Large-scale GPU Deep Learning Platform Design and Case Analysis

Zhang Qing

Alfie Lew

AI Age Has arrived

Steam Age

- In 1760s
- The first technological revolution

Electric Age

- In 1870s
- The second technological revolution

Information Age

- In 1940s~1950s
- The third technological revolution

AI Age

- In 2012
- The fourth technological revolution

AI Application Trend

- More and More Users
  - The Internet
  - Security and Surveillance
  - Finance, health care
  - Car manufacturers
  - Robots, entertainment
- More and more application scenarios
  - Image / video analysis
  - Speech recognition
  - NLP/OCR
  - …

Smart City

Financial

Medical care

Automobile

Household

Entertainment

Deep Learning Process Flow

Deep Learning Computing Characteristics

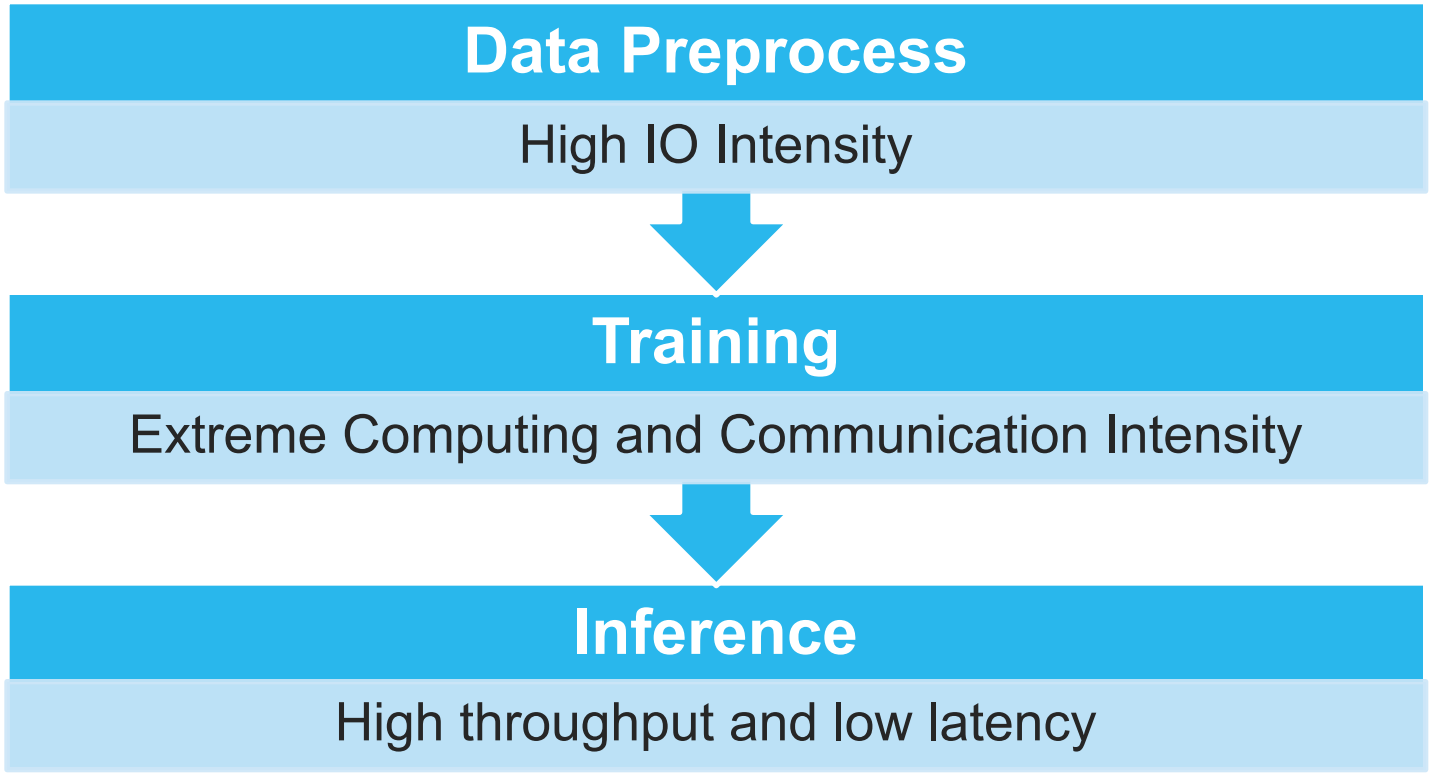

Deep Learning Computing System Trend

- Computing Mode
  - From single node to clusters
  - From local to cloud
- Data Storage
  - From Dedicated (Training and Inference) to Unified Storage
- System Management
  - Development platform
  - Production platform
  - Cloud platform
- Application mode
  - From single user to multi-user
  - From single framework to multiple frameworks

