10 Years of Ceph at Walmart

Meet the Speakers

Anton Thaker

PRINCIPAL SOFTWARE ENGINEER

Email: ceph@gribok.net

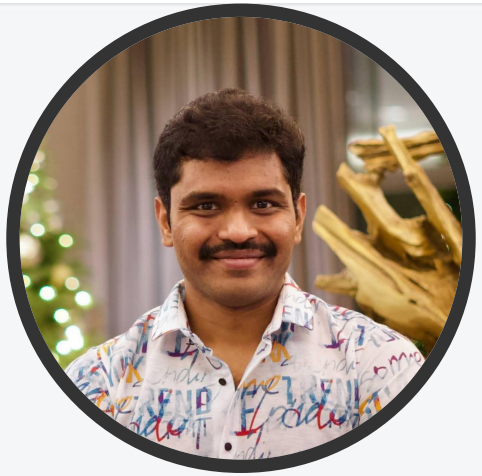

Bharath Krishna Maddi

STAFF SOFTWARE ENGINEER

Email: m.bharathkrishna@gmail.com

Prajwal Kabbinale

SENIOR SOFTWARE ENGINEER

Email: prajwalkabbinale95@gmail.com

Hardware and Scale

Hardware Evolution

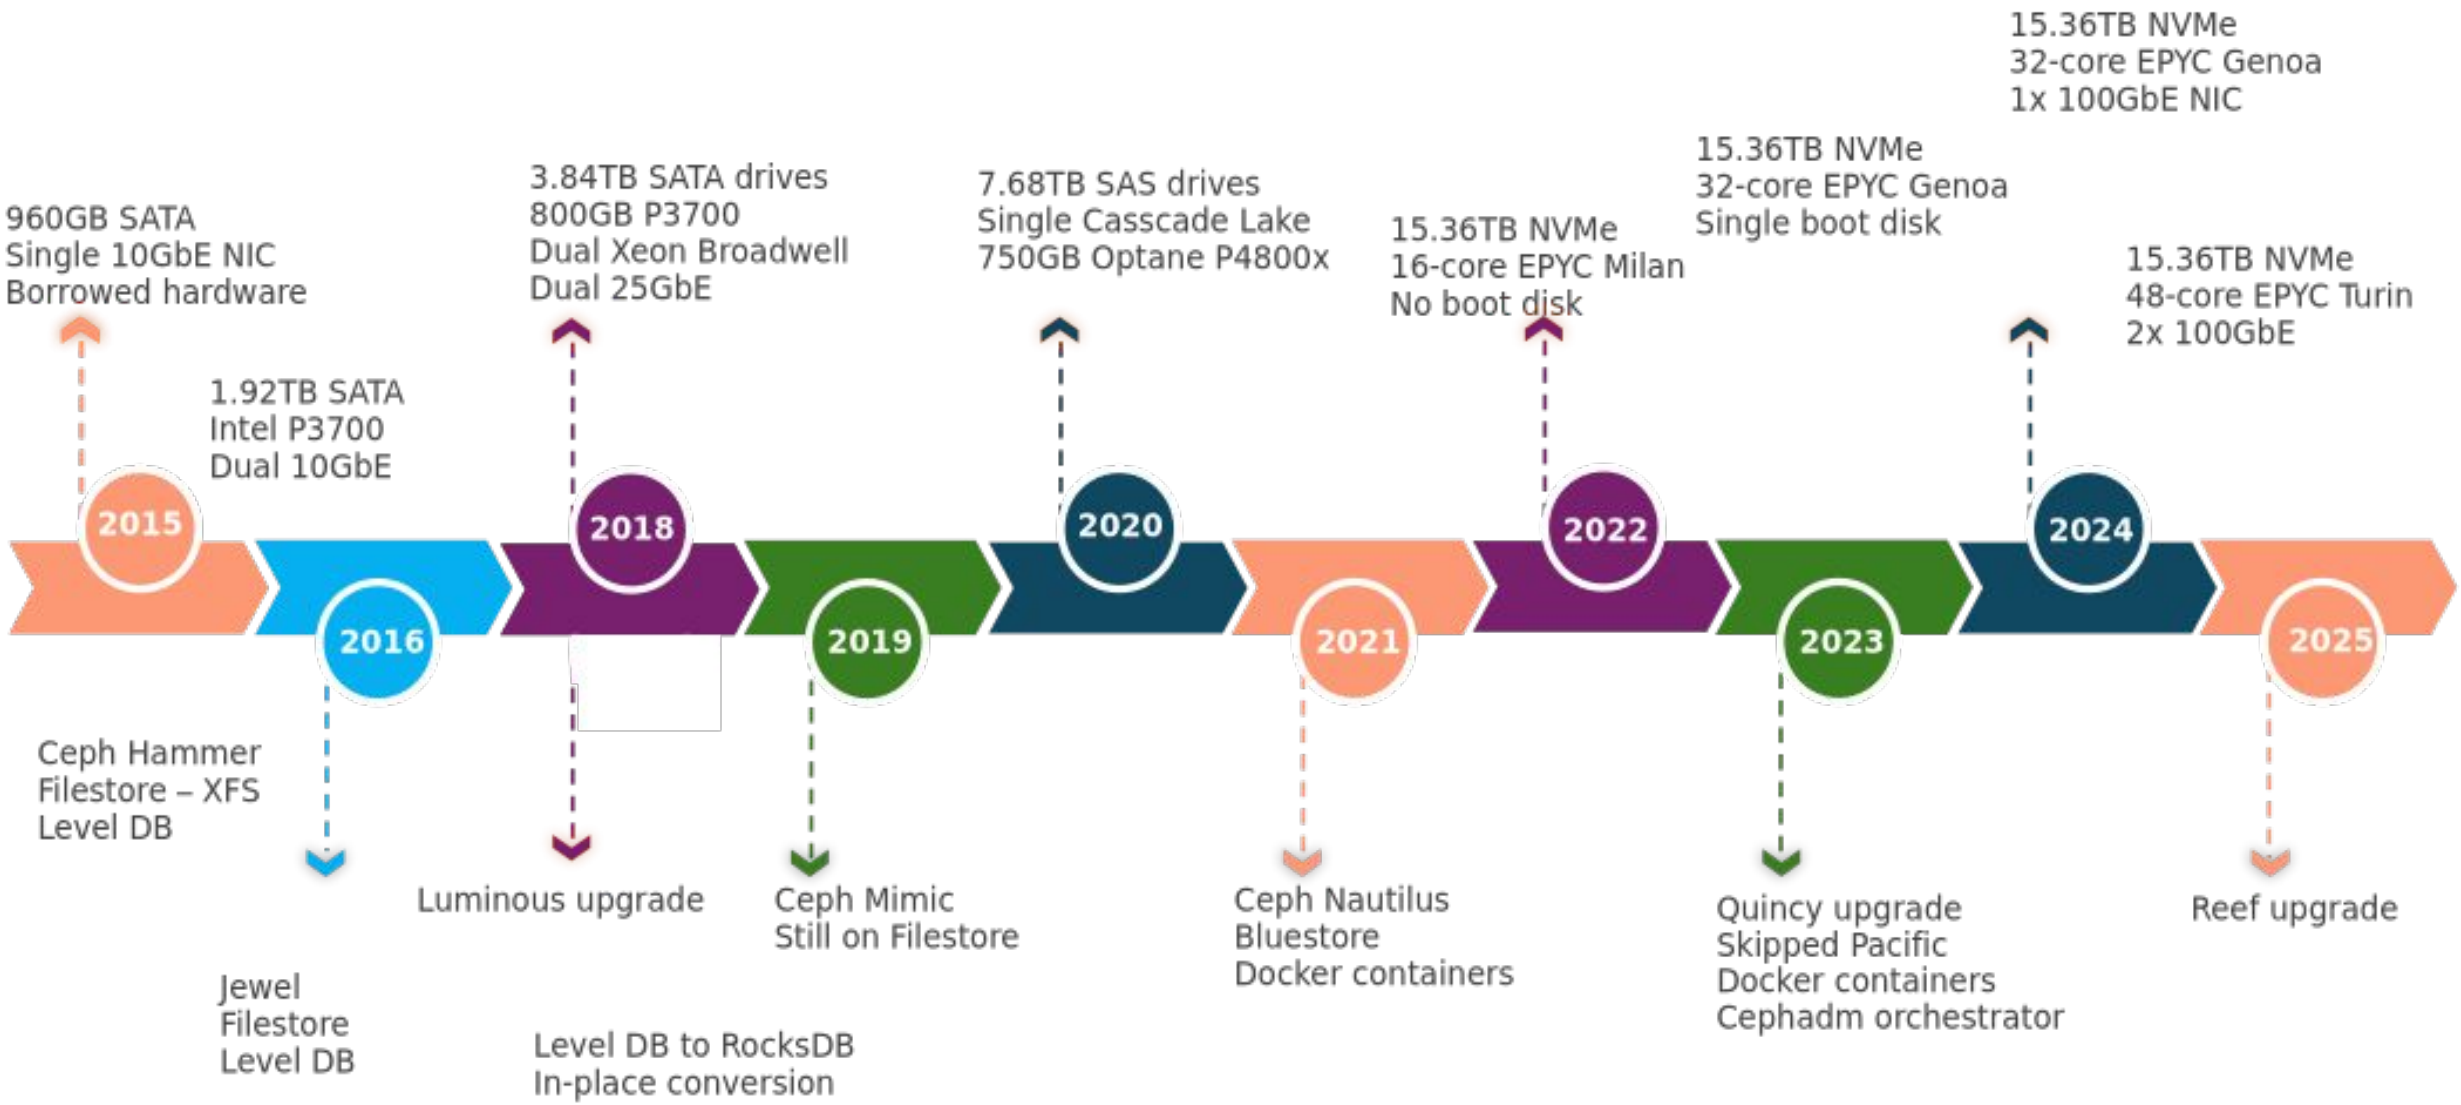

End-to-End Latency Observability

- OSD op & commit latency metrics focus only on server-side storage health
- Missing crucial latency events like peering & network blips

Limitations of Native Monitoring

- Synthetic 'Canary VMs'
- Constant R/W operations
- Reveals the true client experience

