

Query Acceleration of Oracle Database 12c In-Memory using Software on Chip Technology with Fujitsu M10 SPARC Servers

June, 2016
Revision 2.0



Table of Contents

Table of Contents.....	2
1. Introduction	3
2. Oracle Database In-Memory	4
3. Fujitsu M10 Servers	5
4. Oracle Database In-Memory and Software on Chip	7
5. Verification.....	12
6. Summary	19

1. Introduction

In the fast-paced business environment of today, fast and effective management decisions are needed. To accomplish this, high-speed processing of a wide variety of database workloads is required.

Oracle Database In-Memory provided by Oracle Database 12c dramatically accelerates many types of database workloads. The Software on Chip hardware acceleration feature in SPARC64 X and SPARC64 X+ processors in Fujitsu M10 SPARC servers boosts performance further for scan processing and high-speed memory operations.

This technical document describes how to achieve a high-performance query execution when using Oracle Database 12c on the Fujitsu M10 SPARC servers with the Oracle Database In-Memory option and Software on Chip features. This document is aimed at database administrators, operators involved in improving database performance, and customers considering database deployments.

2. Oracle Database In-Memory

Oracle Database 12c provides in-memory processing and a new purely in-memory column format in addition to the existing Oracle row format as an option called Oracle Database In-Memory. Oracle Database In-Memory not only speeds up analytic workloads enabling real time analytics on the most recent transactional data, but also has the potential to improve online transaction processing (OLTP) performance.

Oracle Database In-Memory was designed to be completely risk-free and easy to deploy. Any application that runs on Oracle Database should be able to use the In-Memory option and take advantage of much faster data processing. The design does not require complex data migration or rewriting of applications. Moreover, it works on any hardware platform supported by Oracle Database. By simply enabling Oracle Database In-Memory in an environment where Oracle Database 12c operates, existing applications can be used without any modification.

Oracle Database In-Memory provides both the row format suitable for OLTP and the new column format for analytics processing in memory at the same time; accelerating processing for both workloads. As in conventional systems, OLTP uses a database buffer cache in the SGA (System Global Area), while DWH processing uses an In-Memory Column Store that is a newly provided option in the SGA.

In the In-Memory Column Store, data is converted into columns and the columnar data is loaded into memory. This speeds up analytics processing by accessing only the columns needed by the query and performing the same operations on multiple column values all at once. Because the data type in the Oracle Database column format is uniform, data can be compressed very efficiently. The columnar data compression also accelerates query performance further since the operations can be performed efficiently on the compressed data.

The In-Memory Column Store also improves OLTP performance. In conventional DWH processing, rewriting of large number of indexes when updating table data incurs significant OLTP overhead. When Oracle Database In-Memory is in effect, the need for indexes used in traditional databases for scanning has been largely eliminated because DWH processing uses the In-Memory Column Store, providing improved OLTP performance.

Furthermore, in contrast to other in-memory database systems that execute after selecting either the row or column format, Oracle Database In-Memory allows applications to access both formats. The optimizer automatically determines the best format for the query according to the request from the application, and executes in the column format for analytics processing and in row format for OLTP. In other words, applications do not need to distinguish between row and column formats, and processing appropriate for the request can be executed at high speed.

The row and column formats are synchronized in real time, and consistency is always maintained between the database buffer cache and the In-Memory Column Store data. For example, the latest information updated by scan system processing can be accessed while executing OLTP.

Oracle Database In-Memory provides six different compression levels according to the application, such as those for DML (data manipulation language) or for queries. By using specific compression methods for specific objectives, a limited memory area can be used efficiently.

For more details on Oracle Database In-Memory, see the following URL:
<https://www.oracle.com/database/database-in-memory/index.html>

3. Fujitsu M10 Servers

3.1 Overview

Fujitsu M10 SPARC servers are high-performance, highly reliable UNIX servers suitable for mission-critical workloads such as database and analytics. Fujitsu M10 servers have world record scalability and flexibility due to their advanced technologies and inherit the high reliability and technical capabilities from mainframes and previous generation SPARC Enterprise M-series.

