STEC™ s1120 PCIe Accelerators Increase Performance and Reduce TCO for MySQL Scale-up and Scale-out Configurations

Benchmark tests point to the solution's superiority over HDDs in addressing the challenges of scale for MySQL databases.

STEC®

Contents

Executive Summary

MySQL, the most popular open source database in the world, is frequently used to support Web 2.0 applications and cloud deployments. Scaling MySQL can be a challenge when workload demand starts to strain the database. One option is to scale up by upgrading to a more powerful server, which works well to a certain point of scale. Beyond that, the only option is to scale out by adding more servers to a partitioned MySQL deployment, which is more complex but offers virtually unlimited scalability. The STEC™ s1120 PCIe Accelerator dramatically increases MySQL performance in both scale-up and scale-out configurations. It is an enterprise-class PCIe SSD with exceptional performance, low latency, low power consumption and high durability for hosting business-critical data and mission-critical applications.

To demonstrate its performance advantage, STEC benchmarked an s1120 PCIe Accelerator against six (6) high-performance SAS drives under a MySQL workload using the SysBench test tool. Test results showed that the STEC solution delivered significantly greater performance:

- 9.3x more transactions per second
- 9.2x faster average completion time
- 9.2x faster total execution time
- 11.6x higher application CPU utilization

Therefore, a scale-up MySQL deployment with an STEC s1120 PCIe SSD can support 9x the workload in a server, so organizations maintain this simpler architecture for longer before having to consider scale-out options. Comparing the total cost of ownership (TCO) of the two configurations, the STEC s1120 PCIe Accelerator was found to deliver 5.4x greater MySQL performance per dollar.

In summary, the STEC s1120 PCIe Accelerator is an excellent solution for scaling MySQL databases, and offers important benefits:

- Dramatically faster performance
- Superior price-performance
- Lower IT costs through reduced infrastructure

The Challenges of Scaling MySQL

MySQL is the most popular open source relational database in the world. MySQL is a solid and functional database platform with a development history that stretches for 20 years—all the way back to 1994. The source code is available under the GNU General Public License, so software licensing costs are not an issue. Many Web 2.0 applications and cloud deployments use this database. In fact, Facebook, Twitter, Flickr, Wikipedia and YouTube rely on MySQL.

Scalability is a key challenge for MySQL deployments. As workload demand grows, the size of the database and the number of transactions per second begin to strain the system. Administrators have to find a way to expand performance and capacity, preferably in a manner that minimizes cost and complexity. The two primary methods for scaling MySQL are:

Scale Up

To "scale up" MySQL is to run the database on a more powerful server by adding CPU cores, memory and storage capacity, and I/O. This is a straightforward solution that only involves the cost of hardware. No architecture or code changes are required. It does not entail significantly more power, cooling or rack space. The downside is that MySQL has a limited ability to utilize hardware resources. Adding CPU cores and memory improves performance to a point, and then incremental improvements diminish and eventually turn negative. (Negative returns might seem counter-intuitive. It is due to limitations in the database engine itself. With too many concurrent user sessions, the processes start to interfere with each other and throughput declines.) As such, scaling up works well to a point, but beyond that limit, the only option is to scale out.

Scale Out

To "scale out" is to add more nodes to a system or, in this case, to add more servers to a partitioned MySQL deployment. The primary method for scaling out MySQL is sharding. Database tables are split and distributed among multiple databases on different servers. Each database instance stores a range of rows from the tables, and the nodes are unaware of each other. Sharding is complex because application logic must be modified to account for split database tables. This entails a significant upfront development effort and more complex operational management, but sharding offers virtually unlimited scalability of database performance and size. All of the major MySQL database deployments, such as Facebook and YouTube, use sharding.

STEC s1120 PCIe Accelerators Increase MySQL Performance

In both scale-up and scale-out configurations, the STEC s1120 PCIe Accelerator dramatically increases MySQL performance. The enterprise-class PCIe SSD installs easily in a server's PCIe slot to deliver the exceptional performance, low latency, low power consumption and high durability required for hosting business-critical data and mission-critical applications. Storing MySQL data on STEC PCIe SSDs, rather than on traditional disk drives, will greatly increase the throughput and response time per server.

