using Ceph in a Private Cloud**

Flipkart Internet Pvt Ltd, Bangalore, India





ceph

# Agenda

- About Flipkart
- Flipkart Cloud Platform(FCP)
  - Virtualized Infrastructure
  - Object Store Services
- FCP Storage Team - Tech Stack
- Ceph Configuration and Deployment
- QoS(Server Side, RPS based)
- Cluster Administration and Health Monitoring



# About Flipkart


- Flipkart is India's leading e-commerce marketplace offering
  - Over 80 million products across 70+ product categories
  - 100 million registered users
  - 100,000 Sellers
  - 8 million shipments/month
  - 13 million daily visitors
- Developed a highly scalable and reliable storage infrastructure
- Object storage service backed by Ceph for ~3 years
- Runs two data centers having more than 20,000 hosts
- Completely virtualized infrastructure with over 60,000 VMs



ceph

# Object Storage as a Service

- Scale-out object storage service built on CEPH for Flipkart
- Simple HTTP API - A subset of AWS S3 API
- Multiple clusters behind a single endpoint
- Key features
  - Focus on durability
  - Elastic for performance and capacity
  - Multi-tenant with user SLAs
  - Five 9's of reliability





ceph

# FCP Storage - SLAs Offered

## Shared-Active Workloads

- Latency sensitive e.g. Debian packages, product image uploads & near-term invoices.
- Optimized for small to medium sized objects

## Backup Workloads

- Bandwidth optimized e.g. MySQL DB backup
- Optimized for very large objects.
- Best suited for taking backups with limited retention period.

## Archival Workloads

- Long-term (multi-year) retention optimized e.g. legal compliance data
- Moderate throughput and high latency
- Optimized for medium to very large objects
- Explicit verification targets for durability






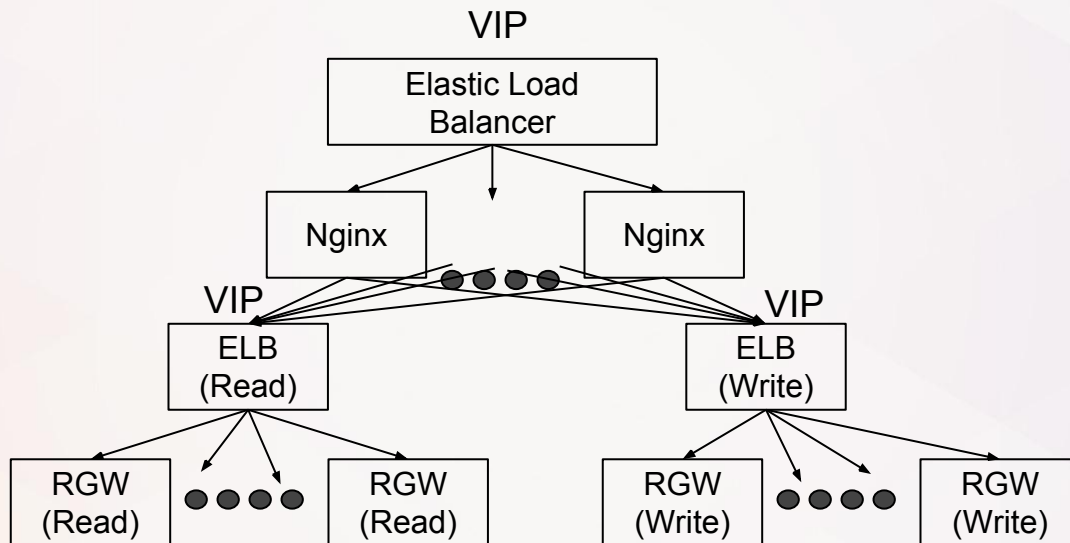
# Shared Active Service

- Stores ~1.5 billion objects
- ~750 OSDs
- Provides ~1 GB/s bandwidth utilization
- Handles ~ 500 RGW requests/sec
- ~2PB of raw storage
- Hybrid deployment with SSDs and HDDs
- 200+ internal customers
- Growing ~50 million objects/month