The Fujitsu M10 server lineup is comprised of three models. Fujitsu M10-1 is an entry-level server suitable for diverse applications, with high performance, reliability, and a virtualization features in a compact 1U form-factor; the Fujitsu M10-4 is a midrange server suitable for data center integration; and the Fujitsu M10-4S is a high-end server that provides immense scalability by connecting units using the flexible Building Block architecture approach (Figure 1).

All models support Software on Chip features in the CPUs to achieve high-speed processing of large amounts of data. In addition, CPU Core Activation enables CPU resources to be added flexibly according to workload demands. Furthermore, the Fujitsu M10-4S uses Building Block architecture to connect multiple four-socket chassis to grow capacity dynamically with the Physical Partition Dynamic Reconfiguration feature.




Figure 1 Fujitsu M10 Server Lineup

3.2 Fujitsu SPARC64 X/X+ Software on Chip

Software on Chip features in the SPARC64 X/X+ processors enhance specific software features by using proprietary Fujitsu instruction sets and microarchitecture. Software on Chip features utilize supercomputer technologies developed by Fujitsu. Software on Chip technology accelerates various applications and operations natively in the Oracle Database and the Oracle Solaris operating environment on Fujitsu M10.

Software on Chip combines four features: extension registers, SIMD (Single Instruction Multiple Data), decimal floating-point calculations, and cryptographic processing. Oracle Database performance is dramatically improved with the following Software on Chip features^{*1} (Table 1).

^{*1} Software on Chip features on Table 1 are available in Oracle Database 12.1.0.2 or later.

Table 1 Software on Chip Features and Operating Environment

Software on Chip Feature	Oracle Database Performance Improvement	Operating Environment
In-Memory scan process	In-Memory scan process	Oracle Database 12c Enterprise Edition Oracle Database In-Memory Oracle Solaris 11.1 or later
In-Memory decompression process	Data decompression process (OZIP) with MEMCOMPRESS FOR CAPACITY LOW	
Decimal floating point calculations	Oracle NUMBER calculations	Oracle Database 12c or later Oracle Solaris 11.1 or later
Encryption calculations	TDE tablespace encryption Network encryption	Oracle Database 12c Enterprise Edition Oracle Advanced Security Oracle Solaris 11.1 or later
HASH	HASH Join	Oracle Database 12c
Copy/Compare	Memory processing	Oracle Solaris 11.1 or later

This document specifically describes the Software on Chip features for Oracle Database In-Memory scan and decompression processing in more detail.

3.3 High-speed, Large-capacity Memory Built In

To maximize the benefits of Oracle Database In-Memory, Fujitsu M10 servers provide large-capacity memory for loading large amounts of data into memory and high-speed memory for fast processing of this data. All Fujitsu M10 models provide up to 1 TB of memory per CPU socket and up to 64 TB in the largest Fujitsu M10-4S configuration. The large memory capacity is ideal for in-memory data processing as it allows a large amount of data to be populated into memory (Figure 2).

Traditionally, increasing memory capacity is thought to cause more memory access delays. However, because the processor and memory are directly connected in the Fujitsu M10 servers, high-speed data transfer occurs. The memory bandwidth has also been expanded by increasing the number of buses between the processors and memory. In this way, the Fujitsu M10 server provides low latency and high bandwidth memory access.




Figure 2 Fujitsu M10 Large-Capacity and High-Speed Memory Access

4. Oracle Database In-Memory and Software on Chip

This section describes how in-memory database processing can be executed faster by combining Oracle Database In-Memory with the Fujitsu M10 server. Oracle Database In-Memory populates data into memory in a format that is suitable for SIMD vector processing. By making use of high-speed vector processing capabilities of the Fujitsu M10 servers, the scan performance can be greatly improved.

4.1 Row Format and Column Format

Oracle Database In-Memory provides a dual format architecture, allowing data to be stored in both row and column formats simultaneously (Figure 3).

