arm

White Paper

Smarter Data Storage

A Guide to Computational Storage on Arm

September 2019

Trusted Partner for Billions of Storage Devices

Smarter Data Storage

A Guide to Computational Storage on Arm

September 2019

White Paper

Contents

1. About This Document

This document provides an overview of computational storage and discusses the benefits of adding computing capabilities to traditional storage devices and systems. It focuses on the Arm-based solutions being developed by our partners and offers practical advice on developing smart storage devices using Arm Cortex processor technology.

This document uses the following terms and abbreviations:

Abbreviation	Meaning
DRAM	Dynamic random-access memory
FPGAs	Field-programmable gate arrays
NAND	Non-volatile type of Flash memory
NN	Neural network
NVMe	Non-volatile memory express
NVMe-oF	NVMe over Fabric
OCI	Open Container Initiative
PCIe	Peripheral component interconnect express
SAS	Serial Attached SCSI (small computer system interface)
SSD	Solid-state drive
TCP/IP	Transmission Control Protocol/Internet Protocol

2. An Introduction to Computational Storage

We have access to more information than ever before. The volume of data we generate is expected to grow by an astonishing 27% a year[1] and that means organizations, from the world's largest to the smallest, will need to find new ways to manage it.

However, business leaders want access to something more than the facts and figures that data itself offers. To be successful in this increasingly digital world, they will need timely analysis and insight that they can convert into added value and services. Therefore, the infrastructure and technology that underpins their operations must be capable of delivering and storing data and analytics in a fast and efficient way.

Computational storage adds computing capabilities to traditional storage devices and systems, which reduces the movement of unencrypted data by enabling it to be processed directly on the drive, with security in place.

[1] The volume of data created, captured or replicated is expected to increase from 33 zettabytes in 2018 to 175 zettabytes in 2025. Source: Data Age 2025. The Digitization of the World From Edge to Core. An IDC White Paper – #US44413318, sponsored by Seagate. https://www.seagate.com/files/www-content/our-story/trends/files/idc-seagate-dataage-whitepaper.pdf

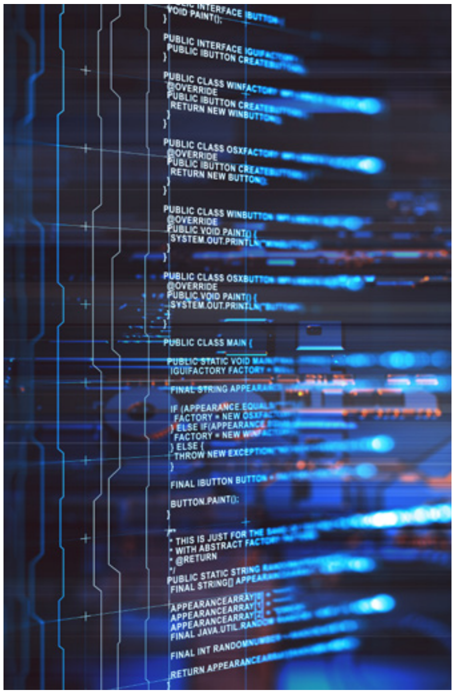

Today, the vast quantities of information we are storing on solid-state drives (SSDs) have to be moved from the device to a server or processor when we want to derive meaning from it. The movement of data introduces a few challenges:

- **Adds latency**
  There are multiple interfaces and protocols that can cause delays. Data is moved from the NAND to the SSD DRAM, packaged in NVMe protocols and then transported over PCIe protocols, or fabric protocols, before it is moved across the system. It then passes through the system DRAM and eventually reaches the host CPU, where it must be unpackaged before the data is available to process. In addition, the size of the stored data is larger than the system DRAM, so this must be repeated many times. All of that takes time and adds latency.

- **Consumes bandwidth and power**
  Transferring data requires energy and generates heat. As more and more data is moved, bottlenecks in the system impede the data transfer and increase demand for power while causing delays.

  In addition, data cannot be processed until it has been transferred from the storage area where it resides. This can stall the main processing device until the data is available – with this gap in use, servers waste a lot of time and power idly waiting for the data to process.