Deep Leaning Challenges

Obtaining large amount of labeled data and preprocessing efficiency

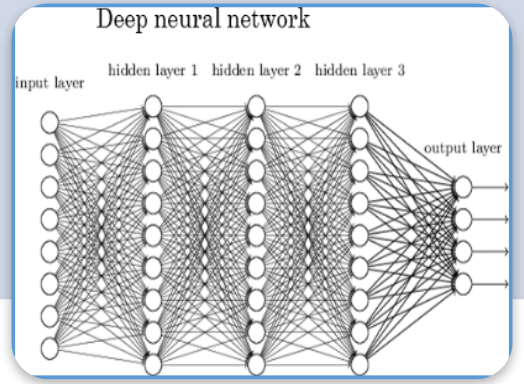

Implementing distributed parallel neural network algorithm for speed, scale, expandability

Large-scale deep learning computing platform

Architecture of Large-scale Deep Learning System

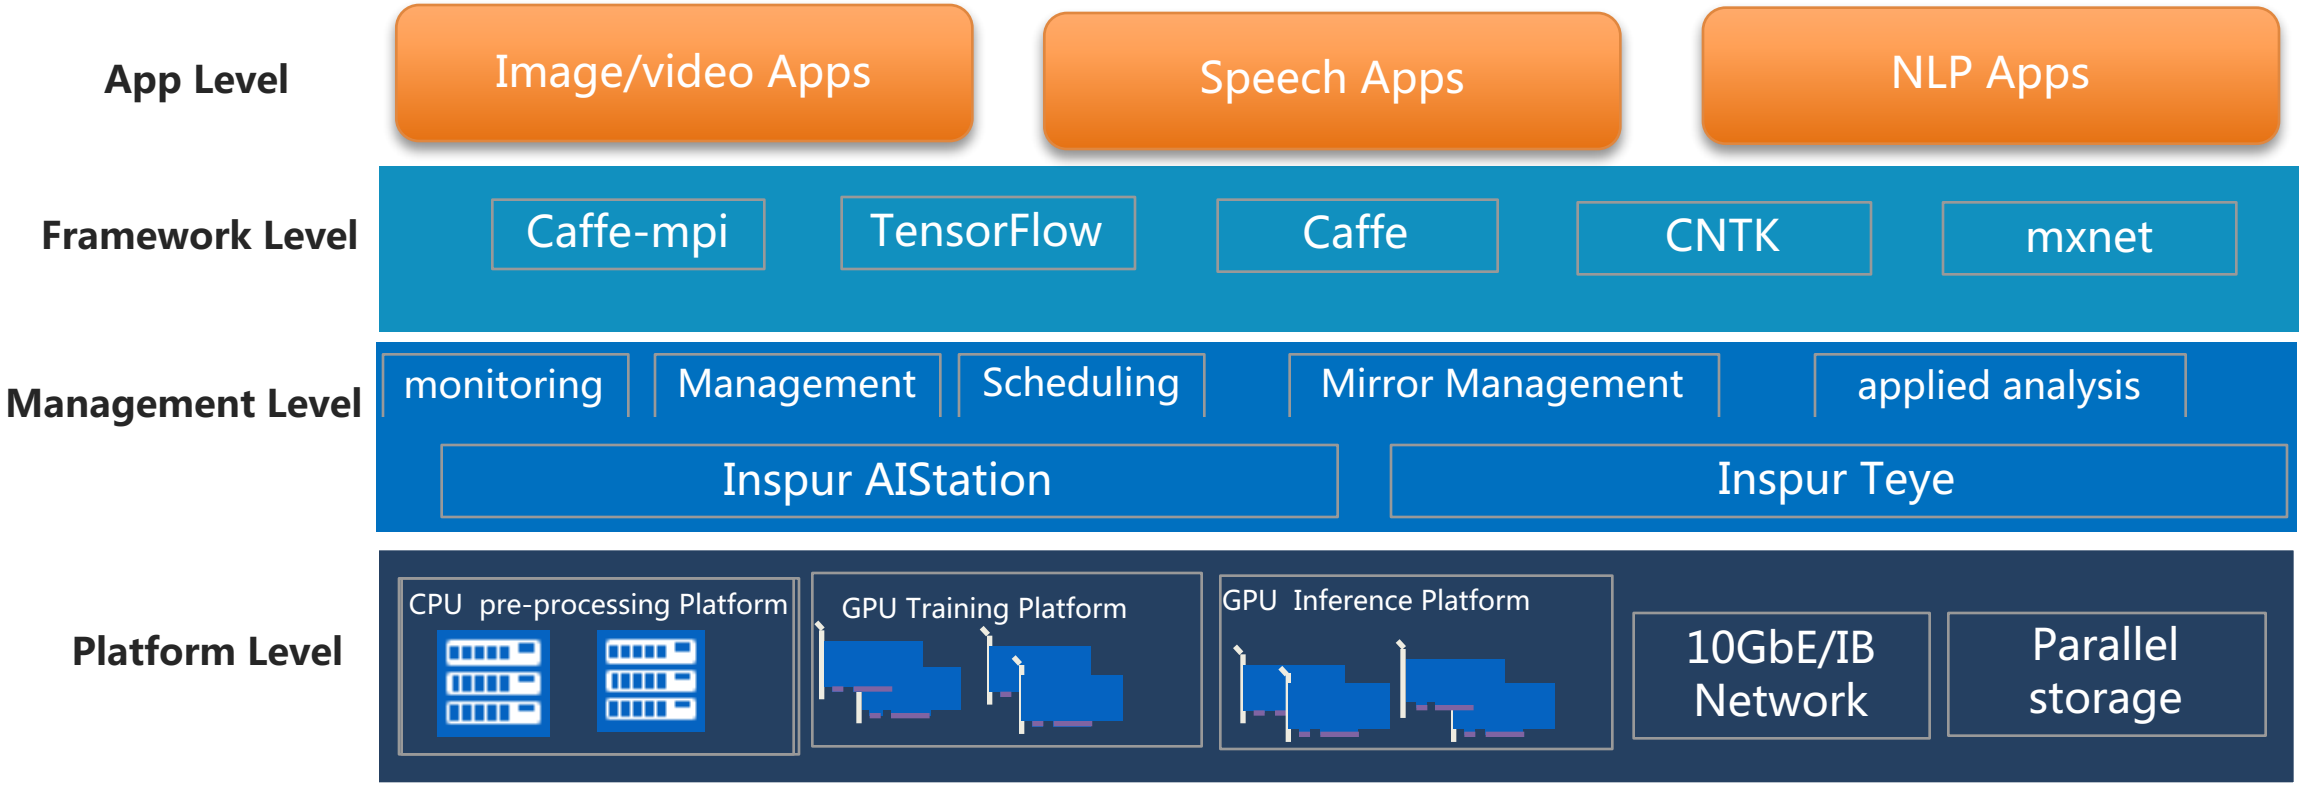

Deep Learning Challenges - Platform Level Design

IO efficiency for data pre-processing

Computing resources required for modeling, tuning and optimization

Inference speed and desired throughput for large amount of sample processing

Architecture of Large-scale Deep Learning Platform

- Computing Architecture
  - Data preprocessing platform
    - CPU cluster
  - Training platform
    - CPU+ P100/P40 GPU
    - HPC Cluster
  - Inference platform
    - CPU+P4 GPU
    - Hadoop
- Data Storage
  - Offline with Lustre
  - Online with HDFS
- Network
  - Offline with IB
  - Online with 10GbE

