ESG Lab Review

Redefining Real-time Database Performance with the SPARC M7 Processor from Oracle

Date: December 2015 **Author:** Mike Leone, ESG Lab Analyst

Abstract: *This ESG Lab review documents the results of recent testing of the Oracle SPARC M7 processor with a focus on in-memory database performance for the real-time enterprise. Leveraging new advanced features like columnar compression and on-chip in-memory query acceleration, ESG Lab compared the in-memory database performance of a SPARC T7 system with a SPARC M7 processor to an x86-based system.*

The Challenges

Databases have long served as the lifeline of the business. Therefore, it is no surprise that performance has always been top of mind. Whether it be a traditional row-formatted database to handle millions of transactions a day or a columnar database for advanced analytics to help uncover deep insights about the business, the goal is to service all requests as quickly as possible. This is especially true as organizations look to gain an edge on their competition by analyzing data from their transactional (OLTP) database to make more informed business decisions. The traditional model (see Figure 1) for doing this leverages two separate sets of resources, with an ETL being required to transfer the data from the OLTP database to a data warehouse for analysis. Two obvious problems exist with this implementation. First, I/O bottlenecks can quickly arise because the databases reside on disk and second, analysis is constantly being done on stale data.

Figure 1. Comparing a Traditional OLTP and Analytics Implementation with In-memory