Canary VM Monitoring

- Ceph + Node Exporter
- Centralized Logging
- Faster RCA

Unified Telemetry

Latency Observability - Continued

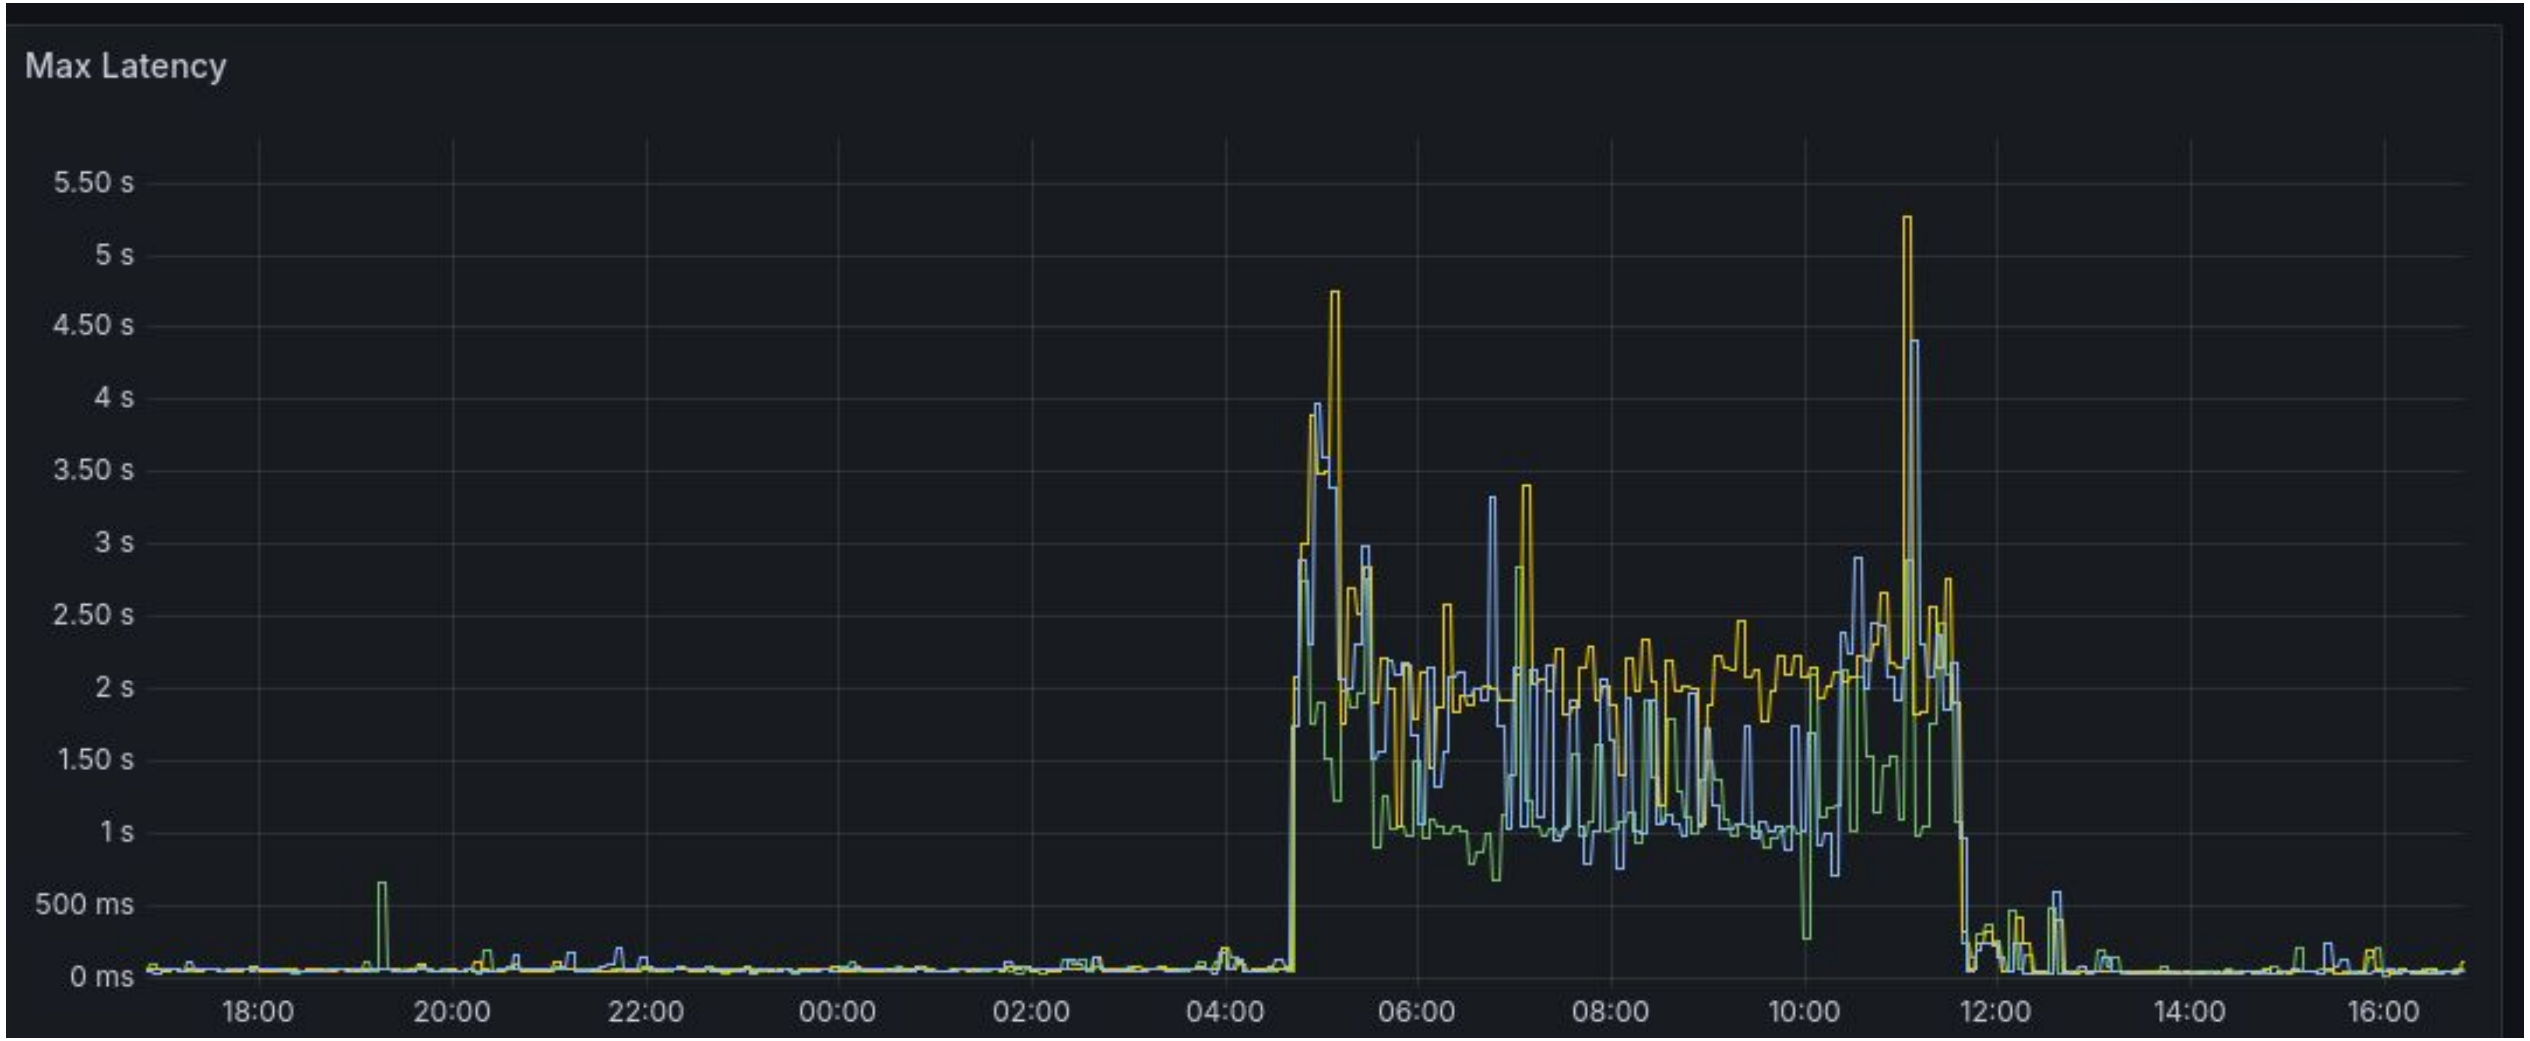

Latency Observability - Continued

