Case Study:
Accelerating virtual machines with Kingston DC500M SSDs

Introduction

At Hardwareluxx, we operate one of the largest German IT websites under the domain www.hardwareluxx.de with news, test reports and an integrated discussion forum. Since 2018, our websites and services have been delivered via two Synology FlashStation FS3017. Within these servers, we operate a mix of SSDs (Kingston DC400) and 2.5-inch hard drives. To increase the performance of the FS3017 for a specific application, an older array of eight Seagate Constellation HDDs were converted and upgraded to the modern Kingston DC500M SSD. The virtual machines (VMs) that operate on them should be greatly accelerated to make the applications run faster.

Challenge

Our websites are heavily visited due to the 2.4 million monthly visitors and the 270,000 forum members. As such, we need to split the workload, by having web server and MySQL servers run separately on different hardware. While the web server can be operated virtualized on the Synology FS3017, the database has been outsourced. After switching to new software for the nearly 27 million-post forums, there was the possibility to use Elastic search as a new option for the post search function, which was previously very resource-intense. For this purpose, a separate VM was created, which, however, requires a correspondingly fast IO background on the server due to the response time of the search.

Attempts to run the Elasticsearch VM sensibly due to the size of the database with an HDD RAID federation caused search delays. With heavy load on the search function, an overload was to be expected. In the absence of capacity on other free volumes that were based on SSDs, the HDD array had to be replaced with SSDs. Instead of the eight 1 TB hard drives, set to eight correspondingly large SSDs to maintain an identical capacity.

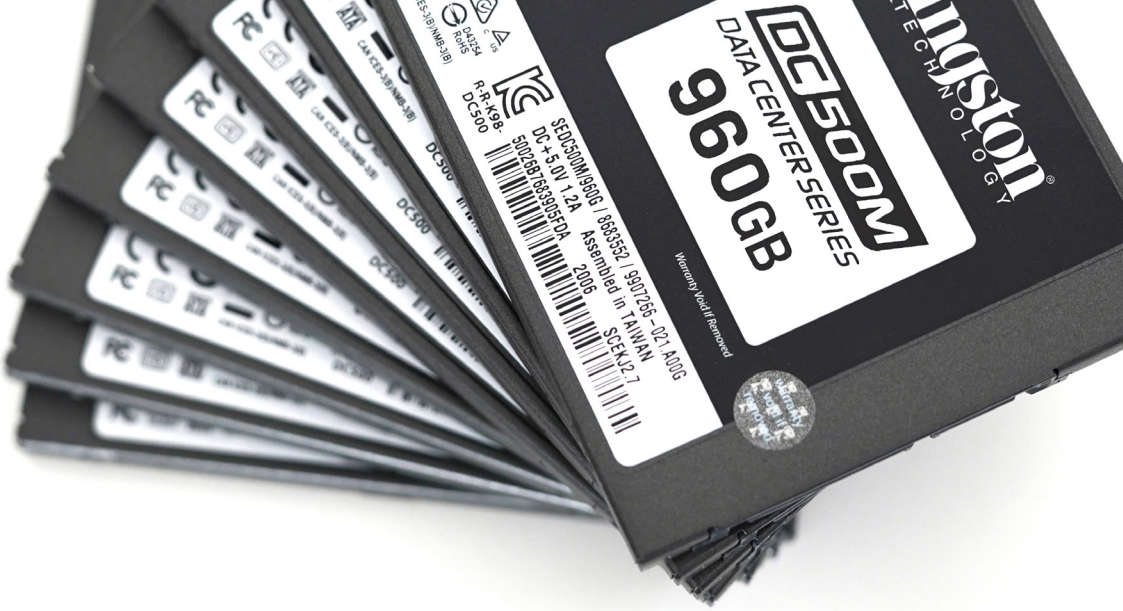

Approach

Kingston was selected to be used as they are a manufacturer that we have had good experience with using their DC400 SSDs. Of the 16 DC400s used, no drive had to be replaced after two years of use, so we confidently returned to Kingston as a manufacturer.
However, we needed to consider the correct SSD selection from their offering.

To select the right SSDs, we must first analyse which workload on the SSDs ends up being the main. If it is mainly read access, it is a good option to choose the Kingston DC500R. In our case, it is assumed that the viewfinder results must be constantly updated, correspondingly both writing and reading accesses take place. That is why we chose the DC500M (Mixed). The M-Series has a higher number of replacement cells due to overprovisioning and thus a significantly higher specification in terms of write access.

