



Building Big Data Analytics Data Lake with All-Flash Ceph

QCT Marco Huang
QCT Amy Chang



Agenda

- Introduction of QCT
- Why data lake architecture
- Brief on Data Lake with All-Flash Ceph architecture
 - Architecture design
 - Hardware selection
 - Testing result
- Conclusion






A leading cloud datacenter
solution provider that delivers

Server, Storage, Networking,


Rack System and

Cloud Solution


under a single, proven roof




Server



Storage



Networking



Rack System



Solution



Data-powered company needs flexible data analytic framework

Popular Hadoop framework is top choice for executing analytical tasks yet can't scale-out on-demand



Data-Generating
Center



Data-Powered
Company

Paradigm Shift for Enterprise:
Need for data analytics increases



Hyper-converged
Hadoop Framework

VS



Disaggregated SDS
Architecture

Prefer Disaggregated Architecture:
Target to scale-out on demand



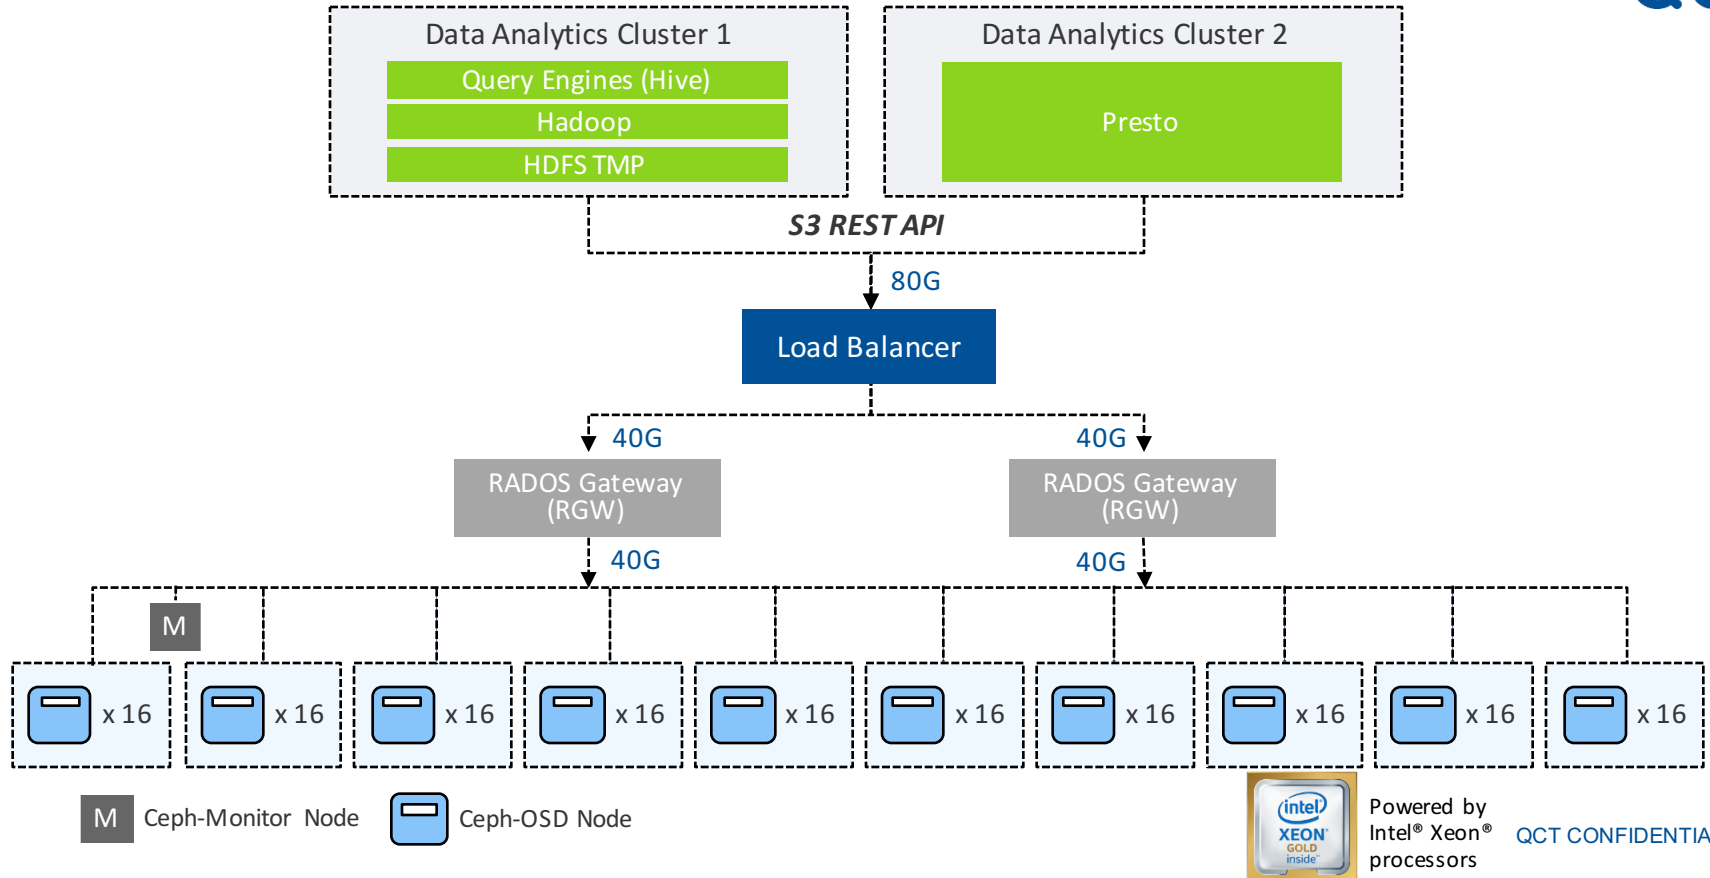
Powered by
Intel® Xeon®
processors

www.QCT.io

QCT CONFIDENTIAL

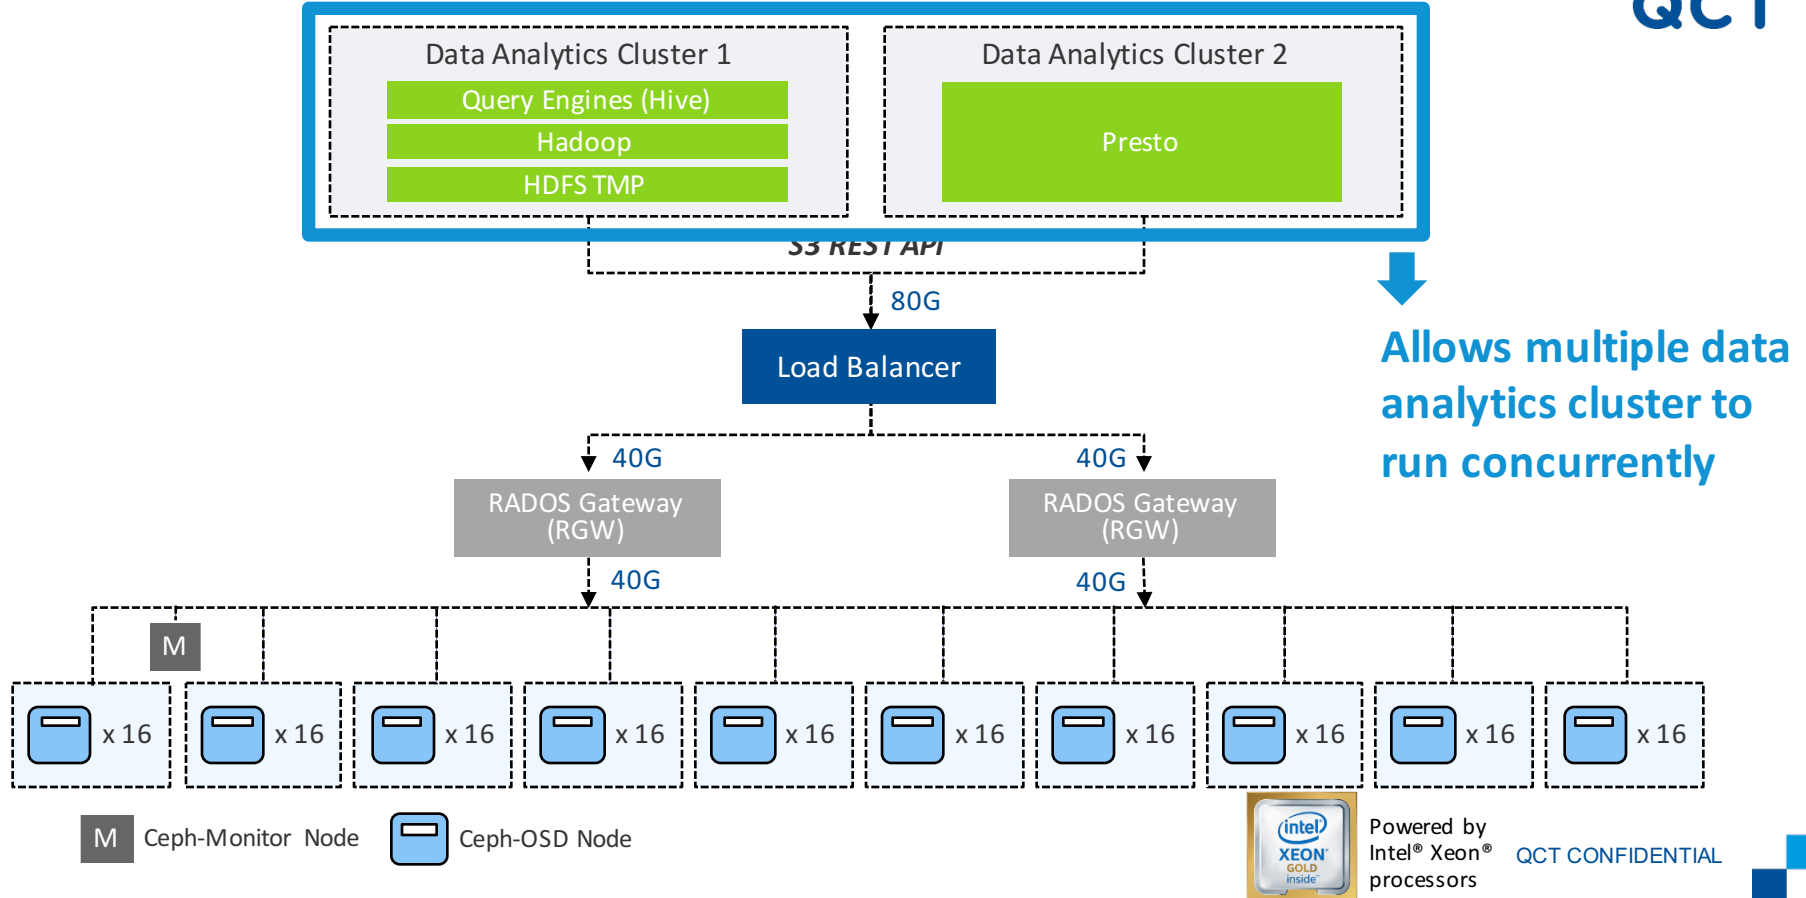
Data Lake with All-Flash Ceph Architecture

Disaggregate data analytics cluster and backend storage to provide higher flexibility



