

FlashGrid Uses NVMe* SSDs to Build High Performance Software-Defined Storage for Oracle* Database Clusters

FlashGrid® Storage Fabric uses locally attached Intel® SSD Data Center Family drives with NVM Express* for lower cost, higher performance software-defined storage for Oracle RAC*.



Executive Summary

Many organizations depend on complex, often expensive database solutions like Oracle Real Application Clusters* (RAC) in order to provide performance and availability for their business-critical databases. Part of both the expense and the complexity of these solutions is driven by their need for high-performance storage, often supplied in the form of storage area networks (SANs).

Storage-management software from FlashGrid, coupled with Intel® technologies, can provide an alternative to a SAN that is more affordable and higher performing, in the form of software-defined storage (SDS) hosted on the same servers that run the database cluster. The deployment of the solution demonstrates not only that SDS is ready for vital database workloads, but also a means by which IT organizations can use fast, affordable Intel® Solid-State Drives (SSDs) with NVM Express* (NVMe*) without being locked into more expensive, proprietary flash appliances or lower-performing SAN implementations.

Meeting the Challenges of Mission-Critical Storage

Providing the storage necessary for business-critical databases like Oracle RAC can be challenging for enterprise IT organizations. Simply getting storage to be responsive enough to meet users' performance requirements can be a significant commitment: Oracle RAC, for example, often requires expensive, complex infrastructure elements, such as SAN storage, to meet these performance targets. And going the SAN route carries challenges all its own. For example, traditional SAN deployments do not scale easily. Beyond the difficulty of scaling, emerging all-flash array-based appliances can be cost-prohibitive, particularly when paired with already expensive Oracle RAC deployments.

Local Storage That Looks Like a SAN

There are alternatives to investing in a SAN storage back end for Oracle RAC, however. FlashGrid delivers a storage solution that uses FlashGrid® Storage Fabric software and Intel SSDs with NVMe to build an SDS solution hosted locally on servers running Oracle RAC.

The Oracle RAC clusters can be connected by a high-speed InfiniBand* network, with each node housing Intel SSDs with NVMe.

8X BETTER PERFORMANCE¹

FLASHGRID® STORAGE FABRIC SOFTWARE
AND INTEL® SSDS WITH NVME*

6X LOWER COST¹

The cluster transparently uses the existing Oracle Automatic Storage Management* (ASM) framework for database synchronization between the nodes. Through the FlashGrid Storage Fabric software, all Intel SSDs appear as local drives to Oracle RAC.

The FlashGrid solution demonstrates that storage clusters using Intel SSDs with NVMe can deliver much higher performance for typical transactional workloads compared to all-flash arrays. Enterprise data centers can use FlashGrid Storage Fabric and Intel SSDs with NVMe in mission-critical database deployments—like Oracle RAC—as a faster, more affordable storage option for online transaction processing (OLTP) workloads.

Higher Performance Storage at a Lower Cost

Locally installed SSDs provide better performance at a lower price and with less complexity than flash-based SAN storage. Using local storage on a cluster's nodes can improve clustered-database performance more than eight times compared to appliance-based all-flash arrays.¹ Part of this performance improvement comes from hosting the storage media locally rather than on the network; local storage means no network latency to worry about. Much of the performance improvement also comes from the NVMe specification; Intel SSDs with NVMe provide six times greater performance than SSDs that use the Serial ATA (SATA) specification.²

Local storage can also improve the total cost of ownership (TCO) and scalability of storage. The overall capital costs for locally installed SSDs with NVMe are up to six times lower than for appliance-based all-flash arrays.¹ In addition, a FlashGrid solution that uses Intel SSDs with NVMe can scale more easily than traditional SAN or appliance-based all-flash array deployments. Each new cluster node comes with its own storage, which eliminates the contention and upgrade complexity typical to SAN storage environments shared between multiple servers.

Solution Details

The locally hosted SDS solution using FlashGrid and Intel compute, network, and storage technologies addresses the need to run Oracle RAC in a software-defined infrastructure (SDI) environment while reducing the cost of deployment. Tables 1 and 2 and Figure 1 provide details about a deployment of the solution at an international cloud-services provider.