  This is expensive (especially if the storage is remote from the processing) and incurs additional operational expenditure (OpEx).

- **Increases security and privacy concerns**
  Unencrypted data is being moved to and from a storage device, which means there is a risk of interference by non-trusted parties.

  Introducing compute capabilities to traditional storage devices and reducing the movement of data means that companies can process trusted information more efficiently and cost effectively, making the most of big data insights through machine learning.

There are many ways to describe intelligent storage solutions. In addition to 'computational storage', architects and developers may also use the following terms interchangeably:

- In-situ processing
- In-storage compute
- On-device processing
- Edge computing or device edge computing

3. How Does Computational Storage Work?

Let's start with a practical example: if we need to search a database, instead of moving a large amount of information from storage to the server to search through the data, the search can be performed on the device itself and only the results returned. When you have several computational storage devices in the system, they can process data in parallel. You can imagine that when performing these functions repeatedly, across large volumes of data, reducing movement and speeding up processing can multiply into significant gains.

This limits data movement, as indicated below.

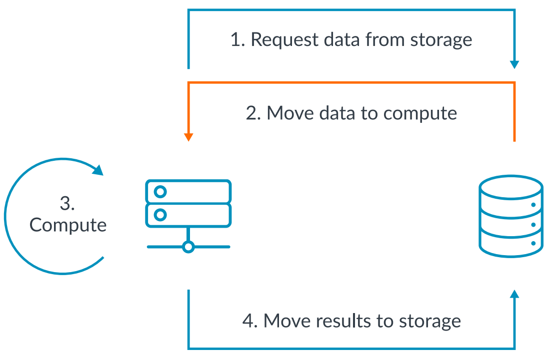

Traditional storage solution

- A request is made to move data from storage
- All the data is moved to compute
- Processing takes place in the server
- Results are moved to storage

In most cases, due to the size mismatch of the system memory and storage (GB to TB), this process must be repeated hundreds of times.

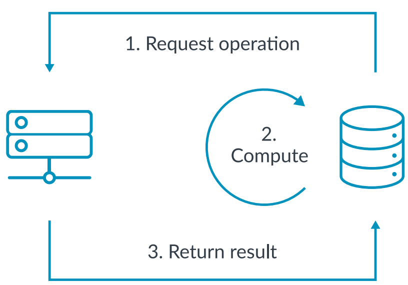

Computational storage

- An operation is requested
- Processing takes place on the storage device(s)
- Just the desired results are returned to the server

According to the Storage Networking Industry Association (SNIA):

"Computational storage solutions typically target applications where the demand to process ever-growing storage workloads is outpacing traditional compute server architectures[2]."

[2] Source: Storage Networking Industry Association. https://www.snia.org/computational

4. The Benefits of Computational Storage

Computational storage will help us realize the potential of emerging technology, across a variety of applications – let's look at a handful that it will be most impacted by:

The Internet of Things (IoT)

Processing data on the drive will help to reduce the cost and complexity of managing the extraordinary amount of data that will be generated by the next trillion connected devices. The data that the end devices create can be processed locally, and only critical data or changes or alerts need to be signaled.

Artificial Intelligence (AI)

Even tasks that rely on large or compressed data sets, for example, image recognition, can be carried out in near real-time rather than being delayed until information has been transferred. In emergency situations this could have life-changing implications.

Machine Learning (ML)

ML can be performed directly on data in the storage drive, increasing the efficiency of searches and returning results to the server, rather than moving terabytes of raw data. ML inference, running a model on data stored on the drive, can be completed quickly and efficiently with standard processing, such as Arm Neon, or with specialized neural network (NN) accelerators or FPGAs that can be built into the drive.

Edge computing

To maximize the benefits of emerging technologies, such as 5G, AI and the IoT, data will need to be processed at the edge of the network or device edge[3]. This aligns with the goal of bringing intelligence to storage devices. The low power requirements mean the devices can sit at a distance from the server without needing to transmit large amounts of data for processing.

Arm's partner, NGD Systems, is exploring the use of computational storage to help airlines analyze flight data more efficiently and effectively. Today, airlines generate multiple terabytes (TB) of data per hour and offloading and analyzing that data can take hours, which is time operators cannot spare. With computational storage, the flight analytics can be provided to the right people at the right time, which will improve safety for the 1.2 million people in the air at any one time.

