ActualTech Media

FORMULUS BLACK

Increase Application Performance and Reduce Costs with Memory-Based Storage

Tom Fenton

CONTENTS

IN THIS PAPER

This tech brief details how Formulus Black's ground-breaking software stack enables any workload to run in memory, without modification, and helps companies find new ways to speed up storage access and harness the full power of today's processors.

INTRODUCTION

Artificial intelligence (AI), machine learning (ML), transactional and analytical databases, high-frequency trading (HFT), fraud detection, and other business-critical applications need to access and process large amounts of data as quickly as possible. Unfortunately, today's traditional computer architecture and storage methods strangle such applications because they can't process the information quickly enough. This can have detrimental consequences.

> Timely access to processed data has a significant impact on how customers and end users perceive an organization.

In brokerage houses that do HFT, for example, a nanosecond of latency can mean the difference between making a profitable trade or not. Similarly, fraud detection applications rely heavily on data processing speed, as they must be able to comb through mountains of data to detect anomalies that might indicate the misuse of a credit card or a fraudulent credit application. Yet another study shows how data analysis and detection need to be extremely performant as bot-driven ad-fraud can occur in less than a millisecond.

Timely access to processed data has a significant impact on how customers and end users perceive an organization. About 10 years ago, we expected inventory stock reports to be generated weekly; five years ago, the same reports were generated nightly. Now, users expect to see the inventory in stock in real-time, and executives want to be able to run ad hoc queries against internal databases and data warehouses instantaneously.

Application and hardware providers have come up with imaginative ways to try and satisfy our need for speed, but new bottlenecks appear with every solution. Currently, one of the biggest problems is that data in secondary storage can't be processed quickly enough. Even non-volatile memory express (NVMe) drives, which connect directly to the PCIe bus and are an order of magnitude quicker than solid-state drive (SSD) and hard-disk drive (HDD) devices, still throttle mission-critical applications.

EARLY ATTEMPTS AT IMPROVING PERFORMANCE FOR DATA-INTENSIVE WORKLOADS

Application and hardware vendors have come up with various ways to deal with performance issues, ranging from using specialized hardware to creating new programs and methodologies. The most exciting of these solutions is the use of memory (either dynamic random-access memory [DRAM] or persistent memory [PMEM]) as the storage media for hyperconverged infrastructure, also referred to as Memory Converged Infrastructure (MCI). MCI, which relies on DIMM-based memory hardware and offers memory bus speeds, bypasses the traditional (and much slower) peripheral storage travel path to reach the CPU. As an example, DRAM memory has a latency of about 51 nanoseconds for reads and writes, PCIe NAND-Flash has a latency of 47 microseconds for reads and 15 microseconds for writes, while traditional enterprise disk storage has a latency of 9 milliseconds for reads and writes. Relative to non-memory channel storage technologies, throughput and bandwidth of even the highest-performing SSD is several orders of magnitude lower.

> Relative to non-memory channel storage technologies, throughput and bandwidth of even the highest-performing SSD is several orders of magnitude lower.

Formulus Black wanted to see just what difference changing from NVMe to memory-based storage would have on database performance. In one test the company found that a database query that took 200 seconds on NVMe took only 2 seconds using memory-based storage; a 100x improvement.

An early example of MCI usage—and a prime example of the hazards involved with specialized hardware to accomplish MCI—is Fusion-io, an innovative startup company that designed and manufactured proprietary products that used flash memory technology to increase the performance of servers. For a while, it was the darling of the data center; in the early 2010s, its ioDrive product was considered to be one of the fastest (albeit most expensive) storage devices on the market.

Recently, other more established companies, such as storage goliath NetApp, have built MCI products, such as MAX Data, that can increase application performance. MAX Data uses a server's DRAM; a non-volatile dual in-line memory module (NVDIMM); or storage-class memory (SCM), such as Intel Optane DCPMM, to create a high-performance auto-tiering filesystem that can survive a system reboot. However, MAX Data has some limitations:

- It has a highly restrictive hardware compatibility list (HCL) for the servers on which it can array
- It requires a NetApp array
- It works only with NetApp ONTAP
- It requires a separate license
- Uses MAXFS—a NetApp proprietary file system

Although some companies sought to improve application performance through proprietary hardware, other vendors found innovative ways to improve performance by having applications directly use the memory on a system. Perhaps the best example of this is SAP HANA, an in-memory relational database system proven to be very performant. However, it requires rearchitecting of applications and migration of existing data. Like the hardware approaches noted earlier, it's expensive and also suffers from vendor lock-in.