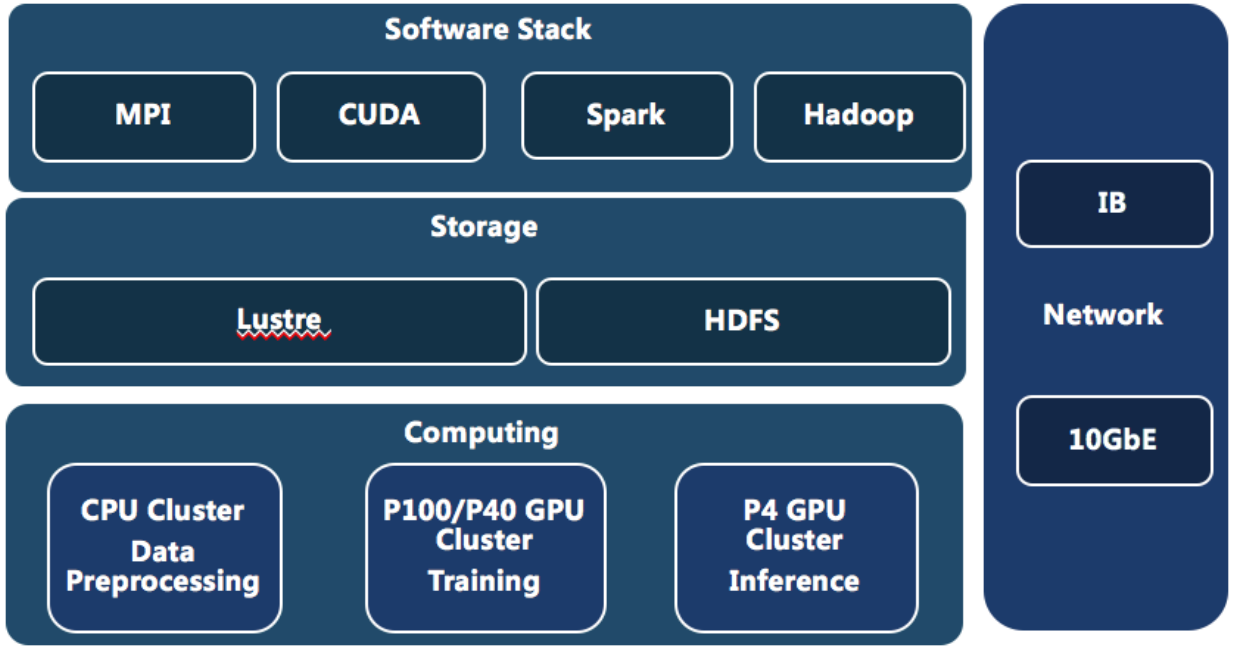

Deep Learning Challenges - Management Layer

- Managing different computing platforms and configurations/devices
- Managing different frameworks for different computing tasks
- Managing the whole system and monitoring different computing tasks

Deep Learning Management System

AIStation is a Deep Learning Cluster and training task management software, which can rapidly deploy training environment for Deep Learning, and comprehensively manage Deep Learning training tasks, providing an efficient and convenient platform for users.

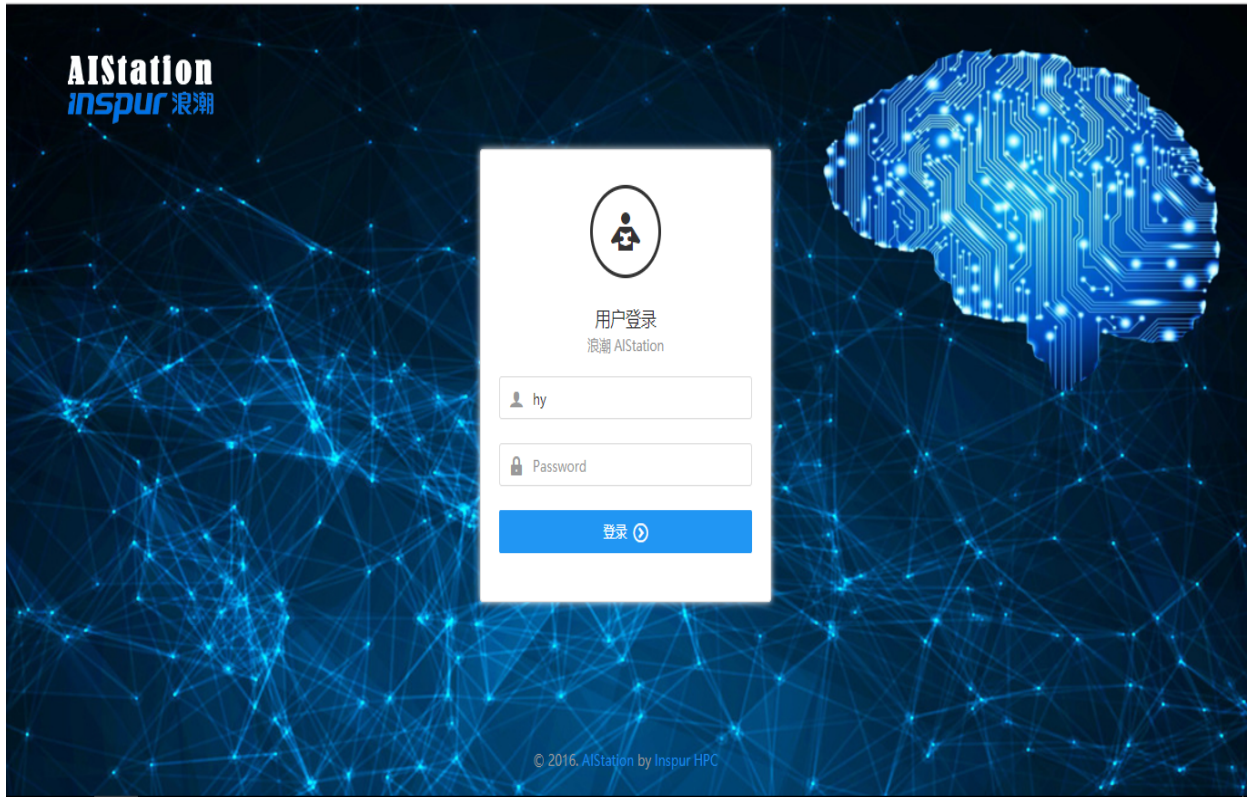

Key Functions

- Deployment of Deep-Learning environment
- Management of Deep-Learning training tasks
- GPU & CPU Monitoring
- GPU resources management and scheduling
- Cluster Statistics & Report

