AI and High-Velocity Analytics Workloads Need a New File System

Storage Switzerland, LLC

WEKA.IO

As Artificial Intelligence, Machine Learning and Deep Learning workloads go mainstream, organizations are struggling with how to develop a storage infrastructure to best meet the unique challenges of these workloads. AI/ML workloads typically include hundreds, if not thousands of servers accessing hundreds of terabytes of data often represented by billions of small files. Since these workloads are trying to simulate or exceed human thinking they need to traverse the data set very quickly.

CHAPTER 1: Why Legacy File Systems Can't Keep up With AI and High-Velocity Analytics

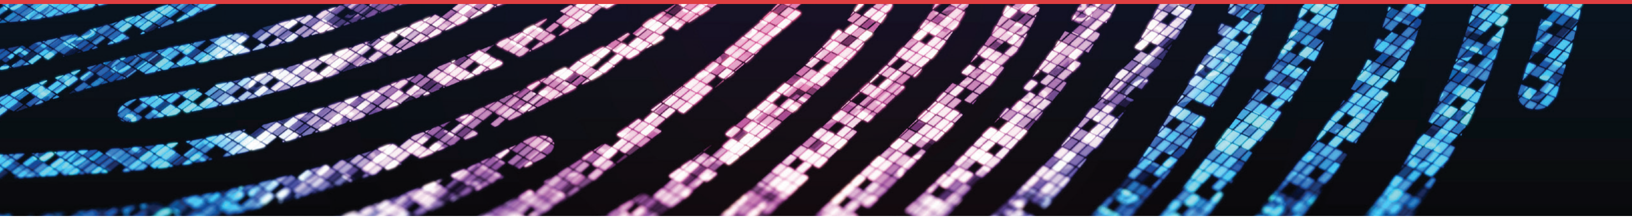

Artificial Intelligence (AI), Machine Learning (ML) and Deep Learning (DL) workloads often start as skunk works projects within an organization. After the proof of concept and testing they move into production, which means storage performance and capacity demands for the solution increase rapidly. Since most of AI/ML architectures are built from dozens if not hundreds of servers simultaneously accessing the same unstructured data set, a shared file system is the most obvious choice. IT planners often try to use legacy NAS systems to support the AI/ML workloads but quickly find legacy NAS lacking.

Why Traditional NAS Falls Short

A NAS has two critical components, the software that actually provides the services and the hardware that delivers those services. The traditional NAS hardware is a single or dual controller system that routes IO requests to several shelves of hard disk or flash media. All data flows through these controllers and in AI/ML workloads they quickly become the bottleneck.

The data that AI/ML workloads use are typically made up of very large quantities of small files. It is not uncommon for file counts to reach into the high millions or even billions. The software of the typical NAS system can also bottleneck in these high file count situations. The NAS file system is bogged down by the metadata required to track these files.

The typical legacy NAS performance work around is to leverage flash media. A hybrid or all-flash NAS provides a performance improvement over hard disk drives but eventually the flash media is also inhibited by the hardware IO bottleneck and the software's inefficiency in managing metadata.

Why Scale-Out NAS Falls Short

In an attempt to address storage performance concerns, many IT professionals attempt to leverage Scale-Out NAS systems which connect dozens of storage servers, called nodes, into a single storage construct. The problem is that most scale-out NAS systems are not truly parallel; they have a single set of control nodes that manage IO movement. A request for data must first go through the control nodes, which then route the IO to the storage nodes. Once the request is received the storage node then routes the IO traffic back through the control nodes before the data is sent to the requesting user or application. There is no way, in most scale-out storage architectures, to route storage traffic directly to the nodes containing the data, instead all data must route through a set of control nodes. These control nodes create a similar bottleneck as scale-up architectures.

Why Legacy Parallel File Systems Fall Short

Most AI/ML storage architectures eventually end up with a parallel file system. These file systems enable compute servers to communicate directly with the node(s) that have the data those servers need. With a parallel file system performance scales as nodes are added, enabling it to keep pace with a rapidly growing compute infrastructure and data set.

However, the problem with most parallel file systems is they were written over a decade ago. Their performance was designed in an era of single core processors and hard disk drive based media. While some parallel file systems dedicate certain processes to certain cores they are not multi-threaded.

