WHITE PAPER

Data Center Transformation:

Why a Workload-Driven and Scalable Architecture Matters

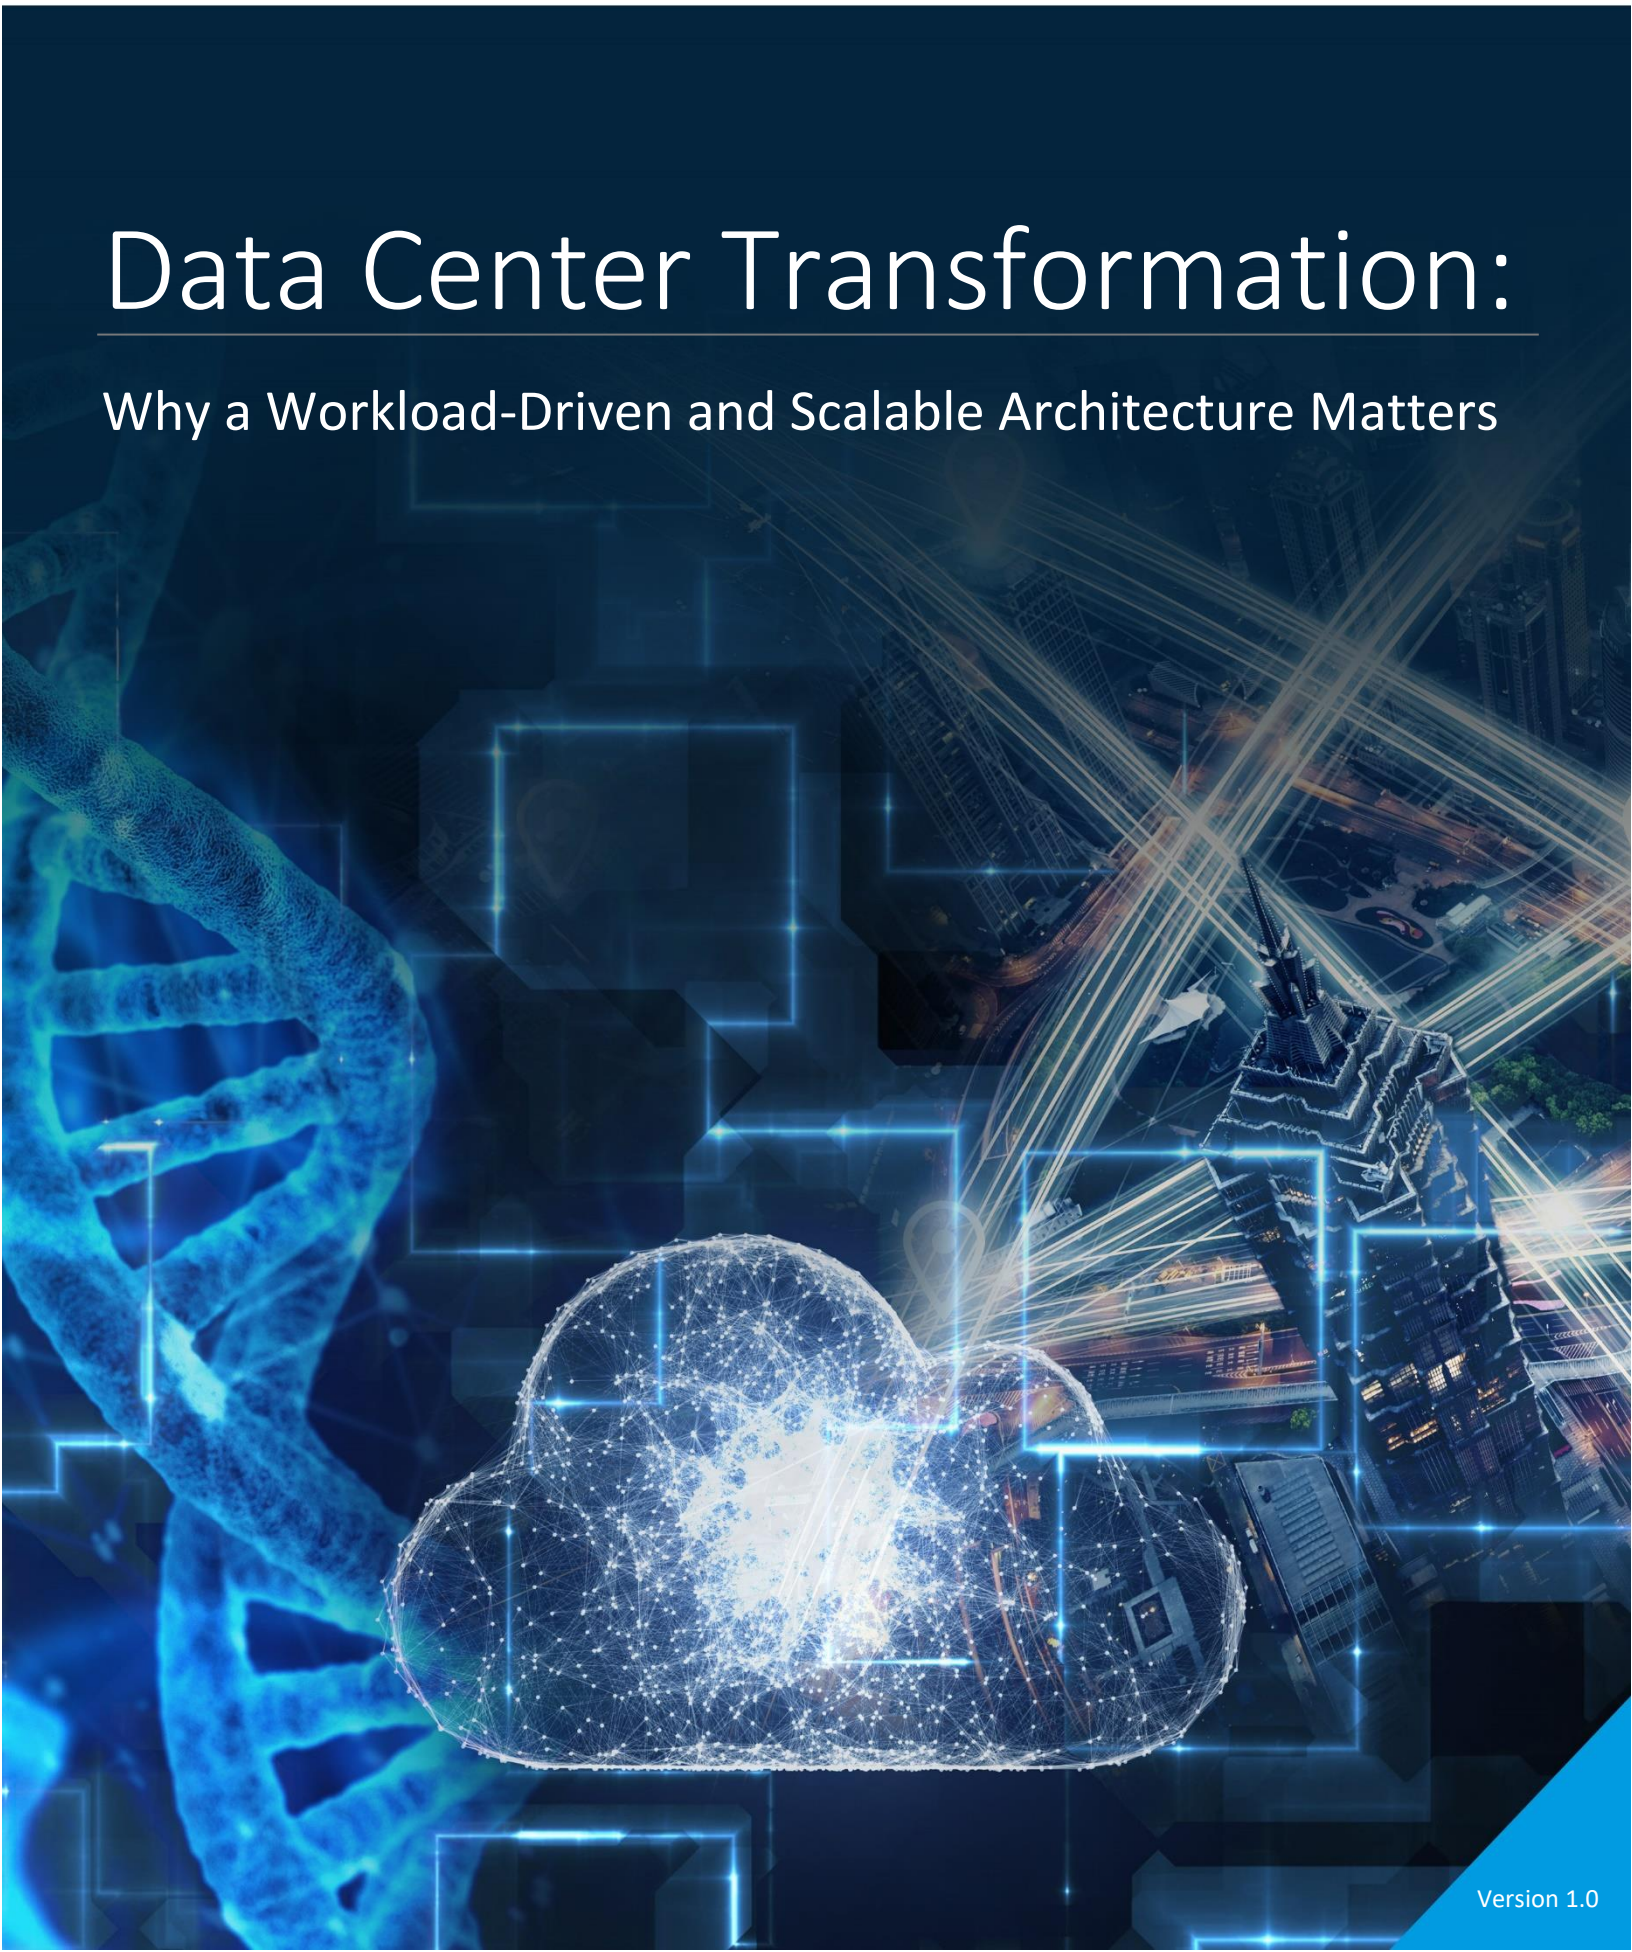

Version 1.0

OUTLINE

TABLES

FIGURES

1. Introduction

Nowadays we see that every company from global giant enterprises to small and medium enterprises (SMEs), are rethinking their digital transformation strategy. Originally, enterprises were inclined to build a data center with a proprietary system with enterprise-grade software all at one time, spending lots of money to make sure the system worked well. Thanks to open source technologies becoming more and more prosperous, parts of the commercial software stack are now available to be replaced by open source ones. Also, due to fast technology updates, enterprises are more willing to purchase on-premise systems gradually according to business growth instead of buying the type of oversized systems that companies forecasted would take 3~5 years to build and reach 100% utilization. By doing this, they can overcome technology changes and reduce TCO, which also allows enterprises to start embracing commercial/open source hybrid solutions with a scalable architecture.

2. What is QCT Platform on Demand (QCT POD)?

QCT is offering its QCT Platform on Demand (QCT POD) solution that empowers enterprises to kickstart their transformation journey. It combines advanced technology with a unique user experience to help enterprises reach better performance and gain more insights. With flexibility and scalability, QCT POD enables enterprises to address a broader range of HPC, Deep Learning, and Data Analytic demands that fulfill various applications.

