HITACHI
Inspire the Next

Enable predictable performance at the petabyte scale and deliver guaranteed efficiency for superior return on investment. Running on Hitachi Virtual Storage Platform (VSP) family systems, Hitachi Accelerated Flash (HAF) accelerates data access and reduces your storage needs.

DATASHEET

Hitachi Accelerated Flash

Intelligent Flash Performance and Data Center Efficiency

Delivering on enterprise demands for real-time customer engagement requires more than a fast storage array. It requires an end-to-end approach to data management that leverages both the storage operating system and flash media, to deliver low latency performance even as data levels grow exponentially.

Current all-flash arrays (AFAs) rely on performance management to be handled in the array controller along with all other operations, such as data reduction. As data levels increase, they can cause controller bottlenecks, sporadic response times and a poor customer experience. To offset this issue, IT organizations have been forced to make tradeoffs in the number of workloads or amount of data they store on an individual AFA. These decisions result in the need to deploy and support more systems, raising costs and increasing management complexity.

Hitachi Accelerated Flash fundamentally changes this paradigm. By combining a flash-aware I/O stack within Hitachi Storage Virtualization Operating System (SVOS) with industry-leading intelligent flash modules (FMDs), HDS enables IT teams to consolidate more workloads and data onto a single system. At the same time, they can provide scalable high performance and sub-millisecond response times. And with no-penalty, inline compression in each FMD, performance gains are achieved at the lowest effective cost[1], supporting greater operational efficiency.

HAF benefits start with SVOS, which includes an enhanced, flash-aware I/O stack designed to improve low latency throughput, even as data scales to multiples of petabytes. Using patented express I/O and flash-specific data reduction algorithms, SVOS accelerates read and write operations for a lower response time. SVOS also includes FMD integration, enabling the offload of metadata management tasks from SVOS to FMD memory for scalable performance. In configurations with a large number of FMDs, SVOS has the ability to scale processing many times further than common AFAs can and deliver more predictable long-term performance.

FMDs are intelligent flash devices engineered to help improve performance and efficiency by actively working with SVOS. The second and third generation of flash modules (FMD DC2, FMD HD) are built from the ground up to support concurrent, large I/O enterprise workloads and enable hyperscale efficiencies. At their core is an advanced embedded multicore flash controller that increases the performance of multilayer cell (MLC) flash. They intelligently manage I/O and have the ability to prioritize application I/O requests over background tasks, such as garbage collection, to minimize potential latency.

Each FMD also includes an inline compression offload engine and enhanced flash translation layer that delivers up to 80% data reduction (typically 2:1) at 10 times the speed of competitors. Because inline compression can be distributed and managed across FMDs, it gives HDS the ability to easily provide five times, even 10 times the data reduction processing capabilities of traditional AFAs for a better long-term performance profile.

Because of this, Hitachi Data Systems guarantees at least a two times savings in flash capacity, but you're likely to get a lot more.

For virtual server environments, HAF can provide additional advantages beyond long-term scalable performance and efficiency. Leveraging SVOS quality of service (QoS) functionality and integration with VVols, HAF can be used to deliver more predictable customer experiences. SVOS QoS and VVols work together to allocate specific resources to a virtual machine (VM), including FMDs. This ensures that an application running on a VVol not only gets a specific level of performance, but also that it has dedicated FMDs for managing compression. By dedicating data reduction resources to an application, you deliver a more predictable level of latency.

For resiliency, the integrated HAF architecture provides full visibility to every NAND, including its current state and wear level. Using flash wear management, data rebalancing and data reduction to minimize writes, HAF helps extend wear limits.

FMDs come in a rack-optimized and highly dense design with greater than 338TB effective capacity[2] per 2U tray. IOPS performance is up to five times better per flash device than that of enterprise solid-state drives.

Business Benefits

Leading Performance, Lowest Bit Cost, Highest Capacity and Extended Endurance

Hitachi built this new solid-state storage offering with a new 2U flash tray that holds up to 12 high-density flash module drives and associated interconnect cables (see Table 1). FMDs deliver outstanding value compared to enterprise solid-state disk (SSD). When compared to a small-form-factor 1.92TB SSD, it delivers:

- Up to three times the per-device random read 8KB I/O performance.
- Up to five times the per-device random write 8KB I/O performance.
- Up to 70% lower effective[1] bit price.
- 60% lower response time at peak load[3].
- 60 times faster formatting speed.
- Superior data integrity.

