FOR THE WANT OF A NAIL –

How to Unleash the Latent Knowledge and Commercial Value of Your Data

WEKA

Author:

Liran Zvibel

Co-Founder and CEO

WekaIO

INTRODUCTION

It does not matter what industry your business is in, AI can help. AI fundamentally changes business processes and re-envisions data usage, which in turn demands a far-reaching reconsideration of data management. If you want to take full advantage of the business transformation that AI affords, then you'll need to think differently about your infrastructure. AI and learning applications (deep learning and machine learning) require massive amounts of compute power, network bandwidth, and fast AI storage.

TABLE OF CONTENTS

CHAPTER 1:

How Infrastructure May Be Limiting AI Adoption

Artificial Intelligence (AI) has come of age. It's all around us. Yet many people don't even recognize it and what it's doing for us every day. There are many examples of cool stuff being powered by AI. Consider smart cars (Waymo, Uber, Mobileye); smart home devices (refrigerators, HVAC, Internet of Things (IoT)); virtual personal assistants (Siri, Cortana, Google Now); and Virtual Reality games (Call of Duty, Far Cry, etc.). There are also more practical AI applications such as expert systems (online customer support; purchase prediction, fraud detection, voice and speech recognition) as well as pattern and image recognition (identity confirmation/facial recognition, security surveillance, genome mapping, digital pathology, and so forth).

But perhaps most important, yet often out of the public view are the business applications for AI. The IoT is more than just remote-control refrigerators. It enables scenario simulation, fault prediction for business processes, monitoring and maintenance of industrial equipment such as early detection of gas leaks in pipelines. Frost & Sullivan estimates that AI clinical support will enhance medical imaging diagnosis to potentially improve patient outcomes by 30%+ while cutting treatment costs by up to 50%[1]. In 2014, an estimated 228 thousand human genomes were sequenced. By 2017 that figure is expected to jump to 1.6 million genomes, each representing hundreds of GB of data[2]. Machine learning can assist in preliminary drug discovery, clinical trial research, and next-generation sequencing[3]. Intelligent agents with machine learning enable companies such as Walmart to process high volume of transaction records within seconds[4]. Companies such as Pacific Specialty garnered new insights through a holistic view of data and analytics to anticipate customer needs and new underwriting opportunities.[5] Simply stated, AI can unleash the latent knowledge deep within an organization's data stores to drive business success.

AI fundamentally changes business processes and re-envisions data usage, which in turn demands a far-reaching reconsideration of data management. Unlike traditional enterprise applications that are largely transactional by nature (they operate in a sequential and discrete fashion on small data sets), AI workloads are highly parallel with continuous-interrelated activities that act as feedback loops to one another. This is the crux of machine learning: continually iterating on data to gain new insight and knowledge. As a result, AI apps require tremendous amounts of compute and data, which challenges existing network, computer, and storage infrastructures. Although AI apps and their implementation will differ, they will have broad impact across all verticals/industries. Yet at present, few businesses have deployed AI at scale. A recent McKinsey report indicated that only 20 percent of surveyed organizations currently use any AI related technology at scale or in a core part of their businesses.[6]

With the potential of AI being so great, how can businesses overcome the infrastructure challenge? AI requires very high performance with low latency. Traditionally, this has demanded expensive investments in legacy technologies; however, the volume of data and the performance required for AI is not viable on these architectures. This is part of the reason that the prevalence and scale of AI deployments have been limited. The good news is that technological advancement means that some needs have already been addressed. Compute challenges have been overcome through parallel workload processing, which is enabled by innovative technology such as NVIDIA® GPUs. Likewise, network performance has been enhanced through solutions such as InfiniBand connectivity as delivered by Mellanox. This leaves Data Management and Storage as two areas where innovative solutions are essential to ensure that AI initiatives reach their full potential.

By solving data management and storage challenges, AI initiatives could become within the economic reach of not only large enterprises but small and mid-sized organizations as well. While the breadth and depth of AI apps will vary by company size, the ability to deploy them economically at an appropriate scale would be a benefit for any organization. But how can data management and the underlying data center storage architecture be aligned with the needs of AI workloads? That's the topic I'll address in the next chapter.