Scott Shadley from NGD Systems explains in more detail: "As a society, we are generating a huge amount of information and having the ability to sort, search and manage that information is beyond the scope of our current infrastructure. Moving compute closer to data is invaluable and being able to accelerate data analytics creates significant benefits, too."

[3] Source: State of the Edge. Data at the Edge. Managing and Activating Information in a Distributed World. State of the Edge Report, in partnership with Seagate. https://www.stateoftheedge.com/

4.1 Benefits of Computational Storage for Architects, Developers and Manufacturers

Computational storage offers many significant benefits, and for SSD manufacturers, here are a few of the key advantages:

1. **Faster response times**
   Moving intelligence to where it is needed allows results to be delivered in near real-time.

2. **Lowest possible latency**
   There are no additional protocols. Data is transferred from Flash to in-drive processing.

3. **Minimum bandwidth and power use**
   Data remains on the drive with only the results delivered to the server. This allows the device to sit at the edge of the network, as needed.

4. **Data-centric processing**
   Workloads that have time-critical compute are deployed to the drive.

5. **Improved security and privacy**
   Unencrypted data does not leave the drive for processing.

6. **Scalability**
   Compute resources are expanded as more devices are added to the system.

7. **Flexibility**
   Compute can be added to an SSD without affecting its current operation. The additional capability that has been built-in can be utilized in the future.

4.2 Benefits of Computational Storage for Organizations

According to an IDC report, in just over five years' time almost 30% of the world's data will need real-time processing. That means that:

"Enterprises looking to provide superior customer experience and grow share must have data infrastructures that can meet this growth in real-time data[4]."

Currently, vast amounts of data are being gathered and stored on hard disk drives, SSDs or in a data lake. However, it will have to be moved for an organization to benefit from it. With computational storage, the time that would be spent moving the data is instead dedicated to processing it.

4 Source: Data Age 2025. The Digitization of the World From Edge to Core. An IDC White Paper – #US44413318, sponsored by Seagate. https://www.seagate.com/files/www-content/our-story/trends/files/idc-seagate-dataage-whitepaper.pdf

Other benefits for organizations include:

1. **Lower costs**
   A reduced workload on the servers in the data center calls for less power, and therefore, lower operating and cooling costs.

2. **Reduced physical footprint**
   Fewer servers are needed to store and manage the data, which results in less infrastructure and lower costs.

3. **Increased efficiency**
   People will have almost instant access to information. Completion time following a request will be faster.

4. **Improves data security in the data center**
   Unencrypted data does not leave the drive.

5. **Greater control**
   Organizations have additional capability without being encumbered by it. They can take advantage of computational storage using ordinary code, and an open-source operating system and container software that can be installed on a desktop or laptop.

5. Computational Storage Architecture

SSDs already have a significant level of computing capability built-in and computational storage capitalizes on this. Moving from a traditional storage solution to a more intelligent device requires minimal changes to the system-on-chip (SoC) controller.

5.1 The Implications: Power and Cost Requirements

- SSDs already have a large amount of DRAM, typically >1GB per terabyte (TB) of NAND. In terms of the overall power and cost of the components in a SSD drive, the requirement of the processing enables computational storage to be deployed without significantly increasing power or cost.

- The amount of compute the drive has available for computational storage tasks can vary. At the low-end, some may rely on the processing that is already available – either working on background tasks or when the drive is less loaded, for example, overnight. For more computational storage performance, additional processing can be added to the drive, from Arm high-efficiency CPUs through to multiple clusters of Arm premium performance CPUs. In addition, specialized processing can also be added to accelerate general AI/ML with dedicated NN processors or by integrating dedicated hardware or FPGAs.

5.2 Alignment with Industry Standards

- The SNIA's computational storage Technical Working Group (TWG), of which Arm is a founding member, has been established to set industry standards and encourage interoperability between devices and systems. More than 40 companies are represented and working together to define relevant approaches for different types of computational storage. In most cases, the server system must be able to deploy workloads to the drive and then invoke these workloads and receive results, although dedicated standalone capabilities also have applications. Methods to provide computational storage drive services and capabilities are being developed to ensure the drives are standardized, and that drives from multiple vendors can be adopted and deployed.

