Data Pipeline Considerations for Hyperscale AI

Santosh Rao

Senior Technical Director, AI & Data Engg

Apr 2019

NetApp
Data Driven

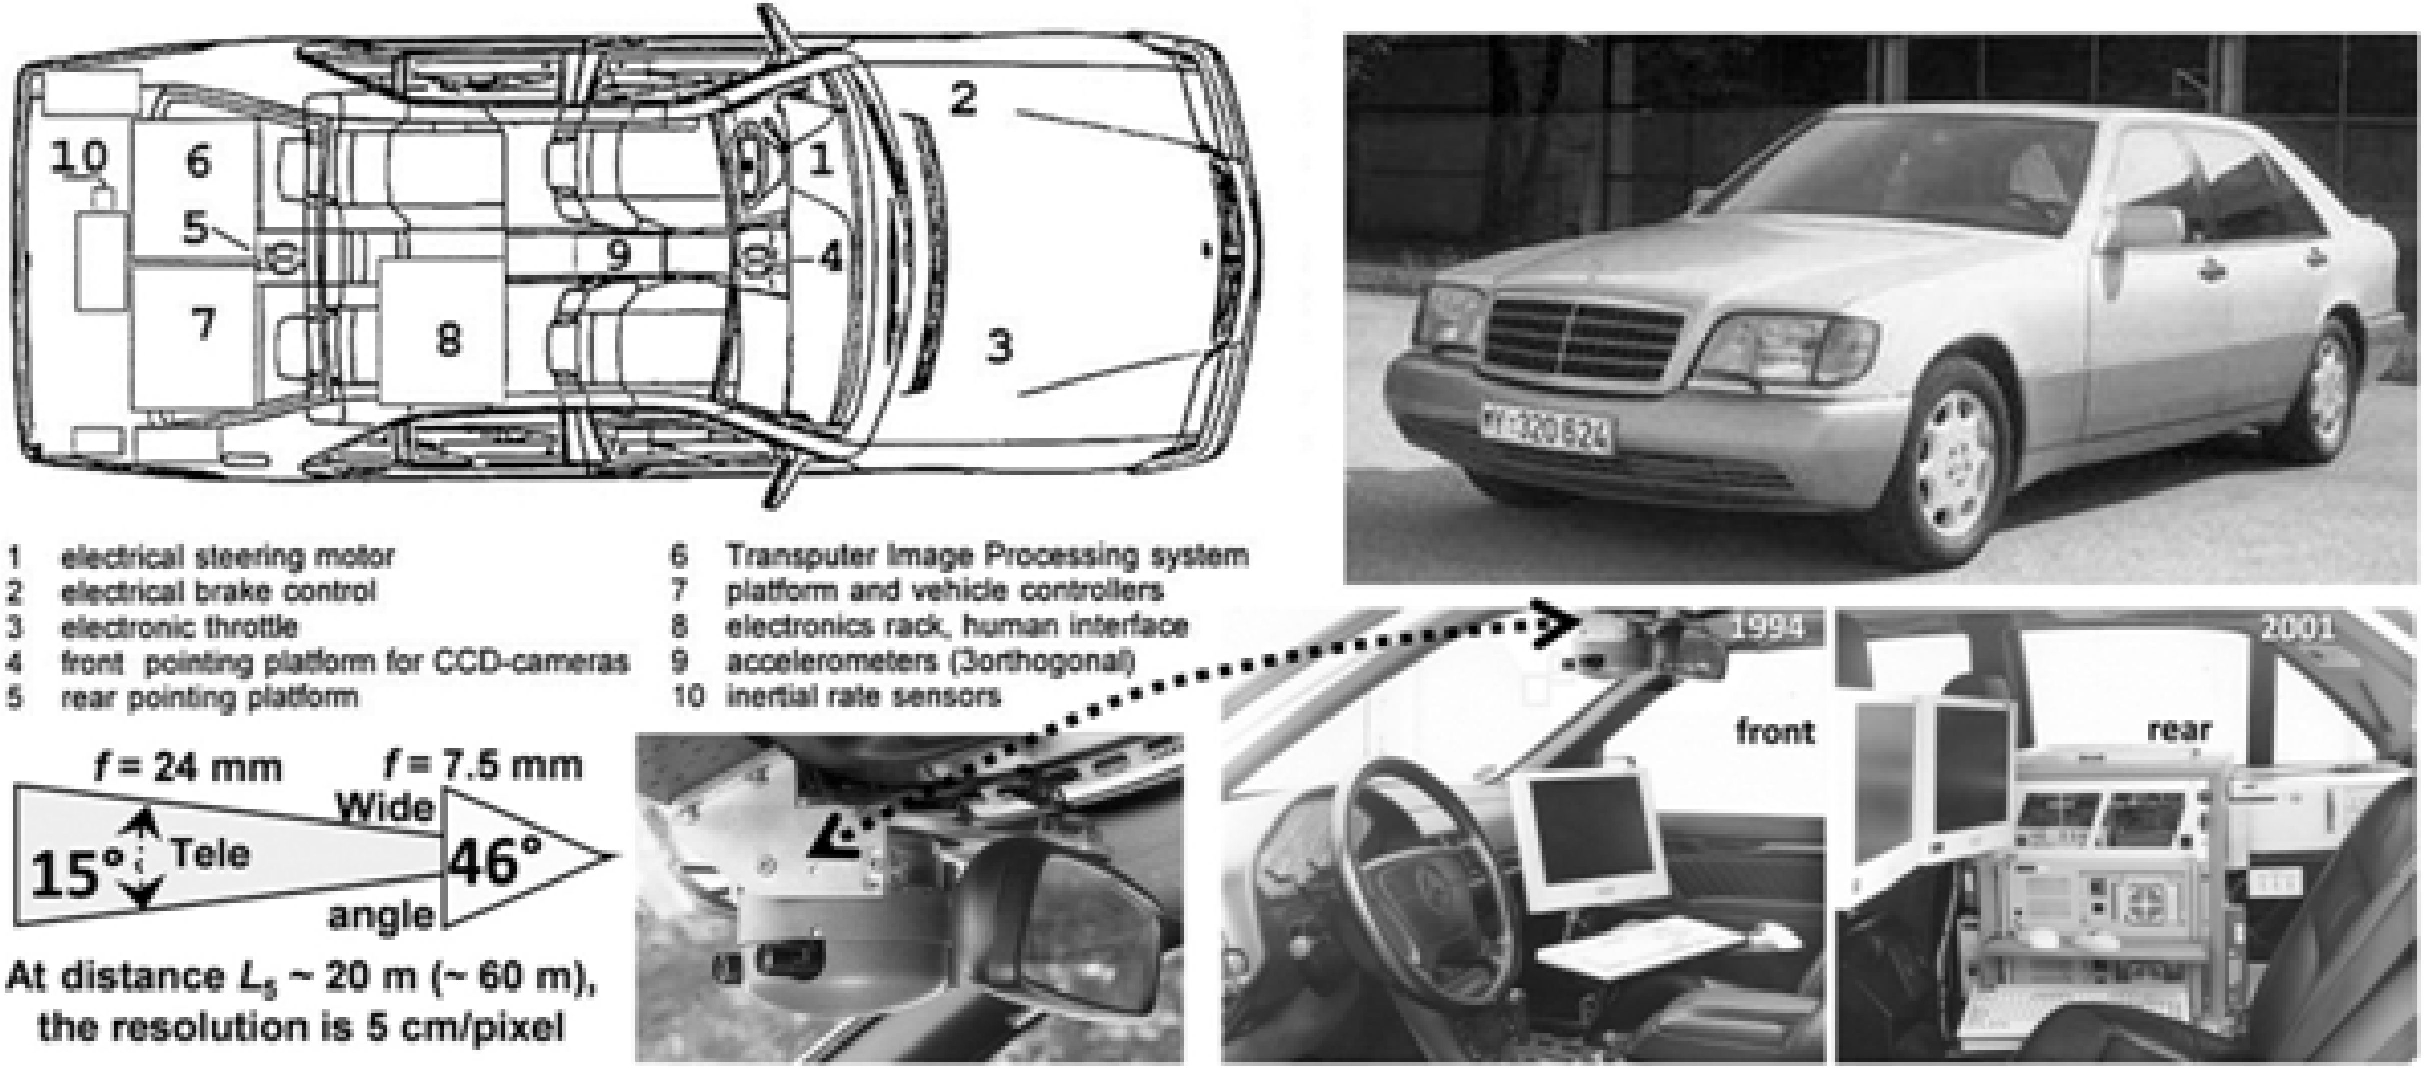

10
6
5
7
8
9
1
2
3
4
1 electrical steering motor
2 electrical brake control
3 electronic throttle
4 front pointing platform for CCD-cameras
5 rear pointing platform
6 Transputer Image Processing system
7 platform and vehicle controllers
8 electronics rack, human interface
9 accelerometers (3orthogonal)
10 inertial rate sensors
f = 24 mm
f = 7.5 mm
Wide
15°
Tele
46°
angle
At distance L_5 ~ 20 m (~ 60 m),
the resolution is 5 cm/pixel
1994
front
2001
rear


STOP
LEFT REARWARD VEHICLE CAMERA
MEDIUM RANGE VEHICLE CAMERA
RIGHT REARWARD VEHICLE CAMERA
MOTION FLOW
LANE LINES
LANE LINES
ROAD FLOW
IN-PATH OBJECTS
ROAD LIGHTS
OBJECTS
ROAD SIGNS