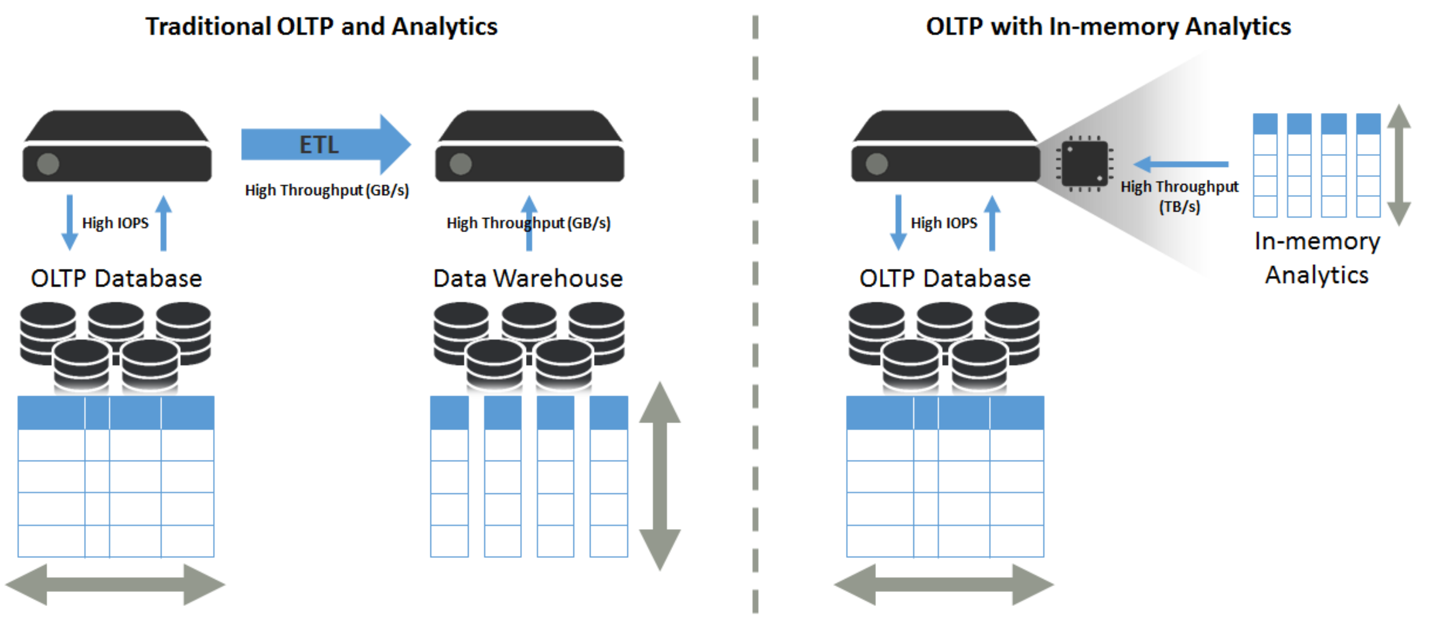

In-memory databases have helped address performance concerns by leveraging main memory to deliver high throughput and low latency when handling analytic workloads. With an in-memory analytics approach, high-throughput ETLs, data warehouse servers, and data warehouses residing on disk can be eliminated, yielding obvious CapEx savings. Real-time analytics can be completed on the transactional data without impacting OLTP performance. But this introduces a new set of problems as the constant growth of data has forced these in-memory databases to be compressed to fit into memory and that means whenever analysis is being done on the data, it must first be decompressed. This has led to bottlenecks at a memory and processor level, with both having to handle compression, decompression, and querying of the analytics database, while continuing to handle other application workloads in need of memory capacity and cycles from the same processor.

The Solution: SPARC M7 from Oracle

The SPARC M7 processor is Oracle's latest hardware innovation that provides a foundation to help organizations build finely tuned systems for the future. At the heart of this forward-looking philosophy is Oracle's Software in Silicon technology, which enables the processor to handle and execute specific software instructions in an accelerated fashion. By allowing software to run on specialized silicon in the processor, organizations gain key benefits related to performance, efficiency, and security.

Building on the existing in-memory features and performance benefits of Oracle Database 12c, the SPARC M7 further improves database performance through the use of on-chip query accelerators that offload database query processing and complete data decompression inline. And with the data accelerators being separate from the cores on the chip, more processing cycles are available to other workloads. This enables organizations to condense large datasets into memory without the performance penalty of decompression, while also freeing up processing power to run simultaneous database applications in real time.

