SDC 20
SNIA INDIA
Storage Developer Conference
December 4-5, 2020
BY Developers FOR Developers

Object Storage Performance: Benchmarking & Troubleshooting

Dheer Moghe (Principal Engineer)
Arth Patel (Sr. Member Technical Staff)
Subin Francis (Sr. Performance Engineer)
Nutanix India

Agenda

- Why Object Storage?
- Benchmarking Object Storage
- Troubleshooting: Tips & Tools
- Performance Results

What is Object Storage?

Block Storage	File Storage	Object Storage
Blocks over Luns	Hierarchical: Files & Directories	Flat: Buckets & Objects
SCSI / iSCSI	NFS / SMB	S3 - HTTP REST
Structured Data	Structured & Semi-Structured Data	Semi-Structured & Unstructured Data

Object Storage Use Cases

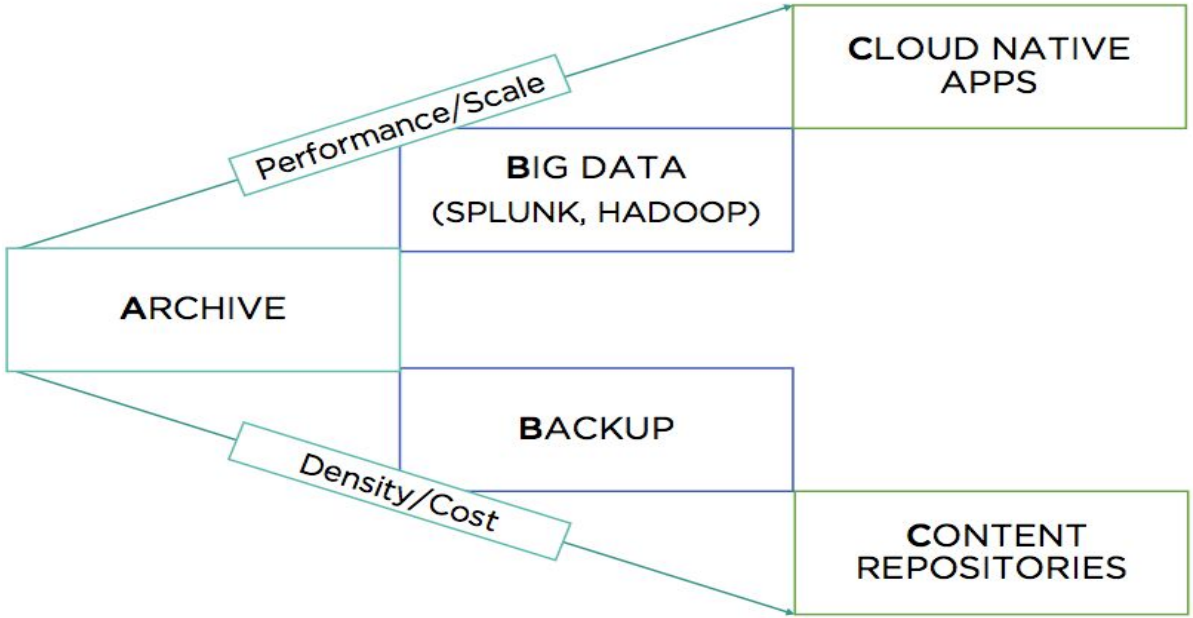

Why Object Storage?

- Enterprise Application natively support object storage
  - Cloud storage, Features like WORM, Legal Hold
- Cloud Native Apps primarily use Object Storage
  - How to run them on-prem?
- Scale & Performance
  - Flat namespace, Multipart Uploads, Object Copy
- Rich & Evolving Feature Set
  - Website hosting, Notifications, Tagging, Rich User Metadata,

Scaleout Object Storage Building Blocks

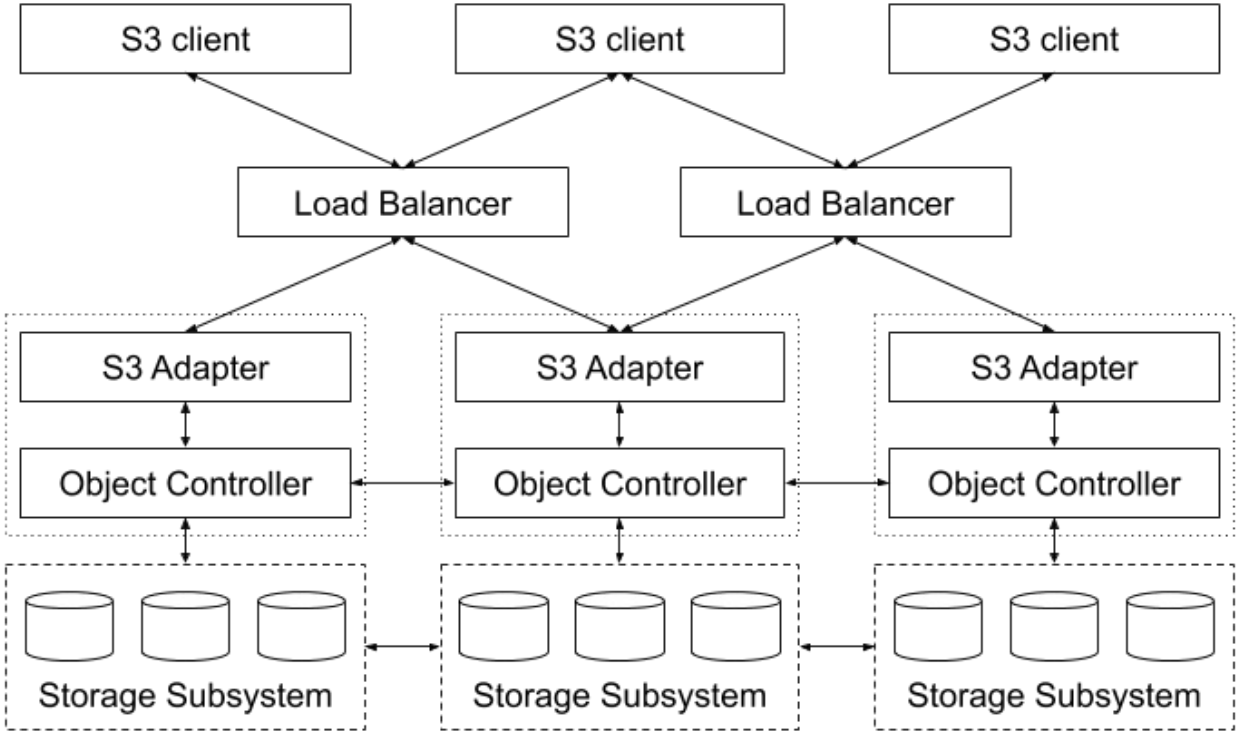

Performance Metric For Object Storage

- Bandwidth
  - Large Object IO
- Latency (**T**ime **T**o **F**irst **B**yte)
  - Small Object IO

