Revolutionizing data center efficiency

Unveiling the superior performance per watt of Kingston's DC600M SSDs in VMware vSAN environments.

In the ever-evolving landscape of data center technology, efficiency and performance are paramount.

This whitepaper presents an in-depth analysis of Kingston's DC600M Solid State Drives (SSDs) in VMware vSAN environments, emphasizing a crucial metric: performance per watt. Through rigorous testing using HCIBench and the SQL tpcc benchmark, the performance of DC600M SSDs in traditional vSAN hybrid and all-flash vSAN datastores is compared, aiming to highlight the performance efficiency of vSAN all-flash datastores built with DC600M.

The findings demonstrate that Kingston Technology's DC600M SSDs not only deliver superior performance in high-demand scenarios but also significantly reduce energy consumption, offering a dual benefit of cost savings and environmental sustainability. This paper aims to provide data center administrators, IT professionals, and decision-makers with comprehensive insights into why the **DC600M SSDs are an optimal choice for modern data centers** seeking to balance high performance with energy efficiency.

Written by: Hazem Awadallah, Senior Systems Engineer, Kingston Technology
Reviewed by: Chris Selden, SSD Product Engineering Manager, Kingston Technology

Introduction

As data centers continue to be the backbone of enterprise IT infrastructure, the quest for more efficient, high-performance storage solutions becomes increasingly critical. The advent of hyper-converged technologies like VMware vSAN has transformed the way storage can be managed, offering scalable, flexible, and relatively easy-to-manage solutions. However, the choice of underlying storage media—SSDs versus traditional hard drives—plays a pivotal role in determining the overall efficiency and performance of these systems.

In this context, a new metric has gained prominence: Performance per watt. This measures how much performance a storage solution delivers for each watt of power consumed, becoming a critical factor in evaluating data storage options. This metric not only reflects the capability of the storage media in handling intensive workloads but also its impact on the overall energy footprint of the data center.

With a focus on VMware vSAN environments, this paper delves into a comparative performance analysis of DC600M SSDs against vSAN hybrid datastores. **Kingston DC600M SSDs—designed for robust enterprise performance and reliability—are introduced, and are put to the test in a series of benchmarks designed to mimic real-world data center workloads.** The objective is to provide a clear, data-driven picture of how DC600M SSDs stand out not only in terms of raw performance but also in their efficiency, offering a compelling argument for their adoption in contemporary data centers.

Introducing Kingston DC600M

Kingston Technology's DC600M SSD is designed for data centers requiring reliable, high-performance storage.

Following on from the success of DC500M, DC600M is the 4th-generation enterprise-class SATA SSD offering from Kingston. Its enterprise-focused firmware is designed to sustain high performance, low latency, and predictable consistency of enterprise workloads conforming with strict Quality of Service (QoS) requirements and encompassing sophisticated ECC algorithms to ensure reliability of enterprise workloads within the entire lifetime of the drive.

It is engineered to provide resilience against power loss, safeguarding data integrity with onboard power loss protection (PLP). With capacities ranging up to 7.68TB, the DC600M is designed to deliver consistent latency and IOPS, making it an ideal choice for high-volume rack-mount servers and demanding data environments. This drive is particularly well suited for system integrators, hyperscale data centers, and cloud service providers seeking to balance performance with durability.

Kingston's DC600M SSD has proudly secured its place on the VMware ESXi compatibility list, up to the most recent vSAN 8.0 Update 2. This endorsement is a testament to Kingston's dedication to delivering enterprise-grade SSD solutions that meet the stringent requirements of leading-edge virtualization environments.

Testing environment

SATA/SAS/HYBRID testing environment (hardware)	SATA testing environment (OS and software)
PowerEdge Dell R740xD 3 Node Cluster supporting with 8 2.5" NVMe and 16 2.5" SATA/SAS drive bays/server	Hypervisor: VMware ESXi, 7.0.3, 20036589
Intel(R) Xeon(R) Silver 4114 CPU (10c/20t) @ 2.20GHz x8	vSAN 7U3f (VMware ESXi, 7.0.3, 20036589 + VMware VirtualCenter 7.0.3 build-20150588)
768 GB 24x32GB Kingston DDR4 Dual Rank ECC Memory @ 2400MHz/Node, 2304GB/cluster	Guest OS: Windows Server 2019 Datacenter, v1809
2xCisco nexus N5K-C5010 20 port 10Gbe data-center-class switches for vSAN network traffic	Microsoft SQL Server 2017 (RTM) - 14.0.1000.169 (X64)
PERC H740P configured in HBA passthrough mode	HammerDB-v3.2
	HCIBench 2.5.3

Figure 1.1 Hardware and software environment used during the tests

Figure 1.1 shows the hardware and software used throughout the tests conducted in this paper. **The testing was conducted on a meticulously configured hardware and software ecosystem, specifically architected to challenge and evaluate the performance of the Kingston DC600M SSD.** The hardware foundation was a Dell PowerEdge R740xD 3 Node Cluster, each node powered by Intel® Xeon® Silver 4114 CPUs, augmented with 768GB of Kingston Dual Rank ECC memory, culminating in a total of 2304GB for the cluster.

