



# HGST Virident™ Solutions 2.0

## Software Modules

- HGST Virident Share: Shared access from multiple servers
- HGST Virident HA: Synchronous replication between servers
- HGST Virident ClusterCache: Clustered caching for SAN or DAS and Oracle RAC
- HGST Virident Space: Clustering & logical volume management

## Key Benefits

- Real-time replication for scale out databases allows for server consolidation and cost savings
- Investment protection for existing SANs with shared caching using SSDs
- Shared flash storage lets you purchase only what you need. Add more as your performance requirements grow

## Overview

Accelerate your applications while minimizing the costs of centralized storage. Protect your investment in existing SAN and NAS while gaining the benefits of blazing speed from PCIe Flash. With HGST Virident™ Software 2.0 you can drive new levels of TCO and ROI using clustering, high-availability, shared storage and caching in a server-side flash network.

Building a flash platform has never been easier. With new features like an enhanced Control Panel, 16-node clustering, and support for Ethernet, the HGST Virident Solutions 2.0 makes deploying PCIe Flash consistent with how you currently manage high-end SAN or NAS arrays.

Built on top of HGST's award-winning FlashMAX® PCIe SSD platform, the software provides specific modules available now. These include HA for synchronous replication; Share for shared PCIe storage; ClusterCache for accelerating existing SAN or DAS; ClusterCache with Share for Oracle RAC and HGST Virident Space for clustering and cluster volume management.

## Flexible Scaling

Scale flash storage with or independent of compute resources. Scale with devices from 550GB to 4.8TB. Scale out by adding server nodes as needed. Use standard servers of your choice. For Linux clusters, leverage HGST Virident Space to define any-to-any clustering and fail-over. Using Space you can also aggregate multiple volumes into a single logical volume up to 38TB in size and then carve this volume into logical units that match your application consumption avoiding the need to buy more flash than necessary.

## Reduce Costs

Reduce both capital and operational expenditures by eliminating proprietary SAN or NAS hardware and maintenance costs. Protect investments in your existing SAN infrastructure by accelerating its performance with Server-side flash cache. With features like HA, you can improve replication performance for scale-out databases and drive server consolidation by as much as 4:1. Using Share, now with Ethernet support, you can unleash the power of flash to servers without having to directly install SSDs on each server bus, delivering better storage utilization and lower cost.

## HGST Virident HA

HGST Virident HA provides a high-throughput, low-latency synchronous replication across servers for data residing on FlashMAX PCIe devices. If the primary server fails, the secondary server can automatically start a standby copy of your application using the secondary replica of the data.

Build high-performance failover clusters for KVM, Oracle, SQL Server, MySQL, and NoSQL databases. Easily integrates with Corosync, CMAN and Pacemaker.

### Cluster Monitor for HA

With an enhanced event manager, upper-layer fail-over tools like Corosync and Pacemaker now have state information for in-sync and out-of-sync notifications to better assure granular recovery.

### HA Provides Flash-Aware Incremental Re-Sync

Automatically resynchronize only the data that has changed. By using the underlying flash management layer to track all data changes, server performance overhead is reduced.

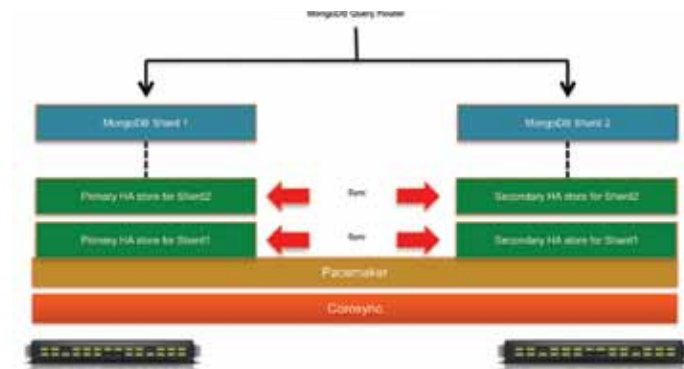
### Network Resiliency

Network multi-path capability embedded in the HA driver protects against network failures and simplifies system configuration.

