SOLUTION BRIEF

Intel® Storage Builders

StorPool Storage
Intel® SSD DC S3510 Series
Intel® Xeon® Processor E3 and E5 Families
Intel® Ethernet Converged Network Adapter X710 Family

Extremely Fast Distributed Storage for Cloud Service Providers

StorPool's 468,000 IOPS, 3.5 GB/s, all-flash storage solution is powered by Intel® technologies

Introduction

With the ubiquity of cloud services and fast Internet services today, customers expect higher and higher performance and have less tolerance for delay caused by latencies in access to data. Cloud Service Providers (CSPs) face the critical decision trade-off between accelerating storage with an all-flash array at increased costs or to curb performance and keep costs low. The former can potentially reduce a CSP's profit margins and the price competitiveness. But, with StorPool's Distributed Storage software on an Intel® architecture-based platform, companies can deliver high-speed, low latency data access while containing costs.

Built with Intel® technologies—including Intel® Xeon® processor E3 and E5 families, Intel® Solid State Drive Data Center (Intel® SSD DC) family of storage devices, Intel® Ethernet Converged Network Adapter, and StorPool software, a Software Defined Storage (SDS) system can deliver up to 8.7 TB of usable storage with 450,000+ IOPS[1] at sub-millisecond latencies, as shown in testing described in this brief.

Scalability, Flexibility, and Performance in an SDS Solution

StorPool is next-generation, high-performance data storage software for enterprise and cloud. It pools the attached local storage media to create a single resource of parallel, shared block storage that scales linearly with each added server or storage device up to multiple terabytes. StorPool's performance is enabled by the following:

- Extreme high efficiency—The storage software stack has been completely re-engineered from the ground up for high-performance, low latency, and redundancy. StorPool uses its own copy-on-write on-disk format, RDMA network protocol, client, etc.—all optimized for fast and efficient distributed storage. Thus, StorPool operates at near 100% efficiency, delivering all the performance available from the underlying drives and network.
- Distributed, parallel architecture—It is a distributed system, which works with many pieces of the data in parallel. This delivers high performance, but also high reliability, and balances the load across storage nodes. Performance increases with every added drive or server.