Figure 2. Simultaneous OLTP and Data Analytics with Oracle

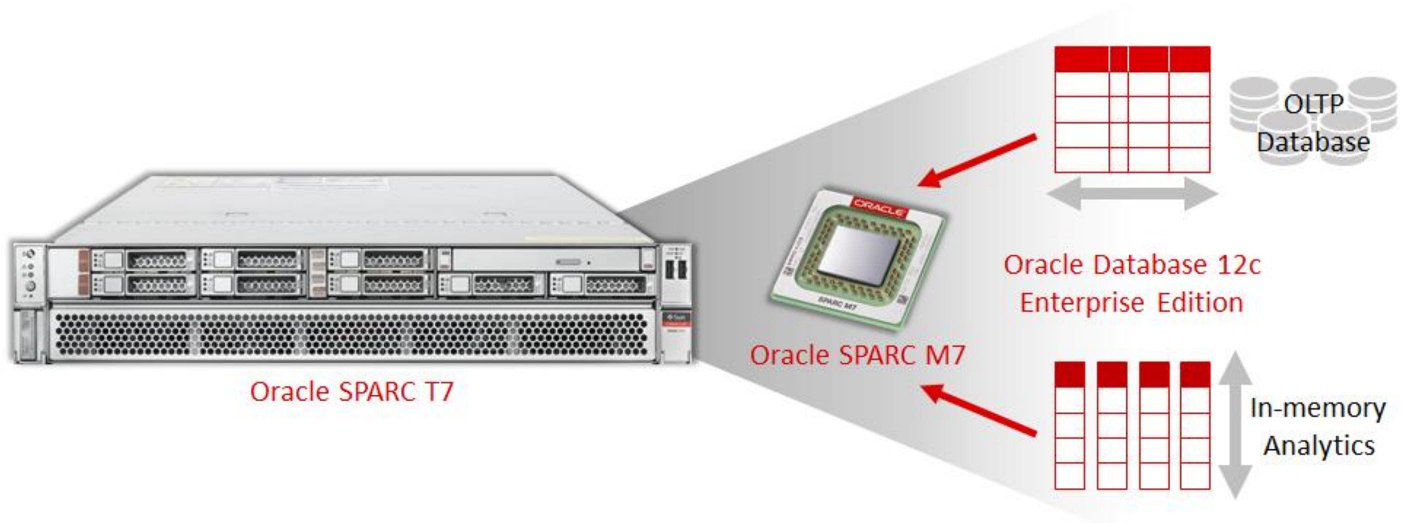

With security being top of mind for all IT organizations, Oracle uses data integrity checking to ensure applications only access their defined region of memory, preventing software errors and security intrusions. By running low-overhead, Software in Silicon, unauthorized memory access can be identified, diagnosed, and corrected in real time. For encryption, crypto instruction accelerators are integrated directly into each of the 32 SPARC M7 processor cores that enable high-speed encryption for over a dozen industry-standard ciphers.

In-memory Database Performance for Advanced Analytics

By combining Oracle's SPARC M7 processors with the in-memory functionality of Oracle Database 12c Enterprise Edition, enterprises can accelerate data analytics to make better business decisions faster. Through the use of columnar compression, a data analytics accelerator co-processor (DAX), and uniform bandwidth between all processors, organizations gain measurable efficiencies that deliver significantly higher levels of performance.

Columnar Compression for In-memory Capacity Savings

Oracle leverages columnar compression to reduce the size of large, disk-resident databases so they can be accessed in memory. The columns are first compressed using dictionary encoding, which identifies the cardinality of each column and constructs a dictionary of numbers. Cardinality can range from an attribute like gender (cardinality of 2) to an attribute like zip code (cardinality of 99,999). As an example, if all the values in a column contained the name of states in the United States, there would be a finite number of 50 total options (0-49). Instead of storing the name of each state, Oracle uses a dictionary ID. This, in turn, reduces the amount of required space to store all 50 states in just 6 bits. For reference, the state name "South Dakota" has 12 characters or 192 unicode bits. That makes the comparison 192 bits for one state versus six bits for all 50 states. The savings is obvious. For the dictionary encoding to occur, Oracle uses a command to compress the column using the oracle setting "memcompress for query low." Additional compression can be done using the query high command, which uses Oracle's proprietary version of zip (OZip). This helps to reduce the number of dictionary ID occurrences by compressing multiple occurrences of the same dictionary ID and storing it once.

ESG Lab audited results that measured the in-memory compression rates of a database that were created using an internal benchmarking tool called the Real Cardinality Database (RCDB) benchmark. A scale factor of 1,750 was used, which translates into a 1.75 TB data warehouse, which gets transformed into a star schema of 1.1 TB. The database consisted of one fact table and four dimension tables with over 10.5 billion rows, 56 columns, and most cardinalities ranged from 5 to 2,000. The setting of "memcompress for query high" was used when loading the database into memory and the size of the compressed tables were compared with their original sizes. As shown in Figure 3 and Table 1, the compression rate for the *lineorder* column was 6.2x, reducing the original 1.1TB database column to just 179 GB. The other four columns contained information about the date, part, supplier, and customer, all of which yielded a 9.8x compression rate.

Figure 3. In-memory Compression Rates of Columnar Database Tables with Oracle SPARC M7

