WHITE PAPER

A SOFTWARE DEFINED STORAGE SOLUTION WITH GIGABYTE & BIGTERA VIRTUALSTOR™ SCALER

GIGABYTE™ BIGTERA

OVERVIEW

This white paper presents an example architecture and performance benchmarking for a software defined storage system with 192 TB of usable capacity, using GIGABYTE server hardware and Bigtera VirtualStor™ Scaler software.

SOFTWARE DEFINED STORAGE

Enterprises are experiencing rapid storage requirement growth due to the expansion of the digital universe. IDC forecasts that by 2025 there will be more than 163 zettabytes (or 163 trillion gigabytes) of data in the world[1]. While the amount of data that businesses are generating is ever increasing, IT budgets are not keeping pace with what traditional storage infrastructure would cost to facilitate handling this data. Traditional storage solutions **scale up** which requires heavy investment and massive over provisioning.

[1] **IDC,** Data Age 2025: The Evolution of Data to Life-Critical, April 2017

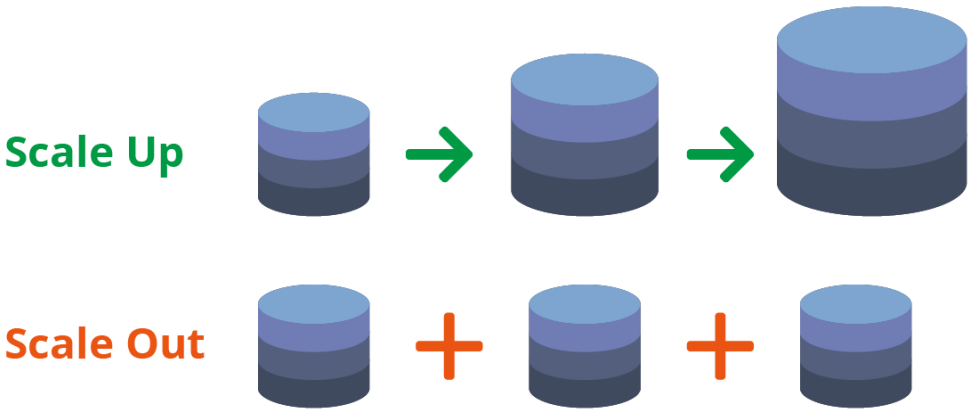

A Software Defined Storage (SDS) solution introduces enterprises to a far more effective and efficient scale-as-you-grow infrastructure model. A SDS system **scales out**, so enterprises can add capacity to an existing system whenever they need it, eliminating over-provisioning. Additionally, the virtualized and disaggregated nature of SDS means that new and legacy storage systems and different storage types such as file, block and object storage can be unified together in a single platform, eliminating "storage islands" that increase the complexity of managing infrastructure as well as wasting time and resources for data duplication.

BIGTERA VIRTUALSTOR™ SCALER

Bigtera VirtualStor™ Scaler is an ideal choice for a software defined storage solution. Not only does it offer all the advantages of an open source software defined storage platform, but also adds many innovative features in **performance** (SSD acceleration technology with **Caching within Object Storage Device**), **unified storage**, **efficiency** and **resilience**.

