

OPEN

Compute Engineering Workshop

March 9, 2015


San Jose



Seagate Kinetic Open Storage Platform

David Burks, Director of Product Management at Seagate

Mayur Shetty, Senior Solutions Architect at Seagate



Agenda

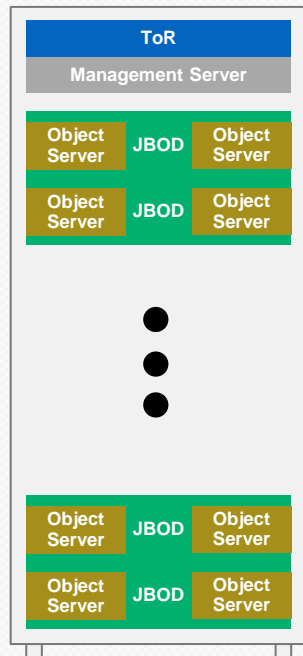
- **Provide information needed to benchmark Kinetic-Swift**
 - Preparing the Cluster
 - Start with a Goal
 - Pick the right Tools
 - Understand the System
 - Monitor for Bottlenecks
 - Analyze the results
- **Benchmarking with knobs**



TCO Case Study for 500 PB Swift Deployment




Conventional



8 x 5U
84-Bay

Kinetic



Price Comparison

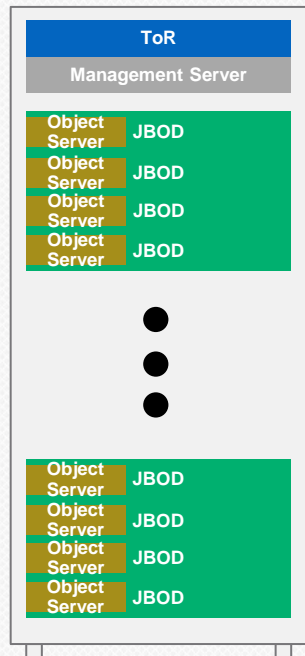
	Conventional	Kinetic	Savings
Total Storage Deployed (TB)	500,000	500,000	0.00%
HDD Quantity	125000	125000	0.00%
Total Capex	\$39,652,824	\$36,813,538	7.16%
Avg. Power per Rack (W)	9392	6320	32.71%
Annual Power Expense	\$21,885	\$14,727	32.71%
TCO Cost/TB	\$30.95	\$26.87	13.19%

“Worse case” TCO comparing very dense conventional system to equally dense Kinetic system

TCO Case Study for 500 PB Swift Deployment




Conventional



40 X 1U
12-Bay
vs
8 x 5U
84-Bay

Kinetic



Price Comparison

	Conventional	Kinetic	Savings
Total Storage Deployed (TB)	500,000	500,000	0.00%
HDD Quantity	125000	125000	0.00%
Total Capex	\$52,204,015	\$36,813,538	29.48%
Avg. Power per Rack (W)	12560	6320	49.68%
Annual Power Expense	\$29,267	\$14,727	49.68%
TCO Cost/TB	\$45.50	\$26.87	40.94%

“Great Case” TCO comparing low density conventional system to very dense Kinetic system



Kinetic Adoption

Solutions & System Verification

Chassis

- Sanmina 4U 60 - Now!
- Supermicro 1U 12 - Now!
- Hyve 1U 16 - Now!
- Rausch 4U7 - Now
- Seagate – Q4, 2015



Kinetic Enabled Software

- Swift/Swift Stack - Now
- Scalify - Q2, 2015
- Nexenta - Q3 2015
- OpenVStorage - Q2 2015
- Ceph – Q3 2015




System Test/Configuration & Benchmarking




Kinetic vs Makara, latest comp

Reads and Writes (MB/s)