# Shared Active Tech Stack

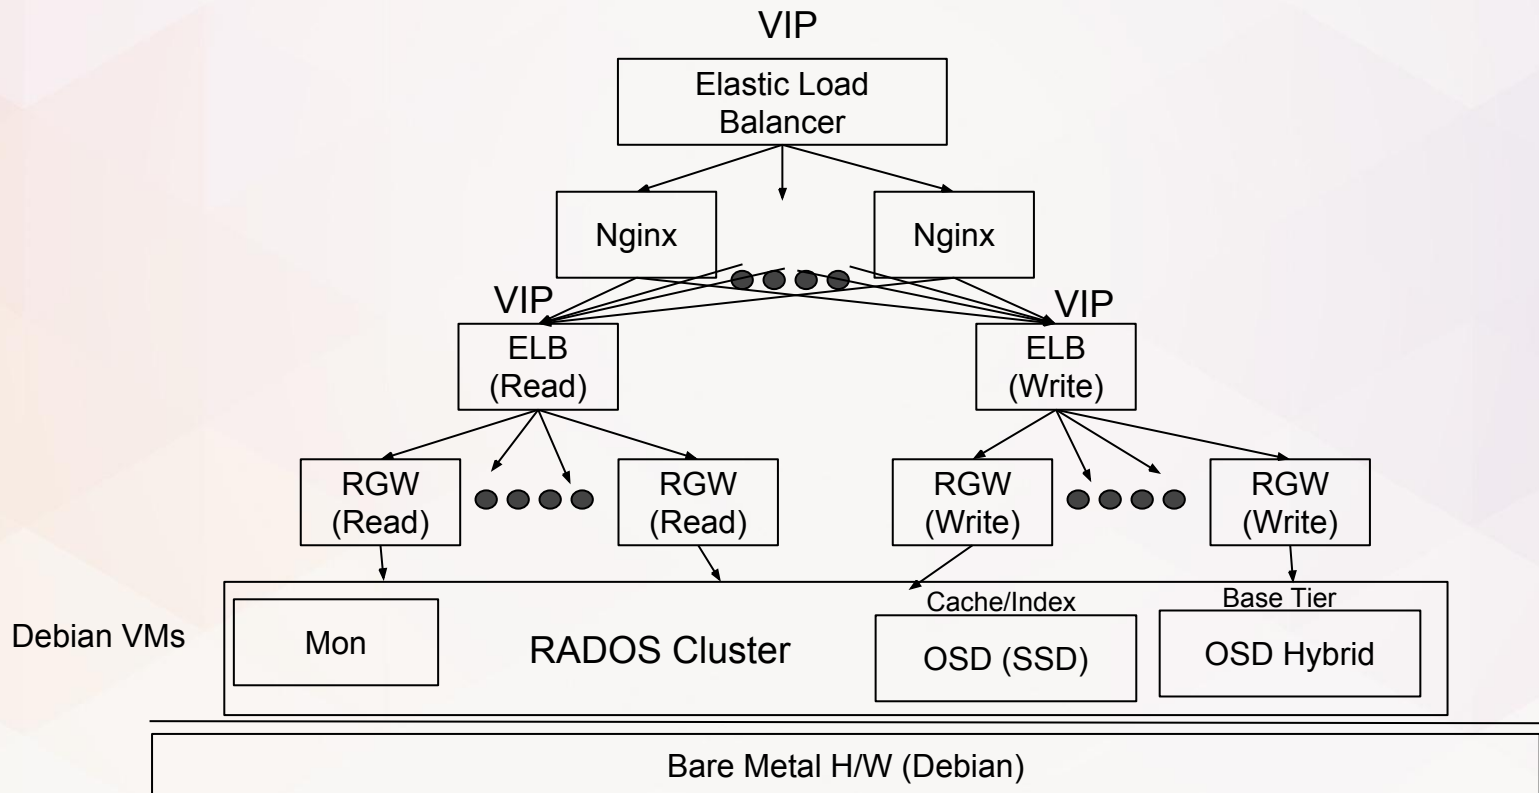


# Shared Active Tech Stack

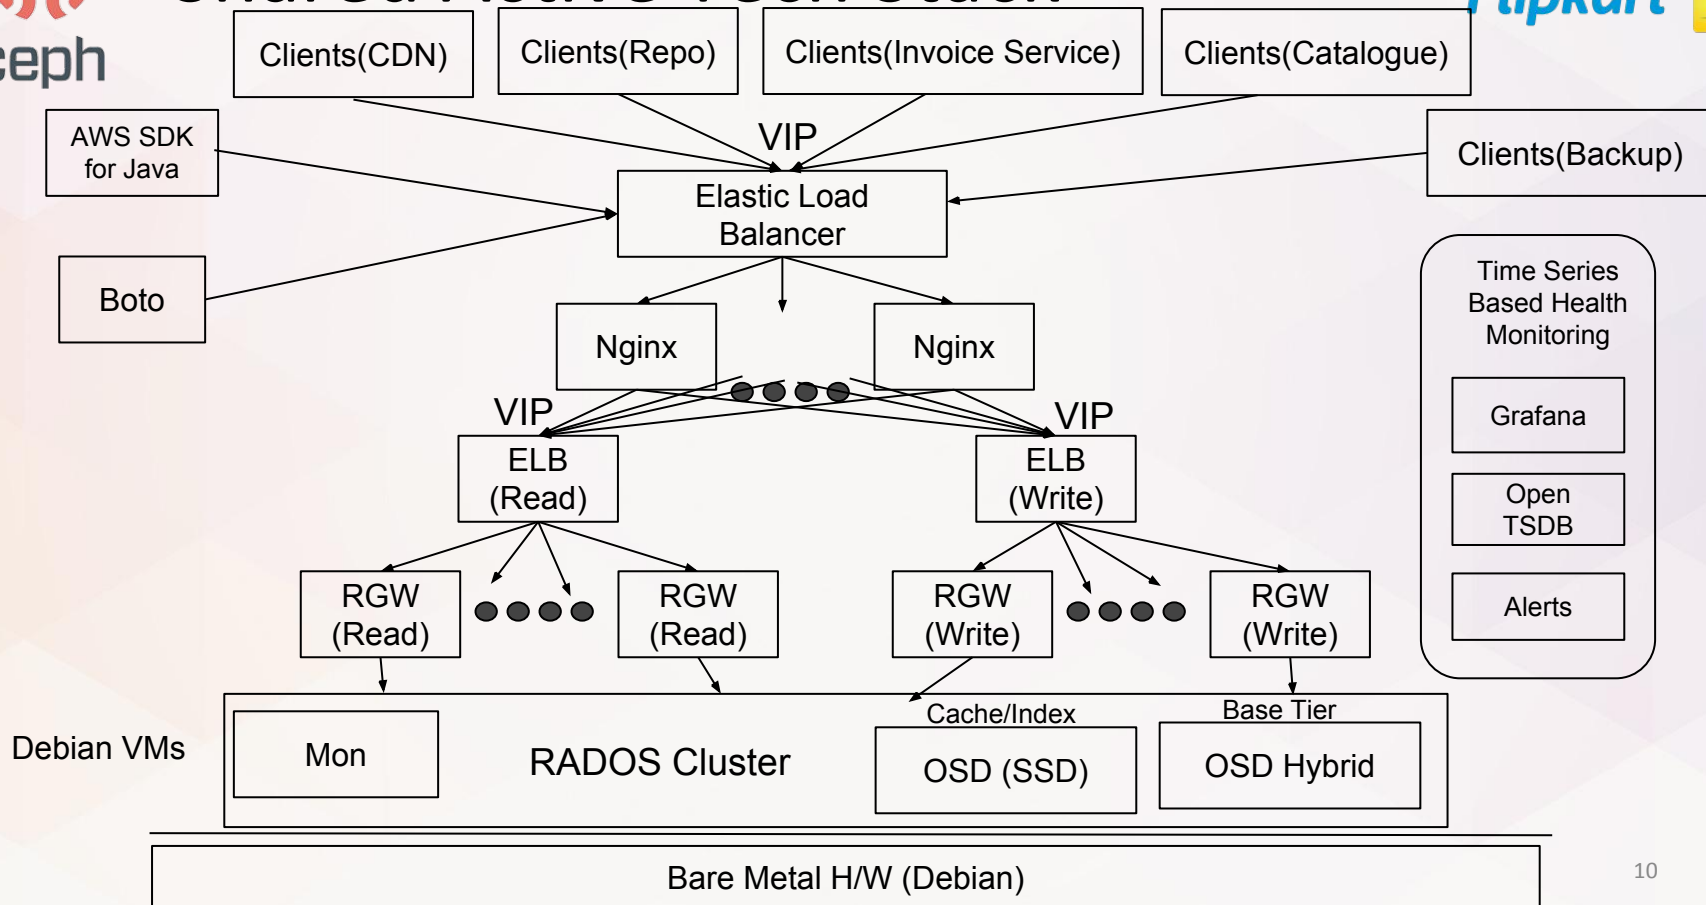




# Shared Active Tech Stack



# Shared Active Tech Stack





ceph

# Ceph Deployment: HW Instance

- HW VM instances for Ceph components
