IBM Power Systems solution for MariaDB

Performance overview of MariaDB Enterprise on Linux on Power featuring the new IBM POWER8 technology

Axel Schwenke
MariaDB Corporation

Hari Reddy
IBM Systems and Technology Group ISV Enablement

Basu Vaidyanathan
IBM Systems and Technology Group Performance Analysis

October 2014

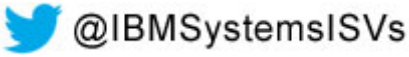

Table of contents

Abstract 1
Introduction 1
Advantages of MariaDB on Power Systems 1
MariaDB architecture 2
Sysbench OLTP benchmark 3
MariaDB performance 3
Relative performance of IBM Power S822L and IBM System x3650 M4 3
Power Systems built with the POWER8 technology 6
Tested configuration details 7
Configuration of the IBM Power S822L server with PowerKVM 7
Configuration of IBM System x3650 M4 with native Ubuntu (base-metal) 7
Configuration of IBM System x3650 M4 with Ubuntu/KVM system 8
Summary 8
Appendix A: Benchmark run details 9
Appendix B: Resources 11
About the authors 12
Trademarks and special notices 13

Abstract

This white paper describes the performance of MariaDB version 10.0.13 on IBM Power Systems servers featuring the new IBM POWER8 processor technology. The target audience is users and system integrators interested in using Linux on Power and MariaDB.

Introduction

The performance results in this paper demonstrate how the MariaDB application and IBM® Power Systems™ perform using the Sysbench benchmark suite. This white paper details the results on the IBM Power® System S822L server running MariaDB 10.0 on Ubuntu 14.0.4 with IBM PowerKVM 2.1.0 and Sysbench[1] as the benchmark test suite.

The innovative design of the IBM POWER8™ processor technology is for very demanding, data-intensive applications making it well-suited for MariaDB customers.

Built with a processor designed for such data workloads, the IBM Power Systems design combines the computing power, memory bandwidth, and I/O in ways that are easier to consume and manage; building on strong resiliency, availability and security, demonstrated by:

- Computing power with 50% more cores and smart acceleration enabled by the Coherent Accelerator Processor Interface (CAPI).
- Massive memory with over twice the bandwidth of prior generation servers to process data faster and achieve greater speed and efficiency for transactional applications such as MariaDB.
- Systems that are easy to deploy and manage with open source technologies such as OpenStack, kernel-based virtual machine (KVM), simplified virtualization management, and flexible capabilities to drive rapid adoption and dramatically simplify IT consumption.
- Better cloud economics for scale-out infrastructures, with price-performance advantages and security to confidently move data-centric applications to the cloud.

The subsequent sections provide clear examples of the advantages of MariaDB, and its performance advantages over other platforms on Linux on Power.

Advantages of MariaDB on Power Systems

The key advantages of MariaDB running on IBM Power Systems include:

- **Performance**: Over 2.2 times the performance of Intel® x86-64 based systems (Ivy Bridge)
- **Economics**: Fewer cores are required compared to other platforms and this in turn means, reduced software and support costs

[1] For more information about Sysbench refer to: https://launchpad.net/sysbench

MariaDB architecture

The MariaDB architecture consists of the following components and all ran in a single hardware partition along with the Sysbench online transaction processing (OLTP) benchmark component used as a workload driver to MariaDB.