[1] "From $600 M to $6 Billion, Artificial Intelligence Systems Poised for Dramatic Market Expansion in Healthcare." Frost.com. January 5, 2016. https://ww2.frost.com/news/press-releases/600-m-6-billion-artificial-intelligence-systems-poised-dramatic-market-expansion-healthcare.

[2] Hiatt, David. "The Next Digital Arms Race In Life Sciences." The Next Digital Arms Race In Life Sciences. August 23, 2017. http://www.bio-itworld.com/2017/08/23/the-next-digital-arms-race-in-life-sciences.aspx.

[3] Mehta, Tapan. "Artificial Intelligence Euphoria in Healthcare." Artificial Intelligence Euphoria in Healthcare. November 13, 2017. https://dminc.com/blog/artificial-intelligence-euphoria-healthcare-life-sciences/.

[4] Ruth, Joao-Pierre. 2017. "6 Examples of AI In Business Intelligence Applications". Techemergence. https://www.techemergence.com/ai-in-business-intelligence-applications/.

[5] "Pacific Specialty Case Study". 2016. Avanade.Com. https://www.avanade.com/~/media/asset/case-study/pacific-specialty-case-study.pdf.

[6] "Artificial Intelligence The Next Digital Frontier?". 2017. Discussion Paper. McKinsey Global Institute. https://www.mckinsey.com/~/media/McKinsey/Industries/Advanced%20Electronics/Our%20Insights/How%20artificial%20intelligence%20can%20deliver%20real%20value%20to%20companies/MGI-Artificial-Intelligence-Discussion-paper.ashx.

CHAPTER 2:

Aligning Data Center Storage with the Needs of AI Workloads

As I discussed in the previous chapter, Artificial Intelligence (AI) fundamentally changes business processes and provides potentially life altering results by unleashing the latent knowledge deep within an organization's data stores. Of course, discovering latent knowledge requires an aggressive approach to find, unlock, and leverage its associated business value. In other words, there is a lot of work associated with AI — it's more than just having the right applications — it's about having the right infrastructure that can cost-effectively support them.

Deploying AI in data centers challenges existing network, computer, and storage infrastructures with workloads that are parallel in nature with interrelated activities that continuously act as feedback loops to one another. According to Topbots[1], there are twelve essential categories of AI vendors: business intelligence, productivity, customer management, HR & talent, B2B sales & marketing, consumer marketing, finance & operations, digital commerce, data science, engineering, and industrials & manufacturing. As you might imagine, unlocking knowledge (and business advantage) requires sifting through lots of data, and perhaps across formerly siloed disciplines. If your storage/data center administrators are already pushed to the limit maintaining the status quo, just imagine their reaction to the needs of AI.

The key design consideration for AI-capable infrastructure is the ability to deliver high bandwidth with low latency for both small and large files. Consider an AI workload that supports the monitoring and maintenance of industrial equipment. This entails periodically reading log files, which themselves are usually small in size. Yet the aggregated log data for all the equipment on the shop floor, if not multiple sites, could yield a huge data set. Conversely, an image/pattern recognition workload might be examining hours of surveillance data searching for the moment when the unauthorized removal of an item from the premises occurred.

Another example involves self-driving cars, which require real-time processing of image and censor data to help navigate common driving hazards such as roadway construction, traffic, emergency vehicles, and poor driving conditions. While much of this data and its processing will occur in-vehicle, a great deal of data will also need to be transferred and stored for use later by manufacturers to improve decision making algorithms and by insurance companies for liability purposes. Consider for a moment the volume of data that would be generated by a fleet of delivery vehicles or rental cars; it's easy to see the impact on such as enterprise.

This raises an issue beyond file size and storage capacity. In traditional environments, compute and data are typically housed in a centralized data center. Therefore, the distance and hence latency of interconnections can be quite low. However, the Internet of Things (IoT) is predicated upon data that can be distributed far away from compute resources. For real-time applications, low latency is essential for smooth operations. The highly distributed and parallel nature of AI applications demand high bandwidth at low latency, regardless of file size and location.