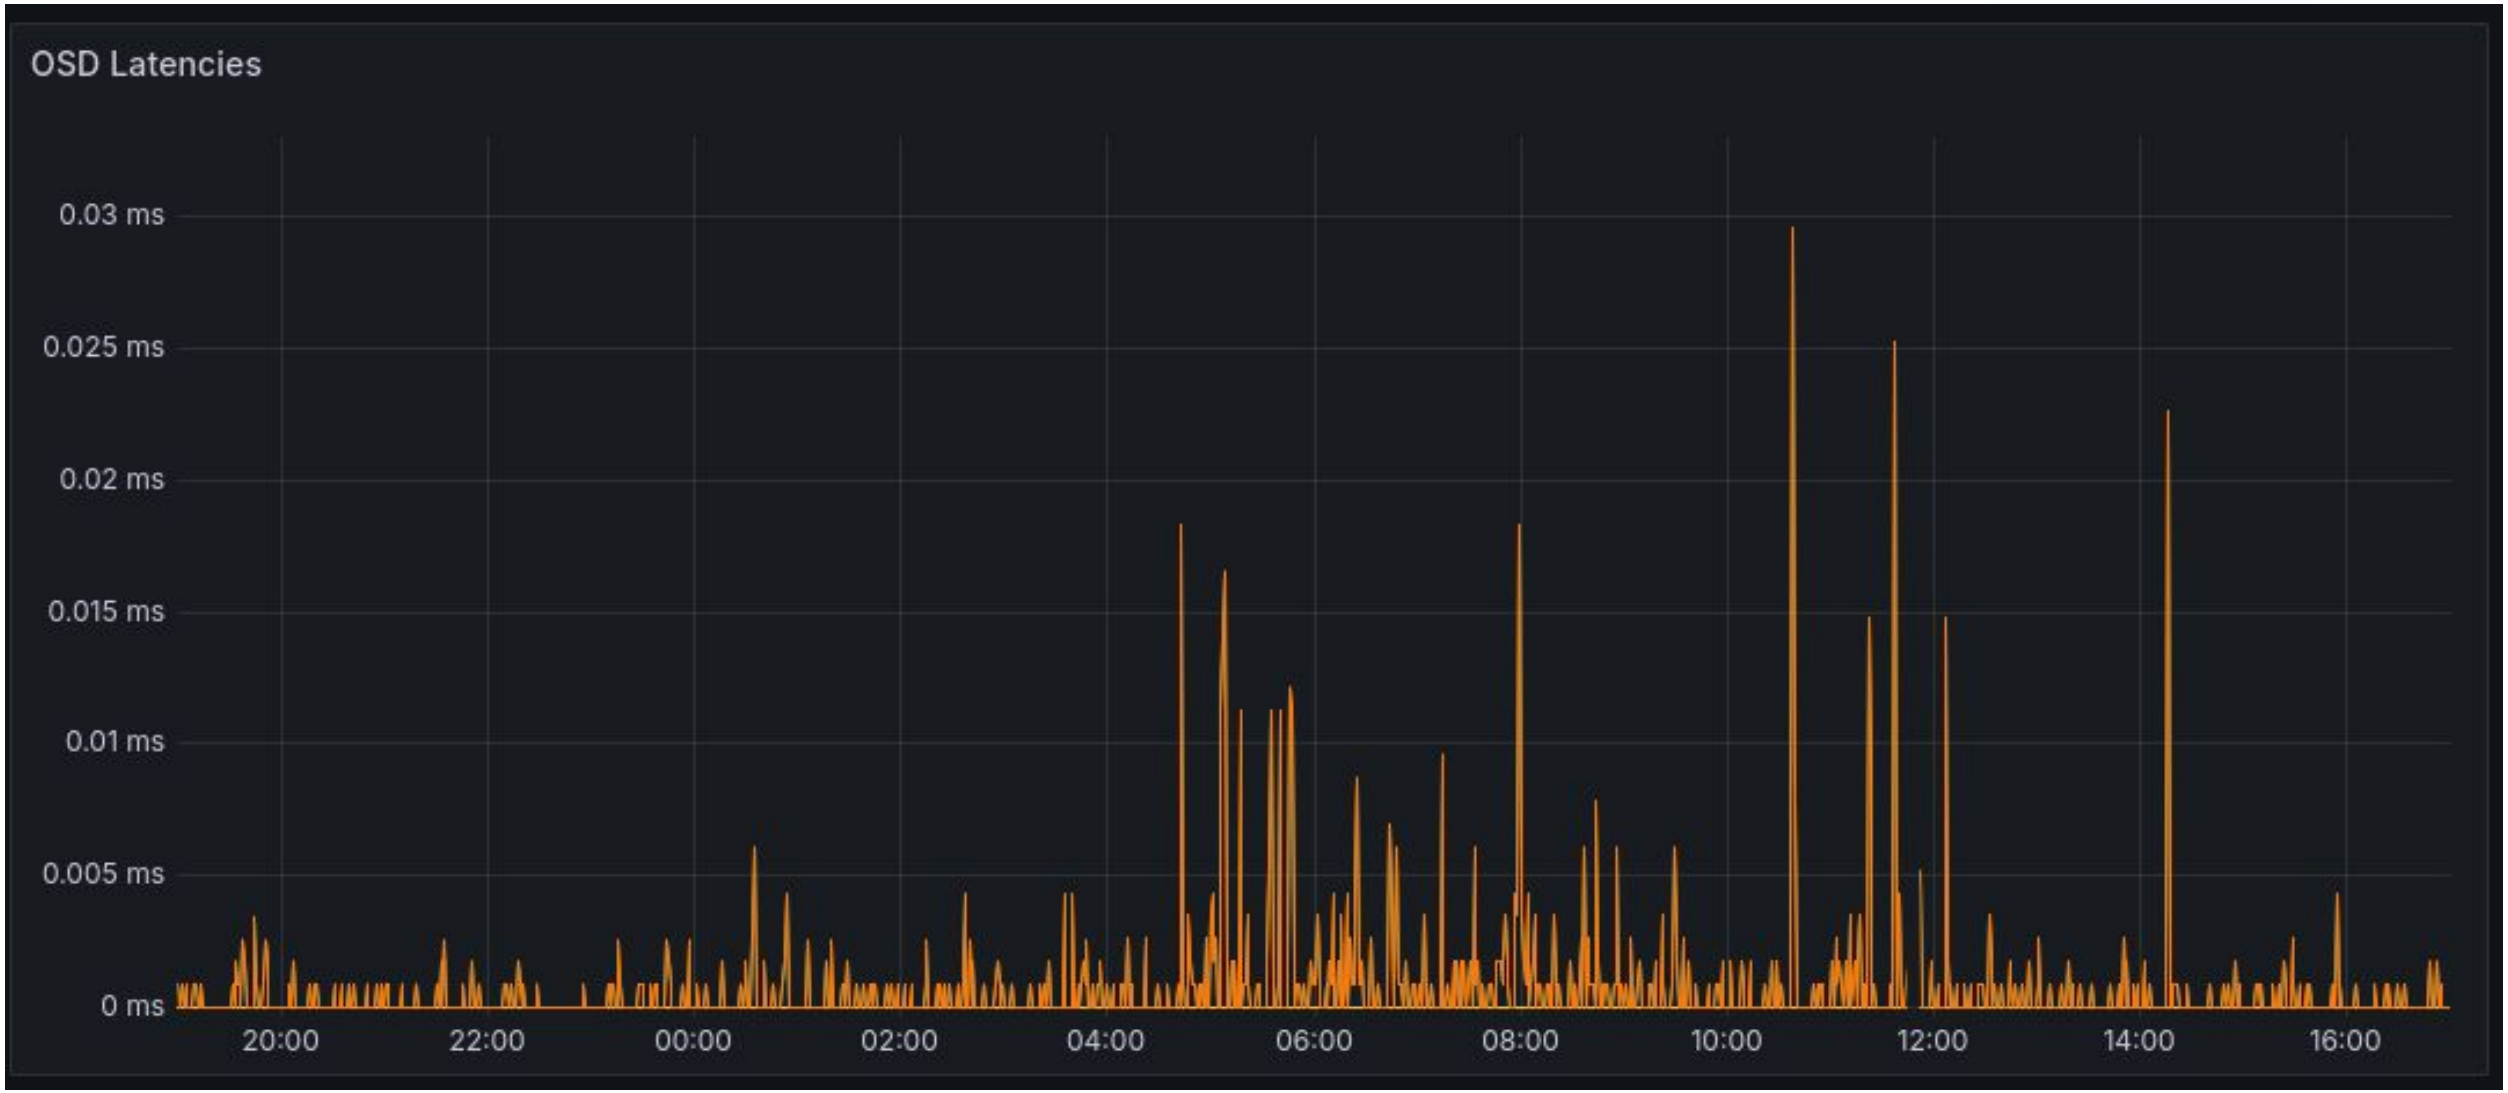

Performance Mathematics

Defining the Performance Cliff

Stable Baseline: Latency remains flat and predictable under normal and moderate IOPS loads.

The Inflection Point: As IO saturation crosses the hardware threshold, queuing delays cause latency to spike exponentially.

SLA Protection: To prevent "hockey stick" degradation, performance capacity limits must be defined and enforced via QoS *before* reaching this inflection point.