BRINGING STORAGE TO THE MEMORY CHANNEL

Over the past decade, multiple factors have come into play that have radically changed the data center. For example, CPUs have seen an explosion in the number of cores, DRAM prices have dramatically dropped, and SCM has become available.

Furthermore, 16-core CPUs have become the most popular servers in the data center, 32-core CPUs are no longer considered exotic, and 64-core CPUs are available. The proliferation of cores within servers means that data can now be processed quicker than ever. The problem is that many companies that have shelled out for these processors find those cores sitting idle, as they can't get data from storage fast enough. And no organization can afford to waste money like that.

As the number of cores per CPU increased, the price of RAM dropped considerably, as well. By capitalizing on this cost reduction, companies can now afford to equip their servers with more memory than ever before. Additionally, in 2019 Intel released SCM devices based on XPoint memory media that, when compared to DRAM, are less expensive, support more capacity, and let data survive a reboot.

What was needed was a way to take advantage of the increased cores and memory on commodity x86 servers to improve application performance—without having to rely on propriety hardware or software, and without having to rearchitect applications. One of the companies leading the way in this transformation is Formulus Black.

FORSA MEMORY-BASED BLOCK-LEVEL DEVICE

Formulus Black's Forsa overcomes many of the limitations of other MCI solutions by utilizing DRAM or SCM to unleash the power of modern servers in the data center.

Forsa takes a portion of the DRAM or SCM on a system and creates Portable Operating System Interface (POSIX)-compliant block devices—called Logical Extensions of Memory (LEM)—that's used like any other block device on which you'd mount a filesystem such as HDFS, XFS, EXT3, and so forth.

> Formulus Black's Forsa overcomes many of the limitations of other MCI solutions by utilizing DRAM or SCM to unleash the power of modern servers in the data center.

Forsa's LEM is extremely robust and flexible, and can be resized on the fly if needed. Memory-based storage devices behave just like any other block-level device (only with much less latency): They can be cloned, snapshotted, and backed up using existing tools and methodologies. Forsa is currently supported on various Linux distributions, and doesn't require any propriety hardware.

You can migrate your applications or data to a memory-based block-level device and enjoy the benefits of having

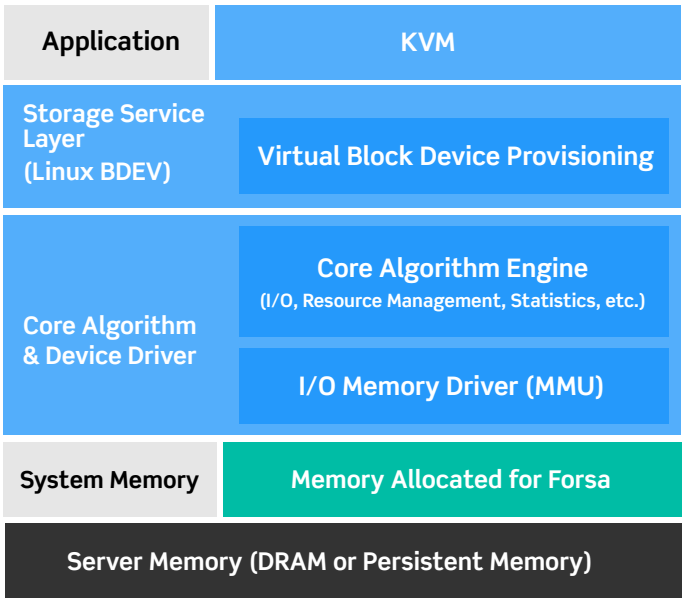

your data processed at far greater speeds than if your data was kept in traditional storage.

By using Forsa, you don't need to rearchitect your applications, except to put the data or program on the memory-based block storage device. Even though the price of DRAM and PMEM has dropped dramatically, it's still a relatively scarce resource that Forsa maximizes using Formulus Bit Markers to expand memory capacity by eliminating duplicate copies of data.

One of the limitations of some of the early solutions that attempted to create a memory-based storage device concerned data protection. Forsa overcame this issue with BLINK, which takes an application-consistent memory state capture and copies it to the SSD, allowing for recovery should the need arise. BLINK even interfaces with UPS battery backup software, allowing you to back up data stored in DRAM to persistent storage, such as an SSD, when a power failure occurs.

Closely related to data protection is data integrity, which Forsa supports with bad block replacement (BBR) to ensure that the data stored on the LEM is error-free.

Because Forsa works with commodity (rather than proprietary) hardware, you can choose the right hardware for your data center. Some companies may choose DRAM because it performs better than SCM, while others may prefer SCM as it supports greater capacity, costs less, and will survive reboots. The point is that Forsa puts the choice in *your* hands—not a specific vendor's.

