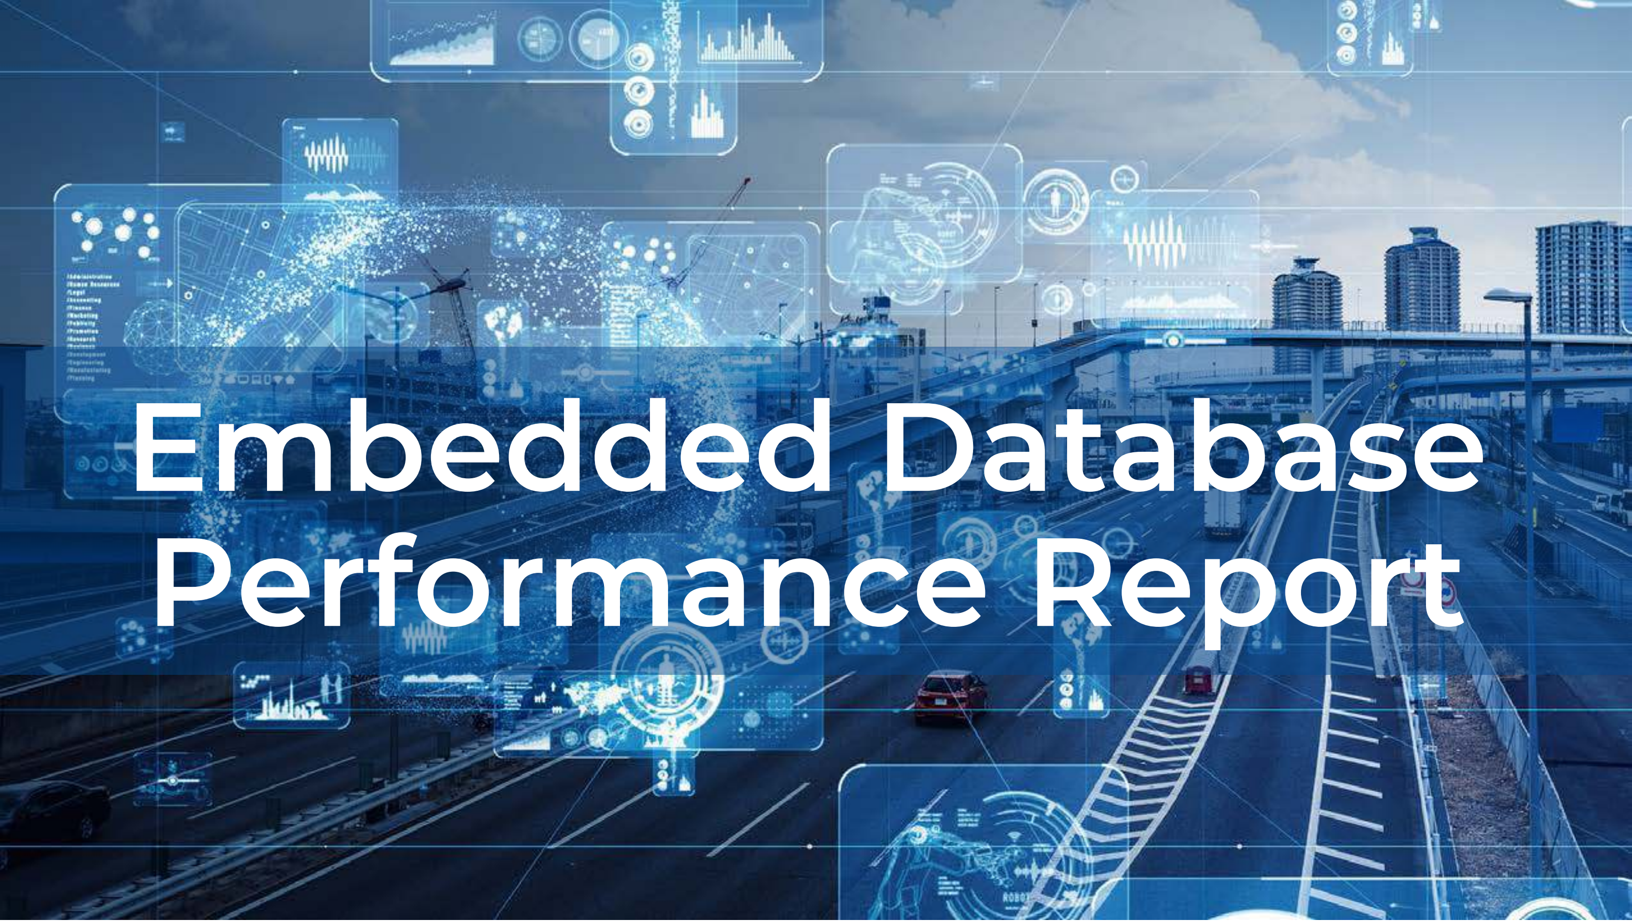

Embedded Database Performance Report

Actian Zen more than 6x faster than MySQL

MCG Global Services Benchmark Results

May 2021

ACTIAN™

Key insights

- This benchmark compared Actian Zen Enterprise Server and MySQL Enterprise, both running on the same Ubuntu Linux in 8- and 16-core VMs as AWS EC2 instances, each using its ODBC driver. The benchmark used is derived from the TPC-C industry standard benchmark.

