

# Hitachi Flash Storage: Flash for Multiple Enterprise Data Center Deployment Strategies

By John Webster

January 2016



**Evaluator Group**

*Enabling you to make the best technology decisions*



**Evaluator Group**

## Economic Benefits of Enterprise Flash Storage

Most enterprise applications can benefit from integration with flash storage. But, in particular, those that are online transaction processing (OLTP)-oriented as well as those hosted on virtualized servers will see accelerated levels of performance. This speed will result in an immediate return on investment (ROI) as measured by increased levels of application user productivity as well as operating expenditure (opex) cost reductions. The simple fact is that more transactions completed in a given period of time yield more revenue. An investment in solid-state storage drives increased levels of application performance, resulting in more revenue generating transactions every business day, and yielding an immediate return on the flash investment that can be used to justify the investment.

A positive ROI can also be seen on the customer side of the equation. For example, business user groups and corporate executives wanting to engage with customers via Web- and cloud-based applications, quickly experience benefits. Flash integration underpins an ability to deliver to these customers a positive application experience, thereby enhancing revenue generation.

## What Does Enterprise IT Want From Flash Now?

At the moment, enterprise IT is generally deploying all solid-state storage in the form of flash to accelerate and/or stabilize specific application workloads, such as database transaction processing (OLTP) and virtualized server deployments, including virtual desktops. The drivers of this trend commonly include:

### **Ability to Take Maximum Advantage of Multicore Processor-Based Servers**

The proliferation of high-performance, multicore servers has created an imbalance characterized by an overabundance of processing power on the one hand and an inability to fully utilize it on the other. The all-flash system introduces a persistent storage option that is an order of magnitude less expensive than dynamic random access memory (DRAM) and at least an order of magnitude faster than disk. By eliminating rotating disk as an I/O bottleneck, server-based CPUs spend less time waiting for data and more time doing actual work; thereby, they put the imbalance between CPU power and actual output back into balance. This means that servers can handle additional workloads without adding CPU performance, yielding ongoing savings in application support costs.

### **Predictable Performance in Virtualized Server Environments**

Virtualized server environments make efficient use of CPU and RAM resources but generate random I/O patterns that can have a negative impact on data throughput from disk. This result is often referred to as the I/O blender effect. Virtual desktop infrastructures (VDI) on top of virtualized servers only exacerbates the problem. The all-flash system neutralizes this effect. All VDI users can be given at least the performance level they were used to with physical desktops, if not better. More importantly, consistent and predictable performance levels can be maintained, a critical issue with VDI deployments, as these environments grow. Upward scale can be achieved without adding more server hardware,

allowing the VDI environment to grow economically and with minimal impact to existing application users.

## Opex Savings

Many software vendors are now basing licensing costs on the number of CPU cores in use. Therefore, getting more work out of each processing core saves considerably on expenses devoted to software licensing fees as the workload grows. These savings, compounded over time, represent a significant return to the opex budget on an all-flash system investment.

In addition, operational expenses unique to database environments can be reduced. Database applications typically generate a random I/O workload. This creates a problem for disk storage systems that is typically solved by increasing the number of spinning media devices for the aggregated number of I/Os. Using flash instead eliminates the need to continually add disk spindles to a storage system to maintain consistent and predictable levels of low-latency performance as the database scales to meet increasing workload demands.

Finally, flash can save significantly on energy costs over an extended time period versus spinning disk. It can reduce the amount of power required by the data center as it scales to meet increasing processing demands. Over several years, these savings contribute to a real reduction in OPEX compared to what would be expected with hard drives.

## Current Flash Deployment Strategies: Add or Replace?

We have seen that flash deployments are now following two basic scenarios:

- *Replace* spinning disk with solid state in the form of all-flash systems.
- *Add* flash to existing environments as flash tiers.

Either way, the need to accelerate application performance, meet service level agreements (SLAs), and improve the user and customer experience is motivating both groups. However, which deployment scenario storage architects adopt depends on their guiding philosophy.

### Why Add?

Some believe that the cost of rotating disk on the basis of cost per unit of capacity will always be lower than solid state, so there will always be a use case for it. Therefore, their current flash strategy is to add flash either to existing storage systems or as all-flash systems. The advantages of this approach include:

- Minimizing disruption to production IT environments.
- Minimizing management change.
- Consolidation of flash and disk-based data to common storage pools.