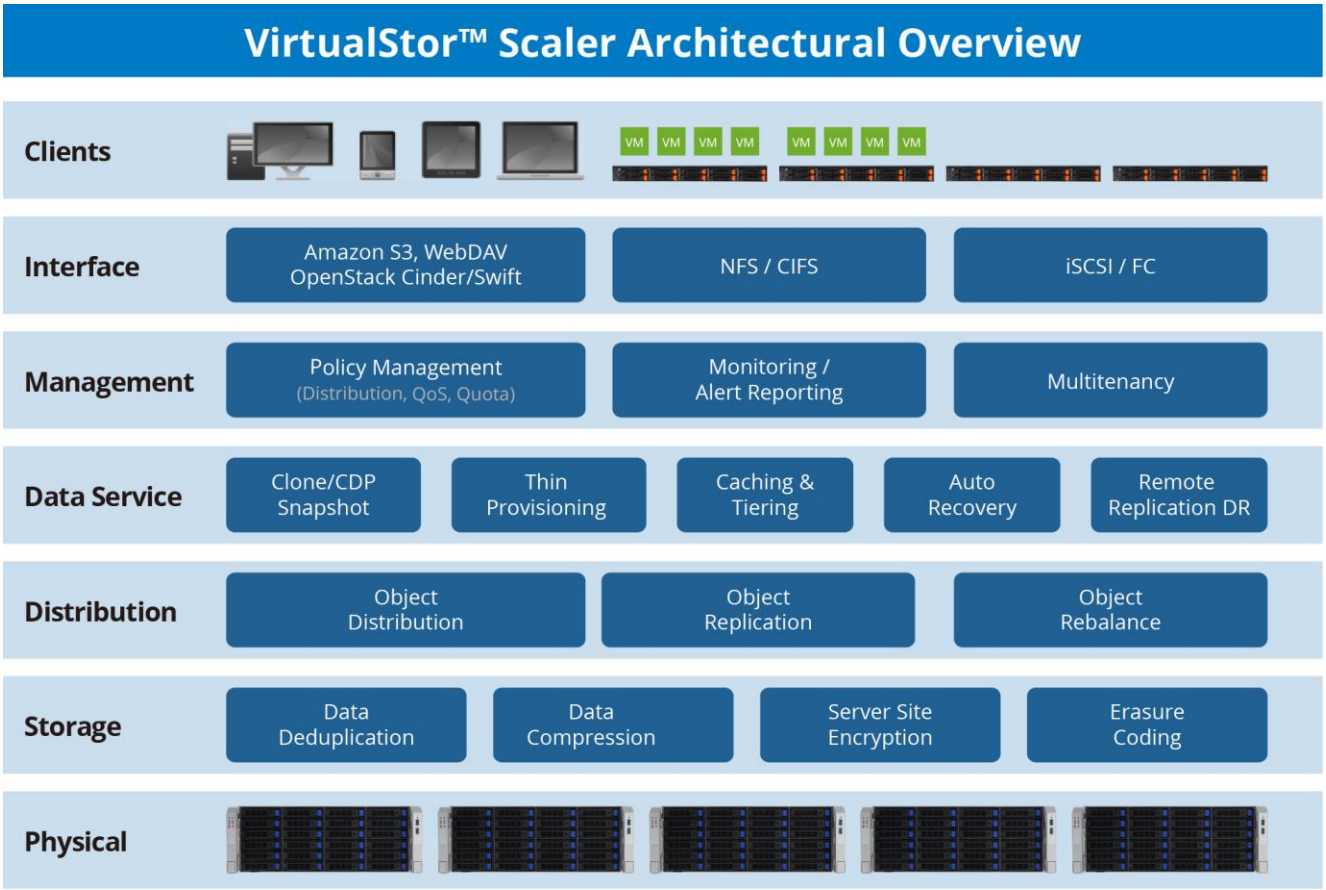

Caching within Object Storage Device (OSD)

As the innovator in software-defined storage technology, Bigtera has revolutionized the storage industry by a new approach to leverage high-performance solid state drives (SSD), optimizing performance for random input and output operations (I/O). This unique backend engine is called **BigteraStore**.

BigteraStore provides SSD acceleration technology such as data caching and merging small I/Os into sequential I/Os to further improve IOPS. Unlike other software-defined storage platforms that put many SSD OSD together into a cache pool, BigteraStore provides an OSD caching ability that supports random I/O caching by taking advantage of high-performance SSDs. In addition, the larger block sizes of sequential I/O will skip the SSD and directly read/write to the hard drives. Usually SSD is by magnitude faster than the spinning hard drives for random I/O, but with less advantage in sequential

read/writes. Therefore, VirtualStor™ Scaler can intelligently serve random I/O and sequential I/O differently, to optimize overall performance. Also, thanks to this intelligent design that will skip certain I/O through SSD, VirtualStor™ Scaler can optimize SSD usage to significantly reduce SSD write wear, so as to increase your SSD lifespan endurance and decrease your CAPEX compared to other software-defined storage platforms.

Unified Storage - VirtualStor™ Scaler provides a unified storage platform so companies do not need to choose between the types of storage they need. As more appliances are added, the appliances seamlessly become part of a single massive decentralized storage entity. VirtualStor™ Scaler can be partitioned into storage of any type, accomplished by abstracting the storage hardware from the control layer, and allowing for network attached storage (NAS) and storage area networks (SAN) that can run simultaneously. These storage types are supported by several storage protocols: SAN (iSCSI / FC), NAS (CIFS / NFS), Object Storage (Amazon S3 / OpenStack Swift) and OpenStack Cinder RBD.