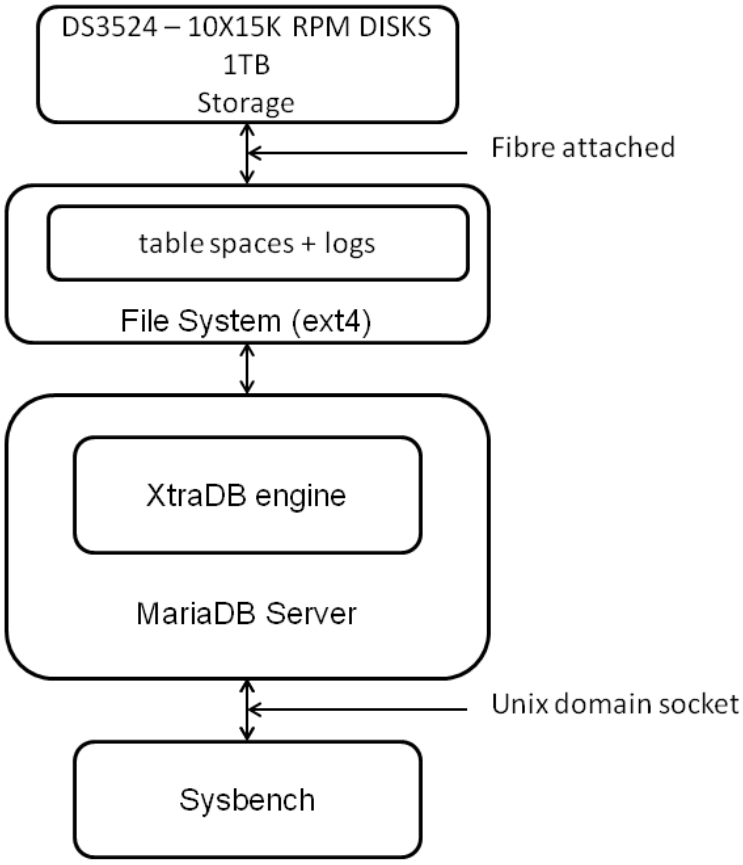

Figure 1: MariaDB and Sysbench architecture and test topology

Sysbench OLTP benchmark

Sysbench is an OLTP application benchmark ran on top of a MariaDB database running the XtraDB storage engine. The storage engine handles all aspects of managing and storing the data that is below the SQL level. The storage engine uses table spaces and log files to store data and redo logs on persistent storage. The XtraDB engine used in MariaDB is a branch of the InnoDB engine from MySQL, and it is Atomicity, Consistency, Isolation, Durability (ACID)-compliant.

Sysbench has two operating modes: a read/write mode to the database and a read-only mode. Sysbench simulates load that is typical for an OLTP application: transactions are short and use well-indexed tables. A read-only transaction in Sysbench contains 14 SELECT statements, a read/write transaction additionally contains two UPDATE, one INSERT and one DELETE statement.

MariaDB performance

The main measure of performance used in the Sysbench benchmark study is the transactions per second (TPS). Additional metrics such as average response time per transaction, 99 percentile response time, and processor utilization are also measured. All these measures can help provide a comprehensive view of the behavior of the implementation of MariaDB on a given architecture.

Relative performance of IBM Power S822L and IBM System x3650 M4

Though both IBM Power S822L and IBM System x®3650 M4 servers are equipped with two sockets each, in this study, one socket was used as this allows you to measure and compare the performance of virtual partitioning that is becoming increasingly common place.

IBM Power S822L with the guest partition running PowerKVM is over 2.2 times faster than IBM System x3650 M4 with the guest partition running Ubuntu/KVM for read-only mode of the Sysbench benchmark.

When considering the performance of these configurations on the read/write performance, which introduces the effects of a virtual hypervisor and its I/O characteristics, the Power S822L server continues to demonstrate a strong performance, which is of two times that of the System x3650 server (refer to Figure 2).

Even when comparing IBM Power S822L with Ubuntu running in a guest partition to the IBM System x3650 M4 server running native Ubuntu, the read-only performance of the Power Systems server remains two times faster.

The conclusion drawn from this testing is that the Power S822L server is a capable and high-performing system that is well suited for MariaDB workloads and the growing data-intensive demands of its customers.

