Red Hat Performance Briefs

Accelerating Oracle Database Performance on Red Hat Enterprise Linux 6

Fusion-io PCIe SSDs deployed as *transparent* front-end cache speed up decision support workloads by up to 2x compared to legacy storage

Version 1.0

September 2012

redhat.

1801 Varsity Drive
Raleigh NC 27606-2072 USA
Phone: +1 919 754 3700
Phone: 888 733 4281
Fax: +1 919 754 3701
PO Box 13588
Research Triangle Park NC 27709 USA

The GPG fingerprint of the security@redhat.com key is:
CA 20 86 86 2B D6 9D FC 65 F6 EC C4 21 91 80 CD DB 42 A6 0E

Send feedback to refarch-feedback@redhat.com

Table of Contents

 4
1 Executive Summary.......... 4
2 Test Configuration.......... 5
3 Testing Methodology.......... 8
4 DSS Performance Results.......... 9
5 Conclusion.......... 13
6 References and Further Information.......... 14
 Appendix A: Revision History.......... 15

Red Hat Enterprise Linux 6
with Fusion-io directCache

1 Executive Summary

This performance brief examines the performance characteristics of Fusion-io directCache modules on Red Hat Enterprise Linux 6 for an Oracle Database Decision Support System (DSS) application. Fusion-io directCache software makes Fusion-io PCIe SSD devices available, not as raw disks, but as a front-end cache to *existing* storage. Fusion-io directCache inserts itself between existing applications and storage and requires no significant modifications to the existing storage, file systems, data, or application while allowing to reap significant performance benefits.

Red Hat partnered with Fusion-io for this effort. Fusion-io provided the hardware and tuning guidelines for testing. These tests show different use cases where the Fusion-io SSDs and directCache software can obtain significant performance gains for Oracle Database 11*g* R2 running on Red Hat Enterprise Linux 6.

Table 1-1 presents a summary of the maximum improvements Fusion-io directCache achieved over legacy Fibre Channel (FC) storage in the testing conducted for this performance brief. Using Fusion-io directCache delivered up to 1.94 times the Aggregate DSS Query Rate (per hour) and 2.06 times the Single-LUN Peak Read Throughout (MB/s) of legacy FC storage.

Performance Metric	Legacy FC Storage	Fusion-io directCache	Improvement
Aggregate DSS Query Rate (per hour)	6,103.20	11,824.00	1.94x
Single-LUN Peak Read Throughput (MB/s)	391.41	805.67	2.06x

Table 1-1: Summary of performance results.

2 Test Configuration

2.1 Hardware

Server	**Platform:**	1 – IBM System x3650 M3 (7945); 2 processor sockets
	CPUs:	2 – Intel Xeon processor X5670; 2.93GHz; 24 cores total (including hyper-threads)
	RAM:	32GB total (16GB per NUMA node)
	HBA:	1 – QLogic QLE2460, single-channel 4Gbps
Fibre Channel Storage	**Device:**	1 – IBM System Storage DS3254; 1.66GB data cache, firmware: 07.83.22.00, 6 Gbps SAS interface
	Shelves:	1 – 24 SAS Disks; 73GB, 15,000 RPM
	Arrays:	2 – 20-disk RAID10 (678.7GB); 4-disk RAID0 (271.5GB)
Fusion-io	**Device:**	1 Fusion-io ioDrive, 640GB, PCIe SSD in one x4 PCIe slot; firmware: v7.0.0, rev 107322

Table 2.1-1: Hardware configuration.

2.2 Software

Operating System	Red Hat Enterprise Linux 6.2 Server, x86_64 with kernel 2.6.32-220.4.2.el6.x86_64
Fusion-io Kernel Modules	IoMemory Virtual Storage Layer: v3.1.5, build 126
	directCache: v1.1.4, build 964
Oracle Database	11*g* Release 2 (11.2.0.1), Enterprise Edition, x68_64
DSS Workload Generator	Quest Benchmark Factory for Databases: v6.6.1, build 437

Table 2.2-1: Software configuration.