QCT POD provides an on-premises rack level system with best practice hardware and software integration for specific workloads. With its in-house designed hardware and system management tools, QCT POD is pre-validated, pre-configured, and Intel Rack Scale Design (RSD) ready, which saves customers' time and resources with its rapid deployment and easy management. To reach greater flexibility and scalability, QCT POD offers common building blocks to meet business demands from different verticals, like Manufacturing, Healthcare, and High Education and Research. QCT ensures the quality and serviceability of the infrastructure, which accelerates time-to-value for enterprises.

3. QCT POD Building Blocks Overview

QCT POD integrates best-in-class hardware with software components recommended by QCT. The architecture could further be divided into 3 parts - Management Building Block, Compute Building Block and Storage Building Block - each of them connected by network fabric. With a modularized design, the configuration of each part would be different depending on the workload to fit user demands.

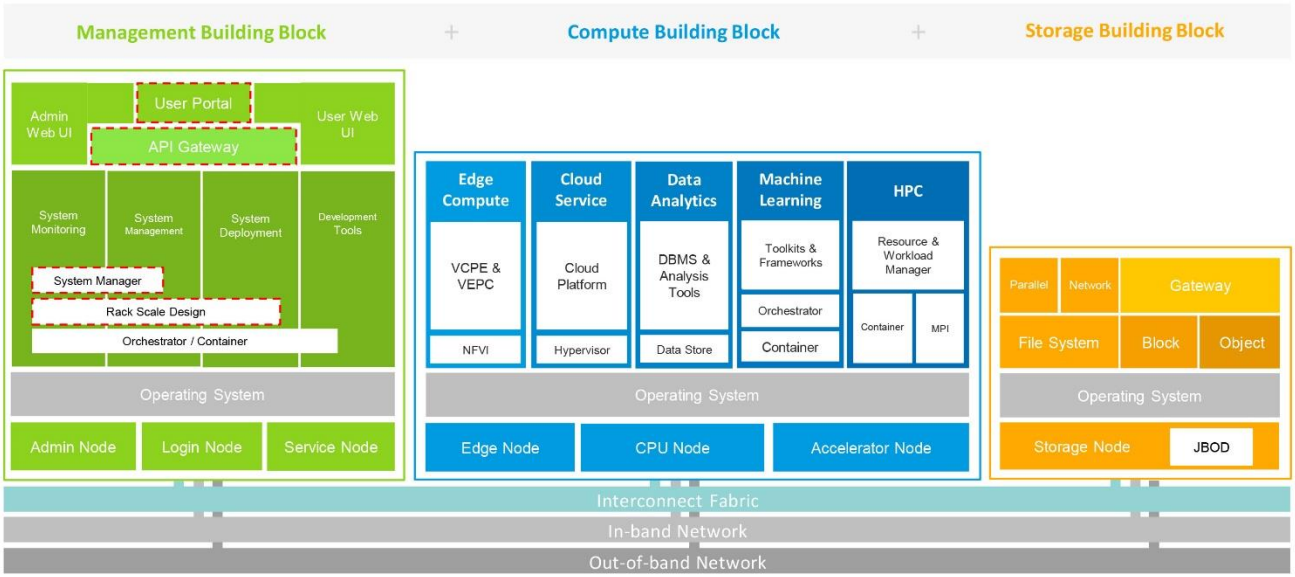

Figure 1. QCT POD Building Blocks Overview

4. Storage Building Block

QCT has partnered with a selection of leading software-defined storage companies to design its portfolio of scale-out and high-performance storage solutions. These technology partners provide the building blocks for the QCT storage offering, which includes file, object, and block storage, along with its hardware components.

Big Data trends have taught us a lot about the value and importance of information. Due to the exponential growth of data, we learned to leverage analytics to get better insights and improve performance in a variety of domains across industries. Think of sales and marketing, research and development, healthcare, energy efficiency, etc. Today, analytics have evolved far beyond Big Data with Artificial Intelligence (AI) and Machine Learning (ML) becoming a reality. With the rise of these new frontiers, AI & ML workloads are having a huge impact on storage requirements, especially for file and block storage.

