SOFTWARE-ENABLED FLASH™

Software-Enabled Flash™ Technology:

Understanding How Open Source Software Works with Purpose-Built Flash Storage Hardware

TECHNICAL BRIEF

In recent years, the exponential growth of data has propelled the need for faster and more efficient storage solutions. Traditional storage technologies have struggled to keep up with the demands of modern applications and workloads. As such, the combination of open source software (OSS) and purpose-built hardware has emerged as a powerful solution to address these challenges. This technical brief presents the concepts behind an emerging solution called Software-Enabled Flash (SEF) technology that leverages vendor-agnostic OSS designed to be compatible with purpose-built flash storage hardware. Use cases and associated benefits of this powerful new technology are included.

Understanding SEF Technology

SEF technology refers to the integration of OSS with purpose-built flash storage hardware creating a high-performance storage solution that is highly customizable through software and an application programming interface (API). The software component of SEF provides advanced features such as data management, caching and workload isolation. The SEF hardware component optimizes the underlying storage media to deliver exceptional performance, flexibility and reliability.

The OSS Component of SEF

OSS plays a pivotal role in SEF by providing the necessary tools and frameworks to manage a storage system efficiently. Open source projects that include the Linux® operating system, zoned file systems and flexible data placement are becoming interesting options to consider when building software-defined storage (SDS) solutions. Based on an open source platform, SEF offers developers flexibility and extensibility, as well as a vibrant development community to help foster continuous SEF innovations and improvements.

A zoned file system is a notable example for which SEF can further leverage its advanced features such as copy-on-write snapshots, data integrity checks and easy-to-use storage pool management. SEF utilizes zoned file systems at the software layer to enable high data integrity and redundancy, and efficient data deduplication.

Purpose-Built Hardware in SEF

While software forms the intelligence behind SEF, purpose-built hardware provides the mechanism to achieve the necessary performance and reliability requirements. In this case, purpose-built hardware refers to storage devices designed specifically to meet the demands of flash-based storage systems and feature high-speed interfaces, advanced memory controllers and custom firmware optimized for flash media.

Solid-state drives (SSDs) are a good example of purpose-built flash storage hardware (in standard interface options) used in SEF, utilizing flash memory as the storage medium. These drives incorporate wear leveling[1] algorithms, advanced error correction coding and power-loss protection[2] circuits, and when combined with OSS, deliver significant improvements in performance, endurance, SSD life expectancy and data protection.

SEF Use Cases

The capabilities of SEF target the following use cases:

High-Performance Databases

Database systems enable businesses with information related to sales transactions, product inventory, customer profiles, marketing activities, etc. These systems store data in a centralized location using searchable and sortable formats. SEF can dramatically enhance database system performance through OSS and purpose-built hardware enabling lower latencies, faster data access and improved overall throughput. This allows businesses to process larger workloads, support more users and deliver real-time analytics, all of which provides enhanced user experiences.

Virtualized Environments / Cloud Computing

Virtualized environments enable multiple virtual machines (VMs) to run on one physical system or cluster, sharing system resources as needed. Cloud computing is the delivery of different services through the internet including data storage, servers, databases, networking and software. These platforms rely heavily on efficient storage solutions. Through OSS and purpose-built hardware, SEF offers enhanced storage performance, fast VM provisioning, low application latency, improved resource utilization, better scalability and simplified management when compared to traditional storage solutions.

Big Data Analytics

Big data analytics describes the process of uncovering trends, patterns and correlations in large amounts of raw data to help make data-informed decisions. From these valuable insights, applications can be enhanced, operations streamlined and efficiencies improved. Big data uses familiar statistical analysis techniques and applies them to more extensive datasets and algorithms. The analysis from these large datasets demands high-speed processing and efficient data retrieval. With the combination of OSS with purpose-built hardware, SEF delivers rapid data ingestion, accelerated data analysis and fast query response times. This facilitates timely decision-making, improves data-driven insights and enhances overall business productivity.