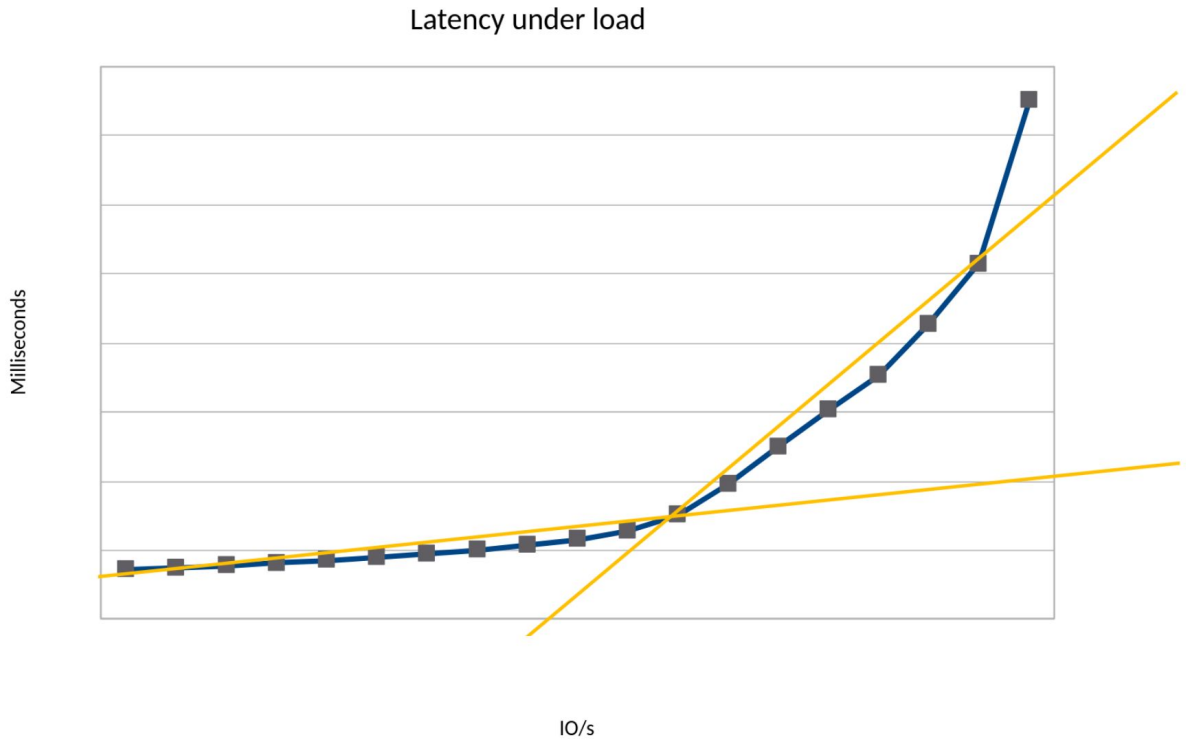

Performance Over Provisioning

Performance Capacity

- Define Aggregate Perf Limits
- Historical Usage Patterns
- Over Provisioning Factor
- Influence Hardware Design

Thundering Herd

- Multi-tenancy
- Absorbing Bursts
- Latency Impact
- Top Talker Observability

Dynamic Throttling

- Inflection Point
- Identify Hot Spots
- Throttling Automation

Know Your Workload

- **Gap:** Ceph metrics show volumes — not who owns them
- **Impact:** Hard to identify top talkers in multi-tenant Environment
- **Approach:** Correlate OpenStack ownership data with Ceph RBD metrics
- **Outcome:**
  - True project-level top talkers
  - Faster RCA & noisy-neighbor detection
  - Workload Characterization
  - Better capacity planning