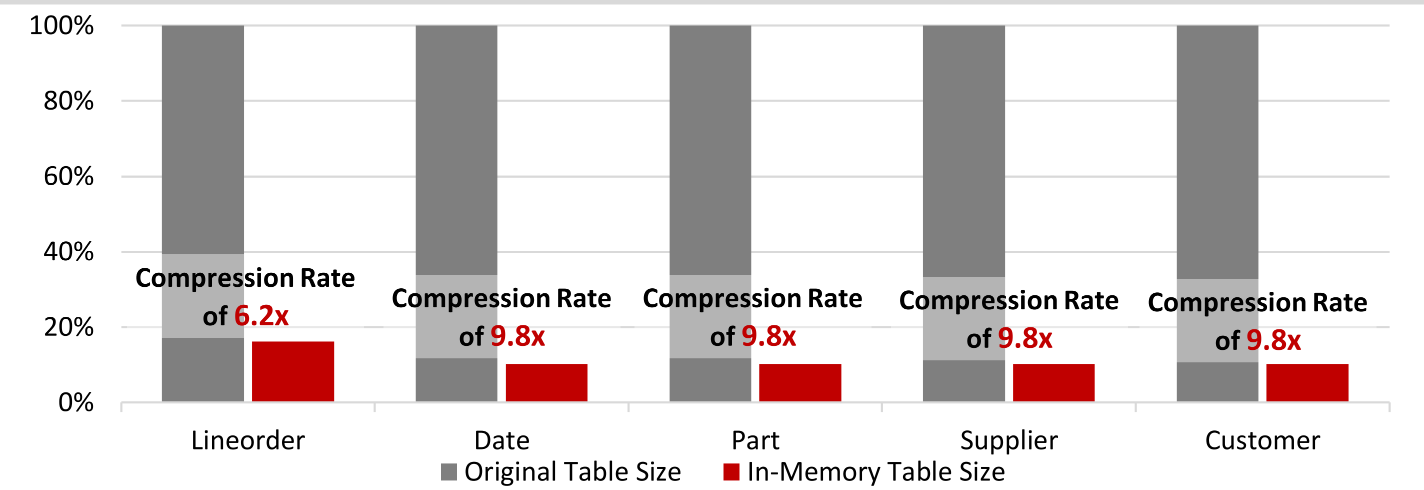

Table 1. In-memory Compression on the Oracle SPARC M7

Column Name	Original Table Size (Bytes)	In-Memory Table Size (Bytes)	Compression Rate
Lineorder	1,103,524,528,128	178,586,451,968	6.2x
Date	11,534,336	1,179,648	9.8x
Part	11,534,336	1,179,648	9.8x
Supplier	11,534,336	1,179,648	9.8x
Customer	11,534,336	1,179,648	9.8x

Data Accelerator (DAX) for Improved Query Performance

Oracle SPARC M7 has created a more efficient process to handle compressed data with the introduction of the DAX, which is a data analytics accelerator co-processor that enables database scan software to run directly in silicon. Traditionally, when data is compressed, there is a three-step process that must occur when doing analytics. First, the data gets uncompressed, then it gets written back to memory in its uncompressed format, and finally the data is scanned and the results are returned. The DAX can take that three-step process and reduce it to a single step by taking the compressed data, directly scanning it, and returning the result. By offloading the work to the DAX, CPU cycles are no longer wasted doing decompression, freeing up the processor to handle other simultaneously running application workloads. And with the DAX having direct memory access, database scans can be run at near full memory bandwidth.

ESG Lab analyzed the query performance of the 179GB in-memory compressed database using RCDB. Testing was done on two separate systems: a SPARC T7 server with a single SPARC M7 chip and an x86-based server with a dual-chip E5-2699 v3. Additional details about the two configurations can be found in the appendix, including information regarding memory, operating system, storage, and processor details. Baseline tests were run on both systems to uncover the optimal number of users before creating artificial bottlenecks. The workload consisted of 2,304 unique database queries issued in succession. An example of a query is shown in Figure 4. This particular query returns the number of distinct customers that have issued a one-line item order. From this query, various predicates can be added and changed to construct unique, subsequent queries that scan different areas of the database. For example, the same question can be asked for different weeks, months, years, varying order prices, and priority levels.

Figure 4. Database Query Example

```
select count(distinct(lo_custkey)) from
(select lo_custkey
 from lineorder, date_dim
 where lo_orderdate = d_datekey
  and d_weeknuminyear = 25
  and d_year = 2015
  and lo_orderpriority <> '2-HIGH'
  and lo_ordtotalprice between 7000 and 15000
  group by lo_orderkey, lo_custkey
  having count(lo_linenumber) = 1 );
```

Three key performance metrics were measured during the tests: elapsed time to complete all queries, queries per minute, and database memory throughput. The results are summarized in Figure 5 and Table 2.