- Actian Zen Enterprise was compared with MySQL across multiple sets of warehouses being queried concurrently in increments of 1, 2, 4, 8, 16, 32, 64, 96 and 128 users. Overall, Actian Zen was 6X faster than MySQL, specifically at:
  - 4 warehouses and up to 32 users, Zen performs 3X – 7X faster
  - 16 warehouses and up to 96 users, MySQL could only support up to 8 users before generating errors and, on average, Zen was 5X faster
  - 32 warehouses and up to 96 users, MySQL could only support up to 4 users and Zen was 6X faster

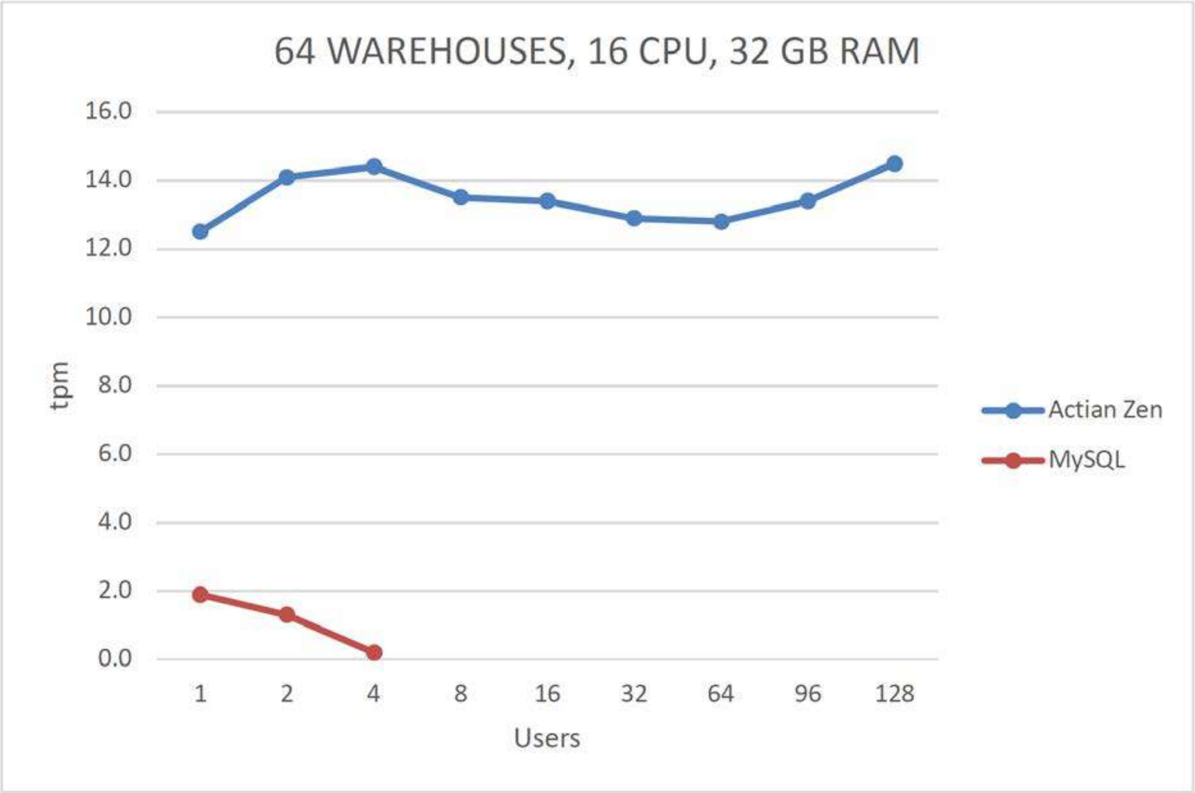

Checkout the Actian Zen performance advantage today!

Visit https://www.actian.com/zen

Embedded Database Performance Benchmark

Product Profile and Evaluation:
Actian Zen and MySQL Enterprise

By McKnight Consulting Group
May 2021

Sponsored by

Executive Overview

Embedded databases are built into software, transparent to the application's end users and require little or no ongoing maintenance. Embedded databases are growing in ubiquity with the rise of mobile applications and internet of things (IoT) giving innumerable devices robust capabilities via their own local database management system (DBMS). Developers can create sophisticated applications right on the remote device. For these uses, the embedded architecture is preferred over client-server approaches which rely on database servers accessed by client applications via interfaces. Today, to fully harness data to gain a competitive advantage, embedded databases need a high level of performance to provide real-time processing at scale.

All these web, mobile, and IoT applications have generated a new set of technology requirements. Embedded database architecture needs to be far more agile than ever before and requires an approach to real-time data management that can accommodate unprecedented levels of scale, speed, and data flexibility. Sometimes legacy databases are unable to meet these new requirements, and developers are therefore turning to embedded database technology.

To quantify embedded database performance, we conducted this benchmark study, which focuses on the performance of application-ready embedded database solutions Actian Zen and Oracle MySQL Enterprise. The intent of the benchmark's design was to represent a set of database transactions that an organization developing edge applications might encounter.

The test methodology was based on a workload derived from the well-recognized industry-standard TPC Benchmark™ C (TPC-C). Our benchmark harness was adapted from the Percona Labs TPC-C driver to fit our use case. We conducted the benchmark on Zen and MySQL installed on the same Amazon Web Services EC2 instance and utilizing each engine's ODBC driver.

Overall, the benchmark results were insightful in revealing the query execution performance of Actian Zen and MySQL revealing some of the differentiators in the two products.

