



Scaling MySQL without Compromise

Introduction

The database market segment continues to grow at a rapid pace spurred by innovative database capabilities and new set of applications. MySQL is the fastest growing database among SQL-based database deployments. MySQL continues to be the choice for a broad range of database developers who want a high performance database that is reliable, affordable and easy to use. Customers are realizing the cost of ownership advantages and platform flexibility that MySQL offers. The open source LAMP (Linux, Apache, MySQL, PHP/Perl/Python) stack is now being embraced by mainstream retailers, banks and telecom companies, going beyond the Internet giants such as Google and Yahoo, who were the early adopters.

Key Challenges

Success with any MySQL deployments requires a careful consideration of operational requirements, performance needs and cost objectives during the design and deployment of the database infrastructure. Data volume along with data access traffic continues to grow exponentially, dictating the need for a scalable infrastructure. Traditional hard-drive based infrastructures are unable keep up with such harsh demands. IT professionals have tried to overcome such scalability challenges with limited success by simply throwing more hardware to the problem or leveraging a variety of operational techniques. Such band-aid solutions increase the capital/operational expenses and the management complexity of a MySQL deployment leading to significant cost of ownership disadvantage.

Increase DRAM on Master Servers

Throwing more DRAM at the server can become an expensive proposition due to DRAM's non-linear price-to-density relationship. Also, MySQL expects data to be persistent. Given that the DRAM is volatile memory, there is always a need to flush the data to backend hard drives periodically which significantly limits the performance gains. So, the steady state performance of the technique may not yield the expected benefits.

Before Virident tachION

- Less scalable
- Complex environment
- More servers
- Inflexible infrastructure

After Virident tachION

- Highly scalable
- Simple environment
- Less servers
- Flexible solution

Implement Master/Slave(s) Configuration

In this configuration, a single Master server is assigned for all write updates and one (or many) Slave servers provide read-only operations. Master updates the data to the Slaves using asynchronous MySQL replication. This technique helps in scaling of read-intensive workloads. However, performance scalability is limited by the write traffic on the Master and MySQL replication's ability to keep up.

Database "Sharding"

The concept of sharding stems from breaking down a large database into smaller chunks spread across a number of distributed servers for performance and scalability. Sharding requires careful architectural design prior to deployment and complex ongoing data management. Achieving SLA reliability can be a challenge due to these distributed "shard" databases.

Deploy FC-SAN Storage

FC-SAN can be a better option than DRAM for larger datasets. Unfortunately, its large latencies pose significant impediments to MySQL performance scaling. Latency is critical for typical management operations such as long queries, replication etc. Moreover, deploying and managing an FC-SAN can be very expensive.

Virident tachION Advantage

Flash-based storage has created a paradigm shift in the way in which data is stored, managed and accessed in MySQL environments by delivering best of both worlds – performance of DRAM and non-volatile nature of disk-drives. The read and write performance issues of the traditional approaches are instantly alleviated, ensuring significant scalability of the database architecture.

Of the various flash-based devices in the market, Virident *tachION* PCIe-based SSD delivers the highest performance – highest bandwidth at lowest latencies. Virident's *tachION* addresses all the IO throughput and latency bottlenecks while meeting data reliability requirements of mission critical MySQL deployments. Benchmarking and customer deployments have shown that by simply moving data to the Virident *tachION* drive from traditional storage, users can get up to 15X performance gain over HDDs. The drive is available in a low profile form factor making it compatible with all servers. The *tachION* drive design is modular with an on-board, flash-aware RAID5 along with flash level ECC, ensuring enterprise class data availability even in the event of a flash module failure.

Blazing Performance

The *tachION* drive connects directly with the CPU using PCIe delivering latencies of 10s of μ s, orders of magnitude better than SAN or SAS-based flash devices. Such low latencies are critical to ensure fast application response times.


Minimize Slave Servers

Leveraging Virident's *tachION* as the primary store in slave servers in MySQL environments results in up to 15X improvement in transactions per second (TPS) compared to HDD arrays. Thus, Virident *tachION* delivers significant read scaling by allowing the dataset size to exceed the size of the DRAM without compromising the performance as is the case with disk-drives. This configuration also enables IT professionals to minimize the number of slave servers deployed for the same level of performance requirements which in

turn reduces both capital outlay and the operational complexity of the MySQL environments.

Scale Without Complex Sharding Technique

The *tachION* drive is very efficient in handling write traffic thus offering the most balanced read-write performance in the industry. By deploying *tachION* drives in the MySQL master server to store all of database dataset, write workload scaling issue is fully addressed as well. This means that there is often no need for complex sharding techniques. Reduction in management complexity of sharding enables the application team to focus on revenue enhancing activities such as introducing newer services etc. If the dataset is larger than the *tachION*'s capacity, striping the data across multiple *tachION* cards on a single server, locating the most IO-intensive data on *tachION* drive, or implementing flash-memory friendly caching solution can deliver similar benefits.



Summary

Virident's *tachION* drive enables MySQL environments to scale without compromise. It delivers superior performance while alleviating the concerns related to data availability of flash-based devices. Thus, it helps customers to adapt to the relentless growth of MySQL data and traffic in the most cost effective way.

For more information on Virident *tachION* drives, please visit www.virident.com