A larger problem is the lack of support of flash media, specifically NVMe media. In the past, using a hard disk driver to interface with a flash drive was acceptable since SAS based flash behaved very similarly to a hard disk drive. NVMe however connects via a different bus (PCIe) and supports much higher queue depths and command counts but the driver software needs updating to take advantage of it. Most parallel file systems run on top of a Linux foundation and don't fully exploit the performance of NVMe. While these systems will see a performance improvement over traditional SAS flash they won't achieve anything close to the per drive potential performance of the hardware.

Why Direct Attached Storage Falls Short

These challenges eventually lead the organization to use direct attached storage for their AI and ML workloads. Direct attached storage eliminates the overhead of the network, but in most cases still uses an inferior NVMe driver. Additionally, a direct attached storage solution inherits all of the challenges common to direct attached solutions which led to the use of shared storage in the first place. Data is rarely in the right place at the right time and IT needs to constantly copy it from one compute server to another. Capacity utilization is also inefficient. Most direct attached AI/ML environments use less than 30% of storage capacity, which means the most premium tier of storage goes largely unused.

Time For A New File System

AI and ML are the definition of modern workloads. It makes sense then that the storage architecture should also be modern. It needs to be parallel in nature but also be optimized for modern multi-core storage servers and have native, built-in support for advanced storage technology like NVMe. The goal of the modern file system should be to extract maximum performance out of the flash media while also delivering scalability, efficiency and ease of use.

In our next chapter Storage Switzerland discusses how to design a file system capable of supporting AI and ML workloads as well as having advanced capabilities, like native cloud integration, that modern data centers require.

CHAPTER 2: Designing a File-System for AI and High-Velocity Analytics

Our previous chapter highlighted the challenges of supporting artificial intelligence (AI), machine learning (ML) and deep learning (DL) workloads with legacy file systems. Control node bottlenecks, inferior (or lack of) non-volatile memory express (NVMe) drivers, and inefficient capacity utilization are among the pain points that come with trying to process the millions or billions of small files that typically comprise AI, ML and DL workloads with legacy network-attached storage (NAS) approaches.

In this installment, we will explore the hallmarks of the modern file system architecture. Notably, it should be optimized to fully capitalize on the performance acceleration offered by NVMe while at the same time optimizing I/O performance. Furthermore, it should offer integration with cloud compute and storage resources for cost efficiency. At the same time, a distributed architecture is critical to optimizing data protection.

Built-in NVMe Support

NVMe is a storage protocol designed to accelerate the transfer of data between host systems and solid-state drive (SSD) storage media, over the server's peripheral component interconnect express (PCIe) bus. NVMe can enable the enterprise to more fully utilize the SSD's maximum performance levels by increasing command counts and queue depth – a key value proposition when it comes to serving AI, ML and DL workloads. To do so, however, the storage infrastructure must be architected correctly; NVMe exposes any storage infrastructure bottlenecks because it is so latency efficient. Legacy NAS architectures were not designed to take advantage of NVMe. For instance, file servers continually request information about metadata before they execute operations, adding significant communication overhead. This precludes the ability to fully exploit potential performance acceleration.

Excellent Performance with Small Files

Another significant change that must occur to the file system architecture in order to support AI, ML and DL is the ability to support rapid inspection of very small files. Legacy file system architectures were designed for workloads such as high-performance computing (HPC) that require rapid processing of large files. AI, ML and DL workloads, on the other hand, require equally fast processing but of millions (or billions) of small files. This creates a situation whereby metadata access requests, which typically account for anywhere from 70% to 90% of data requests being served by a NAS system, become the bottleneck. As a result,

the modern file system must be written to continue to deliver high levels of network bandwidth, but at the same time extreme levels of IO performance, to ensure utilization of central processing unit (CPU) and graphics processing unit (GPU) resources.

Cloud Integration

Storage Switzerland sees the hybrid cloud model as being cost effective, and as a result popular among enterprises, for hosting AI, ML and DL workloads. Organizations may want to temporarily shift workloads to the public cloud for peak processing demands but then later bring those workloads back on-premises for normal conditions. The problem is that legacy NAS architectures were not designed to enable seamless portability of data and workloads between on and off-premises infrastructure resources, or for the ability to run workloads in parallel across these resources – both of which are required for a true hybrid cloud architecture.