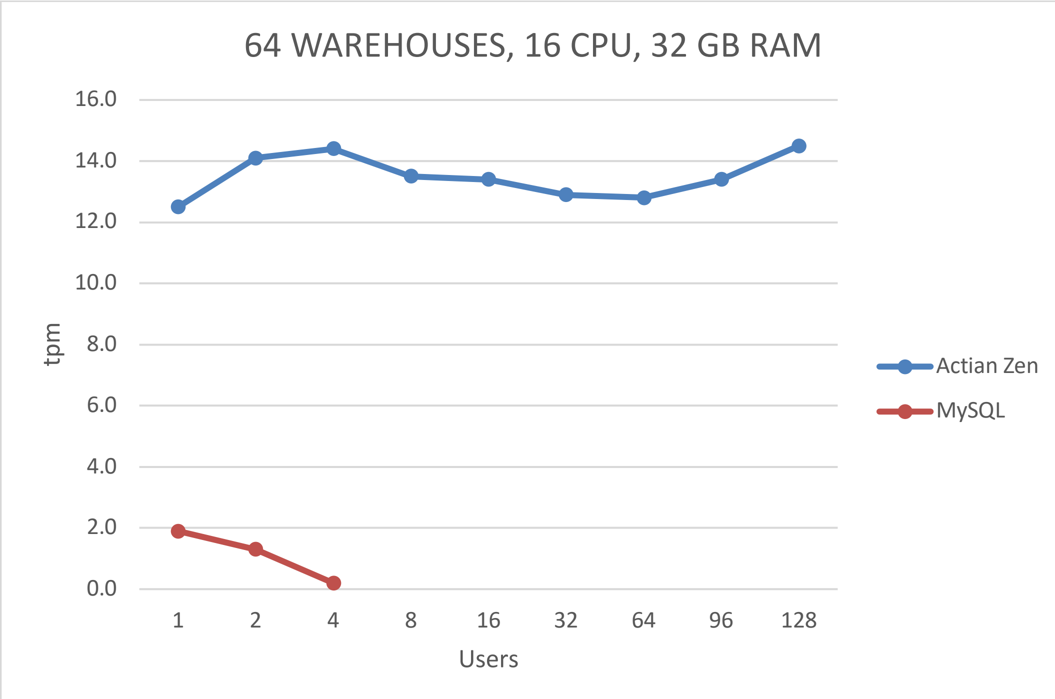

In our final test at 64 warehouses, Actian Zen transactions per minute times were around 6X faster at all user counts 1 through 4, after which we received errors with MySQL.

Actian Zen is a mature platform for embedded database applications with over 30 years of engineering and development behind it. Features that contributed to its extremely fast performance include, but are not limited to, the Btrieve API and Turbo Write Accelerator.

We leave the issue of fairness for the reader to determine. We hope this report is informative and helpful in uncovering some of the challenges and nuances of platform selection.

Embedded Database Selection

Organizations that utilize application-laden smart devices rely on embedded database platforms to process edge data at high speed and bring it in with consistency to harmonize an ecosystem of activity. Volumes for data that can be utilized at the edge is rapidly expanding—placing significant performance demands on embedded architectures. Thus, a key differentiator is the depth by which a database maintains performance to scale with simple queries representative of real-world use cases of embedded databases.

While performance is very important, it is not the only consideration. Developers choosing embedded database must consider data access, scalability, and availability.

Both MySQL and Actian Zen were designed to "set it and forget it," with little-to-no ongoing database administration. Since Oracle acquired MySQL, the engine has been re-engineered to add more enterprise capabilities. However, Actian Zen was engineered purposefully to pare down an enterprise platform to be embedded within environments like OEM and IoT. Therefore, Actian Zen has features that MySQL does not—including auto-reconnect networking and automated defragmentation.

Both platforms offer application support. Zen is a key-value store that operates natively as either SQL or NoSQL and is flexible enough to be relational, document-based. MySQL offers document-based support for JSON documents. Additionally, Zen exclusively offers the high performance Btrieve 2 API. The Btrieve 2 API supports NoSQL and native development support for Java and C/C++ based devices and SWIG for Python, Perl, and PHP—in addition to its SQL support. MySQL offers software development kits for C/C++, .NET, Node.js, Python, and PHP. Finally, we tested the ODBC libraries of both platforms. Rather than have two different applications, each independently using each engine's native libraries, we tested ODBC, so the same application code was portable to either engine without modification.

While the subject of this benchmark is embedded applications, Actian Zen Edge is part of the overall Zen family of Zen Core, Zen Enterprise, and Zen Reporting Engine. When combined, this suite of products enables not only embedded applications, but client-server (with zero ETL) and cloud deployments as well.

In a client-server configuration, Actian Zen comes with the capability to automatically synchronize in real time between Zen Core or Edge on a remote device and Zen Enterprise on a server—without ETL. This capability is critical for today's needs and uses, because the potential number of mobile devices could easily number in the thousands, and all that information may need to funnel into a core database on a server. Actian has a real-time synchronization capability of Actian Zen Edge or Core to Zen Enterprise via the Btrieve API without an intermediary, which can allow you to achieve scale with simplicity.

Platform maturity is also a consideration. MySQL was initially released in 1995. Actian Zen was initially released as Btrieve in 1982 (and later PSQL) and has been in production with many multi-national organizations with nearly 40 years of engineering and enhancement.

This report focuses on the performance of the two embedded database options via ODBC. It is important to get into the right embedded database early in the development cycle when the stakes are less critical.

Benchmark Setup

The benchmark was executed using the following setup, environment, standards, and configurations.

An aim of the benchmark is to simulate a typical real-world scenario and use case for embedded databases. In our benchmark, we derived a workload from the well-recognized industry-standard TPC Benchmark™ C (TPC-C). However, this is NOT an official TPC-C and the results are not comparable with other published TPC-C results. One reason is our choice to utilize the ODBC (Open Database Connectivity) C library to handle query execution, cursors, and handles within the benchmark driver code. In a TPC-C benchmark, you would use the database platform's native libraries to achieve this. However, we desired our code to be portable across these (or any) platforms. The trade-off of this approach is that ODBC creates some additional overhead which impacted overall performance for both platforms. While this approach is good for comparison (apples-to-apples), it is not good for determining the maximum throughput of the engine it could likely achieve with its own native libraries.

Platform Configuration

Our benchmark included two different embedded RDBMS—Actian Zen and Oracle MySQL Enterprise—installed on the same types of EC2 instances. The benchmark test driver was run directly on the same server as the database engines.