Row format is traditionally used in Oracle Database; with transactions or records represented as one row. Because the row format saves all attributes contained in the record as a group in memory or on storage, all attributes can be accessed quickly. For this reason, the row format is suitable for transaction processing.

On the other hand, the new column format pivots the data 90 degrees, and stores each attribute in a separate column structure. The column format is suitable for analytic scans that access only a few columns but many rows in the table.




Figure 3 Row Format and Column Format

4.2 Faster Scan Processing with Column Format

4.2.1 Optimization with a Dictionary

The column format of Oracle Database In-Memory has been optimized to fully utilize memory and CPU performance. Data in each row stored on the disk is converted to short fixed-length data using a dictionary and stored in contiguous areas in memory (Figure 4).




Figure 4 Optimization with a Dictionary

This not only allows large amounts of data to be loaded into memory, but also enables a large number of entries to be loaded efficiently into the register when scanning column data. When combined with the SIMD feature of the SPARC64 X/X+ processors, data can be processed efficiently and at high speed.

4.2.2 Faster Scan Processing with SIMD

A scan example (NATION='JAPAN') using SIMD is described below.

First, the dictionary is referenced, to determine the encoded value for Japan (5 in this case) and then the scan is executed for the encoded value, 5.

Then, the data saved in memory is loaded into the register. In scan processing that does not use SIMD, only one data entry can be loaded into each register. For this reason, only one comparison can be executed at a time. When using SIMD instructions on Fujitsu M10 servers, eight data entries can be loaded into one register because data are stored contiguously in memory as fixed-length data. By comparing this register with the register where the comparison target 5 is stored, eight comparisons can be performed at the same time. By using a dictionary in this way, multiple data entries can be stored in a single register and compared at the same time using SIMD instructions; enabling drastically accelerated scans (Figure 5).




Figure 5 Faster Scan Processing with SIMD

4.2.3 High-speed Scan with SPARC64 X/X+

The SPARC64 X and X+ processors used in the Fujitsu M10 have advanced SIMD vector processing capability, which delivers high performance for scanning. The In-Memory scan process, one of the Software on Chip features, utilizes SIMD vector calculation processing on Fujitsu M10 servers.

In traditional SPARC servers, such as the SPARC Enterprise family (the predecessor of Fujitsu M10), only one data entry could be processed at each instruction, and the processor could execute two instructions at the same time. Therefore, two data entries could be processed in each cycle. In contrast, the SIMD instruction in SPARC64 X can process up to eight data entries per instruction, and the processor executes one instruction per cycle. SIMD is further enhanced in SPARC64 X+ and is able to process up to 16 data entries per instruction when the size of the data is one byte. Also, two of these instructions can be executed at the same time. This means that up to 32 data entries can be processed in each cycle. Oracle Database 12.1.0.2 or later with patch 20342458^{*2} supports four ALUs (Arithmetic-Logic Units) execution for SIMD compare instructions. Each ALU in SPARC64 X+ processors can process up to 8 data entries, and with all four ALUs up to 32 data entries can be processed in parallel (Figure 6). Please note that Oracle Solaris 11.2 or later is required in order to apply patch 20342458.

In addition to the improved calculation capability, the read bandwidth from main memory to cache, the transfer capability from cache to register, and the number of vector registers have also been greatly enhanced in Fujitsu M10. This design ensures that the enhanced calculation capabilities are fully and natively utilized by Oracle Database In-Memory.

Please note that the performance of scan processing in each single core of SPARC64 X+ is 30 billion rows per second when the width of the dictionary is eight bits. The performance of SPARC64 X is

^{*2} Individual patch 20342458 does not need to be applied if Database Proactive Bundle Patch 12.1.0.2.7 or later is used.

14 billion rows per second in the same case.




Figure 6 SIMD Vector Processing in Traditional SPARC Servers and in Fujitsu M10

4.3 In-Memory Storage Index

In the column format, the information for each column is divided into multiple IMCUs (In-Memory Compression Units). By dividing data in this way, processing can be performed efficiently through multiple processes. Also, an In-Memory Storage Index assigned to each IMCU greatly reduces the amount of data processing.