  The TWG is aligned on the definitions of fixed-purpose and programmable computational storage services. As the names suggest, a fixed-purpose computational storage service provides a fixed function that can be configured and used. A programmable computational storage service can be programmed to provide one or more computational storage services.

  Accessing the fixed and programmable computational storage services may involve extending the interface (NVMe, SAS, etc.) protocol to add new commands. This enables the commands to be intercepted when they are received by the drive and specific actions taken. Adding new protocol commands requires industry alignment to ensure that devices from various manufacturers can interoperate.

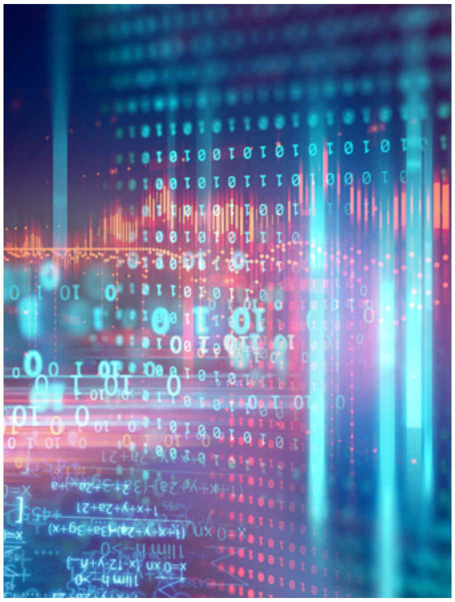

5.3 Utilizing Computational Storage

- The first step to utilize computational storage is to discover computational storage–compliant devices and services.
  - Each device must advertise available services so the host can utilize the capabilities of the drive.
  - Next, the services need to be configured, as appropriate. This includes configuring parameters for proper operation of the service.
  - Once configured, computational storage service drivers may directly use the service, or indirectly use the service, by performing standard storage operations against an associated storage interface (NVMe, SAS, etc.). In the case of a fixed-purpose computational storage service, commands could invoke the computational storage function to perform actions on the data. In the case of a programmable computational storage service, the commands may enable a workload to be downloaded or started.

- NVMe is likely to be one of the first storage interfaces enhanced to support computational storage. NVMe can be carried over PCIe, enabling direct server offload of tasks to the directly attached drive. In another scenario, NVMe over Fabric (NVMe-oF) allows these drives to be remote from the servers and, for example, connected over standard ethernet networking over TCP/IP. As computational storage evolves, a TCP/IP port may be added next to the PCIe/NVMe port on a drive or replace the PCIe/NVMe port altogether.

- One approach to enable programmable computational storage services is to have the open-source operating system (OS), Linux, running on the drive. This makes the drive appear as a standard server to the rest of the infrastructure, so it can be managed in exactly the same way as any other server. For example, computational storage workloads can be developed using the standard Linux open-source environment and then containerized, downloaded and managed on the drive using common open-source applications, such as Docker or Kubernetes. This transforms the drive into a server that can be deployed and managed inside a data center or remotely as a cost-effective edge server.

  To utilize programmable computational storage services running Linux, a user will write ordinary application C code on their desktop or laptop to perform the desired function on the data (search, manipulate, etc.). This C code can then be compiled into an OCI compliant container image using a tool like Docker, which is installed on the user's computer. When it is compiled, the container image can be deployed to the computational storage drive using an orchestrator and executed in an OCI compliant container runtime such as runc. Again, container software is required to be running on the Arm processor on the drive, along with Linux. At that point, the executable functions exactly as it was written, making requests for data from the drive as if it were the host computer, and performing the computations.

- Front-end and back-end compute is typically handled by an Arm Cortex-R processor (Cortex-R8 or Cortex-R5). The Cortex-R series provides real-time capabilities to ensure timely handling of latency-critical control code. They typically have some memory that is tightly coupled to the processor, which is ideal for fast execution and manipulation of control structures. To enable computational storage, a Cortex-A processor needs to be added to the drive, which involves making minimal changes.