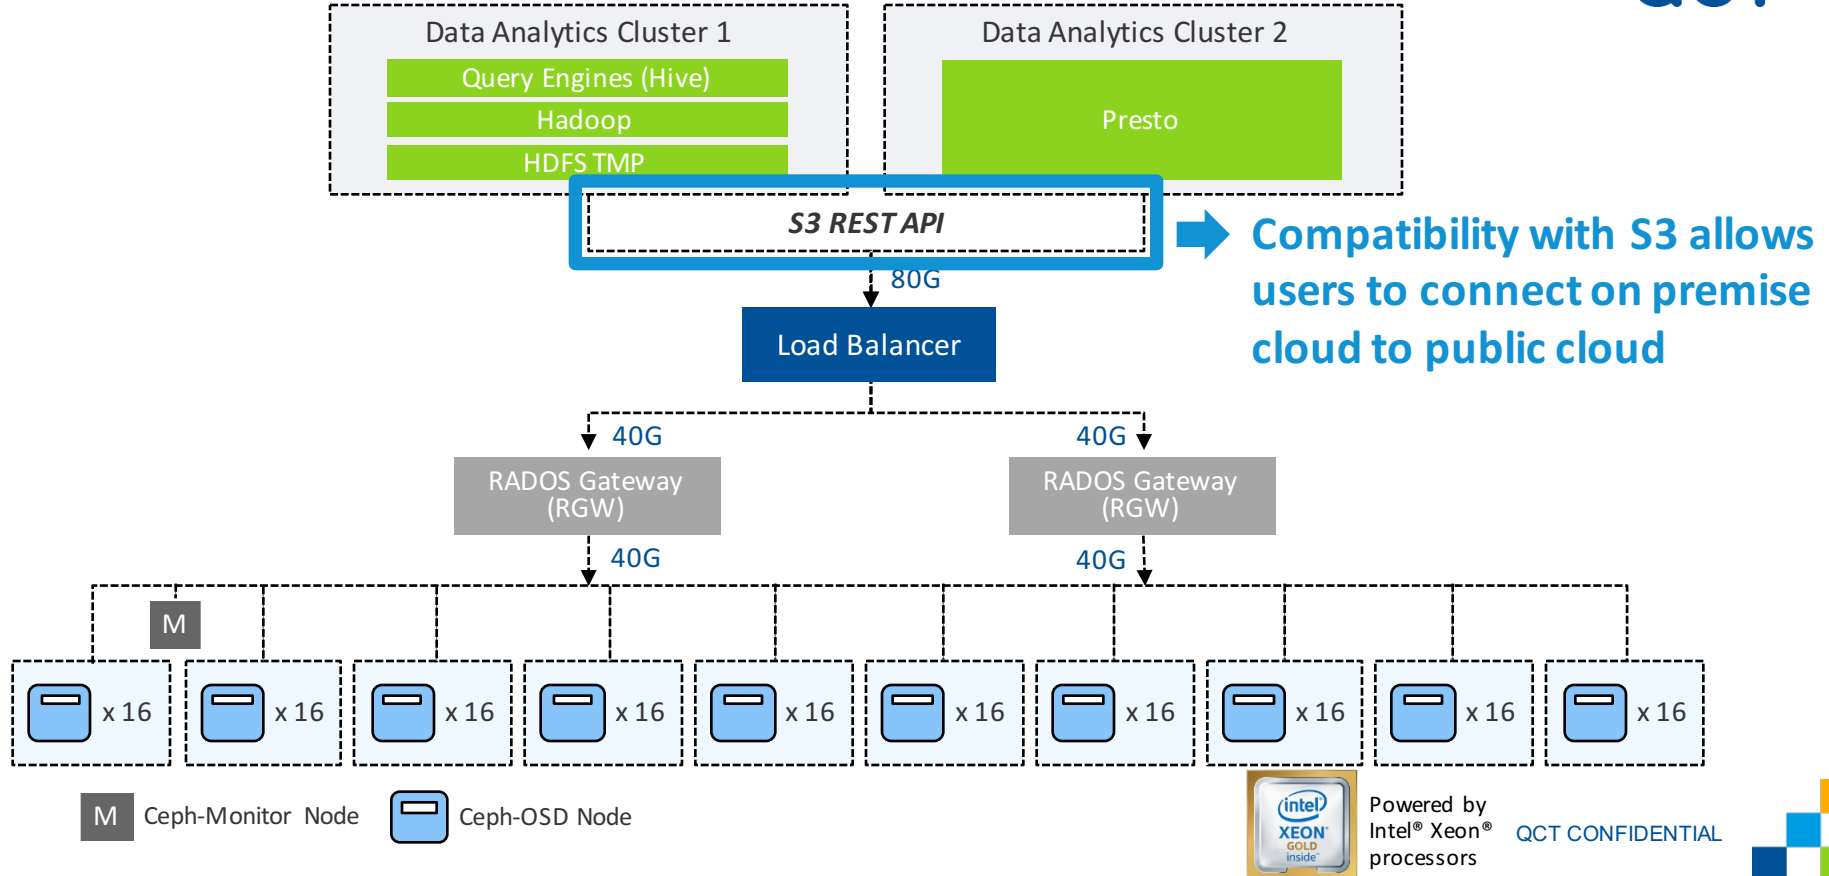
Data Lake with All-Flash Ceph Architecture

