

Predictive Outcomes at Massive Scale

with Amazon Sagemaker and WEKA

Amazon Web Services provides a managed service called Sagemaker for all the lifecycle elements of building, deploying, and maintaining a wide variety of machine learning (AI/ML) models, all without the knowledge of being an AI infrastructure expert. The focus of the service is data first, providing data curation/hygiene tools, transport, and validation tools, such as building, training, and deploying AI models. It also offers a range of built-in algorithms, frameworks, and development tools for machine learning workflows, as well as features for data preprocessing, model optimization, and real-time model monitoring.

Amazon Sagemaker has a wide range of access approaches for ingesting data for processing, with integration to AWS native storage offerings as well as others.

WEKA was founded on the idea that current storage solutions have only provided incremental improvements to legacy designs, allowing for a widening gap between compute performance and data storage performance. Storage remains a bottleneck to application performance, and with the continued densification of compute in areas such as GPU-based applications, has become even more problematic. WEKA has built a software-only, high-performance file-based storage solution that is highly scalable and easy to deploy, configure, manage, and expand. The design philosophy behind the WEKA file system (WekaFS) was to create a single storage architecture that runs on-premises and/or in the public cloud with the performance of all-flash arrays, the simplicity and feature set of network-attached storage, and the scalability and economics of the cloud.

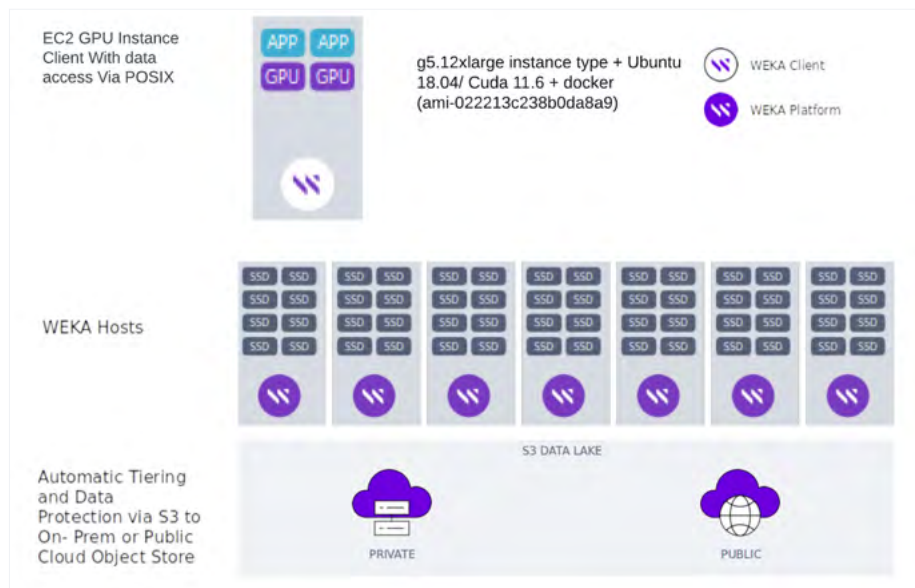
Giving Agile Compute Access to Data

Each of the Sagemaker data access patterns is useful for different requirements, for our solution validation purposes, we will be using the "Amazon Sagemaker Local Mode" configuration to facilitate training using data on a WEKA instance in the cloud.

Currently, Sagemaker is limiting its data sources to mostly only one of the AWS 1st-party solutions. We can connect it to a WEKA cluster using Sagemaker Local Mode.

Sagemaker Local Mode allows using the Sagemaker functionality on a single GPU instance with a single GPU.

When utilizing Sagemaker the data is copied from the input storage to an S3 bucket and to local NVMe's on the GPU instances to be processed. This is intended to compensate for the performance characteristics of the storage used.



WEKA eliminates the need to copy the data to/from an S3 bucket to/from local NVMe thereby saving on the effort and cost it takes and simply providing the data from WekaFS. Due to its performance characteristics the data is served faster than even the local NVMe (as we'll show in the next section)