Existing Object Storage Benchmarks

- Single instance benchmarks
  - Limited network bandwidth
- Distributed benchmarks
  - Complex setup
  - Unstable
  - Limited performance

Benchmarking Object Storage

- Poseidon Benchmark (Homegrown distributed tool)
  - Containerized
  - Configurable data compressibility
  - Object size range and distribution (zipf, linear, …)
  - IO patterns
  - Ability to bypass components from the stack

Benchmarking Object Storage

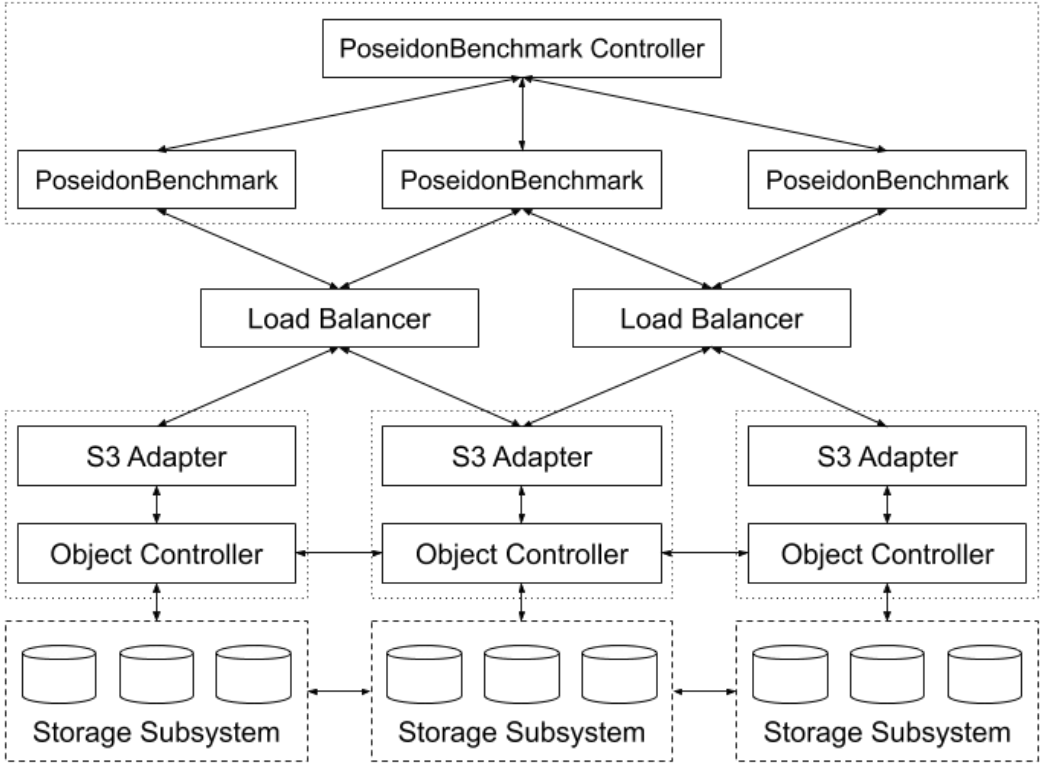

Benchmarking Object Storage

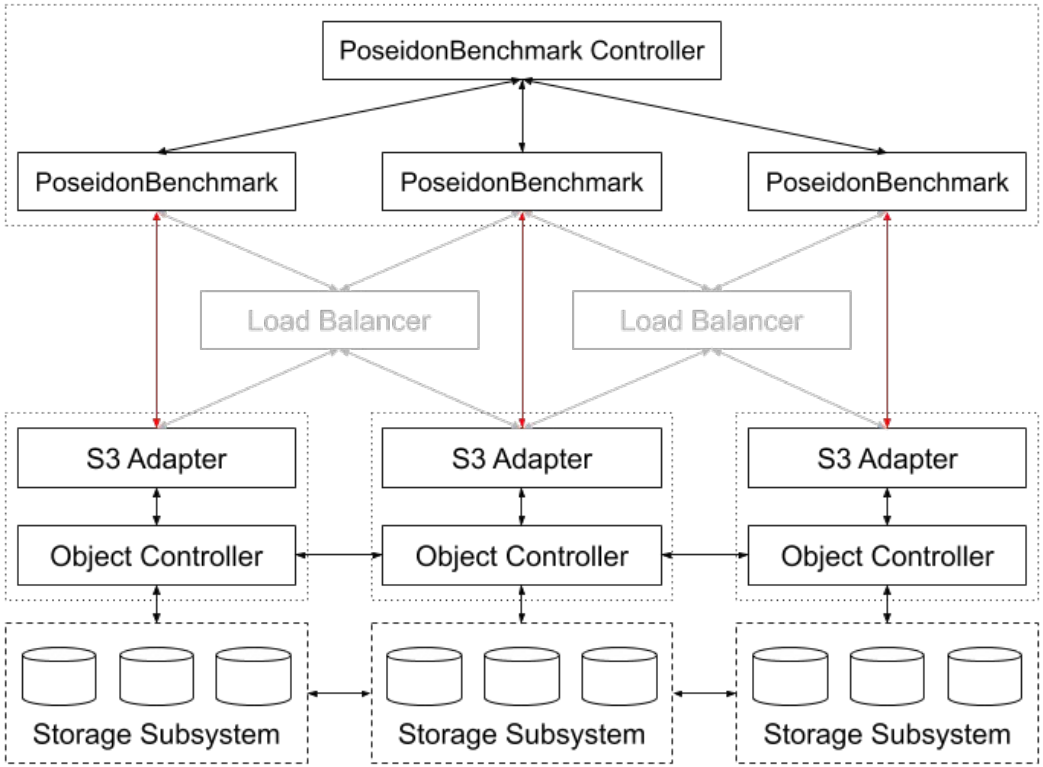

Benchmarking Object Storage

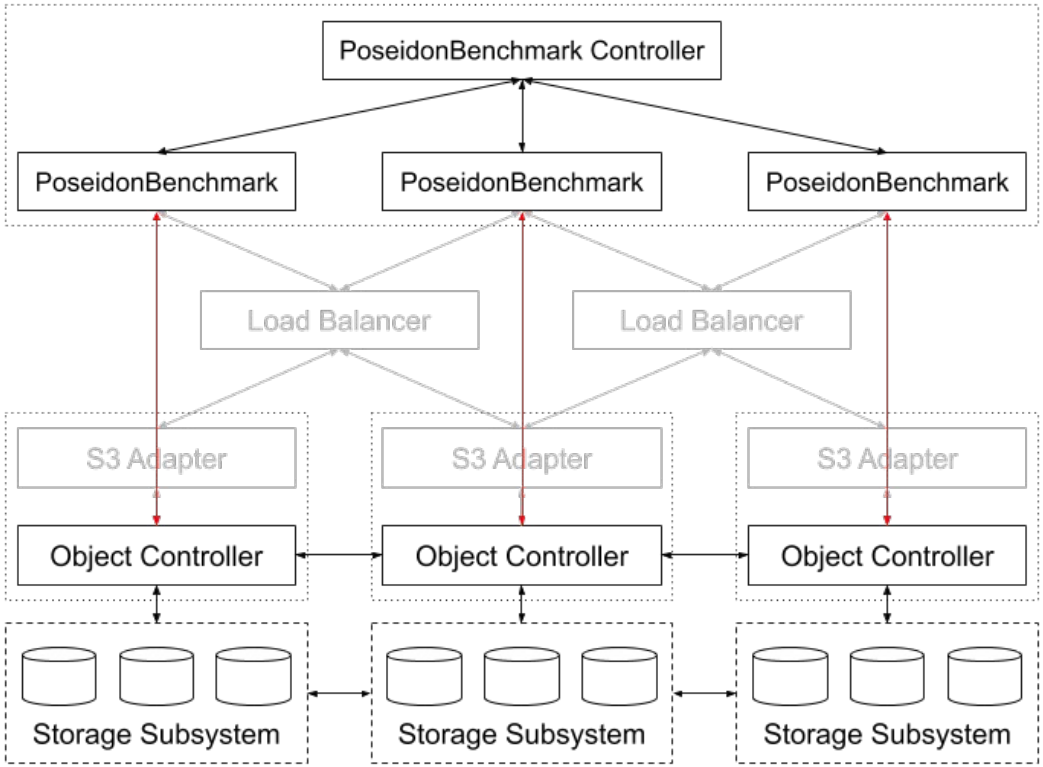

Benchmarking Object Storage

