

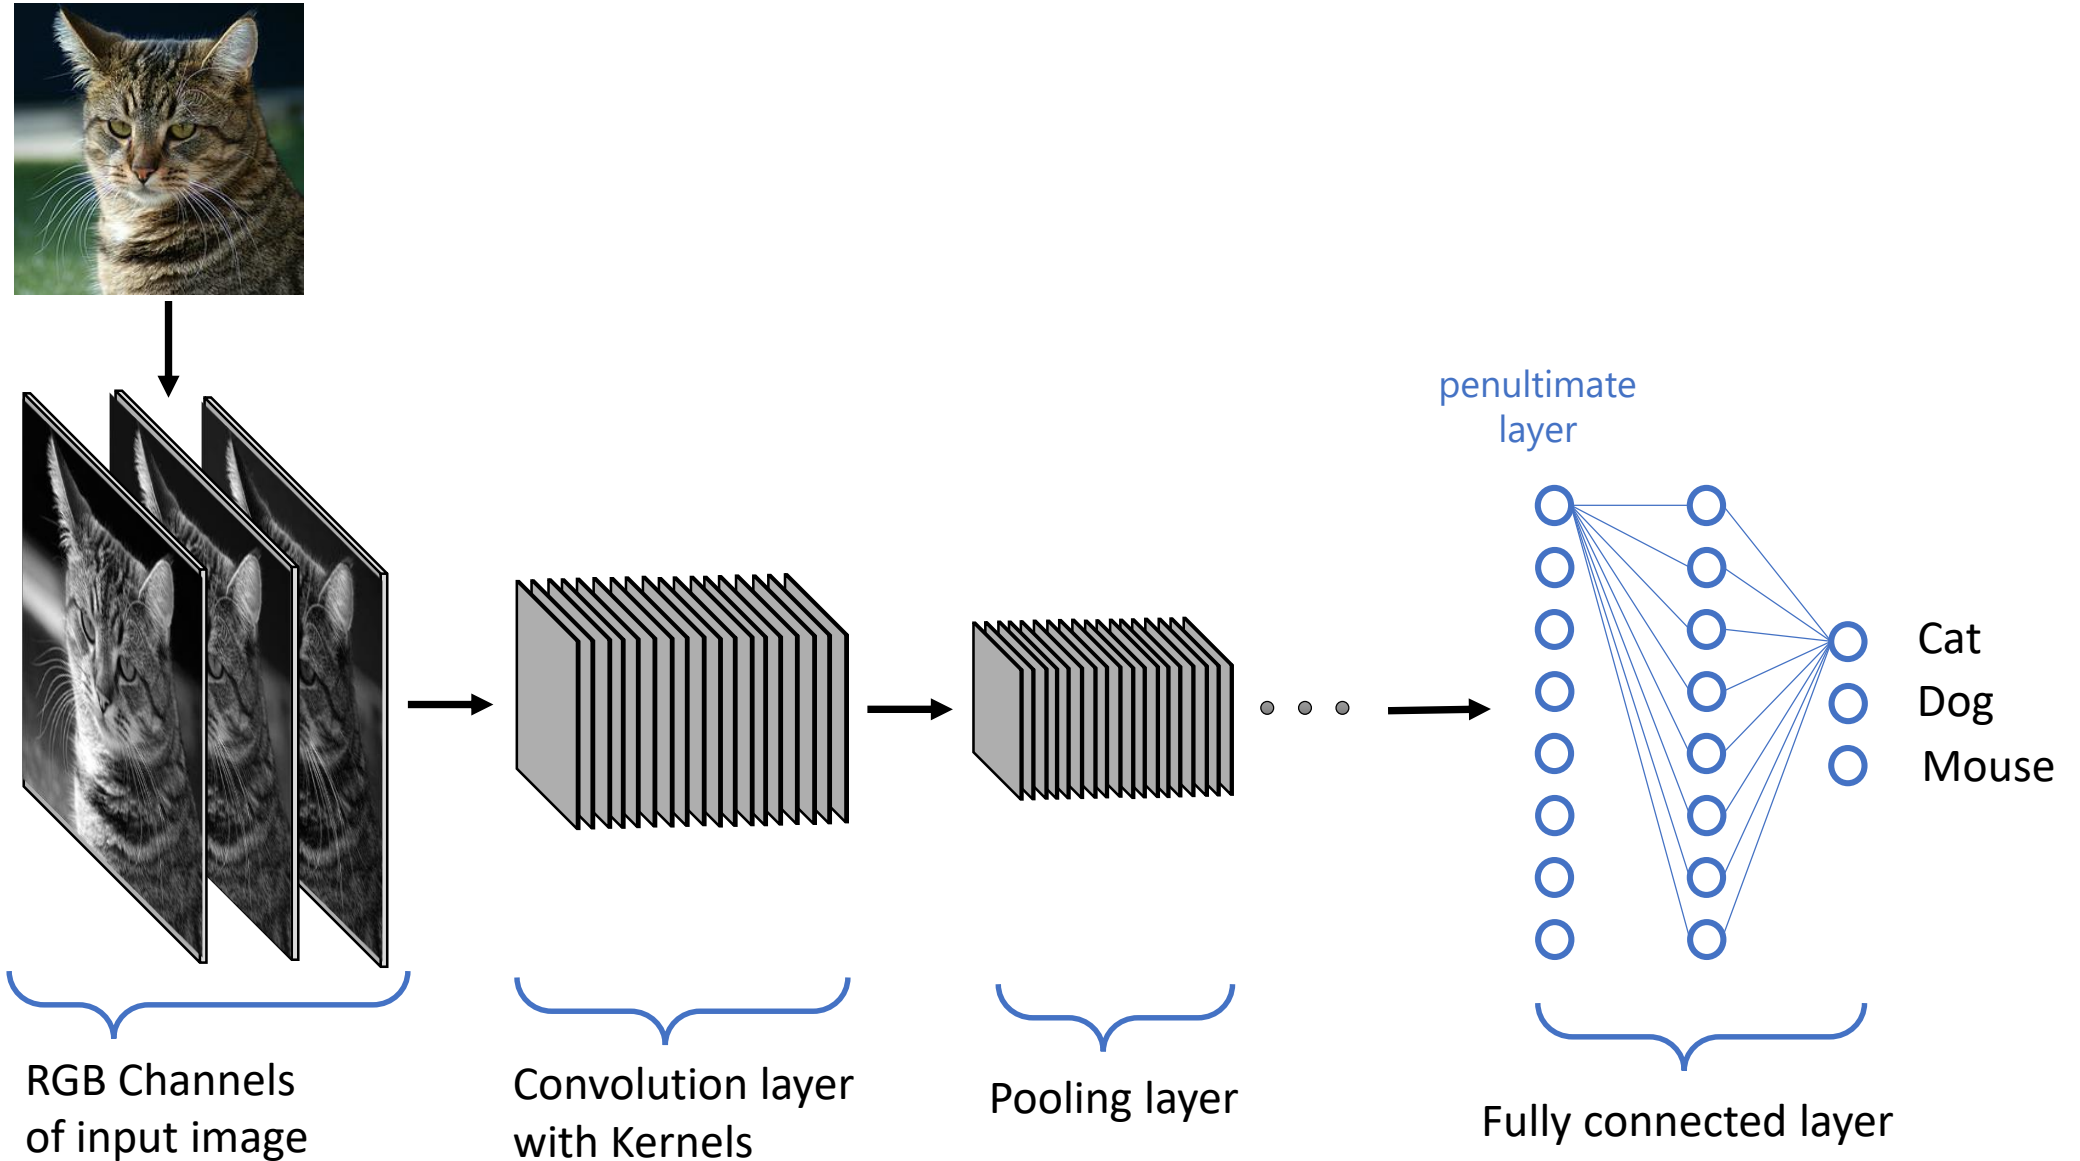
Mathew Salvaris

Distributed Deep Learning

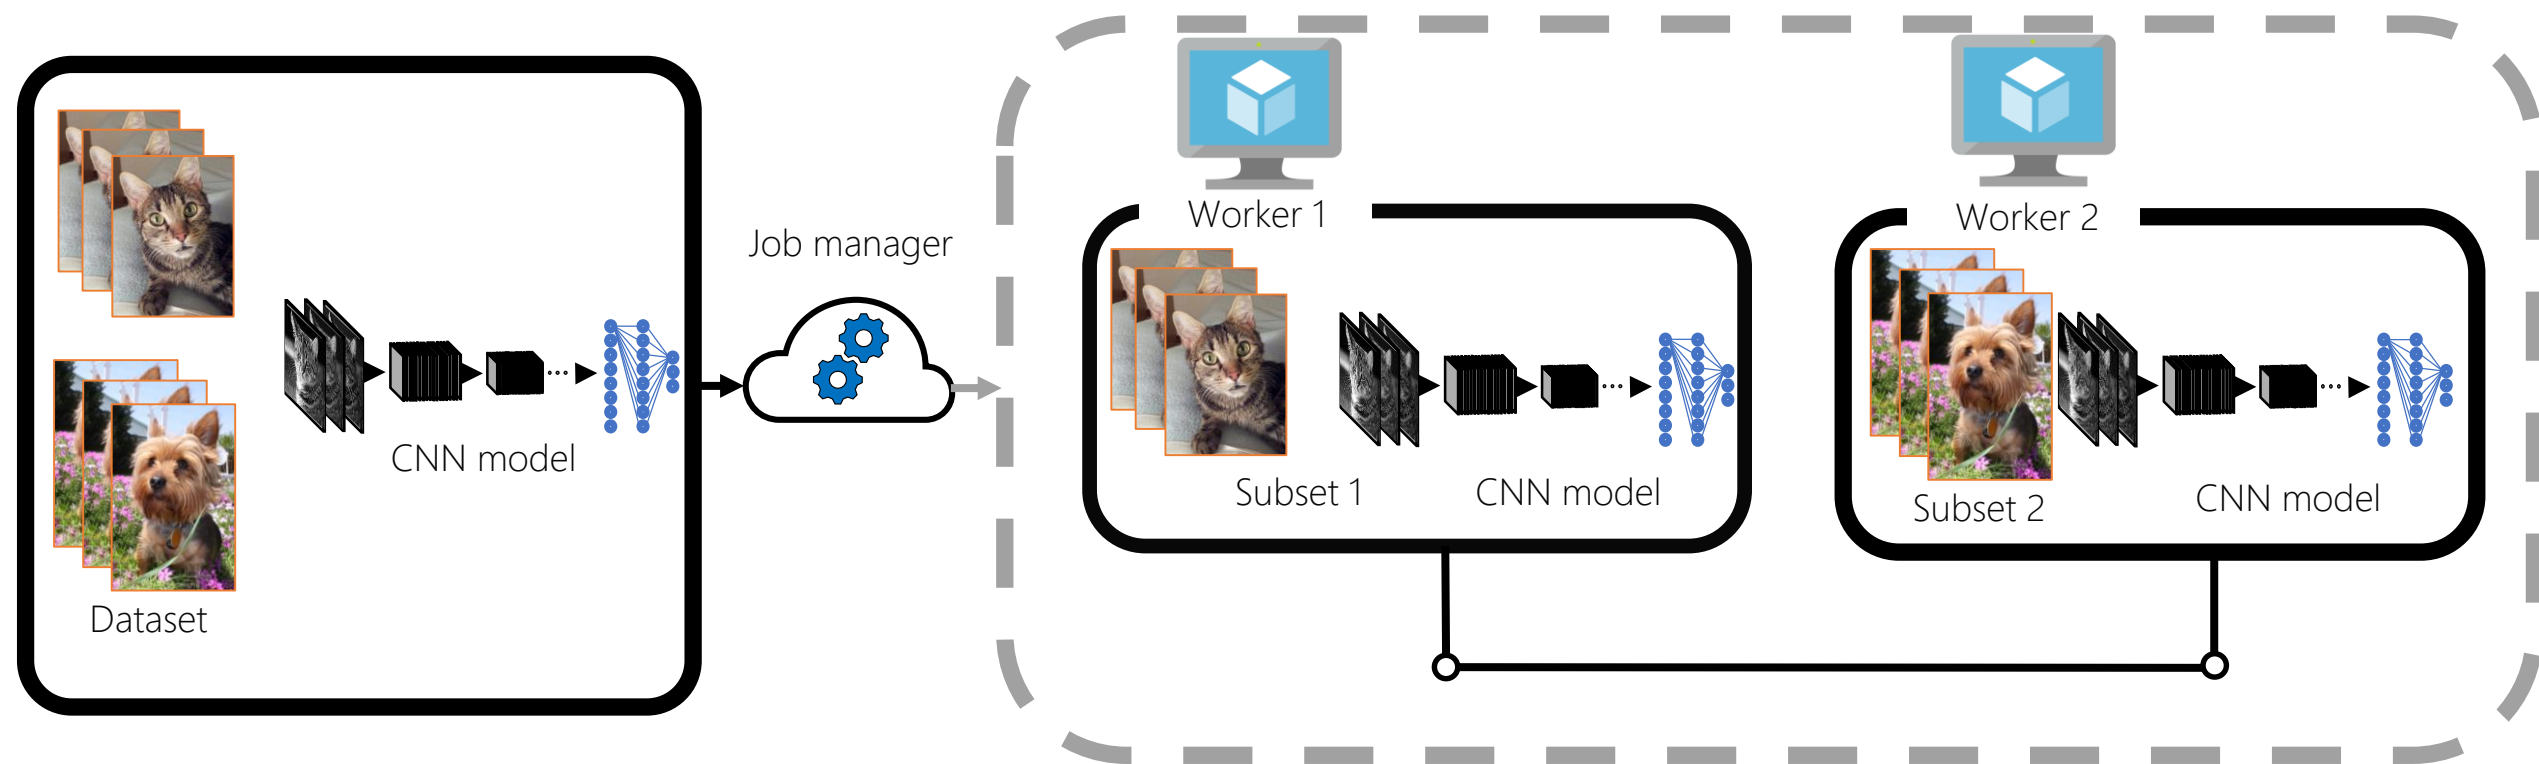
What will be covered

- Overview of Distributed Training