RDBMS

Engine	**Actian Zen Enterprise Server for Linux x86_64**	**MySQL Enterprise for Linux x86_64**
Version	v. 14.20.012.000	v. 8.0.23-commercial

AWS EC2 Instances

8-core Virtual Machine	**c5.2xlarge**
Processor	8x 3.6GHz Intel Xeon Scalable Processors (Cascade Lake)
RAM	16 GB
OS	Ubuntu 20.04 (Focal Fossa)

16-core Virtual Machine	**c5.4xlarge**
Processor	16x 3.6GHz Intel Xeon Scalable Processors (Cascade Lake)
RAM	32 GB
OS	Ubuntu 20.04 (Focal Fossa)

Other Software

We also installed unixODBC 2.3.6 for the management of the ODBC drivers.

Test Use Case

As mentioned, the test was adapted from the Percona Labs TPC-C driver to fit our use case. We conducted the benchmark on Zen and MySQL installed on the same Amazon Web Services EC2 instances and utilizing the unixODBC driver and C libraries. With the Percona Labs code, we kept the structure the same, which follows the TPC-C specification very closely, but substituted ODBC functions for connection and statement handles, parameter bindings, SQL statement preparation, SQL statement execution, results retrieval, and freeing allocated resources.

In our test and the TPC-C specification, the scale factor of the benchmark is the number of "warehouses". The warehouse count, as show in the following figure, determines the cardinality and data volume of the initial database once data is generated and loaded.

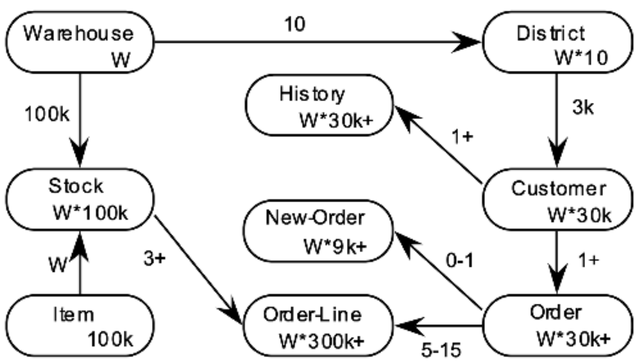

Another test parameter we adjusted and scaled during our testing is the number of concurrent users (as threads within the testing application). Concurrent users and the number of warehouses are not completely independent variables. We found that we could not use more four (4) to six (6) times the number of concurrent users than the number of warehouses. If we used too many concurrent users per warehouse, we began to see threads trying to access the same rows simultaneously and we saw deadlock errors.

The following table shows the different test parameters we used:

Warehouses	1 user	4 users	16 users	32 users	64 users	96 users	128 users
1	✓	✓	✓				
4	✓	✓	✓	✓			
16	✓	✓	✓	✓	✓	✓	
32	✓	✓	✓	✓	✓	✓	
64*	✓	✓	✓	✓	✓	✓	✓

*The 64 warehouses scale factor was only tested on the 16 CPU configuration and not the 8 CPU.

Benchmark Results

The following figures display the average time elapsed for each database transaction for both Actian Zen and MySQL.

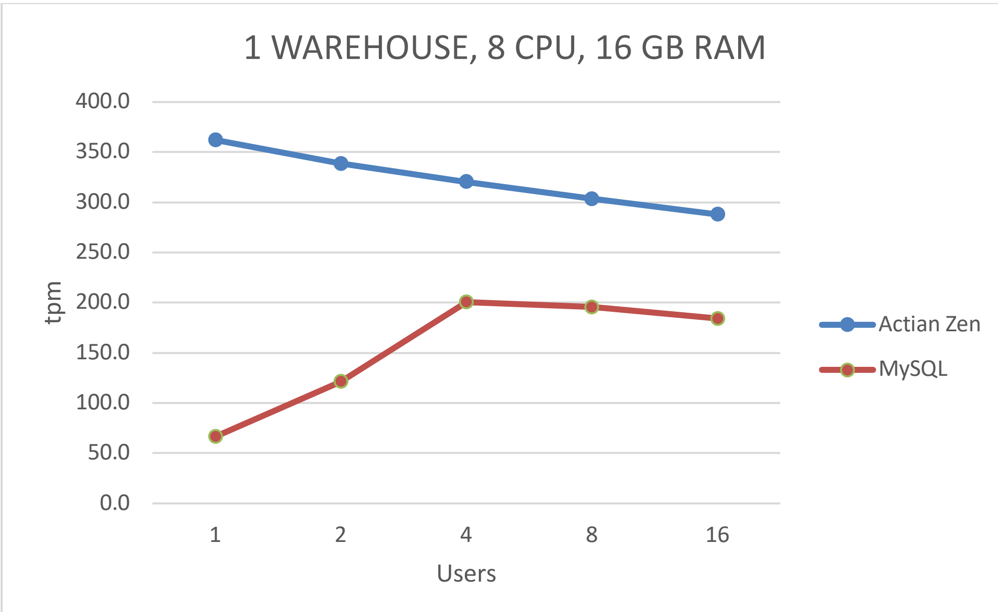

In the case of the 1 TPC-C warehouse (a proxy for workload) with 8 CPU and 16 GB RAM, Actian Zen transactions were 3X – 7X per minute times were faster than MySQL.