Compared to the DC400 series still in use at the same time and the DC450 series available in the meantime, the DC500 series has the general advantage that built-in capacitors ensure safe storage of the data even in the event of a power failure. If written to the drive during a power cut, the mounted capacitors help to safely write the data contained in the caches to the flash devices and enable a secure shutdown of the SSD. For databases, this is an additional security option that can increase the integrity and consistency of the data if there is no other way to prevent a power failure.
Already by using the DC400, we realized that the Kingston SSDs and the firmware included from a QoS point of view are suitable for delivering consistently low response times and high IOPS, even if the drive is filled with a lot of data. This is also a criterion that is very important when selecting the IO subsystem.
More professional solutions, such as a DC1000M or similar, were out of the question due to the necessary SATA protocol of Synology FS3017 servers. As capacities for our SSDs, we chose 960 GB, as we considered this to be the optimum between the price-performance ratio due to the capacity and installation options on the other side within the FS3017, and we were able to cover our storage requirements with eight DC500M in the planned RAID F1 - a special form of RAID5 from Synology.

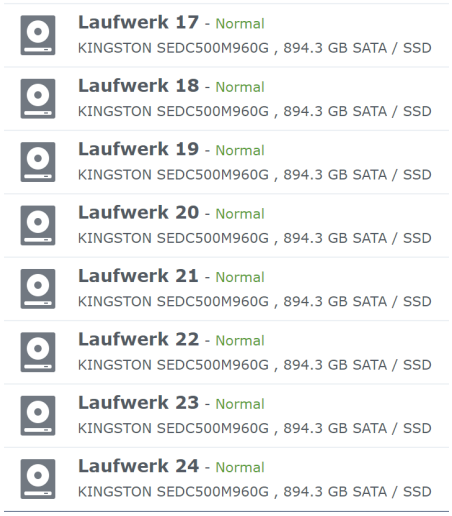

The SSDs are combined in an array in the FS3017.

Execution

Before replacing the hard drives, we conducted a final series of tests on the old HDD RAID5 to measure the performance advantage of the SSDs.

Due to the RAID controller installed in Synology FS3017, it was possible to read the array sequentially at approximately 520 MB/s and write to the array. However, the array also achieved poor access times and low IOPS for 4K data, which is common on disks.

After the conversion to the Kingston DC500M, we tested the performance of the array with a RAID5, a RAID6 and a RAID F1, but also with a RAID0. Interestingly, Synology's FS3017 seems to be partially limiting the sequential accesses already, in reading we reached a maximum of 1,200 MB/s, in the letter 1,900 MB/s. The IOPS, which we were able to increase by a factor of 4, also significantly showed improvement in all configurations.

After the different measurements, we decided on an array of six SSDs in a RAIDF1 - the Synology variant of a RAID5, in which an SSD is more clearly strained by the permanent writing of parity. The RAID6 had minimally poorer writing performance. Two SSDs serve as hot spare to jump in in the event of a possible failure.

Conclusion

By switching the array to Kingston's DC500M, we were able to achieve our goal of running our Elasticsearch VM with very high performance to enable our readers to search quickly, responsively. At the same time, we can benefit from the more modern architecture with greater resilience and can add more VMs to the array due to increased performance. In addition, we save a few watts by lower power consumption of the SSDs compared to the hard drives.

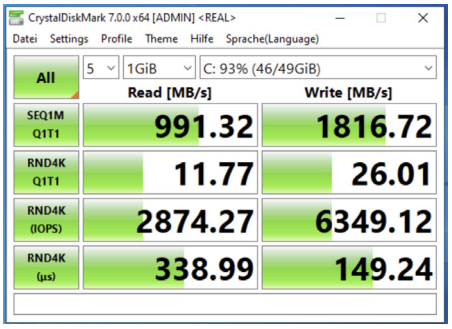

RAID 5: Fast setup with good IOPS with just six SSDs.

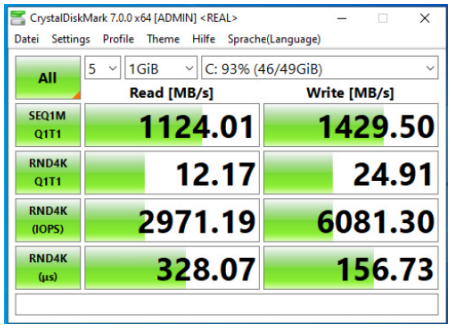

RAID6: Higher reliability, but lower write speeds.

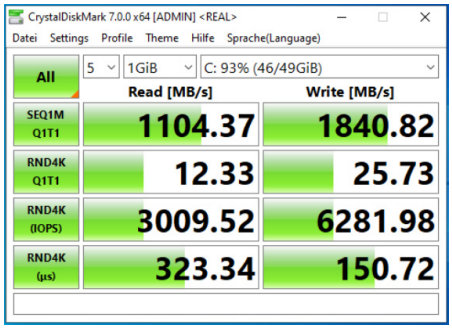

RAIDF1: Special RAID5 configuration from Synology preventing simultaneous failure.