AI - Impactful

Rapid AI adoption world wide

300,000 lives saved by autonomous vehicles per decade

4 Billion AI-powered devices w/ voice-assistants 2019

$59.8 Billion Growth of AI software market in 2025

1 Billion AI-enabled video cameras by 2020

$10B in venture capital for AI

190% AI patents grew by over 5 yrs

71% Automation potential in Manufacturing

4th largest number of AI startups in Berlin

Source: Ihttps://techjury.net/stats-about/ai/

BDA

Designed for Processing Log Files

Evolved as a Dataware house
Big Query
Big Reporting

Limited ML,
No DL

NO RL

HPC

Designed for Particle Physics Simulations

Evolved for ML, Simulation and all kind of science experiments

MPI, ML, DL

Energy Efficiency

AI

Designed for Deep Learning

Evolved to do
Deep RL
Deep GM
Deep Sequence

Cuda,**NCCL**, ML.DL.RL

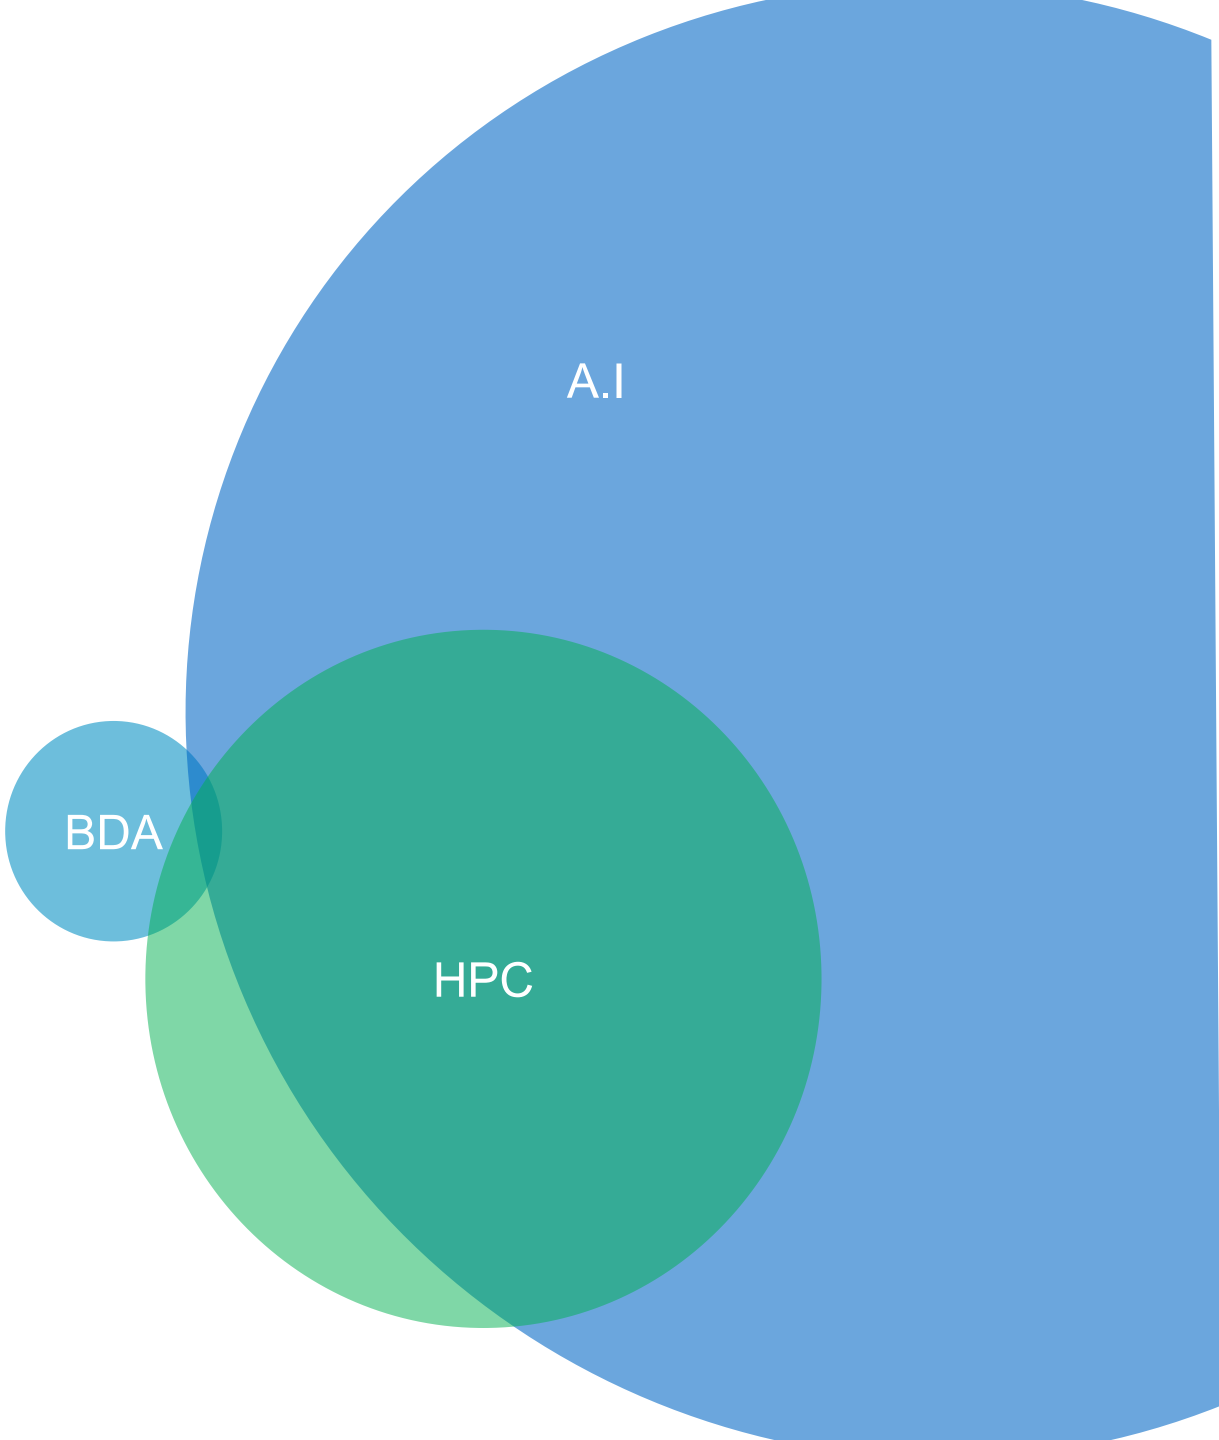

Refining Analytics doesn't lead to Artificial Intelligence

Analytics vs A.I

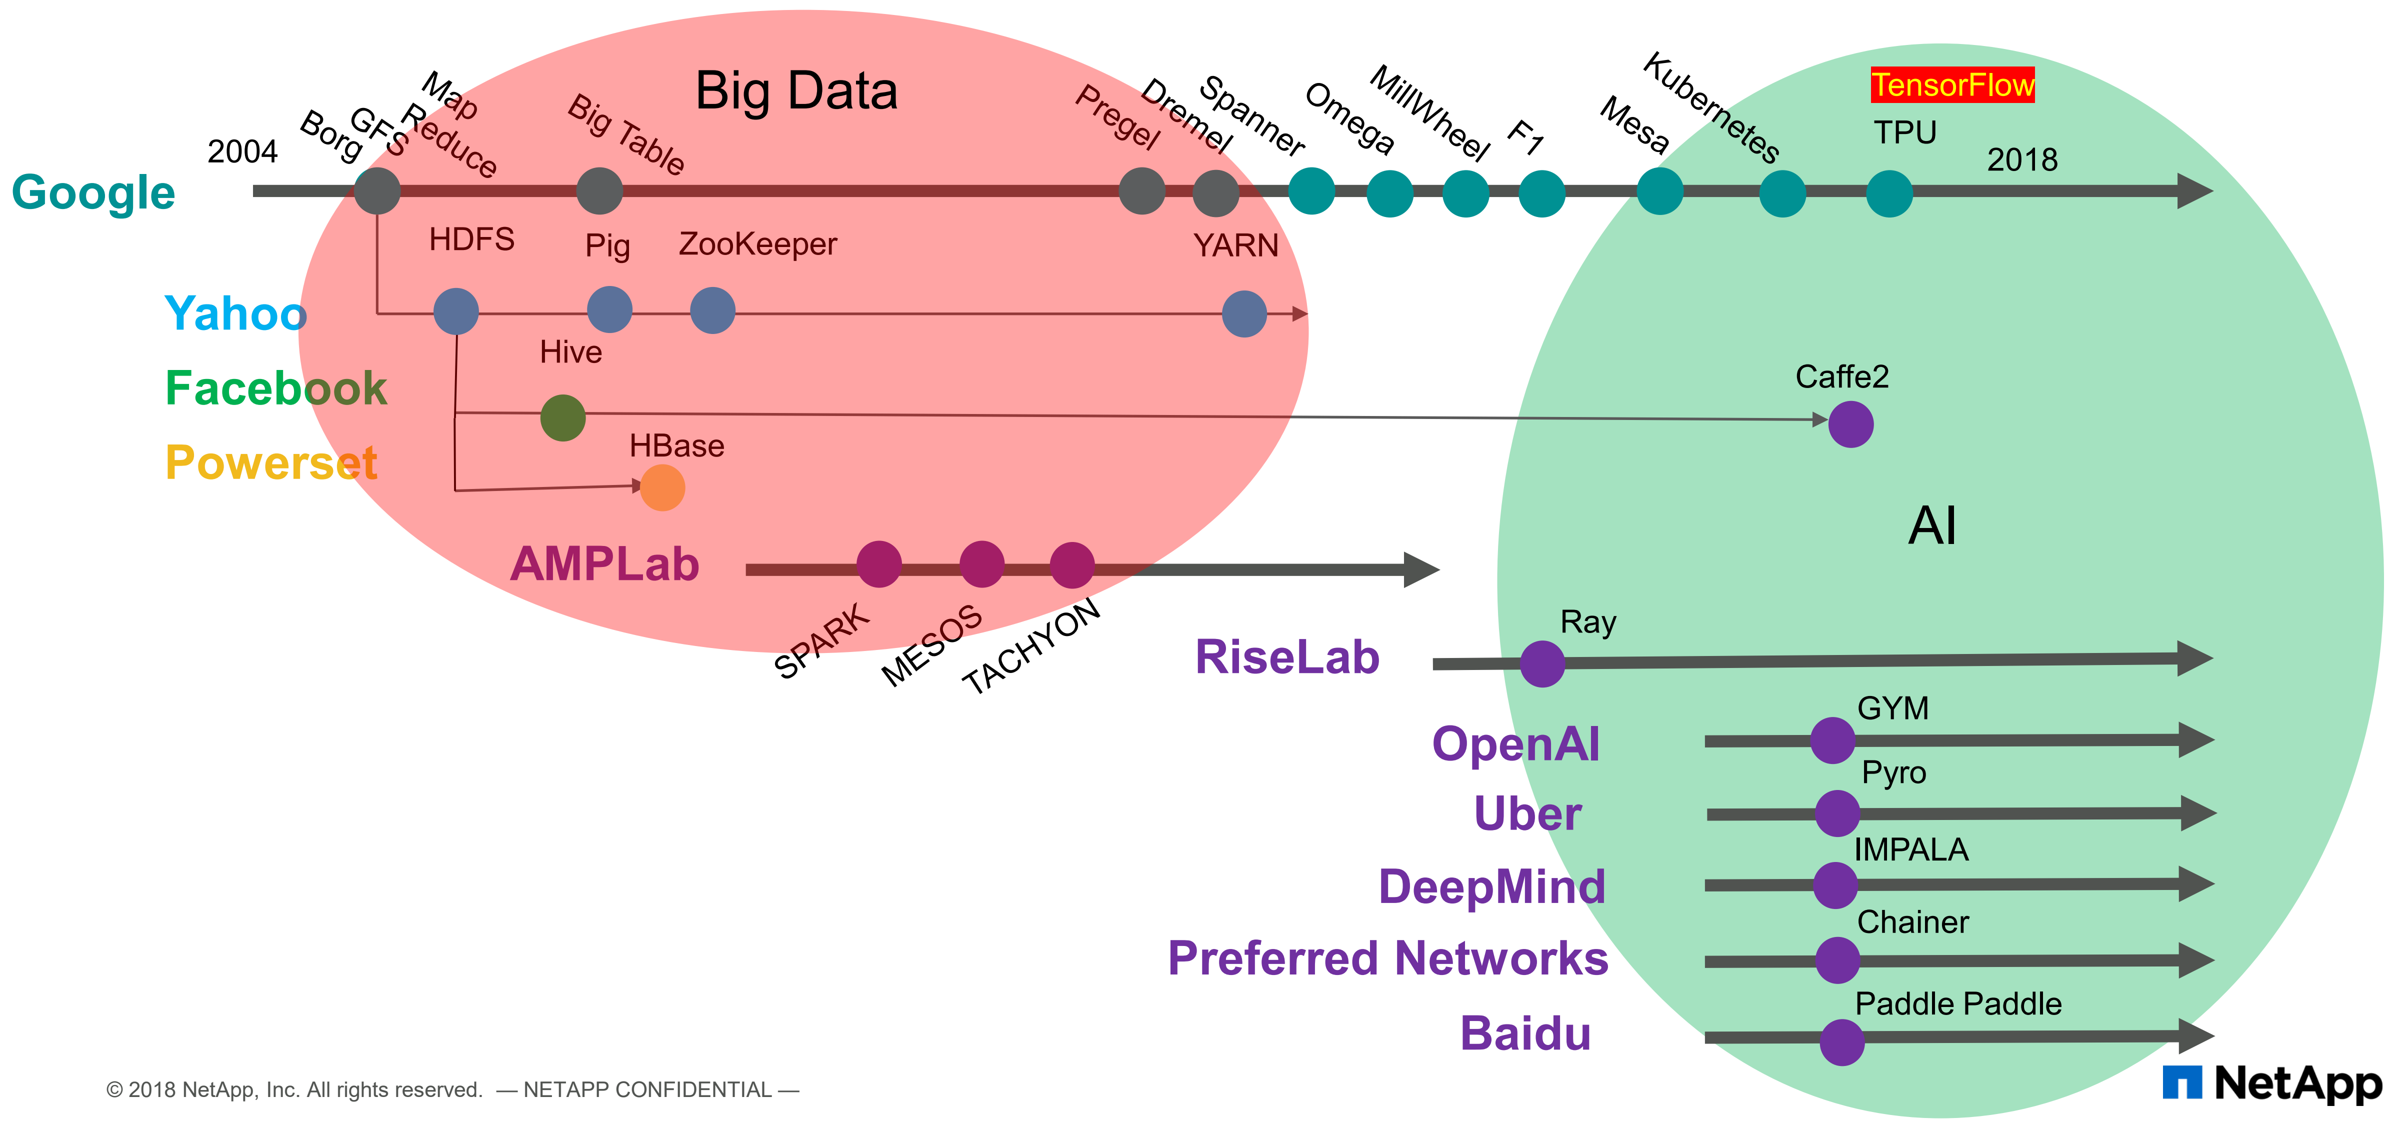

NetApp

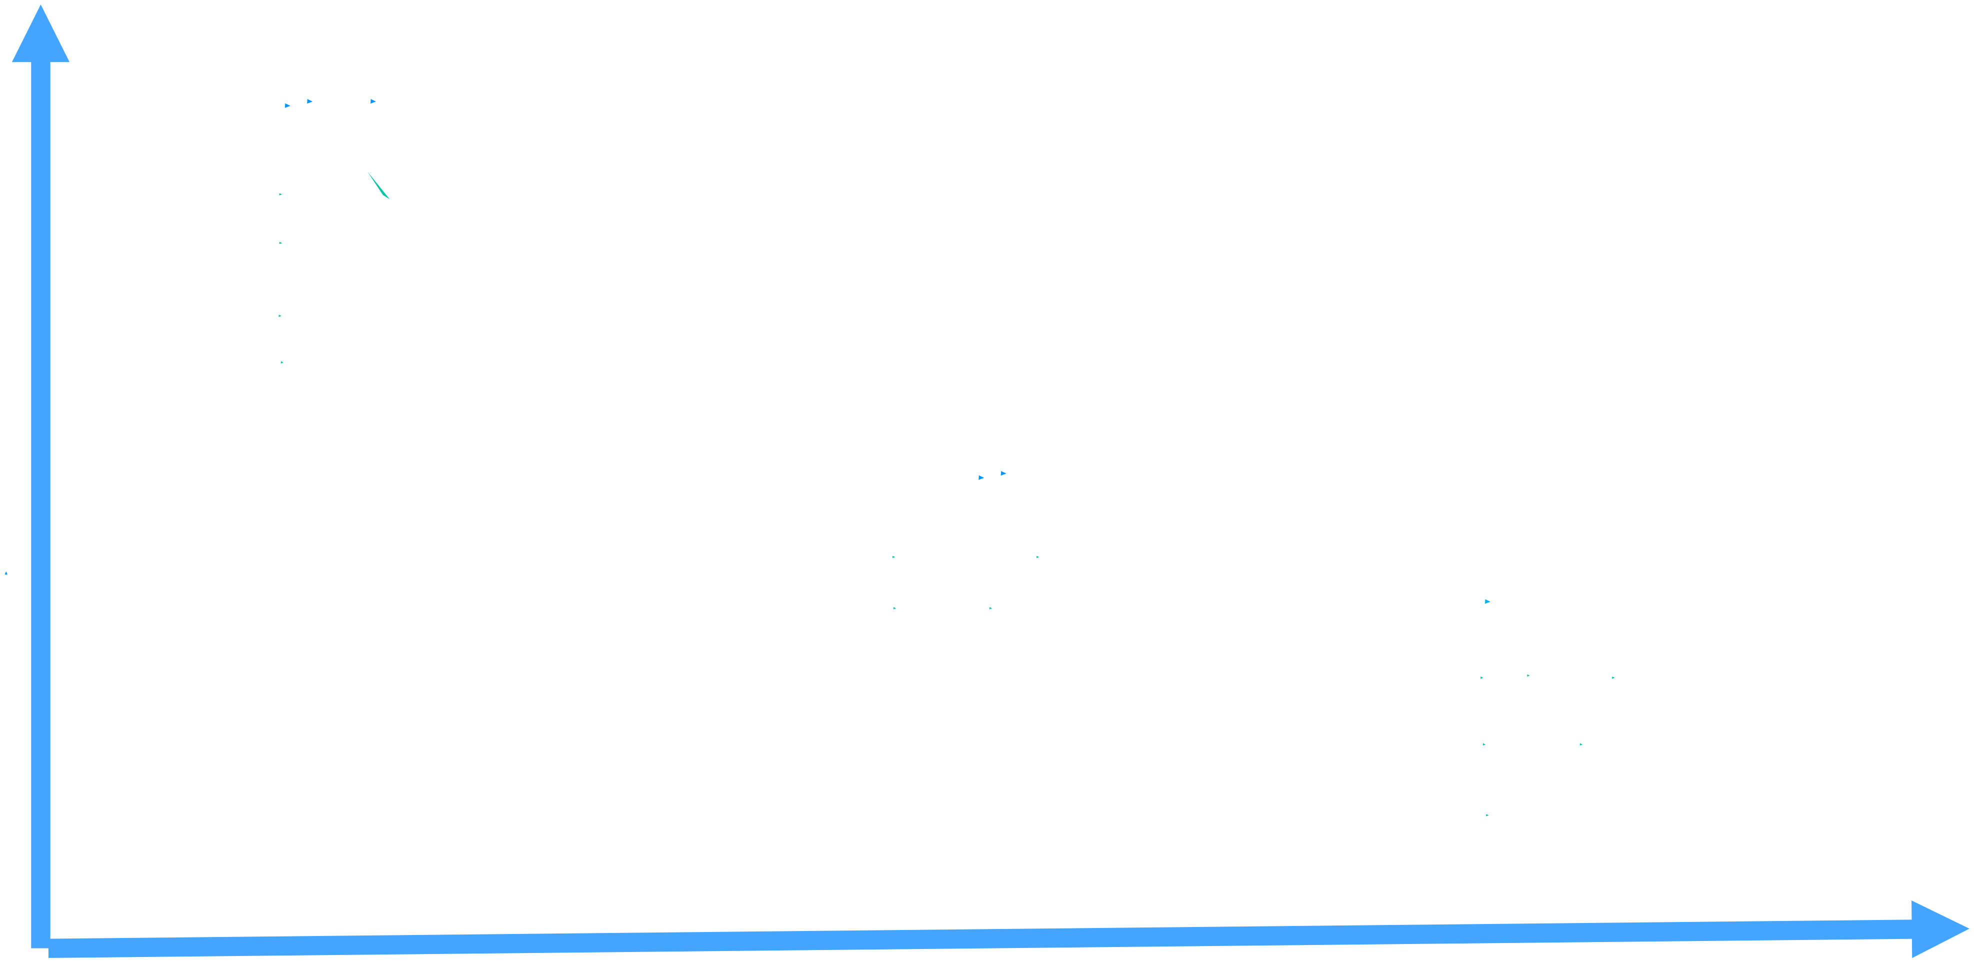

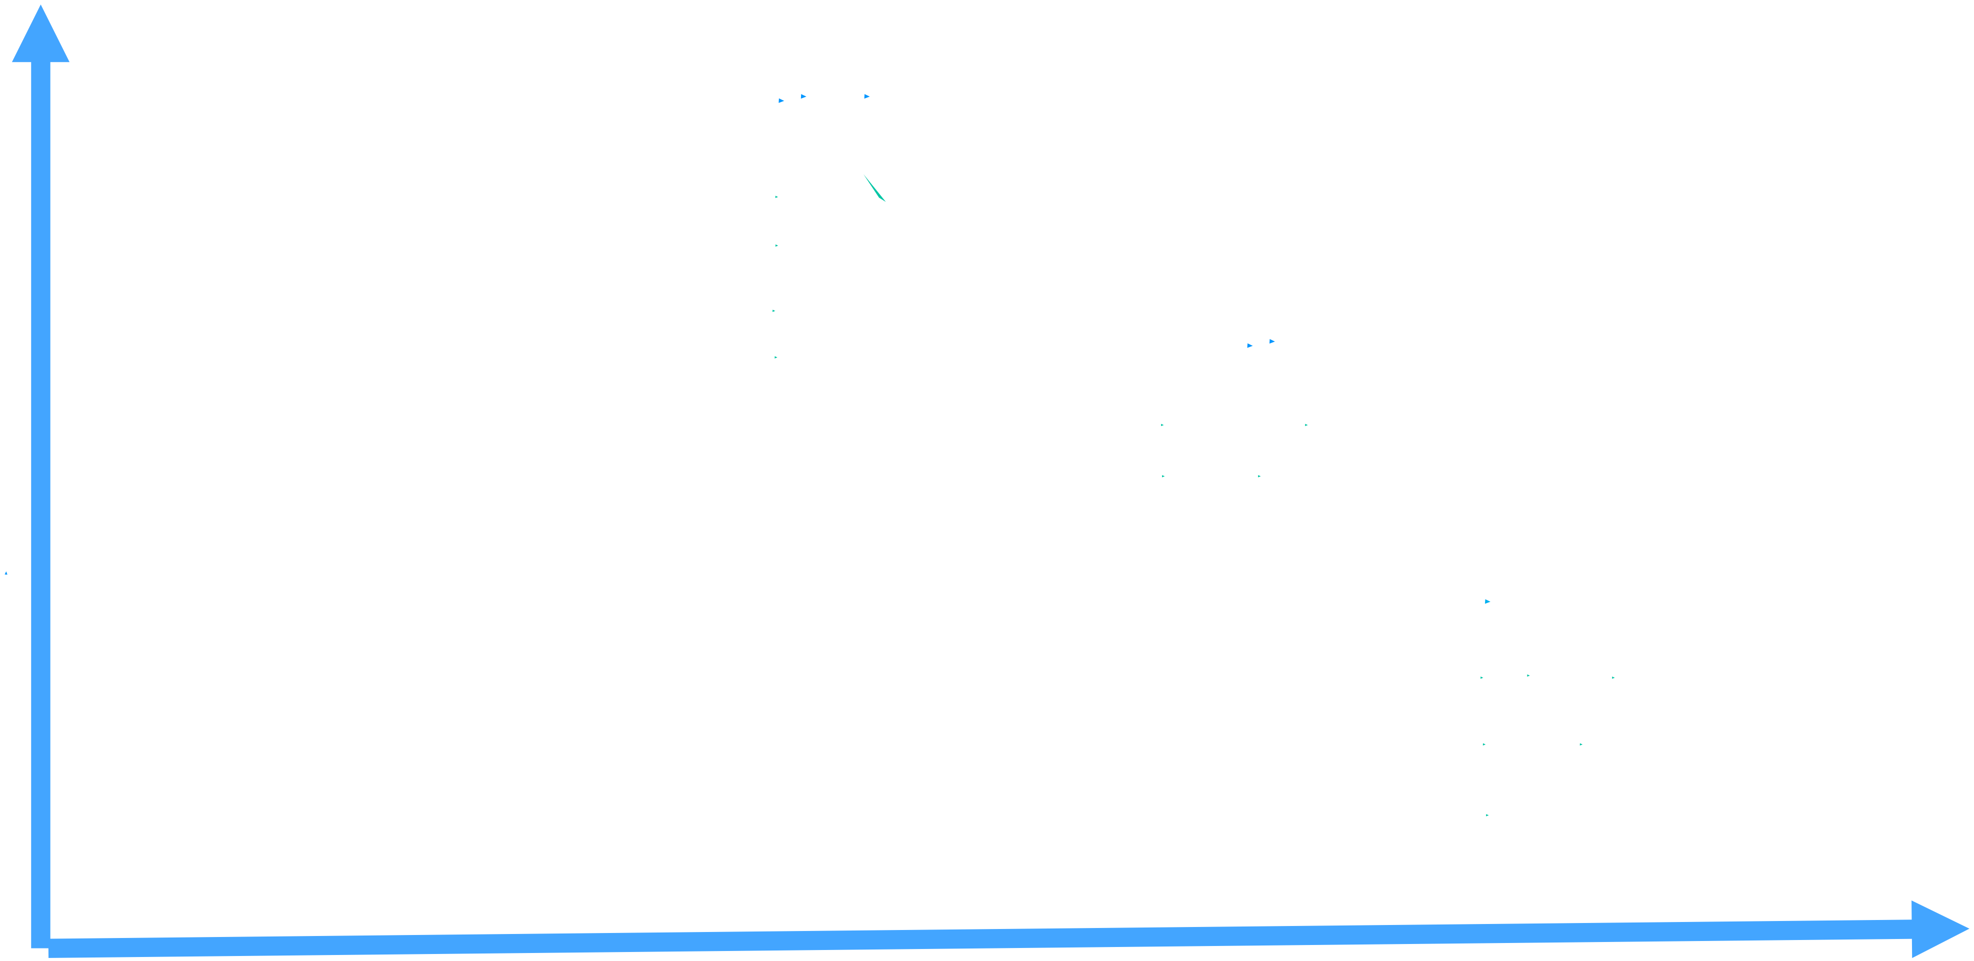

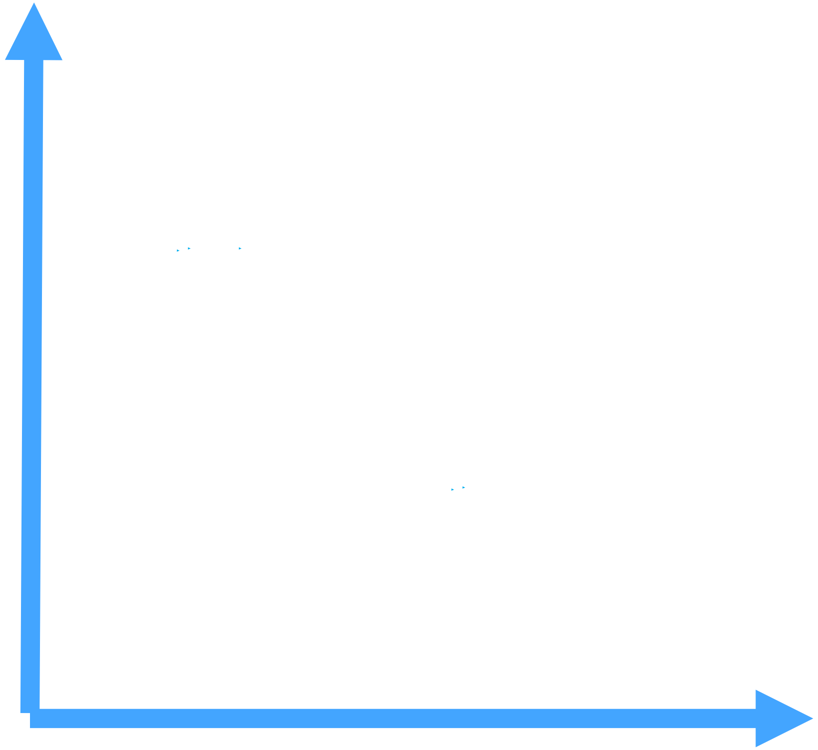

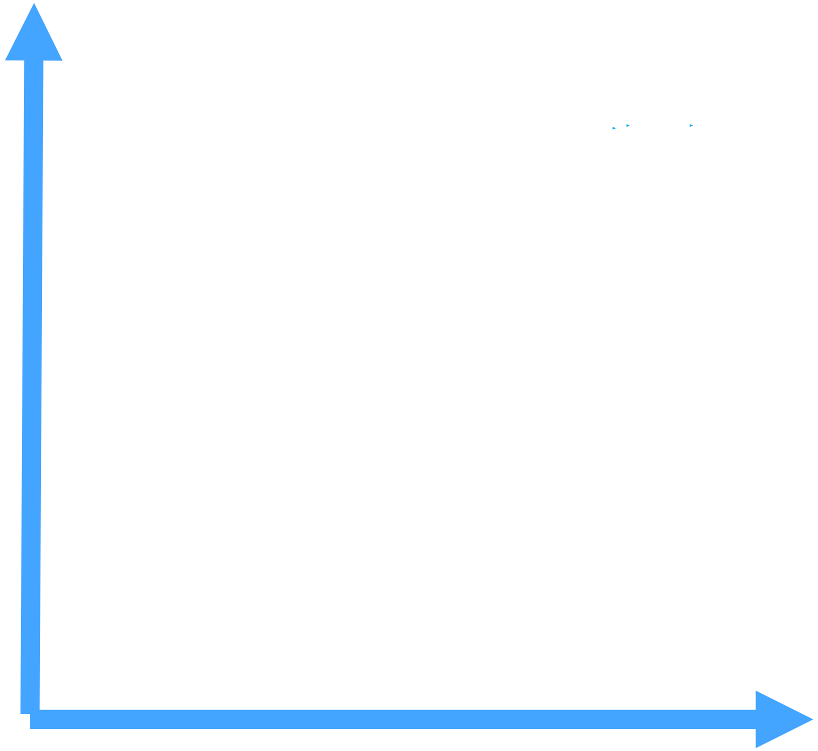

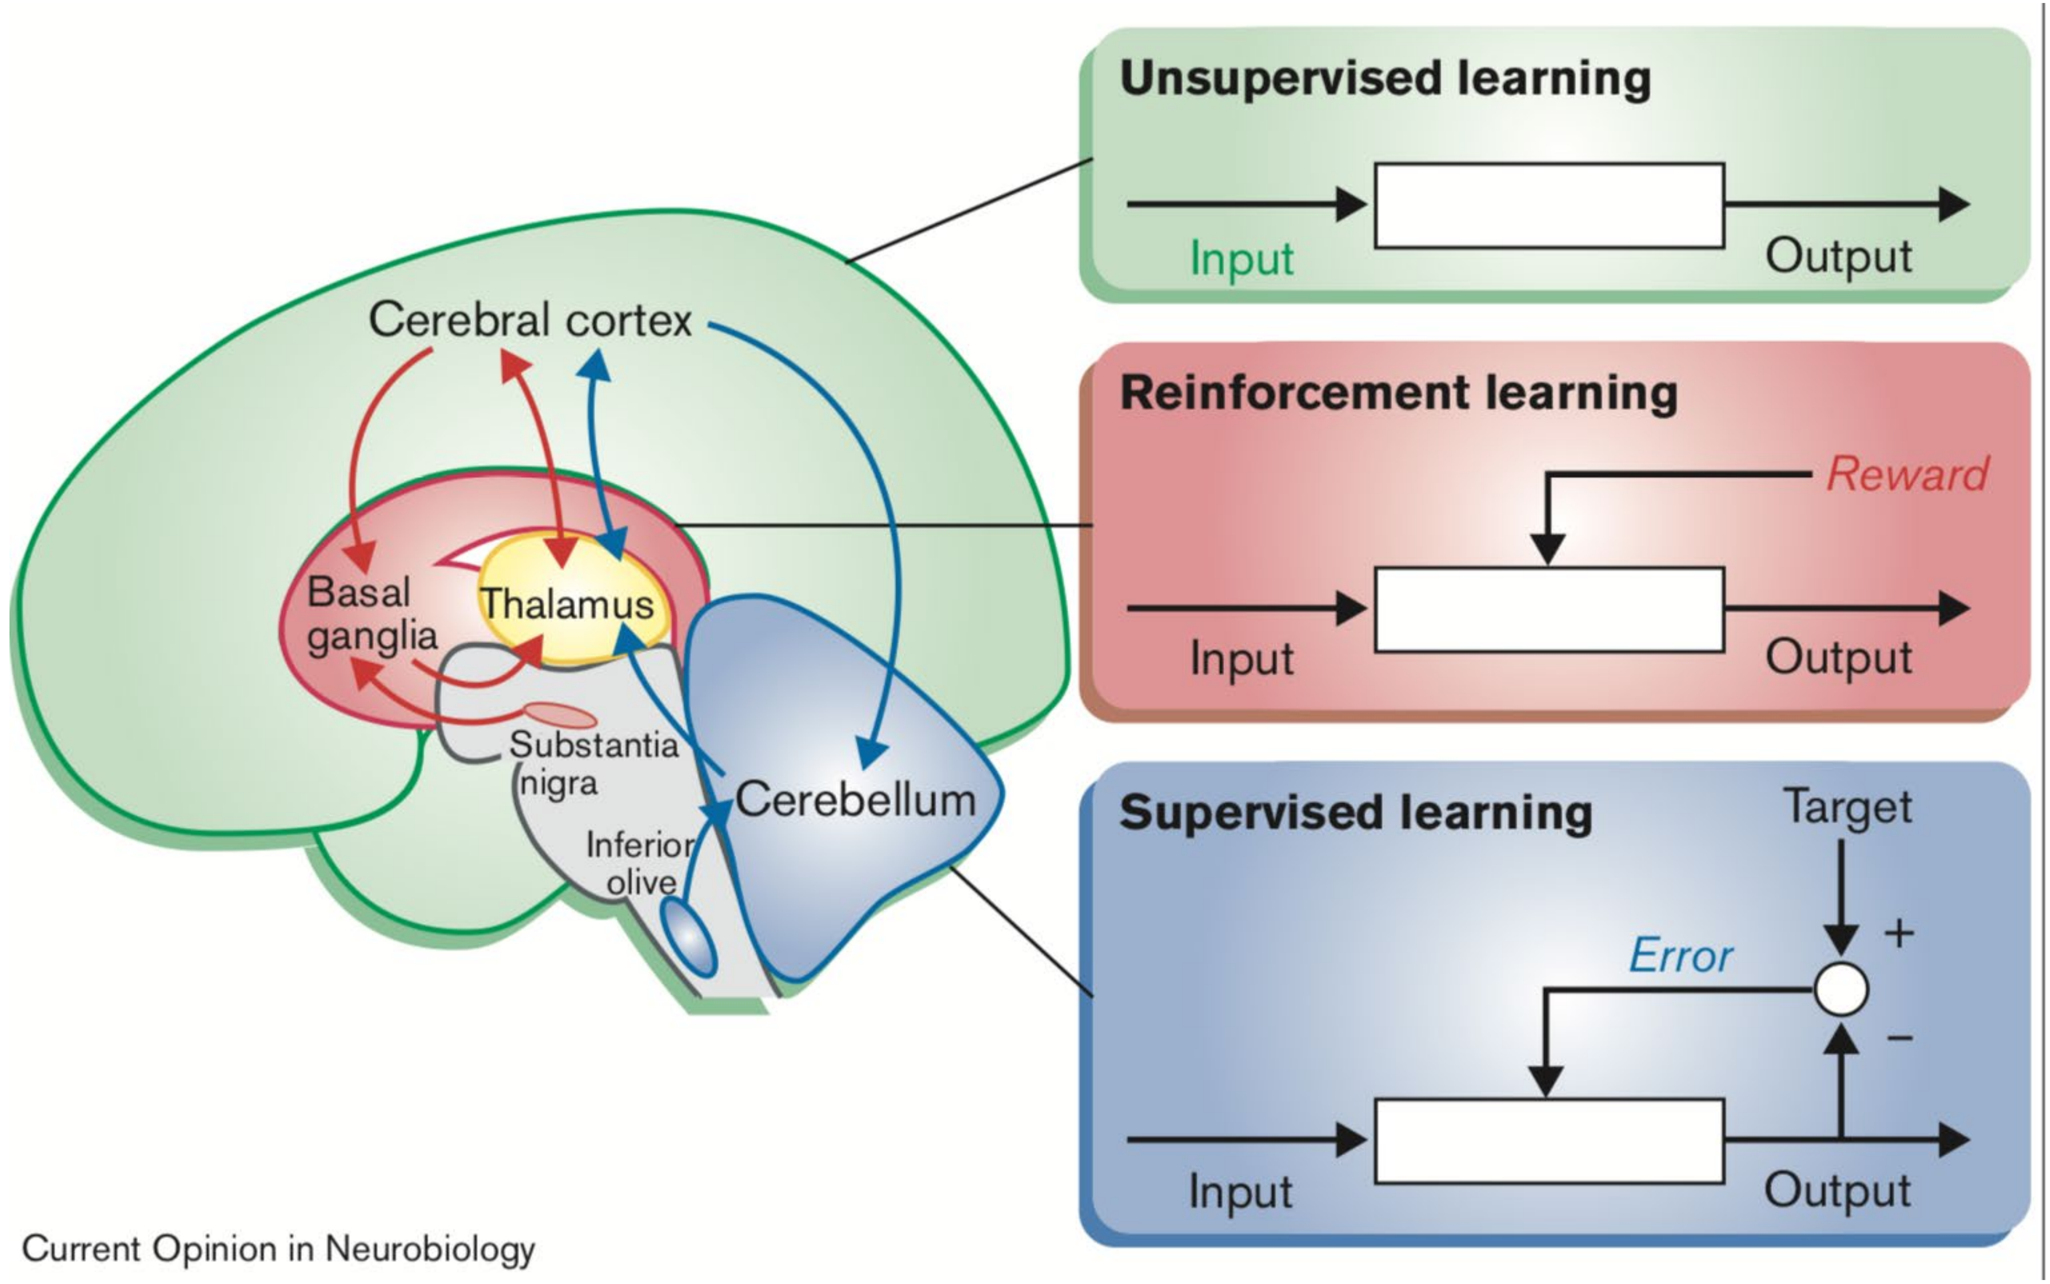

Source: Kenji Doya Complementary roles of basal ganglia and cerebellum in learning and motor control

