

Optimizing Oracle with HP 3PAR StoreServ All-flash Storage

By Russ Fellows & Edward Whalen

August 2015



Evaluator Group

Enabling you to make the best technology decisions



Evaluator Group



Executive Overview

Many applications deemed mission-critical are transaction-processing applications, utilizing underlying database technology. Organizations with these applications are spending a significant portion of their annual IT budget on database-licensing fees and any database efficiency improvements have both productivity and financial benefits. Moreover, there is a significant need to improve both the performance and efficiency of database applications in order to meet increasing workloads and potentially lower costs.

Although computer technology has advanced significantly over the past decade, electro-mechanical devices such as storage hard-disk drives have not kept pace with the advances seen in solid-state technologies. Specifically, CPU, memory and networking speeds have all experienced performance improvements of more than 1,000-fold over the past 25 years. In contrast, hard-disk storage has only seen a 3X performance increase when measured in I/O latency.¹

Improving storage performance is critical to improving the performance of Oracle database applications, which in turn increases operational efficiency. Solid-state storage technologies have been used to enhance the performance of critical database applications for several decades. Early solid-state storage utilized computer memory technologies to provide high speeds, but did so at prices that made these solutions unsuitable for most applications.

With the recent advances in NAND flash technology, solid-state media is now within the same order of magnitude in cost as traditional hard-drive media when measured on a cost per Gigabyte basis. When measured in terms of performance, solid-state delivers significantly lower cost per I/O operation than traditional disk media. However, due to concerns with new all-flash storage systems, coupled with a lack of enterprise capabilities, enterprise IT managers and DBAs have been reluctant to use solid-state storage systems with their databases.

In order to assist organizations in their analysis of options, HP asked Evaluator Group to analyze the specific requirements Oracle-based applications have, particularly in larger enterprise settings. Evaluator Group also worked with Performance Tuning, a respected organization with extensive experience in optimizing Oracle databases in real-world settings. This paper analyzes financial benefits, data availability and other concerns and provides a set of recommendations to optimize Oracle database performance and the HP storage technologies best suited for Oracle.

¹ Moore's Law - http://en.wikipedia.org/wiki/Moore's_law

Storage Performance by Application

Application performance is often not linearly related to improvements within the entire system. Due to the inter-dependency, individual delays cascade, creating end-to-end delays that are greater than expected. Storage and storage performance are frequently cited concerns in database applications, with a significant number of Oracle DBA's citing storage as a leading performance concern. Solid-state storage technology can significantly reduce I/O processing time - or latencies - which in turn benefits database performance.

Evaluator Group comments: In discussions with our clients and based upon lab tests comparing traditional disk-based systems to solid-state, we believe transaction oriented applications can experience up to an order of magnitude performance improvement measured at the application level when using solid-state.

Depicted below in Figure 1 is an overview of workload categories, along with their applicability to the use of high performance solid-state arrays. Over time, as cost for capacity continues to improve, a greater number of workloads become practical to deploy on solid-state systems.

