WHITE PAPER

Oracle HCTS Confirms SanDisk's Lightning SSD Compatibility for Oracle Solaris and SPARC T5 Servers

March 2015

SanDisk®
a Western Digital brand

Western Digital Technologies, Inc.
951 SanDisk Drive, Milpitas, CA 95035

www.SanDisk.com

Table of Contents

Introduction

IT organizations across the world are facing critical challenges in management of data and fast and efficient delivery of services geared around business critical data. One easy and transparent way to improve I/O performance and reliability of any computing platform is to replace existing spinning disk drives with solid state drives (SSDs) . This paper presents an overview of the tests conducted to verify the operational compatibility of Lightning® SSDs from SanDisk® with Oracle® SPARC® T5 server platforms using the Oracle HCTS hardware compatibility test suite.

Due to the diversity in components and variations in hardware platform architecture, it is very important to confirm if a product is compatible with the target server architecture . Oracle HCTS (Hardware Compatibility Test Suite) is a freely downloadable application for verifying whether a system or hardware component is compatible with the Oracle SPARC T5-2 or x86 server platforms. Verification with this tool and audit of test results are important criteria enabling Oracle to list third party systems and components on the Oracle Solaris Hardware Compatibility List (Oracle HCL).

A comprehensive listing of certified compatible components for Oracle Solaris and SPARC T5-2 server platforms is listed at: www.oracle.com/webfolder/technetwork/hcl/index.html.

SPARC Server Platform

Oracle's SPARC T5-2 server with two processors supports 32 DDR3 memory DIMM slots which may be populated with 8GB, 16GB or 32GB DIMMs for a maximum of 1TB of memory per server . There is also eight low-profile PCIe 3 .0 slots and six 2 .5" form factor bays for disks or SSDs .This server is based on the SPARC T5 processor and is optimized for small to midsize web-tier workloads, as well as database and enterprise applications . It is fast, secure, and ideal for cloud infrastructures .

Lightning SSDs

SanDisk offers one of the leading flash memory storage solutions, SSDs that are increasingly used for emerging applications such as cloud storage, big data analytics, and RDBMS (relational data base management systems) . SanDisk's Lightning SAS interface SSDs incorporate a suite of enterprise class features and write endurance enhancements that make cost-effective multi-level cell) (MLC) NAND fully usable for demanding enterprise applications and database workloads.

Test Configuration Details

The main focus of this effort is to confirm functional interoperability of SanDisk SSDs with Oracle T5 platforms. Some basic performance test results with various workloads are also presented. The certification tests cover only 64-bit architectures.

The following server platform was tested:

Server: SPARC T5-2 server, dual processor, 16 cores each, clock 3,599 MHz
Server Firmware: Open Boot PROM (OBP) 4.36.1
Operating system: Oracle Solaris 11.2
Test Application: Oracle HCTS version 5.6

Model Name	Part Number	Server Platform	OS version	Firmware
Lightning WI 200GB, LB206S	SDLB6HS-200G-00	SPARC T5-2	Oracle Solaris 11.2	P325
Lightning MU 800GB, LB 806M	SDLB6HM-800G-00	SPARC T5-2	Oracle Solaris 11.2	P325

Figure 1: SSDs Tested

Two different Lightning family drive models were tested individually as follows: Lightning 200GB[1] Write-Intensive (WI) and Lightning 800GB Multi-Use (MU) . Both drives incorporate a proprietary SanDisk flash controller, based on a parallel processing architecture, to deliver tested high performance and proven reliability . These SAS SSDs have been specifically designed for mission-critical, 24/7 operation .

The Lightning WI drive variant has very high write endurance of 25 drive-writes-per-day (DWPD)[2] endurance, making it ideal for high data rate logging or any highly write-intensive workload . The Lightning MU model, 10 DWPD endurance, is focused on applications with a typical workload of 70/30 read/write mix .