In storage compute:

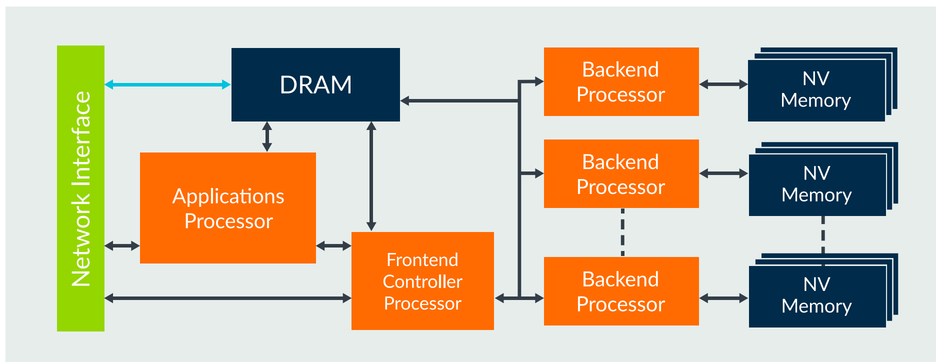

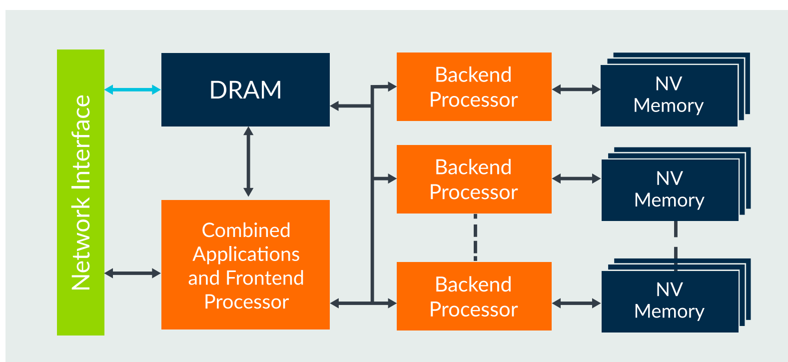

When storage controllers are developed, adding additional compute for computational storage is very cost-effective, especially in the small process geometries now being adopted, and trivial compared to the cost of the large memories in the drive. Open-source software, including Linux and container software, enable the compute to be utilized.

By running Linux on the drive, isolated from the secure standard drive real-time firmware that remains unchanged, the benefits of the huge open-source community are available to develop the computational storage workloads. Arm has partnered with Linaro to ensure all leading Linux distributions and applications have not only been ported, but also optimized for running on Arm processors. Linaro, for example, was founded in 2010 and has been working with many partners to create a thriving Linux ecosystem on Arm.

NGD Systems used an open-source application called OpenALPR, which enables authorities to search and match vehicle license plates, to demonstrate the range of tasks that can be parallelized on computational storage devices. OpenALPR was loaded on to a drive via Docker. The use of containers enabled the company to run non-matched applications in real-time.

6. At a Glance: Adding Computational Storage to Your Device

In summary, there are a number of ways to implement computational storage, including embedding an Arm Cortex-A processor in an SSD controller SoC. This approach has already been adopted by multiple Arm partners.

STEP 1: Applications processor

Cortex-A processors offer a high-performance and power efficient solution and are designed for complex computing tasks. They can be integrated into an SSD controller SoC on the main data path or off to the side. Communication pathways must be created for the core to make data requests from the SSD control path.

STEP 2: Linux-based operating system

A stripped-down version of the Linux OS can be built to run on the Cortex-A processor. Linux has been optimized for this portfolio of Arm processors. Through our ongoing support for and investment in Linaro, one of the top five contributors to the Linux kernel, we are helping developers access cutting-edge code that is freely available and ready-to-run on Arm. Industry representatives and the open-source community work together to develop projects and tools that will reduce fragmentation and redundancy, as well as providing a common foundation.