Table 1. Hardware specifications for the FlashGrid and Intel deployment

Storage	Sixteen 2.5-inch Intel® SSD DC P3700 Series drives with NVMe*, 400 GB each
Network	Two ConnectX-4 VPI MCX456A-ECAT dual-port*
Server	Two Aquarius Server T50 D20 platforms*, 256 GB RAM each
Processor	Four Intel® Xeon® processor E5-2690 v4, 2.6 GHz, 14 cores/28 threads, 35 MB, LGA2011-3

Table 2. Software-stack details for the FlashGrid and Intel deployment

Operating System	Oracle* Linux 7.2*
SDS Software	FlashGrid® Storage Fabric 16.11
Database Software	Oracle Grid Infrastructure* and Oracle RAC*
Performance-Testing Tools	Calibrate_IO* HammerDB*



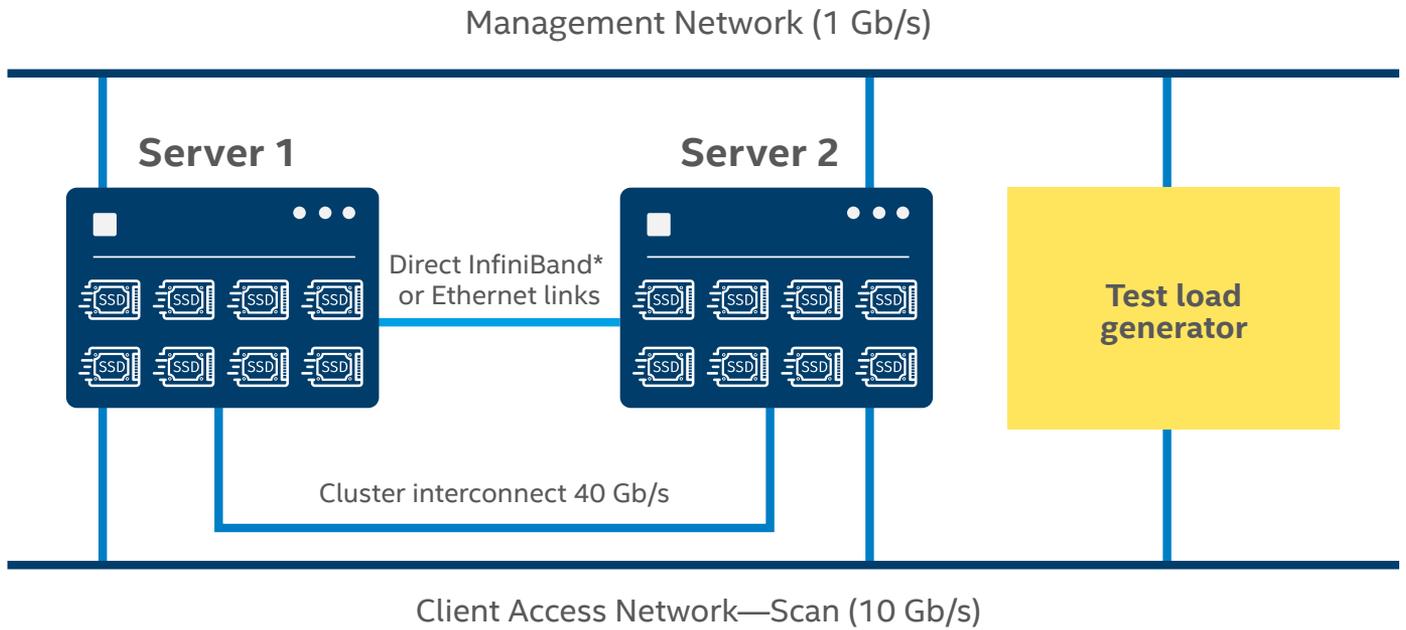


Figure 1. Architecture diagram of two servers, each with eight Intel® SSDs with NVMe*, connected by an InfiniBand* network connection over FlashGrid® Storage Fabric to support Oracle RAC*

Intel's Role

In the FlashGrid deployment, Intel technologies boosted the performance and scalability and helped reduce costs across all aspects (see Table 3).