For more information on Fusion-io directCache software, see reference 2 below.

2.3 Platform and Hardware Setup

The operating system and database were installed following the detailed instructions in the Red Hat Reference Architecture for Oracle Database 11*g* R2 (see reference 1 below). In particular, EXT4 file systems were created on each LUN so that Oracle could access data using both asynchronous and direct I/O methods. The adaptive tuning system (tuned) was configured to use the “Enterprise Storage” profile.

The DSS database platform, as described above, was the baseline configuration during testing. (See Figure 2.3-1.)

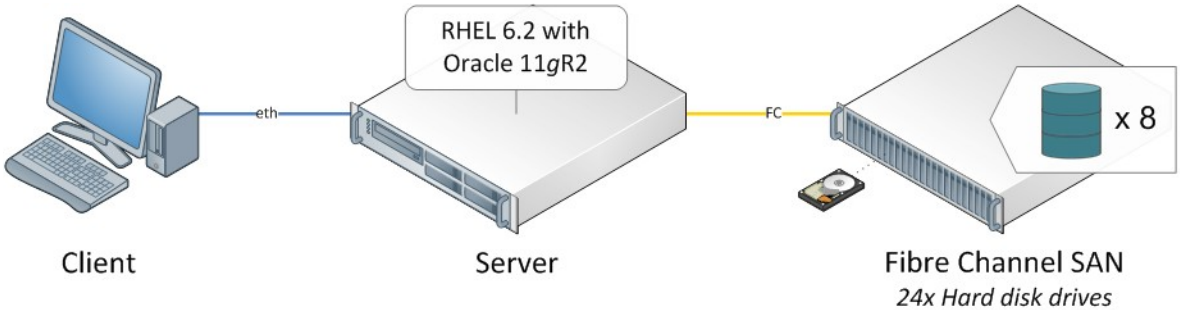

Figure 2.3-1: Baseline configuration: Database LUNs on Fibre Channel storage.

The tests with the transparent SSD cache were conducted using a Fusion-io ioDrive (640GB) plugged into the available x4 PCIe slot and Fusion-io directCache software (see below). Four of the five Fibre Channel LUNs used by the database were assigned to the front-end cache, but the Oracle log-file LUN was not cached, as it was used exclusively for sequential writes. The Oracle data, underlying file systems, and physical storage required no reconfiguration: each file system was simply “plugged into” the cache. (See Figure 2.3-2.)