StorPool
DISTRIBUTED STORAGE

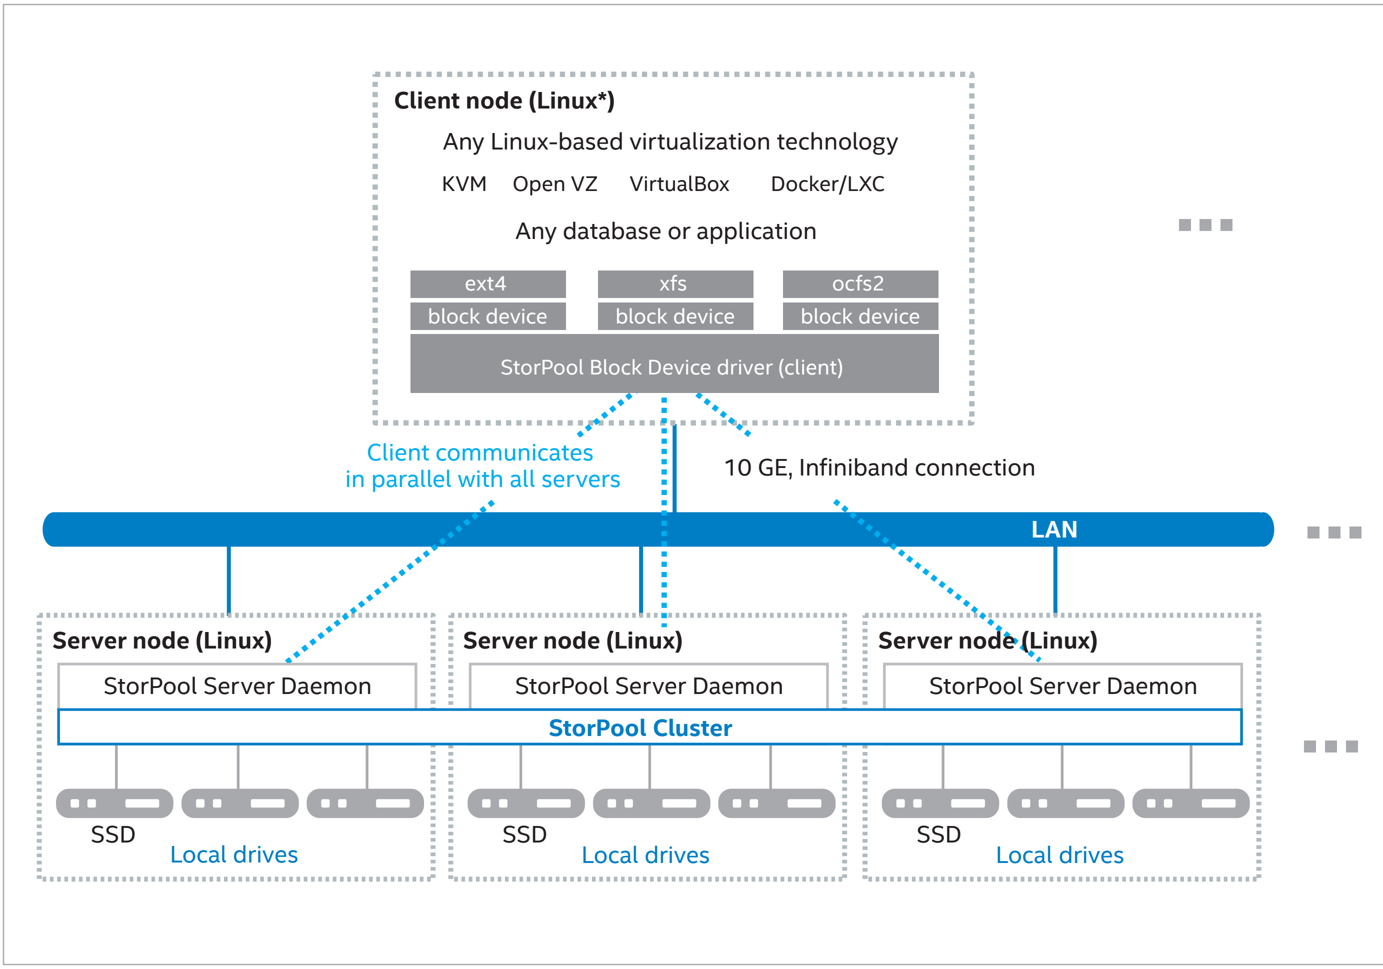

StorPool architecture: Logical network diagram.

Architecture Overview

StorPool runs on a cluster of industry high-volume servers running GNU/Linux* in a fully-distributed, shared-nothing architecture (Figure 1). Because it uses very few server resources (CPU, RAM)—typically about 5-10%—it leaves the remaining 90-95% for compute work-loads on the same nodes.

StorPool presents standard block devices to the OS. One or more volumes can be created through StorPool management interfaces, and multiple copies (replicas) of the data are written synchronously across the cluster to help ensure redun-dancy. Users can set the desired number of replication copies.

To scale capacity, IT simply adds storage media to the servers or adds servers to the cluster. StorPool does not impose any strict hierarchical storage structure that links and reflects to the underlying disks. It simply creates a single pool of data storage (global namespace), that utilizes the full capacity and performance of a defined set of commodity drives.

StorPool natively supports Linux with KVM*, LXC*, LVM*, Docker*, and any other virtualization technology compatible with the Linux storage stack. The StorPool cluster appears as a standard Storage Area Network (SAN) to other hypervisors/operating systems. The software is inte-grated with OpenStack*, CloudStack*, and OpenNebula*, and it supports OnApp*, libvirt* and Proxmox*, as well as custom cloud management solutions.