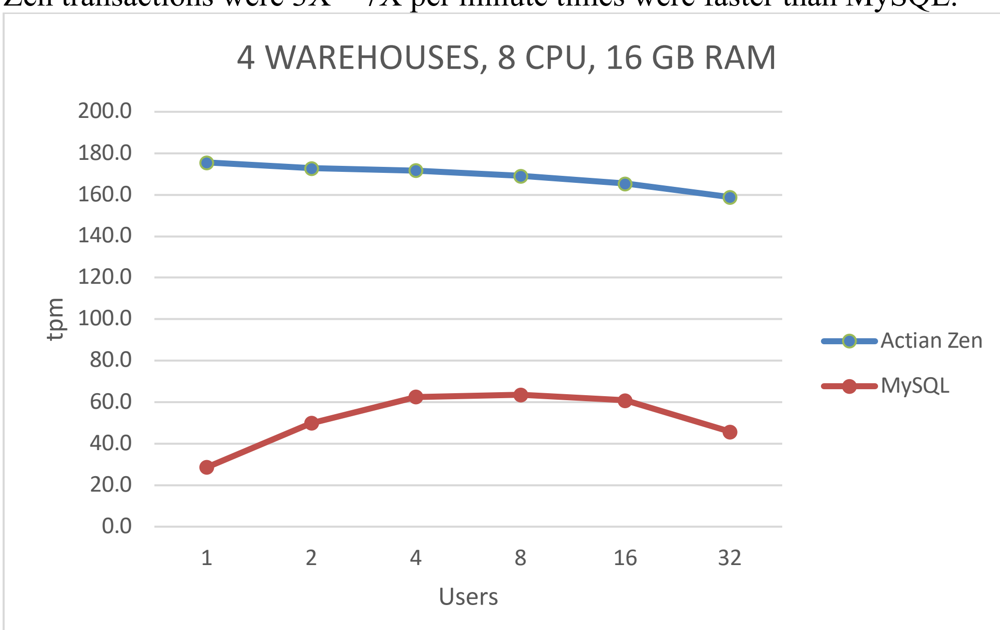

In the case of the 4 TPC-C warehouses, Actian Zen transactions per minute times were around 3X faster at all user counts than MySQL.

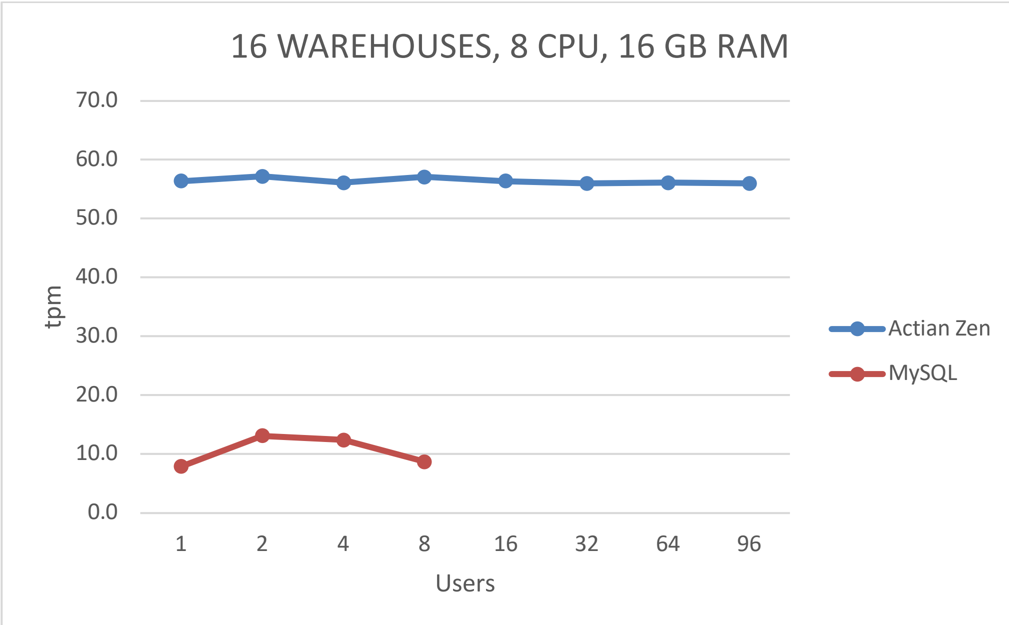

In the case of the 16 TPC-C warehouses, Actian Zen transactions per minute times were around 5X faster at all user counts 1 through 8, at which point we received errors with MySQL. We never hit the upper limits of Zen in our tests.

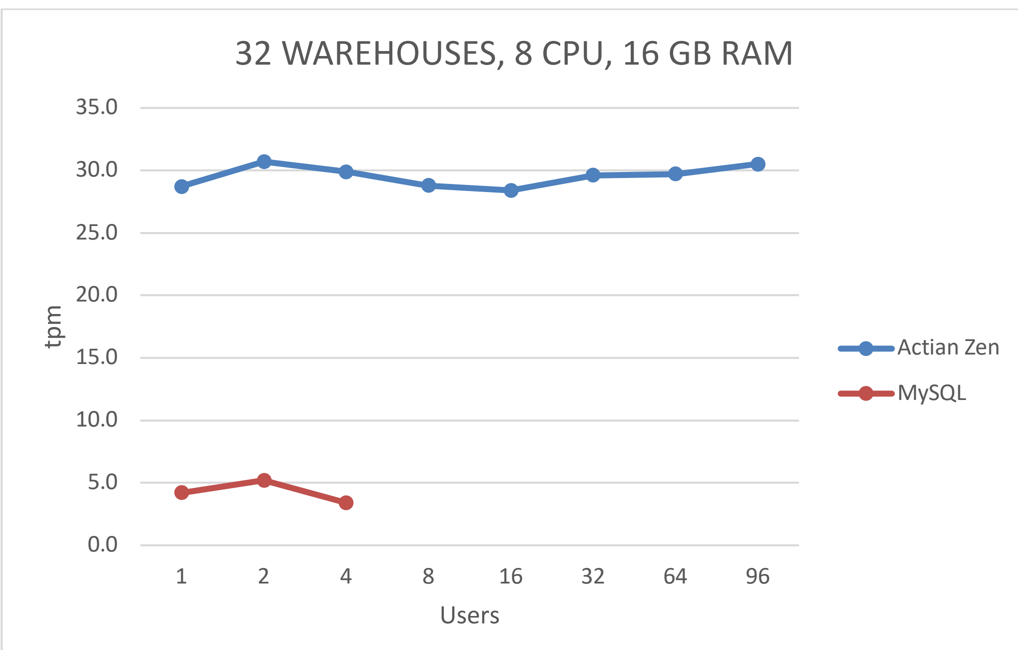

In the case of the 32 TPC-C warehouses, we received errors with MySQL at 8 users.