### Why Replace?

Others believe that at some point in the future, all enterprise storage available from vendors will be solid state and that the cost of solid state, all cost factors considered, will be lower than rotating disk. Therefore they're starting the conversion in their data centers now by replacing disk storage systems with solid state. These factors include:

- Long-term total cost of ownership (TCO) of flash versus disk that includes a fully weighted cost calculation (power, cooling, maintenance and so forth, over a three-to-five-year period).
- The desire to make a strategic investment now.
- The requirement to establish a consistent way to accelerate all existing and new primary applications.

## Hitachi Flash Storage

Hitachi has announced a new all-flash system: Hitachi Flash Storage (HFS). HFS is built to HDS design requirements, which include high density and capacity efficiency, management simplicity, energy savings and selectable data services. It is positioned as a high-density all-flash offering for VDI and smaller production environments where server virtualization is now common. It is also appropriate for test and development for databases where performance is critical and deduplication of copies is desired. It offers performance commensurate with or exceeding that of competing all-flash systems and will be price-competitive in small-to-medium-enterprise (SME) segment of the market. However, because it can be integrated with Hitachi Virtual Storage Platform G series (VSP G series) systems, it can also function as an all flash storage module, or tier, in a larger hybrid (flash plus disk) storage system.

From a hardware perspective, HFS features dual active-active controllers and up to 60 solid-state disk (SSD) slots. Additional high-availability features include hot-swappable field replaceable units (FRUs)<sup>1</sup> and dual 1 + 1 redundant high-efficiency power supplies.

Features and attributes include:

- High-density flash, incorporating a 2U enclosure with 60 SSDs and 100TB of raw capacity.
- Fibre Channel and iSCSI host interface protocol choices.
- 1.6TB SSD serial-attached SCSI (SAS) back-end interface.
- High availability, dual active-active controller design with no single points of failure. Each controller has RAID groups mapped to it, and the other controller has secondary access.
- Writes go directly to SSD, bypassing cache. Cache is used for metadata.
- Simplified management interface with optional integration into existing VSP G series management platform.

---

<sup>1</sup> Customer-replaceable FRUs include SSDs, controllers, fans, power supplies, DIMMS, and host interface cards.

- User-selectable data services including in-line data deduplication, compression, thin provisioning native data replication (snapshots and remote replication) and multitenancy controls (see below for more detail).
- Data analytics reporting.
- Encryption (in a later release).
- Open APIs (OpenStack Cinder).
- VMware vCenter integration and VMware vStorage APIs for Array Integration (VAAI) support.
- Online microcode upgrades and online model upgrades.

HFS will be introduced in three model variations (see Table 1). All models will be generally available starting January 2016 and pricing will include five years of standard support.

TABLE 1. HITACHI FLASH STORAGE MODEL SPECIFICATIONS

	Small (HFS A220)	Medium (HFS A250)	Large (HFS A270)
SSDs	10 x 1.6TB	30 x 1.6TB	60 x 1.6TB
Maximum Capacity	16TB raw 12.8 useable <sup>2</sup> 64 TB effective <sup>3</sup>	48TB raw 38 useable 192 TB effective	96TB raw 76 useable 384 TB effective
Host Port Options	8 x 16Gb/sec Fibre Channel <sup>4</sup> 8 x 10Gb/sec iSCSI 8 x 40Gb/sec iSCSI 16 x 8Gb/sec Fibre Channel and 8 x 56Gb/sec InfiniBand (iSER) in a future release	Same	Same
Cache	64GB cache per controller, 128GB total	64GB cache per controller, 128GB total	128GB cache per controller, 256GB total
RAID Levels	1, 1 + 0, 5, 6	Same	Same
Device Interface	6Gb/sec SAS	Same	Same
Chassis Height	2U	Same	Same
Maximum Power Consumption	570W	750W	880W
Weight	74.0 lbs. / 33.6 kg	80.8 lbs. / 36.7 kg	91.0 lbs. / 41.3 kg
Model Upgrades	Add 20 x 1.6GB disk pack	Add 30 x 1.6GB disk pack	

<sup>2</sup> Useable capacity assumes 20% of total raw capacity is used for RAID overhead.

<sup>3</sup> Effective capacity assumes 5:1 deduplication/compression ratios.

<sup>4</sup> Available 1 - 2 months after GA.

	yields medium model  Add 50 x 1.6GB disk pack and 4 x 32GB DIMMS (cache) yields large model	and 4 x 32GB DIMMS (cache) yields large model	
--	---	--	--

## Manageability

One of the most striking attributes of HFS is the number of interfaces that can be used to manage it. Yes, there are native graphical (GUI) and command line (CLI) interfaces, but other Hitachi and third-party management options are also supported. These include:

- Hitachi Infrastructure Director and Hitachi Storage Virtualization Operating System (SVOS) via HFS attachment to VSP as an all-flash storage component within a virtualized storage environment: Supported on first release.
- VMware vCenter and OpenStack Cinder: Available upon first release with additional open REST APIs to be supported in future releases.