### Efficient Scale-Out

NoSQL databases like MongoDB, Cassandra, Hbase and Oracle NoSQL use sharding, a technique to spread data arbitrarily across a number of inexpensive servers populated with SSDs. One of the challenges of supporting sharded databases is achieving real-time consistency of data across sharded servers. In most cases, NoSQL sharding implementations use an 'eventual consistency' approach that causes the overall performance of the architecture to slow down. One way to address this bottleneck is to add more RAM and more servers.

HA delivers a synchronous replication between sharded servers using over Infiniband or Ethernet that significantly accelerates the performance of shared consistency driving down the number of servers and the amount of RAM required. Not only does the overall performance of the scale out improve, but recovery time in the event of a server failure also improves.



### HGST Virident HA for KVM

Kernel-based Virtual Machine (KVM) is a popular open source solution for running multiple virtual Linux or Windows machines on Linux x86 servers. Using KVM each VM has its own private virtualized network card, disk, graphics adapter, etc. One of the challenges for KVM is redundancy. Leveraging HA, a fully redundant KVM environment may be achieved similar to VMware's VSAN.

## HGST Virident Share

Today's server-side flash is available as a direct attached storage device only. This means that the flash device is dedicated to a single server and often grossly under-utilized. The effective cost/GB for a dedicated card is often high and gets higher the lower the utilization.

HGST's breakthrough technology, Share, uses Remote Direct Memory Access (RDMA) over Infiniband or Ethernet to provide a server-side storage network. The benefits include high availability through replication and clustering along with improved resource utilization through sharing.

### Ethernet Support

Share now provides Ethernet-based remote hosts to be initiators to the shared cluster target, opening up the possibilities for any host to access shared flash.

### Low Latency

Share uses Remote Direct Memory Access to minimize latency overhead of data transfer over the network.

### Simplicity for Oracle RAC


Achieve high-availability for an Oracle RAC environment using Oracle's native ASM tool.

Scale Oracle RAC clusters without adding to back-end SAN.

### Cost Savings with Share

To consider the financial benefits of shared storage using Share, we calculate the desired flash capacity that will be allocated per server. Then we purchase the total number of cards required. In the case of current flash competitors, this is one card per server. In HGST's case with Share, fewer cards are required, as they will be shared across the server-side network.

Since a single HGST FlashMAX SSD can deliver more than 350,000 random 4K Random Reads, and hold 4.8TB of MLC flash, we need only two SSD devices to service a network of 6 servers requiring 1.6TB of flash each.



Share enables block-level shared access to FlashMAX devices from other cluster members.

Build high-performance active-active clusters, such as Oracle Real Applications Cluster (RAC), keeping the throughput and low-latency benefits of PCIe-attached flash.

Give access to flash resources for servers that do not have local flash, with minimal performance overhead.

### Oracle All Flash Scaling

Share works with Oracle RAC to deliver high-performance Flash leveraging Oracle's Automatic Storage Management (ASM). Using ASM, you can scale up from 2 to 64 nodes, each with 4.8TB FlashMAX SSDs to achieve an all-flash platform of more than 500TB.

## HGST Virident ClusterCache

Accelerate SAN or DAS storage by using FlashMAX devices as a high-performance PCIe-attached cache. ClusterCache improves applications like Single Instance Oracle with more than 40x improvement in IOPs and 7x increase in throughput. This means faster response times for business applications. Quicker access to data can give an organization a competitive advantage and result in greater productivity and profits.

### Investment Protection

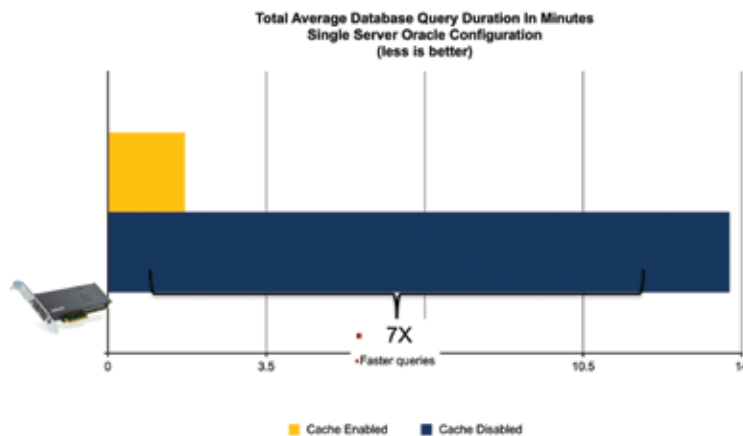
Use ClusterCache to achieve higher performance and longer flash lifetime compared to third-party caching solutions built on SSDs. Further boost caching efficiency by blacklisting processes, such as a backup application, that should not use the cache resources.