  - Monitors
    - Higher CPU cores and RAM
  - RGWs
    - Higher Bandwidth and CPUs
    - An instance group of N RGWs
  - Base Tier OSDs
    - Higher Capacity, Hybrid(SSD+HDD)
  - Cache Tier OSDs(SSDs)
  - Index OSDs (SSDs+More RAM)
    - Low latency, large capacity SSDs

- An instance has the following:
  - Host OS (Debian)
  - Network
  - CPU (Hyper Thread cores)
  - Disks (Direct attached)
  - RAM



ceph

# QoS



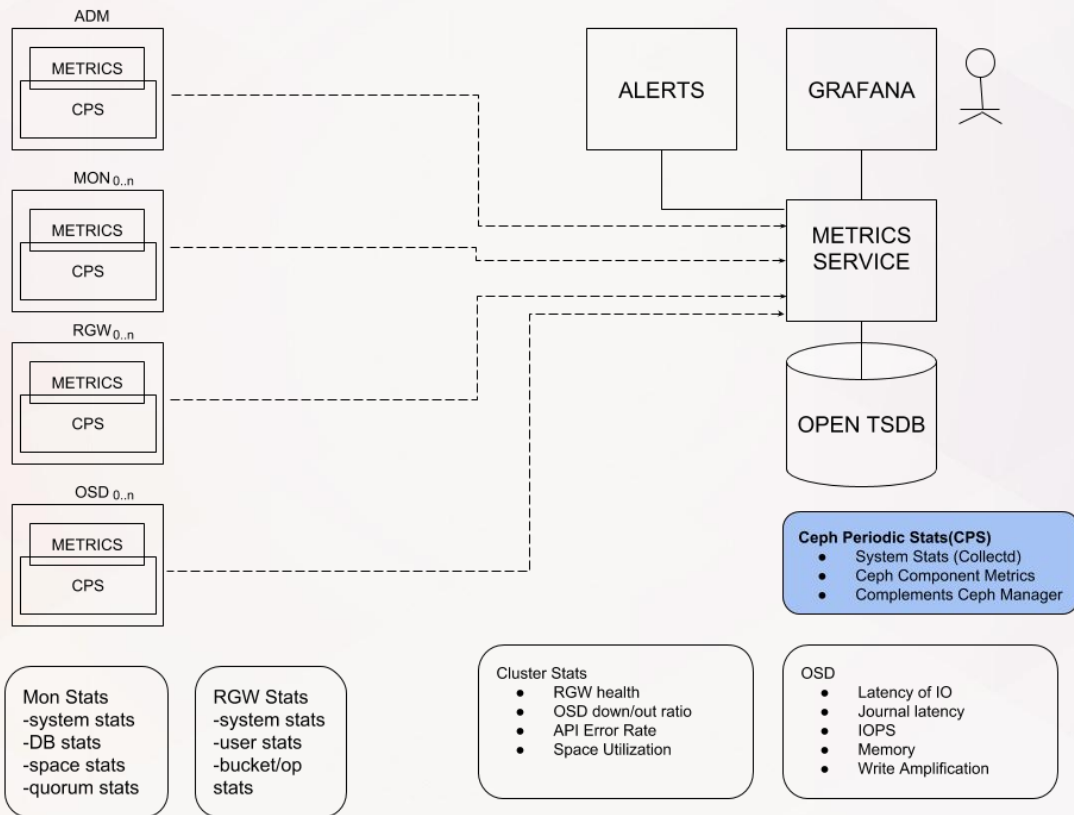
**Elastic Load Balancer**  
Spread load across RGW/Nginx