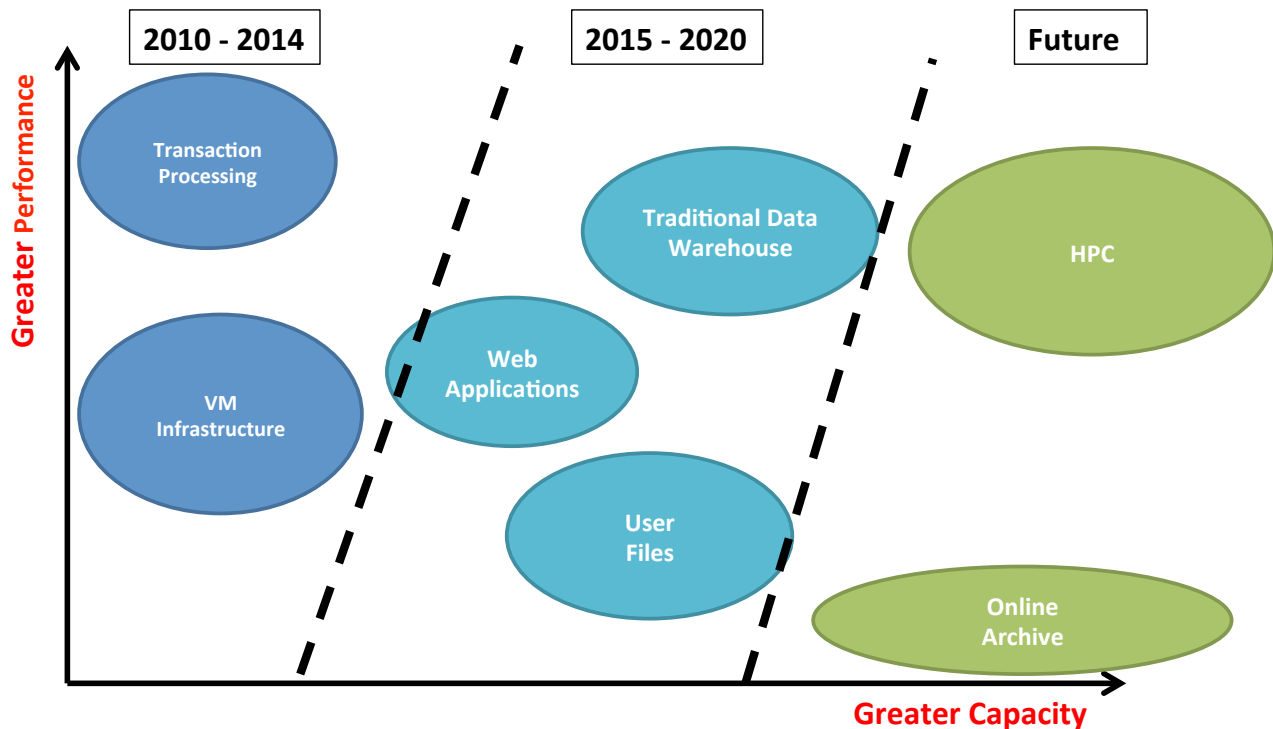


Figure 1: Enterprise Deployment of Solid-State Storage Over Time (Source: Evaluator Group)

As seen above, currently solid-state storage is applicable to transaction processing workloads and is well suited to a growing number of workloads as prices continue to decrease. Solid-state systems are significantly less expensive than traditional storage systems when measured on a cost per I/O basis. As NAND flash prices continue to decline, coupled with greater storage efficiency technologies, solid-state systems are beginning to compete with traditional arrays in other application environments. This trend will continue with solid-state technologies gradually displacing rotating media as the primary storage type for many applications.

Solid-State Storage Considerations

There are multiple aspects that organizations consider when evaluating specific technologies. Ultimately, three broad categories are the primary aspects of consideration:

- Cost
- Reliability
- Performance

Each of these is examined in order to understand the total set of considerations typical organizations use when making a technology evaluation.

Financial Benefits of Solid-State

Measuring the value of storage is a difficult task, and should include multiple aspects. As outlined above, cost, reliability and performance are all important considerations. The most common method of measuring storage is capacity, due in part to the relative ease of measurement.

However, capacity is only one metric, and is quickly becoming the least interesting aspects of storage systems with the advent of data efficiency technologies, solid-state performance and enterprise availability features all important items of consideration.

Evaluator Group Comments: Enterprise storage systems provide value to applications in multiple ways, including, performance and availability in addition to the capacity provided to applications. Moreover, the value of storage must include each aspect, moving beyond a simple \$ / GB comparison. While cost and capacity are important considerations, the value of performance and enterprise reliability are also important, particularly for business critical database applications.

A primary consideration for database applications in general and Oracle specifically is the efficiency of the application processing. The costs associated with Oracle licenses can be a significant portion of IT budgets for enterprises using Oracle databases. As a result, well-run organizations optimize their IT equipment and operations in order to maximize the utilization of their strategic investment in Oracle.

Licensing costs can appear as either a capital expense, or operational expense depending upon a firm's accounting standards. For a high-level analysis of CAPEX and OPEX implications of Oracle, we utilize current Oracle pricing available online.² Although prices vary substantially, Oracle Enterprise Edition lists for \$47,500 per processor for a perpetual license. Oracle counts processor families differently, with current Intel x86 cores counting as 0.50 - or one half - of a processor. Additional options, like Real Application Clusters, Active Data Guard, Advanced Compression and others each incur additional costs. As a result, it is common for customers to pay a fee of over \$50,000 per x86 processor for each perpetual license with support an additional fee.

Optimizing an application begins with an analysis of inefficiencies or bottlenecks that are inherent in any system. One of the most common limitations with databases is storage, or more generally the I/O path from the database to the long-term storage media. To Oracle and other DBAs, "storage" is typically viewed as the entire I/O path below Oracle Automatic Storage Manager (ASM), which includes the OS, storage network and storage system.