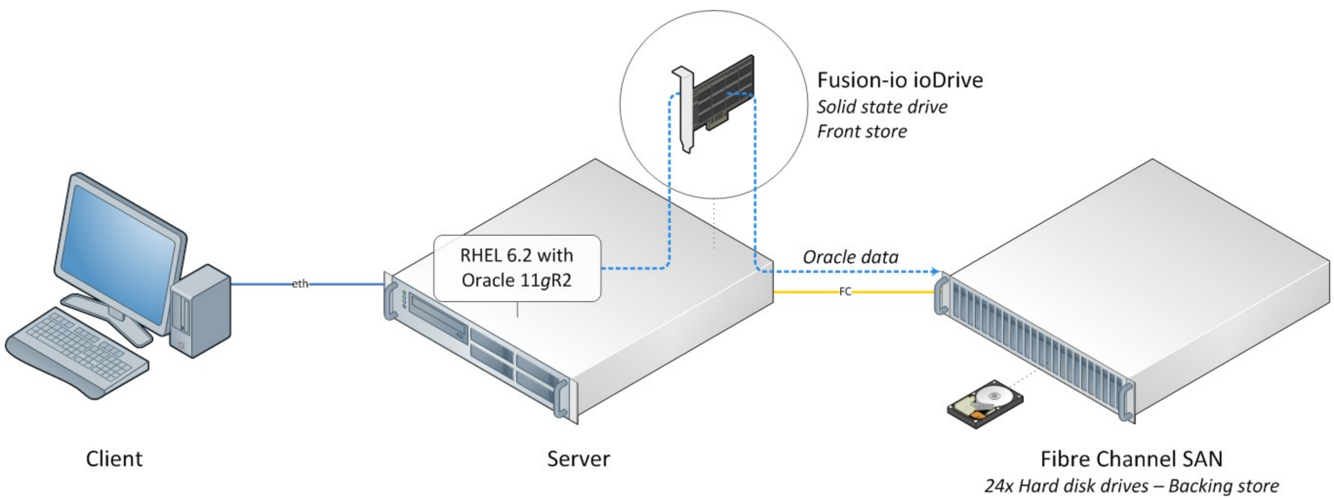

Figure 2.3-2: Fusion-io directCache: ioDrive used as transparent front-end cache for Database LUNs on Fibre Channel storage.

The Fusion-io directCache software is available from the Fusion-io support site as binary or source RPMs. The software was installed and configured as follows (see also, reference 3):

1. Log in as root.
2. Install the Fusion-io ioMemory VSL RPM.
3. Update the ioDrives firmware to the current version.
4. Install the Fusion-io directCache RPM.
5. Load the Fusion-io kernel modules: `iomemmory-vsl` and `fio-directcache` with `modprobe`. The bare ioDrive is available as the device `/dev/fct0`.

```
modprobe iomemmory-vsl
modprobe fio-directcache
```

6. Stop the database and unmount the database LUNs if necessary.
7. Add each LUN to directCache with the command `dc-bind`. For example, to bind the LUN `/dev/mapper/data01` to the ioDrive prepared above, use the command

   ```
   dc-bind /dev/mapper/data01 /dev/fct0.
   ```

   The command will list the new block device that should be used to access the LUN through the front-end cache. For example, `/dev/fdc0a`, `/dev/fdc0b`, etc.
8. Mount the cached LUN for this example at `/u02/data01` with

   ```
   mount /dev/fdc01 /u02/data01
   ```

3 Testing Methodology

The Decision Support System database platform was used for this round of testing. To produce the baseline results, the workload applied stress to the legacy storage until I/O operations were becoming the bottleneck. The 100GB Oracle Database 11*g* Release 2, Enterprise Edition database was deployed on a server with 32GB RAM and connected to back-end storage via one 4Gbps Fibre Channel link. The back-end storage consisted of two RAID arrays: the first array was configured as RAID0 and used for temporary space and indices. The second array, configured as RAID10, was used for data and logs. Five LUNs were created on these RAID arrays: the DSS data was split between two LUNs, and the temporary tablespace, logs, and indices each were assigned their own LUN.

The DSS workload compromised 10 queries from the Benchmark Factory DSS scenario without the refresh functions scaled to one 100GB database. Two levels of tests were performed: one 10-query stream and five simultaneous 10-query streams. In each case, the order of the queries was randomized. Quest Benchmark Factory for Databases generated the simulated user load.

To gauge the efficiency of Fusion-io directCache, the tests examined the DSS performance of the ioDrives for both cold and warm starts. For the former, the cache was cleared before each DSS test, which forced directCache to load most of its data from the physical Fibre Channel LUNs. In the latter case, the test started with a cold cache that was warmed up during three back-to-back DSS runs. The warm-start performance results were taken from the fourth run.

This procedure tested the ability of directCache to retain and present data from its cache as well as refresh data from the legacy Fibre Channel storage.

4 DSS Performance Results

Table 4-1 summarizes the one- and five-stream DSS performance of the baseline system (legacy Fibre Channel storage arrays) as well as Fusion-io directCache for both cold and warm starts. The Fusion-io ioDrives with directCache software provided a 41.2 percent to 50.6 percent improvement over runs using only legacy Fibre Channel storage.

Fusion-io Cache State	One DSS Stream		Five DSS Streams	
	Query Rate	Benefit over FC	Query Rate	Benefit over FC
No Cache (FC)	6,103.2	N/A	1,392.9	N/A
Cold Start	10,385.2	41.2%	2,817.7	50.6%
Warm Start	11,824.0	48.4%	2,561.8	45.6%

Table 4-1: DSS performance in queries/hour. Higher rates are better.