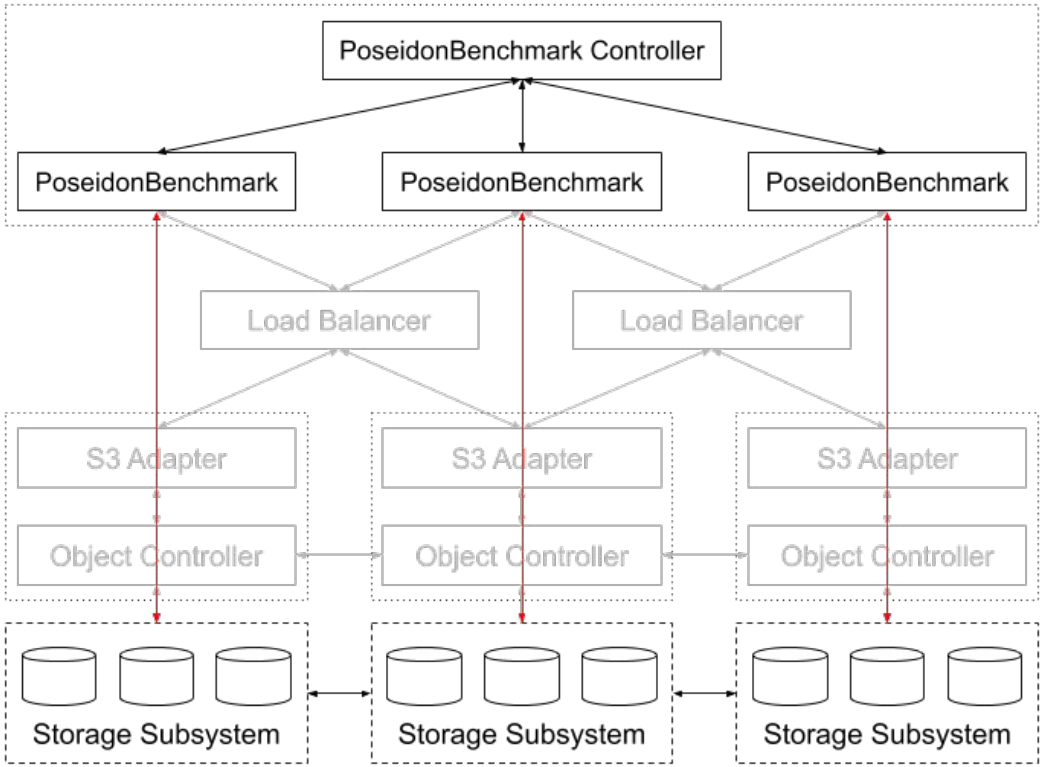

Optimizations : Pipelining

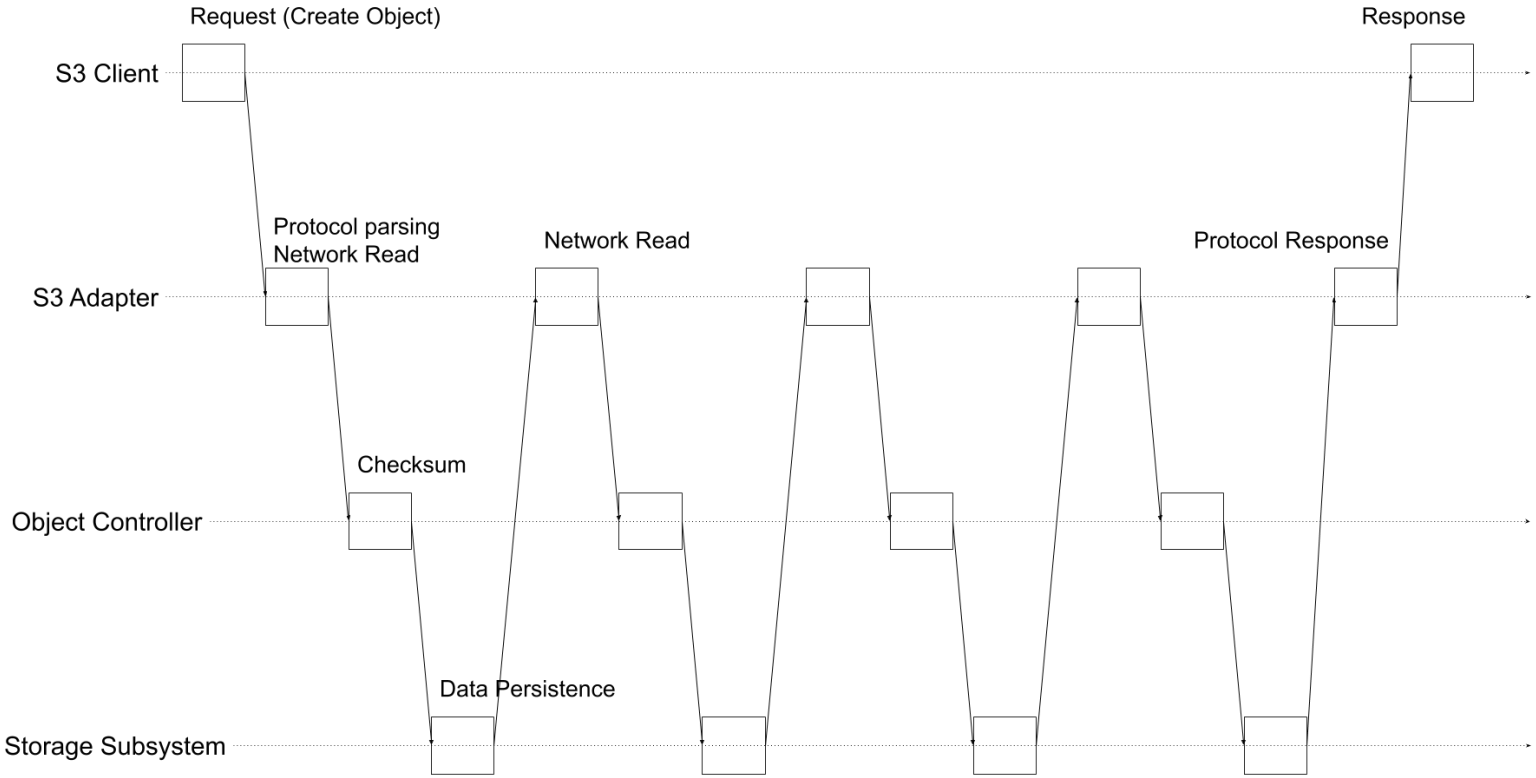

Optimizations : Pipelining

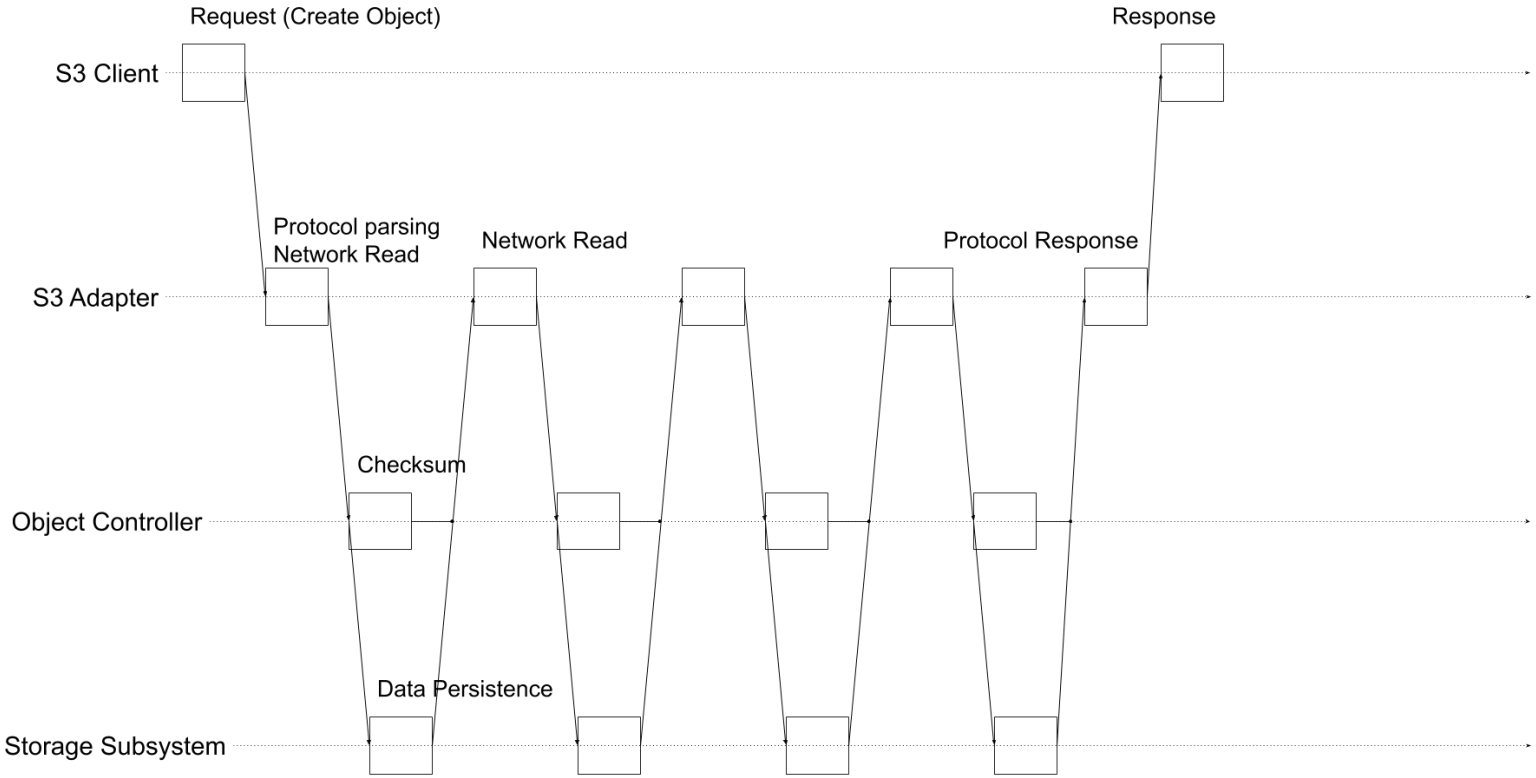

Optimizations : Pipelining

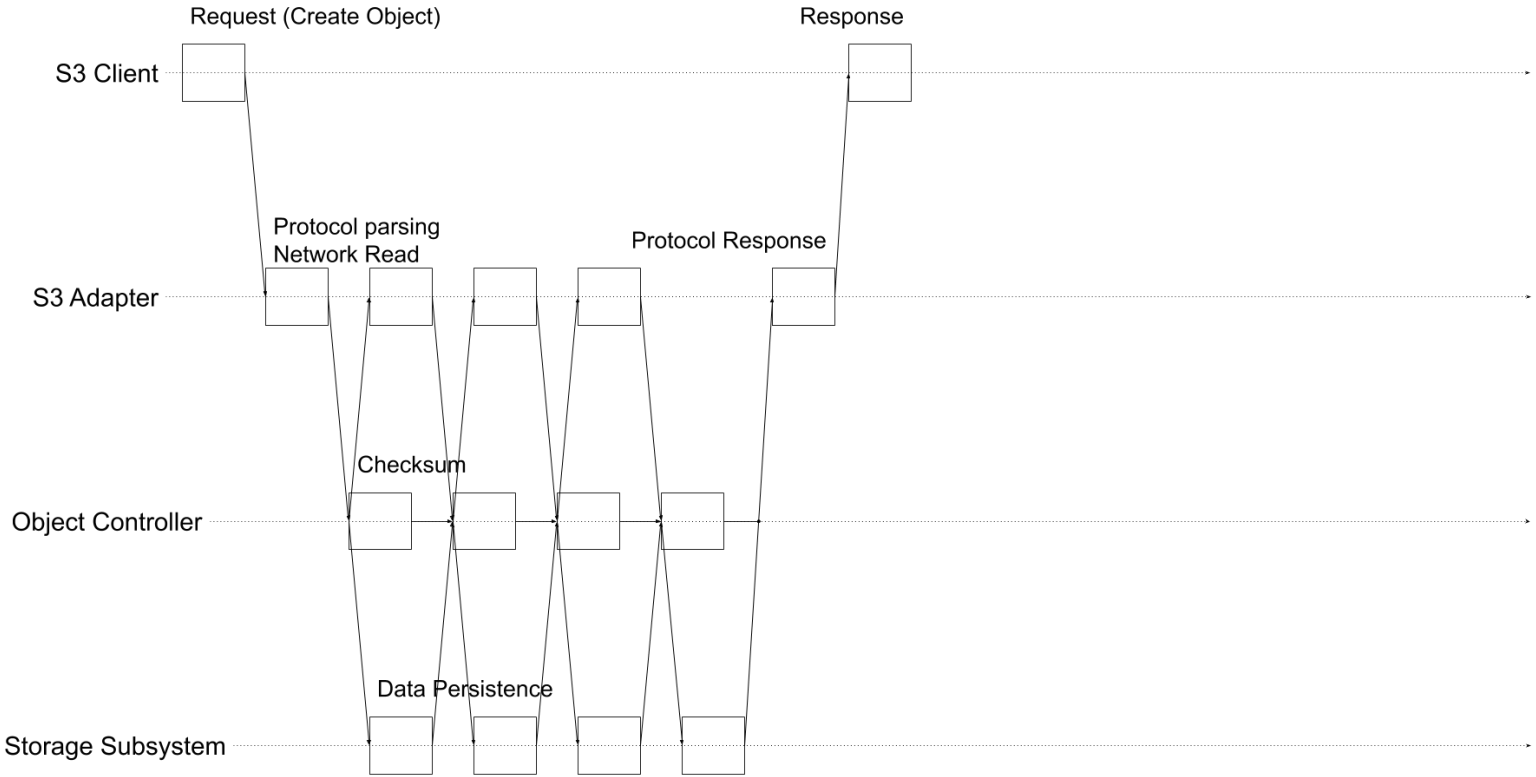

Optimizations : Reordering

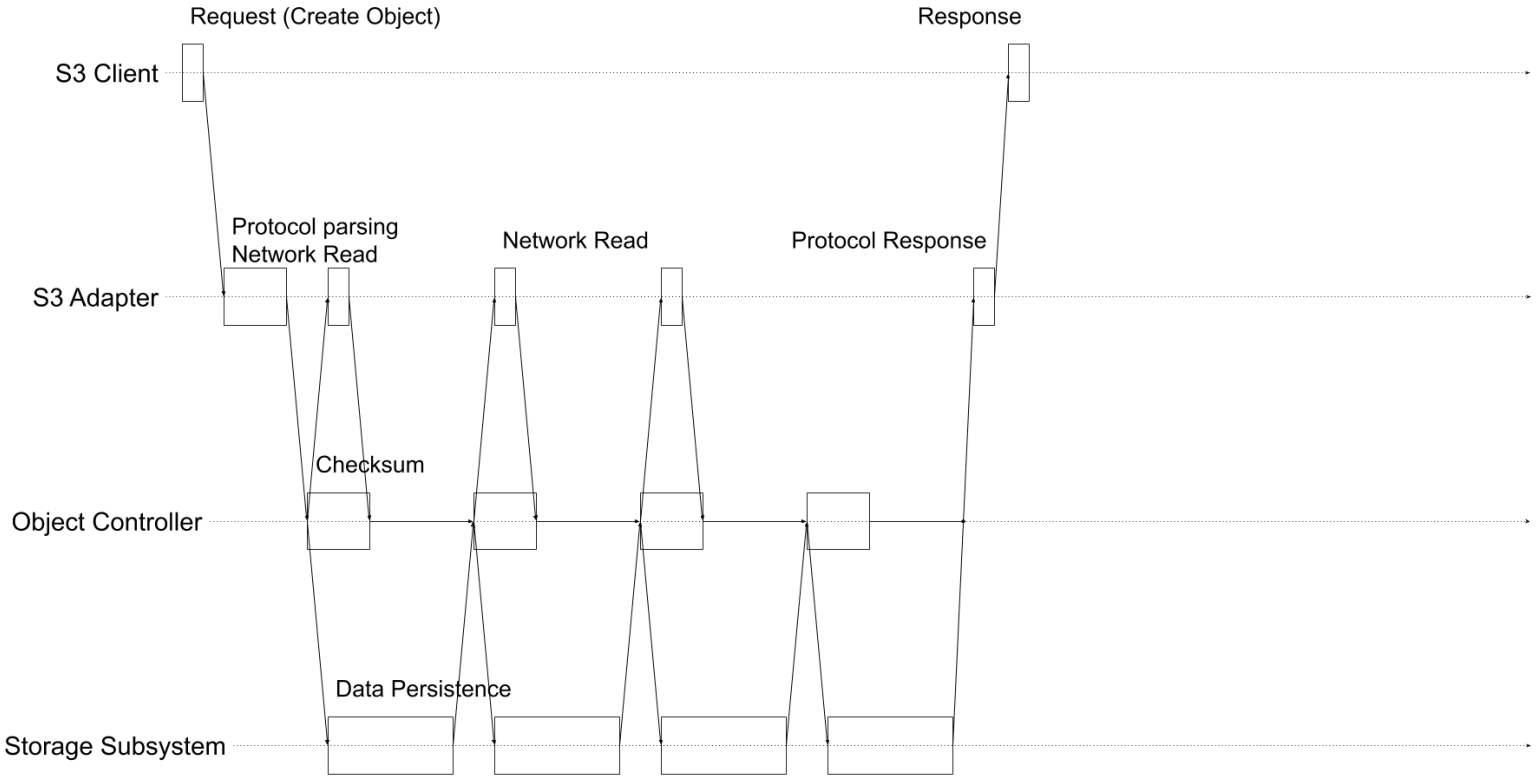

Optimizations : Reordering