Network connectivity was managed through dual Cisco Nexus N5K-C5010 switches, ensuring seamless vSAN network traffic. The tests were conducted on vSAN 7U3f (VMware ESXi, 7.0.3, 20036589 + VMware VirtualCenter 7.0.3 build-20150588). On the guest OS side, Windows Server 2019 Datacenter served as the operating platform, with Microsoft SQL Server 2017 handling database operations. Performance benchmarks were carried out using HammerDB and HCIBench, offering a comprehensive and rigorous assessment of the SSDs under test.

Three physical drives with the same capacity per disk group were used for both SATA SSD and hybrid testing. For the hybrid-tier testing, two Dell branded Seagate Exos 10k RPM 1.2 TB SAS drives (ST1200MM0099) per server were used for the vSAN capacity tier, and 1 DC600M 960GB for the vSAN cache tier.

For the SATA SSD all-flash vSAN testing, 3 Kingston DC600M 960GB were used (test 2) and 3 Kingston DC600M 3840GB drives (test 1 and 3), with 1 drive for the vSAN cache tier and 2 drives for the capacity.

The vSAN default storage policy was used throughout the tests conducted in this paper. The vSAN Default Storage Policy is the standard policy applied to the VMs provisioned from vSAN datastores, ensuring data resilience through a RAID-1 mirroring configuration that can tolerate a single failure (host, disk, or network). It uses thin provisioning to optimize space utilization and sets no specific IOPS limit for objects, allowing for flexible performance. This policy does not reserve flash read cache (although this is possible for hybrid tiers), ensuring all-flash performance is available to all data as needed, and maintains data integrity with checksums while avoiding forced provisioning to ensure storage allocation occurs only when resources are sufficient.

For the later tests in this paper, the racadm tool incorporated within Dell's srvadmin v11.0.0 package (srvadmin-idracadm8) was used to collect power telemetry from each of the vSAN nodes via IPMI out-of-band ssh connectivity.

For these database tests, a Server 2019 Guest VM with SQL server 2017 was used and a separate vmdk provisioned from the vSAN datastore for data, log, and backup. Hammer DB, a free, open-source database load-testing application was used to run the TPCC benchmark for OLTP applications and TPC-H benchmark for data analytics workload. Throughout the various tests in this paper, the TPCC benchmark specification is chosen here to simulate OLTP transactional workloads and ensure conformance, repeatability, and reliability of testing results.

Test 1: Assessment of raw storage subsystem performance—HCIBench

To assess the raw performance of the I/O subsystem, VMware's recommended tool for benchmarking the vSAN datastore-HCIBench v2.5.3, was used. This automation toolkit deploys multiple VMs spread across all the hosts in the vSAN cluster while running specific workloads using vdbench on all guest VMs in parallel. Results from the run with 6 VMs (2VMs/host) on the DC600M 4TB vSAN datastore are presented.

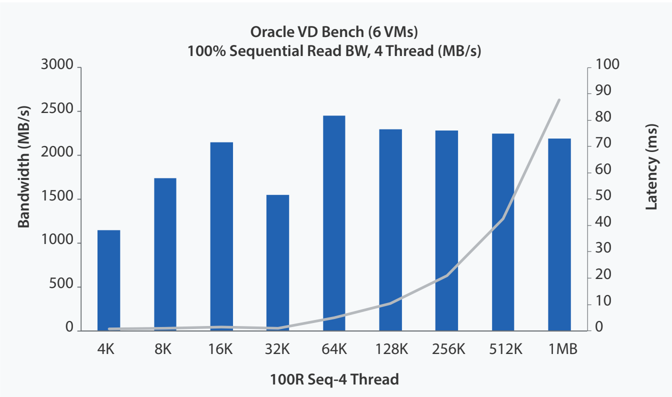

Bandwidth (MB/s)	1147.55	1737.82	2146.07	1550.28	2450.64	2296.39	2279.54	2249.3	2189.88
Latency (ms)	0.6765	0.8685	1.4067	0.9663	4.9445	10.459	21.0433	42.6583	87.7075

Figure 1.2 Sequential Read Performance, Kingston DC600M 3840G 9 Drive vSAN datastore.

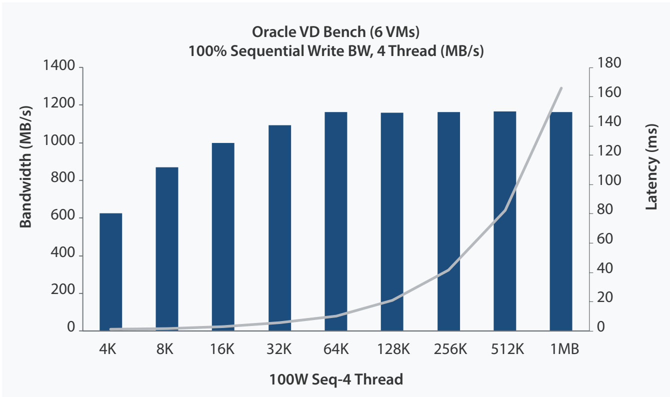

Bandwidth (MB/s)	626.02	868.58	999.01	1094.44	1162.47	1158.63	1162.73	1167.12	1161.6
Latency (ms)	1.2042	1.7343	3.0128	5.4915	10.3533	20.775	41.4505	82.6272	166.094

Figure 1.3 Sequential Write Performance, Kingston DC600M 3840G 9 Drive vSAN datastore.