- What affects distributed training
 - Network
 - Model
 - Data location
 - Data format

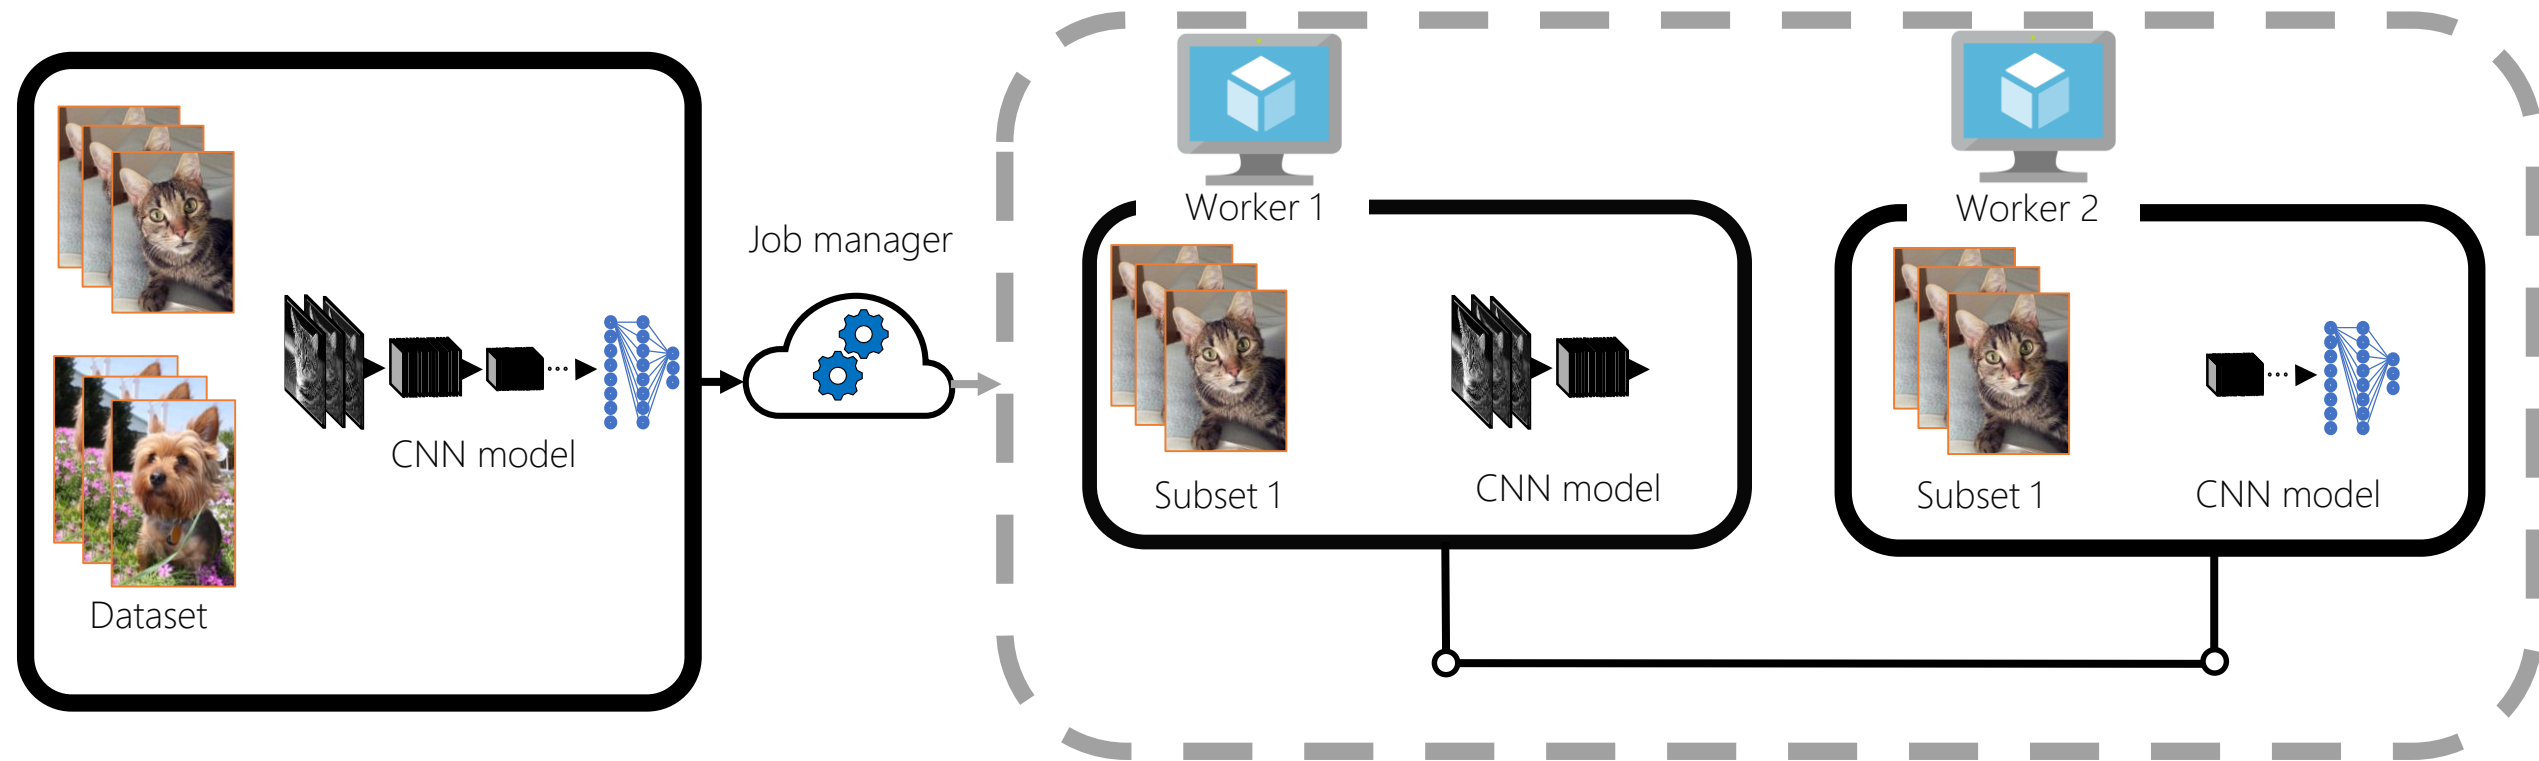
Deep Learning Model (CNN)