Traditional parallel file systems such as Spectrum Scale, Lustre, Hadoop, and others support highly coordinated concurrent access from multiple clients while featuring optimized I/O paths for maximum bandwidth. As a result, these files systems can increase throughput and scalability for large files. However, they are not designed with smaller files in mind and generally do not offer the same performance associated with large files. In addition, these file systems are 20 years old and do not take advantage of the benefits offered by advanced technologies such as GPU processors and NVMe flash memory. At the same time, parallel file systems are extremely complex because there are many moving parts (metadata servers, multiple storage targets, tunable system parameters, etc.) that require ongoing optimization to run at peak efficiency. Data management in such environments is a nuanced and ongoing specialized task that is typically beyond the capability of a traditional storage administrator. As a result, such installations require dedicated, skilled architects and administrators, which can be a non-trivial expense.

As you are probably starting to realize, a successful AI deployment will require a rethinking of existing infrastructure in order to maximize your AI efficiency and ROI. But this is just the start. Artificial Intelligence is driven by large amounts of data, yet the result of AI is even more data. The ability to scale with the growth of AI data is essential for success. In the next chapter I'll discuss how AI depends upon large scale storage.

[1] Jia, Marlene. " The Essential Landscape of Enterprise AI Companies". March 31, 2017. Accessed November 07, 2017. https://www.topbots.com/essential-landscape-overview-enterprise-artificial-intelligence/

CHAPTER 3:
AI Depends on Large Scale Storage

As I discussed in my previous post, AI workloads are data driven. Today's analytics-based AI requires tremendous amounts of AI storage. Without this capacity, you would not be able to benefit from information and knowledge unleashed by AI workloads. Whether you are sequencing human genomes, reading medical imaging, or performing clinical trial research, the amount of data to process is huge. The same is true with machine learning monitoring the Internet of Things (IoT); intelligent agents enabling customer support, purchase prediction, and fraud detection; and business intelligence and other analytical applications. In each case, you can see the commonality of storage at a large, if not enormous scale. Plus, AI deployments result in even more data to digest and reprocess ad infinitum; hence the ever-growing need for more capacity.

Unlike centralized workloads where the bulk of data resided neatly in the datacenter, AI data comes from a variety of locations. Consider IoT workloads where data tends to originate on the outer edge of the network; after all, this is where the "things" reside and the action takes place. However, the processing takes place at the central hub rather than out on the spokes. A genome comprises terabytes of raw data, and is likely a shared resource that resides away from the application server. Managing data access in these environments requires not only massive scalability but also high I/O rates combined with low latency. Without these, processing would grind to a halt as workloads would be severely I/O bound, leaving expensive GPU based servers idling as they wait for data.

Traditionally, a storage solution that would meet these criteria required expensive investments to achieve the requisite scalability and performance. As a result, all but the largest enterprises or most focused startups could not economically justify deploying AI to harvest the commercial benefits that AI could deliver. You might be thinking, "Yes, perhaps, but there are many alternatives today that were not available in the past." That's true. Let's consider some commonly thought of alternatives.

Cloud Service Providers (CSP) are often viewed as an option; after all, they require minimal or no capital investment in new hardware and scale easily with a flexible pay-as-you-go pricing scheme. However, these low-cost approaches lack the configuration flexibility of onsite equipment and therefore can rarely meet the specialized needs of AI workloads. Additional problems associated with CSPs include noisy neighbors who may be co-resident on the same physical infrastructure, negatively affecting application performance, limited network or storage bandwidth that impacts latency and throughput, and the requirement to move large amounts of data to and from the cloud, which is both time consuming and costly.

Off-the-shelf NAS solutions might seem to be an easy-and-immediate solution. However, while perfectly acceptable at smaller scale, NAS has limited capacity scalability along with decreasing performance and increased latency at the scale required by AI workloads. How about a hybrid approach? Theoretically one could devise such a solution. However, remember that a hybrid approach combining multiple components would result in many moving parts including metadata servers, multiple storage targets, etcetera, each of which would require ongoing tuning to run at peak efficiency. Does this seem like a simple, cost-effective, and scalable solution?

After a closer examination, these choices aren't as viable as first thought. But what is viable alternative that is also affordable? A successful solution needs to be able to prioritize fast and efficient data storage while being able to scale to petabytes of data capacity. Remember, too that scalability is bidirectional. Also, keep in mind that files are not one size fits all. A medical image is very different from a word processing document, log file, or database table. Consistent performance regardless of file size or type is an essential part of any data management and storage solution.