STEP 3: Containerized workloads

Open-source container software, including Docker and Kubernetes, can be used to containerize workloads. Packages can be installed on the Linux installation and compute functions can be built on a standard PC. The containers for these compute functions can be downloaded to the Linux installation and run on the Linux container.

Docker and Kubernetes are part of our diverse ecosystem of thousands of partners who use their expertise to help manufacturers reduce the cost and risk associated with Arm-based development and minimize time to market.

7. Innovation in Storage Solutions

California-based NGD Systems has been working with Microsoft Research on a proof-of-concept to demonstrate the benefits of computational storage. The preliminary results of a trial that performed an image search suggest that adding intelligence to an SSD could deliver significant throughput improvement and increase the number of image queries that can be handled per second by more than four times, compared to processing the data on a host device.

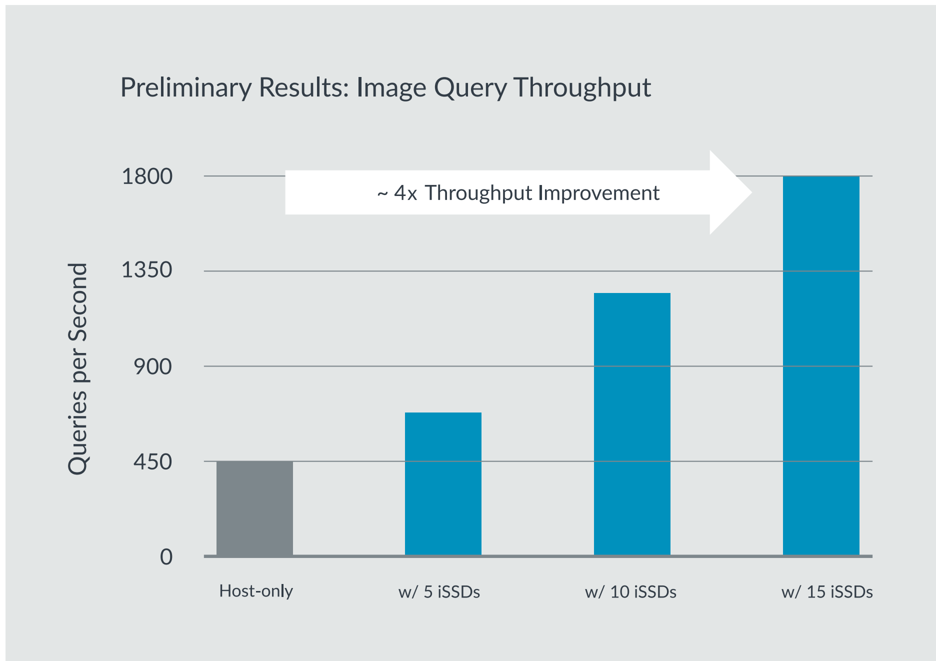

Source: NGD Systems / Microsoft Research

The team used NGD Systems' Newport computational storage platform to eliminate the need to move data before processing it. Its in-situ processing capability is powered by Arm Cortex-A53 processors running a 64-bit operating system, a version of Linux, which is embedded in the SSD controller.

VP of Marketing, Scott Shadley explains: "We have been working with Microsoft Research to demonstrate that computational storage on Arm makes analyzing large datasets more efficient and cost-effective. There are also more compute resources available when the quantity of processors within each drive, rather than just the capability within the server, is taken into account."