Figure 4-1 and Table 4-2 summarize the performance of individual queries for one DSS stream. Six queries for the cold start and seven queries for the warm start improved by over 50 percent over legacy Fibre Channel storage. Two queries for the one-stream case showed little or no improvement.

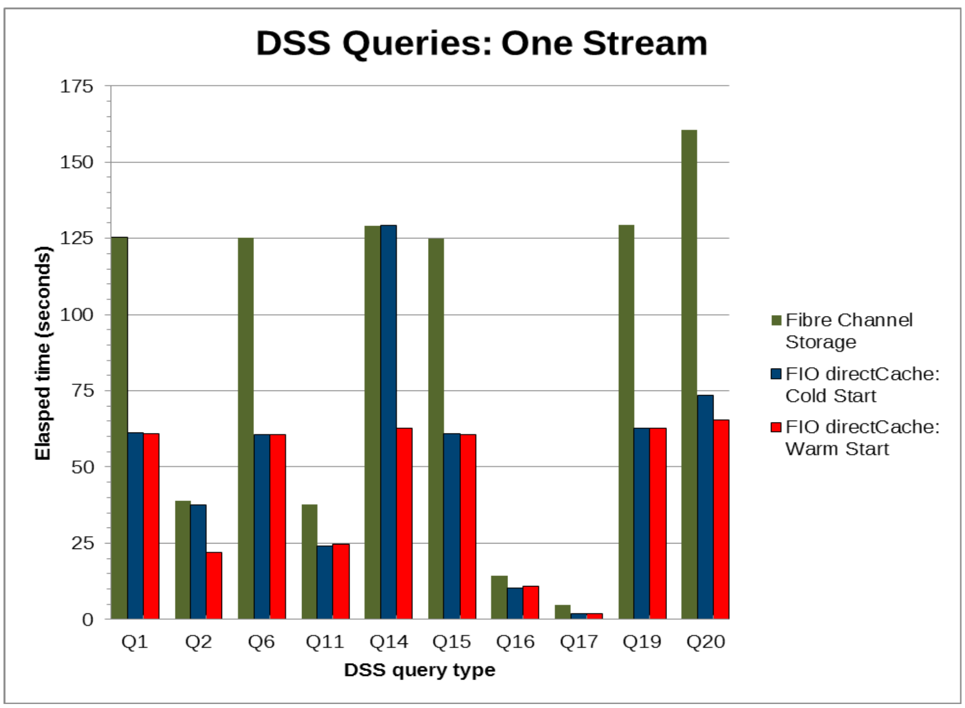

Figure 4-1: Individual query times in seconds for one DSS stream. Lower times are better.

Query	No Cache (FC)	FIO directCache: Cold Start		FIO directCache: Warm Start	
	Duration	Duration	Relative to FC	Duration	Relative to FC
Q1	125.31	61.31	51.1%	61.05	51.3%
Q2	38.80	37.42	3.6%	21.88	43.6%
Q6	125.28	60.67	51.6%	60.63	51.6%
Q11	37.78	23.97	36.6%	24.57	35.0%
Q14	129.04	129.19	-0.1%	62.81	51.3%
Q15	124.87	60.79	51.3%	60.73	51.4%
Q16	14.32	10.20	28.8%	10.99	23.3%
Q17	4.62	2.03	56.1%	2.00	56.7%
Q19	129.41	62.80	51.5%	62.64	51.6%
Q20	160.61	73.53	54.2%	65.50	59.2%

Table 4-2: Performance of DSS queries for one stream in seconds. Lower times are better.

Figure 4-2 and Table 4-3 summarize the performance of individual queries for five DSS streams. Six queries for the cold start and five queries for the warm start improved by over 50 percent over legacy Fibre Channel storage. Performance of two queries decreased during cold-start runs and three during warm-start runs.