Parallel file systems have become extremely popular in industries that process large amounts of data. High Performance Computing users rely on high-performance file systems such as IBM Spectrum Scale, ThinkParQ BeeGFS, and Lustre. But as AI & ML compute farms become more powerful through GPU innovations, it becomes more and more challenging for these file systems to effectively feed the GPU's and storage risks becoming a bottleneck. How can this be avoided?

QCT has partnered with ThinkParQ and Excelero to solve exactly that problem. The joint solution was designed to meet any HPC & AI storage requirements and maximize GPU ROI by ensuring unlimited data feeds.

ThinkParQ's BeeGFS parallel file system was designed for I/O intensive workloads that require massive performance and capacity scalability. BeeGFS has been deployed by hundreds of scientific and research organizations as it is ideal for demanding, high-performance workloads found in HPC and life sciences. Think of Artificial Intelligence, deep learning, cognitive computing or other data-intensive analytics.

BeeGFS customers benefit from NVMe flash as it provides lower latency and higher performance. The challenge is, however, to deploy NVMe in a responsible and efficient manner: the NVMe protocol was designed to be used locally, in-server, so when you deploy NVMe in large storage clusters under a shared file system, such as the ones that run BeeGFS, you risk having a stranded capacity or unrealized performance. By running BeeGFS on NVMesh, you can leverage the full potential of your NVMe devices with a shared namespace across your entire application cluster.

BeeGFS was designed to scale simply and granularly: by increasing the number of servers and drives in the system, performance and capacity of the file system scales out to the desired level; seamlessly from small systems to massive clusters with thousands of nodes. NVMesh scales in the very same way: you simply add capacity or performance by adding more NVMe devices to your servers or adding more servers to your environment.

By deploying BeeGFS and NVMesh on All-NVMe QCT Storage Servers customers combine the power of the fastest parallel file system with the fastest block storage. This end-to-end scale-out, high-performance storage solution brings impressive performance benefits for mission-critical workloads across industries.

Early benchmark tests for the joint solution show impressive results. See below for the MDTest benchmark:

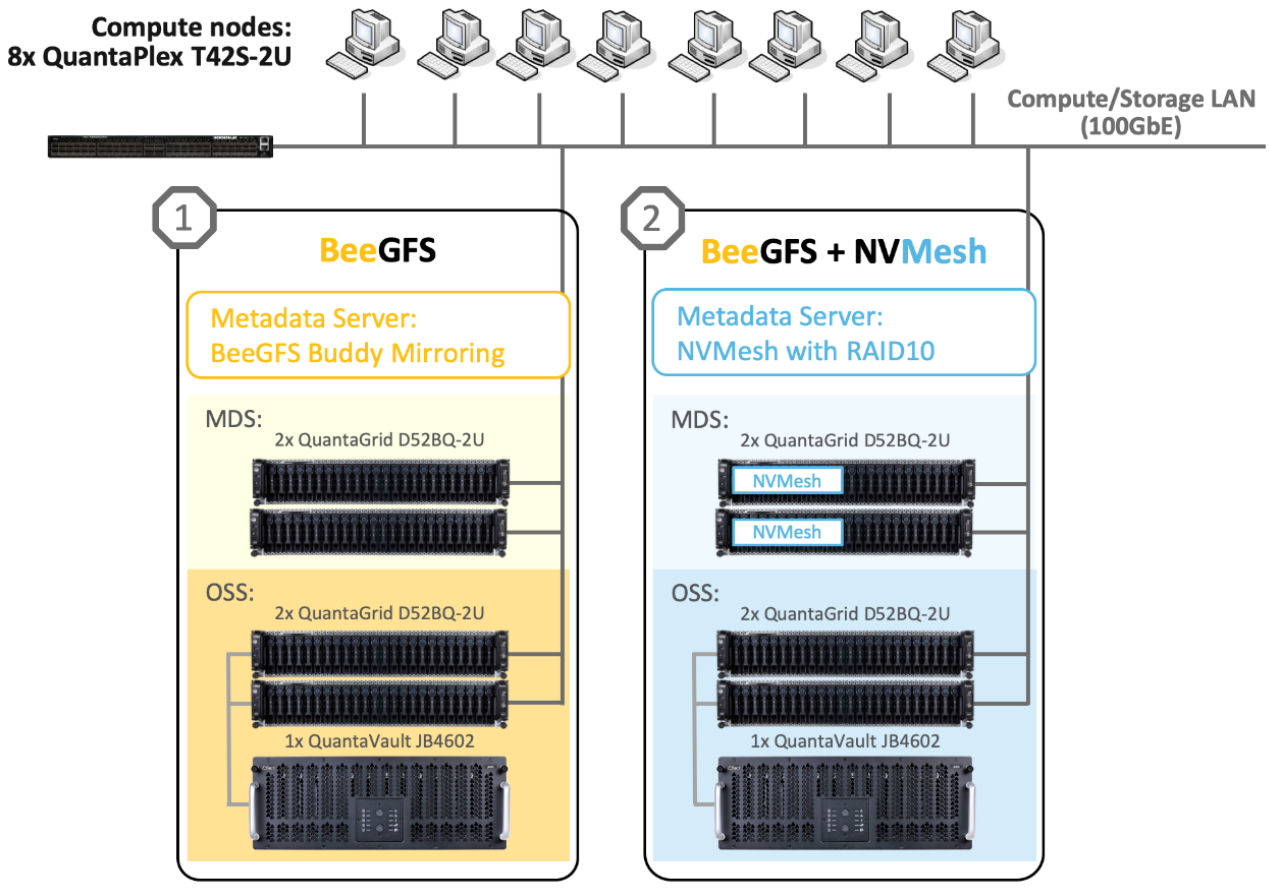

Figure 2. The architectures of MDTest benchmark

Scenario 1	Model	Qty.
MDS (BeeGFS Buddy Mirroring)	QuantaGrid D52BQ-2U w/ 1.6TB NVMe*8	2
OSS	QuantaGrid D52BQ-2U	2
	QuantaVault JB4602	1
Scenario 2	**Model**	**Qty.**
MDS (NVMesh with RAID10)	QuantaGrid D52BQ-2U w/ 1.6TB NVMe *8	2
OSS	QuantaGrid D52BQ-2U	2
	QuantaVault JB4602	1

Table 1. The configuration of MDTest benchmark scenarios

QuantaGrid D52BQ-2U *2			
Component	**Detail**	**Qty. per node**	**Total Qty.**
CPU	Intel Xeon Gold 6130 CPU	2	4
DIMM	16GB 2666MHz DDR4	12	24
Boot	480GB SSD	2	4
Storage	1.6TB NVMe	8	16
Management Network	Mellanox ConnectX®-4 MCX4421A (dual port)	1	2
Data Network	Mellanox ConnectX®-5 MCX516A-CCAT (dual port)	1	2

Table 2. The configuration of Metadata Server (MDS)

QuantaGrid D52BQ-2U *2			
Component	**Detail**	**Qty. per node**	**Total Qty.**
CPU	Intel Xeon Gold 6130 CPU	2	4
DIMM	16GB 2666MHz DDR4	12	24
Boot	480GB SSD	2	4
Management Network	Mellanox ConnectX®-4 MCX4421A (dual port)	1	2
Data Network	Mellanox ConnectX®-5 MCX516A-CCAT (dual port)	1	2
QuantaVault JB4602 *1			
Component	**Detail**	**Qty. per node**	**Total Qty.**
Disk	3.5” 14TB SAS HDD (RAID 6)	60	60

Table 3. The configuration of Object Storage Server (OSS)

The MDTest benchmark shows overall better performance for file system operations, here, in the example you can see the file stat operations per second. Scenario 2 (BeeGFS + NVMesh) shows the significant performance boost compared to Scenario 1 (BeeGFS Buddy Mirroring without NVMesh), and as we add nodes, the performance gap becomes larger. At a scale of 32 processors, the performance of Scenario 2 is 2.4x higher than that of Scenario 1.