In the sequential throughput tests, the 9-drive 4TB DC600M vSAN array achieved a robust peak of 2.468GB/s read bandwidth while keeping latency under 5ms per I/O. For writes, it reached a peak of 1.16GB/s, with latency staying below 10ms. As the I/O block size increased, a corresponding rise in latency was observed, which aligns with expectations given the higher data transfer rate. Notably, the absence of significant tail latency spikes highlights the DC600M's excellent QoS and firmware optimization, reinforcing its capability to handle large-scale data transfers with efficiency.

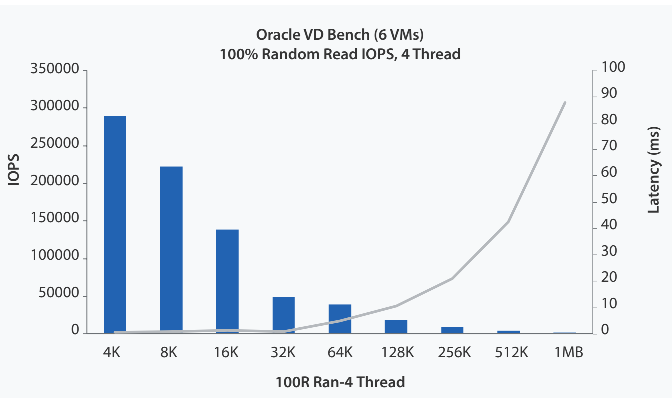

IOPS	289176	221857.6	138910.1	49390.3	39084.8	18300.5	9107.8	4501.4	2189.4
Latency (ms)	0.6853	0.8715	1.3862	0.9707	4.954	10.4912	21.0627	42.6298	87.7118

Figure 1.4 Random Read Performance, Kingston DC600M 3840G 9 Drive vSAN datastore.

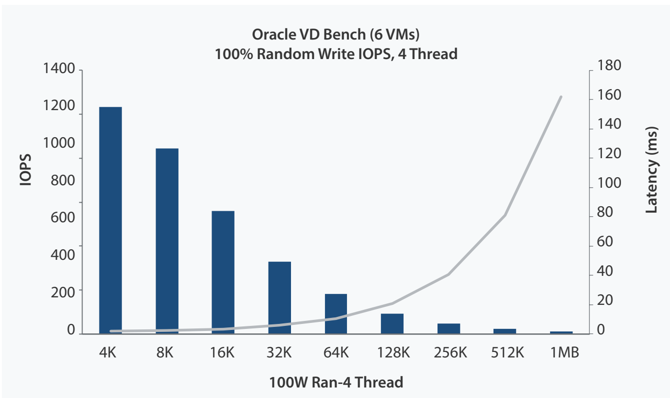

IOPS	103247	84369.1	56032	32944.7	18174	9331.4	4726.5	2377.8	1189.1
Latency (ms)	1.8557	2.2713	3.4233	5.8155	10.5637	20.6062	40.738	81.045	162.1052

Figure 1.5 Random Write Performance, Kingston DC600M 3840G 9 Drive vSAN datastore.

In the random read IOPS benchmarks, the DC600M SSDs achieved a peak of 289,176 IOPS at 4K, with a remarkable latency of 0.68ms. The random write tests showed a strong performance of 103,247 IOPS at 4K, with less than 2ms latency.

During the mixed workload scenarios, combining 30% write and 70% read operations, the SSDs scaled up impressively to 215,660 IOPS while maintaining sub-millisecond latency, demonstrating their high efficiency and responsiveness.

It will be seen later how this raw performance directly correlates to enhanced transactional application capabilities, ensuring rapid processing in database environments and supporting a high volume of concurrent transactions without compromising response times.

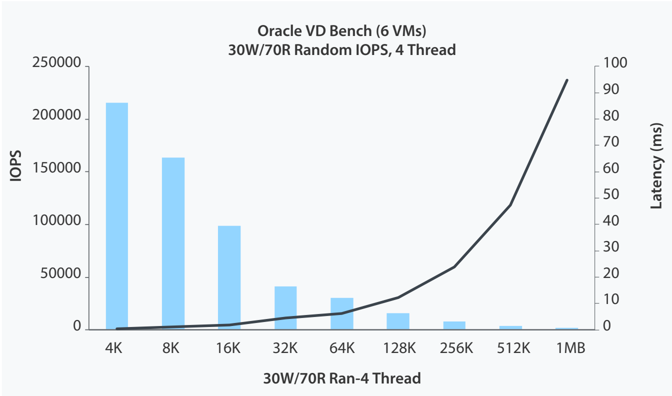

IOPS	215660.9	163572.1	98525	41529.8	30699.9	15736.8	8001.3	4049.9	2028
Latency (ms)	0.55	1.1753	1.9448	4.6138	6.2448	12.2023	24.0198	47.4783	94.8558

Figure 1.6 Random Mixed (70R/30W) Performance, Kingston DC600M 3840G 9 Drive vSAN datastore.