Distributed training mode: Data parallelism




Distributed training mode: Model parallelism




Data parallelism vs model parallelism

Data parallelism



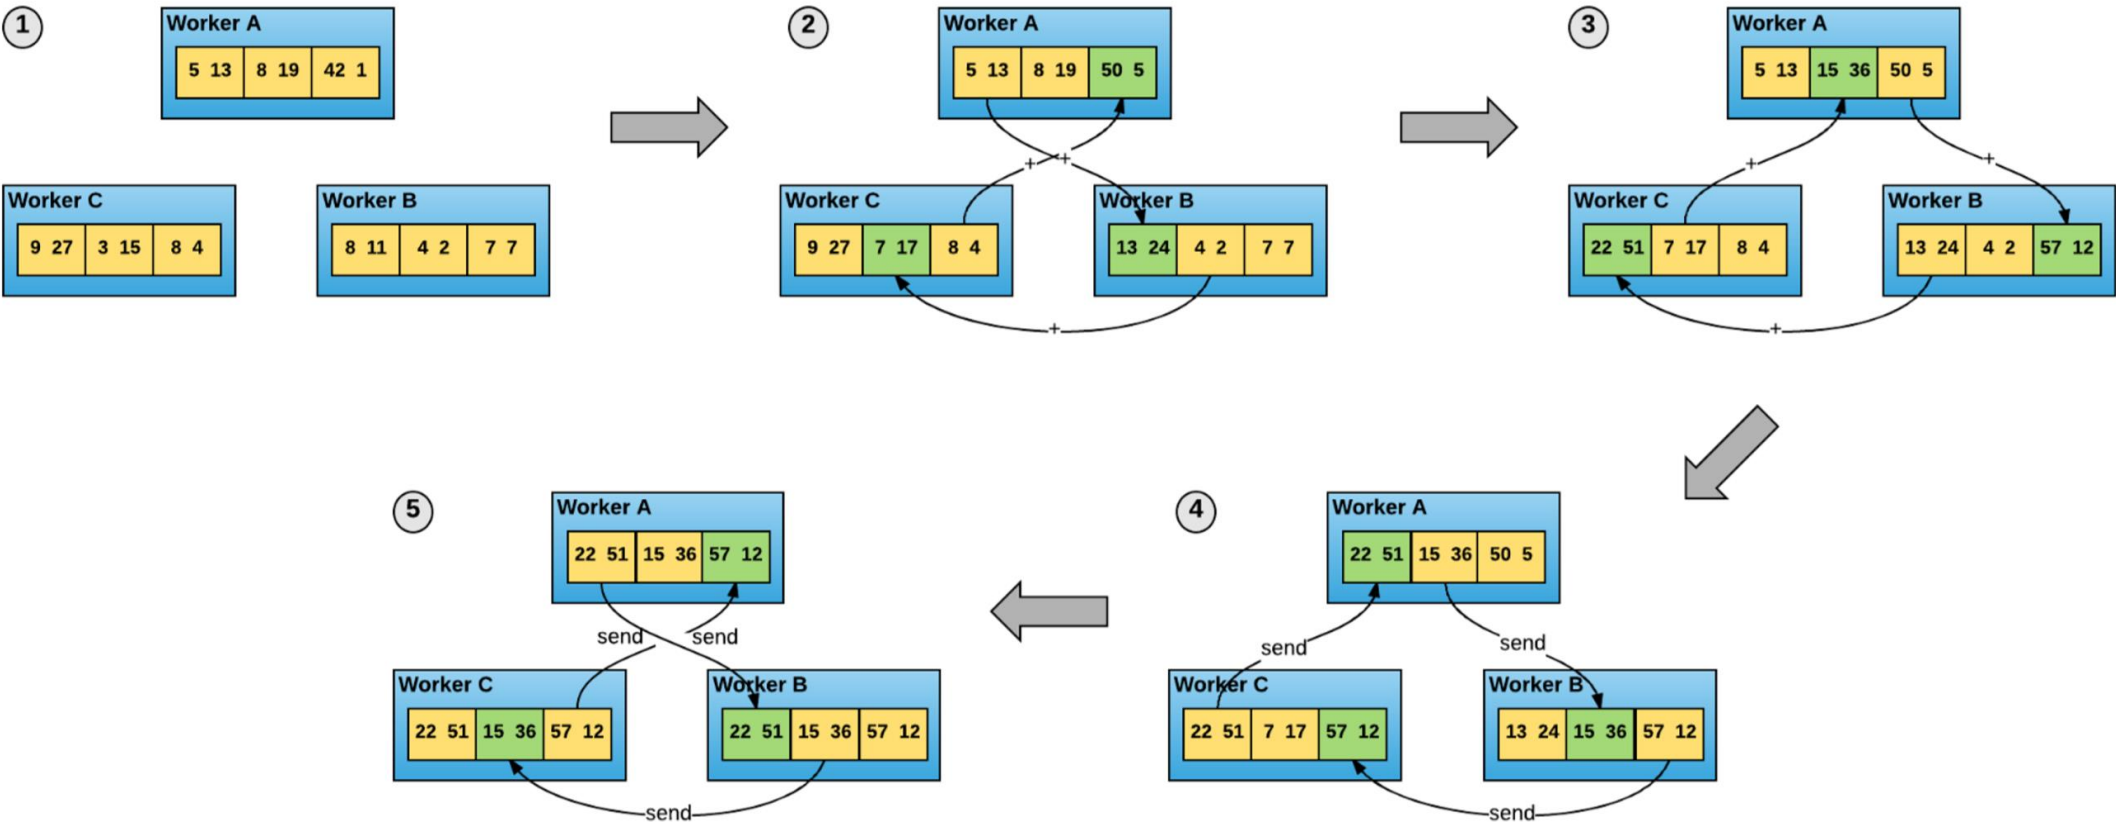
- Easier implementation
- Stronger fault tolerance
- Higher cluster utilization

Model parallelism



- Better scalability of large models
- Less memory on each GPU

Horovod: Ring All Reduce



Effects of Network, Model and Precision

Clusters of 8 nodes using **K80, P40, P100**
and **V100** (4 GPUs per node+Infiniband)

Two MPI configurations **OpenMPI+NCCL**
and **IntelMPI**

Setup

345 experiments across many different models including **ResNet50**, **MobileNet V2** etc.