AIStation - Workflow

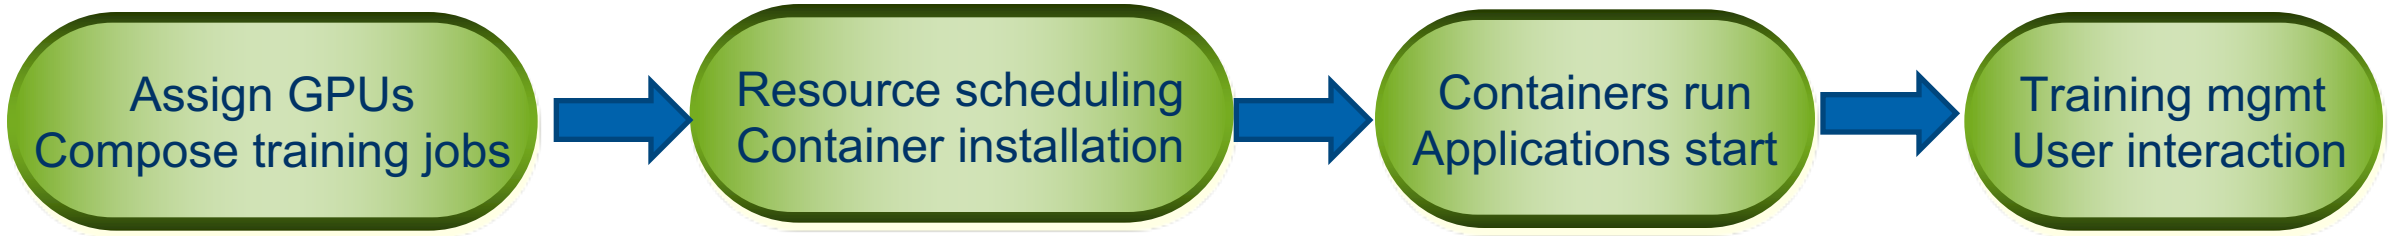

1. Resources：GPU
2. Templates：TF1
3. Images：TF/v1.0
4. Parametes：ps,ws…
5. Data：volume

1. Job starter
2. TF1.yaml

1. Run Contariners
2. Execute Job commands

1. Shell access
2. VNC access
3. Training visualization

AIStation - Integrating Deep Learning Frameworks

Integrate Deep Learning Frameworks

- Supports Multiple Deep Learning Frameworks: Caffe, TensorFlow, CNTK, etc.
- Support various models: GoogleNet, VGG, ResNet, etc.
- One-Key deployment of the Deep Learning environment
- Training jobs submit & schedule
- Training process management & visualization