By now, I suspect that you are starting to appreciate the different data management and storage needs of AI workloads. These workloads represent a change for most any organization, and with change comes the challenge of meeting the need while maximizing ROI. In the next chapter, I'll discuss the characteristics of the new approach to data management and storage infrastructure that you will need if you want to cost-effectively deploy AI.

CHAPTER 4:

Want AI? You'll Need a Modular Approach to Maximize GPU Performance

Previously I've discussed how AI is all around us. It's obvious that storage scalability is essential, perhaps at a level that you might not have experienced in the past. AI workloads are highly parallel with continuously-interrelated activities that act as feedback loops to deliver more refined and actionable data. However, simply expanding your existing infrastructure is not going to enable a successful AI deployment. While you can achieve scale through Cloud Service Providers, their offerings are rarely tailored to a specific need, and WAN connectivity is not cost effective at the requisite I/O bandwidth and latency. Alternatively, onsite high performance NAS isn't the answer either. Whereas NAS might appear to be an easy-and-immediate solution, design limitations severely limit its scalability, and hence suitability. Scaling is more than just raw capacity, it's the ability to handle millions of directories with billions of files per directory, while delivering consistent performance/latency regardless of whether a directory contains five 10 TB files or 7 million 4 KB files. And consideration needs to be paid to optimizing GPU performance–don't let these expensive resources sit idle.

Data management is an important part of any storage solution. While data layout is not as critical as in the past (for example locality has been overcome by flash and low latency interconnects such as InfiniBand), it remains a consideration in overall design. Effective management includes appropriate storage tiering that aligns with performance and cost considerations, data protection commensurate with data value, support for concurrent access and multiple protocols as well as differing data types such as file/OBS and structured/unstructured. Further, disparate workloads require support for multiple file systems. For example, logical partitions can have different QoS needs such as high performance, data protection level (N+2, N+4, etc.), and use cases (scratch vs. library vs. working data set, etc.).

It's time for something new. But what is a viable alternative?

A successful solution needs to prioritize fast and efficient data storage and scale to petabytes of capacity. This implies a modular approach, one that balances capacity and performance while maintaining the ability to scale either independently of the other. Innovative data management tools that can enable dynamic and tailored high-performance storage that operates in a multi-petabyte namespace are essential. Keep in mind that scalability is a two-way street. Unused capacity represents money and resources that could be better spent elsewhere. Therefore, an efficient data protection scheme is critical. Why settle for the cost of triple replication (300% overhead!) when you could achieve the same protection at significantly lower capital and operating expense?

Critically, any solution must be able to deliver high bandwidth and low latency while maintaining a consistent performance profile across all file sizes. AI relies heavily on GPU based servers that can cost between ~$50k and ~$80k each. Inadequate performance from the storage systems results in poor GPU performance, because the GPUs are starved for data, with idle compute cycles, which is a very inefficient use of such expensive resources. Since AI data can be found along the edges of the network, support for distributed environments including parallel access for POSIX, NFS, SMB, and HDFS protocols is key. Coherent and consistent performance regardless of file location is also an essential part of any data management solution. "What about data locality?" you might be thinking. Data locality is irrelevant since local copy architectures (e.g. Hadoop, or caching solutions) were developed in a 1GbitE and HDD era. Modern 10GbitE networks are 10x faster than single SSDs, and it is much easier today to create distributed algorithms where locality is not important. You see, with right networking stack, shared storage is faster than local storage.

Furthermore, the solution must be able to cost-effectively achieve these goals so that organizations of most any size can economically justify investments in AI. This precludes simply investing in additional equipment and personnel as in the past. Instead, why not rethink how the current resources within an organization are deployed to maximize the ROI on existing investments? After all, AI is all about unleashing latent knowledge and commercial value. By solving data management and storage challenges in a new way at a fraction of the cost of traditional approaches, AI initiatives could become within the economic reach of not only large enterprises but small and mid-sized organizations as well.

Rather than investing in additional hardware or CSPs, consider leveraging your existing infrastructure and spare resources to develop a software-centric data management approach that can meet your AI needs. Does this sound too good to be true? Well, it's not. In the final chapter I will show exactly how WekaIO is enabling AI for organizations of all sizes.

