



Integrating Flash Memory into Popular Applications

Server-Side Flash and Software Use Cases

Earle F. Philhower, III <earle.philhower.iii@hgst.com>

Sr. Manager, Solutions Lab

HGST, a Western Digital Company


Introduction - Who is HGST?



- Founded in 2003 through the combination of the hard drive businesses of **IBM**, the inventor of the hard drive, and **Hitachi**, Ltd (“Hitachi”)
- Acquired by **Western Digital** in 2012
- More than **4,700** active worldwide patents
- Develops innovative, advanced **hard disk drives**, enterprise-class **solid state drives**, external storage solutions and services
- Delivers intelligent **storage devices** that tightly integrate hardware and **software** to maximize solution performance


Deploy Islands of Flash or Server-Side Software?

Traditional Flash Deployments



No shared access
No built-in high availability
Higher costs

Server-Side SW Deployments



Shared access
Built-in high availability
Scalable aggregate bandwidth
Lower costs



Local Flash Simply Cannot Solve Some Business Problems...


- **Server Consolidation:**
“How can I use flash to deploy fewer servers?”
- **Cost Avoidance:**
“How can I share the data on flash inside a single server?”
- **High Availability:**
“How can I get my data when a server fails?”
- **Investment Protection:**
“How can I maximize the investment in my legacy storage?”

ANSWER: Flash + Software = Virident

Virident = HGST Server-Side Software



HGST's Virident Server-Side Software



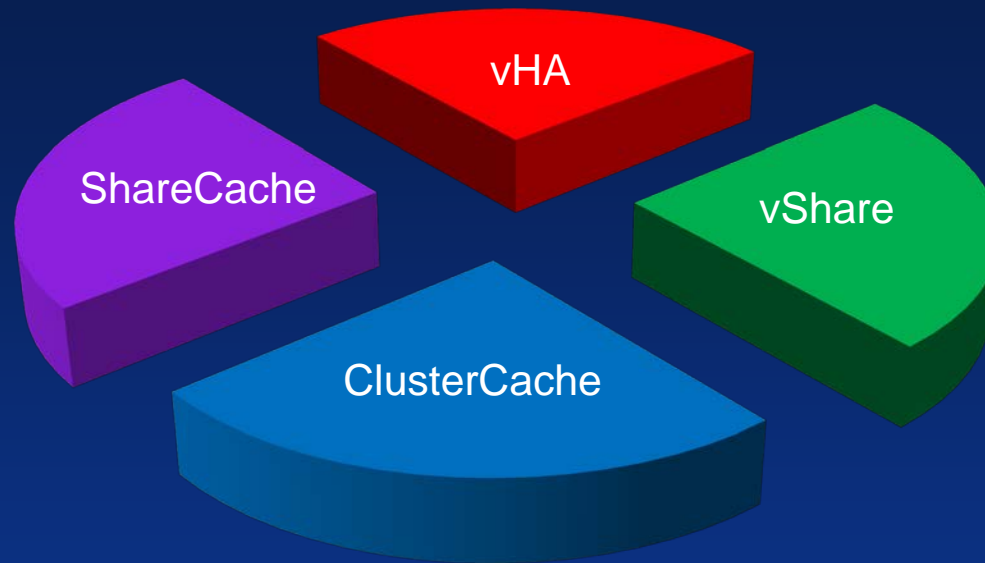
Flash-Aware Performance Tier Optimized for Databases and Virtualization Applications

- Performance Storage Close to Apps
- Highest Performance
- High Efficiency
- Lower Costs

- Direct Attach Storage
- Shareability
- Capacity
- Manageability
- Availability



Virident Server-Side Software Components



Data Protection: vHA enables high-performance synchronous mirroring



Shared Storage Resources: vShare offers shared access to partitions of a remote PCIe SSD



Investment Protection: ClusterCache provides transparent block cache for SAN & DAS