Figure 5. In-memory Analytics Query Performance Analysis

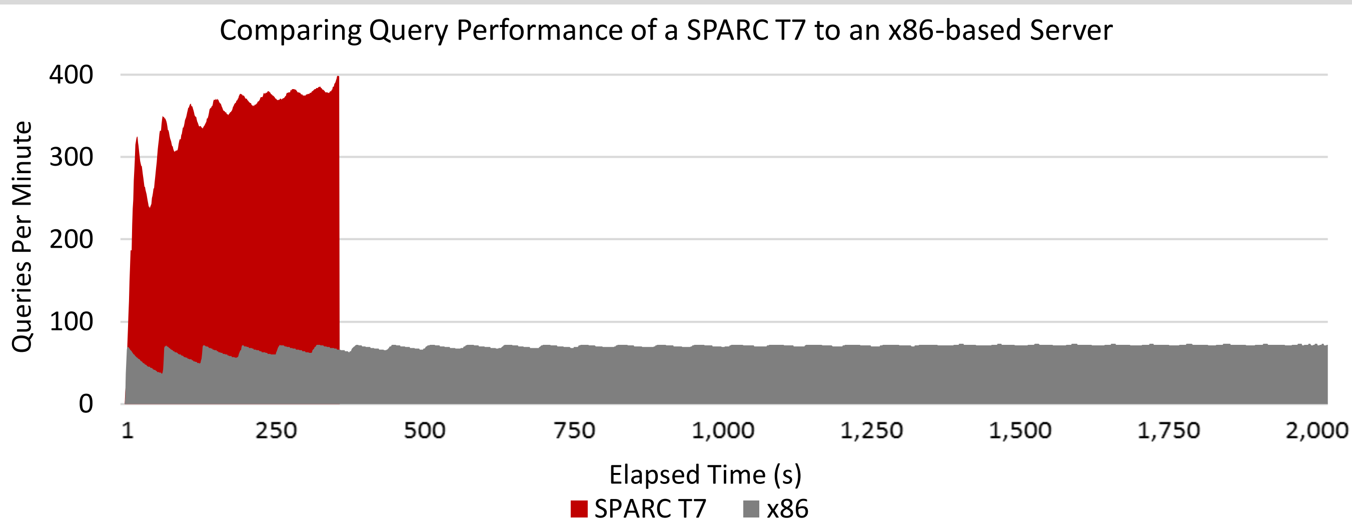

Table 2. In-memory Analytics Query Performance Analysis

System	Elapsed Time (s)	Queries Per Minute	Database Memory Throughput
SPARC T7-1 (32 core SPARC M7)	381	363	143 GB/s
X86 server (2 x 18 core E5 v3)	2,059	67	20 GB/s

What the Numbers Mean

- The same set of analytical queries were run against the same database on a SPARC T7 server with a SPARC M7 processor and an x86-based E5 v3 server. It should be noted that the x86 system that was tested contained two chips, which is one more than the SPARC T7 system.
- The SPARC system completed the entire set of queries in just 381 seconds, which was 5.4x faster than the x86-based system (2,055 seconds).
- Another metric of measurement, queries per minute, yielded the same 5.4x speed up with the SPARC system when compared with the x86 server. On a per chip level, this performance advantage improves even further to 10.8x for the SPARC system.
- With memory bandwidth and achievable throughput being a key to in-memory performance, the SPARC system scanned over 48 billion rows per second, achieving 143 GB/sec of throughput, more than 7x faster than the 20 GB/sec achieved by the x86 system.

Why This Matters

Operating more efficiently is an ongoing mission for IT in any organization. In a perfect world, IT would spend significantly less, achieve higher levels of performance, and use less resources to do it while continuing to gain a competitive advantage. In-memory database analytics has enabled organizations to operate in real time, but this has led to processor and memory bandwidth bottlenecks. The easy response is to just buy more, but that is costly and not a sustainable business model, especially as datasets continue to grow.