STEC's superior PCIe-based solid-state architecture is purpose-built to reduce storage access latencies and speed up application performance. By connecting server CPUs directly to the flash storage using the most efficient, reliable and scalable architecture, STEC PCIe SSDs provide low-latency data access for host applications while eliminating the tax on host CPU and DRAM—accelerating high levels of application performance.

In addition, STEC's PCIe SSD comes with built-in enterprise-class reliability that makes it a solution suitable for storing business-critical data. The s1120 PCIe Accelerator protects all data in flight or at rest (on the card) from loss or corruption in the event of flash media errors or unforeseen power loss or power failure events, providing a highly reliable solution for enterprise deployments. This means that STEC PCIe SSDs can be used for primary storage—not just for caching applications. PowerSafe™ Technology ensures that data on board PCIe SSDs is protected against unscheduled power outages. Using patented CellCare™ Technology, STEC delivers unprecedented endurance while using multi-level cell (MLC) media—enabling the use of less-expensive MLC-based STEC PCIe SSDs, even in write-intensive caching and database applications. Finally, STEC PCIe SSDs are backed by a five-year warranty to ensure their robustness in business-critical enterprise applications.

Test Methodology and Configuration

To demonstrate the performance advantage of the s1120 PCIe Accelerator, STEC benchmarked an s1120 800GB MLC PCIe SSD against six (6) high-performance 450GB SAS drives under a MySQL workload. The PCIe SSD was pre-conditioned to reflect real-world operating conditions and performance. The test tool was SysBench, a multi-threaded benchmark for evaluating servers running under a heavy database load. A test was run for each storage configuration on an x86 Linux server running MySQL with 200GB database under steady-state conditions. A variety of OLTP performance statistics were collected and compared.

Storage configuration

1 x STEC s1120 PCIe Accelerator MLC SSD
800GB usable capacity
Pre-conditioned to reflect real-world deployment

6 x Hitachi 450GB 15K RPM SAS disk drives
Software RAID controller
RAID 10 (mirrored, striped)
1.2 TB usable capacity

Server configuration

2 x Intel Xeon CPU E5620 @ 2.4GHz - 8 cores
8GB DDR3 memory
Red Hat Enterprise Linux Server release 6.1/ 2.6.32-131.0.15.el6.x86_64
Ext3 file system
200GB MySQL database file size

Test tool configuration and method

SysBench 0.4.12.5
MySQL package version: mysql-5.1.61-1

Test method for disk drives

/sysbench -- debug=off -- test=oltp -- mysql-user=mysql -- oltp-table-size=800000000 -- mysql-db=test -- oltp-table-name=stest -- num-threads=20 -- max-requests=1000000 -- oltp-auto-inc=off -- mysql-engine-trx=yes run

Test method for PCIe SSD

/sysbench -- debug=off -- test=oltp -- mysql-user=mysql -- oltp-table-size=800000000 -- mysql-db=test -- oltp-table-name=stest -- num-threads=20 -- max-requests=1000000 -- oltp-auto-inc=off -- mysql-engine-trx=yes run

Results

The server configuration with the STEC s1120 PCIe Accelerator demonstrated a performance improvement of approximately 9x over the configuration with six (6) HDDs in the MySQL SysBench test (see Chart 1 and Figure 1). Transactions per second and read/write requests were 9.3x higher for the STEC configuration. These metrics illustrate the IOPS workload each system was able to support. Average request completion time and total execution time both showed more than a 9x improvement. Average CPU utilization was 11.6x higher for STEC—indicating that storage I/O was, in the case of HDDs, the bottleneck in system performance. Since the PCIe SSD boosts storage I/O, the system is able to more fully utilize other resources like CPU and memory, rather than sit idle for much of the time. By improving server hardware utilization, PCIe SSDs enable organizations to derive a higher return on that investment.

The test results show that a server using an STEC s1120 PCIe Accelerator performed more than 9x the workload as one with six (6) high-performance SAS HDDs. Therefore, a scale-up MySQL deployment can support 9x the workload in a server, so organizations maintain this simpler architecture for a longer duration before having to consider scale-out options.