Disaggregate data analytics cluster and backend storage to provide higher flexibility



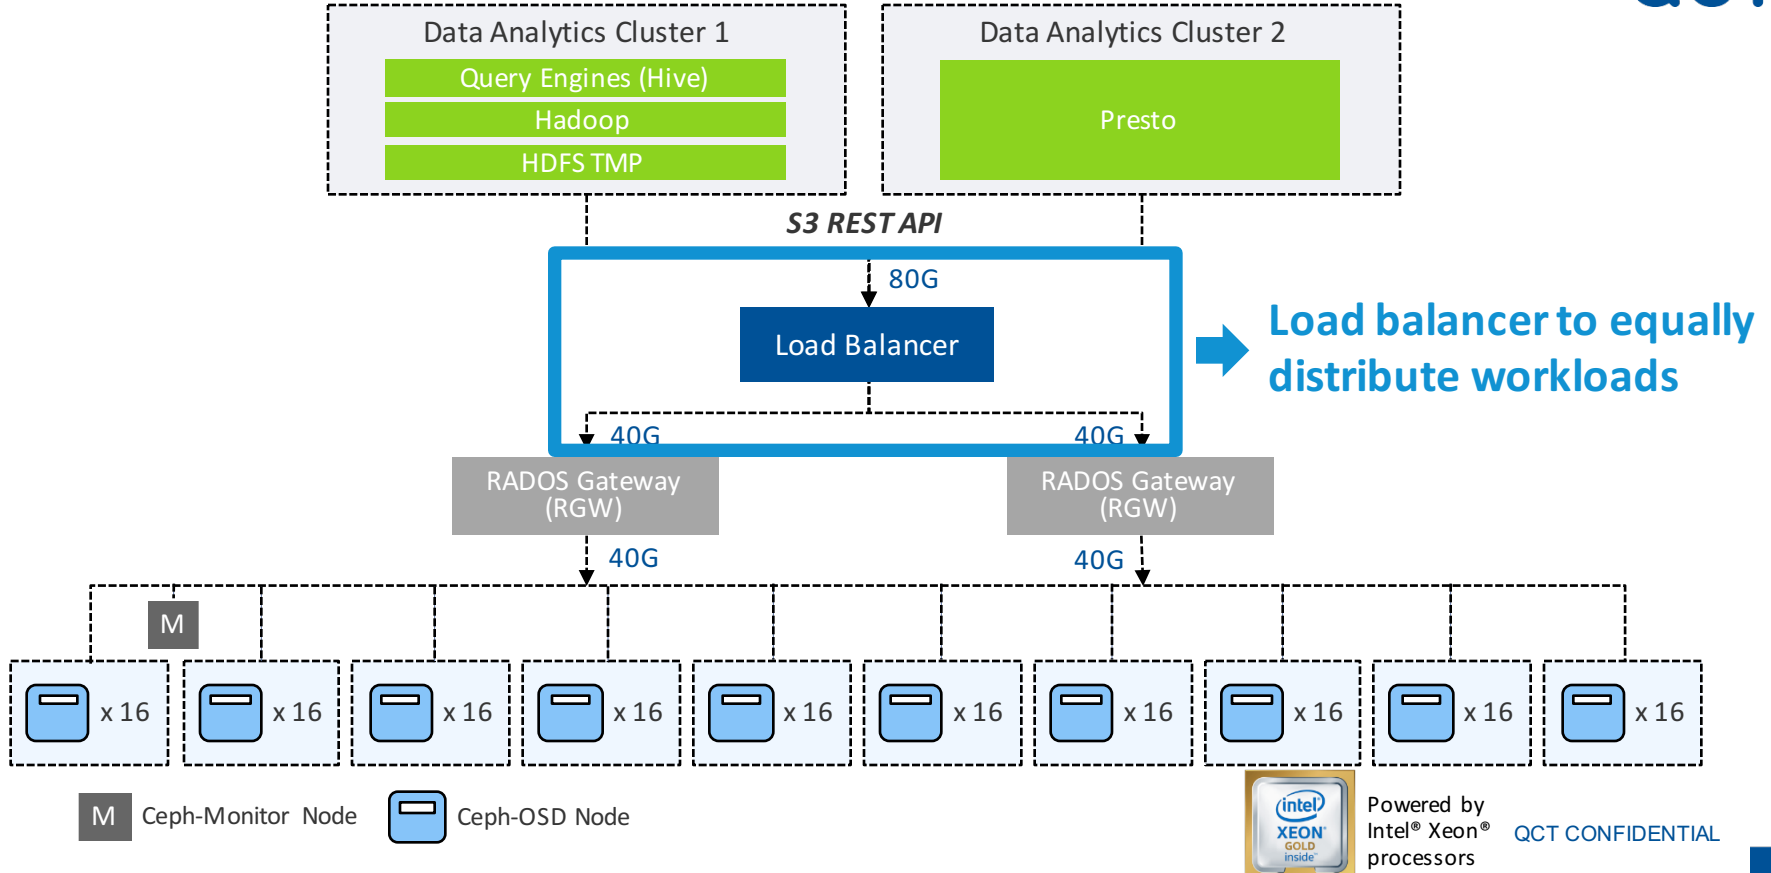
Data Lake with All-Flash Ceph Architecture

Disaggregate data analytics cluster and backend storage to provide higher flexibility



Data Lake with All-Flash Ceph Architecture

Disaggregate data analytics cluster and backend storage to provide higher flexibility



Data Lake with All-Flash Ceph Architecture

Disaggregate data analytics cluster and backend storage to provide higher flexibility