Using **synthetic** data


Batch size remains **64** across all models and GPUs

Use the benchmarking scripts from TensorFlow

Experiments


Distributed training with synthetic data







Single GPU





32 GPUs






GPU K80 P40 P100 V100

Batch Execution

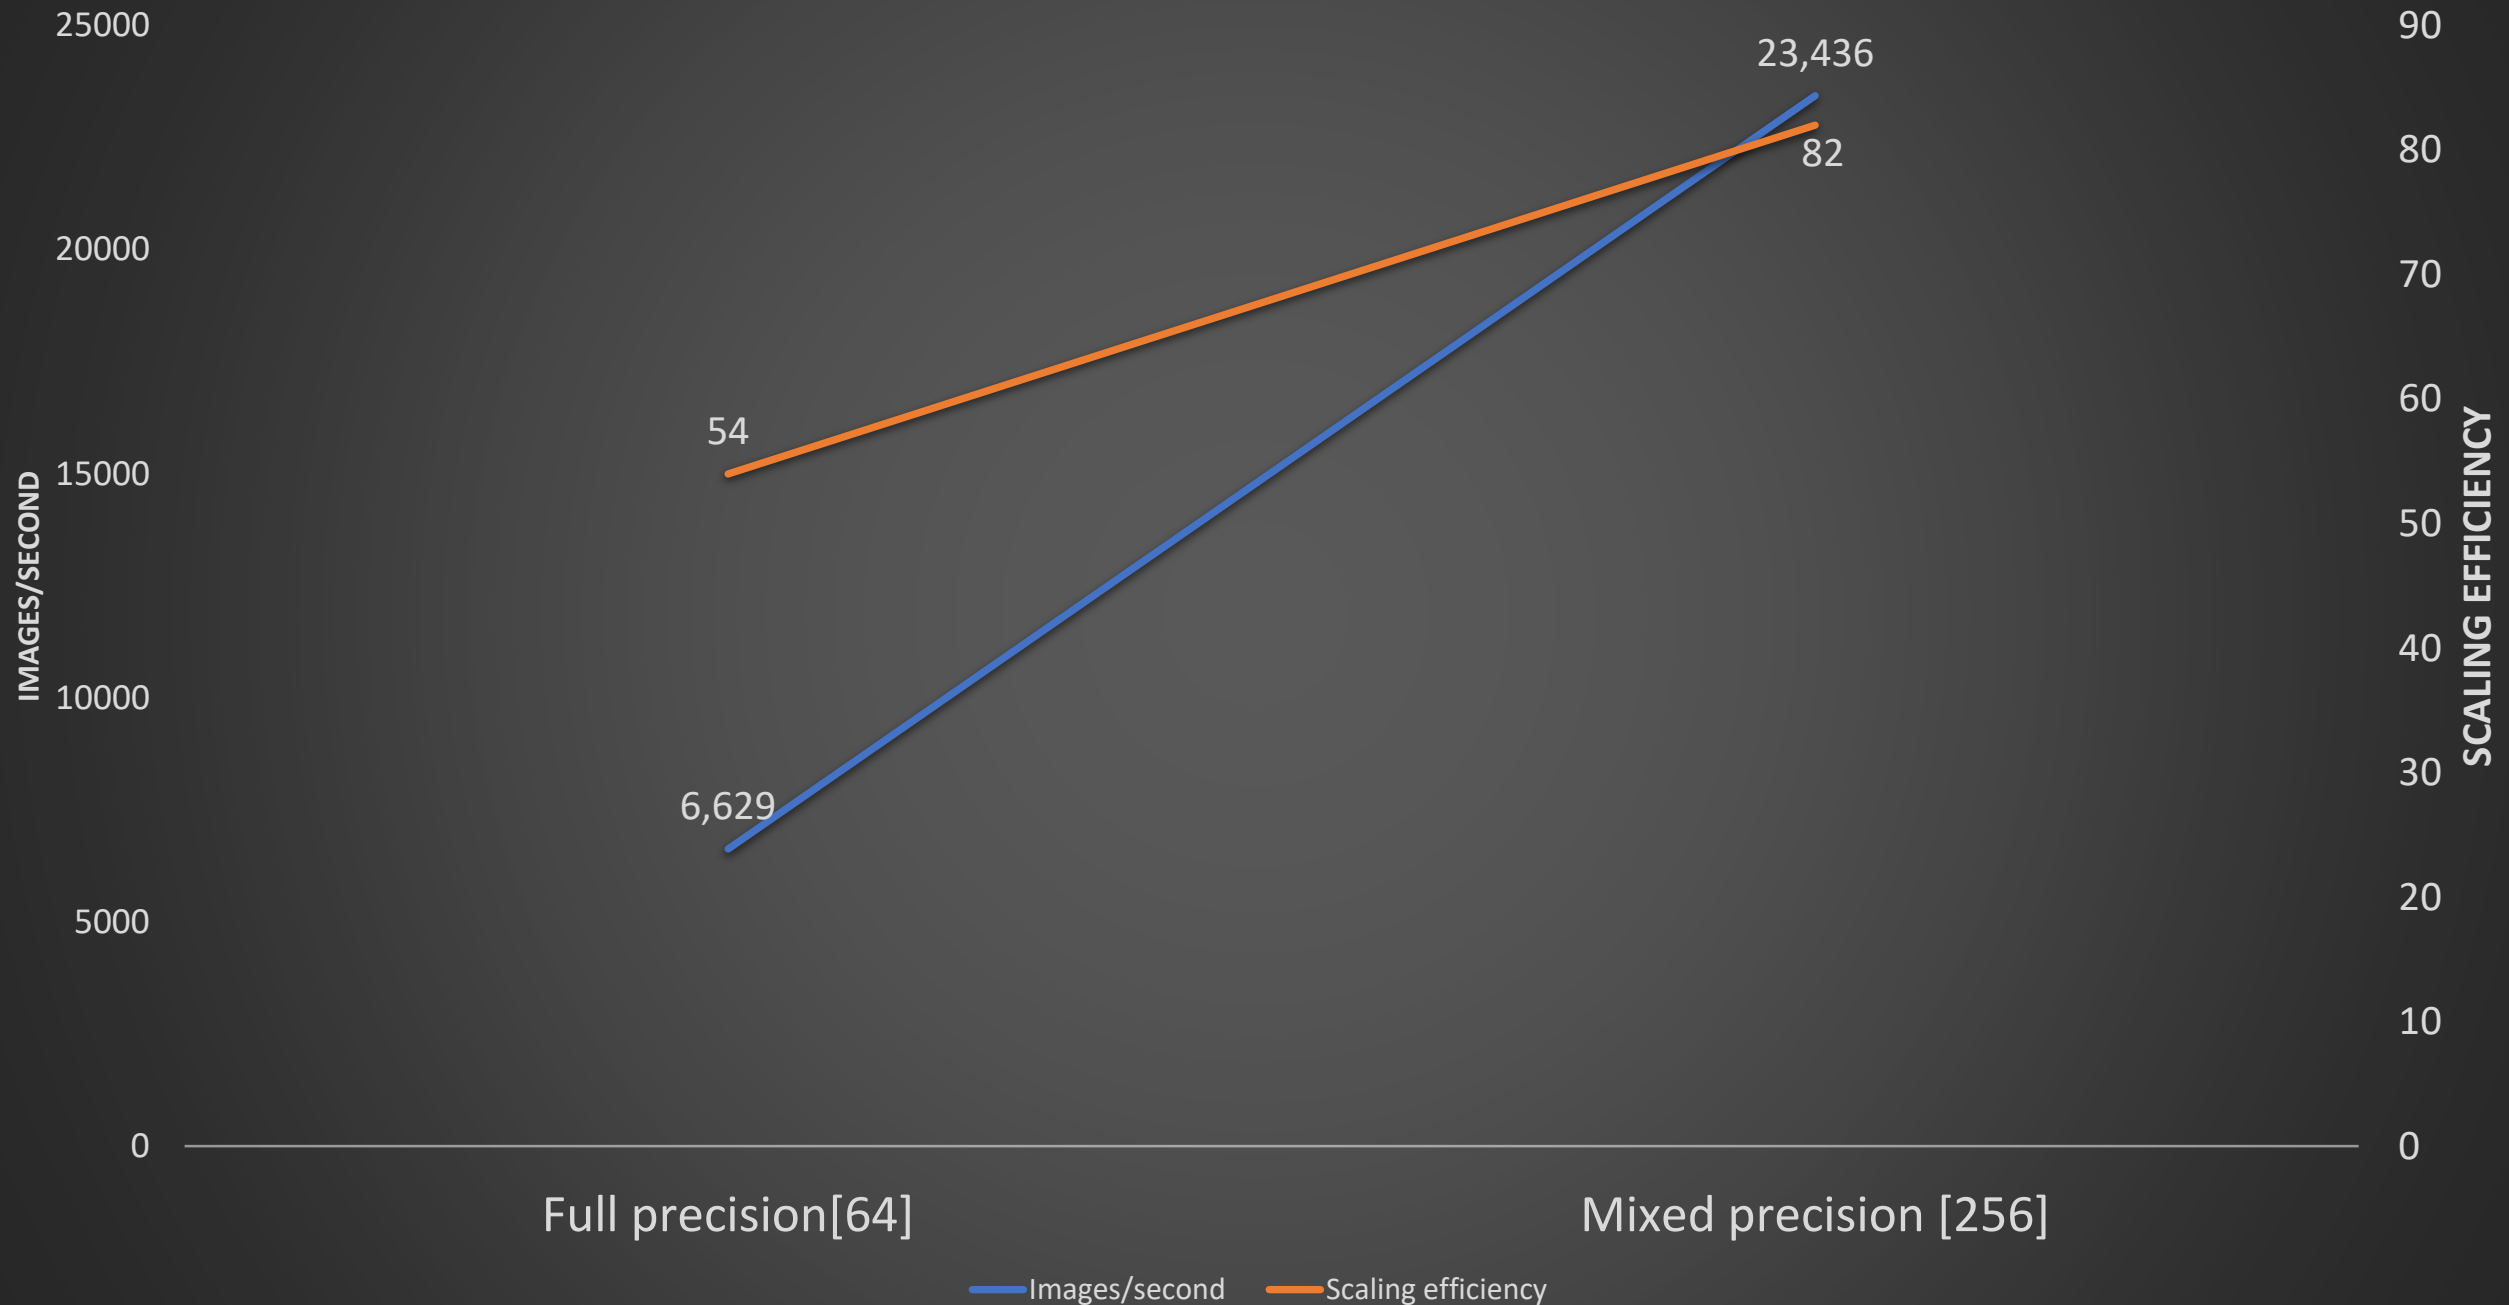
Data Transfer

Time it takes to
process batch on
GPU

Time it takes to
transfer weights
between GPUs



ResNet50 Full Precision vs Mixed Precision [32 V100s]





Effects of Storage

Using **ResNet50** across three frameworks [**PyTorch**, **TensorFlow**, **Keras**]