## Selectable Data Services

Also worth noting (as compared to competing all flash storage systems) is the ability that administrators have to select the data services they want running on HFS. Because use of advanced data services in an all-flash system has performance implications, an administrator can configure only the data services needed for the application. These include:

**In-line data deduplication and compression** to reduce data before it's written to the SSDs. Optional tuning of deduplication block sizes is available to balance data reduction rates with performance. Hitachi guidelines lead to an average of 5:1 capacity savings.

**Thin provisioning** allocates capacity on-demand from a pool of logical volumes.

**Snapshot and clone copying** provide support for up to 64 copy-on-write snapshots per logical volume. Full clones of logical volumes can also be copied.

**Native remote replication** allows logical volumes to be copied to a second HFS system with either asynchronous or synchronous writes.

**Quality of Service controls** allow maximum IOPS and bandwidth consumption to be set for logical volumes to prevent "noisy neighbor" problems. These include I/O throttling, volume management, cache management and other administrative controls.

## Hitachi Flash Storage: Responding to Both Deployment Strategies

HFS can be deployed as a standalone flash storage system to replace hard disk drive (HDD)-based storage systems and can be managed by either its simplified native software platform or by Hitachi Infrastructure Director. The TCO of HFS makes this a viable option, thanks to long-term opex savings coming from its high-density footprint, low power consumption and management simplicity. The high-performance capability of HFS enables additional savings via a reduction in software licenses.

However, for enterprise data center administrators wishing to add flash to an existing HDD environment, HFS can be virtualized as an all-flash tier of a VSP G series HDD system running SVOS. Manageability across the SVOS “stack” brings management efficiency versus all-flash system point solutions.

- Common and automated management.
- Common data protection and business continuity processes.
- Common security model.
- Long-term data lifecycle management.

## Evaluator Group Assessment

As mentioned earlier, HFS is positioned as a high-density all-flash offering for VDI and smaller production environments. We note that virtualized server environments make efficient use of CPU and RAM resources but generate random I/O patterns that can have a negative impact on data throughput from disk, a result often referred to as the I/O blender effect. Virtualizing desktops (VDI) on top of virtualized servers only exacerbates the problem. The all-flash system neutralizes this effect. All VDI users can be given at least the performance level they were used to with physical desktops, if not better. More importantly, consistent and predictable performance levels can be maintained, a critical issue with VDI deployments, as these environments grow. Upward scale can be achieved without adding more server hardware, allowing the VDI environment to grow economically and with minimal impact to existing application users.

Virtualized servers also pose challenges for IT administrators responsible for data protection and business continuance. The all-flash system that is optimized for data resiliency and availability within virtualized server environments can not only address these issues but also potentially accelerate the performance of these processes.

With HFS, Hitachi is delivering the ROI and TCO benefits of flash now. HFS allows administrators to start their strategic investments in all-flash data IT environments or deploy what they believe to be the best storage media for the application. In addition, the ability of HFS to be a virtualization target of an SVOS-based controller allows it to respond to both add and replace deployment strategies. Either way, HFS represents a production data center implementation of high-performance flash storage. We expect all of the HFS models to be popular with channel resellers. HFS will find homes in small-to-medium scale enterprises as well as larger organizations as departmental storage systems or integrated with Hitachi Virtual Storage Platform G series systems.



## About Evaluator Group

*Evaluator Group Inc. is a technology research and advisory company covering Information Management, Storage and Systems. Executives and IT Managers use us daily to make informed decisions to architect and purchase systems supporting their digital data. We get beyond the technology landscape by defining requirements and knowing the products in-depth along with the intricacies that dictate long-term successful strategies. [www.evaluatorgroup.com](http://www.evaluatorgroup.com) @evaluator\_group*

**Copyright 2016 Evaluator Group, Inc. All rights reserved.**

*No part of this publication may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying and recording, or stored in a database or retrieval system for any purpose without the express written consent of Evaluator Group Inc. The information contained in this document is subject to change without notice. Evaluator Group assumes no responsibility for errors or omissions. Evaluator Group makes no expressed or implied warranties in this document relating to the use or operation of the products described herein. In no event shall Evaluator Group be liable for any indirect, special, consequential or incidental damages arising out of or associated with any aspect of this publication, even if advised of the possibility of such damages. The Evaluator Series is a trademark of Evaluator Group, Inc. Hitachi is a trademark or registered trademark of Hitachi, Ltd. All other trademarks are the property of their respective companies.*