ESG Lab validated that the Oracle M7 processor delivers high levels of efficiency and performance by leveraging software in silicon to handle the demands of large, in-memory analytic database workloads. Unique compression algorithms take row-based OLTP databases and converts them to dense, columnar analytic databases that fit in memory. When doing analytics, the SPARC M7 leverages an on-chip data accelerator that handles decompression, scanning, and result return in a single step. Compared with an x86-based E5 v3 system that contained more chips and more cores, the Oracle SPARC T7 with an M7 processor yielded significant performance improvements. When running an identical workload on both systems, the M7 achieved 5.4x faster query execution times at a system level, 10.8x faster queries per minute on a per chip basis, and over 7x more throughput.

Simultaneous OLTP and Analytics Performance for the Real-time Enterprise

As in-memory analytics help transform organizations to be more data-driven, so grows the desire to be more real-time. Analyzing stale data should be a thing of past. With the new SPARC M7 processor, Oracle looks to push the limits of real-time by leveraging the same server to handle the analytics in memory, while also handling the processing of the OLTP database. Now transaction processing and analytics can be run simultaneously without sacrificing performance.

Using Oracle Database 12c, a dual-formatted database is maintained to enable organizations to run both OLTP and analytics at the same time. The databases are simultaneously active and consistent from a transactional standpoint. As a transaction is processed in the row-formatted OLTP database, the columnar analytics database is also updated with the same information. Columnar compression ensures the analytics database can fit into memory, and the DAX not only efficiently handles analytic requests by offloading in-memory queries, but also frees up cores to handle the transaction processing of the OLTP database.

ESG Lab audited performance results of testing that simulated a real-time enterprise. Using Oracle's real-time enterprise benchmark, both OLTP transactions and analytic queries were run simultaneously against two schemas. The first was a real-time online orders system (OLTP database). The OLTP workload simulated an order inventory system with a large number of users that performed reads and writes that stressed lock management, connectivity, and database access. The benchmark leverages a 100GB dataset, with 15 million customers, 600 million orders, and up to 580 concurrent users. The workload consisted of multiple transaction types. Examples of these transactions include updating customer information, updating prices, processing new orders, and tracking supplier inventory.

The second schema was the analytics database, which was the related historical order's schema that was configured as a real cardinality database (RCDB) star schema. A 1.05TB data warehouse (scale factor 1,050) was transformed into a star schema of 1 TB, which was reduced to 110 GBs once loaded into memory. As discussed earlier, the database consisted of one fact table and four dimension tables with over 10.5 billion rows, 56 columns, and cardinalities that ranged from 5 to 2,000. The RCDB benchmark was used to simulate up to 2,304 queries issued in succession with varying predicates.

Tests were run on two systems with a goal of comparing the collective performance of the simulated real-time enterprise: a SPARC T7-1 server with a single SPARC M7 chip and an x86-based server with a dual-chip E5-2699 v3. Additional details about the two configurations can be found in the appendix, including information regarding memory, operating system, storage, and processor details. The tests yielded two performance reports. The OLTP workload report returned the measured number of transactions/sec and average transaction response time, while the analytics workload report returned queries/min and average query elapsed times. The results are shown in Figure 6 and Table 3. Two reports are generated: one for the OLTP-Perf workload and one for the RCDB DSS workload.