Cellular Network Traffic Scheduling with Deep Reinforcement Learning

Sandeep Chinchali [1], Pan Hu [2], Tianshu Chu [3], Manu Sharma [3], Manu Bansal [3], Rakesh Misra [3]
Marco Pavone [4] and Sachin Katti [1,2]
[1] Department of Computer Science, Stanford University
[2] Department of Electrical Engineering, Stanford University
[3] Uhana, Inc.
[4] Department of Aeronautics and Astronautics, Stanford University
{csandeep, panhu, pavone, skatti}@stanford.edu, {tchu, manusharma, manub, rakesh}@uhana.io

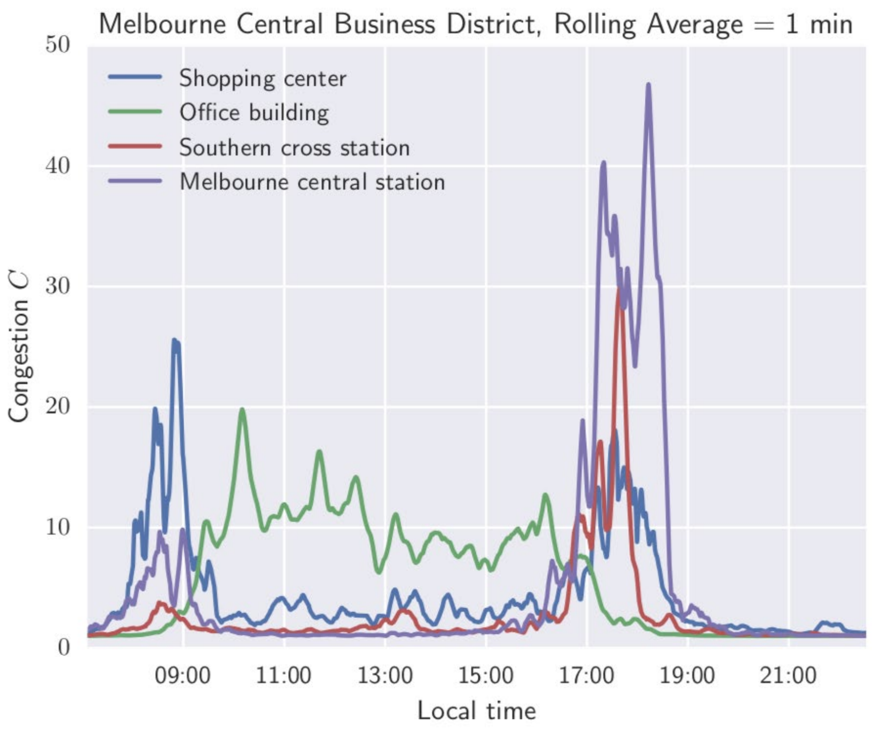

Figure 1: Time-variant congestion patterns in Melbourne.