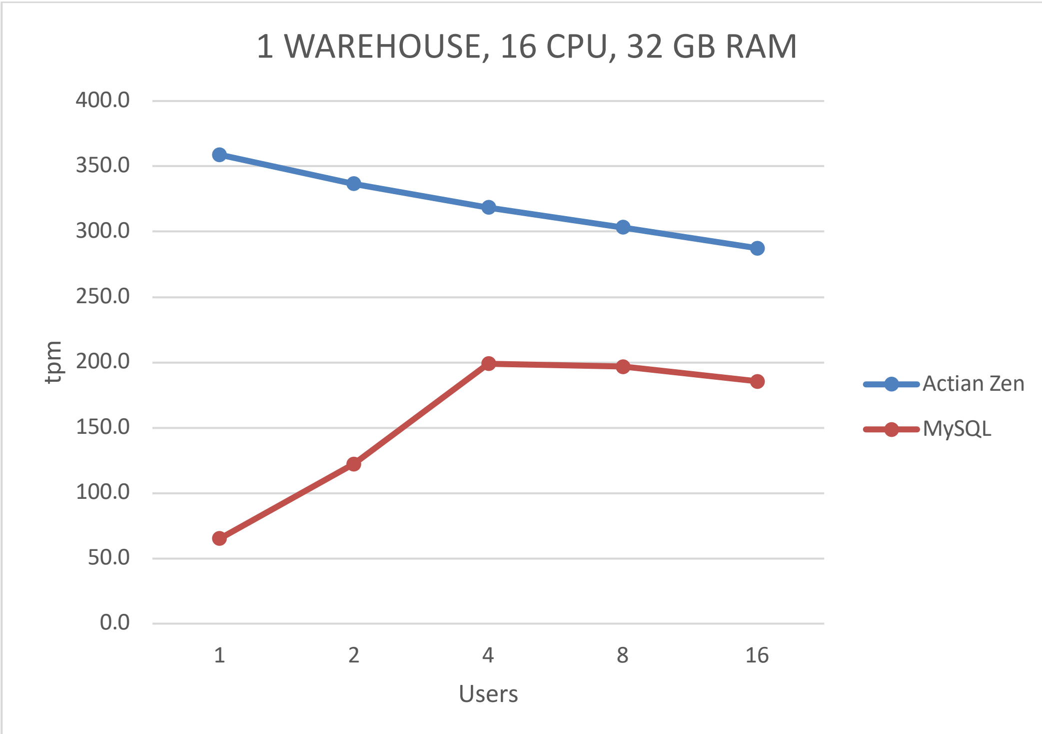

We moved up to 16 CPU and 32 GB RAM and the numbers were similar to the 8 CPU, 16 GB RAM test.

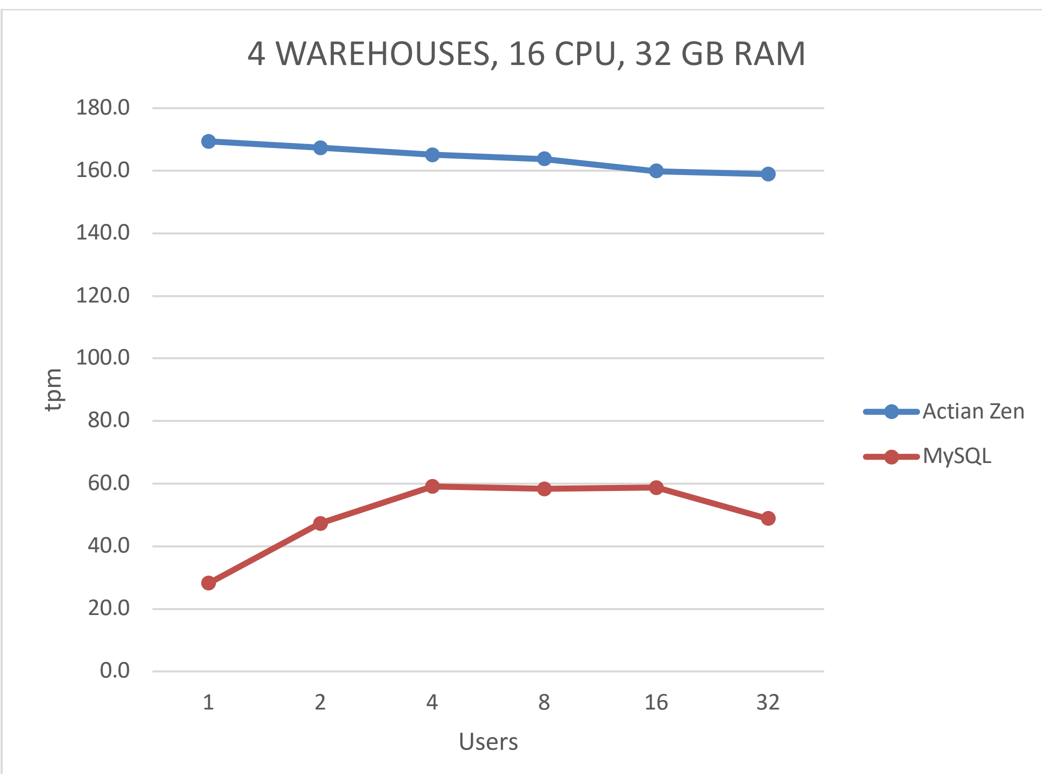

In the case of the 4 TPC-C warehouses, Actian Zen transactions per minute times were again around 3X faster at all user counts than MySQL.