Content Delivery Networks

Content delivery networks (CDNs) is a group of geographically distributed and interconnected servers responsible for communicating with users in their geographic vicinity. Their main function is to reduce round trip time by bringing the content closer to the website's visitor. CDN PoPs (Points of Presence) are strategically located data centers that contain numerous caching servers and require low latency storage to distribute content efficiently across geographically dispersed networks. SEF provides the necessary performance and scalability to handle CDN infrastructure demands, enabling fast content delivery, reduced network congestion and a seamless user experience especially for media-rich applications.

Storage-as-a-Service

Storage-as-a-Service (STaaS) enables organizations to manage storage capacity and workloads without the overhead costs associated with upfront capital for storage hardware, software or human resources. STaaS is a more cost-efficient way of using public cloud storage resources for data storing store versus building a private storage infrastructure. This service requires workload isolation to avoid noisy neighbors, data interference and the intermingling of data within a flash super block.

SEF features two forms of workload isolation for STaaS requirements. The first is flash die level isolation using Virtual Devices that assigns sets of flash dies to specific workloads, effectively isolating them from other workloads and associated interferences (Figure 1). This approach provides high tenant and workload isolation.

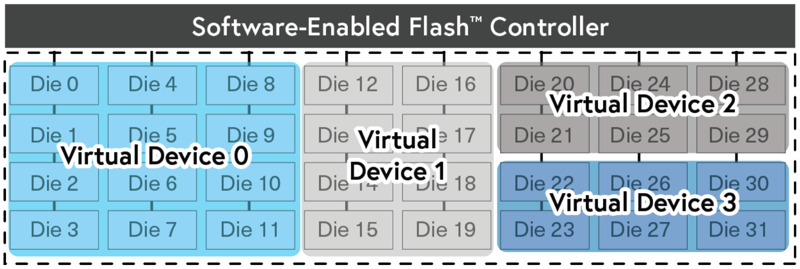

Figure 1: SEF Virtual Devices block diagram
(Used with permission from KIOXIA America, Inc.)

The second approach provides a finer grain of workload isolation through Quality of Service Domains that subdivide Virtual Devices, imposing a secondary level of isolation between different workloads. Workloads in different Quality of Service Domains may utilize the same flash dies though their data is never intermingled within a flash super block (Figure 2). Many Quality of Service Domains can occur within a single Virtual Device simultaneously, allowing for many more simultaneous workloads than flash dies in the system.

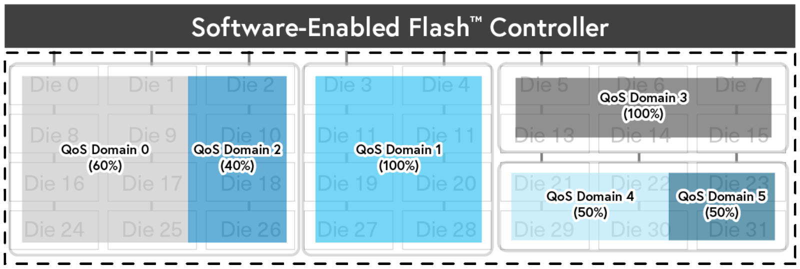

Figure 2: SEF Quality of Service Domains block diagram
(Used with permission from KIOXIA America, Inc.)

Summary

Software-Enabled Flash technology, utilizing OSS and purpose-built hardware integration, represents a powerful approach to meet the evolving storage needs of modern applications. Organizations can achieve higher storage performance, improved data management and enhanced reliability via a fully customizable API. As the demand for faster and more efficient storage solutions continues to grow, SEF is poised to play a pivotal role in driving innovation and enabling new flash capabilities across a wide variety of industries and applications.

The Software-Enabled Flash Project, managed under the Linux Foundation®, provides a vendor neutral collaboration environment with a repository of documentation, source code and discussions related to this new technology. For more information, including the complete SEF technology stack and its benefits, visit https://softwareenabledflash.org.

FOOTNOTES:

[1] Wear leveling is a method for distributing program and erase (P/E) cycles uniformly throughout all of the memory blocks within an SSD. This prevents continuous P/E cycles to the same memory block resulting in greater extended life to the overall NAND flash memory.

[2] Power Loss Protection (PLP) records data in buffer memory to NAND flash memory, utilizing backup power of the solid capacitor in case of a sudden shut down or power outage.