By improving the performance of Oracle based applications, organizations are able to choose to deliver more business results in the same period of time, or reduce the cost of processing the same amount of data. Each option can provide benefits for a company and is dependent upon the nature of their operation. As shown, Oracle fees on a per processor basis can exceed \$100,000 per server and may be a significant portion of annual IT budgets. Optimizing this component can result in significant cost savings and efficiency gains.

Optimizing Workloads

For companies that typically operate with the same amount of data, such as those creating monthly bills, processing payroll, or performing similar operations, the amount of data is often fixed. The value of faster application processing comes from the ability to perform the same amount of work in less time while using less capital or lowering operational expenses.

Another common fixed workload application is backing up, or copying data to another location. Optimization occurs not from doing more, but from doing the same amount of work in less time. The so-called "backup window" is a critical time period and can be open for many hours, or in some cases only for a few minutes or seconds. This then directly corresponds with another key measurement, data reliability and availability. The fewer data protection points, the greater the exposure and potential loss of data. The financial cost for loss of data varies by industry and company size but is always a significant business consideration.

² Oracle Technology Pricing - April 9, 2015. <http://www.oracle.com/us/corporate/pricing/technology-price-list-070617.pdf>

Faster Application Processing

Without current, up-to-date information, businesses are by definition operating inefficiently. The business value of near real-time reports or ad-hoc queries can provide substantial business value, yet are typically not done due to the significant time and I/O impact to existing application workloads. Every aspect of a business can suffer from inaccurate database information, including order inventory and distribution, to marketing, advertising and sales.

Reports are often run as batch processes, typically during periods of low transaction processing, while ad-hoc queries may be restricted. Both of these operations require complex database queries that lead to significant I/O activity of data that may not be active.

Evaluator Group comments: Computer and storage systems all cache recently accessed data. However it is the ability to quickly access cold data that defines the performance of complex query operations. For this reason, tiering or caching storage systems are significantly less effective in accelerating ad-hoc queries and reports than all-flash arrays.

Therefore, many companies typically restrict the ability to run these types of operations in order to minimize potential disruption to daily operations that involve transaction processing. The lack of up-to-date information, or not having answers to current conditions, can affect business efficiency, and ultimately both reduce revenue and increase operational costs.

Data Protection with Solid-State

Protecting data is a prime consideration for all IT administrators and is of particular concern to Oracle DBAs. The lack of enterprise class data protection features has been an inhibitor to adoption of solid-state storage systems. Many all-flash storage systems began shipping to customers without the data protection and disaster recovery features typically found in high-end enterprise class arrays. This has been one reason for the cautious adoption of all flash storage for critical applications such as Oracle.

In addition, storage must be aware of an application in order to enable the creation of application-consistent data protection points using array-based technologies. Thus, storage systems must provide tools that support snapshot and replication of application-consistent data. Not only is creating and managing protection points important, but so too is the ability to rapidly recover an application when an incident occurs. Therefore, true application integration consists of all of these capabilities.

Evaluator Group Comments: The ability to create data protection points with application consistency, and support the rapid recovery of an application from a protection point is what defines true application integration for a storage system. Perhaps nowhere in the enterprise is application integration with storage more critical than with Oracle databases.

A new consideration with solid-state storage is the performance impact of remote replication technology. Many organizations rely upon synchronous replication to protect their critical data with no data loss. Due to latencies with traditional storage media, it is possible to include a synchronous, round-trip I/O operation in nearly the same time as a local I/O operation. With solid-state storage this is no longer true, since replication I/O delays may be significantly slower than local I/O. This results in synchronous replication slowing all-flash systems down to the speed of traditional arrays.

An alternative is to utilize asynchronous replication technologies. However, this method can add significant exposure to data-loss, something that critical applications are unable to tolerate. For this reason, neither synchronous nor asynchronous replication appears to be appropriate. Clearly, solid-state systems must provide remote data protection and offer a method to balance data loss with transaction processing speed. New techniques to remedy this issue exist and are examined subsequently in this paper.

Solid-State Storage Performance

The relatively high performance of NAND flash media in comparison to its cost has had a dramatic impact on storage system designs. With costs falling rapidly, many workloads that were previously not cost effective to deploy on an all-flash system have now become economically feasible. However, sustained performance for thousands of consolidated workloads requires more than simply adding SSD media into an existing storage array.