Chart 1: MySQL Performance Statistics - Solution with STEC s1120 PCIe Accelerator vs. HDDs

Comparison	6 x SAS HDD 15K RPM 450GB	STEC s1120 PCIe Accelerator MLC 800GB	Improvement Factor
Transaction per sec	27	251	9.3
Read/write requests per sec	517	4760	9.2
Average request completion time (ms)	735	79.8	9.2
95th percentile request completion time (ms)	1168	153	7.6
Total execution time	10.2 hrs. (612 mins.)	1.1 hrs. (66.5 mins.)	9.2

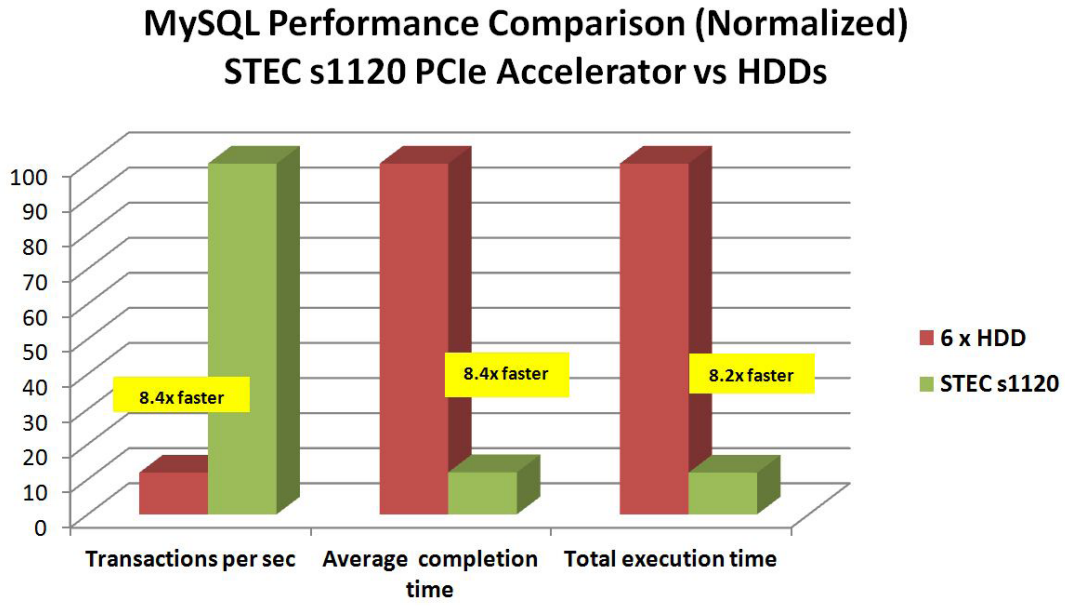

Figure 1: *MySQL Performance Comparison (Normalized) - STEC s1120 PCIe Accelerator vs. HDDs*

STEC Price-Performance Advantage

The STEC s1120 PCIe Accelerator delivered 5.4x greater MySQL performance per dollar than the HDDs (see Chart 2 and Figure 2). The TCO of the two configurations were compared, including hardware cost and estimated three-year operating cost. Since the STEC solution delivered 9x greater performance, it would take nine (9) servers with HDDs to equal the performance of one with the PCIe SSD. Since most MySQL deployments use a mirror PCIe SSD configuration for redundancy, the TCO was performed by assuming a mirror PCIe SSD configuration.

Chart 2: MySQL Price-Performance - STEC s1120 PCIe Accelerator vs. HDDs

	6 x SAS HDD 15K RPM 450GB	STEC s1120 PCIe Accelerator MLC 800GB
List price of base server (dual Intel Xeon E5620)	$9,999	$9,999
List price of storage (HDDs or SSD)	$1,570	$11,414
Est. 3-year operating cost per server	$3,455	$3,455
Total server cost	$15,024	$24,869
Cost per transaction	$442	$87