SanDisk also markets Lightning RI (Read-Intensive) which was not tested for this paper . The Lightning RI model is targeted at high capacity, mostly read content applications (1 DWPD or less) such as video on demand, with cost savings compared to the two higher write endurance models previously mentioned.

Test Setup and Results

The SPARC T5-2 server was factory installed with Solaris 11.2 and Oracle HCTS 5.6 was downloaded and installed on the system to exercise storage tests validating Lightning SSD compatibility .

The form factor of the drives was checked and found compatible with the Oracle T5-2 SPARC T5-2 server .

The **`prtconf -pv`** command was run to determine if the SanDisk SSDs could be successfully recognized by Oracle Solaris 11.2. The logs were checked and found to show proper device driver attachment for the drives .

Figure 2 shows the test sequence .

HCTS Test Type	Test Description	Test Result
mpflip	Detects all processors that are on-line	Passed successfully
Storage probe test	Probes all storage devices connected to the T5-2	Detected target Lightning drives
misabuf	Reads and writes to raw disk	Passed successfully
dd	Writes and reads records to device	Passed successfully
Bonnie test	Runs the disk and file systems I/O benchmark	Passed successfully
Fs_Stress test	Stresses file systems by writing and reading various block sizes to file system	Passed successfully
vdbench tests	Performs disk and file system tests	Passed successfully

Figure 2: Test Sequence

ZFS File System Performance Tests

Performance measurements were done using SanDisk test tools . The drives under test were partitioned and mounted with a ZFS file system using the entire user capacity of the drives . The purpose of the testing was to fully characterize realistic drive performance under varying workload conditions, using the native file system as opposed to raw device I/O .

The following example steps prepare a drive to begin performance testing:

1 . The SSD to be tested was physically inserted into the server system .
A ZFS file system and storage pool was created, for example, by running the following command:

```
# zpool create -f test c3t5001E82002637BEAd0
```

This setup is extremely straightforward and you don't have to do anything else to immediately start using the storage pool. The pool just created will be automatically mounted by default to /test . You also do not have to add an entry to something like the fstab file to make the mount persistent across boot cycles .

2 . The following command was used to confirm that the pool was created and mounted:

```
# df -h /test
```

Output:

```
Filesystem Size Used  Available Capacity Mounted on
test       183G 31K   183G      1%       /test
```

I/O Workload Summary

The following is a description of the I/O workload that was used:

- Precondition the drive
  - Two complete drive capacity writes, 128KB block sequential 100% write
  - One hour 4KB 100% random writes
- Performance measurements begin from here:
  - Random workload sweeps with 12 parallel I/O threads
    - 0% Write, 100% read access mode
      - Queue depths 1, 2, 4, 8, 16, 32, 64, 128
        - For each queue depth above, run block sizes 512B, 2KB, 2KB, 4KB, 8KB, 16KB, 32KB, 64KB, 128KB, 256KB, 512KB, 1024KB
    - 30% write, 70% read access mode
      - Repeat same queue depths and block sizes as above
    - 100% write, 0% read access mode
      - Repeat same queue depths and block sizes as above
- Sequential workload sweeps
  - Repeat same pattern as random workload above except for sequential access mode
    - Representing three additional access modes

The above workloads result in a total of six basic access modes:

- Three random modes
  - 0% write 100% read
  - 30% write, 70% read
  - 100% write, 0% read
- Three sequential modes
  - 0% write 100% read
  - 30% write, 70% read
  - 100% write, 0% read

Performance Test Script Details

Performance measurements were done using tools developed by SanDisk. This was done to validate performance of the SSD in the SPARC T5-2 server with the Solaris 11.2 OS. The basic logic of the test is implemented as a bash script and uses the utility fio as the workload generator to incrementally sweep block sizes, queue depths, and access modes as described in the I/O Workload Summary section above.