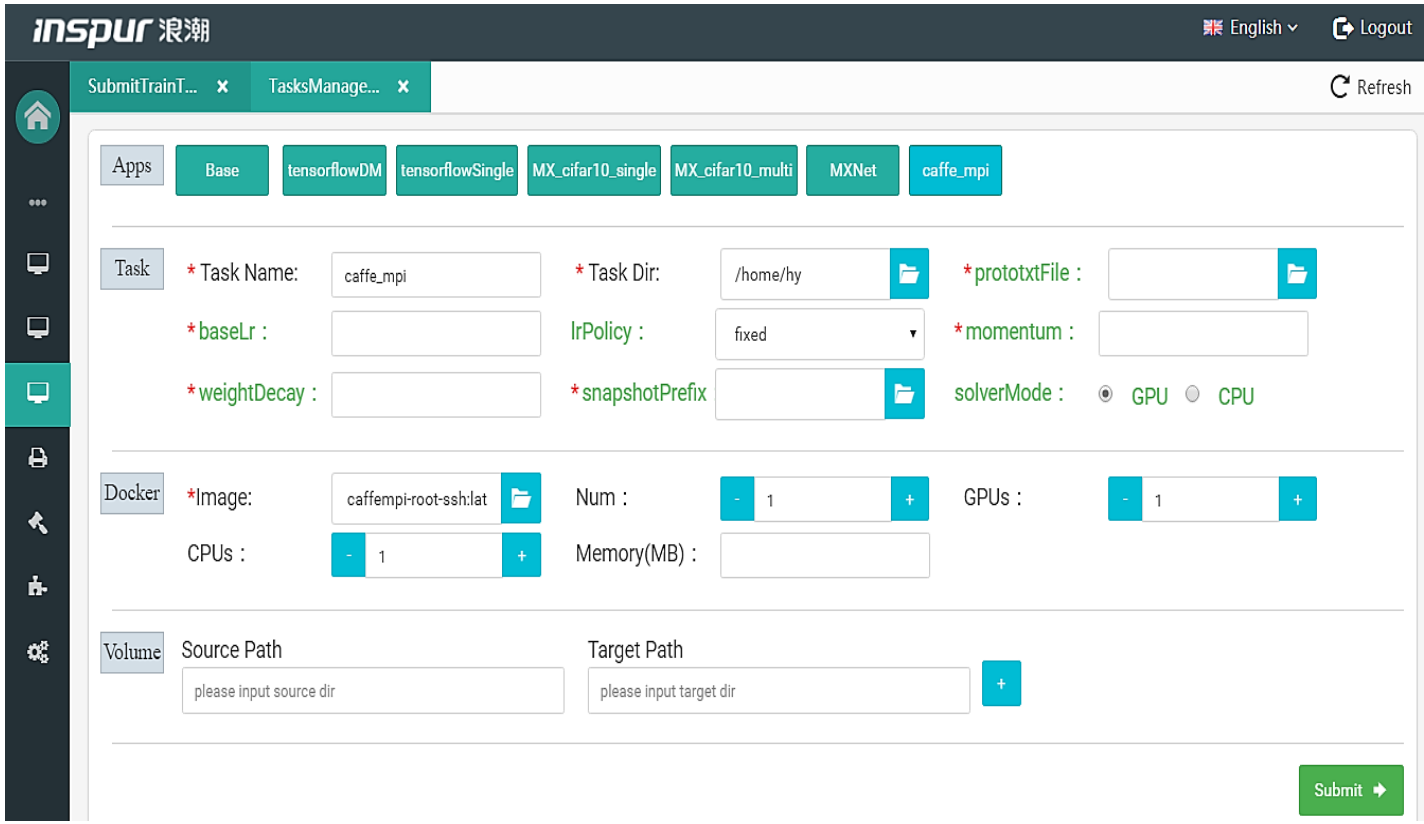

GPU Resource Utilization 20% ↑

Training Jobs Throughput 30% ↑

Teye : Application Optimization Analysis Tool

- Analyzing the bottleneck and characteristics of Applicatation
  - GPU driver data：clock，ECC，power
  - GPU runtime data：memory util，memory copy，cache，SP/DP Gflops
  - CPU runtime info： AVX，SSE，SP/DP Gflops, CPI

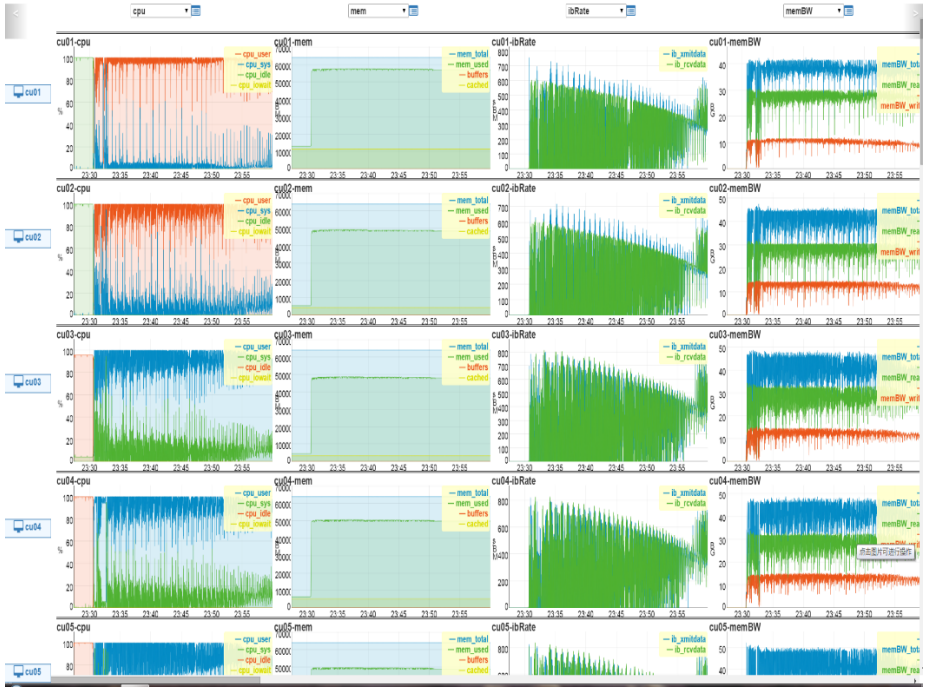

Deep Learning Challenges - Framework

- How to select from many Deep Learning Framework?

  Caffe, TensorFlow, MxNET, CNTK, Torch, Theono, DeepLearning4j, PaddlePaddle …

- What framework to use for a given scenario and model?
- Using a single framework or multiple frameworks?