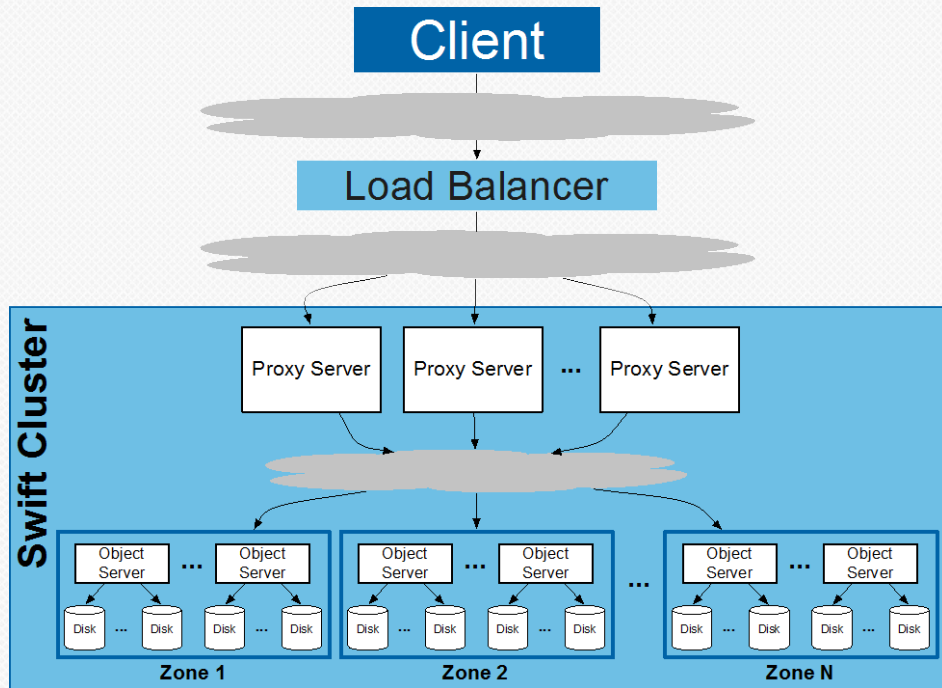
Kinetic vs Makara, latest comp

Reads and Writes (90% resp time in Mliseconds)

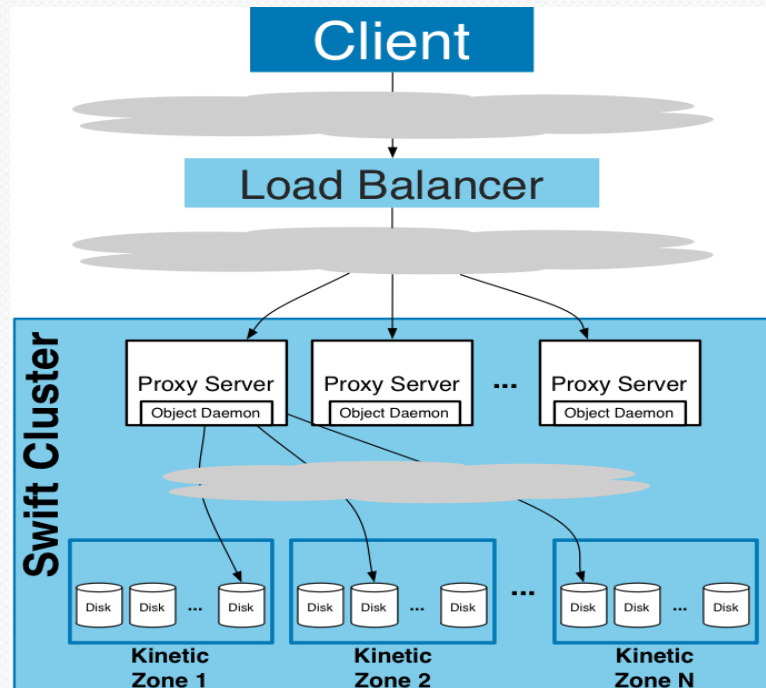


OpenStack Swift

Conventional Swift



Kinetic Swift



Preparing the Cluster

- **Networking**
- **Storage**
- **Kinetic Drives**
- **Kinetic-Swift**
 - <https://github.com/swiftstack/kinetic-swift/wiki/Deployment>