Efficient Large-Scale Fleet Management via Multi-Agent Deep Reinforcement Learning

Kaixiang Lin
Michigan State University
linkaixi@msu.edu

Renyu Zhao, Zhe Xu
Didi Chuxing
{zhaorenyu,xuzhejesse}@didichuxing.
com

Jiayu Zhou
Michigan State University
jiayuz@msu.edu

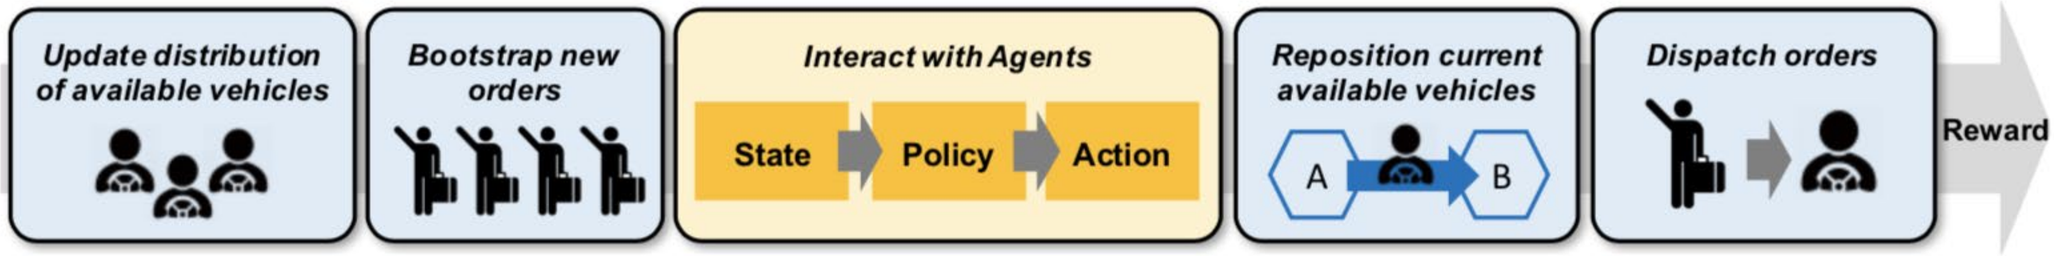

Human-level control through deep reinforcement learning

Volodymyr Mnih[1]*, Koray Kavukcuoglu[1]*, David Silver[1]*, Andrei A. Rusu[1], Joel Veness[1], Marc G. Bellemare[1], Alex Graves[1], Martin Riedmiller[1], Andreas K. Fidjeland[1], Georg Ostrovski[1], Stig Petersen[1], Charles Beattie[1], Amir Sadik[1], Ioannis Antonoglou[1], Helen King[1], Dharshan Kumaran[1], Daan Wierstra[1], Shane Legg[1] & Demis Hassabis[1]

Curriculum Learning

Yoshua Bengio[1] YOSHUA.BENGIO@UMONTREAL.CA
Jérôme Louradour[1,2] JEROMELOURADOUR@GMAIL.COM
Ronan Collobert[3] RONAN@COLLOBERT.COM
Jason Weston[3] JASONW@NEC-LABS.COM
(1) U. Montreal, P.O. Box 6128, Montreal, Canada (2) A2iA SA, 40bis Fabert, Paris, France
(3) NEC Laboratories America, 4 Independence Way, Princeton, NJ, USA

When a large language model is trained on a sufficiently large and diverse dataset it is able to perform well across many domains and datasets. GPT-2 zero-shots to state of the art performance on 7 out of 8 tested language modeling datasets. The diversity of tasks the model is able to perform in a zero-shot setting suggests that high-capacity models trained to maximize the likelihood of a sufficiently varied text corpus begin to learn how to perform a surprising amount of tasks without the need for explicit supervision.[5]

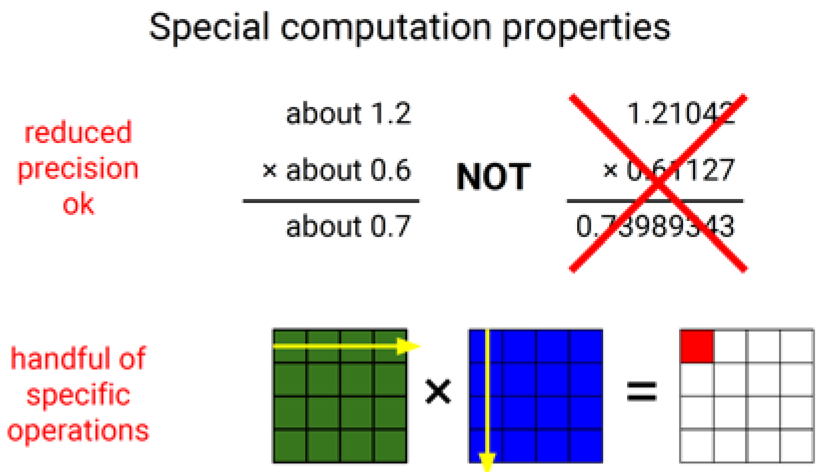

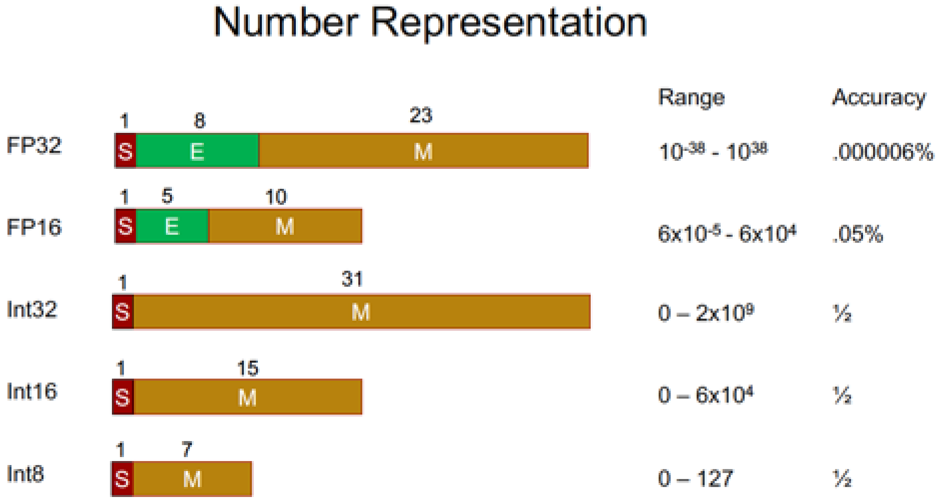

- Deep Learning is Empirically Scaleable (Baidu)
- Computationally Homogenous
- Constant runtime & memory use
- Highly Portable
- Easily Baked into silicon → “Semiconductor Rennaisance”

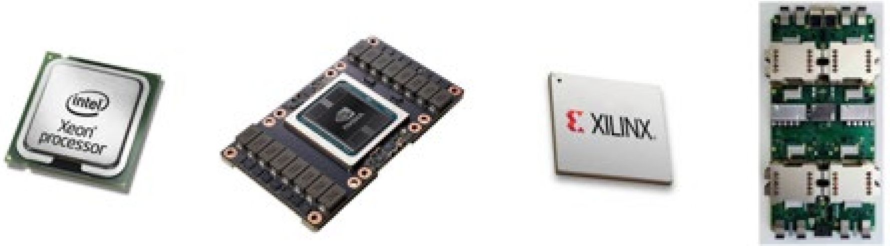

CPU
- Threads
- SIMD

GPU
- Massive threads
- SIMD
- HBM

FPGA
- LUTs
- DSP
- BRAM

TPU
- MM unit
- BRAM

- **Relax Precision** : Small integers are better
- Relax Synchronization : data races are better
- Relax Communication : sparse communication is better
- Relax Cache Coherence : incoherence is better

“Olukutan,Stanford Neurips Keynote 2018”