In order to deliver consistent storage performance, the entire system must be designed to support thousands of simultaneous operations, with the ability to isolate and prioritize specific workloads when required. Mid-tier arrays using dual-controller architectures are unable to provide the reliable performance levels required for Tier-1 environments, where a controller failure can significantly reduce performance. In contrast, true Tier-1 storage systems commonly use four, eight or even more controllers to balance workloads under all scenarios.

Evaluator Group comments: Performance improvements vary substantially, however; our clients report latency sensitive application performance improvements ranging from 2X to 10X by moving to all solid-state storage. Performance improvements depend upon multiple factors, although higher performing storage is a significant element.

One approach to scaling performance is a scale-out system, which adds independent nodes to a cluster. This approach has benefits, but these designs typically have limitations due to system overhead that limits the size. Federating multiple systems is a concept that enables systems to effectively grow beyond the limitations of their cluster interconnects. Storage federation enables transparent workload migration for load balancing, disaster failover scenarios or to enable on-line upgrades.

Oracle Database Considerations

With Oracle often being the database behind critical enterprise applications, protecting Oracle data is imperative. Traditionally, disk drives have been a principle source of failure in database systems. Because of the mechanical nature of disks drives, parts wear out and cause failures. Although, RAID technology reduces the possibility of data loss, RAID rebuild times and degraded operations can still impact overall availability and performance.

Using solid-state storage can enhance both reliability and performance. By understanding Oracle I/O performance it becomes clear how solid-state storage is a benefit. An Oracle database derives most of its performance from the use of a memory cache, known as the Oracle Buffer Cache. The Buffer Cache works off of a least-recently-used list. A key design aspect of Oracle is the System Global Area (SGA), which is effectively an in memory cache. Using a large SGA enhances performance by caching data to speed read operations and buffer writes.

Updates to the Oracle database done by users all occur in the Buffer Cache. Writes are deferred to a background process called the DB-Writer. Because of this, as long as the DB-Writer is keeping up, write performance will not result in performance degradation for applications. However, read performance is a cause of delays, since processing cannot occur without the data. Thus, optimizing read performance is more important than optimizing write performance. Solid-state technology provides the read performance needed to provide enhanced performance to Oracle databases.

In addition, Oracle has provided additional features to take advantage of solid-state disk technology, such as the Database Smart Flash Cache. The Oracle Database Smart Flash Cache provides a way of using SSD storage to create a second-level Oracle database cache. In order to use the Database Smart Flash Cache, a cache file is created on SSD storage. As data is aged out of the Oracle in-RAM cache, it is written to the Smart Flash Cache. This provides a level of cache that is slower than RAM but faster than database storage. When data is accessed, the Buffer Cache is checked first, followed by the Flash Cache, and finally data is read from disk. This optimization was introduced in Oracle version 11.

Optimizing Oracle with Solid-State Storage

Optimizing Oracle performance requires an end-to-end, optimized I/O path that includes Oracle, the Operating System and its I/O stack, the storage network and storage systems. With the dramatic adoption of solid-state technologies by storage vendors over the past several years, nearly every vendor provides some SSD-enhanced offering.

It is important to start with the bottom of the I/O stack, since bottlenecks at higher layers may not become apparent until the underlying subsystem has been optimized. The areas of focus include the following items, in order:

1. Optimize the storage system
2. Optimize the host OS, I/O stack
3. Optimize Oracle I/O subsystem (ASM)

As outlined previously, it is important to use storage that meets the business availability and data protection needs first, before focusing on the performance aspects. Specific recommendations for using solid-state technologies with Oracle include the following:

Storage Optimization

Storage performance and reliability can be the difference between a stable, well performing database and a problem database. Database outages can be devastating and costly; careful planning and design is critical. By carefully choosing a storage platform with enterprise features, DBA's can optimize their applications. Here are a few recommendations for storage consideration:

1. Choose a storage platform with enterprise class availability and management features.
2. Ensure storage system has key Oracle integration points
 - a. Use thin provisioning with Oracle ASM - if and only if storage system retains thin volumes
 - b. Utilize storage system integration with Oracle protection tools such as RMAN
 - c. Use storage integrated snapshots to quiesce Oracle before making protection points

Select a storage platform with all-flash, tiering and high-capacity options for secondary copies

OS Optimization

A common inefficiency in Operating Systems is the use of I/O schedulers, which were designed for slow, spinning media. However, with high-speed storage and solid-state in particular I/O schedulers reduce performance and should be eliminated. As with any I/O, using more logical devices reduces queuing and therefore improves performance. Finally, it is not an Oracle best practice to utilize file systems for Oracle databases.