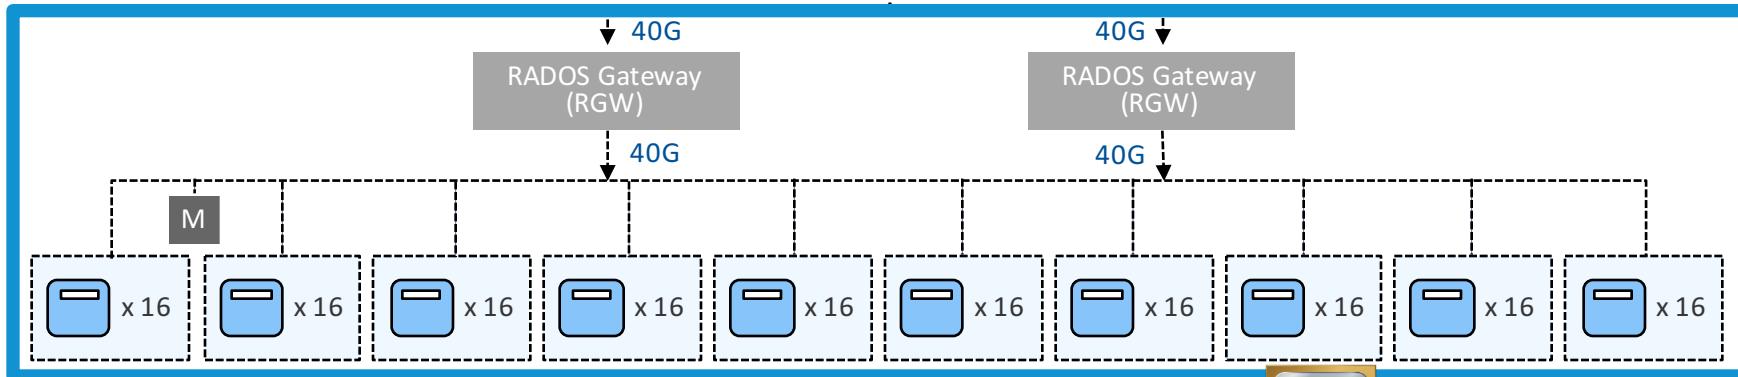


S3 REST API

80G

Load Balancer

Ceph as backend storage to scale-out on demand



M Ceph-Monitor Node

Ceph-OSD Node



Powered by Intel® Xeon® processors

QCT CONFIDENTIAL

QuantaGrid D52BQ-2U – Scale Along with Your Business

Intel Purley platform with up to 24 2.5" bays with SATA/SAS/NVMe support



Top shelf Xeon® P processor¹

Up to **10x** PCIe expansion slots

Up to **26x** hot-swap drive bays

Up to **3TB** memory capacity²

All screw-less, hot-swappable!



*As many as 24x SFF + optional extra 2x rear SSD bays
(SATA/SAS/NVMe support)*



*12x LFF + optional extra 2x rear SSD bays
(SATA/SAS/NVMe support)*



Powered by
Intel® Xeon®
processors

1. With limited conditions
2. With specific CPU

QxStor Ceph – Know your Demand, Easy to Configure

QCT QxStor Big Data Analytics Data Lake with All-Flash Ceph Solution



Powered by Intel® Xeon® processors

Throughput Optimized

For Streaming Media



QxStor RCT-200

QxStor RCT-400



D51PH-1ULH



T21P-4U

+25% In **Total Storage Capacity** available¹

-33% In sequential writing **Latency** Testing²

+50% In sequential writing **Throughput**²

Cost/Capacity Optimized

For Archiving



QxStor RCC-400



T21P-4U

Up to 560TB Per chassis

Up to 63% Cost down

IOPS Optimized

For Mission Critical App



QxStor RCI-300



D51BP-1U

Purley Available!



D52BQ-2U

+100% Improve in **IOPS** performance³

1.6M/s **Highest IOPS**³

-50% Reduce in **Latency**³

¹ SKU statistics of RCT-200 ² Test result of RCT-400 ³ Test result of RCI-300

All-Flash Ceph is preferred for data analytic workloads

NVMe is preferred both from the business and performance perspective



Business Perspective

NVMe is no longer a luxury device for enterprise with IO intensive workloads



High Performance-Optimized Storage



Cost Efficient Compared to HDD



Suitable for Mission Critical App

Performance Perspective

NVMe exhibits exceptional results on system metrics than conventional disks

CPU Utilization	Network Traffic	Disk Read Throughput	Disk Read Latency
x9.24	Incoming: x3.9	x2.81	x9.77
	Outgoing: x16.1		