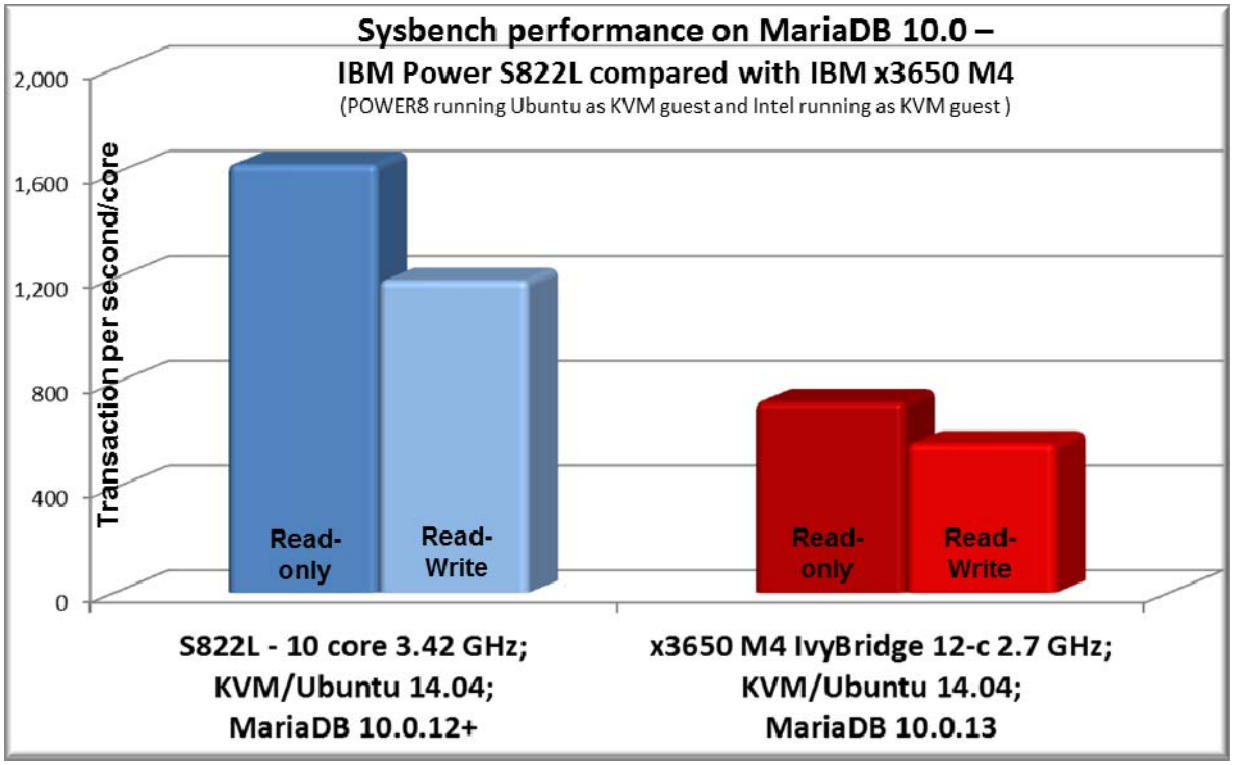

Figure 2: Relative performance of IBM Power S822L and IBM System x3650 M4

Processor utilizations of both these systems were similar and ranged from 93% to 99%. The duration of the measurement was 600 seconds for read-only transactions and 3600 seconds for read/write transactions. Those times include the database warm-up phase.

Refer to the “Tested configuration details” section and the “Summary

The performance results in this paper demonstrated how the MariaDB application and IBM Power Systems servers perform using the Sysbench benchmark suite.

The performance testing result of IBM Power S822L indicates that the Power S822L server is a capable and high-performing system that is well-suited for MariaDB workloads and the growing data-intensive demands of its customers.

The results also demonstrate that IBM Power Systems running under a virtualized environment can perform very well and use resources efficiently. These systems can deliver over two times the performance of similarly configured virtualized servers on x86 systems. This directly translates into significant reduction in the total cost of ownership for the customer.

Appendix A: Benchmark run details” section for more details.

Power Systems built with the POWER8 technology

In this section, after presenting an overview of the POWER8 technology, the configuration of the systems used for this MariaDB benchmark is provided.

POWER8 is a multicore, multichip (node), and a multisocket processor technology. The number of chips and sockets available vary with the model purchased. A representative layout of the POWER8 processor is shown in Figure 3 with double the memory bandwidth when compared to the IBM POWER7+™ processor.