Test 2: SQL TPCC performance, DC600M all flash and hybrid

The goal for test 2 was to get a baseline on the level of performance expected with the TPCC benchmark under a prolonged I/O-bound stress test on VMware vSAN with an all-flash datastore provisioned with DC600M 960GB and a hybrid datastore provisioned with DC600M 960GB and 1.2TB 10K RPM hard drives.

A schema of 2000 warehouses resulting in a tpcc database size of 157 GB was created. 40 virtual cores for each SQL server VM were used to allocate enough CPU resources to saturate the transactional throughput, but only assigned 32 GB of RAM to make the test I/O bound. **The virtual user sequence was tuned to scale up from 1 to 512 users and allowed each virtual user sequence to run for a long time (20 minutes, with a 10-minute ramp-up time)**. This allowed collection of disk latency metrics during the entire duration of the test run.

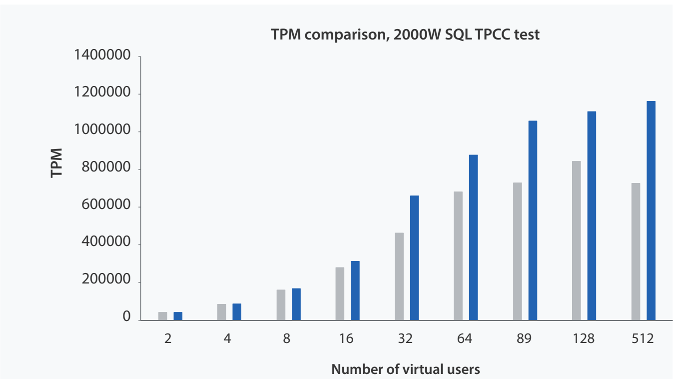

DC600M Hybrid vSAN datastore	42624	84966	160858	279419	462117	681638	729318	842809	728178
TPM DC600M 960G vSAN datastore	40867	86690	168764	313511	659598	876430	1058328	1106963	1163048

Figure 2.1 DC600M vSAN all-flash vs. hybrid transactions per minute autopilot test with 1-512 users

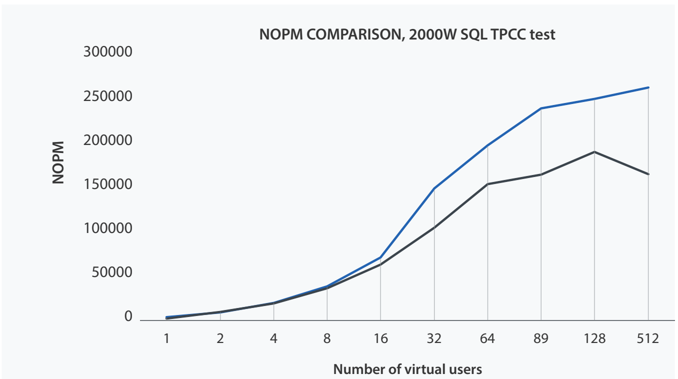

NOPM DC600M 960G vSAN datastore	3748	8916	18857	36727	68193	143376	190438	230173	240757	252858
NOPM DC600M Hybrid vSAN datastore	1858	9288	18489	34952	60691	100553	148302	158555	183263	158643

Figure 2.2 DC600M vSAN all-flash vs. hybrid orders per minute autopilot test with 1-512 users

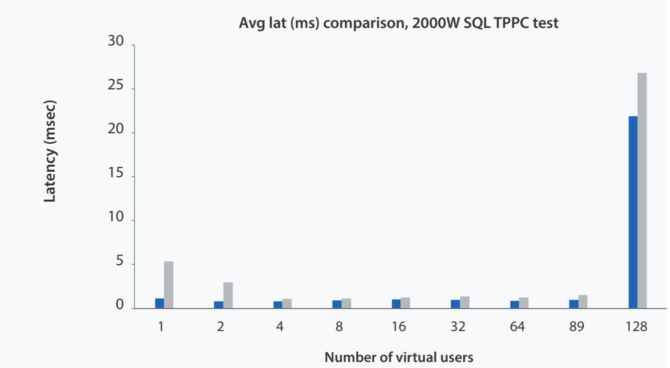

	1	2	4	8	16	32	64	89	128
DC600M 960G vSAN datastore	1.16	0.83	0.79	0.90	1.03	0.98	0.85	0.96	21.86
DC600M Hybrid vSAN datastore	5.35	2.99	1.07	1.15	1.22	1.35	1.26	1.51	26.84