1. Use block storage devices, and use Oracle ASM
2. Use NOOP I/O scheduler for Linux devices. This will reduce "double sorting" in the OS and then in the physical device
3. Use multiple logical devices to reduce I/O queuing. The OS will create a queue for each filesystem
4. If using filesystem, set the NOATIME setting to disable updating access times

Solid-State with Oracle Recommendations

1. Optimize Storage

- a. Best Option: All-flash systems for highest performance. Allows for optimal access to all data, thereby accelerating transaction processing and reporting consistently for all application data
- b. Second Option: Utilize SSD acceleration for portion of database. This provides enhanced performance while still using spinning HDD for less critical data
 - i. Start with 20% of capacity as SSD, then analyze AWR and Statspack reports
 - ii. Implement SSD as a Tier with Auto-tiering technology for primary table spaces
 - iii. No need to optimize Redo or Archive Logs since they are primarily write-only

2. Optimize Oracle

- a. Analyze Statspack or AWR reports looking for I/O and read-related wait events. This is the primary indication that storage latency is an issue.
- b. Utilize SSD disk groups in Oracle ASM in prioritized order
 - i. DB flash cache, selective tables or partitions on Flash by looking for tables with high reads
 - ii. SSD for TEMP and UNDO tables, particularly for multi-pass sorts or joins. These tables can see both high write and read access.
 - iii. Oracle 12C – Automatic Data Optimization (ADO) allows for storage tiering as well as different levels of compression, depending on the data.

By optimizing the storage subsystem, the performance of the database is enhanced. Some operations in Oracle, such as large reports or decision support queries, can involve millions or sometimes billions of I/O operations. Reducing I/Os by a few milliseconds can result in cutting hours off of a report.

HP Solid-State

HP provides a variety of storage products that utilize solid-state technology. Delivering consistent high-performance requires more than adding SSDs to a storage platform designed for traditional hard drives. Initially, when NAND flash in disk drive formats became economically feasible, adding SSDs was common. However, system inefficiencies quickly became apparent, thereby limiting the usefulness of this approach.

The second wave of solid-state storage deployments includes systems optimized or designed to utilize solid-state media, most commonly in the form of SSDs. These systems have either been designed with solid-state, or optimized such that controllers and other elements of the array do not impede the inherent speed of the media. Delivering consistent performance is another differentiating factor with this generation, accomplished by minimizing the write penalty with NAND flash and efficiently managing the utilization of the media.

Another key consideration when optimizing storage is how efficiently systems store data. There are a number of data reduction technologies that can be applied to increase efficiency, although each technology has an affinity for specific workloads and data patterns. Storage efficiency technologies include thin provisioning, compression and data deduplication.

Evaluator Group Comments: For Oracle data, thin provisioning and compression provide storage efficiency while data deduplication typically provides no benefits due to how Oracle stores data. Although compression within the array can provide benefits, these may be limited if Oracle DBAs set compression within the database. Thus, thin provisioning is often the primary efficiency technology for reducing Oracle storage capacity needs.

HP has worked with Oracle on thin provisioning technologies, specifically evaluating thin-deduplication in a joint project with Oracle.³ The result is the ability for HP 3PAR storage systems to retain thinly provisioned volumes when used by Oracle. Whereas storage not specifically integrated with Oracle will grow fat over time, 3PAR thin volumes remain thinly provisioned through the use of the Oracle ASM reclamation process, which frees up unused portions of a thinly provisioned volume.

HP 3PAR Storage

The HP StoreServ platform is a storage family that scales from 2U midrange hybrid configurations of the 3PAR 8200, to the all flash 8450 and Tier-1 enterprise class 20850 systems all using a common architecture and set of data protection technologies. The additional 3PAR Federated Storage

³ HP 3PAR - Oracle Thin Provisioning: <http://www.oracle.com/technetwork/database/oracle-asru-3par.pdf>

capabilities enable an even greater number of enterprise level high-availabilities for use cases such as online migration during system upgrades, transparent system failover and remote BC/DR capabilities.