Using **real and synthetic** data. Real data on local, NFS and Blob storage

Batch size remains **64** across all configurations


Uses **V100** GPUs

Experiments


Distributed training with NFS



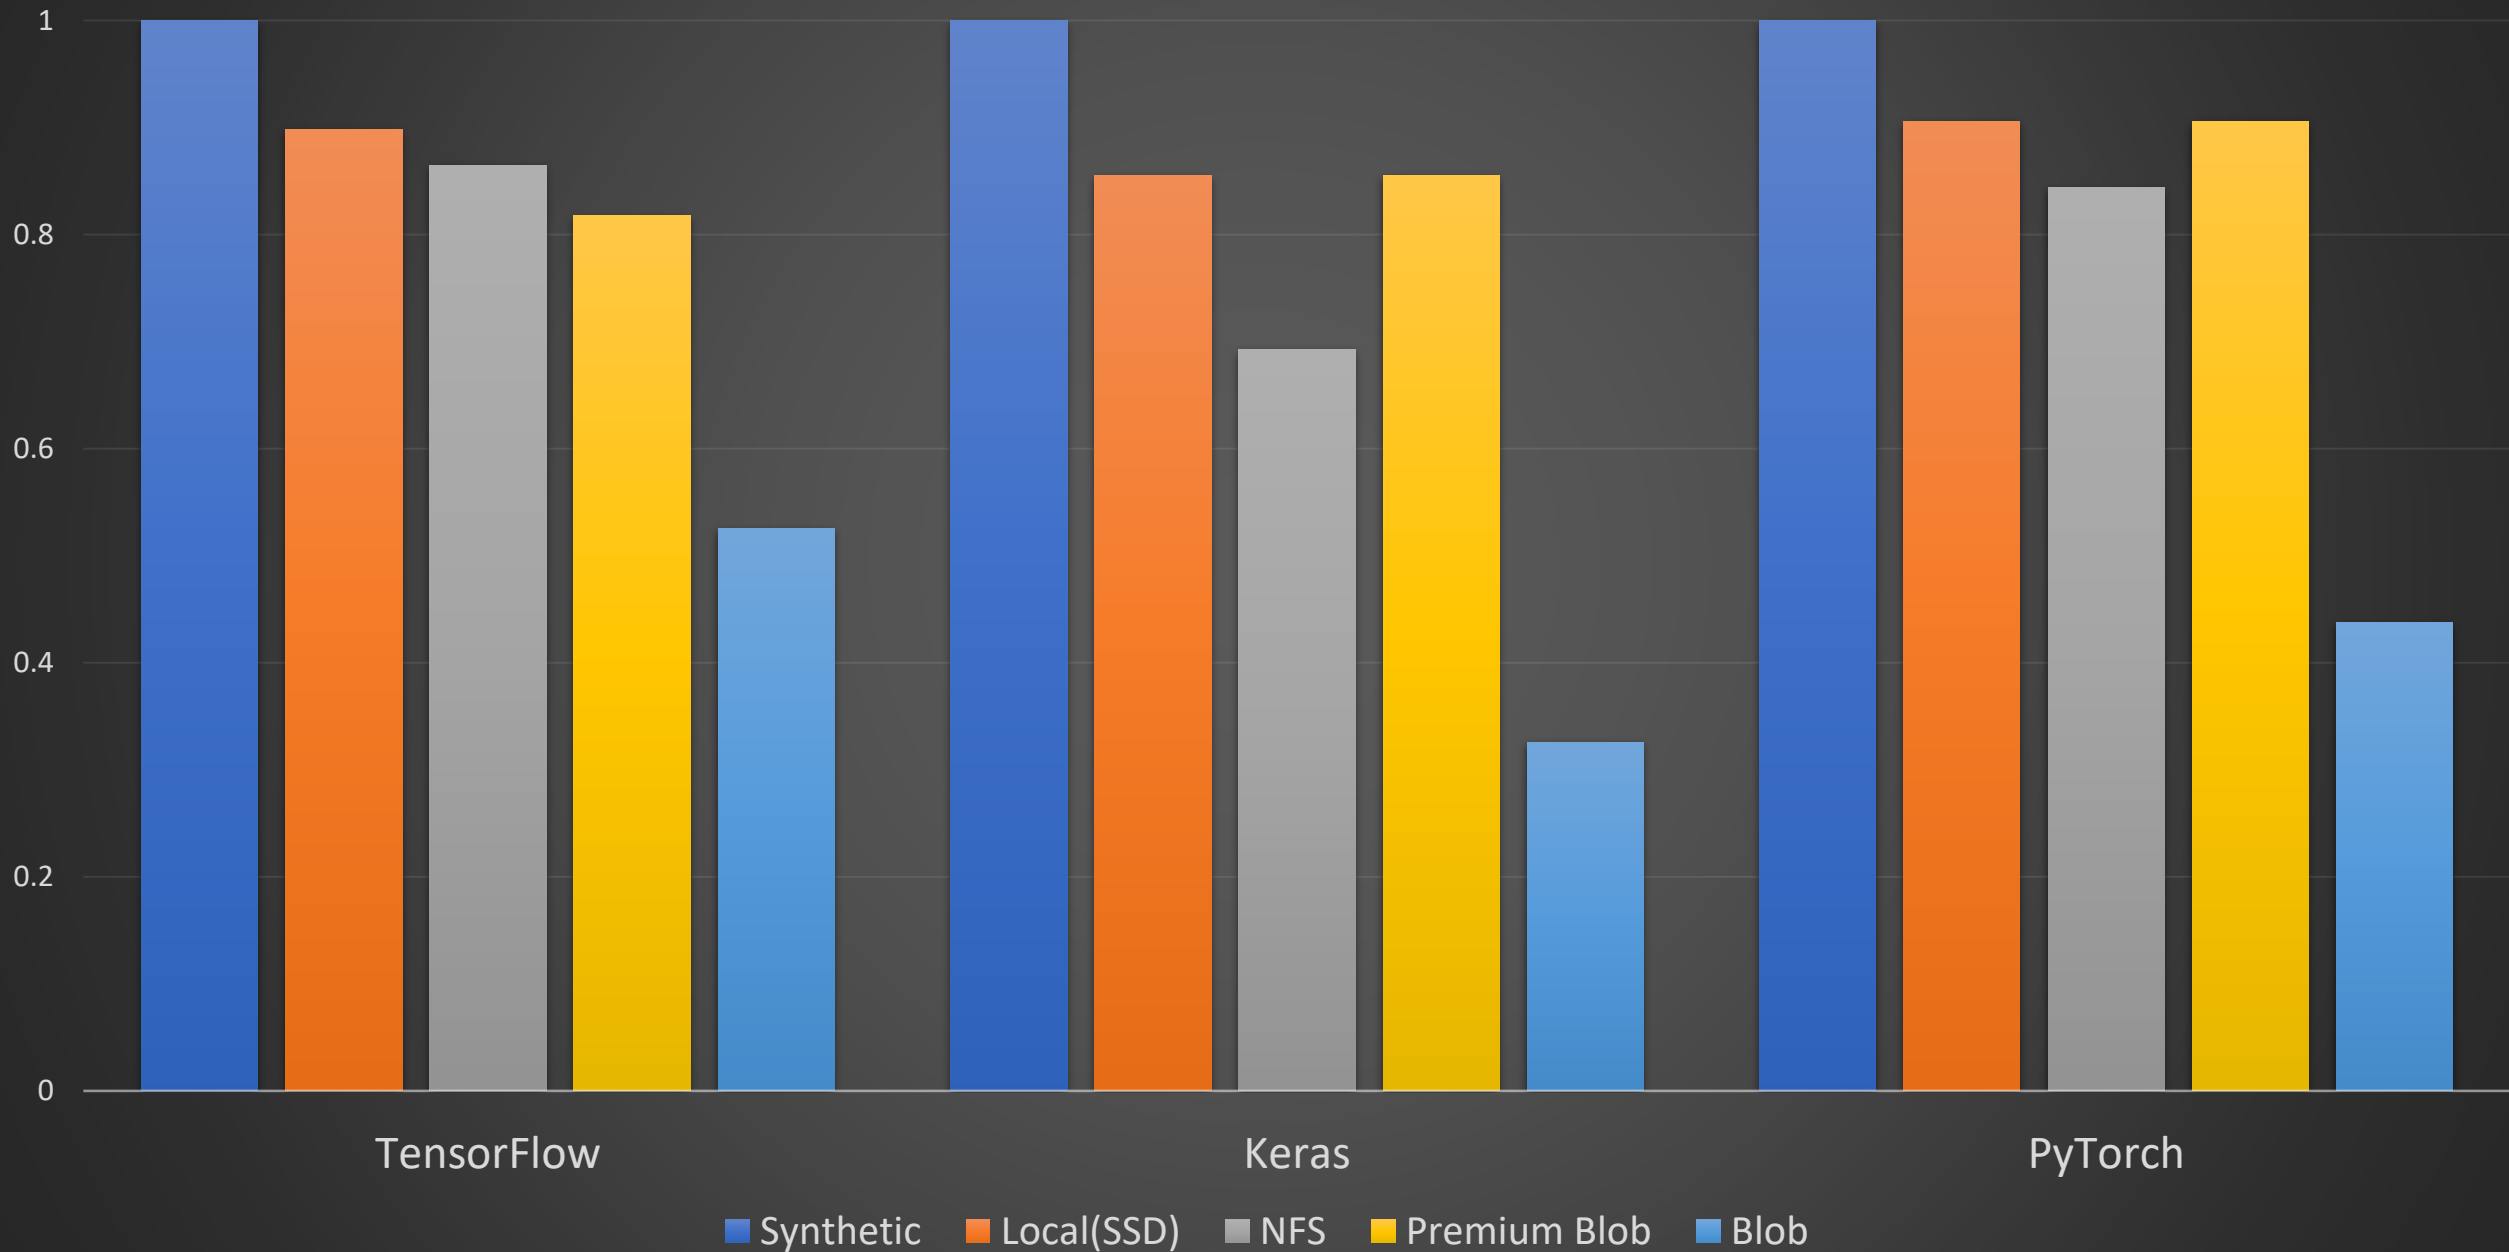
Distributed training with blob storage



Distributed training with local storage