The In-Memory Storage Index is created and maintained automatically for the In-Memory Column Store, which has the effect of reducing the amount of data processing. The In-Memory Storage Index saves a maximum value and minimum value for each column included in the IMCUs, and enables the processed data to be reduced based on filter predicates in SQL statements (Figure 7).

For queries, when the predicate of a WHERE clause is specified, the In-Memory Storage Index is scanned first to determine whether the specified value is included in the corresponding IMCU. If this shows that the specified value is not included, the scan for the corresponding IMCU can be omitted.




Figure 7 Scan Processing Using In-Memory Storage Index

4.4 In-Memory Compression

By default, data populated into the In-Memory Column Store is compressed. To efficiently use memory, Oracle Database In-Memory provides compression levels to offer flexibility when balancing performance and data size.

- **NO MEMCOMPRESS**
Data is not compressed.
- **MEMCOMPRESS FOR DML**
Optimizes data for DML operations.
- **MEMCOMPRESS FOR QUERY LOW**
Default compression level. Optimizes for query performance. Uses the previously mentioned dictionary encoding, run length encoding and bit-packing as compression techniques.
- **MEMCOMPRESS FOR QUERY HIGH**
Performs more compression than QUERY LOW and less compression than CAPACITY LOW.
- **MEMCOMPRESS FOR CAPACITY LOW**
Applies proprietary compression technology called OZIP that is specialized for Oracle Database and can perform high-speed decompression processing.
- **MEMCOMPRESS FOR CAPACITY HIGH**
Applies high-compression technology similar to gzip or DEFLATE.

4.5 Combination of In-Memory with Software on Chip SIMD

By combining the Oracle Database In-Memory Storage Index performance improvements with the benefits to scan processing from Software on Chip SIMD vector processing, acceleration of more than 100x over conventional database processing can be expected.

Of the six compression levels, OZIP used for MEMCOMPRESS FOR CAPACITY LOW can achieve high-speed decompression by utilizing the In-Memory decompression process, which is one of the Software on Chip features provided by SIMD vector calculation processing on Fujitsu M10 servers. Please note that *all* scan processes in any compression option utilize the In-Memory scan process (Table 2).

Table 2 Software on Chip SIMD Corresponding to Each Compression Level

Compression Level	Software on Chip SIMD		
	Compression	Decompression	Scan
NO MEMCOMPRESS	—	—	Yes
FOR DML	—	—	Yes
FOR QUERY LOW	—	—	Yes
FOR QUERY HIGH	—	—	Yes
FOR CAPACITY LOW	—	Yes	Yes
FOR CAPACITY HIGH	—	—	Yes

Hereafter, the In-Memory decompression and scan process are referred to as "Software on Chip SIMD".

The next section describes how the Fujitsu M10 servers are suitable for high-speed database processing based on performance verification benchmarks.

5. Verification

The Fujitsu M10 speeds up database processing by combining Oracle Database In-Memory and SPARC64 X/X+ Software on Chip features. This section provides evidence to show how the Fujitsu M10 is the optimum server for high-speed databases.

5.1 Verification Environment

For these verifications, a Fujitsu M10-4S server running Oracle Solaris 11.1 or 11.2 and Oracle Database 12c Enterprise Edition is used as the database server, with Sun Flash Accelerator F80 PCIe cards as storage (Figure 8).



Figure 8 Verification System Configuration

5.1.1 Hardware and Software

Database Server

Fujitsu M10-4S	(1unit)
Processor	3.7GHz SPARC64 X+, 4 sockets (16 cores per socket)
Memory	1TB
Operating System	Oracle Solaris 11.1.16.5.0 and 11.2.15.5.1

Storage System

Sun Flash Accelerator F80 PCIe Cards (4 cards in the Fujitsu M10-4S Database Server)

Oracle Database

Oracle Database 12c R1	(12.1.0.2)
OPatch version	12.1.0.1.10

5.1.2 Storage Settings

Four Sun Flash Accelerator F80 PCIe cards are used to create a single ZFS storage pool, a ZFS file system is created, and database files and REDO logs are placed on the ZFS file system.