- Nginx Rd/Wr Separation**
- Independent sets of RGW instances
  - Handle GET and PUT in isolation
  - Maintained in Nginx conf


QoS

- RGW Rate Limit**
- Per user quota policy
  - Requests are rate limited per GET, PUT, CREATE, DELETE
  - Deployed per RGW instance


**CRUSH**  
Spread load across physically separated tiers across different PDU



# SA Service: QoS(Grafana)



# SA Service: QoS(Grafana)



# Challenges faced so far...




Issue	Timeouts due to lack of iops
Reason	<ol style="list-style-type: none"><li>1. Increased client and rebalancing load</li><li>2. Split and merges in filestore at the same time</li><li>3. Unable to serve I/O requests</li></ol>
Mitigation	<ol style="list-style-type: none"><li>1. Migrate journals to SSDs</li><li>2. Introduce SSD writeback cache tier in front of HDD OSDs to absorb iops</li><li>3. Introduce rate limits for clients</li><li>4. Two set of RGWs to serve read and write operations</li></ol>





# Challenges faced so far...




Issue	Service outage due to rebalancing in large metadata pools (index, log etc)
Reason	<ol style="list-style-type: none"><li>1. Buckets with more than 100 million objects</li><li>2. RGW usage pool had high number of entries and was enabled</li><li>3. Unable to write when rebalancing happens on large metadata pools</li></ol>
Mitigation	<ol style="list-style-type: none"><li>1. Manual sharding of large bucket indexes</li><li>2. Disable RGW log pool</li><li>3. Introduce bucket quota</li></ol>





# Challenges faced so far...

Issue	Write timeouts due to high and concurrent bucket create calls
Reason	Contention in Leveldb to due high number of concurrent updates to a bucket's metadata
Mitigation	Enable quota on bucket creation calls





# Challenges faced so far...

Issue	Data leak
Reason	Leaks with incomplete multipart uploads
Mitigation	<ol style="list-style-type: none"><li>1. Fixed : <a href="https://github.com/ceph/ceph/pull/15630">https://github.com/ceph/ceph/pull/15630</a></li><li>2. radosgw-admin orphan find did not scale with millions of objects</li><li>3. Developed an offline tool to delete leaked objects</li></ol>



# Challenges faced so far...

Issue	Data loss
Reason	Race in RGWCompleteMultipart
Mitigation	<ol style="list-style-type: none"><li>1. Fixed : <a href="https://github.com/ceph/ceph/pull/16732">https://github.com/ceph/ceph/pull/16732</a></li><li>2. Accounted loss by scanning all objects for missing parts.</li></ol>



# Challenges that still persist...

- Upgrades are still a pain
  - Upgrade from Hammer to Jewel, issue with *chown*.
  - Upgrading to Luminous:
    - Filestore has split and merges problem.
    - BlueStore solves them but upgrade path is not smooth
    - Adding new osds with BlueStore backend.
- Occasional PG incomplete issue and resulting data loss to bring PGs back online.
  - <http://tracker.ceph.com/issues/17002>
- End-to-End QOS



Questions?