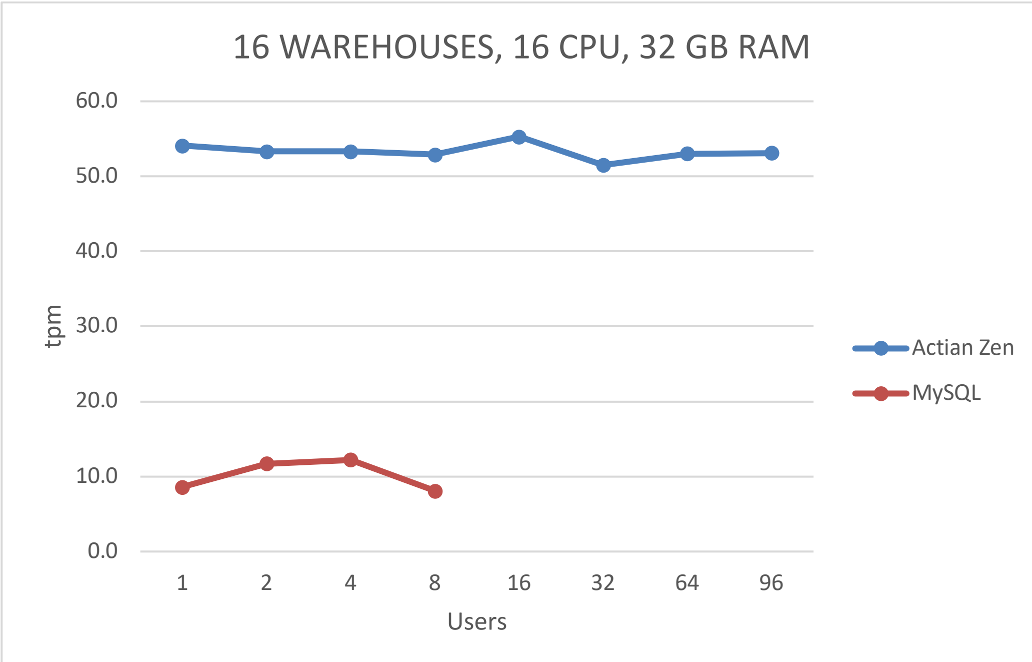

In the case of the 16 TPC-C warehouses, Actian Zen transactions per minute times were around 5X faster at all user counts 1 through 8, after which we received errors with MySQL.

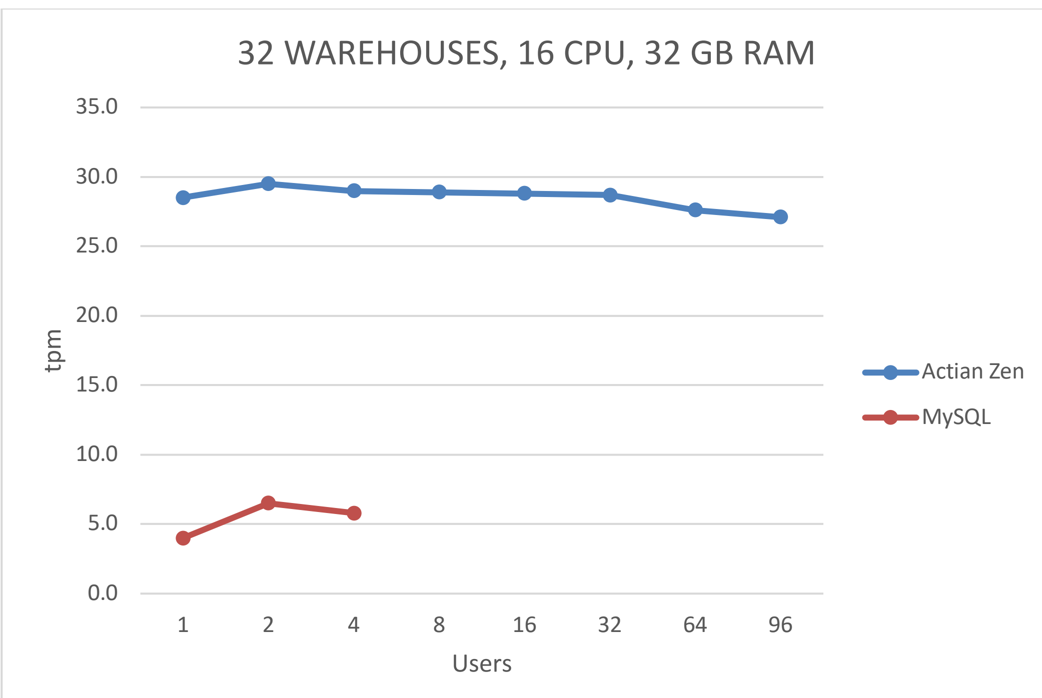

In the case of the 32 TPC-C warehouses, we received errors with MySQL at 8 users.