Application Acceleration: ShareCache delivers a common flash pool



vHA - Low Latency Synchronous Replication



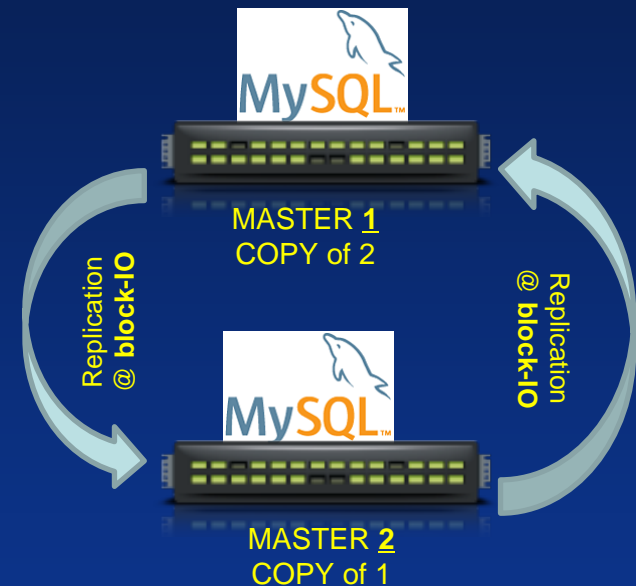
- High performance **synchronous mirroring** optimized for **PCIe SSDs**
- **Full and incremental sync capabilities**
- **Integrate with Cluster managers to enable automatic failover**
- Use with scale-out databases to facilitate **server consolidation**

vHA + MySQL = Active :: Active Database Replication, Consolidation

Before



After



- Consolidate HA servers 2:1, significant CAPEX reduction
- Much faster replication thanks to block-level vHA replication
- All nodes are active masters and passively replicate volumes of peer at block level
- On failure of a node, paired node runs 2 instances until recovery



Virident vHA Business Impact



Zero Downtime. vHA delivers granular recovery points which means no lost transactions where time is money.



Server Consolidation. RDMA-based synchronous replication delivers faster updates for scale-out servers and databases reducing need for extra RAM and extra servers.



High Availability. vHA delivers higher availability avoiding costly downtime for mission critical applications.



vShare – Shared Flash Building Block



- High performance remote data access optimized for PCIe SSDs
- Low-latency RDMA-capable networks

- Allows for consolidation flash devices in a central high capacity server
- Ability to share flash devices across multiple servers
- vShare + vHA = High Performance Server-SAN

vShare - Oracle RAC with only Local Flash Storage

Before

High-End SAN
Fibre Channel



Fibre Channel



Database stored on SAN

After

Database stored locally on flash
SAN repurposed



1.1TB FlashMAX II in each server
Using vShare to make globally visible


- 5X improvement in Application performance
- No server upgrade required, same base hardware
- SAN upgrade savings on hardware and software licenses



vShare - Oracle RAC Validation

HGST FlashMAX II and vShare received Oracle Validation on April 8, 2014. The validated stack includes:

- Oracle 12c (Single Instance & Real Application Clusters)
- Oracle Linux 6 Update 5 using ASMLIB
- 2x HP ProLiant DL380p Gen8
- 2x HGST FlashMAX II 1.1TB
- Mellanox IS5023 QDR Infiniband Switch
- 1 X Dual Port Mellanox MT27500 [ConnectX-3] 40GB



The solution enables shared server-side flash storage for Oracle RAC. This is the first configuration validated by Oracle using PCIe Flash.



Virident vShare Business Impact



Cost avoidance. Share flash in a common storage pool and buy flash only when you need it.



Investment protection. Accelerate Oracle RAC without costly server and SAN upgrades.



Improved Server Utilization. Powerful CPUs can leverage flash storage using Ethernet or Infiniband.



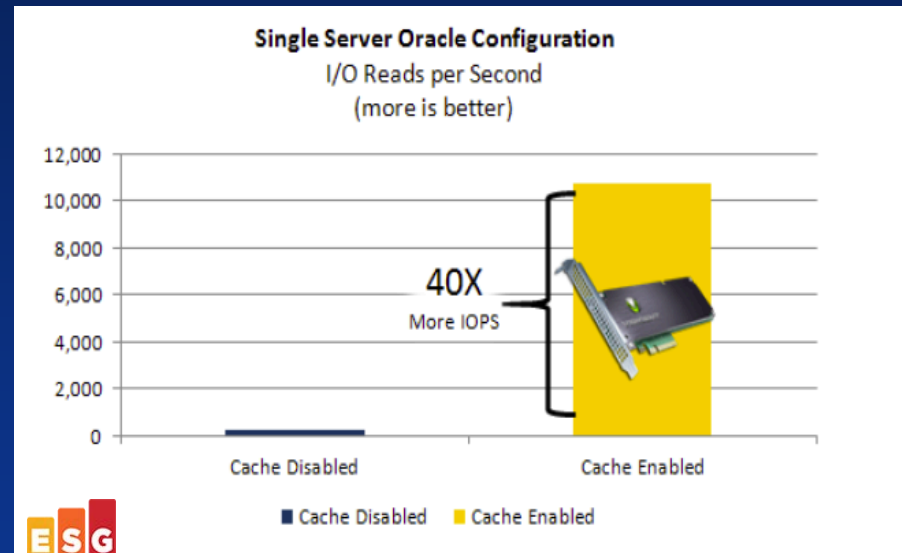

ClusterCache – Flash-aware Integrated Block Cache

- Flash-aware block caching module with multiple modes: **Write-through, Write-around, Write-back**
 - **Flexible sharing** of flash capacity among multiple caches
 - **Thin provisioned**, resizable at run time
-
- Flash-aware means **higher performance and higher endurance**
 - **Co-exists** with existing backend storage and IT storage management policies
 - Leverage across **all applications and hypervisors**