Figure 2.3 DC600M vSAN all-flash vs. hybrid avg latency(ms.) autopilot test with 1-512 users

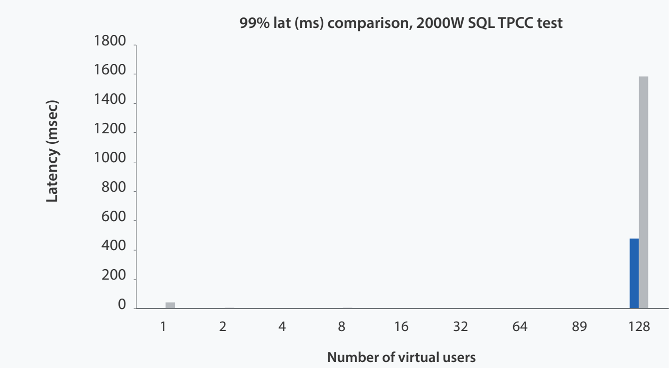

	1	2	4	8	16	32	64	89	128
DC600M 960G vSAN datastore	4.41	2.41	2.01	2.95	5.01	3.29	2.65	2.49	479.09
DC600M Hybrid vSAN datastore	44.29	8.36	3.51	5.95	4.20	4.04	2.95	3.69	1583.39

Figure 2.4 DC600M vSAN all-flash vs. hybrid 99% latency autopilot test with 1-512 users

Figure 2.1-2.4 showcases a detailed performance comparison between DC600M vSAN hybrid and all-flash datastores under a SQL TPC-C benchmark, with particular focus on TPM (Transactions Per Minute), NOPM (New Order Transactions Per Minute), average latency, and 99th percentile latency across different numbers of virtual users.

In the TPM comparison, the all-flash datastore exhibits a significant lead in transaction throughput, consistently outperforming the hybrid datastore as the number of virtual users increases, reaching a peak of 1.16M TPM and 252,858 orders per minute at 512 virtual users.

Comparatively, the hybrid vSAN datastore scales up to a peak of 842,809 TPM and 183,263 orders per minute at 128 virtual users. This trend highlights the superior scalability of the DC600M all-flash vSAN datastore and ability to handle higher transaction volumes as the number of users increases. From a business perspective, if there are 89 users sending transactions to the database simultaneously, each user can process 145% more transactions (translating into more orders per minute) (Fig 2.2) if the vSAN hybrid infrastructure is upgraded to DC600M all-flash.

Latency metrics provide additional insight into system performance. The average latency remains lower for the all-flash datastore across all user counts, suggesting that not only can the system process transactions more rapidly, but it also does so with quicker response times. This is particularly critical for time-sensitive transactional applications where even minor delays can have a significant impact.

The 99th percentile latency comparison reveals that under the highest stress—with 128 virtual users—the all-flash datastore maintains lower latency, whereas the hybrid datastore experiences a substantial increase. This indicates that the all-flash configuration provides not only better average performance but also greater consistency, ensuring that even the slowest transactions are completed in a timely manner.

Collectively, these results demonstrate the tangible benefits of DC600M-built all-flash vSAN datastores in handling the demands of OLTP workloads, showcasing their ability to deliver high transactional throughput with low latency, even as the number of virtual users escalates. **This performance differential underscores the all-flash datastore's suitability for environments where efficiency and speed are paramount.**

Test 3: SQL TPCC stress test, DC600M all flash and hybrid with power telemetry and slot tracing

In test 3, the performance efficiency of vSAN hybrid vs. all-flash datastores is assessed, and a new metric for this assessment—Number of orders per average watt of power consumed—is derived.

For this test, an all-flash vSAN datastore provisioned with 9 DC600M 3840GB and a hybrid datastore provisioned with 1 DC600M 960GB/and 2 1.2TB 10K RPM drives are used.

A comprehensive test utilizing a 2,000W database, with the number of users set at 89 and a fixed duration of two hours, including a 20-minute ramp-up period, is conducted. Real-time power consumption (in watts) from each vSAN node is meticulously monitored. To accomplish this, the racadm command-line tool, part of Dell's srvadmin version 11.0.0 package (srvadmin-idracadm8), via IPMI out-of-band SSH connectivity, is employed.

In parallel, dpmstat—an advanced tracing feature native to the H740P RAID controller—is leveraged to accurately record the total gigabytes read and written, as well as the maximum latency per slot. This allows analysis of performance patterns across both all-flash and hybrid vSAN datastores, providing detailed insights into data transfer volumes and latency on the cache and capacity tiers.