### Greater Flash Endurance

ClusterCache directly accesses and manages physical flash memory and eliminates double mapping of block addresses.

Create separate caches on a single FlashMAX device for multiple back-end storage LUNs. Choose between write-back, write-through, or write-around caching modes to fit the needs of your applications.

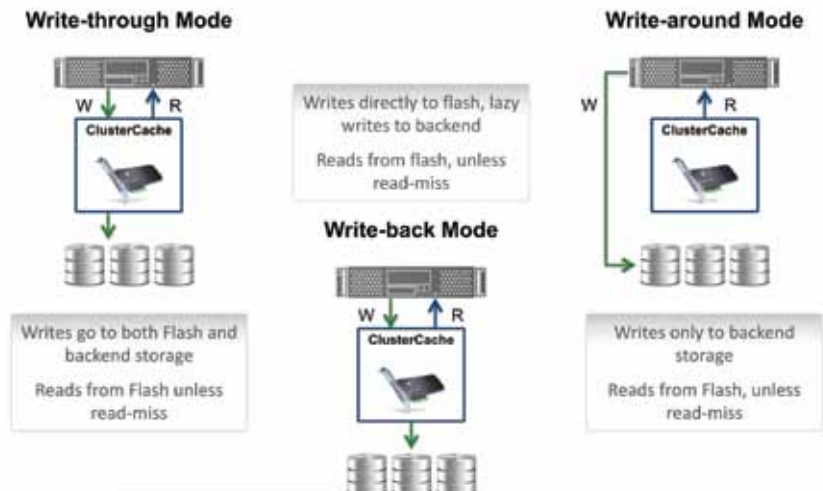
Working with the Enterprise Strategy Group (ESG) HGST was able to illustrate significant performance gains with ClusterCache enabled for Single Instance Oracle.



### Protect Existing Investments

ClusterCache works transparently in front of existing SAN or DAS storage, accelerating application performance without the need for costly infrastructure upgrades.

In addition to a 7x improvement in throughput and a 40x increase in IOPS, ESG and HGST showed that database query time went from approximately 14 minutes for the SAN to less than 2 minutes with ClusterCache.



## HGST Virident ClusterCache + Share for Oracle RAC

HGST offers the ability to share the resources of a FlashMAX II device in caching mode. ClusterCache combined with Share may be thin provisioned for use at run time, effectively making all of the capacity of the PCIe Flash device available for use as a cache. For customers looking to best optimize the resources of a PCIe card for caching, ClusterCache with Share provides a compelling solution.

### ClusterCache with Share Advantage

Leverage ClusterCache with Oracle RAC to accelerate an existing SAN without expensive SAN upgrades. ClusterCache allows for scaling up to 64 nodes and works with existing Oracle Automatic Storage Management (ASM).

### ClusterCache Compatibility

- Redhat Enterprise Linux 5, 6
- CentOS 5, 6
- Oracle Enterprise Linux 5, 6 (RH and UEK kernels)
- SUSE Enterprise Linux 10 SP3, SP4, 11 SP1, SP2, SP3

