LITE-ON Storage is now Solid State Storage Technology Corporation

The Streams Directive

NVMe™ Hints for Optimal Writes

Initial support for Streams as part of the Directives feature was introduced in the non-volatile memory express (NVMe™) 1.3 specification in 2017. Although support exists also for SCSI and SATA drives, those

standardizations are based on the NVMe™ model. Directives as a feature enable better communication between the host and a SSD controller through the passing of relevant metadata, or data about data, through existing NVMe™ commands. Metadata can be useful in input/output (I/O) decision-making by transferring information such as workload type while also allowing the host to specify more educated placement for the storage media.

Streams specifically work within the header of write commands as I/O-based Directives for passing allocation metadata. This means that streams may be utilized to associate logical blocks into groups of data based on processes, source machines, etc. This allows for optimization of writes in multiple ways such as performing operations at the block level which improves the performance and endurance of flash-based devices. As a result, garbage collection is also improved as data can be grouped more intelligently including by anticipated lifetime, all with less fragmentation.

According to the Storage Networking Industry Association (SNIA) the stream operation is ideal when different lifetime data is intermixed which normally causes an increase in garbage collection overhead and higher write amplification. The Streams Directive allows write operations to be grouped within a given stream which then ultimately places related data in a contiguous location. This reduces the write amplification factor (WAF) by allowing writes to occur more intelligently. All data within a stream should also be deallocated at the same time, that is trimmed and erased, thus also reducing maintenance overhead.