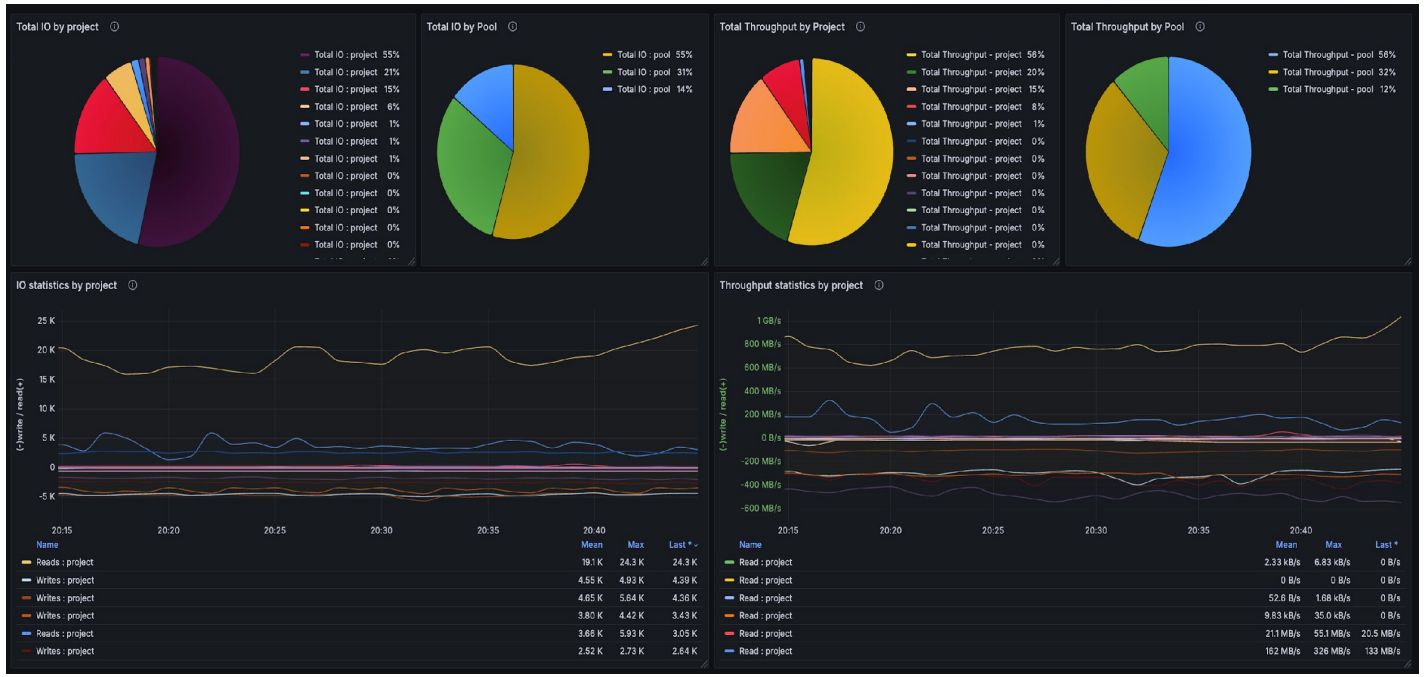

Operational Realities

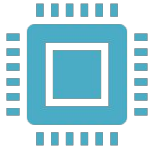

Orchestrator Limitations

Rack-Level Maintenances

Configuration Drift

RGW QoS Limitations

Ceph vs Enterprise Storage TCO