High Efficiency - VirtualStor™ Scaler can use RAID-5 together with two replica copies or N+M Erasure Coding, so as to provide more efficient space utilization than other software defined storage with 3 replica copies. Administrators can also assign various services on Bigtera's unique multi-tenant storage technology "Virtual Storage" to virtually extend the available space, and enable compression and data deduplication for backup or archiving. VirtualStor™ Scaler automates efficient optimization of storage resources in several ways. First, Thin Provisioning functionality provides resources just as they are needed. Secondly, storage resources are balanced across storage nodes so no single node carries more than their fair share of the load.

Resilience - VirtualStor™ Scaler data availability functions include data replication, erasure coding, self-repairing, and RAID features. Erasure coding offers administrators an alternative to data replication, to ensure that there is no single point of failure for any of the data blocks. VirtualStor™ Scaler uses round-robin DNS and IP takeover services. Round-robin DNS uses a list of IP addresses for workload balancing , if any of the appliances encounter issues, the remaining appliances take over application and workload services seamlessly by taking over the IP of the appliance that encounters issues. VirtualStor™ Scaler can protect data stored with Amazon S3 API using Intel® AES-NI encryption technology. Encryption can be enabled for critical data or applications, while data that has a lower level of confidentiality can be left unencrypted.

Easy Data Migration & Management - VirtualStor™ Scaler can seamlessly migrate existing data from a customer's legacy storage with minimum downtime, and then consolidate the legacy storage in to the storage pool. And in addition to optimizing SSD usage to significantly reduce SSD write wear, VirtualStor™ Scaler can also notify the administrator when an SSD life is running out, and predict storage capacity and performance. These features help an administrator to plan ahead and decreases management headaches.

GIGABYTE SERVER HARDWARE

GIGABYTE servers are optimized for a software defined storage system. Using a standard x 86 architecture allows interoperability with all major and well-known storage software & applications for easy deployment and scale-out. GIGABYTE servers also all feature rich expansion capacity and configuration flexibility on each model. This allows storage system builders to add on capacity (storage drives, memory etc.) as needed only, with no over-provisioning necessary. Our servers also support flexible networking connectivity choices for different storage needs – high speed options for high-performance hot storage, or cost efficient options for archiving and warm / cold storage. For this whitepaper we have used the following two GIGABYTE server products for our software defined storage system architecture:

COMPUTE NODES: H261-H61

Our Intel® Xeon® Scalable 2U 4 Node Density Optimized Server has an innovative architecture and excellent performance numbers to help you benefit from higher space utilization, enhanced thermal performance, reduced power consumption and better serviceability.

- 2U - 4 nodes with rear access to the node trays
- Intel® Xeon® Processor Scalable Family
- 8 x LGA 3647 sockets
- Intel® C621 Express Chipset
- 6-Channel RDIMM/LRDIMM DDR4, 64 x DIMMs
- 4 x dedicated management ports
- 1 x CMC global management port
- 8 x 2.5" U.2, 16 x SATA/SAS hot-swappable HDD/SSD bays
- 8 x low profile PCIe Gen3 expansion slots
- 4 x OCP Gen3 x16 mezzanine slots
- Aspeed® AST2500 remote management controller
- 2200W 80 PLUS Platinum redundant PSU

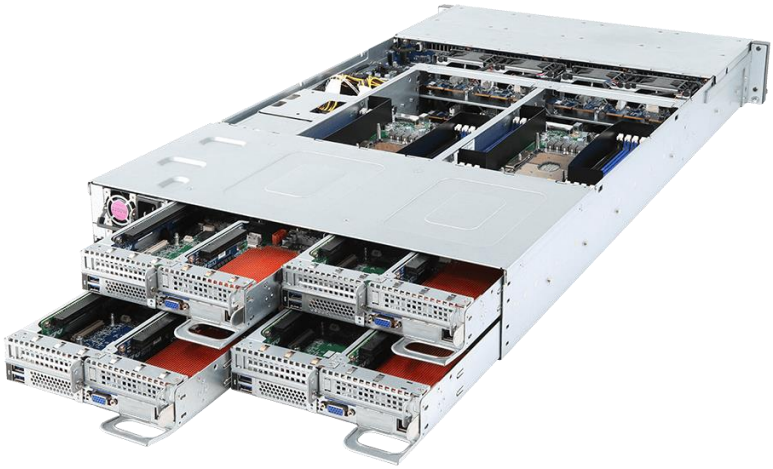

STORAGE NODES: D120-C21

Our Intel® Xeon® D-series 1U Storage Server can support 16 HDD or SSD devices. This dense 1U unit offers users the right balance between performance, cost-per-gigabyte, building block size, and failure domain size.