Data Pipeline for AI Workflow

Scale and Optimize Each Stage of the Data Pipeline

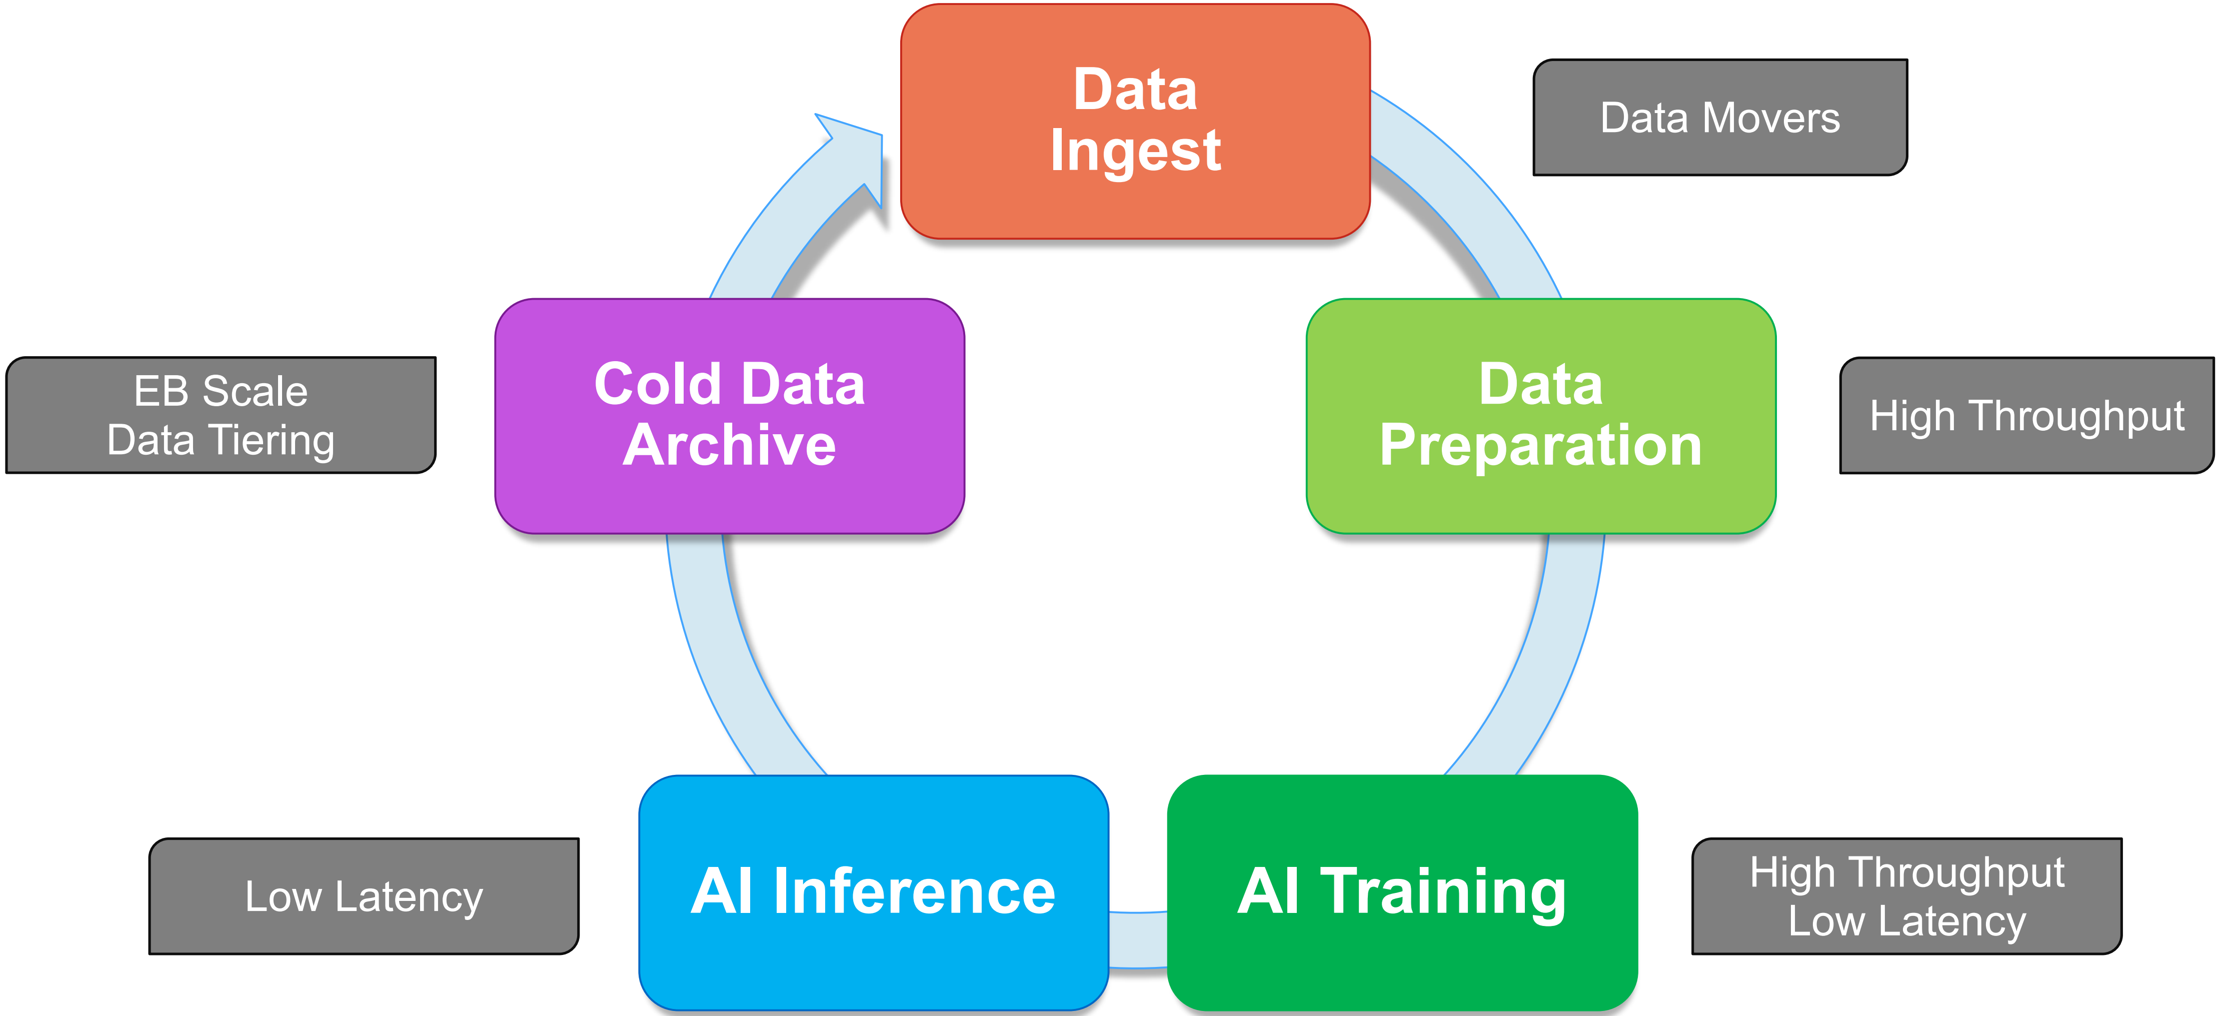

Edge to Core to Cloud

Seamless data management

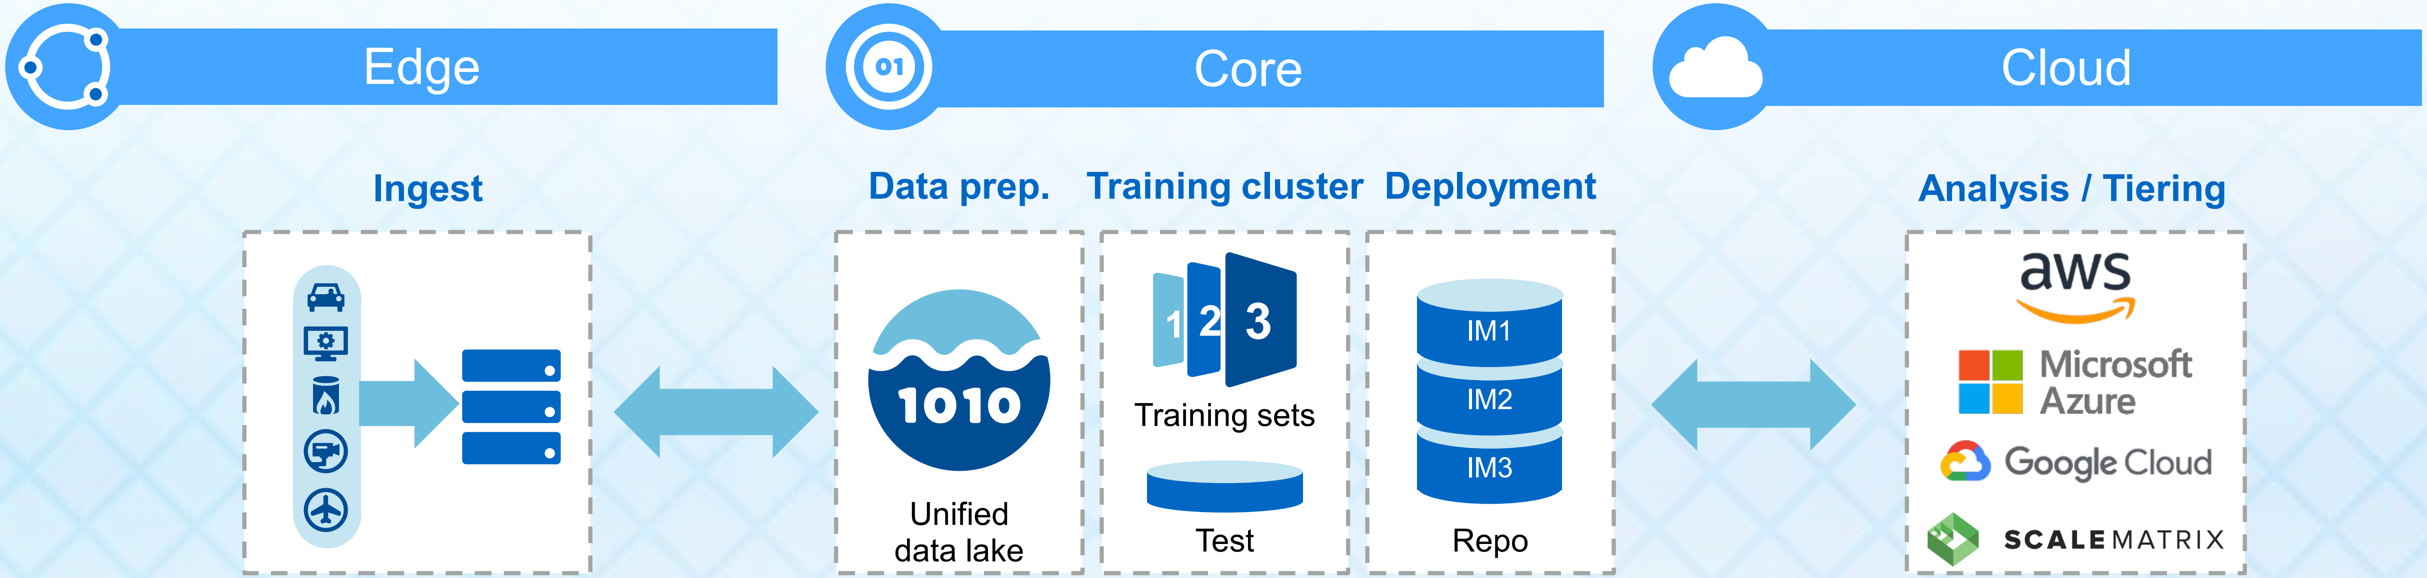

Edge – Ingest
- Data collection
- Edge-level AI

Core – Data prep.
- Aggregation
- Normalization

Core – Training cluster
- Exploration
- Training

Core – Deployment
- Deployment
- Model serving

Cloud – Analysis / Tiering
- Cloud AI (GPU instances)
- Data tiering

ONTAP Select
Vector Data

CloudSync
SnapMirror
Kafka, NiFi

ONTAP AI
FlexPod AI
FlexPod Inference

FabricPool
CloudSync
NDAS

Cloud Volumes
ScaleMatrix ONTAP AI

Full Scale-out ONTAP AI with DGX-1

24-node A800 cluster, driving 108 DGX-1's

NetApp

Full Scale-out ONTAP AI with DGX-2

24-node A800 cluster, driving 36 DGX-2's

NetApp

Dense Cabinet for AI

Deliver Dense Cabinet for Exascale AI

Challenge:

- Data centers facilities lack the power and cooling for the latest high-performance AI computing infrastructure.

New Capability:

- Dense and Modular Dynamic Density Control (DDC) liquid-air cooled cabinet for AI
- This cabinet combines the efficiency of water with the flexibility of air, cooling up to 52kW of power load in a 45U cabinet.
- Cabinets can be deployed in any environment.
- Provides clean-room environment, guaranteed air flow, integrated security, and fire suppression.

SCALEMATRIX & DDC CABINETS SUPPORT
45U & 52kW

2 x 1U Cisco® 3232C	2kW
1 x 4U NetApp® A800	2kW
3 x 10U NVIDIA® DGX-2™	36kW
Total kW	**40kW**

ARTIFICIAL INTELLIGENCE

Trends >

Artificial Intelligence Timeline

Why Now?