Storage Switzerland views the ability to access and pay for compute resources on demand as one of the most effective use cases of off-premises cloud services. When it comes to AI, ML and DL, many enterprises are looking to get started quickly and for as limited an overhead (including upfront investment in infrastructure and ongoing management of infrastructure) as possible. This makes the ability to burst AI, ML, and DL workloads to the cloud on a temporary basis for processing appealing. Other enterprises have invested in some on-premises CPUs and GPUs, but have the need to also run some workloads in parallel in the cloud due to temporary spikes to compute and storage needs, and to then compare these results to analytics jobs conducted on-premises. Finally, the ability to tier data across on and off-premises storage resources can help to control costs; data should be tiered from most expensive and fastest-performing on-premises SSD media, to lower-cost object storage services, depending on how frequently it is accessed by the enterprise.

Cost-effective Data Protection

Finally but far from least importantly, the need to serve millions or more than a billion of very small files increases the risk that storage media will fail or data might be corrupted; these realities coupled with the need to provide demanding performance levels during a failed state creates new data protection and disaster recovery requirements. The scale-up architectures of traditional NAS systems mean that it will take a long time for the system to return to acceptable levels of performance in a failed state, because all processes must flow through the single head node. This under-utilization quickly becomes very expensive, with the introduction of more expensive processors and storage media and networking.

As a result, a distributed approach that enables rebuilds to be spread out across nodes is needed. Also required is an Erasure Coding type of protection scheme that optimizes capacity utilization while increasing resilience. Erasure Coding requires more processing however so the modern file system much also distribute the data protection load across node computing power.

In our next installment, we will explore the challenges of relying solely on benchmarks when evaluating next-generation file architectures.

CHAPTER 3: Debunking AI and High Velocity Analytics Benchmark Results

Benchmarks are necessary when trying to understand the performance characteristics of a particular storage system in a particular environment. The problem is they are susceptible to manipulation by vendors. The Standard Performance Evaluation Corporation (SPEC) reduces some of this manipulation by enforcing standardized testing and results submission. Vendors have to clearly document their test configurations so that unrealistic designs are easily exposed. The problem with benchmarks gets worse as an increasing number of organizations begin selecting storage systems for artificial intelligence (AI) and machine learning (ML) workloads.

The Proof of Concept Challenge

AI and ML workloads are very difficult to set up in a proof of concept, testing environment. Part of the problem is that understanding what the AI/ML project will look like three to five years from now, when it is in full production, is hard to determine. Another part is that gathering the hardware and software needed to test a potential new storage system is very expensive. Finally, there is also the time involved in configuring, and reconfiguring, the test environment as each storage system candidate shows up.

The organization is stuck at a crossroads. There are no AI/ML specific benchmarks but internally testing every system is almost an impossible task. Organizations need to consider a blended strategy where they intelligently dissect benchmark results to develop a (very) short list of storage candidates to test.

Dissecting Benchmark Data

While SPEC is doing an amazing job standardizing test results and providing transparency into the configurations used, organizations still need to be careful when they interpret the results. Storage vendors still use unrealistic hardware configurations in an attempt to achieve a top spot.

Another variable to consider is that many of the vendors submitting results are primarily software companies. They are limited by the configuration they used.

In some cases these configurations are valid, as they are attempting to show that their software is not the limiting factor, and that they can max out the hardware configuration. In other cases the configurations are suspect and should be looked at with some level of skepticism. If a system is able to deliver an unprecedented SPEC SFS score but the configuration to achieve that score is 10X the organization's budget then it doesn't have much value. Some vendors will submit multiple configurations so that customers can see the performance difference at different price bands.

Ideally, organizations should use the benchmark as an initial cut list to narrow down the field of potential vendors to two or three systems that are brought in-house for on-premises testing.

Test Equal to Your Budget

When it comes time to test, make sure that during the proof of concept the vendor sends a storage system configuration that is within budget. Know exactly what the test configuration costs. Simulating the workload to perform the actual test is difficult, again especially with AI/ML workloads. The best case scenario is to use an application, server and storage configuration that duplicates production as closely as possible. An alternative is to use a workload generator solution that can capture realtime IO from production and play back that IO on the test configurations. A final option is to use standard testing tools on the equipment, tweaked to simulate the workload's IO pattern. Each of these options gets steadily worse in terms of accuracy.

Leverage the Cloud

An increasing number of modern file systems can run equally well in the cloud as they can on-premises. Leveraging the public cloud may be the ideal test environment. Compute power and storage IO can be "rented" as needed during the test and then "torn down" after the test is complete. The organization is only paying for the test environment while an actual test is in progress. Even if the organization still decides to test on-premises, leveraging the cloud may lower the short list to a single candidate.