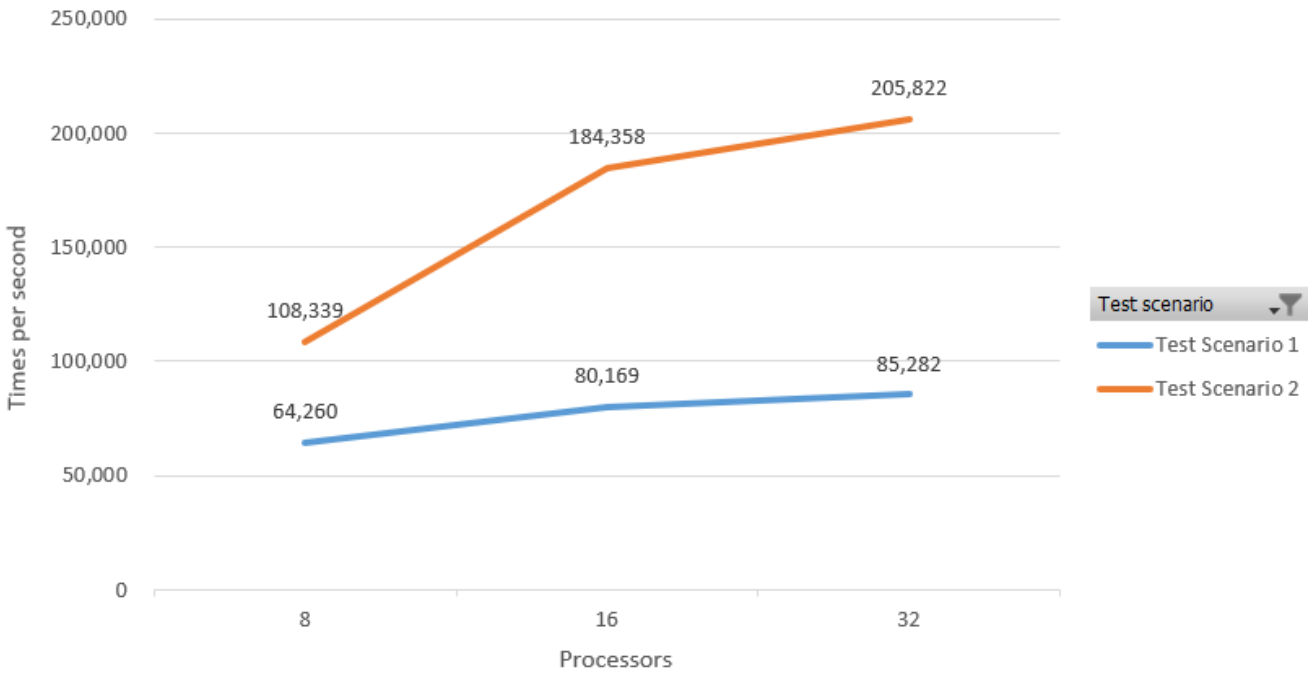

Figure 3. The performance result of MDTest benchmark

With this design, QCT is raising the bar for high-performance storage: this solution with BeeGFS and NVMesh gives users the scalability of the fastest parallel file system with the low latency and high performance of distributed NVMe. This gives customers the benefits brought by the performance of local flash with the convenience of centralized storage, while avoiding proprietary hardware lock-in and reducing the overall storage TCO.

While the new solution was primarily designed with the strong efficiency requirements of HPC & AI in mind, it is now available to customers from all industries with similar demands for high IOPS and low latency, such as finance, engineering, life sciences, energy, and many others.

5. Management Building Block

The Management Building Block plays an important role in the QCT POD, which is the place for users and administrators to access and manage the system. It includes related software tools combined with QCT hardware to provide a stable environment. With the Management Building Block, administrators can streamline the deployment process and simplify the monitoring, which highly improves efficiency for customers.