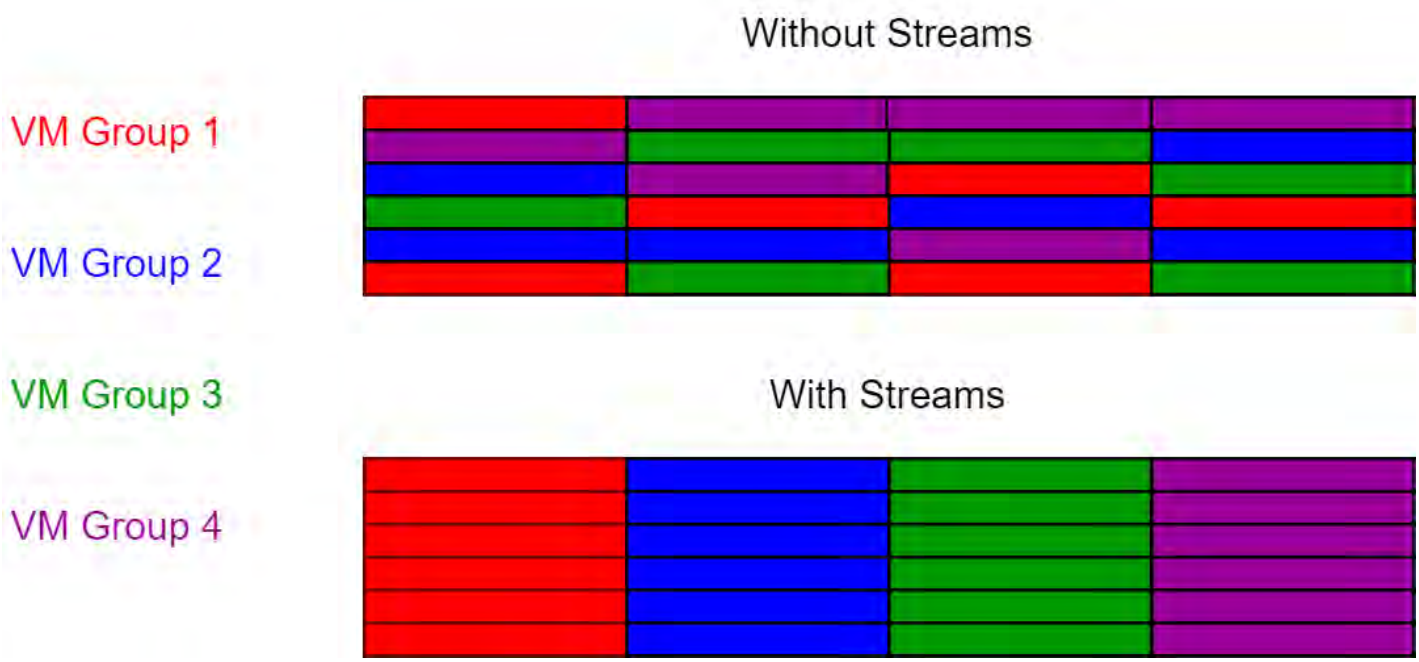

Organization of writes based on lifetime VM workload groups

Commands

There are two primary Directives commands, Directive Send and Directive Receive, beyond the existing Write command. Within the Streams Directive, there is also an identifier, that is a host-associated namespace, as well as resource allocation dependent on a namespace. In general, the Identify Directive will use the Receive command to return bitmaps of all directives, such as via Get Status, while the Send command will enable or disable a directive. Streams are specifically associated with writes and therefore the Write command as related to an identifier. This combination allows the Streams Directive to allocate per namespace and NVM subsystem, as by for example the Allocate Resources operation with Namespace Streams Requested (NSR) or Allocated (NSA).

The Receive command for the Streams Directive allows for the reporting of properties for both the NVM subsystem and namespace. For the subsystem this includes information such as the maximum amount of streams or Max Stream Limit (MSL), the amount free or available, the amount currently open or NVM Subsystem Streams Open (NSSO), and the total NVM Subsystem Streams Available (NSSA). For the namespace this includes the stream write size (SWS), stream granularity size (SGS), the amount of total streams allocated for the namespace or Namespace Streams Available (NSA), and the number of streams currently open via the Namespace Streams Open (NSO) parameter. Receive may also report the status of open streams and request resources for a specific namespace and host as detailed above.

The Send command for the Streams Directive can release identifiers and resources for specific namespaces and hosts. This identifier is related to the Write command whereby a namespace is associated with a host unless it is specified to be zero where it is treated as a normal write. The identifier also opens or releases the associated stream, for example with the Release Identifier or Release Resources operations, depending on the amount of control requested by the host. Although Directives can work with other I/O commands, the Streams Directive specifically works for writes within this framework.

Summary

Through the use of the Streams Directive, it is possible to give hints with metadata via write headers to generally improve the performance and endurance of non-volatile media as with NVMe™ solid state drives (SSDs). These hints link hosts to namespaces within the greater NVM subsystem structure, allowing for multiple streams at once which are segregated based on performance characteristics of the source data. For example, workloads with the same determined lifetime – that is, data that should exist and then be erased after a comparable, finite amount of time – can be written within the same stream and allocated rather than interleaved to a specific physical location on the storage media.

This allows the flash to operate on a more organic block-sized level which reduces unnecessary writes, and thus write amplification, while also improving performance by allowing for the organization of writes with predictable performance characteristics. Further, as this data is written, trimmed, and then erased as a block, the amount of overhead for maintenance or garbage collection is reduced, especially when combined with communication between the host and NVM subsystems for workload scheduling. This means more efficient timing of maintenance operations.

Streams, under the Directives heading, involve the Receive, Send, and Write commands along with an associated identifier. This coupled with information such as the Stream Write and Granularity Sizes enable the host to prepare the storage media for the grouping of incoming writes through size-based allocation. Operations and resources are also tracked by namespace and NVM subsystem to ensure sources, such as virtual machines, are appropriately serviced and grouped in a transparent manner. This gives the host as much or as little control as desired by the system administrator, ensuring optimal performance and endurance through selective and intelligent stream management.

ABOUT US

A subsidiary of KIOXIA Corporation, **Solid State Storage Technology Corporation** is a global leader in the design, development, and manufacturing of digital storage solutions. We offer a comprehensive lineup of high-performance customizable SSDs for the Enterprise, Industrial, and Business Client markets. With various form factors and interfaces, our SSD solutions help businesses simplify their storage infrastructures accelerating variable workloads, improving efficiency, and reducing total cost of ownership.

Learn more at www.ssstc.com