### Application Support

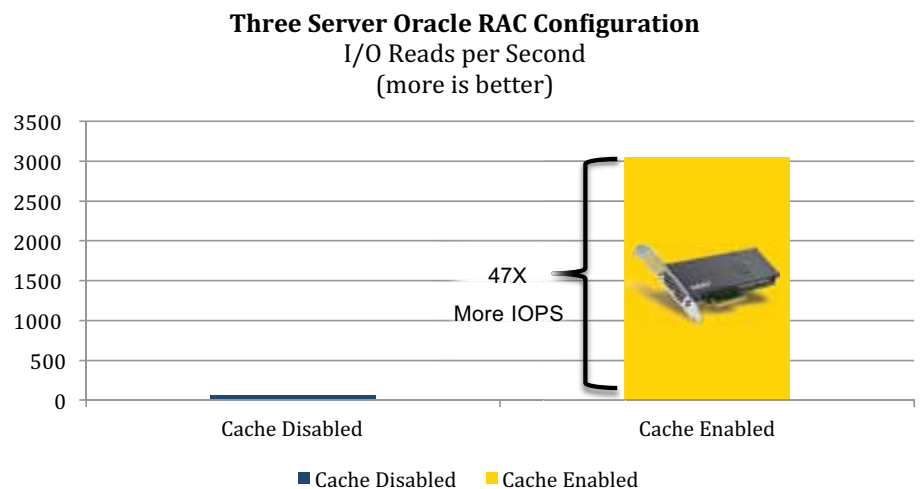
Oracle RAC 11.2 and higher

### Boost Performance or Oracle RAC

Oracle Real Application Clusters (RAC) environments cannot benefit from PCIe Flash since RAC requires shared storage. With ClusterCache and Share, Oracle RAC can take full advantage of the high performance and low latency of FlashMAX PCIe SSDs. ClusterCache with Share helps minimize the cost of high availability Oracle RAC clusters while providing the highest available flash performance and reducing the need for expensive SAN storage solutions. It effectively converts the HGST FlashMAX SSDs into a shared storage infrastructure capable of full integration with Oracle RAC and managed by Oracle Automatic Storage Management.

Using high-speed RDMA over Infiniband or Ethernet, multiple servers can access a shared cache on a central server where Oracle sees a shared storage device. This architecture can be configured with a single server cache or multiple server caches as needed for capacity.

ClusterCache with Share delivers significant performance advantages compared to a configuration where no PCIe Flash is resident. Using Oracle Enterprise Manager it is possible to see 10X improvement in throughput and more than 47X increase in IOPs as certified by the Enterprise Strategy Group and illustrated below.



## HGST Virident Space

HGST Virident Space is the newest member of the software family. Space is a Logical Volume Manager similar to the Volume Managers that are part of high-end SAN or NAS storage arrays. With Space you can cluster up to 128 servers and 16 FlashMAX SSD devices into a single logical volume and then carve up smaller volumes to the exact size needed by each application, optimizing flash resources for the most cost-effective flash deployment.

### Flexible Networking

Space works with Ethernet or Infiniband to create clusters giving you a choice for network infrastructure.

### Space for Linux Clusters

With the introduction of HGST Virident Space, revolutionary new capabilities for Server-SAN can be achieved. Never before has there been technology for PCIe SSDs that allows for clustering and aggregation of Flash that resembles a SAN Array Volume Manager. By layering this volume management capability with sharing and high availability, the next generation performance tier of storage is emerging.

With Space for Linux, you can define an any-to-any relationship between 128 hosts and 16 FlashMAX SSDs. Each of these nodes can have a FlashMAX 4.8TB Capacity SSD. Replicating any-to-any delivers a massive 38.4TB all-flash cluster.



### Investment Protection

Drive flash costs down by purchasing on the amount of flash you need and add FlashMAX devices as your requirements grow.

Space allows the cluster to grow dynamically by adding a new FlashMAX SSD into a server or by adding a new server with an SSD. The devices can move between servers for easy upgrades.

### Space for KVM

Failover is automatic and transparent to any application using the cluster. For Kernel-based Virtual Machine (KVM) environments, Space turns the KVM environment into a virtual SAN where a pool of Flash may be used for application acceleration.

### Space for GFS

Use Space in conjunction with Red Hat Global File System to create an all-flash shared file system with blazing performance and granular volume management.



© 2014 HGST, Inc. 3403 Yerba Buena Road, San Jose, CA 95135 USA, Produced in the United States. 8/14 All rights reserved.

FlashMax is a registered trademark and Virident is a trademark of HGST, Inc. and its affiliates in the United States and/or other countries.

Contact HGST for additional information. HGST shall not be liable to third parties for unauthorized use of this document or unauthorized use of its trademarks.

References in this publication to HGST's products, programs, or services do not imply that HGST intends to make these available in all countries in which it operates.

Product specifications provided are sample specifications and do not constitute a warranty. Information is true as of the date of publication and is subject to change. Actual specifications for unique part numbers may vary.

For more information, visit [www.hgst.com](http://www.hgst.com).