Datasheet

SANtricity SSD Cache

Boost application performance and automate performance tuning with flash

KEY FEATURES

Fast Data Delivery
Accelerates read I/O by up to 100 times compared with standard disk drives alone

Analytics-Based Automation
Automatically identifies hot data to increase application performance while reducing administration costs

Optimized Cost and Performance
Optimizes your storage purchase by leveraging SSDs to meet your overall IOPS performance requirement and by using lower-cost HDDs for capacity

The Challenge

Transactional applications are the core of the customer-facing systems for most businesses, and IT departments constantly try to find ways to decrease response times, handle more users, or add new capabilities to these systems. Slow transactions cost businesses productivity, lowering profits and affecting customer service.

Application performance is often restricted by the performance of individual disk drives. The transactional performance of disk drives themselves has not changed much over the past 20 years, so the main method of increasing transactional performance is to increase the number of disk drives. Flash technology is changing everything, however. Solid-state drives (SSDs) provide dramatically faster response times than spinning drives do and offer close to 100 times the transactional performance, but they are more expensive. Organizations must optimize their storage costs and administrative demands to meet their application response times and capacity requirements as their businesses evolve.

The Solution

Many transactional applications are very read intensive, meaning that they have a high percentage of read I/O compared with writes. For these applications, deploying NetApp® SANtricity® SSD Cache on E-Series delivers the performance benefits of flash without the high cost of all-flash capacity. Data is not uniformly hot; mixed workloads have a small subset of the overall dataset that accounts for most of the performance needs of the workload. This subset is called the working set, active data, or hot data. If the working set data is mostly read, and it all fits within the SSD Cache, your application realizes the full benefit of flash performance, and transactions are correspondingly much faster.

SSD Cache accelerates data access through the caching use of solid-state disks in the drive trays, and it is expandable to 5TB per storage system. SSD Cache is assigned to specific volumes and acts as a memory extension to significantly reduce latency and to speed execution of random reads. After SSD Cache has been configured, there is no ongoing management or tuning; with SSD Cache, you can set it and forget it.

NetApp®

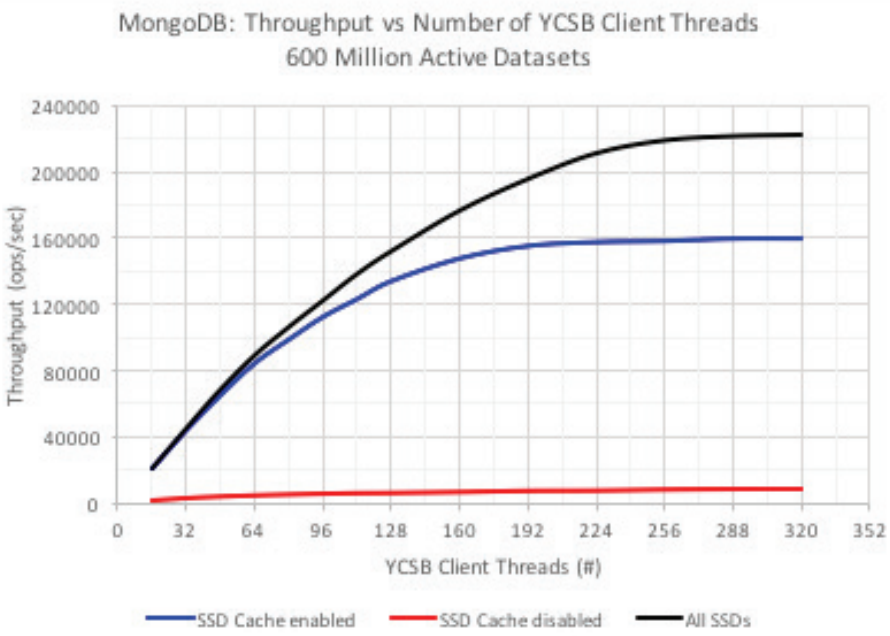

Figure 1) Throughput versus client threads with SSD Cache enabled, SSD Cache disabled, and all SSDs.

Analytics-Based Cache

SSD Cache provides intelligent read caching capability to identify and host the working set that is hot on the SSDs. Because this caching approach works in real time and uses pattern-recognition algorithms, users do not need to set up complicated policies to define the trigger for data movement between tiers.

Dynamic changes in the workload are common in shared storage infrastructures that host a wide variety of applications on the same physical infrastructure. SSD Cache reacts to dynamic changes in the various workloads in real time, rather than waiting for the next data movement window to arrive, as is the case with some automated tiering solutions.

SSD Cache also employs pattern-recognition algorithms to preload cache with additional data that is likely to be read next. Pattern-recognition algorithms also avoid loading data that is unlikely to be read and that would "pollute" the cache. SSD Cache identifies large sequential writes as data that is unlikely to be read thousands of times in the next few minutes. As a result, it does not cache the I/O—it lets the writes be serviced as high-throughput writes directly to the disk drives. Predictive algorithms prevent the cache from affecting performance in write-intensive workloads, and these algorithms work even better in predictable read-intensive workloads.

Optimized Cost and Performance

Because SSD Cache provides analytics-based read caching, the vast majority of the workload capacity can be hosted on HDDs, the least-expensive media. This feature reduces overall cost while improving the overall performance with higher throughput and lower response times. SSD Cache optimizes the right media for the right use at the right time. It does so by leveraging the superior performance of flash media for the most frequently accessed blocks while using lower-cost HDDs for capacity requirements.

SSD Cache is included with the SANtricity operating system; the only extra purchase is for the SSDs that are to be used as cache. Your NetApp representative will help you determine the performance benefits that you can expect for your applications, based on their specific I/O profile.

About NetApp

Leading organizations worldwide count on NetApp for software, systems and services to manage and store their data. Customers value our teamwork, expertise and passion for helping them succeed now and into the future.
www.netapp.com