ResNet50 - Relative performance across storage



Data Loaders and Preprocessors



Keras Data Loader

Simple with no parameters for buffering and parallelizing



PyTorch Data Loader

Specify number of workers with *num_workers*

TensorFlow

Highly configurable

Many options : buffer, shuffle, cache
and shard

Daunting and easy to get wrong

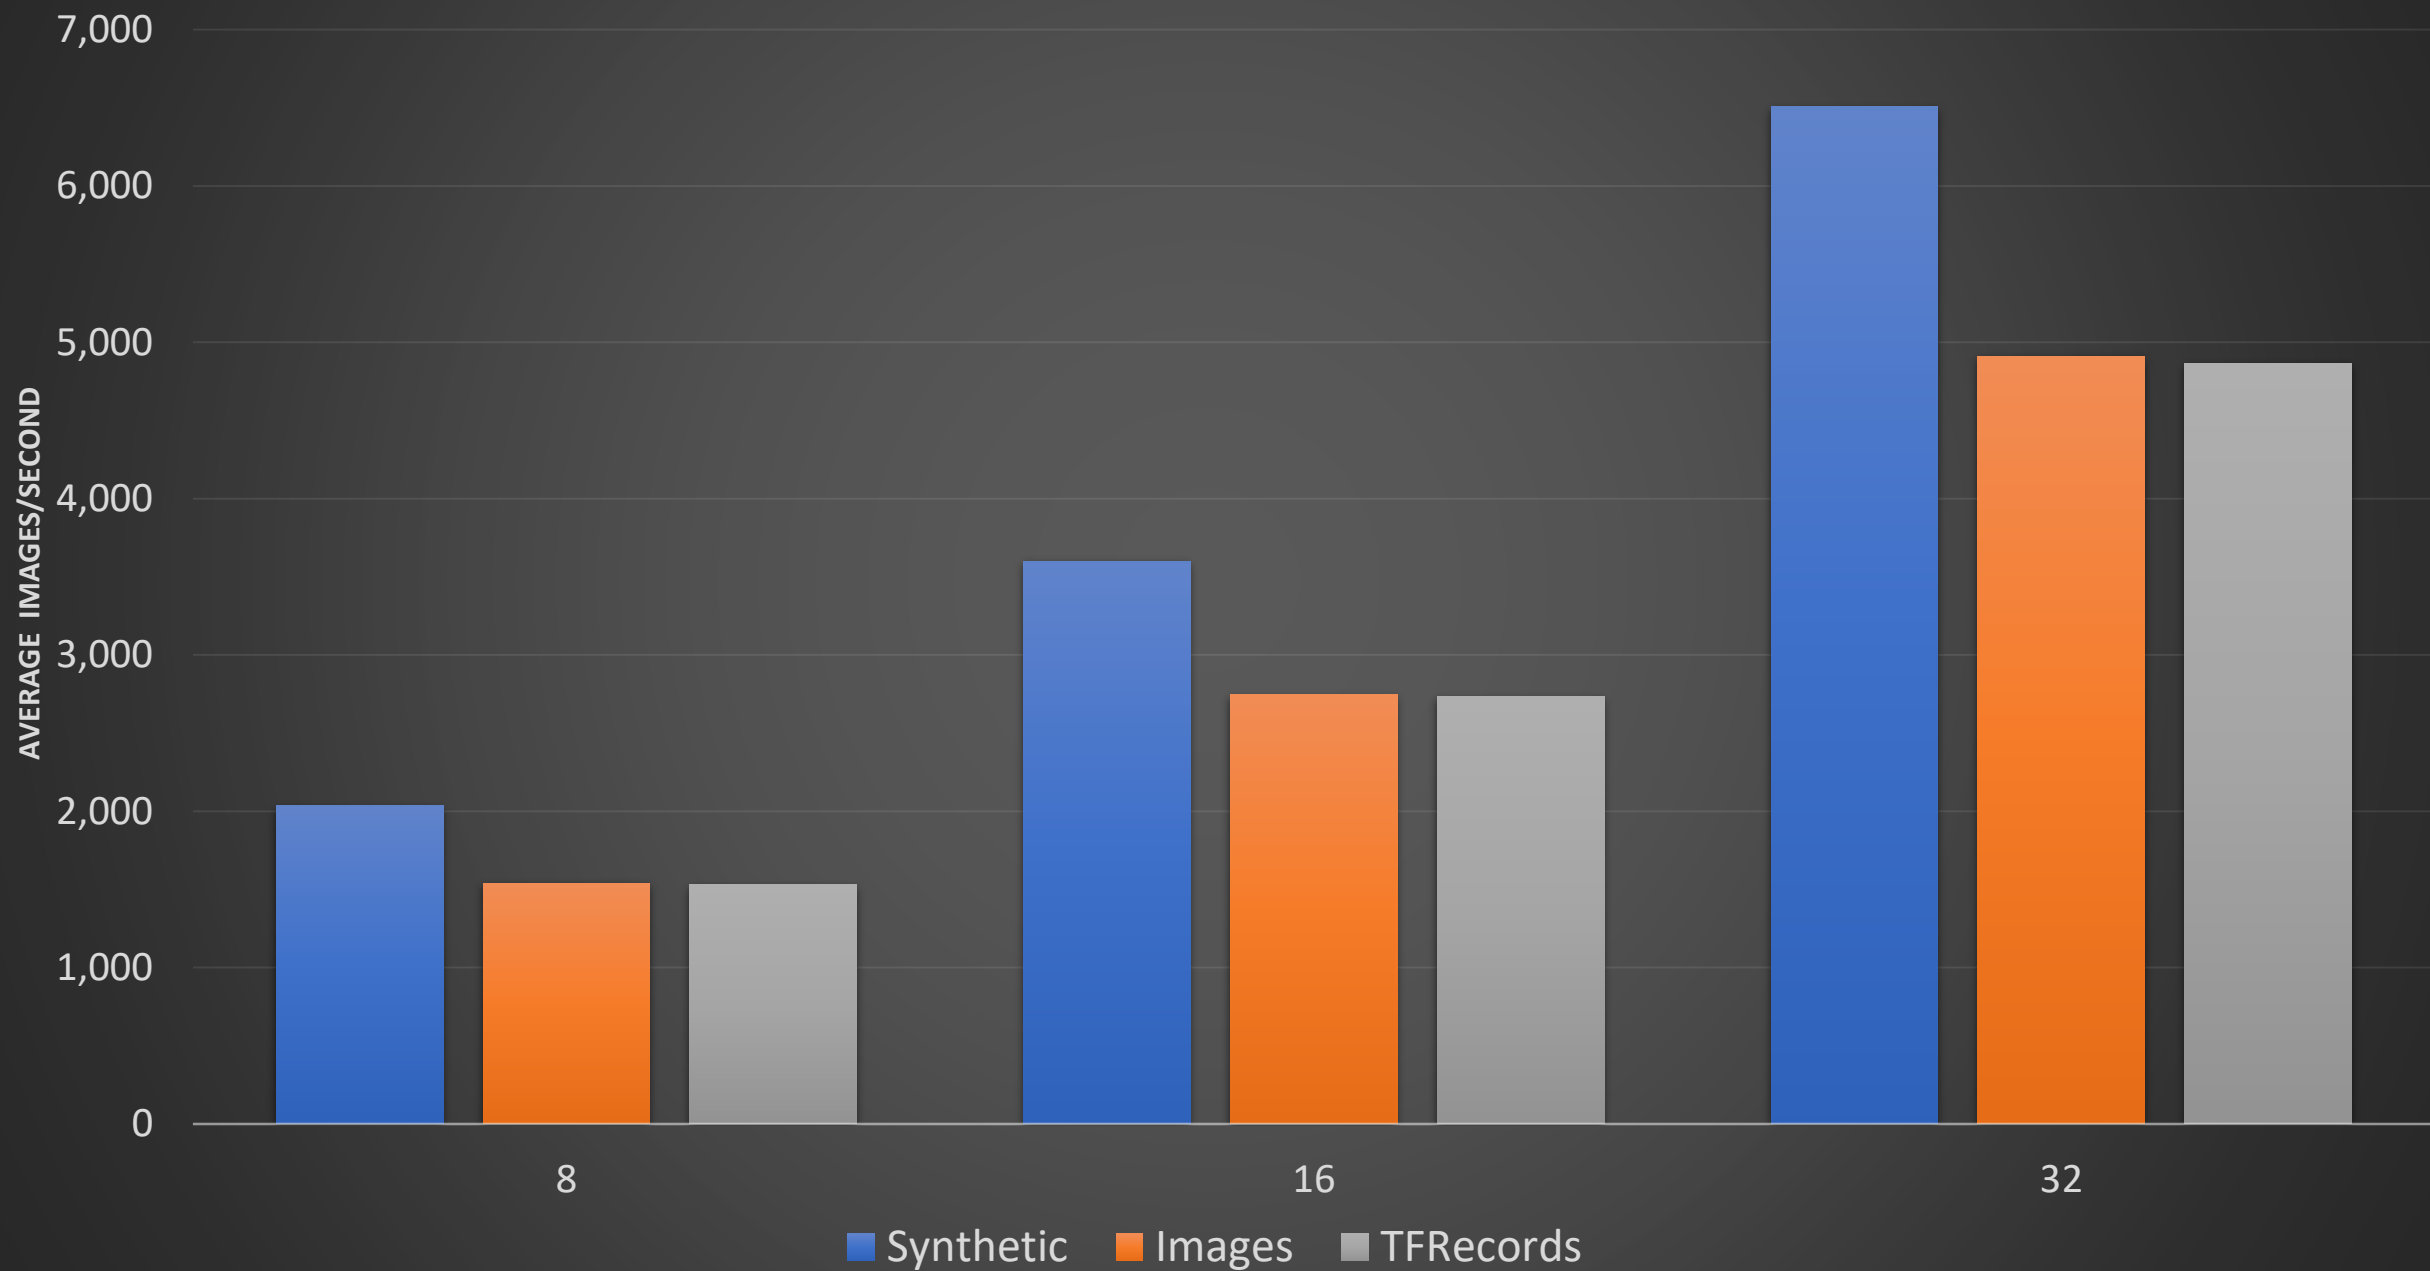
Effects of Data Type

A vinyl record is shown in the background, slightly out of focus. The center label is visible, with the word 'STEREO' printed on it. The record's grooves are clearly visible, radiating from the center.