StorPool is compatible with many file systems, such as ext4 and XFS, and with any system designed to work with a block device. These include databases and cluster file systems, like OCFS* and GFS*. In the case of VMware* and Windows*, StorPool appears as one large block, and it is formatted with VMFS, NTFS, or FAT.

Intel® SSD DC Family Accelerates StorPool SDS

Optimized for the reliability, resiliency, and consistent performance that data centers demand, the Intel SSD DC family (SATA) of storage devices deliver the high throughput and low latency that today's high-performance workloads require in enterprise and cloud. This class of SSDs from Intel accelerates data center performance with read/write throughput speeds up to 550/520 MB/s and 4kB random read/write IOPs/second up to 85,000/45,000.[2] Applications benefit from 55 µs typical latency with maximum read latencies of 500 µs 99.9 percent of the time.[2] Combining performance with low typical active power (less than 6.9 watts), the Intel SSD DC family for SATA improves data center efficiency with superior quality of service, reduced energy costs, and is ideal for server or application upgrades.

Intel® Xeon® Processors E3 and E5 Families

Intel Xeon processors E3 and E5 families are architected for today's data centers—both cloud and enterprise. Each generation of Intel Xeon processor E5 family offers more cores, enhanced technologies, and increased performance. These data center workhorse server processors are optimized for orchestration, with enhanced features, such as Intel® Resource Director Technology, which offers cache monitoring and allocation technology, code and data prioritization, and memory bandwidth monitoring. With accelerated cryptography performance, Intel Xeon processor E5 family enables encrypted data to move fast over secure connections.

Intel Xeon processor E3 family delivers reliability and performance designed for small-footprint and high-density servers at a low cost.

Intel® Ethernet Converged Network Adapter X710

The Intel Ethernet Converged Network Adapter X710 addresses the demanding needs of the next-generation agile enterprise and cloud data centers, by providing unmatched features for both server and network virtualization, flexibility for LAN and SAN networks, and proven, reliable performance. The 10/40 Gbps X710 adapter family delivers superior performance with a theoretical throughput of 80 Gbps bidirectional throughput (quad adapter required).

Proven Price-Performance Efficiency

In a recent CSP installation, a StorPool cluster was built with a usable space of 8.7 TB of all-flash Intel SSD DC S3510 (SATA) devices at an estimated total end-user cost over three years of $35,000. The costs included $15,000 in hardware (except networking fabric) and $750 per month software licensing, which included 24/7 enterprise-grade support and proactive monitoring of the software.[3]

The installation was benchmarked to characterize the performance of the installation. It delivered impressive peak storage performance of 468,000 IOPS[1] with latency under load of 0.2 to 0.4 ms.

Performance Testing and Results

For the performance benchmark, five client hosts exercised the storage system. The test configuration is shown in Table 1.

CPU	MEMORY	STORAGE	NETWORK ADAPTER	OS	StorPool SW
Storage Servers: Dell* R330: 3X with the following configuration in each					
Intel® Xeon® processor E3-1230 v5	2X 16 GB DDR4 UDIMM @ 2133 MT/s	4X Intel® SSD DC S3510 @ 1.6 TB (6.4 TB per server; 19.2 TB total)	2X Intel® Ethernet Converged Network Adapter X710 Family @ 10 Gbps	Linux* 3.10.0-514.2.2.el7.x86_64	16.01.248.9d12c0a
2X IBM/Blade Networks G8124 - 24-port 10GbE switch; Direct-attach SFP+ cables					
Clients (hypervisor hosts): Dell* R630 servers: 5X with the following configuration in each					
2X Intel® Xeon® processor E5-2620 v4	16X 16 GB DDR4 RDIMM @ 2133 MT/s; 256 GB total		4-port Intel Ethernet Converged Network Adapter X710 Family @ 10 Gbps (only two ports used for storage)	Linux 2.6.32-642.11.1.el6.x86_64	

Table 1. Test System Configurations.[4]

Test System and Parameters

- Testing was performed from all five client hosts in parallel to measure total system throughput.
- Testing was performed on the all-SSD storage pool with two copies and usable capacity of 8,128 GB (~8.73 TB).
- Testing was performed on five 100 GB volumes.
- Volumes were configured to store two copies on SSDs, taking total approximate 1,100 GB of raw space (including copies and protection). This allocation policy provided a maximum of 8,128 GB usable space from the available SSDs.
- Tests were performed from five initiators in parallel to measure total system throughput.
- Performance was measured using FIO, version 2.0.13, AIO, Direct.