A Frameworks Comparison

- Compute Platform: Inspur SR-AI Rack(16 GPUs) + AIStation+Teye (management)
- Framework: Caffe, Tensorflow, MxNet
- Model: Alexnet, Googlenet
- Performance
  - Alexnet: 4675.799 Images/s (16 GPUs/GPU = 14X) → Caffe is best
  - Googlenet：2462Images/s (16 GPUs/GPU = 13X) → MxNet is best

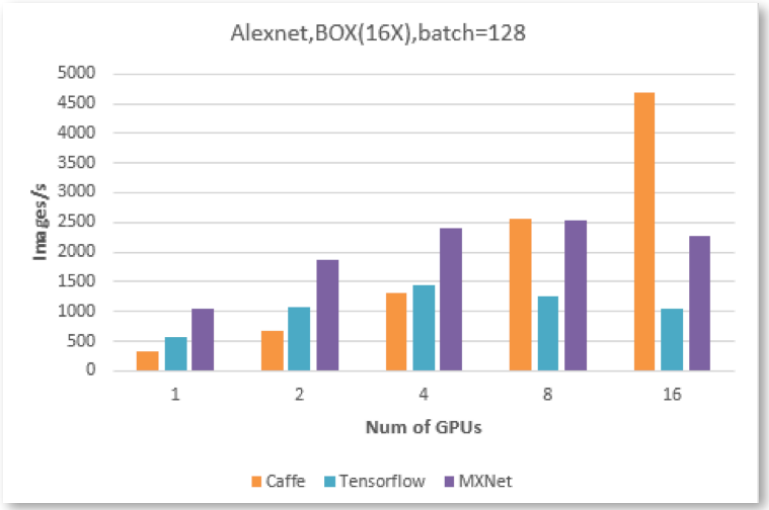

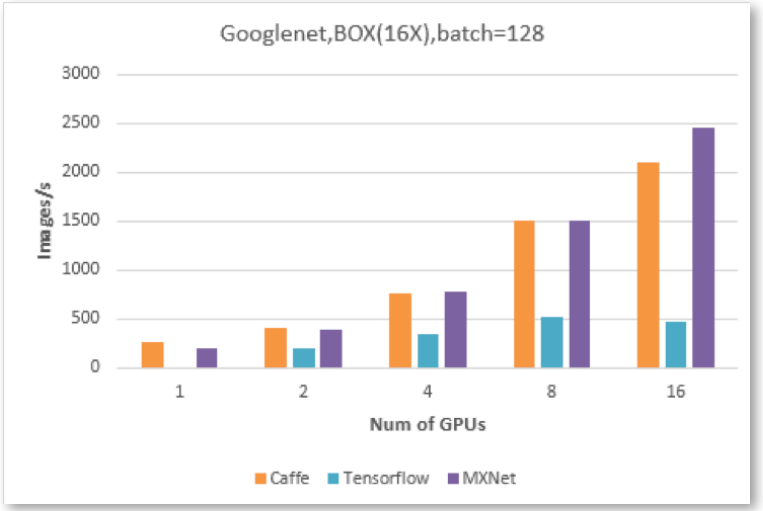

Factors to Consider when Selecting Framework

- Based on model size and complexity
- Based on different application scenarios
  - Image
  - Speech
  - NLP
- Based on data size to select distributed framework
  - Caffe-MPI
  - Tensorflow
  - MxNet

Deep Learning Challenges - Applications Layer

- How to improve the recognition accuracy?
  - Model design
  - Data pre-processing
- How to improve Training performance?
  - CUDA Programming for half Precision (pascal)
  - CUDA Programming for mixed Precision
- How to improve Inference performance?
  - CUDA Programming for Int8