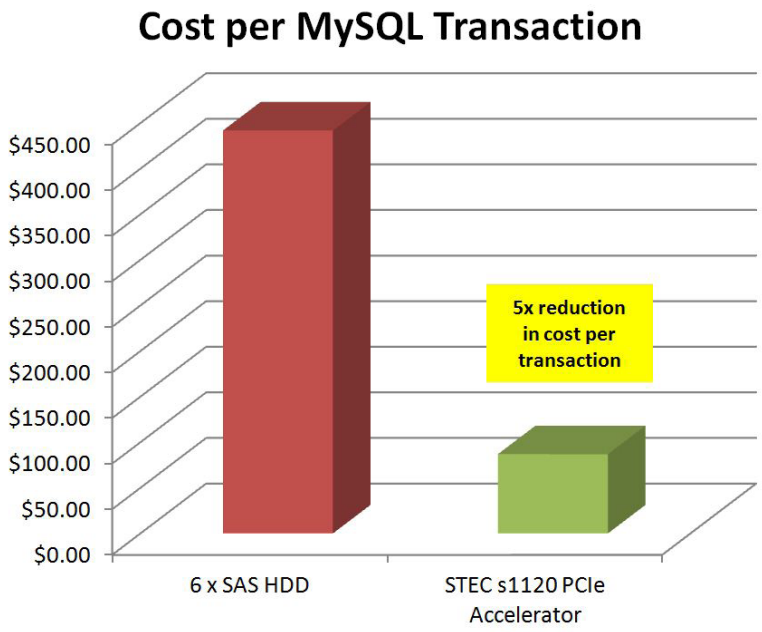

Figure 2: *Cost Per MySQL Transaction - STEC s1120 PCIe Accelerator vs. HDDs*

In a scale-up implementation, the STEC solution saves money by delivering greater price-performance as well as deferring, perhaps indefinitely, the costly and complex change to a scale-out architecture. For scale-out, it saves money by reducing the number of servers needed and eliminating the operating expenses associated with more complex environments (see Figure 3).

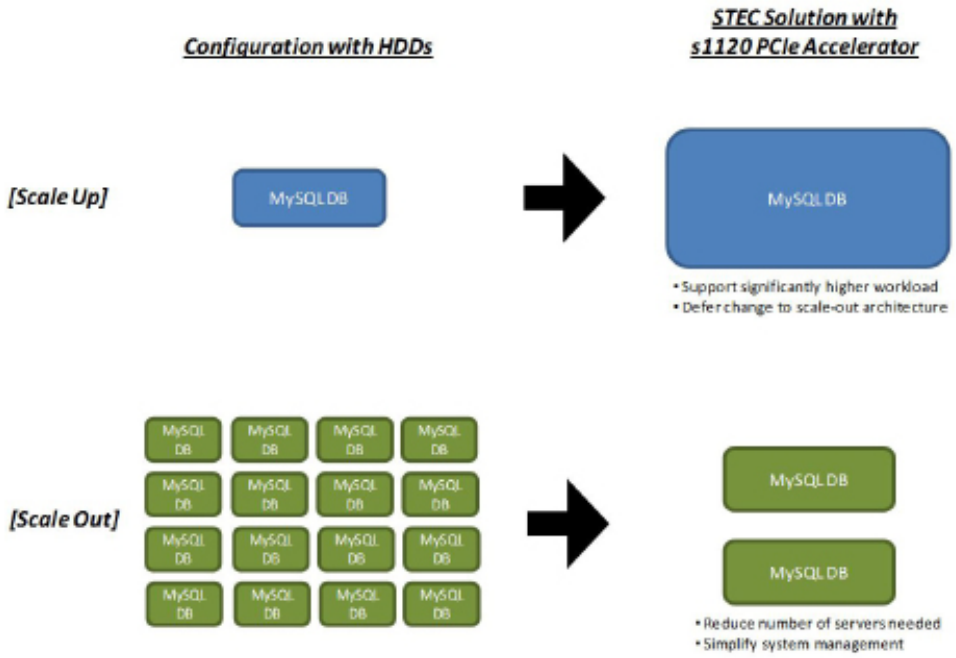

Figure 3: *MySQL Server Consolidation with STEC Solution*

Conclusion

The STEC s1120 PCIe Accelerator is an excellent solution for addressing the challenges of scale for MySQL databases. Whether your organization is employing scale-up or scale-out architecture for MySQL, storing data on STEC PCIe SSDs instead of HDDs offers numerous benefits:

- **Dramatically faster performance**—9x greater IOPS performance
- **Superior price-performance**—5.4x greater MySQL performance per dollar
- **Lower IT costs**—Scale-up MySQL configurations can continue to support rising application demands before having to consider more complex architecture, and scale-out configurations only need a fraction of the number of servers

For more information on STEC products, solutions and technology, please visit www.stec-inc.com

facebook.com/userstecinc twitter.com/stec_inc youtube.com/user/stecincssd

+1.949.476.1180
3001 Daimler Street, Santa Ana, CA 92705