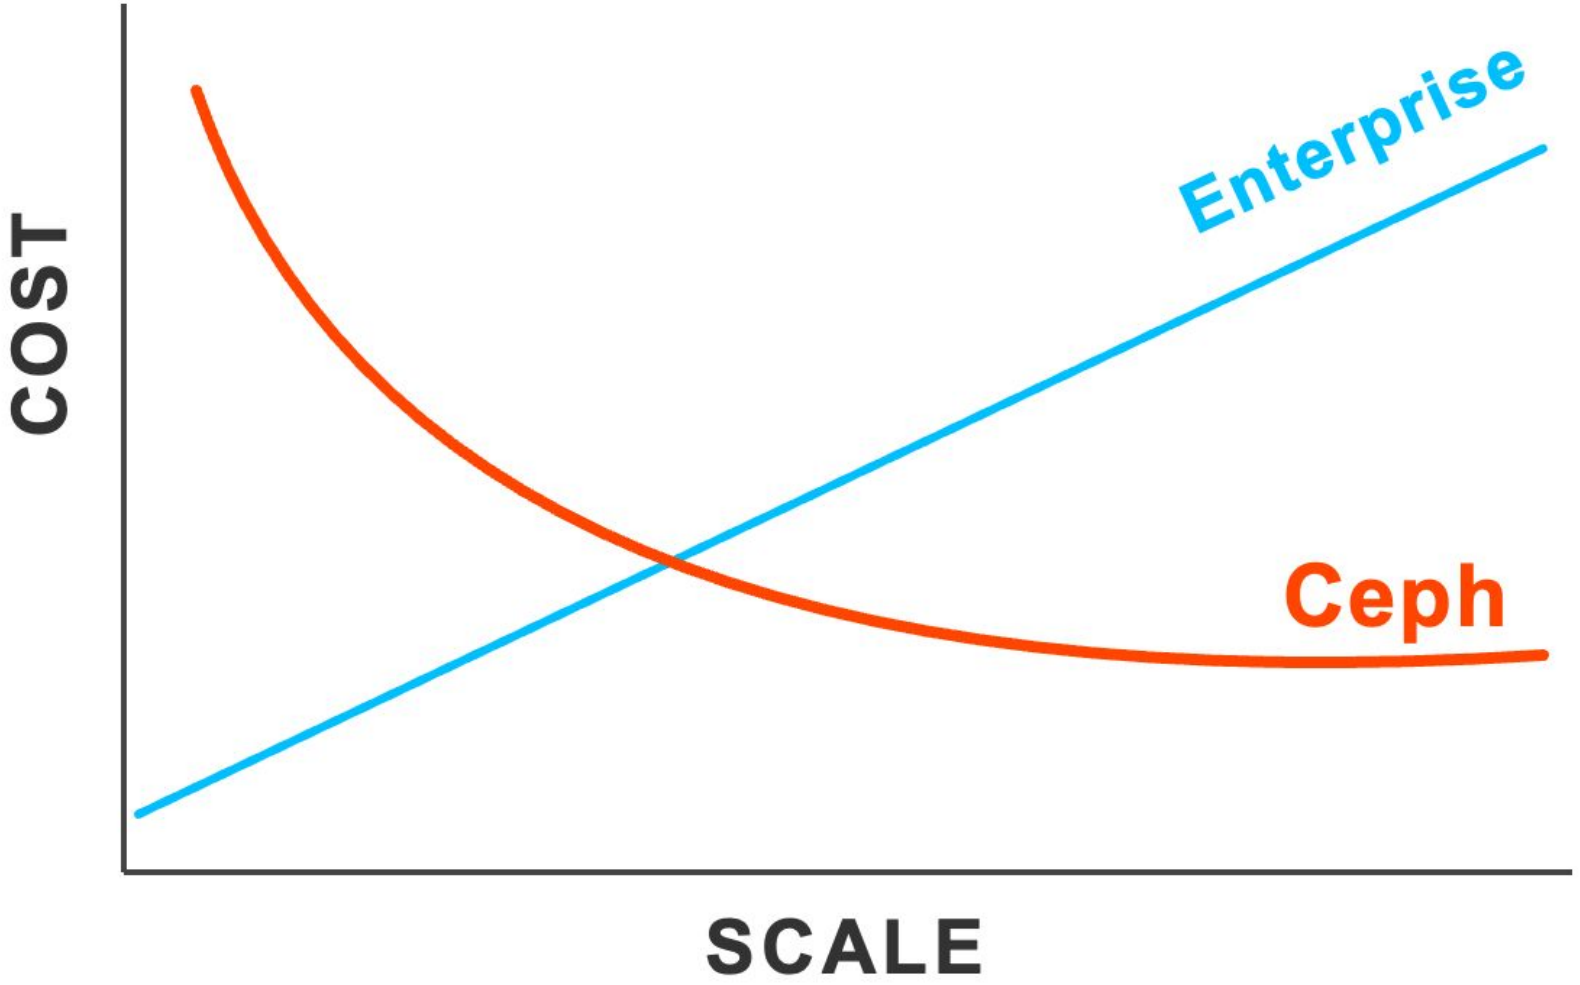

The Road Ahead

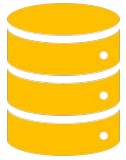

Persistent Storage for Kubernetes

Object Storage Classes

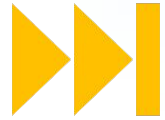

New Features Adoption
Fast EC
Deduplication

Questions?