Start with a Goal

- Capacity
- Throughput
- Object Sizes
- Latency



Benchmark options

- COSBench by Intel
- getput by HP
- swift-bench
- ssbench
- knobs by Seagate



Benchmarking with knobs

- **Installing knobs**

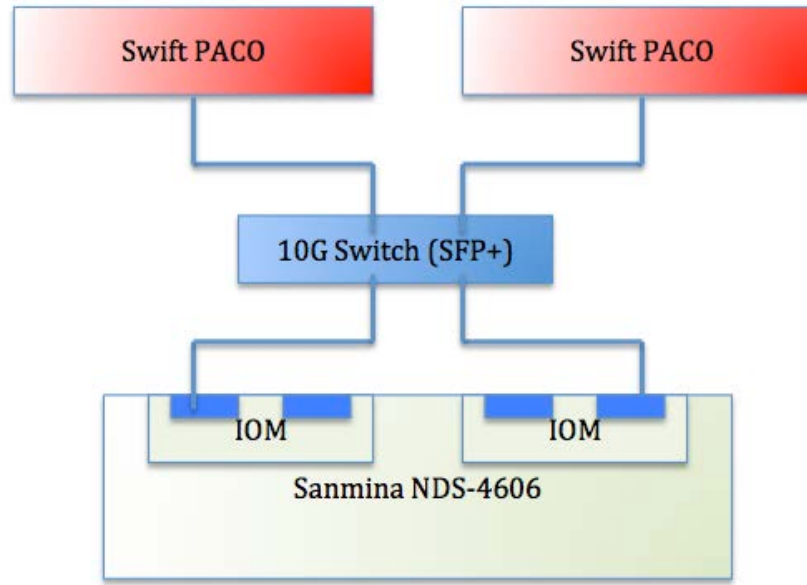
- <https://github.com/Seagate/knobs>

- **knobs Usage**

- `benchmarker-swift-rwrite.py -total 1000 10 http:127.0.0.1:8080/auth/v1.0 -size 8000000`
 - `benchmarker-swift-read.py <timestamp.knobs>`
 - `tochart.py`



Understand the System



- 2 x 10 GbE ports from the Kinetic chassis, hence a max bandwidth of ~ 2.5GB/s
- 2 x 10 GbE port from the Swift PACO node, max bandwidth of ~2.25 GB/s



Understand the System

**With 60 drives in the chassis, and each drive delivering 50 MB/s we get:
 $60 \times 50 \text{ MB/s} = 3000 \text{ MB/s}$**

Drive Utilization	Chassis perspective (x 3)	Both PACO (App/Benchmark)	Single PACO (App/Benchmark)	Network utilization of single PACO (1 x 10 GbE)
25%	750 MB/s	250 MB/s	125 MB/s	10%
50%	1500 MB/s	500 MB/s	250 MB/s	20%
~60%	1750 MB/s	583.33 MB/s	291.67 MB/s	23.33%



Knobs 8 MB Object test

Knobs 8 MB object test (put):

Rate	Throughput (Ingress rate)	Throughput (drive perspective x 3)	% Network Utilization (chassis)	Average Latency	Throughput delivered
31.25	250 MB/s	750 MB/s	30%	531.41ms	234.32 MB/s
62.5	500 MB/s	1500 MB/s	60%	1.41s	447.73 MB/s
72.91	583.28 MB/s	1750 MB/s	70%	9.04s	530.41 MB/s

Knobs 8 MB object test (get):

Rate	Throughput	% Network Utilization	Average Latency	Throughput delivered
93.75	750 MB/s	30%	233.075ms	678.04 MB/s
187.5	1500 MB/s	60%	334.837ms	1325.72 MB/s
218.75	1750 MB/s	70%	419.4ms	1492.4 MB/s



Thank You!

More information at developers.seagate.com

Mayur Shetty
Senior Solutions Architect
mayur.shetty@seagate.com