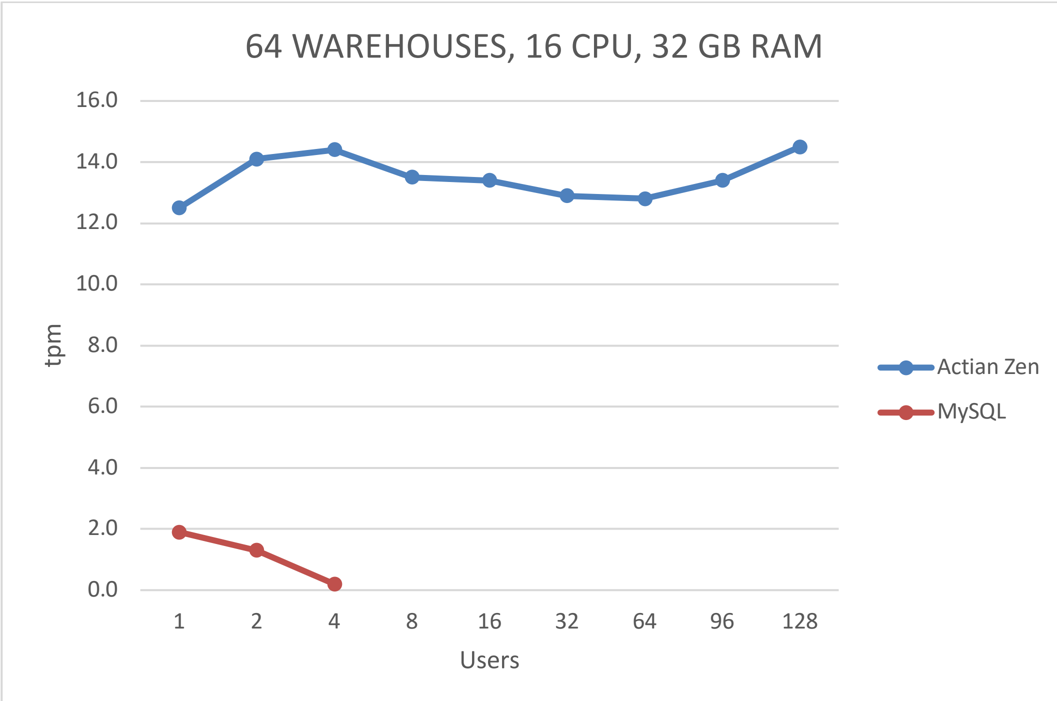

In our final test at 64 warehouses, Actian Zen transactions per minute times were around 6X faster at all user counts 1 through 4, after which we received errors with MySQL.

Conclusion

Actian Zen consistently outperformed MySQL in our tests, with the transactions per minute gap expanding at higher numbers of warehouses (a proxy for workload).

For both CPU/RAM combinations we tested, at 16 TPC-C Warehouses, MySQL gave errors at 16 users. At 32 users, MySQL gave errors at 8 users.

In the case of the 1 TPC-C warehouse with 8 CPU and 16 GB RAM, Actian Zen transactions per minute times were faster than MySQL. In the case of the 4 TPC-C warehouses, Actian Zen transactions per minute times were around 3X faster at all user counts than MySQL. In the case of the 16 TPC-C warehouses, Actian Zen transactions per minute times were around 5X faster at all user counts 1 through 8, at which point we received errors with MySQL. We never hit the upper limits of Zen in our tests. In the case of the 32 TPC-C warehouses, we received errors with MySQL at 8 users.

We moved up to 16 CPU and 32 GB RAM and the numbers were similar to the 8 CPU, 16 GB RAM test.

In the case of the 4 TPC-C warehouses, Actian Zen transactions per minute times were again around 3X faster at all user counts than MySQL. In the case of the 16 TPC-C warehouses, Actian Zen transactions per minute times were around 5X faster at all user counts 1 through 8, after which we received errors with MySQL. In the case of the 32 TPC-C warehouses, we received errors with MySQL at 8 users.

In our final test at 64 warehouses, Actian Zen transactions per minute times were around 6X faster at all user counts 1 through 4, after which we received errors with MySQL.

These tested operations underlie nearly all operations that occur on an embedded database for an IoT or mobile implementation. In a real world equivalency comparison, the number of our benchmark warehouses is equivalent to the number of sites or locations in real environment. The number of benchmark users simulates the same number of downstream devices with active, simultaneous, concurrent connections to the database. Thus, the goal of our benchmark was to show how these embedded databases perform as complexity and concurrent usage expands.

Actian Zen is a mature platform for embedded database applications with nearly 40 years of engineering and development behind it. Also, Zen's Turbo Write Accelerator is part of its performance advantages. Since it costs much less to continue writing than to stop and restart, contiguous writes are significantly faster than non-contiguous writes, a common scenario in a multitude of time series use cases with IoT. The Turbo Write Accelerator (TWA) pre-allocates open slots within the physical file so that multiple pages can be written as a single coalesced page—improving I/O performance and reducing the overhead of interaction with the operating system.

The result of the application of the methodology to the architecture, both explained herein and replicable, shows a marked performance advantage to Actian Zen. Overall, Actian Zen is an excellent choice for companies needing high performance and a scalable embedded database.

About The Authors

William McKnight is President of McKnight Consulting Group (MCG) (http://www.mcknightcg.com). He is an internationally recognized authority in information management. His consulting work has included many of the Global 2000 and numerous midmarket companies. His teams have won several best practice competitions for their implementations and many of his clients have gone public with their success stories. His strategies form the information management plan for leading companies in various industries.

Jake Dolezal has 25 years of experience in the Information Management field with expertise in business intelligence, analytics, data warehousing, statistics, data modeling and integration, data visualization, master data management, and data quality. Jake has experience across a broad array of industries, including: healthcare, education, government, manufacturing, engineering, hospitality, and gaming.

MCG services span strategy, implementation, and training for turning information into the asset it needs to be for your organization. We strategize, design and deploy in the disciplines of Master Data Management, Big Data Strategy, Data Warehousing, Analytic Databases and Business Intelligence.

About Actian

Actian, the hybrid data management, analytics and integration company, delivers data as a competitive advantage to thousands of customers worldwide. Through the deployment of innovative hybrid data technologies and solutions Actian ensures that business critical systems can transact and integrate at their very best – on premise, in the cloud or both. For more information about Actian Vector and the entire Actian portfolio of hybrid data management, analytics and integration solutions on-premise or in the cloud, visit www.actian.com.