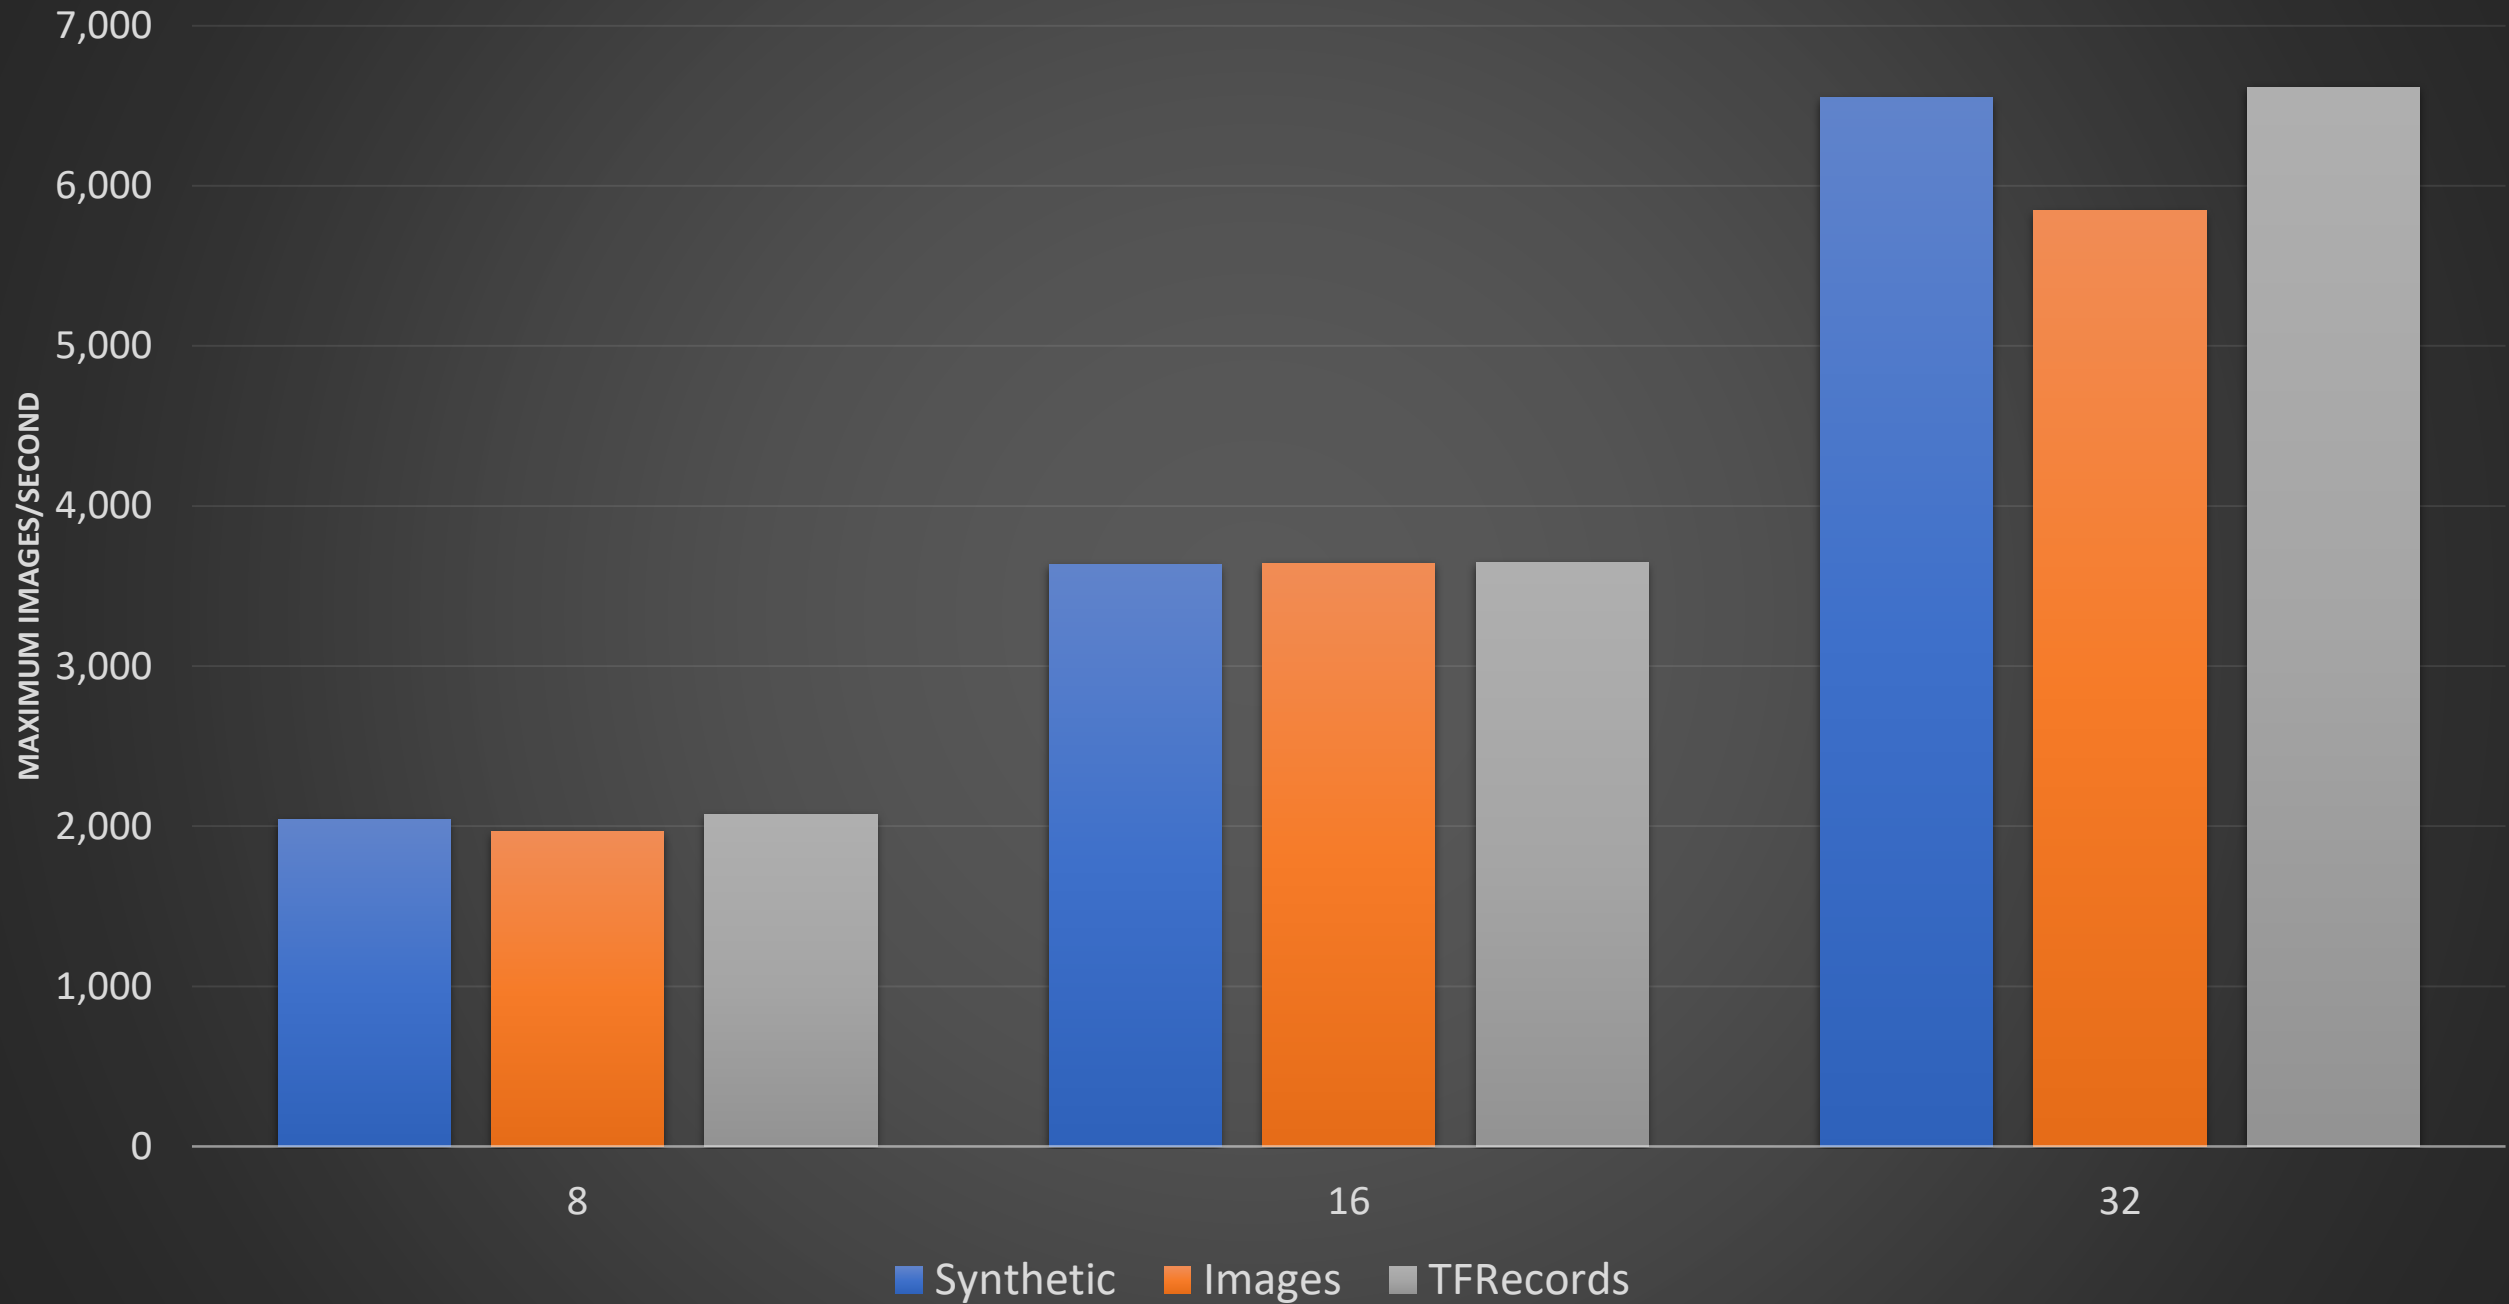
TensorFlow Records

- Binary data format created for TensorFlow – Recommended format for TensorFlow
- Can aggregate number of examples to smaller number of TFRecords – efficient for transferring and reading in the cloud
- Have to export data to format - Has to be tailored to use case

ResNet50 – Data Type Performance [Average]



ResNet50 – Data Format Performance [Maximum]



Things not discussed

Asynchronous distributed training

Tradeoff between batch size and other parameters

Optimization of TensorFlow pipeline

Other data formats such as Parquet (Petastorm)

Transform libraries [augmentations]

Distributed file systems BeeGFs and other storage
GlusterFS, Lustre etc.

Models other than CNN

Summary

Do try to use enhanced networking wherever possible especially for the latest GPUs

Training small models using distributed training is not recommended

Do use TFRecords or other columnar or row based data formats

Not all data loaders are equal

Thanks
&
Questions?