Based on RedHat Enterprise Linux or CentOS, this solution integrates open source technologies to form the software stack. To be more specific, Prometheus works as a system monitor, Alert Manager is for system management, xCAT and Ansible represent system deployment, and a wide selection of development tools is available as well. To manage a cluster is always not an easy task, but QSM (Quanta System Manager) and Intel RSD (Rack Scale Design) are the remedy to relieve these pain points. By adopting QCT POD, administrators can simply leverage QSM and RSD through a Web UI to do system management.

There is more than just the integrated software stack, on top of the management building block, there're dashboards for administrators and users' ease of use:

- Grafana: a popular and easy to use dashboard which helps administrators to manage the performance of each node
- Kibana: dashboard for administrators to audit and monitor security
- Jupyter Hub: a common dashboard for multi-users to login and gain output

QCT POD is a scalable architecture, and allows users to leverage an all-in-one server for the management building block at small scale. As the cluster expansion demand is driven by business growth, the company is flexible to scale its data center by categorizing the management building block into a login node, admin node, and service node.

6. Compute Building Block

QCT POD contains diverse categories of compute building blocks to fulfill different vertical workloads, such as HPC, ML, Data analytics, Cloud service and Edge Compute. Based on the demand from each field and industry, QCT delivers the specific hardware and software combinations to solve challenges and gain better performance for customers. For example, within the HPC compute building blocks, there could be Singularity and MPI as options for the orchestrator. On top of that, several job schedulers like Altair PBS Pro and Slurm could be the choices for users to manage their workloads. As for the ML building blocks, Kubernetes and Docker would be the base for the orchestration, and users could further utilize other frameworks like TensorFlow, Keras, and PyTorch. With a flexible building blocks design, QCT POD would be able to meet diverse demands and solve challenges for customers.

7. Learn More

See related solutions at https://go.qct.io/solutions/data-analytic-platform/qxsmart-hpc-dl-solution

See QCT partner BeeGFS at https://www.beegfs.io/content/

See QCT partner Excelero at https://www.excelero.com/

ABOUT QCT

QCT (Quanta Cloud Technology) is a global datacenter solution provider extending the power of hyperscale datacenter design in standard and open SKUs to all datacenter customers.

Product lines include servers, storage, network switches, integrated rack systems and cloud solutions, all delivering hyperscale efficiency, scalability, reliability, manageability, serviceability and optimized performance for each workload.

QCT offers a full spectrum of datacenter products and services from engineering, integration and optimization to global supply chain support, all under one roof.

The parent of QCT is Quanta Computer Inc., a Fortune Global 500 technology engineering and manufacturing company.

http://www.QCT.io

UNITED STATES

QCT LLC., Silicon Valley office
1010 Rincon Circle, San Jose, CA 95131
TOLL-FREE: 1-855-QCT-MUST
TEL: +1-510-270-6111
FAX: +1-510-270-6161
Support: +1-510-270-6216

QCT LLC., Seattle office
13810 SE Eastgate Way, Suite 190, Building 1,
Bellevue, WA 98005
TEL: +1-425-633-1620
FAX: +1-425-633-1621

CHINA

云达科技, 北京办公室（Quanta Cloud Technology）
北京市朝阳区东大桥路 12 号润诚中心 2 号楼
TEL +86-10-5920-7600
FAX +86-10-5981-7958

云达科技, 杭州办公室（Quanta Cloud Technology）
浙江省杭州市西湖区古墩路浙商财富中心 4 号楼 303 室
TEL +86-571-2819-8650

JAPAN

Quanta Cloud Technology Japan 株式会社
東京都港区芝大門 2-5-8 芝大門牧田ビル 3F, 105-0012
TEL +81-3-5777-0818
FAX +81-3-5777-0819

GERMANY

Quanta Cloud Technology Germany GmbH
Hamborner Str. 55, 40472 Düsseldorf
TEL +492405-4083-1

TAIWAN

雲達科技（Quanta Cloud Technology）
桃園市龜山區文化二路 211 號 1 樓
1F, No. 211 Wenhua 2nd Rd., Guishan Dist., Taoyuan City 33377, Taiwan