The fio utility, for each workload, spawns 12 parallel I/O threads to exercise the storage device and record several performance parameters including input/output operations per second (IOPS), bandwidth, read and write latency, and CPU utilization. Each basic workload run is 60 seconds in duration, with an average value log file saved for each. This results in 576 individual performance logs that are consolidated into six summary files, one for each access mode.

The entire test process takes from 16 to 24 hours depending on several factors such as the capacity and performance of the storage devices, performance of the test system as well as other factors. When the workload phase is completed, the six summary files are plotted using a simple bash script and the gnuplot utility. This results in the output of 12 chart files in the .png format, representing detailed IOPS and bandwidth results for the six access modes. Each of the six access mode charts show performance data for 96 combinations of block size and queue depth.

The following selected chart samples show the IOPS or MB/s bandwidth performance obtained from a single Lightning SSD setup as a zfs storage pool. In this case the charts represent the tested performance of Lightning MU 800GB, model LB 806M for the selected workloads.

The chart in Figure 3 shows peak sequential bandwidth at an astonishing 2 .25GB/s at all queue depths and 512KB block sizes . The ZFS file system has a very sophisticated set of caching algorithms that try to cache both the most frequently used data, and the most recently used data, while adapting the balance between the two in real time .

This is a graphic example of the read performance that can be achieved using SSDs as the Level 2 Adaptive Replacement Cache (L2ARC) for ZFS file systems . SSDs form an excellent layer of cache between the ZFS RAM-based Adaptive Replacement Cache (ARC) and the actual storage devices and are much cheaper on a per-GB basis than DRAM.

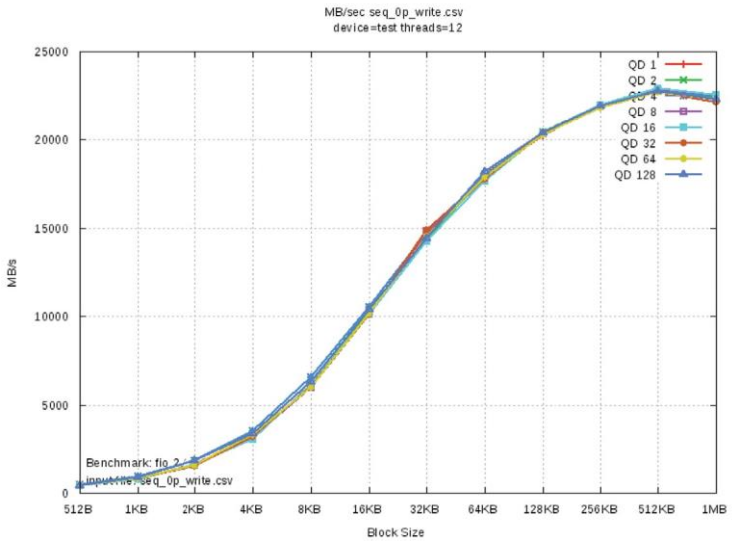

Figure 3: MB/s performance: sequential access 0% write, 100% read

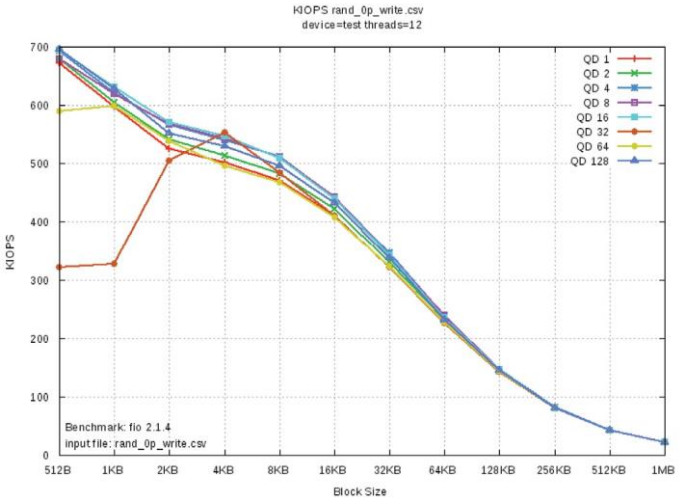

Figure 4: IOPS performance: random access 0% write, 100% read

Figure 4 shows the IOPS performance obtained during random access (0 percent write, 100 percent read) from a single SSD (SanDisk's Lightning MU 800GB model LB 806M) set up as a ZFS storage pool .