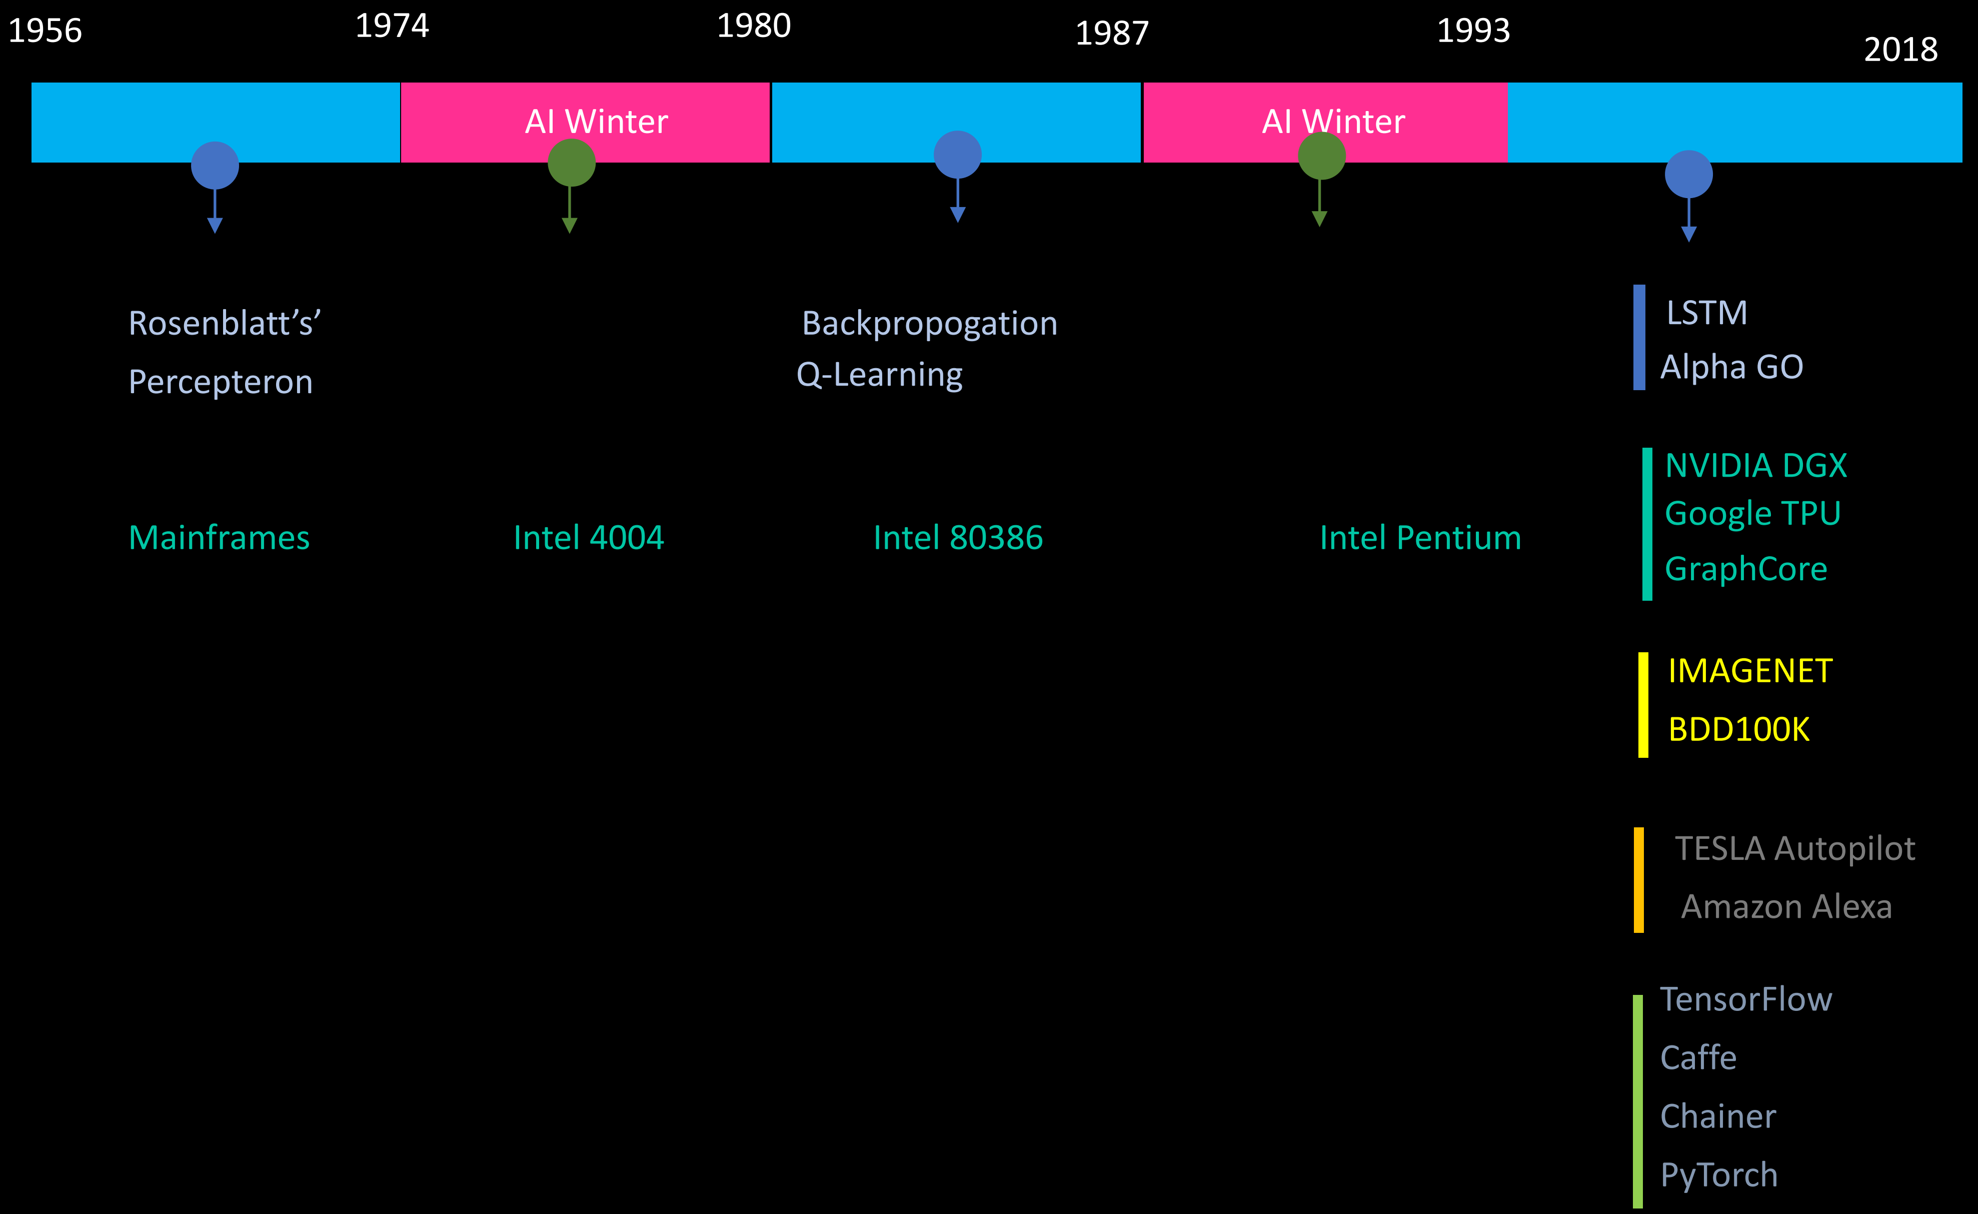

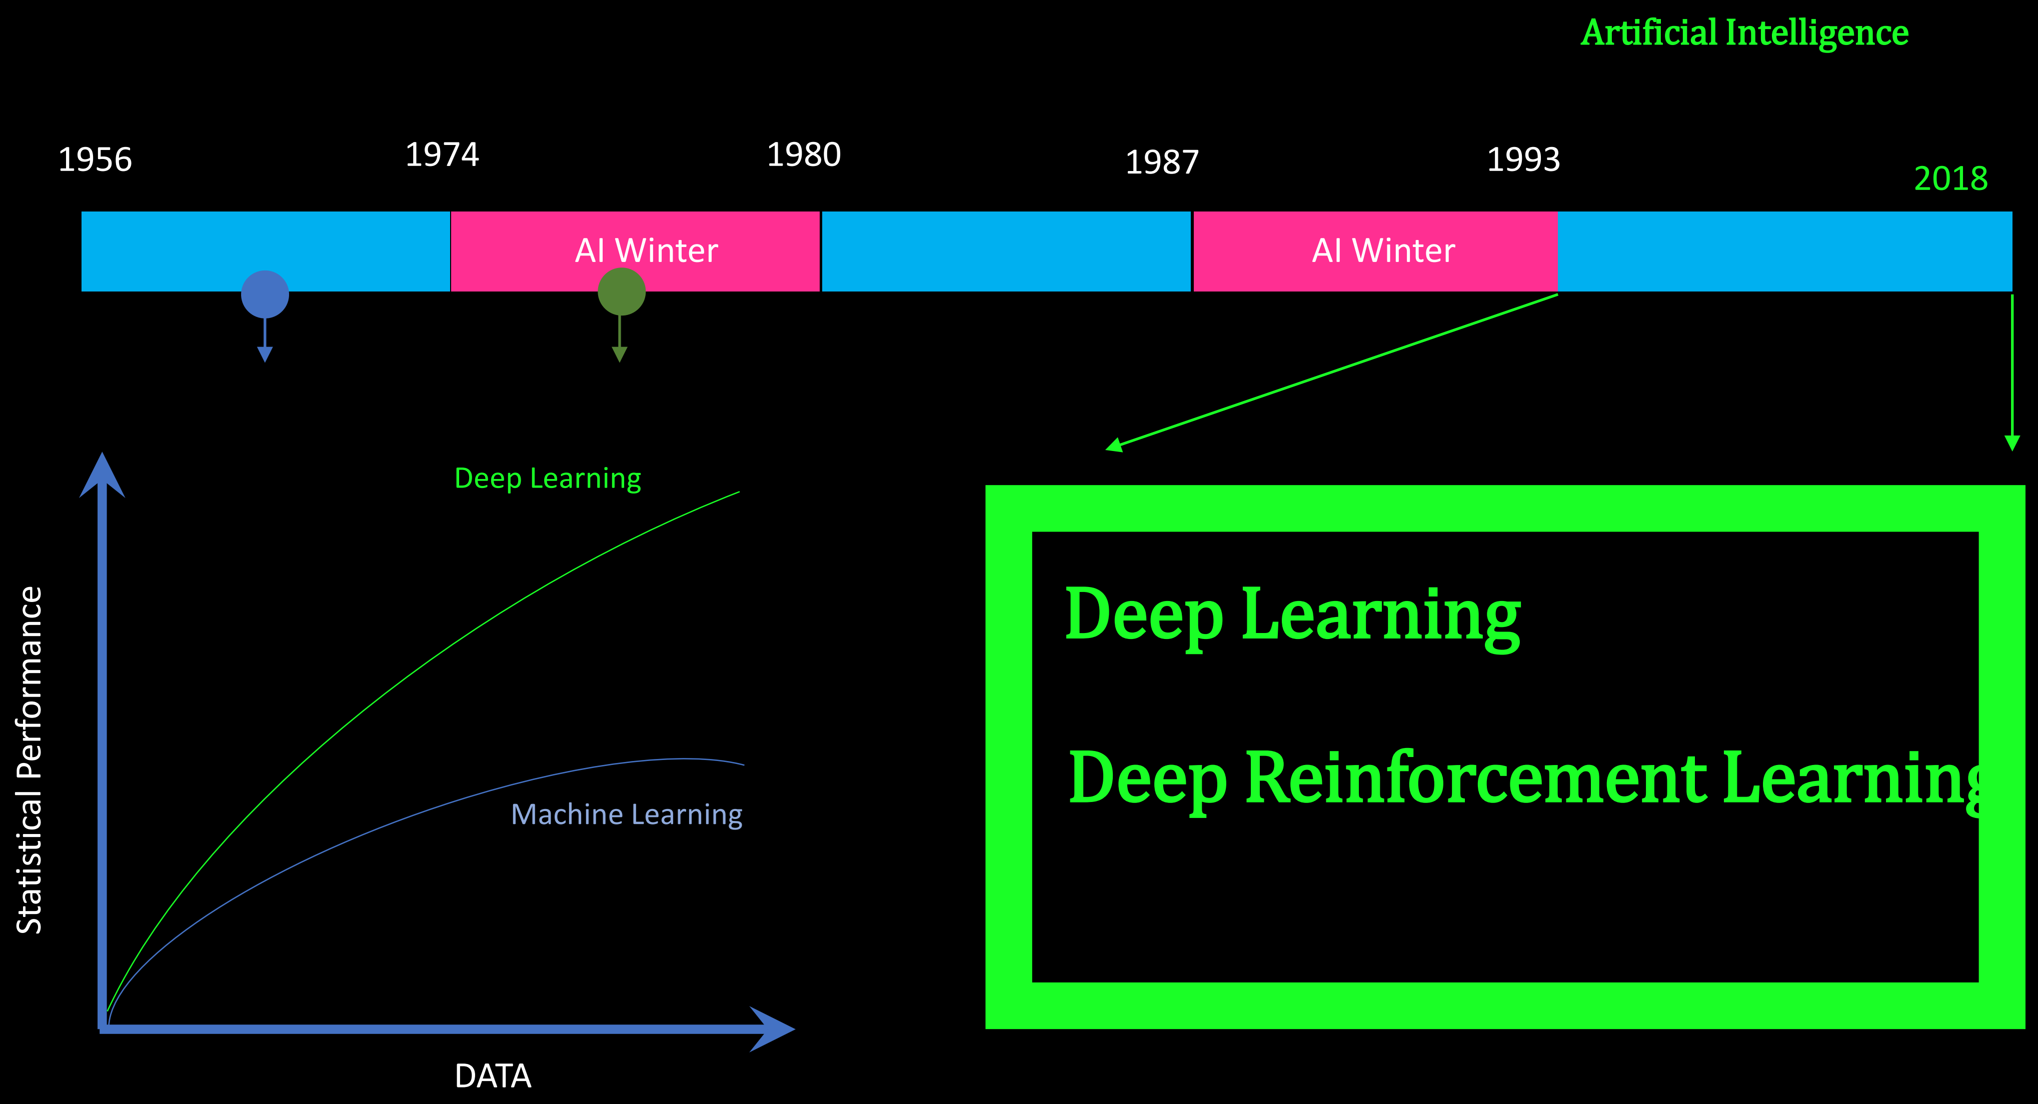
Artificial Intelligence
1956
1974
1980
1987
1993
2018
AI Winter
AI Winter
Deep Learning
Machine Learning
Statistical Performance
DATA
Deep Learning
Deep Reinforcement Learning