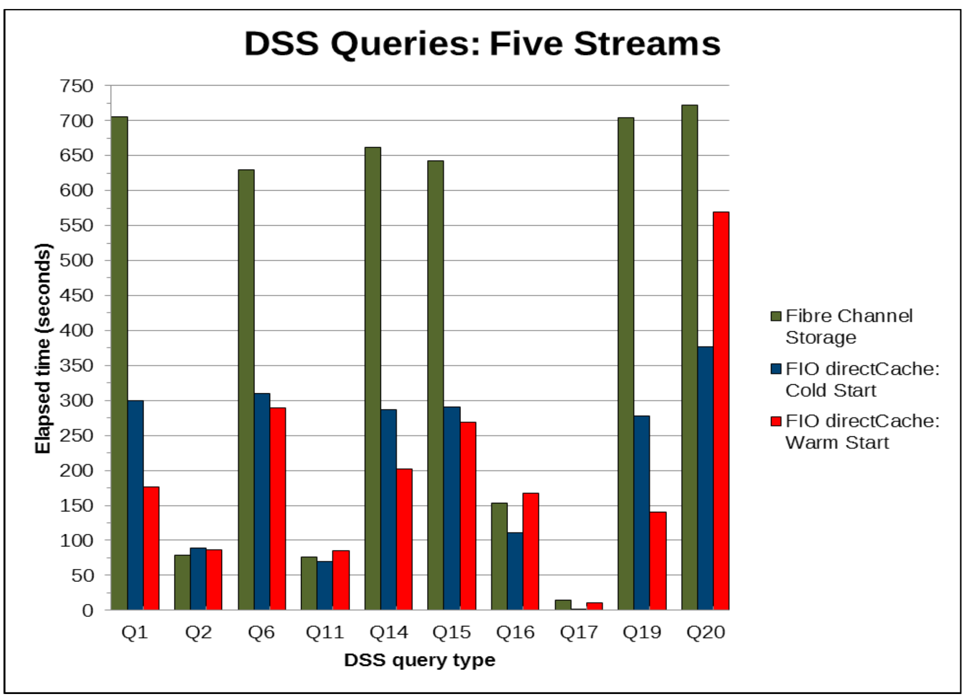

Figure 4-2: Individual query times in seconds for five DSS streams. Lower times are better.

Query	No Cache (FC)	FIO directCache: Cold Start		FIO directCache: Warm Start	
	Duration	Duration	Relative to FC	Duration	Duration
Q1	705.10	299.56	57.5%	175.88	75.1%
Q2	78.25	89.36	-14.2%	86.27	-10.2%
Q6	630.11	310.36	50.7%	289.89	54.0%
Q11	76.24	69.55	8.8%	84.88	-11.3%
Q14	661.13	287.19	56.6%	201.44	69.5%
Q15	642.47	290.98	54.7%	269.38	58.1%
Q16	153.31	111.35	27.4%	166.99	-8.9%
Q17	15.16	2.06	86.4%	11.08	26.9%
Q19	704.54	277.99	60.5%	140.67	80.0%
Q20	721.56	376.37	47.8%	569.51	21.1%

Table 4-3: Performance of DSS queries for five streams in seconds. Lower times are better.

During experiments, the disk activity was measured with the Linux utility *iostat* (note: directCache presents each LUN as a disk to Linux). Used as a micro-benchmark in these tests, *iostat* demonstrates the ability for each storage configuration to deliver data to OracleDatabase 11*g* R2.

Table 4-4 presents an excerpt of the *iostat* data at 15-second intervals. It focuses on the read performance: I/O transaction rate (IOPs), total read rates, and disk utilization.

During the first two minutes of the run, the read throughput for Fibre Channel and directCache storage were similar, about 390 MB/s. Then, while the throughput rate for Fibre Channel storage remained the same, the rate for directCache increased by a factor of two.