Table 3. Intel® technologies' benefits for the FlashGrid deployment

Compute	Storage	Network
<p>Intel® Xeon® processor E5-2690 v4, 2.6 GHz (CM8066002030908)</p> <ul style="list-style-type: none"> Designed for architecting next-generation data centers running on SDI Supports increased scalability, automation, and orchestration capabilities across compute, storage, and networking workloads Meets performance and cost requirements for the required server configuration¹ 	<p>2.5-inch Intel® SSD DC P3700 Series, 400 GB (SSDPE2MD400G401)</p> <ul style="list-style-type: none"> Provides enterprise-grade reliability and 2.7 GB per second (GB/s) of read bandwidth per SSD, which is critical for running read-heavy analytics workloads that perform full table scans Helps lower IT costs and provide greater responsiveness¹ 	<p>Intel® I/O Expansion Modules, 10 gigabit Ethernet (GbE), with two enhanced small form-factor pluggable (SFP+) transceivers (AXX10GBNIAIOM)</p> <ul style="list-style-type: none"> High-performance network-infrastructure component Reduces latency and improved data throughput for increased efficiency

Moving from Proprietary Hardware to a Software-Defined Architecture for Mission-Critical Databases

The deployment detailed here demonstrates FlashGrid Storage Fabric software installed and configured locally on the same servers that run Oracle RAC. This storage-management software facilitates communication between Intel SSDs with NVMe and Oracle ASM without requiring any specialized configuration of the Oracle deployment. Deployments like this one, with FlashGrid software making use of Intel storage, compute, and network technologies, can be a viable option for businesses interested in investigating or transitioning to SDS. Moreover, this deployment demonstrates how IT organizations can take advantage of the performance and affordability of Intel SSDs with NVMe without being obligated to use siloed, appliance-based flash arrays or lower-performing SAN solutions.

Beyond providing a new option for database storage, this solution also addresses several existing difficulties typically associated with traditional all-flash-array SAN implementations for mission-critical databases. Using local SDS can reduce the overall complexity of the architecture

supporting a company's transactional workloads. Hosting storage locally on the same servers that run databases can significantly reduce implementation costs. And, because the database storage is local and uses the NVMe standard, this solution can provide substantial increases in overall input/output (I/O) throughput.

About FlashGrid

FlashGrid software allows large and small enterprises to run mission-critical databases on commodity compute, storage, and network infrastructures of their choice—in the data center or in the public cloud.

Learn More

To learn more about how Intel SSDs with NVMe can increase storage performance and utilization for your business while simultaneously helping to lower IT costs, visit intel.com/nvme.

The Solutions Library on the Intel® Builders home page in particular can help you find reference architectures, white papers, and solution briefs like this one that can help you build and enhance your data center infrastructure: <https://builders.intel.com/solutionslibrary>.

You can also follow us on Twitter by using [#IntelBuilders](https://twitter.com/IntelBuilders).

To learn more about cost-efficient, high-performance storage solutions for running mission-critical databases on standard hardware infrastructure, visit the FlashGrid home page at flashgrid.io.



¹ Based on internal FlashGrid testing conducted in May 2016 comparing the cost of 12 Intel® SSD DC P3700 Series 2-TB drives (three nodes in a cluster) with three-way mirroring across nodes and Oracle ASM* against a single Dell EMC XtremIO 4.0* 10-TB X-Brick* flash storage array. Source: FlashGrid. "Hyper-converged storage for Oracle RAC based on NVMe SSDs and standard x86 servers." May 2016. flashgrid.io/wp-content/uploads/resources/FlashGrid_Intel_P3700_SSD_wp.pdf.

² Intel. "NVMe vs SATA: Intel® SSD with NVMe Express* outperforms SATA SSD's." intel.com/content/www/us/en/solid-state-drives/intel-ssd-dc-family-for-nvme.html.

Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors. Performance tests, such as SYSmark* and MobileMark*, are measured using specific computer systems, components, software, operations and functions. Any change to any of those factors may cause the results to vary. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products.

Cost reduction scenarios described are intended as examples of how a given Intel-based product, in the specified circumstances and configurations, may affect future costs and provide cost savings. Circumstances will vary. Intel does not guarantee any costs or cost reduction.

Intel does not control or audit third-party benchmark data or the web sites referenced in this document. You should visit the referenced web site and confirm whether referenced data are accurate.

Intel technologies' features and benefits depend on system configuration and may require enabled hardware, software or service activation. Performance varies depending on system configuration. No computer system can be absolutely secure. Check with your system manufacturer or retailer or learn more at intel.com.

Intel disclaims all express and implied warranties, including without limitation, the implied warranties of merchantability, fitness for a particular purpose, and non-infringement, as well as any warranty arising from course of performance, course of dealing, or usage in trade.

Intel, the Intel logo, and Xeon are trademarks of Intel Corporation in the U.S. and/or other countries.

FlashGrid is a registered trademark of FlashGrid, Inc.

*Other names and brands may be claimed as the property of others.

© 2017 Intel Corporation.