As outlined originally, the primary considerations of an Oracle DBA or application owner are minimizing risk while cost-effectively improving business results. Although there are costs associated with any technology, the financial benefits of solid-state are numerous, particularly with critical Oracle database applications. Two common hurdles limiting adoption of flash arrays for Oracle have been the perception of high cost and concerns around the lack of enterprise reliability and data protection with new storage vendor's products.

Shown below in Table 1 is an overview of Oracle DBA storage considerations, with HP 3PAR features and benefits in Oracle environments.

Consideration	HP Feature	Benefit
High Availability	3PAR Oracle Recovery Manager	Oracle RMAN integration and recovery
Enterprise Reliability	Tier-1 storage, 3PAR platforms scale from 2 – 8 controllers	Tier-1 - Enterprise HA platform
BC/DR Capabilities	3PAR Application Suite for Oracle and replication suite	Enterprise replication supports multiple BC/DR options
High Performance	All Flash, Adaptive Optimization and Priority Optimization	Accelerate entire application or specific portions
Enterprise Scale	3PAR Federated Storage	Enables scale-out enterprise class storage
Risk of Change	3PAR Online Import	Transparent migration to HP 3PAR from HP, EMC and other arrays
SSD Optimized Replication	3PAR Async Streaming Replication	Minimize recovery procedures without reducing performance
Storage Efficiency	3PAR Data Reduction Technologies	Oracle thin reclamation and data deduplication

Table 1: Oracle DBA Storage Considerations

In addition to the general capabilities of the HP 3PAR platform, there are several specific points of integration between HP 3PAR and Oracle, including:

- Oracle ASM Storage Reclamation Utility
 - Enables Oracle ASM to reclaim unused storage on HP 3PAR systems
- HP 3PAR Thin Persistence
 - Zero write auto detection with 3PAR ASIC assist, eliminates writing 0s
- HP 3PAR Recovery Manager for Oracle
 - Enables creating and managing application-consistent, array based snapshots of Oracle and Oracle RAC for rapid recovery
- HP 3PAR Application Suite for Oracle
 - Integration with Oracle Active Data Guard
 - Integration with HP Data Protector and Symantec NetBackup to use array snapshots as Oracle DR protection points

Summary

There should be no debate over whether solid-state technology can accelerate Oracle applications. The hesitations to date have typically centered on two perceptions: first, that solid-state is expensive and second, that solid-state systems lack enterprise capabilities. The idea that solid-state is more expensive than traditional storage is inaccurate when the financial value is measured correctly. The concern over lack of enterprise capabilities in new solid-state systems is a consideration that can be addressed by using enterprise class storage systems with all-flash optimizations.

Evaluator Group Comments: Minimizing the leading cause of inefficient database utilization, namely storage, is a project that can provide significant ROI benefits. In order to minimize risks, storage should provide enterprise capabilities in addition to high performance. Every organization should have a portion of their Oracle databases on solid-state storage and should consider if all-flash storage can provide additional benefits.

HP 3PAR storage systems are able to provide high-performance coupled with data efficiency and data protection features required in Oracle environments. With support for flash, hybrid and capacity systems, DBAs and architects can optimize deployments to use all-flash for primary databases, with secondary recovery points replicated to capacity based storage for greater efficiency. 3PAR technologies such as federation, thin provisioning along with Oracle integration for fast recovery, Oracle ASM thin persistence and optimized replication provide measurable value to Oracle DBA's.

IT administrators evaluate storage platforms based on their ability to deliver enterprise level data protection at the most cost effective price. DBAs and application owners whose business operations depend on Oracle databases, value enterprise class features, but also demand performance. HP 3PAR's all flash systems are able to provide the features IT and database admins require, while enabling application owners to achieve higher productivity with all-flash performance that is cost effective.

HP Publication: 4AA6-0829ENW

About Evaluator Group

*Evaluator Group Inc. is dedicated to helping **IT professionals** and vendors create and implement strategies that make the most of the value of their storage and digital information. Evaluator Group services deliver **in-depth, unbiased analysis** on storage architectures, infrastructures and management for IT professionals. Since 1997 Evaluator Group has provided services for thousands of end users and vendor professionals through product and market evaluations, competitive analysis and **education**. www.evaluatorgroup.com Follow us on Twitter @evaluator_group*

Copyright 2015 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, inconsequential 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. All other trademarks are the property of their respective companies.

This document was developed with HP funding. Although the document may utilize publicly available material from various vendors, including HP, it does not necessarily reflect the positions of such vendors on the issues addressed in this document.