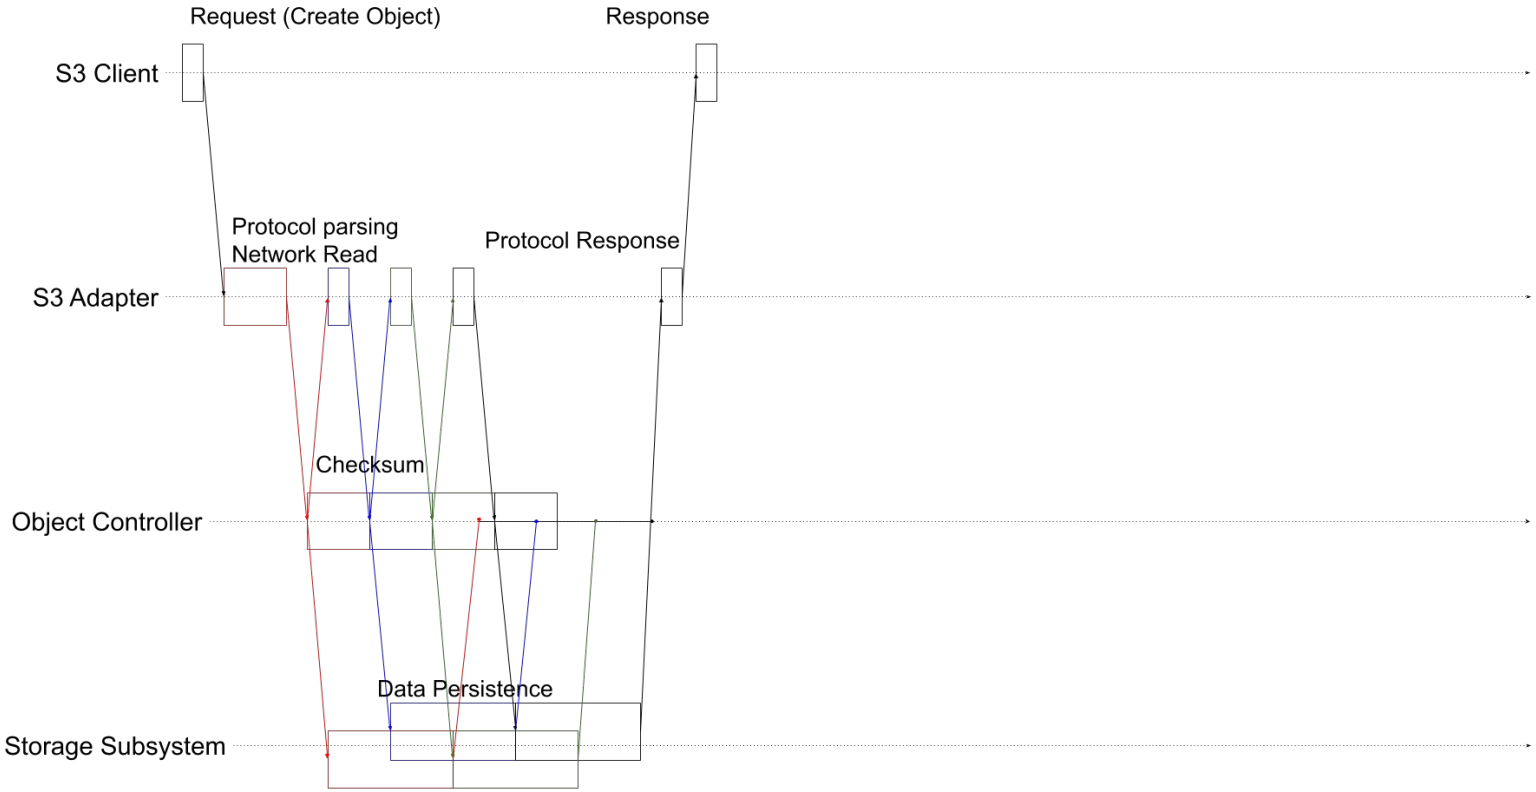

Object Store Network Architecture

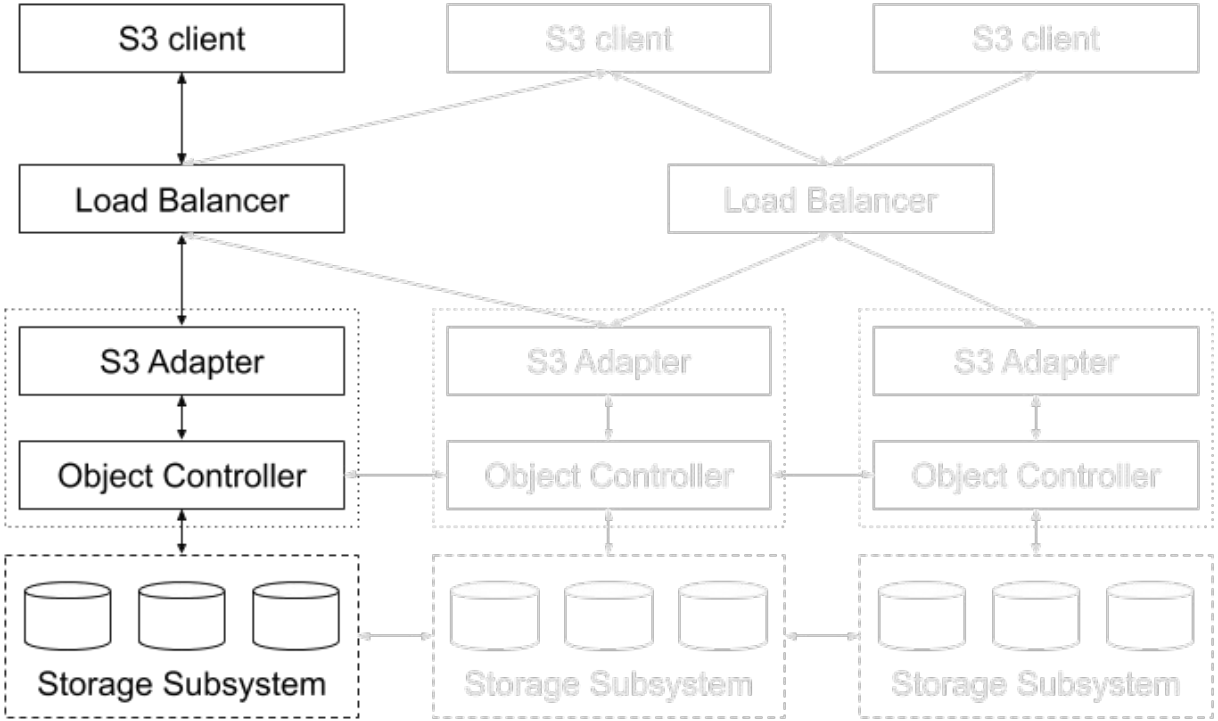

Object Store Network Architecture

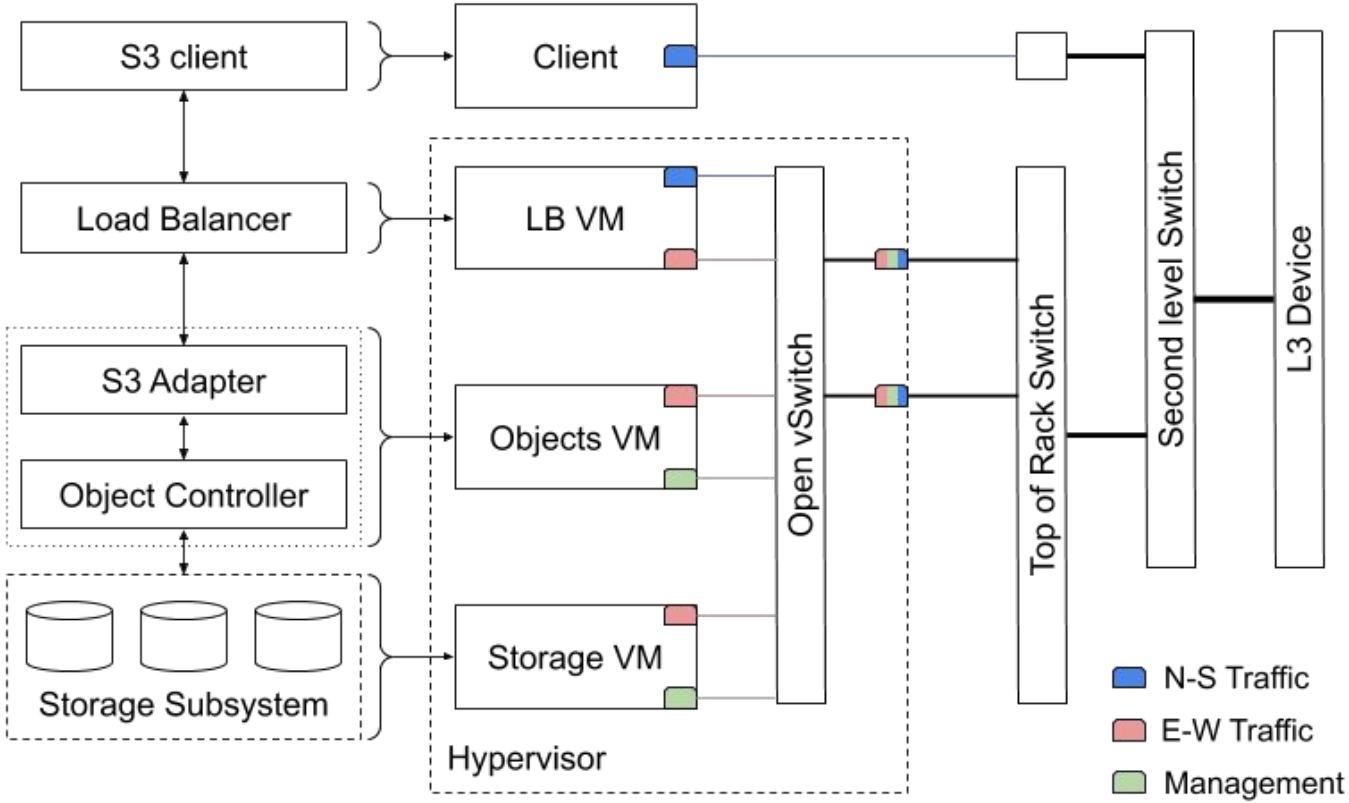

Troubleshooting: Networking Issues

- Vhost-net cpu saturation
- Cross network communication
- Physical link saturation

