

Accelerate MySQL with SanDisk Memory File System NVMFS

Thomas Rochner
SanDisk
Germany

Schlüsselworte

MySQL, Flash, Performance, Atomic Writes, Compression, Endurance, NVMFS, Non-Volatile Memory File System, Double-Buffer Writes

Einleitung

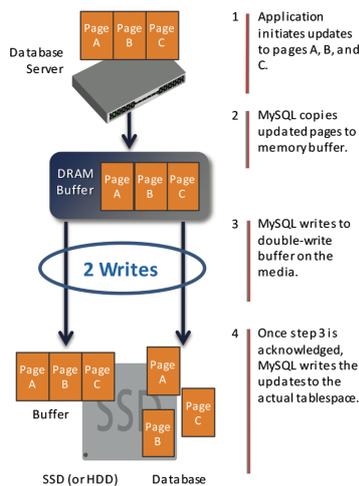
Working with the Oracle MySQL, Percona and MariaDB communities, SanDisk has created an optimized solution that focuses on solving specific performance problems in the database. This solution uses a new flash-aware file system called NVMFS (Non-Volatile Memory File System), which is uniquely capable of translating standard Portable Operating System Interface (POSIX)-compliant file interfaces into “flash-aware” calls to the underlying device. With NVMFS, open source databases like Oracle MySQL, Percona Server and MariaDB can become “flash-aware” and solve many daunting performance problems.

Atomic Writes and Compression

While the performance of open source databases has improved in many areas over many releases, there are two problem areas in MySQL that can significantly benefit from the use of these “flash-aware” interfaces:

Oracle MySQL, Percona Server and MariaDB write all table space data twice to ensure data integrity during system failure, which leads to twice the amount of actual data per write. This “double-write” problem exists for both spinning media and flash.

Traditional MySQL Writes



MySQL with Atomic Writes

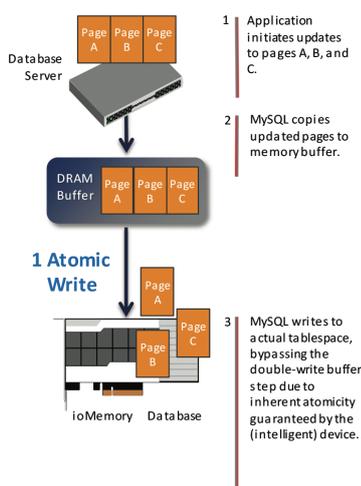


Abb. 1: Traditional writes vs. Atomic writes

These databases have supported compression for a long time, but the implementations have been shown to perform poorly on flash. Generally, most customers do not run compression with production workloads on flash, due to severe latency and throughput penalties.

These problems can be addressed by using both SanDisk NVMFS and ioMemory PCIe devices with a combined approach of hardware and software that uniquely enhance these open source databases. This significantly improves I/O performance, provides consistently low latency, reduces latency variation, and enhances the life expectancy of flash devices.

The SanDisk NVMFS Atomic Writes feature solves the double-write problem. That results in a significantly increased life expectancy of devices at similar throughputs, while providing consistent and lower I/O latency. The Atomic Writes feature uses flash-aware interfaces to write an entire database commit in one operation.

Working with the MySQL community, SanDisk has helped to enhance compression in these databases to enable customers to get the benefit of compression with little performance impact. With this NVMFS compression, applications see up to a 50% increase in capacity while driving transactions per

second at a rate within 10% of the uncompressed rate.

Combining Atomic Writes with NVMFS compression reduces write operations by close to 75%¹, compared to running uncompressed and without atomic writes. This significantly enhances the life expectancy of flash devices and reduces latency.