Deep Learning Applications on GPU

Image Search

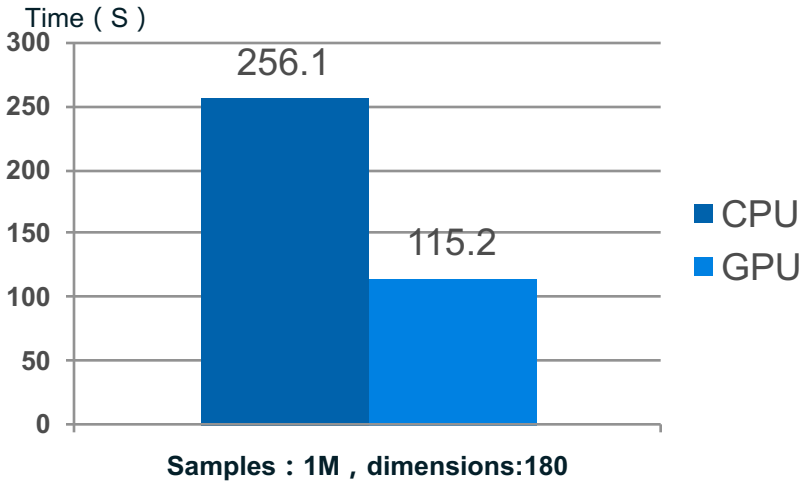

Samples : 1M , dimensions:180

Speech Training

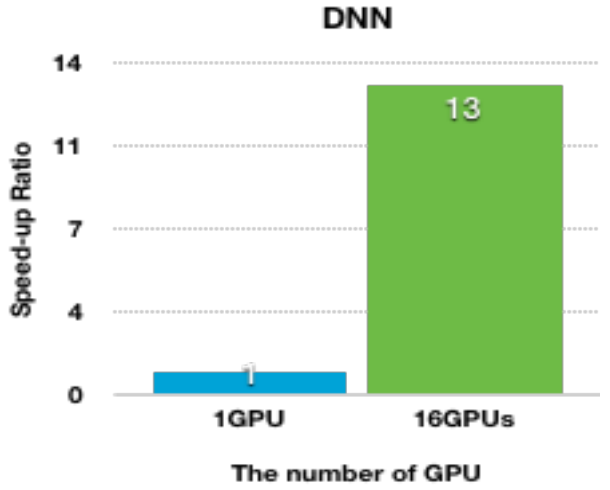

Image Training

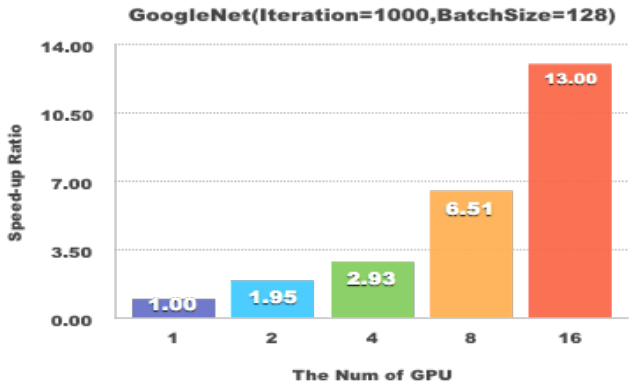

Network Security

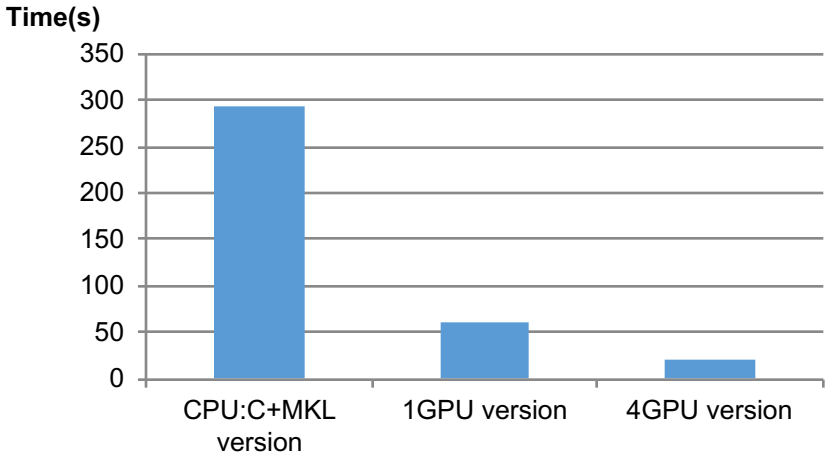

Deep Learning Platform End-to-End

Deep learning training platform

Alexnet/GoogLenet/Resnet | CNN/RNN/LSTM | Model & Algorithm

Caffe-MPI | TensorFlow | MXNET | PaddlePaddle | DL Framework

AIStation management system | T-Eye Tuning Tool | DL Management

Speech recognition

Face recognition

Video monitoring

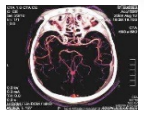
Medical imaging

Personal assistant

NF5280M4

Training Data

10G/IB Network

Model

16 Card GPU Box | 2U8 Card | 2U4 Card | 4U4 Card

Terminal

GPU AI Cloud

P8000 Wrokstation

"Big Win ! "

"This is Daniel wu"

"Pursuit staff"

"Retinopathy"

"Have booked G6"

speech/image/video / natural language

data processing Cluster

Storage

Flash Storage

AS5600/13000

Training GPU Clustre

Inference

AI recognition processing

Inspur Deep Learning GPU Servers

2GPU Server	4GPU Server	8GPU Server	64GPU Server
NF5280M4 Inference	NF5568M4 Training	AGX-2 Training	SR-AI Rack Training

Inspur is a leading AI computing providers:
Supply >60% AI HW to CSP in China

Thank You

Visit us in Booth #911