5.1.3 Verification Tool

The SSB (Star Schema Benchmark) benchmark was used in these verifications. SSB is a TPC-H based benchmark developed for measuring data warehouse performance, and it is used to evaluate the performance of numerous commercial database products.

Tables

The SSB table configuration is used, as it is a typical data model for data warehouses (Figure 9). For details on SSB, see the following URL: <http://www.cs.umb.edu/~poneil/StarSchemaB.PDF>




Figure 9 SSB Table Configuration

In these verifications, in order to confirm the effects of Oracle Database In-Memory, the scale factor is set to 1,400 to make the table size approximately equal to the memory size. The table has the following number of records with the sizes shown.

TABLE_NAME	REC_COUNT	TABLESIZE (MB)
-----	-----	-----
LINEORDER	8399997494	951955
PART	2200000	224
SUPPLIER	2800000	332
CUSTOMER	42000000	5171
DATE_DIM	2556	0

Queries

When data is stored in the In-Memory Column Store, performance is greatly improved when the following operations are performed.

- Queries that scan a large number of rows and apply filter predicates that use operators
- Queries that select a small number of columns from a table or materialized view with a large number of columns, such as a query that selects five columns from a table with 100 columns
- Queries that integrate a small-scale table into a large-scale table
- Queries that aggregate data

Note that the measurement program uses the following proprietary queries, for which query Q1 of the SSB benchmark including the conditions above has been customized.

```
select count(lo_discount)
  from lineorder, date_dim
   where lo_orderdate = d_datekey
   and d_year = 1994
   and d_weeknuminyear = 7
   and d_daynuminweek = 1
   and lo_discount between 2 and 8
   and lo_tax between 2 and 6
   and lo_quantity between 10 and 45;
```

The execution plan selected automatically by the optimizer in the query is as follows.

Id	Operation	Name
0	SELECT STATEMENT	
1	SORT AGGREGATE	
* 2	HASH JOIN	
3	JOIN FILTER CREATE	:BF0000
* 4	TABLE ACCESS INMEMORY FULL	DATE_DIM
5	JOIN FILTER USE	:BF0000
* 6	TABLE ACCESS INMEMORY FULL	LINEORDER

The execution plan performs a full scan of the In-Memory Column Store for two tables, LINEORDER and DATE_DIM, and then the tables are joined via a HASH JOIN operation.

5.1.4 Database Settings

During verification, the size of the SGA and In-Memory Column Store is changed in order to store all the tables used in SSB in the In-Memory Column Store. Note that here the size of the In-Memory Column Store has been set to a high level for verification, but in an actual production environment, consideration must be given to factors such as the memory area for user processes and sorting. The In-Memory option is disabled by default. Set the INMEMORY_SIZE parameter to enable the option as shown below.

```
SQL> ALTER SYSTEM SET SGA_TARGET=900G SCOPE=SPFILE;
SQL> ALTER SYSTEM SET INMEMORY_SIZE=800G SCOPE=SPFILE;
SQL> SHUTDOWN IMMEDIATE
SQL> STARTUP
```

5.2 Verification Contents

The effects of Software on Chip are verified by performing the following two verifications:

- 1) Effects of In-Memory and Software on Chip SIMD using patch 20342458
- 2) Effects of Software on Chip SIMD in Compression Mode using patch 20342458

5.2.1 Verification for Effects of In-Memory and Software on Chip SIMD

The following four modes are compared in order to measure the effects of In-Memory and Software on Chip SIMD using patch 20342458. To see the effects of patch 20342458, the final test is

conducted in Oracle Solaris 11.2 when patch 20342458 is applied and In-Memory and Software on Chip SIMD are used (Table 3).