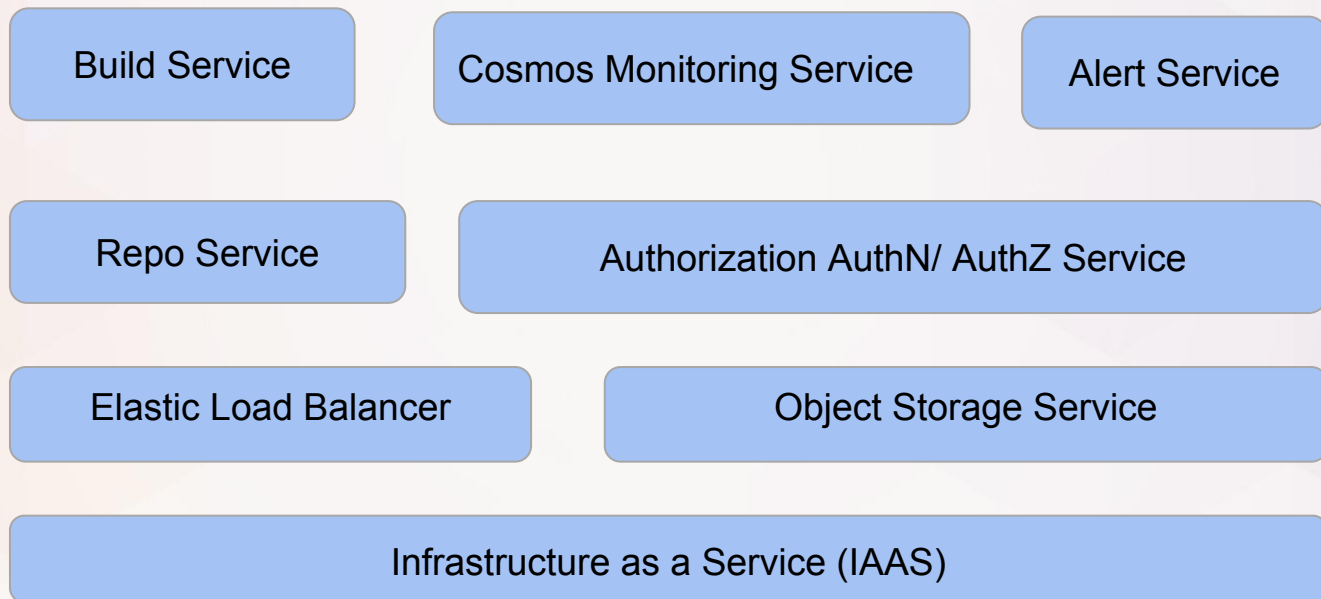
Thank You



# Backup



# Flipkart Cloud Platform: Building Blocks





# Object Storage As a Service

- Policy based services
  - Shared Active
    - Low latency, replicated, high throughput
    - Hybrid Deployment of HDDs and SSDs
  - Data Backup and Disaster Recovery Service
    - Replicated on HDDs
  - Data Archival Service
    - Erasure Coded on HDDs
- *Multi-tenant*
- *Five 9's reliable*
- *200+ internal customers*



# Cluster Administration and Monitoring

- Nagios alerts
- Ceph Periodic stats
  - Has more than 100 metrics
  - Integrates to Cosmos service
  - Stores data on Cosmos's OpenTS DB
  - Complements Ceph manager
- Grafana
  - Visual representation of Ceph Periodic Stats
  - Historic and near real-time status of the service

# Use Cases