Results

Testing revealed high random and sequential read and write IOPs with high throughput (Figure 2). Test data is shown in Table 2.

Latency consistently remained less than 1 ms, with many tests exhibiting values below 0.5 ms as shown in Table 3 and Figure 3.

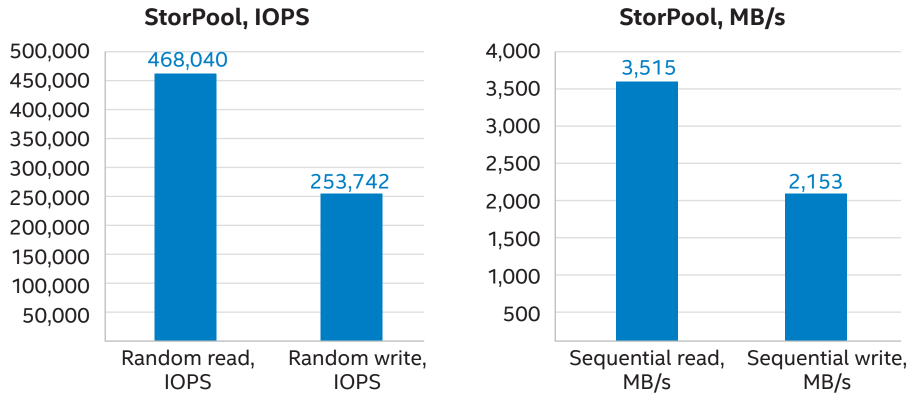

Figure 1. IOPS and Throughput.

TEST	RESULT	TEST PARAMETERS		
		Block Size	Queue Depth	
			Per Client Host	Total
Random read, IOPS	468,040	4k	64	320
Random write, IOPS	253,742	4k	64	320
Sequential read, MB/s	3,515	1M	64	320
Sequential write, MB/s	2,153	1M	64	320

Table 2. Performance Metrics Summary.

TEST	RESULTS		TEST PARAMETERS	
	Average Latency (ms)	IOPS	Block Size	Queue Depth
Random read + write 50/50	0.234	4,245	4k	1
Random read + write 50/50	0.253	8,040	4k	2
Random read + write 50/50	0.268	15,168	4k	4
Random read + write 50/50	0.295	27,952	4k	8
Random read + write 50/50	0.343	49,358	4k	16
Random read + write 50/50	0.417	85,246	4k	32
Random read + write 50/50	0.469	115,386	4k	38
Random read + write 50/50	0.584	122,906	4k	64
Random read + write 50/50	0.615	144,454	4k	80
Random read + write 50/50	0.764	139,412	4k	96
Random read + write 50/50	0.966	147,788	4k	128

Table 3. Latency Test Results.

Interpretation and Recommendations

The impressive throughput and latency results are due to the outstanding performance and latency metrics of the Intel SSD DC S3510 Series and Intel® Ethernet Converged Network Adapter X710 10GbE family. For the storage system, StorPool recommends that customers use Intel® Xeon® processor E3-1230 v5. With just four cores and the efficiency of StorPool's software, more cores are not needed. On the compute side, StorPool recommends using powerful Intel® Xeon® processor E5-2620 v4-based two-socket servers to increase the efficiency and utilization of the compute system. Using standard components from Intel allowed StorPool to build a solution that costs just $4 per usable gigabyte for three years and $0.08/IOPS.