FMD Feature Highlights

The Hitachi multicore embedded flash controller introduces multiple unique capabilities, purpose-built for enterprise workloads, lower cost and leading quality of service.

Inline compression offload engine is a very-large-scale integration (VLSI) engine that enables lossless compression based on a derivative of the LZ77 algorithm. It performs at 10 times the speed with none of the burden associated with typical software implementations or journal-based file systems.

Workload priority access is an embedded flash controller feature that integrates with SVOS. Together, they ensure end-to-end application workload quality of service, from a virtual machine down to an individual NAND flash chip, over typical flash background tasks, such as garbage collection and data refresh.

TABLE 1. HITACHI ACCELERATED FLASH: SPECIFICATIONS

Models	FMD		FMD DC2			FMD HD	
Capacity TB (TiB)	1.7 (1.6)	3.5 (3.2)	1.7 (1.6)	3.5 (3.2)	7 (6.4)	7 (6.4)	14 (12.8)
Form Factor	High-density rack form factor						
Interface	SAS 6Gb/sec		SAS 12Gb/sec				
Data Protection	Full data path protection, end-to-end T10 Data Integrity Field support						
Device Tray	2U, 12 devices						
Thin Provisioning Integration	Yes						
Total Tray Capacity TB (TiB)	21.1 (19.2)	42.2 (38.4)	21.1 (19.2)	42.2 (38.4)	84.7 (76.8)	84.7 (76.8)	169.4 (153.6)
Flash Chip Technology	25nm 32GB MLC NAND	25nm 32GB MLC NAND	19nm 64GB MLC NAND	19nm 64GB MLC NAND	19nm 64GB MLC NAND	15nm 128GB MLC NAND	15nm 128GB MLC NAND
Number of Flash Chips (flash memory chips x bus)	16 x 4	32 x 4	8 x 4	16 x 4	32 x 4	16 x 4	32 x 4
Maximum Operating Temperature	35 degrees C		40 degrees C				
Power, Including Tray (Watts)	354	366	450				
Mean Time Between Failures (MTBF)	2 million hours						
Data Eradication Support	Supported						
Inline Compression Offload Support	N/A		Yes				
Workload Priority Access Support	N/A		Yes				
Single Device Sustained Performance							
Quality of Service at 1ms (60/40 read/write 8KB)	97.2%		99.6%				
Random Reads 8KB (kIOPS)	100	100	150	150	150	150	100
Random Writes 8KB (kIOPS)	70	70	80	80	80	80	80
Sequential Reads (GB/sec)	1.0	1.0	2.0	2.0	2.0	2.0	2.0
Sequential Writes (GB/sec)	0.8	0.8	1.0	1.0	1.0	1.0	1.0

Note: All capacities are based on 1TB = 10^{12} bytes and 1TiB = 2^{40} bytes.

Enhanced data integrity corrects up to 59 bits of errors per 2KB of data, delivering better data integrity than serial-attached SCSI (SAS) enterprise disk drives.

Adaptive data refresh dynamically optimizes page refresh based on applied error correction. The program lowers background task overhead and improves sustained performance while extending flash cell longevity.

Data eradication is a service that will physically erase all the user data on the device, including any bad blocks and the overprovisioning of spaces. After the eradication is done, every cell in the FMD is read to ensure that the eradication was completed successfully.

Next Steps

For more information, please visit www.HDS.com.

[1] Based on 1.92TB SSD OEM price and 2:1 typical compression ratio.
[2] Based on typical 2:1 compression ratio.
[3] Based on comparison with Intel SSD DC S3610 Series, March 2015 Product Specifications.

Hitachi Data Systems

Corporate Headquarters
2845 Lafayette Street
Santa Clara, CA 95050-2639 USA
www.HDS.com | community.HDS.com

Regional Contact Information
Americas: +1 866 374 5822 or info@hds.com
Europe, Middle East and Africa: +44 (0) 1753 618000 or info.emea@hds.com
Asia Pacific: +852 3189 7900 or hds.marketing.apac@hds.com

DS-334-G DG October 2016