- Intel® Xeon® processor D-1541
- 4 x DIMM slots, dual channel, DDR4 2400/2133/1866/1600MHz
- 16 x 3.5" HDD slots
- Supports SATA III 6Gb/s via LSI SAS3216 controller
- Supports software RAID 0,1,1E,10
- 2 x 10GbE SFP+ LAN ports (Intel® X552)
- 2 x GbE LAN ports (Intel® I210)
- Aspeed® AST2400 management controller
- 1 x 400W 80 PLUS Gold 100~240V AC PSU

TEST CONFIGURATION

This section presents the configurations of the Bigtera VirtualStor™ Scaler test cluster. We designed Bigtera VirtualStor™ Scaler in a 4-node cluster as shown in the diagram below.

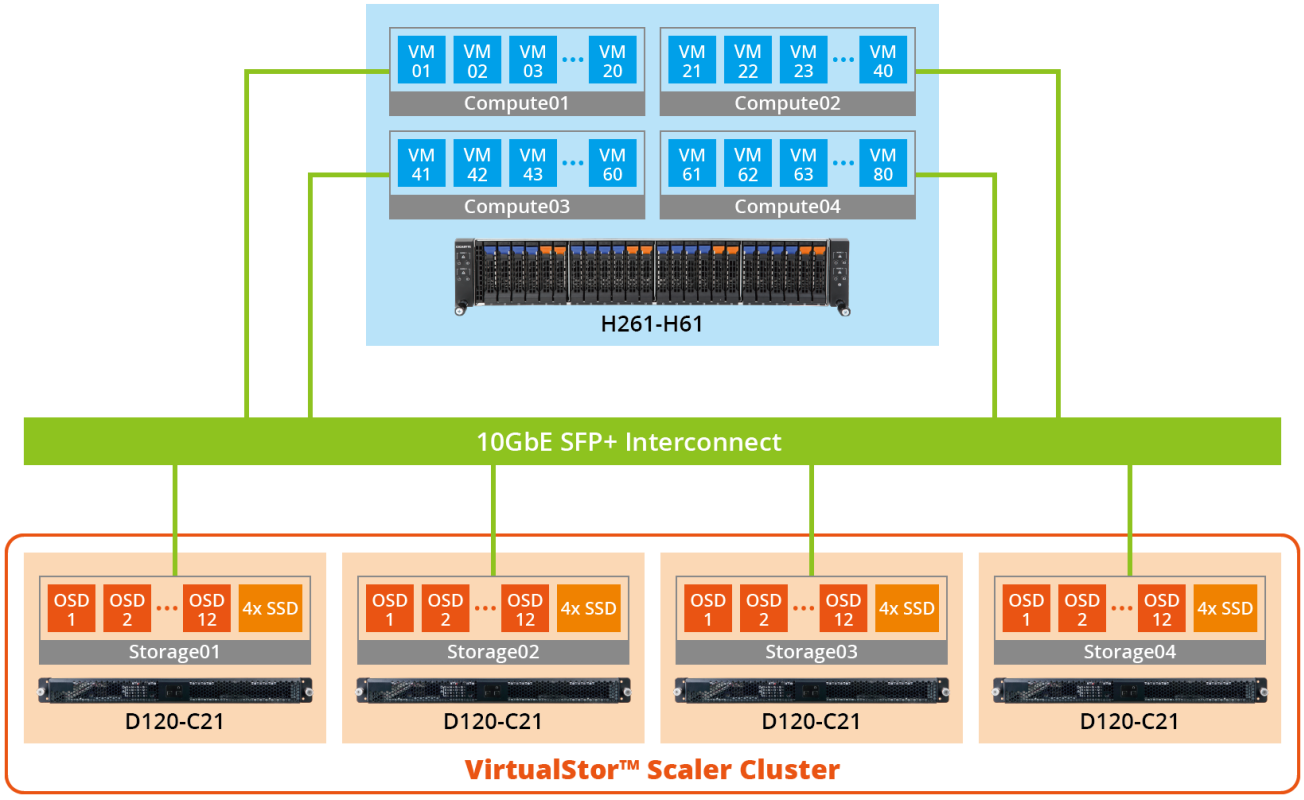

Storage Nodes

Four GIGABYTE D120-C21 storage servers were used to create four storage nodes. Each storage node had 12 HDDs and four SSDs, and one 10GbE link to the network switch. There were a total of **48 OSDs** with **192 TB of usable capacity** in the VirtualStor™ Scaler cluster. The table below summarizes the details of the storage node specifications.