Fibre Channel Storage: Oracle Data LUN		
Reads IOPs	Reads MB/s	Usage (%)
787.73	391.25	100.0
785.87	391.40	100.0
788.07	391.41	100.0
787.93	391.39	100.0
786.00	391.39	100.0
305.33	150.97	38.7
800.67	391.40	100.0
665.40	325.09	83.1
2,665.13	84.31	64.8

Fusion-io directCache: Oracle Data LUN		
Reads IOPs	Reads MB/s	Usage (%)
3,518.53	391.17	100.0
3,518.93	391.19	100.0
3,517.20	391.03	100.0
3,518.53	391.12	100.0
3,518.60	391.16	100.0
1,797.93	199.96	100.0
3,522.00	390.81	100.0
1,736.53	192.63	0.0
11,126.67	739.86	100.0

Fibre Channel Storage: Oracle Data LUN		
Reads IOPs	**Reads MB/s**	**Usage (%)**
829.80	326.50	100.0
786.00	391.41	100.0
787.87	391.37	100.0
787.40	391.40	100.0
786.47	391.38	100.0
787.93	391.39	100.0

Fusion-io directCache: Oracle Data LUN		
Reads IOPs	**Reads MB/s**	**Usage (%)**
7,238.47	804.69	100.0
7,243.53	805.31	100.0
7,244.47	805.31	100.0
5,590.33	620.78	100.0
7,237.40	805.11	100.0
7,247.20	805.67	100.0

Table 4-4: Comparison of individual disk performance during a portion of a single-stream DSS run. The iostat data is taken at 15-second intervals.

5 Conclusion

Red Hat Enterprise Linux 6 is a great foundation for database deployments and with numerous record-breaking industry-standard benchmarks results on DSS workloads (see reference 4 below), it demonstrates unparalleled stability, scalability, and performance when stressed to the limit.

The test results outlined in this brief clearly demonstrate that Red Hat Enterprise Linux 6 enables Fusion-io PCIe SSDs with directCache to deliver as much as 1.94 times the DSS performance of Oracle Database 11*g* R2 of legacy Fibre Channel storage, while achieving double the peak read throughput.

The conversion from legacy storage to directCache is straightforward; there is no need to modify existing Oracle Database files, Linux file systems, or storage, making the transition a relatively inexpensive process that minimizes down time and quickly frees up the resources to tackle other projects.

Moreover, in previous testing Red Hat examined the database performance of Fusion-io PCIe SSDs when used as database disk storage. In that scenario, the Red Hat Enterprise Linux 6 also delivered significant performance improvements (see reference 5 below).

6 References and Further Information

1. For details on how to deploy Oracle Database 11*g* R2 on Red Hat Enterprise Linux 6 for several types of back-end storage, including Fusion-io ioDrives, see "Oracle Database 11*g* Release 2 on Red Hat Enterprise Linux 6: Deployment Recommendations," March 2012, http://www.redhat.com/resourcelibrary/reference-architectures/deploying-oracle-11gr2-on-rhel-6

2. For more information on Fusion-io directCache technology as a front-end cache for legacy storage, see http://www.fusionio.com/products/directcache/.

3. For details on how to install the Fusion-io directCache software on Red Hat Enterprise Linux, see "directCache User Guide – Linux for Version 1.1," June 2012.

4. For additional information on Red Hat Enterprise Linux 6 record-breaking performance on DSS workloads, see "Leadership TPC-H Benchmark™ Performance & Price/Performance using Red Hat Enterprise Linux 6", version 5.0, June 2012, http://www.redhat.com/resourcelibrary/reference-architectures/TPC-H-on-RHEL-6-V3.

5. For a recent study on the benefits of using Fusion-io ioDrives with Red Hat Enterprise Linux for OLTP databases, see "Optimizing Fusion ioMemory on Red Hat Enterprise Linux 6 for Database Performance Acceleration," August, 2011, http://www.redhat.com/resourcelibrary/reference-architectures/Optimizing-Fusion-ioMemory-on-Red-Hat-Enterprise-Linux-6-for-Database-Performance-Acceleration.

Appendix A: Revision History

Revision 1.0	09/28/12
Initial Release	