Table 3 In-Memory and Software on Chip SIMD Settings

Scenario	In-Memory	Software on Chip SIMD	Patch 20342458	Oracle Solaris version
Traditional	Off	Off	Not Applied	11.1
In-Memory	On	Off ^{*3}	Not Applied	11.1
In-Memory + Software on Chip SIMD	On	On	Not Applied	11.1
Patched In-Memory + Software on Chip SIMD	On	On	Applied	11.2

5.2.2 Verification for Effects of Software on Chip SIMD with FOR CAPACITY LOW

Oracle Database In-Memory provides six compression modes. Of those, the FOR CAPACITY LOW option applies compression specialized for Oracle data format called OZIP. In Fujitsu M10 servers, Software on Chip SIMD is used to accelerate decompression processing. To measure the effects of Software on Chip SIMD using patch 20342458, the following four tests are carried out. The final test is conducted in Oracle Solaris 11.2 when patch 20342458 is applied and In-Memory and Software on Chip SIMD are used (Table 4).

Table 4 OZIP and Software on Chip settings

Scenario	In-Memory	OZIP	Software on Chip SIMD	Patch 20342458	Oracle Solaris version
Traditional	On	Off	Off	Not Applied	11.1
OZIP	On	On	Off ^{*3}	Not Applied	11.1
OZIP + Software on Chip SIMD	On	On	On	Not Applied	11.1
Patched OZIP + Software on Chip SIMD	On	On	On	Applied	11.2

5.3 Verification Results

5.3.1 Effects of In-Memory and Software on Chip SIMD

With the In-Memory feature enabled, scan performance is improved through the use of dictionary encoding and In-Memory Storage Indexes. Furthermore, when Software on Chip is in effect, SIMD vector calculation makes processing faster. On top of that, when patch 20342458 is applied, the processing time is even faster.

The verification results show that the processing time when the Oracle Database In-Memory feature is used compared to when In-Memory is not used is improved by 16 to 18 times, and when both In-Memory and Software on Chip SIMD are used, processing time is improved by 49 to 190 times, moreover, when patch 20342458 is applied and In-Memory and Software on Chip SIMD are used, processing time is improved by 73 to 278 times (Figure 10 and Table 5).

^{*3} In these verifications, Software on Chip SIMD was disabled by an internal and proprietary method.




Figure 10 Effects of In-Memory and Software on Chip SIMD

Table 5 Processing Time in Each Scenario (Degree of Parallelism = 1)

Scenario	Processing Time [seconds]	Comparison with Traditional [times]
Traditional	1691	1
In-Memory	96	18
In-Memory + Software on Chip SIMD	9	190
Patched In-Memory + Software on Chip SIMD	6	278

These results demonstrate that combining the Oracle Database In-Memory feature of Oracle Database 12c and the Software on Chip feature of the Fujitsu M10 servers is effective for creating a database that achieves high-speed scan processing. In addition, using patch 20342458 is highly recommended for faster processing.

5.3.2 Effects of Software on Chip SIMD with FOR CAPACITY LOW

With the FOR CAPACITY LOW option, the Software on Chip SIMD feature makes both scan processing and decompression processing of compressed data faster.

The verification results show that processing performance is improved by 1.2 to 1.6 times by enabling OZIP. Moreover, processing is improved by 2.0 to 3.5 times when OZIP and Software on Chip SIMD are used. Furthermore, when patch 20342458 is applied and OZIP and Software on Chip SIMD are used, processing time is improved by 3.4 to 3.7 (Figure 11 and Table 6).




Figure 11 Effects of OZIP and Software on Chip with FOR CAPACITY LOW

Table 6 Processing Time in Each Scenario (Degree of Parallelism = 1)

Scenario	Processing Time [seconds]	Comparison with Traditional [times]
Traditional	85	1
OZIP	54	1.6
OZIP + Software on Chip SIMD	24	3.5
Patched OZIP + Software on Chip SIMD	23	3.7

The verification results prove that the Software on Chip SIMD feature enhances the decompression processing of data compressed by OZIP, leading to an overall performance improvement when the FOR CAPACITY LOW compression option is used. Moreover, applying patch 20342458 increases the overall performance.