Storage Nodes	
Server	GIGABYTE D120-C21
CPU	1 x Xeon D-1541 8 cores 2.1GHz
Memory	128GB DDR4
NIC	Onboard 2 x 10GbE SFP+
Disks	1 x M.2 4 x 960GB SSD 12 x 8TB 7200RPM SATA HDD

Compute Nodes

One GIGABYTE H261-H61 (2U 4 node Density Optimized Server) was used for the test client (compute nodes). Twenty virtual machines (VM) instances were created on each physical compute node to emulate the concurrent access cases. Each VM had one 2 GB RADOS (Reliable Autonomic Distributed Object Storage) block device (RBD) disk. The benchmarking tool **fio - Flexible IO Tester** (http://git.kernel.dk/?p=fio.git) was used to measure the aggregated performance of Bigtera VirtualStor™ Scaler nodes. The table below summarizes the details of the compute node specifications.

Compute Nodes	
Server	GIGABYTE H261-H61 (2U 4 nodes)
CPU	2 x Xeon Scalable Silver 4110 8 cores 2.1GHz
Memory	128GB DDR4
NIC	2 x 10GbE SFP+ (Add On Card)
Disks	1 x SATADOM 128GB 2 x NVMe SSD 400GB 4 x SATA HDD 2TB

RESULTS - THROUGHPUT PERFORMANCE

Performance of throughput was evaluated for sequential I/O workloads of a large block-size. All the tests were run with two-replicas. Journal and cache data were stored in four SSDs and OSD data was stored on 12 HDDs. Throughput and latency across the tests was measured by **fio**. Associated CPU, memory and network utilization was also observed.

Throughput Close to Hardware Limit

For **sequential reads of a large block size (1MB)**, the four node VirtualStor™ Scaler cluster keeps the consistent performance gains at **4003 MB/s**. On the other hand, for **sequential writes of large block size (1MB)**, the four node VirtualStor™ Scaler cluster kept consistent performance gains at **1250 MB/s**. Sequential read / write performance was close to the raw hard drive performance, which means the throughput is optimized for the hardware.

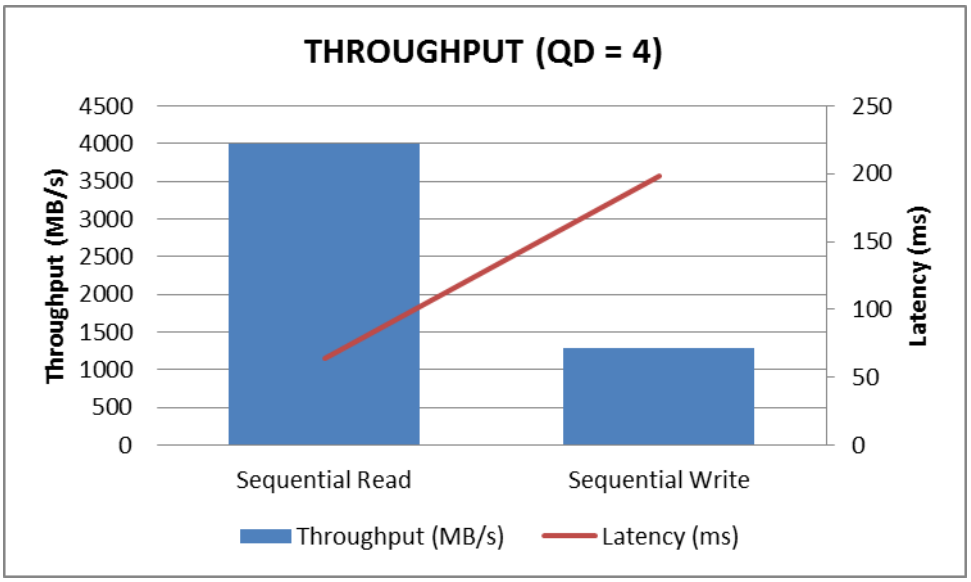

	Read Throughput	Write Throughput	Read Average Latency	Write Average Latency
QD=4	4003 MB/s	1290 MB/s	64.2ms	198.4ms