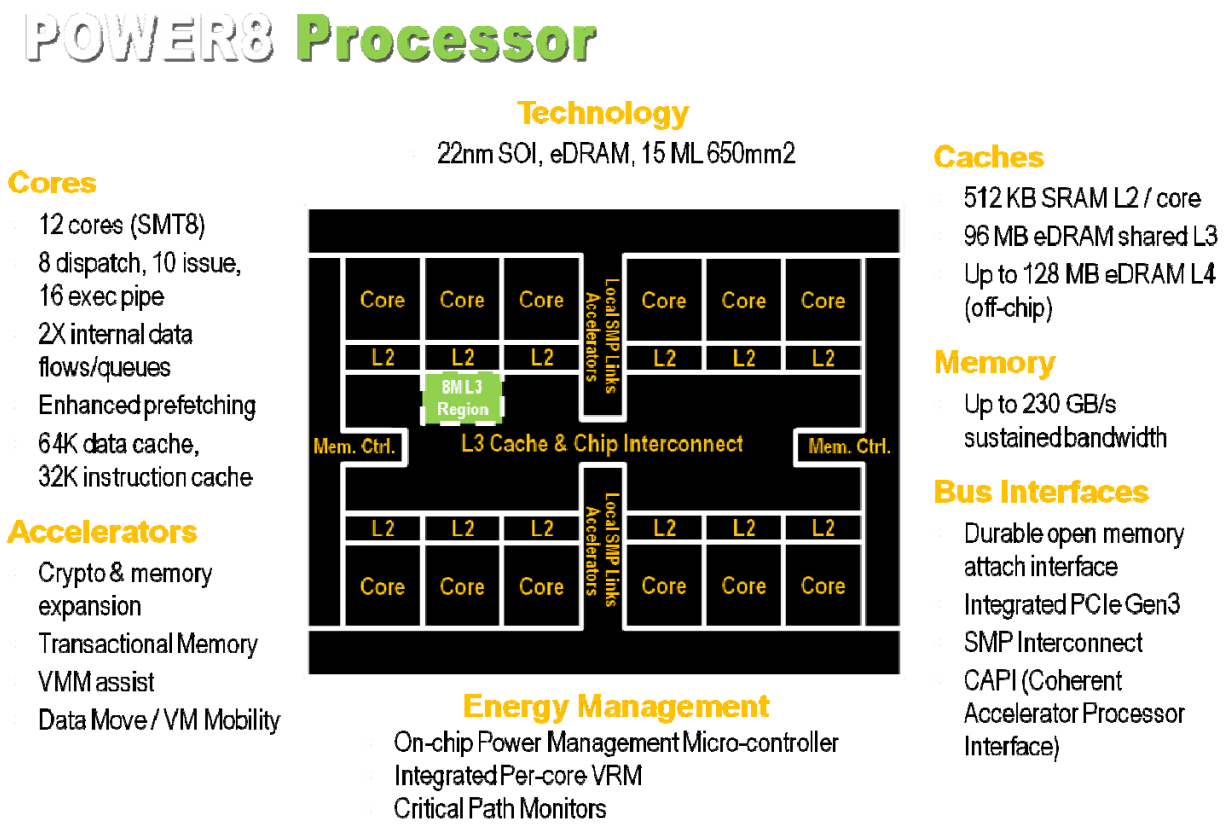

Figure 3: POWER8 processor

Tested configuration details

In this section, the hardware and software configurations used for the testing are described.

Configuration of the IBM Power S822L server with PowerKVM

Hardware

- Processor: IBM POWER8, DCM, 10 core, 4.16 GHz, 96 MB L3 Cache, SMT 8
- Memory: 256 GB, sixteen 16 GB 1866 MHz DIMMs
- Storage: IBM System Storage® DS3524 Express Dual Controller Storage System
- One logical unit number (LUN) of 10 15K RPM disks, RAID5
- Storage connectivity: QLogic Corp. ISP2532-8Gbs Fibre Channel to PCI Express (PCIe) HBA (rev 02)

Base software

- OS: Ubuntu 14.04 IBM; Linux® kernel: 3.13.0-35-generic
- Virtualization: PowerKVM 2.1.0
- File system: ext4

Application software

- Application: MariaDB 10.0.14

Configuration of IBM System x3650 M4 with native Ubuntu (base-metal)