WEKA data in place vs S3 file with local storage Comparison

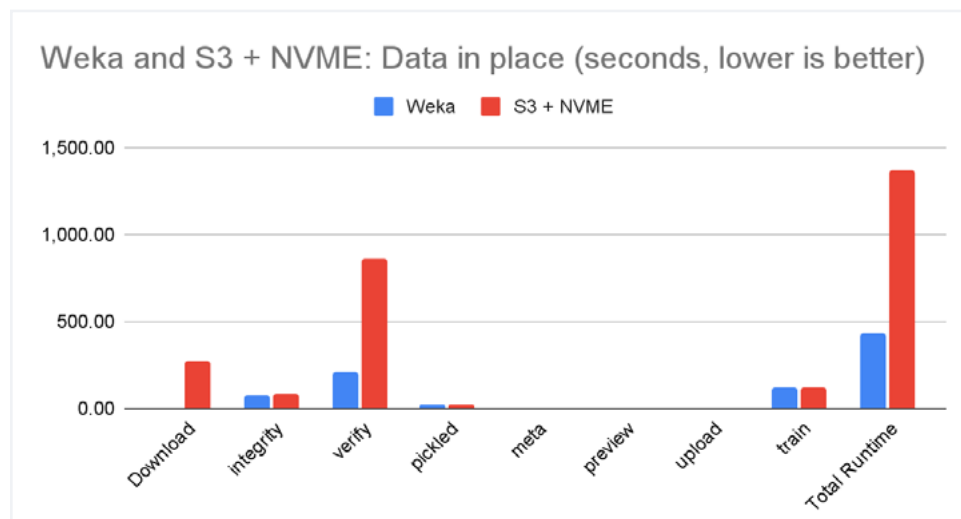
The AWS Sagemaker pipeline is composed of multiple steps as described below, and as mentioned above, since our data set is resident on WekaFS, we can eliminate the download, extraction, and operations associated with moving data, giving immediate advantage to getting the data to a GPU for training opposed to the traditional method of copying the data from the S3 bucket to the local NVMe of the instance.

The following are the steps involved in the Sagemaker pipeline

- Download an augmented CIFAR10 dataset from a local webserver for S3 + NVMe configuration only - This is intended for the comparison of S3+NVMe to WekaFS performance
- Do an md5 hash operation to validate integrity on the downloaded tarball archive for both WEKA and S3 file using a local NVMe.
- Extract the dataset to a local filesystem for both WEKA and to S3 file using a local NVMe
- Do an md5 hash operation to validate individual files containing training data for both WEKA and S3 file using a local NVMe
- Read data from the pickled (python binary encoded data record format) into memory for both WEKA and S3 file using a local NVMe
- Read metadata about the individual records to provide an effective classification for both WEKA and S3 file using a local NVMe
- Upload the uncompressed pickled records to an S3 standard bucket for Sagemaker distributed session management for both WEKA and S3 file using a local NVMe.
- Read a random batch of 4 images from the data set
- Train using the data set using a single local A10G Nvidia GPU

We will take key patterns from the Amazon Fsx for Lustre filesystems, where the primary copy of data resides.

As data is resident on WekaFS, the download element for the S3 + NVMe is the only one tracked. In this scenario, it took Sagemaker 4 minutes to download the data from the S3 bucket to the local NVMe - as the size of the data increases this time will increase as well. With WekaFS this will always remain 0 seconds since the data is already resident and does not need to use local NVMe or AWS S3 storage.



	WEKA	S3 + NVMe
Download	0	275.55
integrity	77.98	81.01
verify	206.48	858.26
pickled	23.89	23.09
meta	0.01	0.02
preview	4.11	4.46
upload	5.44	5.33
train	117.45	117.42
Total Runtime	435.37	1,365.13

Testing Conclusions

The overall pipeline on WEKA was 3X faster than compared to the S3+Local NVMe option (which is one of the standard high-performance options for Sagemaker on AWS).

The Verify step on WEKA was ~4X faster compared to the S3+Local NVMe option.

With these metrics, providing a similar pattern to the first experiment, we can confidently conclude that the WEKA filesystem's performance is at least at parity for ingesting a new dataset, verifying integrity for the tarball archive and the contents of the archive, and driving a training pipeline.

Eliminating the download phase which is normally required for every S3+Local NVMe usage is a key factor for decreasing the overall run time. With WEKA, this process is not required.

Final thoughts on access to data

Without access to key data assets, small scoops, and digests only provide tiny insights, however, when attached to a petabyte scale platform, the real strength of the solution enhances not just time to results, but expedites access to key features previously undiscovered.

Sagemaker Local Mode has the potential for producing impressive outcomes via connecting a rich library deployed in docker containers, but it is enhanced when attached to a massive amount of data on the WEKA AI-optimized data platform.

weka.io

844.392.0665