Figure 6. Real-time Enterprise Performance Analysis

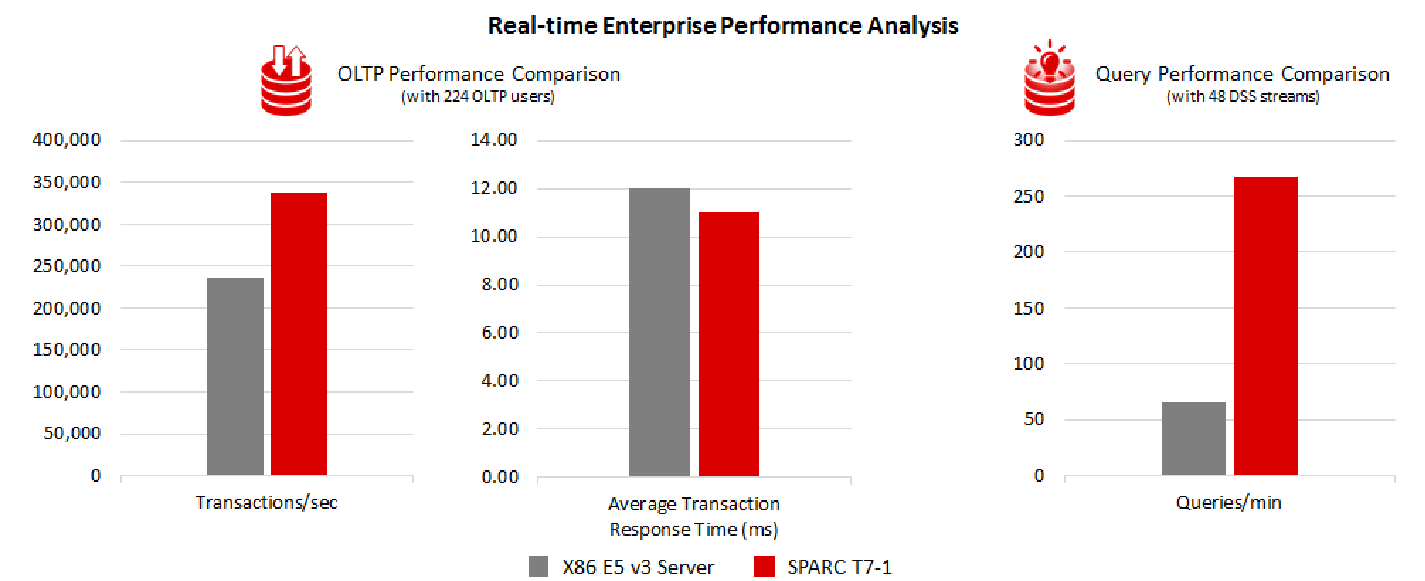

Table 3. Real-time Enterprise Performance Analysis

System	OLTP Transactions		Analytic Queries
	Transactions/sec	Average Response Time (ms)	Queries/min
SPARC T7-1 (32 core SPARC M7)	338,000	11	267
X86 server (2 x 18 core E5 v3)	236,000	12	65

What the Numbers Mean

- The performance of a SPARC T7-1 server with a SPARC M7 processor and an x86 E5 v3 server were compared while running a simultaneous OLTP and in-memory analytics workload against tables in the same database instance. The same set of transactions and queries were executed on each system. It should be noted that the x86 system that was tested contained two chips—one more than the SPARC T7 system.
- For the OLTP workload, the SPARC T7 server achieved 338,000 transactions/sec, which was 43% more transactions/sec than the x86-based server.
- The SPARC T7 server completed each transaction about 10% faster than the x86-based server with average transaction response times of 11ms with SPARC compared with 12ms with the E5 v3 processor.
- When comparing the two systems for analytic querying, the SPARC T7 system completed 267 queries/min, while the x86 system achieved just 65 queries/min—a 4.1x advantage with Oracle.

Why This Matters

With organizations looking for any way to gain a competitive advantage, the need for real-time analytics has never been greater. By leveraging in-memory computing, insights can be achieved faster, but this comes at a cost, especially in infrastructures that have shared resource environments. This has led to organizations de-prioritizing data analytics as to not impact the mission-critical transactional databases that serve as the lifeline of the business and use the same processing power and memory as the in-memory analytics platform.

ESG Lab confirmed that Oracle's SPARC M7 processor can handle the demands of a real-time enterprise by leveraging software in silicon to support both OLTP workloads and DSS workloads running on a single Oracle SPARC T7 system. With the on-chip data analytics accelerator handling the in-memory analytic workloads, more processing power can be dedicated to handling the simultaneously running OLTP workload. When compared with an x86 E5 v3 server, the Oracle SPARC T7 yielded 2.9x more transaction/sec at nearly 10% faster and 8.2x more queries/min on a per chip basis.

The Bigger Truth

IT organizations are at an impasse because of the volume, variety, and velocity of constantly growing datasets. Do they continue down a traditional path of cost and complexity to try and reign in the uncontrollable growth of data, or do they look for a new way to embrace the modern, always-connected world? The ability to both handle and harness the power of data has become a necessity for all organizations and new technology innovations are enabling it to happen at the speed of the business, without having to make tradeoffs that prioritize one workload over another.