Hardware

- Processor: Xeon E5-2697 v2, 12 core, 2.7 GHz, 30 MB L3 Cache, HT
- Memory: 384 GB, twenty-four 16 GB 1866 MHz DIMMs
- Storage: System Storage DS3524 Express Dual Controller Storage System
- One LUN of 10 15K RPM disks, RAID5
- Storage connectivity: QLogic Corp. ISP2532-8Gbs Fibre Channel to PCIe HBA (rev 02)

Base software

- OS: Ubuntu 14.04 IBM; Linux Kernel: 3.13.0-32-generic
- Virtualization: none
- File system: ext4

Application software

- Application: MariaDB 10.0.14

Configuration of IBM System x3650 M4 with Ubuntu/KVM system

Hardware

- Processor: Intel Xeon E5-2697 v2, 12 core, 2.7 GHz, 30 MB L3Cache, HT
- Memory: 384 GB, twenty-four 16 GB 1866 MHz DIMMs
- Storage: IBM System Storage DS3524 Express Dual Controller Storage System
- One LUN of 10 15K RPM disks, RAID5
- Storage connectivity: QLogic Corp. ISP2532-8Gbs Fibre Channel to PCIe HBA (rev 02)

Base software

- OS: Ubuntu 14.04 IBM; Linux Kernel: 3.13.0-32-generic
- Virtualization: Ubuntu/KVM
- File system: ext4

Application software

- Application: MariaDB 10.0.14

Summary

The performance results in this paper demonstrated how the MariaDB application and IBM Power Systems servers perform using the Sysbench benchmark suite.

The performance testing result of IBM Power S822L indicates that the Power S822L server is a capable and high-performing system that is well-suited for MariaDB workloads and the growing data-intensive demands of its customers.

The results also demonstrate that IBM Power Systems running under a virtualized environment can perform very well and use resources efficiently. These systems can deliver over two times the performance of similarly configured virtualized servers on x86 systems. This directly translates into significant reduction in the total cost of ownership for the customer.

Appendix A: Benchmark run details

The following three architectures were used in the benchmark runs:

- Ubuntu PowerKVM guest
- Intel system running native Ubuntu (bare metal)
- Ubuntu/KVM guest

The following are the descriptions of the tables containing the results of the benchmark runs.

- Table 1: Results of running the Sysbench benchmark in the read-only mode with varying thread counts
- Table 2: Results of running the Sysbench benchmark under read/write mode with varying thread count
- Table 3: Peak performance for read-only runs from Table 1 for each architecture
- Table 4: Peak performance for read-only from Table 2 for each architecture

Note: Some software thread counts were not tested in middle range as they did not represent peak values for that specific hardware configuration and, therefore, did not offer additional technical meaning, and are represented with an asterisk (*) symbol in the following tables.

Details of runs to determine peak performance for each system (in TPS – higher values are better)						
	IBM Power S822L – PowerKVM		IBM System x3650 M4 – bare metal		IBM System x3650 M4 – Ubuntu/KVM	
	per socket	per core	per socket	per core	per socket	per core
Threads						
16	9439	944	6546	545	6924	577
32	11990	1199	9450	787	8773	731
40	*	*	**9482**	**790**	8757	730
44	*	*	*	*	**8784**	**732**
80	**16367**	**1637**	*	*	*	*
88	16361	1636	*	*	*	*
128	16181	1618	9258	772	8631	719

Table 1: MariaDB 10.0 read-only mode run details

Details of runs to determine peak performance for each system (in TPS – higher values are better)						
	IBM Power S822L – PowerKVM		IBM System x3650 M4 – bare metal		IBM System x3650 M4 – Ubuntu/KVM	
	per socket	per core	per socket	per core	per socket	per core
Threads						
16	7375	738	5747	479	5370	448
32	9420	942	7342	612	6815	568
40	*	*	7383	615	6858	572
44	*	*	*	*	**6855**	**571**
64	*	*	**7398**	**617**	*	*
80	11913	1191	*	*	*	*
88	**11957**	**1196**	*	*	*	*
128	11687	1169	7388	616	6800	566