Deciding on the storage platform for the organization's AI/ML initiatives is not a task to be taken lightly. Selecting the right storage solution lays a foundation for future AI/ML investment and keeps the organization from buying a new system as each AI/ML project spins up. Testing and evaluating these systems is difficult but leveraging published benchmarks, dissecting them for reality and then performing limited internal testing can lead the organization to the right choice. If the file system has native cloud functionality, that makes the internal testing easier and much less expensive.

CHAPTER 4: WekaIO for AI and High-Velocity Analytics

Storage Switzerland has previously discussed the problems that legacy storage file systems have when it comes to serving modern workloads such as artificial intelligence (AI) and high-velocity analytics. We have also explored the qualities that a modern file system requires. In this chapter, we will evaluate WekaIO as a solution to the storage IO challenges modern workloads create.

The Requirements of AI and High-Velocity Analytics Workloads

AI and high-velocity workloads impose new demands on storage and compute infrastructure alike. They require that unprecedented volumes of data (on the scale of terabytes and, increasingly frequently, petabytes) be processed, to ensure the most accurate response to analytics queries and training of neural networks that fuel AI. This data typically has variable and unpredictable access patterns. These workloads necessitate high-end graphics processing units (GPUs) to enable this data to be processed quickly and accurately. Because this compute infrastructure is expensive, optimal utilization is critical. However, most storage infrastructures lack the levels of bandwidth and low latency that are required to fully saturate the GPU clusters.

WekaIO Matrix File System

WekaIO dubs this storage bottleneck “I/O starvation.” Its Matrix file storage architecture was designed to be massively scalable and parallel in nature, so that large amounts of random data from a centralized shared pool can be fed instantaneously and continuously to multiple GPUs, whether on or off-premises. The Matrix architecture can achieve performance of over 10 Gigabytes per second per GPU node, which is ten times that of traditional network file systems and three times that of a local non-volatile memory express (NVMe) solid-state disk (SSD), according to WekaIO. It

> *“…users may scale their namespace up to an exabyte of capacity…”*

distributes data and metadata across the infrastructure for parallel access, and employs an InfiniBand or 10Gbit and above Ethernet network stack to facilitate rapid and predictable performance without the complexity associated with copying data between direct-attached storage nodes. Performance scales linearly to further support GPU utilization.

The Matrix architecture facilitates a centralized, singular global namespace across performance SSD storage capacity and lower-cost object storage for long-term retention. According to WekaIO, users may scale their namespace up to an exabyte of capacity, and any Amazon S3 compatible object store (whether on or off-premises) is supported.

The solution can be deployed on-premises on pre-validated industry-standard servers, as well as in the public cloud on Amazon Web Services EC2 P3 GPU instances. Cloud bursting via a hybrid model is supported to address peak workload periods. Users with an on-premises footprint may scale on demand to cloud GPU clusters as needed, and then migrate their data back on-premises when processing is complete. The user either creates a snapshot of the file system that runs in the cloud environment, or they store a backup copy of the file system, in the cloud environment, that can be rehydrated on demand.

Conclusion

As discussed throughout this eBook, AI and high-velocity analytics are among the growing number of modern workloads that require parallel processing of large volumes of data. These workloads require a re-think of file storage architectures to deliver required levels of performance and storage capacity without breaking the budget. For its part, WekaIO’s architecture offers levels of linearly scalable performance, parallel processing and cloud bursting that can help to optimize utilization of expensive on-premises infrastructure deployments, and to facilitate agile responsiveness to data intensive application requirements. Enterprises should consider the Matrix solution as a path to faster and more cost-effective analytics and AI workloads storage infrastructure.

The Firm

Storage Switzerland is the leading storage analyst firm focused on the emerging storage categories of memory-based storage (Flash), Big Data, virtualization, and cloud computing. The firm is widely recognized for its blogs, white papers and videos on current approaches such as all-flash arrays, deduplication, SSD's, software-defined storage, backup appliances and storage networking. The name "Storage Switzerland" indicates a pledge to provide neutral analysis of the storage marketplace, rather than focusing on a single vendor approach.

About Our Partner

WekaIO helps companies manage, scale and futureproof their data center so they can solve real problems that impact the world. WekaIO Matrix™, the world's fastest shared parallel file system and WekaIO's flagship product, leapfrogs legacy storage infrastructures by delivering simplicity, scale, and the best performance density per U, for a fraction of the cost. In the cloud or on-premises, WekaIO's NVMe-native high-performance software-defined storage solution removes the barriers between the data and the compute layer, thus accelerating artificial intelligence, machine learning, genomics, research, and analytics workloads.