As shown in Figure 4, peak IOPS is just under 700,000 at all queue depths except 32 and 64, at 512B block sizes . SSDs can deliver orders of magnitude better IOPS than traditional hard disks, and they're much cheaper on a per-GB basis than RAM . They form an excellent layer of cache (L2ARC) between the zfs RAM-based ARC and spinning disk storage.

Figure 5 shows the MB/sec performance obtained during sequential access (30 percent write, 70 percent read) from a single SSD (SanDisk's Lightning MU 800GB model LB 806M) set up as a ZFS storage pool .

Peak bandwidth is between 700 MB/s to 830 MB/s for all queue depths and block sizes greater than 256 KB . This mixed read/write workload is typical of most on-line transaction processing (OLTP) applications . The ZFS file system, with some SSD storage available, provides unprecedented caching benefits for both reads and writes without the need for lots of hard disk drive (HDD) spindles as demonstrated by traditional all spinning disk architectures .

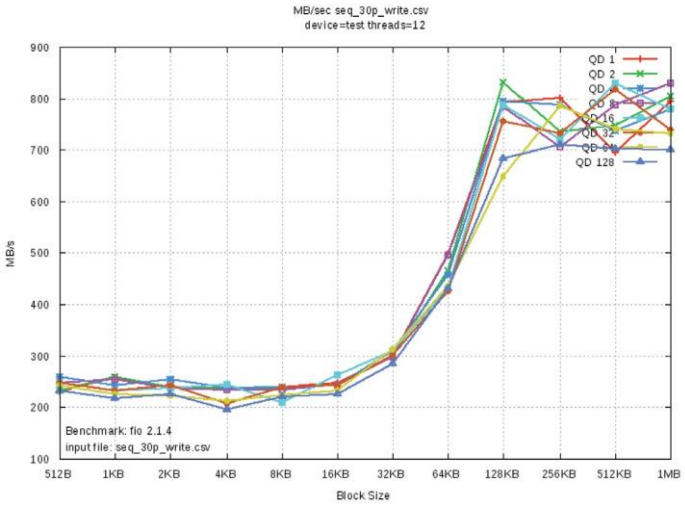

Figure 5: MB/s performance: sequential access, 30% write, 70% read

Conclusions

SanDisk's Lightning SSDs were tested with the Oracle T5 SPARC T5-2 server and Oracle Solaris 11 .2 and were found to be functionally interoperable . All the required HCTS tests passed successfully and the drives were submitted to the Oracle for listing as a certified component on the Oracle Solaris HCL. The specific HCL entries for the tested SanDisk's Lightning SSD products are listed under the "Hard Drive" category and can be viewed at: www .oracle .com/webfolder/technetwork/hcl/data/s11ga/components/views/hard_drive_all_results .page1 .html

Additionally, Oracle Solaris OS 11 .2, with the use of strategically placed SSD storage devices such as SanDisk's Lightning products, can provide substantial application performance improvements along with cost savings compared to alternatives like large numbers of spinning disks and maximum population of DRAM .

1. 1 GB = 1,000,000,000 bytes. Actual user capacity less.

2. The earlier of (I) three (3) years or (II) the time at which such Product unit's total cumulative data recorded on such unit at its S.M.A.R.T. attribute (log page 2Fh, parameter 25) or log page 11h, parameter 1 (Percentage Used Endurance Indicator) reaches 100% of the product units endurance value. For more warranty information, please visit kb.sandisk.com.