Troubleshooting: Debugging Tools

- Nutanix X-RAY
  - HCI performance testing app, automate performance scenarios, support custom scenarios
  - Downloadable virtual machine image, ansible automation, partly open-sourced
  - With x-ray integration, large scale iperf scenarios can be driven from UI

Troubleshooting: Iperf Scenarios

Node level iperf chain

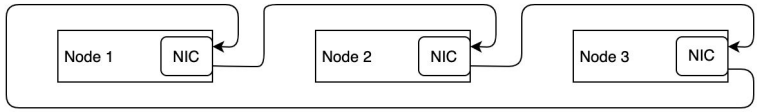

VM level iperf chain - single network

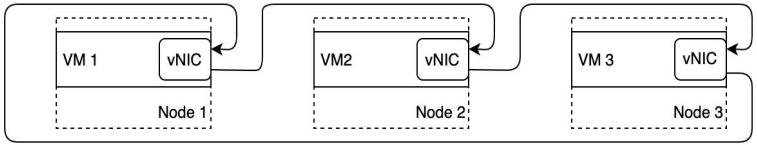

VM level iperf chain - dual network

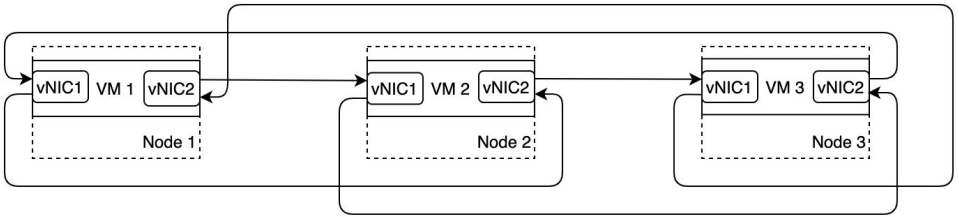

Troubleshooting: Network Optimisations

- Multi queue vNICs
- Storage VM and Object store in same network
- vSwitch bonded uplinks