CHAPTER 5:

Enabling AI For Organizations Of All Sizes

Over the last chapters I've discussed how Artificial Intelligence (AI) is all around us and its potential to unleash latent knowledge deep within your organization's data stores. It does not matter what industry your business is in, AI can help.

If you want to take full advantage of the business transformation that AI affords, then you'll need to think differently about your infrastructure. AI and learning applications (deep learning and machine learning) require massive amounts of compute power, network bandwidth, and fast AI storage. Fortunately, GPU based servers are ideally suited for AI and learning type applications. Infiniband is well suited to deliver extremely low latency and high network bandwidth, and parallel file systems make all the server centric storage and data shareable.

However, not all parallel file systems are created equal. In fact, the traditional file systems used in high performance computing (Lustre and Spectrum Scale) were not designed to take advantage of the performance and low latency of NVMe flash. This is important because AI is one of the most demanding workloads today; it consists of both large and small files, random and sequential access, and structured and unstructured data. AI applications are also very metadata intensive, so the file system must be able to consistently deliver very high metadata performance – not an easy task. For these legacy file systems to perform, AI systems must be over-engineered and augmented with large caching devices to provide decent small file and metadata performance. The result is an overly expensive solution.

GPU servers are quite expensive because they can process data hundreds of times faster than a similar CPU based server. The table below from an article in The Next Platform illustrates this point well (Table 1). Note the extreme difference between the performance of a Xeon CPU based server and that of Nvidia's DGX-1 GPU server. This difference in performance puts a huge demand on the supporting network and storage infrastructure.

	DUAL XEON	DGX-1
FLOPS (CPU + GPU)	3 TF	170 TF
AGGREGATE NODE BW	76 GB/s	768 GB/s
ALEXNET TRAIN TIME	150 HOURS	2 HOURS
TRAIN IN 2 HOURS	>250 NODES*	1 NODE

Source: The Next Platform, https://www.nextplatform.com/2016/05/31/age-gpu-upon-us/

A GPU server consumes data at a rate of 3-4 gigabytes per second, so a 10-node GPU cluster requires an interconnect and storage system that can sustain 30-40 gigabytes per second. Such an infrastructure would be quite expensive using legacy storage solutions. However, it doesn't have to be.

You can position your organization for the future while protecting your existing investments by taking a software-centric approach to AI, learning systems, and data management. WekaIO has developed a storage solution well-suited to AI that includes the world's fastest file system. When coupled with an Infiniband network, it provides over 6 gigabytes per second of bandwidth per GPU server, more than enough performance for any AI application. In fact, this combination provides performance that is over 2x faster than a local file system with a direct attached all-flash array.

As a shareable file system, WekaFS™ is also cloud native, meaning that you can easily burst your AI workloads to a Matrix enabled GPU cluster in AWS using the Snap-to-S3 feature. This allows you to eliminate the investment in a huge AI cluster. Simply spin up a GPU cluster in AWS on-demand. WekaFS leverages S3 compatible object storage to cost-effectively scale as your training data sets grow, and data management is point and click easy, or run your automated scripts using our CLI. A single admin without any special training can easily manage petabytes of data.

SUMMARY

Overcoming the infrastructure challenge means that access to AI is no longer just for the big guys, but indeed within cost-effective reach for organizations like yours.

If this sounds intriguing to you, or if you'd just like to learn a bit more, I suggest you check out these resources or our website in general. You can learn in detail how WekaFS fundamentally changes AI and data management for the better. In addition, you can see real-world applications of WekaIO technology and how we partner with some of the leading supercomputer centers and server and networking vendors to build out an AI optimized solution. Better yet, if you are ready to embrace the future of AI, give us a call, we'd be happy to discuss your needs further.

You, too, can maximize your business potential through AI applications. You can do this with WekaFS, the fastest, most scalable file system storage for compute intensive applications.

To request a free trial of WekaFS, go to: https://www.weka.io/get-started

---

Additional resources:

- Valohai and Weka demonstration [video]
- WekaFS for AI and Analytics
- Weka AI Reference Architecture
- Gartner Glossary: DataOps
- Lenny Liebmann, Contributing Editor, InformationWeek, DataOps blog post, 2014

WEKA 910 E Hamilton Avenue, Suite 430, Campbell, CA 95008 T: 408.335.0085 E: info@weka.io www.weka.io

W02EB202008