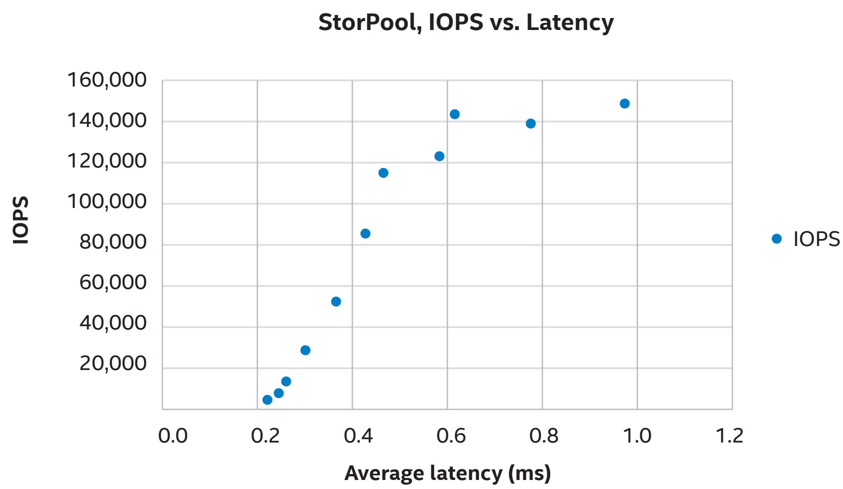

Figure 2. Latency

Conclusion

When a cloud company is looking to differentiate its services with high-performance low-latency storage while containing costs, the biggest competitive advantage it can leverage is its infrastructure layer. It is at the core of its business. StorPool's Software Defined Storage solution can build extremely fast and cost-efficient cloud storage systems, which can set a CSP apart from its competitors.

To learn more, or request a tailored solution, please contact StorPool storage at **info@storpool.com**, or visit **www.storpool.com**.

To learn more about Intel SSDs and storage technologies, visit **www.intel.com/storage**.

The Solutions Library on the Builders home page in particular can help you find reference architectures, white papers, and solution briefs like this one that can help you build and enhance your data center infrastructure: **https://builders.intel.com/solutionslibrary**.

You can also follow us on Twitter, by using **#IntelBuilders**.

[1] Test configuration: Servers: 3x Dell R330, each with the following: CPU: Intel Xeon E3-1230v5; RAM: 2x 16 GB DDR4 UDIMM; 2133 MT/s; 4x Intel S3510 1.6 TB; NIC: 2-port Intel X710 10GbE NIC; HBA: Dell PERC H730 adapter in JBOD mode (Avago 3108 w/ Dell firmware); Linux kernel 3.10.0-514.2.2.el7.x86_64; StorPool software version: 16.01.248.9d12c0a. Network: 2x IBM/Blade Networks G8124 - 24-port 10GbE switch; Direct-attach SFP+ cables. Total storage system resources as tested: 3 storage nodes; - 6x 10GE ports; Total 19.2 TB raw space, 8.7 TB usable space. 5 clients (hypervisor hosts), each with the following: Dell R630 server; CPU: 2x Intel Xeon E5-2620v4; RAM: 16x 16GB DDR4 RDIMM @2133 MT/s; 4-port Intel X710 10GbE NIC, 2 ports used for storage; Linux kernel 2.6.32-642.11.1.el6.x86_64. Full details of test at https://storpool.com/blog/storpool-performance-test-12-ssds-468000-iops.

[2] See http://www.intel.com/content/www/us/en/solid-state-drives/ssd-dc-s3x10-series-brief.html.

[3] Cost reduction scenarios described are intended as examples of how a given Intel- based product, in the specified circumstances and configurations, may affect future costs and provide cost savings. Circumstances will vary. Intel does not guarantee any costs or cost reduction.

[4] See the full testing report at https://storpool.com/blog/storpool-performance-test-12-ssds-468000-iops.

Printed in USA 0817/YMB/HBD/PDF Please Recycle 336278-001US