Additionally, to capture disk latency and throughput metrics, the built-in performance counters available within PowerShell's Get-Counter are utilized. **This gives a granular view of the system's performance, enabling meticulous assessment and comparison of the efficiency of the storage solutions under test.**

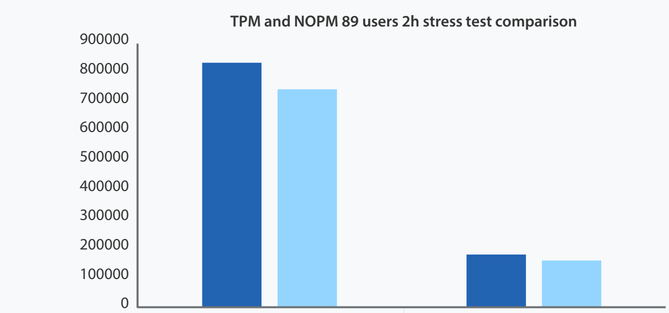

	TPM 89 Users	NOPM 89 Users
DC600M 4TB vSAN all-flash datastore	835043	181563
DC600M Hybrid vSAN datastore	745861	162135

Figure 3.1 TPM and NOPM stress test 89 users, all-flash and hybrid DC600M vSAN datastore

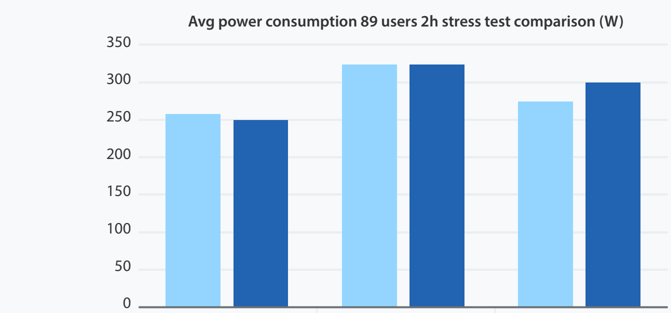

	Server 1 Hybrid 89-user stress test	Server 2 Hybrid 89-user stress test	Server 3 Hybrid 89-user stress test
DC600M all flash vSAN	256.9785276	322.7852761	274.1319018
DC600M Hybrid vSAN	249.0965732	322.5714286	298.5780731

Figure 3.2 Average power consumption 89-user stress test hybrid vs. all flash vSAN DS

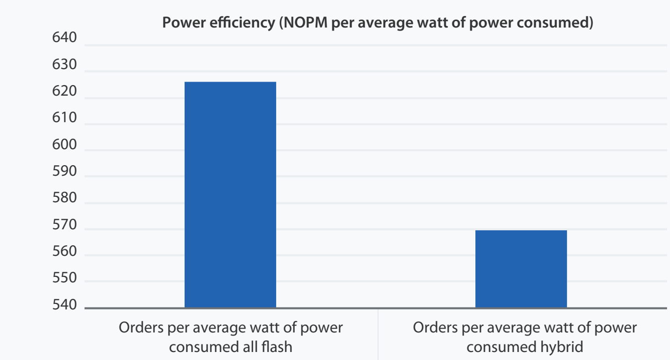

Figure 3.3 Power Efficiency 89-user stress test hybrid vs. all flash DC600M vSAN DS

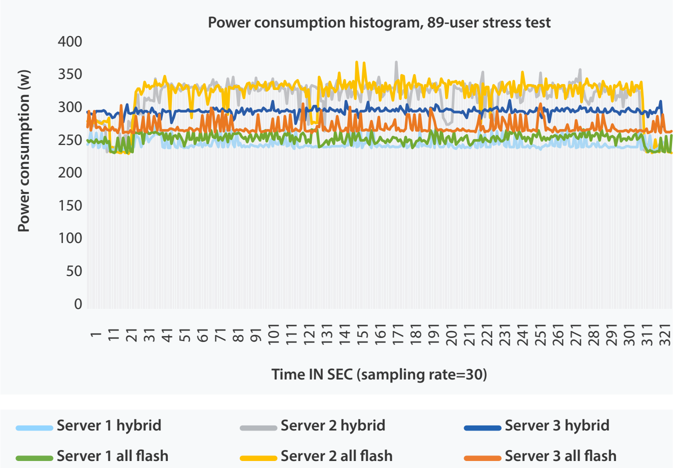

Figure 3.4 Power Consumption histogram 89-user stress test hybrid vs. all flash vSAN DS