USE CASES FOR MEMORY-BASED STORAGE DEVICES

MCI solutions, such as Forsa, have some very interesting use cases, and companies both big and small have exploited them to maximize their data center's value.

Databases seem to be the first application IT managers think of when considering an MCI solution, and many have seen dramatic improvements in database performance after moving them to Forsa. What surprised one IT manager was that they were able to reduce costs; specifically, they reduced their database licensing cost by moving their database from a 16-CPU license to a 4-CPU license—and still had better performance.

> Edge computing is becoming increasingly essential to enterprises, and it's often more important to be able to process data at the edge than in the data center.

Enterprises are embracing the software defined data center and new applications often run on virtual machines (VMs) or in containers. Forsa's LEM supports both KVM-based VMs and Docker applications. These technologies benefit from the same types of performance gains as non-virtualized technologies and Formulus Black was able to document impressive TPS/TPM figures on small (8 vCPUs/ 4GB RAM) and medium (32 vCPUs/ 16GB RAM) sized KVM VMs running under Ubuntu.

Not all Forsa uses cases reside in the data center, as more often we're seeing servers running outside of the data center collecting information from a variety of sources: video cameras, radar, high-definition microphones, and a multitude of sensors. This creates a torrent of data that information must be extracted from locally or passed on to another application or end user who wants to analyze or build reports off that data. Forsa can dramatically improve the efficiency and performance of these edge servers.

The consequences can be much worse with other edge-computing scenarios, such as autonomous vehicles or edge devices used in healthcare settings. Poor latency here can have more serious and impactful affects.

Forsa's benefits in these situations become more obvious. Formulus Black has documented actual use cases in which Forsa has been able to solve difficult problems or increase the effectiveness of IT resources.

REAL CUSTOMERS, REAL ADVANTAGES

Packet is a bare-metal cloud provider with more than 20 locations around the world that's used by developers, SaaS companies, and others who need to deploy physical infrastructure on a global scale. With support from Intel's Non-Volatile Memory Solutions Group (NSG), Packet has teamed up with Formulus Black to use Forsa to create ultra-low latency, high-bandwidth solutions using Intel Optane DC PMM. This makes it easier for developers and IT leaders to adopt persistent memory solutions and supercharge the performance of databases, ERP systems, and custom-built data-intensive applications without having to use exotic software or expensive proprietary solutions.

> The appliance model appeals to users who want a single SKU and support model, as well as to those who have existing hardware or whose preferred vendor doesn't yet offer Forsa.

Because Forsa runs on commodity x86 hardware and doesn't require rearchitecting of applications, custom solutions that have been developed in-house, and which would be difficult and costly to rewrite or port to other systems or the cloud, have seen substantial performance gains by running on Forsa. A prime example of this is etherFAX.

Many people think that faxed documents have been completely replaced by e-mail, but the reality is that heavily regulated institutions such as healthcare, financial firms, and the education sector, often need to fax information to stay legally compliant.

etherFAX fills this niche and transmits millions of pages a day on behalf of these institutions. By using Forsa to provision high-performance memory-based storage and VMs for its document-rendering application, etherFAX effectively reduced the time needed to render documents by more than 70%—without having to rewrite code or rearchitect its application.

Formulus Black sells Forsa as a software package that can be licensed either on a per-socket basis, or, for users who prefer the appliance model, through selected resellers that sell systems with it preinstalled on reference architecture.

The appliance model appeals to users who want a single SKU and support model, as well as to those who have existing hardware or whose preferred vendor doesn't yet offer Forsa.

Additionally, many vendors offer Forsa using DRAM for those who want the most performant block storage device or a PMEM solution, like Intel Optane DC persistent memory for those who want more memory-based storage capacity at a lower cost relative to DRAM, and who want the added peace of mind that comes with using nonvolatile memory. Finally, those who prefer using cloud services can use a provider like Packet to provision Forsa's memory-based storage.

TAKING STORAGE TO THE NEXT LEVEL

It's becoming increasingly clear that storage has developed into the new bottleneck hampering performance for the most demanding workloads. CPUs are fast enough, with enough cores. The problem is that the data can't be served up quickly enough to satisfy the application's needs.

This issue will only continue to get worse, as digital transformation upends previous data center paradigms, forcing companies to find new ways to speed up storage access and harness the full power of today's processors. Formulus Black's Forsa is a ground-breaking software stack that enables any workload to run in memory, without modification. Unlike legacy storage solutions, Forsa enables memory to be provisioned and managed as a high-performance, low-latency storage media for AI, ML, transactional and analytical databases, HFT, and other demanding operations.

If that's what you need, your solution has arrived. Visit formulusblack.com/see-it-work/ to request a free trial.