Performance Results

User VM network performance default vs optimised network

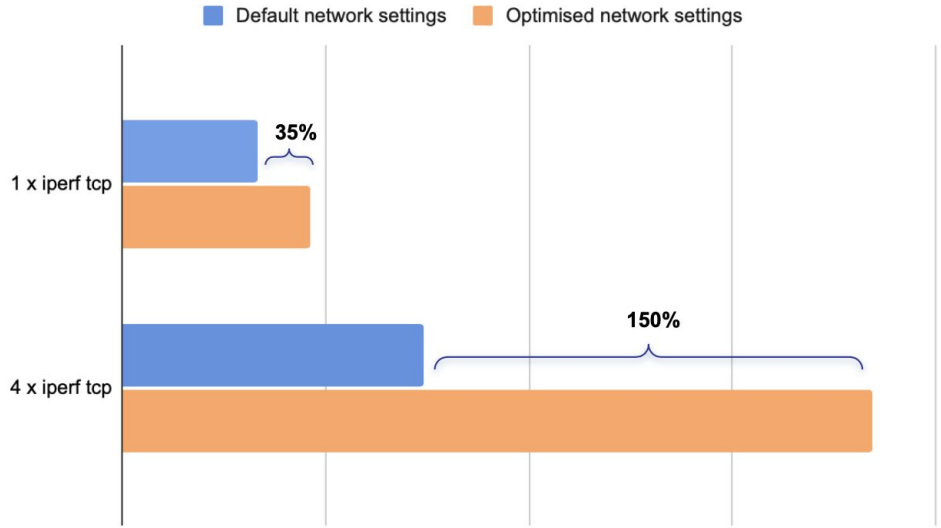

Performance Results

Object storage performance default vs optimised network

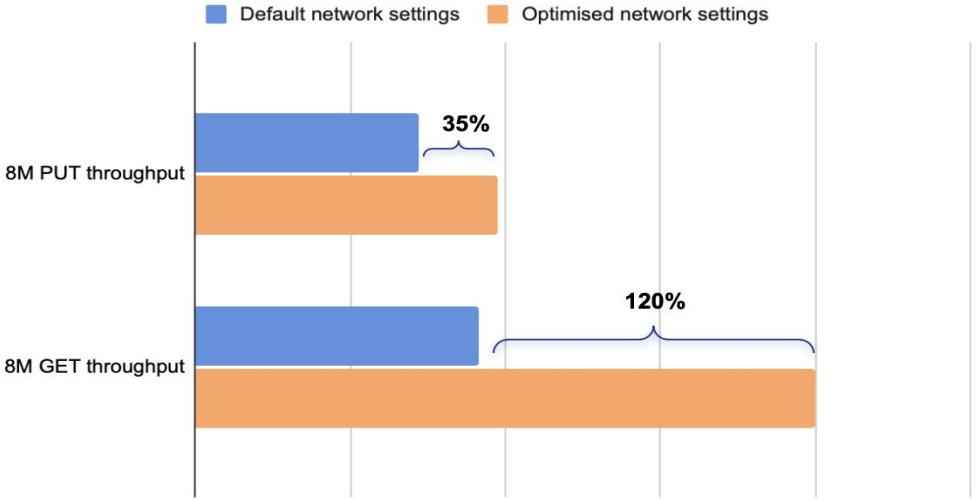

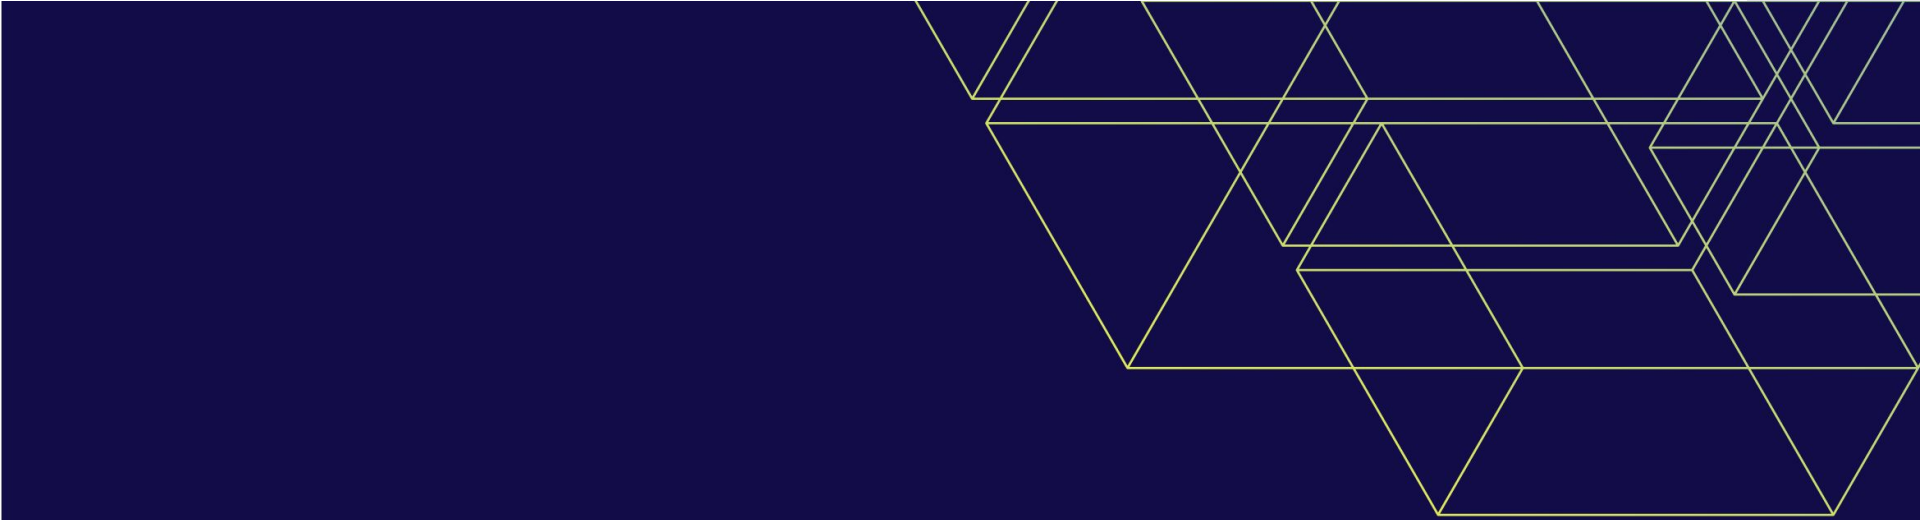

Please take a moment to rate this session.

Your feedback matters to us.