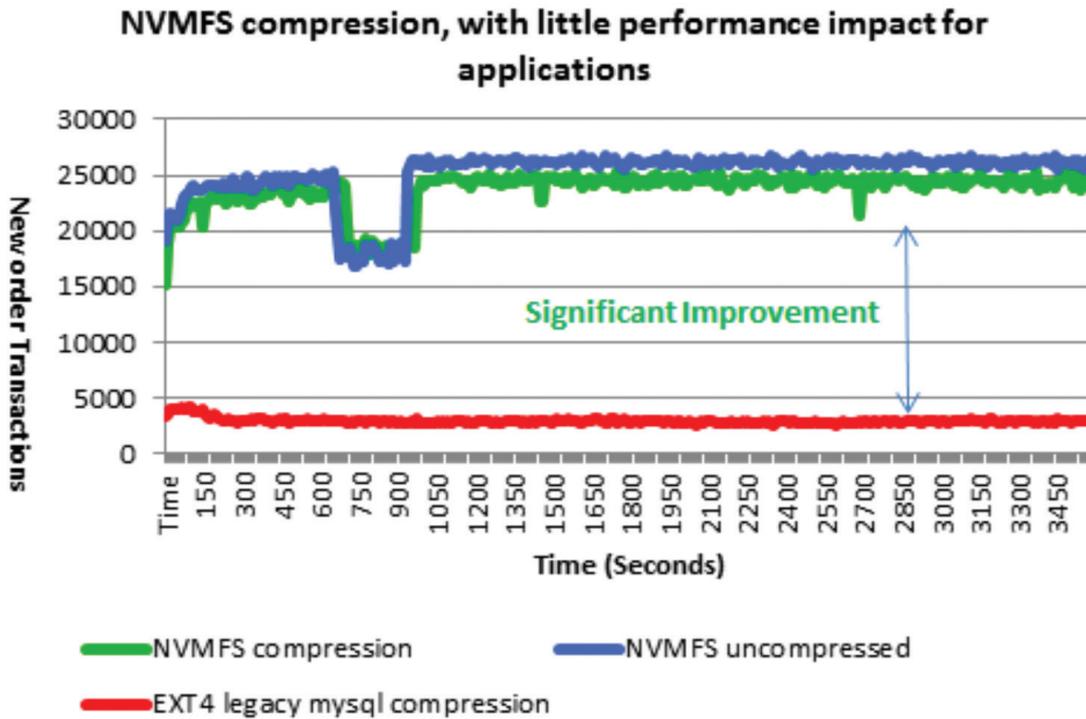


Abb. 2: NVMFS Compression vs. Legacy compression . TPC-like benchmark : 1000 warehouses – 75 GB MySQL Buffer pool, Maria DB 10.0.15

About NVMFS

NVMFS is a new POSIX-compliant file system designed and developed by SanDisk. It is designed from the ground up for the newest generation of flash and non-volatile memory devices, enabling them to overcome many of the suboptimal features of traditional file systems that were not designed for these technologies. It eliminates duplicate logic and leverages new primitives for optimal flash performance and efficiency.

NVMFS provides low latency I/O and high throughput, leveraging the full capabilities of the underlying flash block device. This makes it an ideal match for high performance database environments. Popular relational databases such

as Oracle MySQL, MariaDB, and Percona Server can use the flash-aware interfaces exported by NVMFS (Atomic Writes, PTRIM, etc.) to bypass I/O inefficiencies in the database itself. This in turn achieves much greater throughput and much lower transaction latencies.

The modern database bottleneck is I/O latency – the time it takes to submit and complete one or more operations in parallel. Databases typically need to log small amounts of data quickly during transaction commits, making the time it takes to write data to a persistent state important. The low-latency architecture of NVMFS makes it a perfect fit for ACID (Atomicity, Consistency, Isolation, Durability)-compliant updates.

Die Seitenzahl wird von uns eingefügt!

Bitte fügen Sie Ihre Kontaktadresse hinzu.

Kontaktadresse:

Thomas Rochner

SanDisk GmbH

Kaiserswerther Str. 115

D-40880 Ratingen

Telefon: +49 (0) 800-183-0258

E-Mail thomas.rochner@sandisk.com

Internet: www.sandisk.de