Table 2: MariaDB 10.0 read/write mode run details

Peak performance summary – Read-only mode			
	IBM Power S822L – PowerKVM	IBM System x3650 M4 – bare-metal	IBM System x3650 M4 – Ubuntu/KVM
Threads	80	40	44
Per system transaction rate (TPS)	16367	9482	8784
Per core transaction rate (TPS)	1637	790	732
Average Response time (ms)	4.88	4.11	4.78
99% percentile response (ms)	5.42	7.66	8.25
Processor utilization (in percentage)	98	98	98

Table 3: Peak performance summary for read mode

Peak performance summary – read/write mode			
	IBM Power S822L – PowerKVM	IBM System x3650 M4 – bare-metal	IBM System x3650 M4 – Ubuntu/KVM
Threads	88	64	40
Per system transaction rate (TPS)	11957	7398	6858
Per core transaction rate (TPS)	1196	617	572
Average Response time (ms)	5.38	8.65	6.42
99% percentile response (ms)	10.38	28.8	20.25
Processor utilization (in percentage)	93	96	97

Table 4: Peak performance summary for read/write mode

Appendix B: Resources

The following websites provide useful references to supplement the information contained in this paper:

- IBM Systems on PartnerWorld
  ibm.com/partnerworld/systems

- IBM Power Systems
  ibm.com/systems/in/power/?lnk=mhpr

- IBM Linux on Power – resources
  ibm.com/systems/power/software/linux/resources.html

- MariaDB Foundation
  www.MariaDB.org/

- MariaDB Corporation official website
  www.MariaDB.com/

- IBM Power Systems Hardware Documentation
  http://publib.boulder.ibm.com/infocenter/powersys/v3r1m5/index.jsp

About the authors

Hari Reddy is a technical consultant in IBM Systems and Technology Group ISV Enablement organization. You can reach Hari at hnreddy@us.ibm.com.

Axel Schwenke is a performance analyst in MariaDB Corporation. You can reach Axel at axel@mariadb.com.

Basu Vaidyanathan is a performance analyst in IBM Systems and Technology Group Performance Analysis organization. You can reach Basu at basu@us.ibm.com.

Mark Nellen is a program manager in IBM Systems and Technology Group, ISV Enablement organization. You can reach Mark at mnellen@us.ibm.com.

Maya Pandya is a technology manager in IBM Systems and Technology Group, ISV Enablement organization. You can reach Maya at maya@us.ibm.com.

Sergey Vojtovich is a software developer in the MariaDB Foundation. Sergey can be reached at svoj@mariadb.org.

Michael (Monty) Widenius is an advisor and a board member in MariaDB Corporation and is also a member of the board of directors in MariaDB. You can reach Monty at monty@mariadb.org.

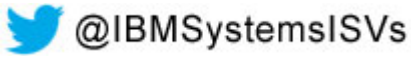

Trademarks and special notices

References in this document to IBM products or services do not imply that IBM intends to make them available in every country. Information is provided "AS IS" without warranty of any kind.

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. If these and other IBM trademarked terms are marked on their first occurrence in this information with a trademark symbol (® or ™), these symbols indicate U.S. registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at "Copyright and trademark information" at www.ibm.com/legal/copytrade.shtml.

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

Intel, Intel Inside (logos), MMX, and Pentium are trademarks of Intel Corporation in the United States, other countries, or both.

Other company, product, or service names may be trademarks or service marks of others.

Information concerning non-IBM products was obtained from a supplier of these products, published announcement material, or other publicly available sources and does not constitute an endorsement of such products by IBM. Sources for non-IBM list prices and performance numbers are taken from publicly available information, including vendor announcements and vendor worldwide homepages. IBM has not tested these products and cannot confirm the accuracy of performance, capability, or any other claims related to non-IBM products. Questions on the capability of non-IBM products should be addressed to the supplier of those products.

Some information addresses anticipated future capabilities. Such information is not intended as a definitive statement of a commitment to specific levels of performance, function or delivery schedules with respect to any future products. Such commitments are only made in IBM product announcements. The information is presented here to communicate IBM's current investment and development activities as a good faith effort to help with our customers' future planning.

Performance is based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput or performance that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput or performance improvements equivalent to the ratios stated here.

Any references in this information to non-IBM websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.