RESULTS - IOPS PERFORMANCE

Performance of IOPS (input and output operations per second) was evaluated for random I/O workloads of small block-size. All the tests were run with two replicas. The same configuration as the previous test was used: journal and cache data were stored in four SSDs, and OSD data was stored on 12 HDDs. IOPS and latency across the tests was measured by **fio**. Associated CPU, memory and network utilization was also observed.

Random Read 250K IOPS

For **random reads of small block size (4KB)**, the four node VirtualStor™ Scaler cluster peaked at **253266 IOPS** with 80 concurrent VMs. For **random writes of small block size (4KB)**, the four node VirtualStor™ Scaler cluster maintained a consisten This white paper presents an example architecture and performance benchmarking for a software defined storage component of a private cloud, using GIGABYTE server hardware and Bigtera VirtualStor™ Scaler software.
t performance around **28266 IOPS** with 80 concurrent VMs. CPU utilization was almost full during the test, which means it has already reached the limit of computing power.

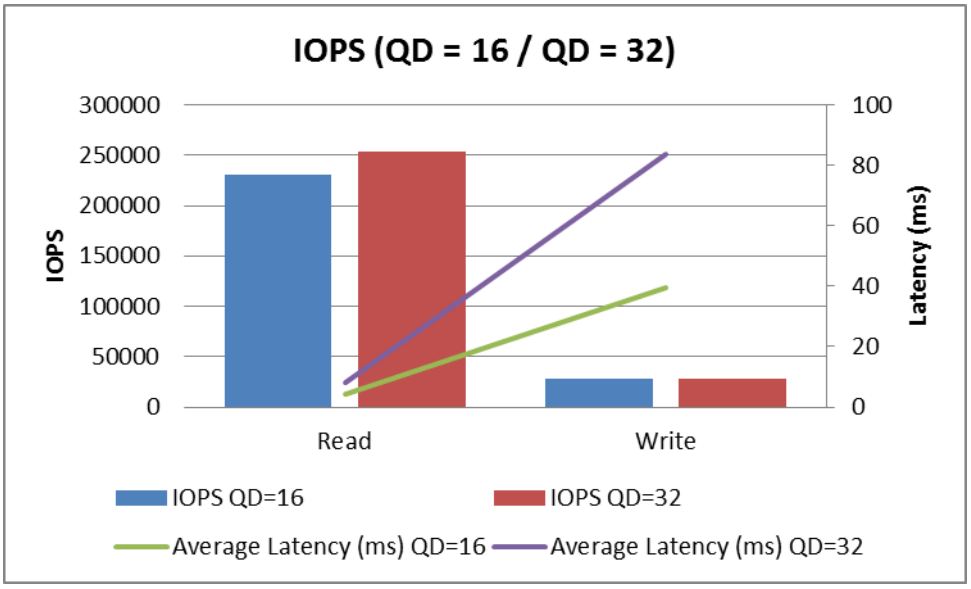

	Read IOPS	Write IOPS	Read Average Latency	Write Average Latency
QD=16	230847	28266	4.435ms	39.565ms
QD=32	253266	28202	8.089ms	83.554ms

SUMMARY

The reference architecture outlined in this document offers optimized throughput and IOPS with a cost-effective hardware platform. In keeping with the spirit of software-defined infrastructure, the architecture's hardware and software can be flexibly customized to fit various demands. For example, if the customer needs more storage capacity, additional storage nodes can be installed, or hard disks of higher capacity can be replaced in existing nodes. If the customer needs greater performance, traditional spinning HDD can be replaced with SSD, and additional storage nodes also can help.

It's time for modern organizations to eliminate the pain points associated with traditional storage. Legacy storage infrastructures not only fail to meet business needs, but also unnecessarily increase IT costs. A software defined storage system, using Bigtera VirtualStor™ Scaler and GIGABYTE server hardware, offers a high-performance, scale-out and cost-effective storage virtualization platform, satisfying business needs without pain.

CONTACT INFORMATION

GIGABYTE™ http://b2b.gigabyte.com

BIGTERA http://www.bigtera.com