Oracle continues to be a leader in the database market and with Oracle Database 12c, organizations have leveraged the in-memory computing capabilities to meet the real-time demands of the business. Data analytics can be done on massive datasets and yield better decisions faster than ever before. With the introduction of the SPARC M7 chip, software is run directly in optimized silicon, enabling higher levels of efficiency and significant performance improvements for in-memory Oracle databases. By converting row-store OLTP databases to real cardinality databases using Oracle compression algorithms, massive analytic datasets can be stored in-memory for faster analytics. Then, when analysis is to be done on the columnar database in-memory, Oracle's new data analytics accelerator (DAX) handles the decompression, processing, and returning results in a single step. This not only improves query performance, but also frees up cores on the chip to handle other workloads, like a simultaneously-running OLTP database workload.

ESG Lab audited performance tests that were run on both an Oracle SPARC T7 system with a SPARC M7 processor and an x86-based E5 v3 system with a goal of highlighting the efficiency and performance advantages Oracle provides in real-time, in-memory data analytic environments. By accelerating in-memory query execution through inline data decompression, the SPARC M7 processor reaches a new level of performance that drastically improves data analytic workloads. In fact, when compared with the x86 system, ESG Lab witnessed queries/min performance improvements of as much as 10.8x on a per-chip basis with the SPARC M7.

The latest Oracle technology also helps organizations consolidate database workloads by being able to simultaneously handle transactional workloads and data analytic workloads in real-time without impacting performance. ESG Lab witnessed a single database instance that housed the same set of data in both row-based and columnar format use a single SPARC M7 processor to drive an OLTP and DSS workload at the same time. This real-time enterprise simulation yielded impressive performance results, especially when compared with the x86-based system. On a per-chip basis, Oracle handled transactions nearly 10% faster, achieved 2.9x more transactions/sec, and completed 8.2x more queries/min.

The world runs in real time and therefore, your organization should, too. With a combination of the latest Oracle SPARC hardware and proven Oracle software, organizations can now operate at the speed of the business without having to worry about the underlying infrastructure. Instead of having to prioritize one database application workload over another, organizations can focus on the processes that use those database applications. If you are looking to modernize your infrastructure with forward-thinking technology that enables higher levels of efficiency with more computational power, ESG Lab suggests checking out Oracle's SPARC M7 processor.

Appendix

Table 4. Test Bed Configurations

In-memory RCDB Benchmark Hardware Configurations	
Oracle SPARC T7-1	Processor Type: SPARC M7 (32 core) Clock Rate: 4133 MHz RAM: 512 GB Operating System: Oracle Solaris 11.3 Database Software: Oracle Database 12.1.0.2.13 Internal Storage: 6 x 600 GB 2.5" SAS-3 External Storage: 180 x 300 GB 2.5" SAS-3
2 socket x86 (Haswell)	Processor Type: E5-2699 v3 (2 x 18 core) Clock Rate: 2300 MHz RAM: 512 GB Operating System: Oracle Linux 6.5 Database Software: Oracle Database 12.1.0.2.13 Internal Storage: 8 x 600 GB 2.5" SAS-3 External Storage: 180 x 300 GB 2.5" SAS-3
Real-time Enterprise Hardware Configurations	
Oracle SPARC T7-1	Processor Type: SPARC M7 (32 core) Clock Rate: 4133 MHz RAM: 256 GB Operating System: Oracle Solaris 11.3 Database Software: Oracle Database 12.1.0.2.10 Internal Storage: 6 x 600 GB 2.5" SAS-3 External Storage: 180 x 300 GB 2.5" SAS-3
2 socket x86 (Haswell)	Processor Type: E5-2699 v3 (2 x 18 core) Clock Rate: 2300 MHz RAM: 256 GB Operating System: Oracle Linux 6.5 Database Software: Oracle Database 12.1.0.2.10 Internal Storage: 8 x 600 GB 2.5" SAS-3 External Storage: 180 x 300 GB 2.5" SAS-3