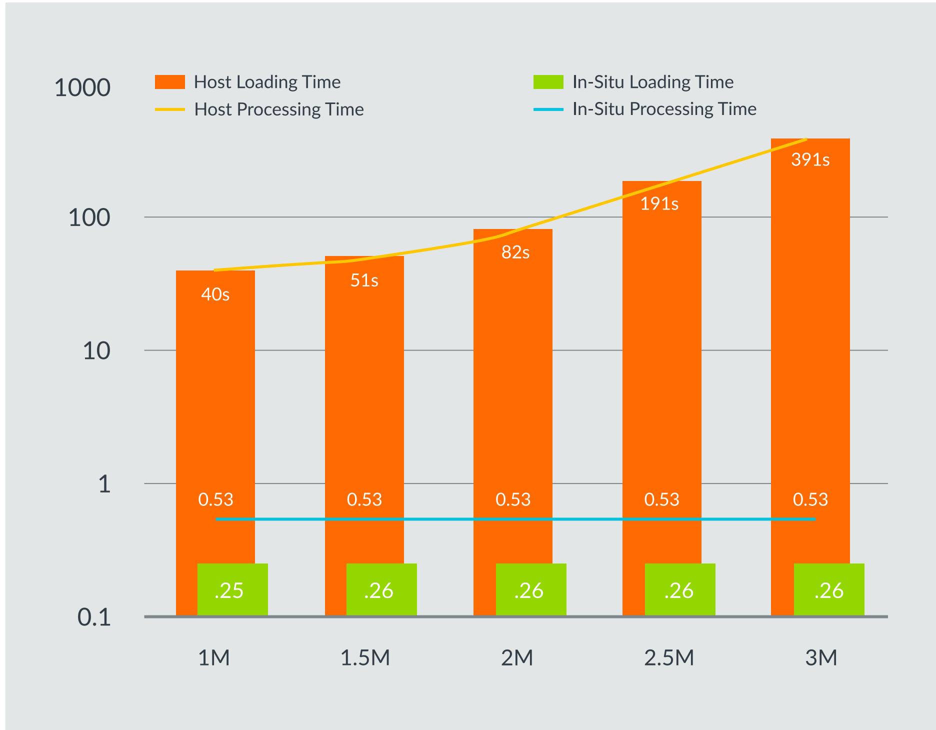

Similar studies have also pointed to the potential benefits. NGD Systems recorded a significant improvement in queries per second using the Facebook Artificial Intelligence Similarity Search (FAISS). During the test, the completion times for a traditional server were over six minutes for the selected dataset. When computational storage was utilized, completion times did not exceed half a second, and quality of service was maintained.

8. Securing Your Storage Device

As the volume and value of data grows, so do attempts to compromise the security of the devices that manage and store it.

Security should be designed-in to devices from the ground up, with the ability to update security features as threats evolve. The Platform Security Architecture (PSA) provides a framework to ensure that all devices have a foundational level of security and achieve a set of fundamental security goals. It reduces the cost and complexity of developing a secure product through four stages:

- Analyze: assess the value of your assets and the threats to them.
- Architect: hardware and firmware specifications.
- Implement: PSA compliant open-source reference code and APIs.
- Certify: an independent certification scheme.

We have a range of security IP that support our security goals, such as Arm TrustZone: proven technology that is already protecting high-value code and data in billions of SoCs to date. It provides isolation between the Secure and Non-secure worlds, which ensures the standard SSD firmware is isolated and protected from the new workloads. Linux has security features built-in and is a trusted operating system.

9. The Next Steps for Computational Storage

Computational storage enables us to maximize the benefits of big data to organizations and to society. It puts processing power where it is needed and gives us quick and easy access to vital information. We are working with our partners to ensure that, as an industry, we design the right storage solutions and that everyone benefits from their development.

Data centers

"For a hyper-scale storage architect, finding a way to avoid moving data from device to device, even in the same rack, will become a paramount need, as will reducing the power consumed to process this data."

– Scott Shadley, VP of Marketing, NGD Systems

Smart cities

"Computational storage devices will be used to build more power and space efficient systems at the edge. This is vital for the delivery of services over 5G networks and for initiatives like IoT and smart cities."

– Stephen Bates, CTO, Eideticom

Automotive

"We believe that more than 4TB of data will be generated daily by just one autonomous vehicle. Some of this data would be stored in the car and need to be uploaded to the cloud for further analysis for applications such as an AI model or 3D map updates. Pre-processing will be essential in order to decide which data should be uploaded, as transmitting all of it is not an option given the network strain. To that end, computational storage can be a great tool to manage this pre- processing in the most effective and productive way."

– Noam Mizrahi, Fellow, VP Architecture and Solutions at Marvell CTO Office, Marvell

For More Information

Visit our website to learn more about computational storage or to start developing your own intelligent storage devices. Alternatively, contact Neil Werdmuller, Director of Storage Solutions at Arm on neil.werdmuller@arm.com for more information.

arm