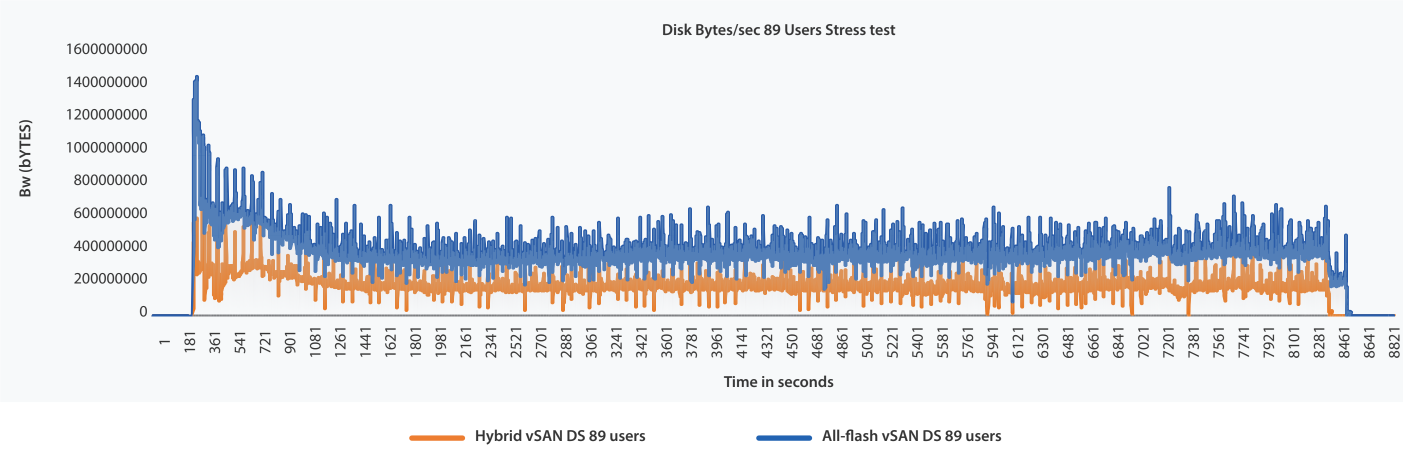

Figure 3.5 Disk bandwidth histogram 89-user stress test hybrid vs. all-flash vSAN DS

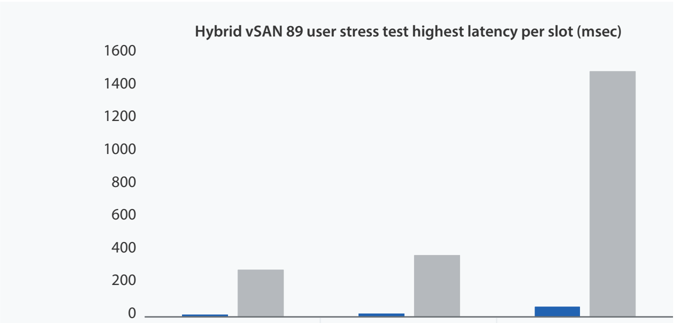

	Server 1	Server 2	Server 3
Highest cache latency	14.4	17.5	58
Highest capacity latency	287	374	1500

Figure 3.6 DPMstat highest LCT latency(ms) 89-user stress test hybrid

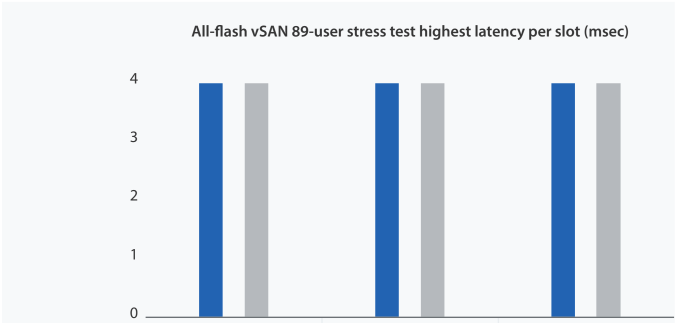

	Server 1	Server 2	Server 3
Highest cache latency	4	4	4
Highest capacity latency	4	4	4

Figure 3.7 DPMstat highest LCT latency(ms) 89-user stress test all-flash vSAN DS

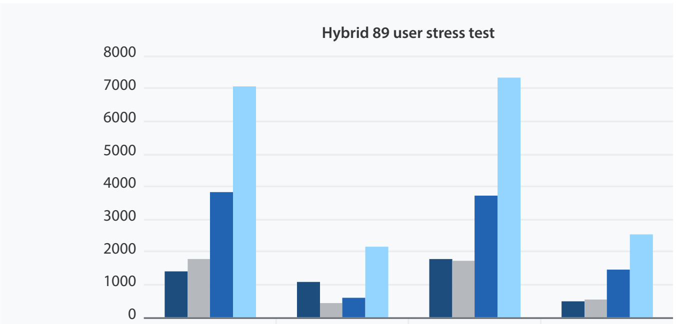

	Total GB read cache	Total GB read capacity	Total GB written cache	Total GB written capacity
Server 1	1436	1126	1817	498
Server 2	1810	464	1756	585
Server 3	3827	604	3748	1488
Total	7073	2194	7321	2571

Figure 3.8 DPMstat GB Read and Written Cache/Capacity Hybrid vSAN DS

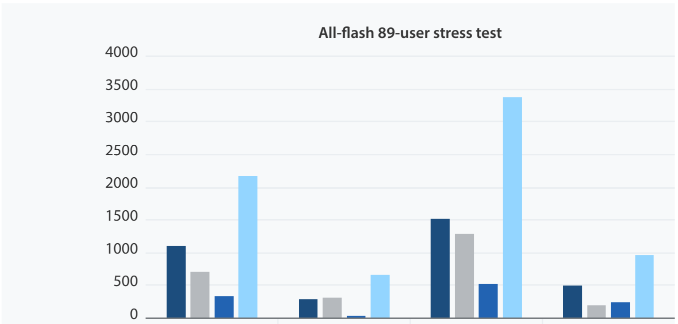

	Total GB read cache	Total GB read capacity	Total GB written cache	Total GB written capacity
Server 1	1114	297	1533	506.7
Server 2	713	326	1291	220
Server 3	339	45	545	245.42
Total	2166	668	3369	972.12