ClusterCache for Oracle Single Instance

ClusterCache





Virident ClusterCache Business Impact



Investment Protection. Longer life for flash media due to optimized algorithms avoiding double mapping of block addresses




Cost Avoidance. Drive significant improvements in throughput and IOPs for Oracle Single Instance avoiding SAN and server upgrades.



Server Consolidation. Accelerate existing databases avoiding purchases of more servers or more RAM.

Virident ShareCache

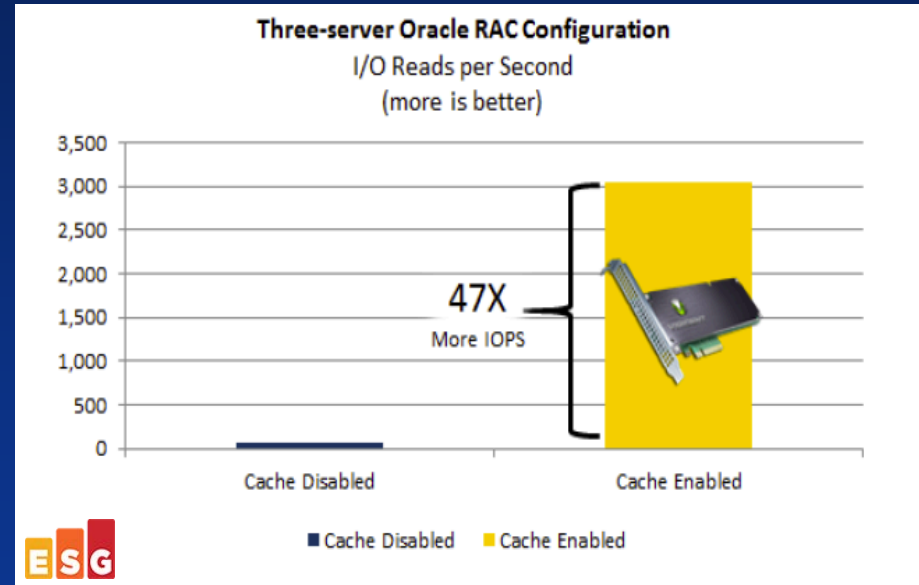



Customer Benefits:

- Single SSD gives caching for multiple servers
- Reduced capital costs
- Reduced footprint
- Optimized use of FlashMAX SSD capacity
- Unique server-side caching for Oracle RAC

ShareCache for Oracle RAC

ShareCache





Virident ShareCache Business Impact



Save on SAN upgrades. 10x improvement in throughput and 47x increase in IOPs using a 3-node ShareCache to front-end an Oracle Real Application Cluster (RAC).



Investment Protection. ShareCache works transparently with existing SAN or DAS storage.



Cost Avoidance. Buy only the flash that you need. Manage it as a common pool. Add more as necessary, but don't over-purchase.



Summary

Full flash benefits at the application and business level need SSD + software

- Virident Server-Side software provides:
HA, Sharing, Clustering, and Caching
- Resulting in application and business benefits:
Server Consolidation, Service Availability,
Investment Protection, at Highest Performance.

Stop by our booth for more information, or visit

<http://www.hgst.com>



Capacity & Trademark Statements

- One gigabyte (GB) is equal to one billion bytes, one terabyte (TB) equals 1,000GB (one trillion bytes), and one petabyte (PB) equals 1,000TB (one quadrillion bytes) when referring to hard drive or solid state drive capacity. Accessible capacity will vary from the stated capacity due to formatting and partitioning of the drive, the computer's operating system, and other factors.
- Ultrastar, MegaScale, HelioSeal, 7Stac, Virident, and FlashMAX are trademarks or registered trademarks of HGST, Inc.
- HGST trademarks are intended and authorized for use only in countries and jurisdictions in which HGST has obtained the rights to use, market and advertise the brand. HGST shall not be liable to third parties for unauthorized use of this document or unauthorized use of its trademarks.
- All other company names, products or trademarks used in this presentation are the marks of their respective owners, and not HGST.