6. Summary

These verification results show that query processing acceleration can be obtained by combining the in-memory column processing of Oracle Database 12c In-Memory with the high-speed vector calculation processing in the Fujitsu M10 Software on Chip SIMD feature.

The verification showed a performance improvement of up to 278 times by enabling both Oracle Database In-Memory and Software on Chip SIMD when patch 20342458 is applied.

The verification tests which examined the effects of Software on Chip SIMD with the FOR CAPACITY LOW option demonstrated that the performance was improved up to 3.7 times when patch 20342458 is applied and both Software on Chip SIMD and OZIP are enabled.

Performance improvements are realized by using the Oracle Database In-Memory option and the Software on Chip SIMD feature of the Fujitsu M10 servers and by applying patch 20342458. A database environment with a Fujitsu M10 server running Oracle Database 12c combines these two complementary technologies to greatly accelerate the query execution times when scanning large amounts of data, supporting the business needs of the customer with efficient use of their investments.

The use of Oracle Database In-Memory with the Fujitsu M10 Server Software on Chip SIMD feature does not simply provide an integrated database server platform, but creates an impressive server that transforms the business of customers from the perspective of database processing performance, efficiency, flexibility and investment cost control.

Oracle Database 12.1.0.2 or later supports and benefits from the Software on Chip features introduced in SPARC64 X+. SPARC64 X+ based Fujitsu M10 systems support the Software on Chip features found in SPARC64 X. Oracle Database 12.1.0.2 or later with patch 20342458 supports the parallel execution by multiple ALUs of the SIMD compare instruction features in SPARC64 X+ (as mentioned in section 4.2.3). The patched Oracle Database provides a further boost in database processing performance. Apply patch 20342458 to enable SPARC64 X+ based Fujitsu M10 systems to fully benefit from the SPARC64 X+ SIMD functionality.

Revision History

Revision date	Revision No.	Revision contents
October, 2015	1.0	New whitepaper created.
March, 2016	2.0	Tests using a patch to enable the Software on Chip features introduced in SPARC64 X+ added.



Fujitsu Limited

105 -7123
1-5-2 Higashi-Shimbashi,
Minato-ku, Tokyo
Shiodome City Center

All rights reserved, including intellectual property rights. Technical data subject to modifications and delivery subject to availability. Any liability that the data and illustrations are complete, actual or correct is excluded. Designations may be trademarks and/ or copyrights of the respective manufacturer, the use of which by third parties for their own purposes may infringe the rights of such owner.

UNIX is a registered trademark of The Open Group in the United States and other countries.

SPARC64, the SPARC64 logo, and all SPARC trademarks are trademarks or registered trademarks of SPARC International, Inc. (U.S.) in the United States and other countries and are used under license from SPARC International Inc.

All other product names and brand names may be trademarks or registered trademarks of the respective companies.

Trademark symbols ((R), TM) are not necessarily added to system names, product names or other trademarks described in this



CONNECT WITH US



blogs.oracle.com/oracle

facebook.com/oracle

twitter.com/oracle

oracle.com

Oracle Corporation, World Headquarters

500 Oracle Parkway
Redwood Shores, CA 94065, USA

Worldwide Inquiries

Phone: +1.650.506.7000
Fax: +1.650.506.7200

Hardware and Software, Engineered to Work Together

Copyright © 2016, Oracle and/or its affiliates. All rights reserved. This document is provided for information purposes only, and the contents hereof are subject to change without notice. This document is not warranted to be error-free, nor subject to any other warranties or conditions, whether expressed orally or implied in law, including implied warranties and conditions of merchantability or fitness for a particular purpose. We specifically disclaim any liability with respect to this document, and no contractual obligations are formed either directly or indirectly by this document. This document may not be reproduced or transmitted in any form or by any means, electronic or mechanical, for any purpose, without our prior written permission.

Oracle and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

Intel and Intel Xeon are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. AMD, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. UNIX is a registered trademark of The Open Group.0115