Powered by Intel® Xeon® processors

QCT CONFIDENTIAL

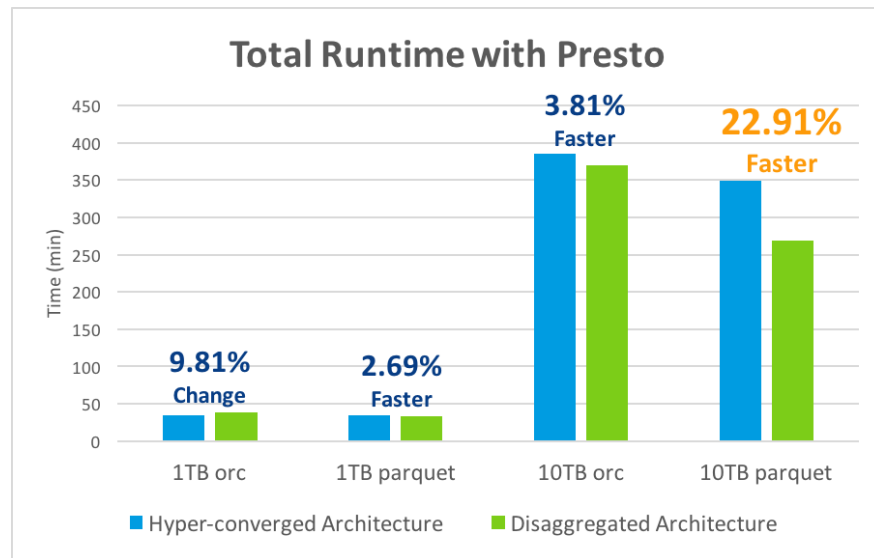
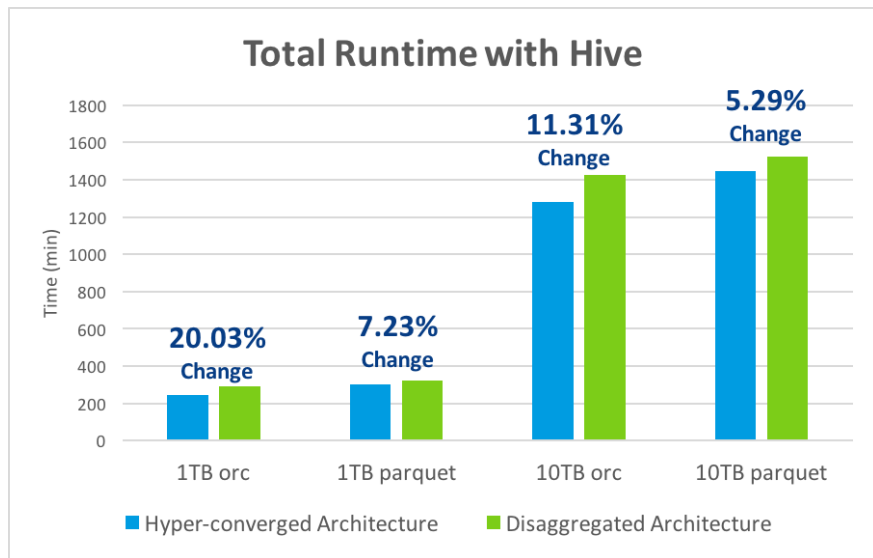


Test Result – Hive and Presto

Assured performance for data lake architecture when compared to HDFS hyper-converged architecture

Minor changes are observed in total runtime comparison between HDFS hyper-converged and Ceph disaggregated architecture using Hive.

Up to 22.91% faster in total runtime for Ceph disaggregated architecture using Presto, the effect is especially notable for large data size.



Disaggregated architecture is suitable for data analytics

Meet the demand for big data frameworks while providing higher flexibility



Scale-Out According to Need

Scaling components independently reduces cost & management complexity



Cost-Efficient Architecture

Lower storage required for data durability than HDFS or RAID based systems



Assured Performance Level

Comparable test results to hyper-converged architecture for data analytics





**Looking for
innovative cloud solution?
Come to QCT, who else?**

www.QCT.io