Figure 3.9 DPMstat GB Read and Written Cache/Capacity All-flash vSAN DS

Figures 3.1-3.8 highlight the results of the power efficiency tests for the all-flash vSAN and hybrid-vSAN datastores. The query answered is, how much performance can be derived for every watt of power consumed? A simple equation used to derive the power efficiency difference is presented:

PPW = (NOPM achieved) / (average power consumed from all 3 servers)

ΔPower efficiency = ΔPPW %

PPW for test 3 is highlighted in Figure 3.3. It was possible to achieve 625 orders per watt for the all-flash vSAN datastore compared to 569 orders per watt for the hybrid datastore, a power efficiency gain of ~10%.

A more empirically accurate method was used to determine the performance efficiency of vSAN all flash datastores. First, disk bandwidth metrics vs. time throughout the test were collected, using the Windows performance monitor shown in Figure 3.5. Then, the dpmstat tracing tool was used to determine how many GBs were read and written to the cache and capacity tiers, as well as the highest latency achieved by the cache and capacity tiers in either scenario.

Figure 3.5, the bandwidth histogram, shows the clear performance advantage of all-flash vSAN datastores at delivering higher throughput, offering a 40% improvement throughout the test. **The Hybrid vSAN datastore shows more variable performance with significant peaks, which may correspond to cache misses where data must be retrieved from the HDD capacity tier.** In contrast, the all-flash vSAN presents a more consistent and higher baseline performance, emphasizing its ability to handle reads from both the cache and capacity tier.

Figure 3.8 and Figure 3.9 illustrate the total gigabytes (GB) read from and written to the cache and capacity tiers in both hybrid and all-flash vSAN datastores during an 89-user stress test, based on data from the dpmstat EXT log. The hybrid vSAN configuration, which leverages SSDs for the cache and HDDs for capacity, demonstrates a marked increase in GBs read and written in the cache layer, particularly on server 3. This indicates substantial cache utilization to facilitate read and write operations, a hallmark of hybrid setups where the SSD cache serves as a performance buffer. This buffer mitigates latency by temporarily storing data before it is transferred to the slower HDD capacity layer.

The hybrid vSAN experiences a notable read-modify-write overhead, necessitated by the process of fetching data into the cache for modification before writing it back to the capacity tier. **This can be a time-consuming task due to the mechanical nature of HDDs.** These spikes in the dpmstat LCT log for the capacity tier are visible in Figure 3.6.

In contrast, the all-flash vSAN datastore exhibits lower total GBs read and written in the cache layer across all servers, and consistent latency, (figure 3.7) signaling more streamlined cache usage attributed to the swift DC600M SSDs employed for both caching and capacity. This efficiency gain is because all-flash storage can manage in-place reads more effectively, forgoing the need for pre-emptive read operations and bypassing the caching tier for reads, thus eliminating the read-modify-write cycle that burdens hybrid setups.

In hybrid vSANs, the system promotes frequently accessed data to the cache tier for rapid retrieval, while relegating less frequently accessed data to the capacity tier. The mechanical latency of HDDs, however, introduces a performance penalty during this promotion and demotion activity. All-flash vSAN datastores, by contrast, capitalize on the consistent high I/O capabilities of flash storage across both tiers, minimizing the necessity for data movement. Consequently, all-flash datastores streamline storage management by reducing the complexities associated with cache tier operations, yielding more predictable performance profiles, particularly in scenarios with high user concurrency.

Conclusion

In conclusion, the evidence presented throughout this study highlights the sophisticated performance capabilities of DC600M SSDs within all-flash vSAN datastores. They offer speed, resilience, consistency, and power efficiency, which are paramount in today's data-centric landscapes. For organizations that priorities seamless operation and robust data handling, these SSDs offer a compelling proposition, providing a balanced profile of durability and performance efficiency.

It's not merely about the immediate gains in throughput and reduced latency; it's about the long-term vision for your infrastructure. As data demands grow and evolve, the adaptability and forward compatibility of your storage solutions become critical. **In this light, DC600M SSDs stand out, offering a platform that not only meets the current benchmarks but also anticipates the needs of tomorrow.**

Choosing the right components for your data storage is a strategic decision that echoes through the operational pillars of your organization. With DC600M SSDs, that decision leans towards a future where data is not a hurdle but a catalyst for growth and innovation.

Consider this analysis and how the integration of DC600M SSDs into all-flash vSAN could align your with goals for efficiency, reliability, and preparedness in an era that demands nothing less.

Visit our website to learn more about Kingston's data center solutions. If you have a project, our Ask an Expert team is here to guide and assist you in achieving your goals.