Artificial Intelligence

1993 2018 2020

AI Winter

Deep Learning

Deep Reinforcement Learning

Bayesian Deep Learning BNP

Meta-Learning Imitation-Learning

Continual-Learning Causal-Learning

Relational Representation-Learning

Smooth GamesOptimization

Probabilistic reinforcement learning

Conversational AI Systems ML Emergent communication

Modelling Physical World Modelling Spatio Temporal domain

Physics,Geophysics,Geochemical Molecules & Materials

Financial Services Intelligent Transport Systems

Medical Imaging & Healthcare

Robustness Compactness Interoperability Security

Social,Political,Ethics

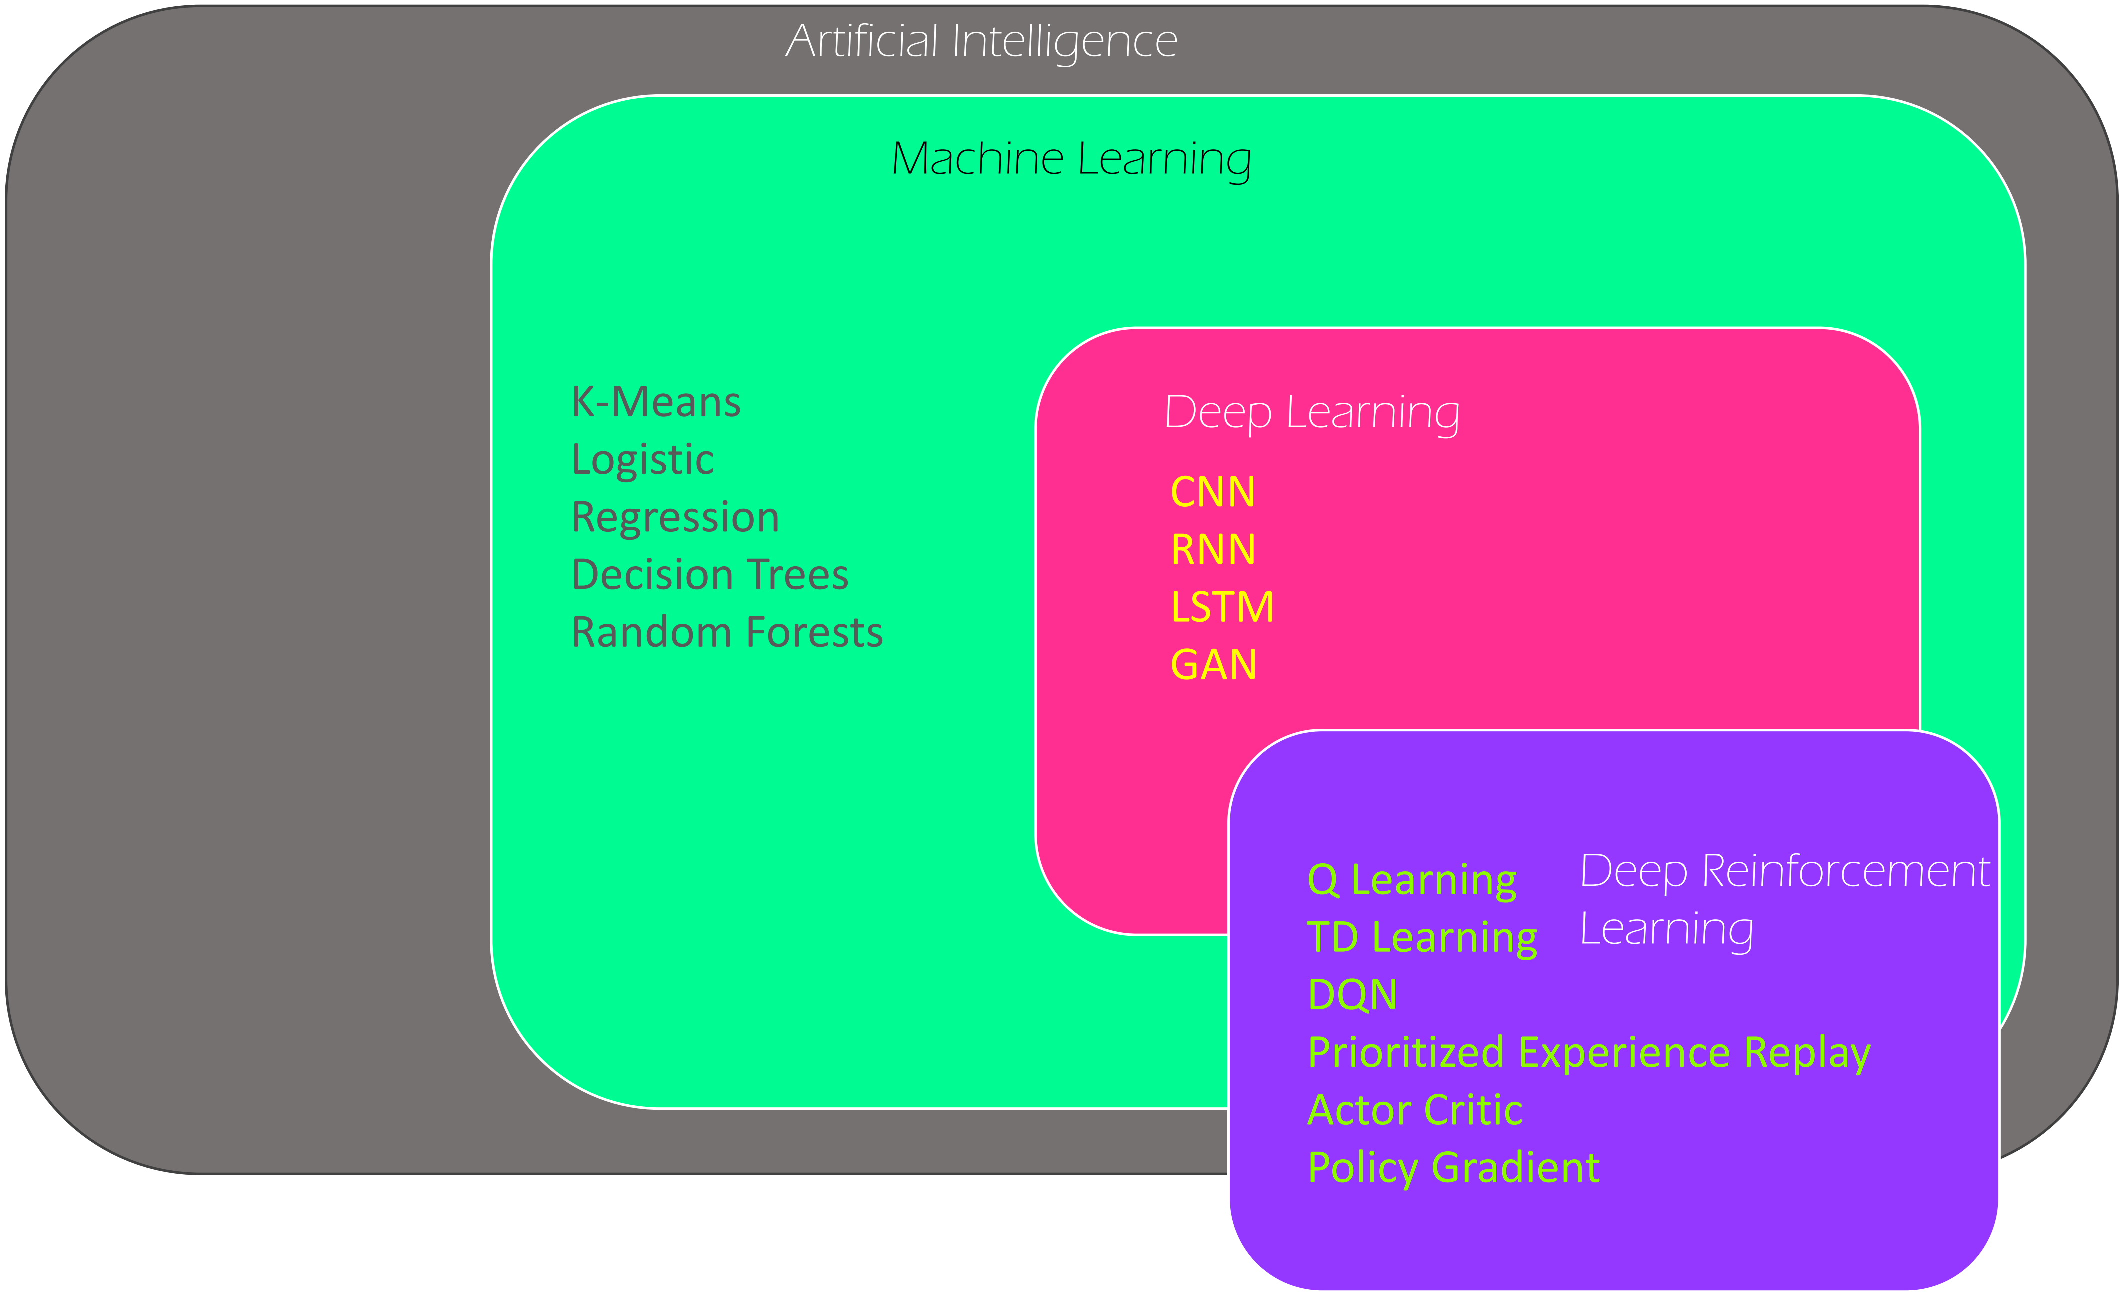
Artificial Intelligence
Machine Learning
K-Means
Logistic
Regression
Decision Trees
Random Forests
Deep Learning
CNN
RNN
LSTM
GAN
Deep Reinforcement
Learning
Q Learning
TD Learning
DQN
Prioritized Experience Replay
Actor Critic
Policy Gradient

AI Use Cases

AI Requirements Discovery

NetApp

WHAT PROBLEM ARE YOU SOLVING?

Defining the AI/DL Task

INPUTS	BUSINESS QUESTION	AI/DL TASK	EXAMPLE OUTPUTS		
			HEALTHCARE	RETAIL	FINANCE
Text Data, Images, Video, Audio	Is "it" present or not?	Detection	Cancer Detection	Targeted ads	Cybersecurity
	What type of thing is "it"?	Classification	Image Classification	Basket Analysis	Credit Scoring
	To what extent is "it" present?	Segmentation	Tumor Size/Shape Analysis	Build 360º Customer View	Credit Risk Analysis
	What is the likely outcome?	Prediction	Survivability Prediction	Sentiment & behavior recognition	Fraud Detection
	What will likely satisfy the objective?	Recommendations	Therapy Recommendation	Recommendation Engine	Algorithmic Trading

SOME KEY DECISIONS TO MAKE

FACTOR	DESCRIPTION
DL Challenge	Supervised or unsupervised, classification or regression, # of labels?
Architecture	What is the simplest architecture I can use?
Training Model	How am I going to tune my neural net? Kinds of non-linearity, loss function and weight initialization? Best training framework?
Data Quantity	How much data will be sufficient to train my model? How do I go about finding that data and is it evenly balanced?
Data Quality	Is my data directly relevant to the problem & real world data.
Data Labels	Is training data is labeled same as raw data sets, how do I 'featurize'?
Data Similarity	Is data same length vectors or does it require pre-processing?
Data Storage & Access	Where is it stored, locally and on network Data pipeline? How do I plan to extract, transform and load the data (ETL)?
Infrastructure	Cloud, On-premise, Hybrid. GPUs, CPUs or both? Single or distributed systems? Integration with languages, ent. apps/ databases.

AI Requirements Discovery - 101

Basics

3 year AI & Data science plans?

Any AI/DL consultants on staff?

Have a data science team on staff?

Buyer vs Influencer?
- Core IT
- Data science team
- AI/DL consultants

Data science use cases?

Conversation w/ each of these functions?

Data

Current strategy?
- Cloud
- On-prem
- Hybrid

Data sources into the data lake?

Type of data:
- Time series
- Audio / video
- Image
- Other

Compute

Compute vendor?
- DGX-1
- White-box GPU
- CPU

AI Inference deployment?
- DGX-1
- Existing Server + GPU
- Existing CPU server

How many compute nodes in the cluster?

GTM

AI budget for the next 2 years?

Storage req's:
- Capacity
- Performance
- Tiering to Cloud

Channel partners, VAR's, consulting vendors?

Edge

Edge use cases?

NetApp

AI Requirements Discovery - Data Science

- What are your top AI use cases / applications?
- Do you use ML or DL or both?
- Data set footprint and growth pattern?

- Data types used – image, video, text, time-series, …
- How big is the data set?
- Composition of the data set – small files, large files, or a combination?

- How do you pre-process your data set?
- Time spent to process data?
- Tools & vendors used?

- Where is the data stored?
- Where/how will training read the data from?
- Any read/write/latency performance targets?
- Maintain versions of datasets?

- Foot print of your compute cluster – # of GPUs, CPUs
- Physical memory of the m/c
- Cached copy of data?

- Are your GPUs fully utilized?
- Bottleneck in the workflow?
- Use distributed training?
- If yes, environment setting?

- ML/DL framework used?
- Software/ framework to speed up data pre-process?
- Software/ framework to control the cluster / environment / storage?

- Version control of library / API for ML/DL?
- NGC or vendor native?
- Container orchestration at scale?

- Success metrics of training?
- Model selection.
- Model footprint?
- # of concurrent models.
- Model deployment platform for Inference.

AI Requirements Discovery - Data Engineer

Data Aggregation

- Edge data type & data footprint?
- Hierarchical edge strategy?
- Tools used to move data from edge to core data lake?

Data Sources

- Capacity of data lake?
- Where does the data land? - HDFS, GPFS, Lustre, NFS, NoSQL, …
- Data source to GPU cluster data movement.

Environment

- Data Tier policy – hot, warm, cold?
- What do you leverage cloud for (if)?
- Hybrid cloud requirements?

Others

- Data protection requirements ?
- Data compliance requirements?
- S/W apps & platforms?
- Inference requirements?

DL Model Training Flow

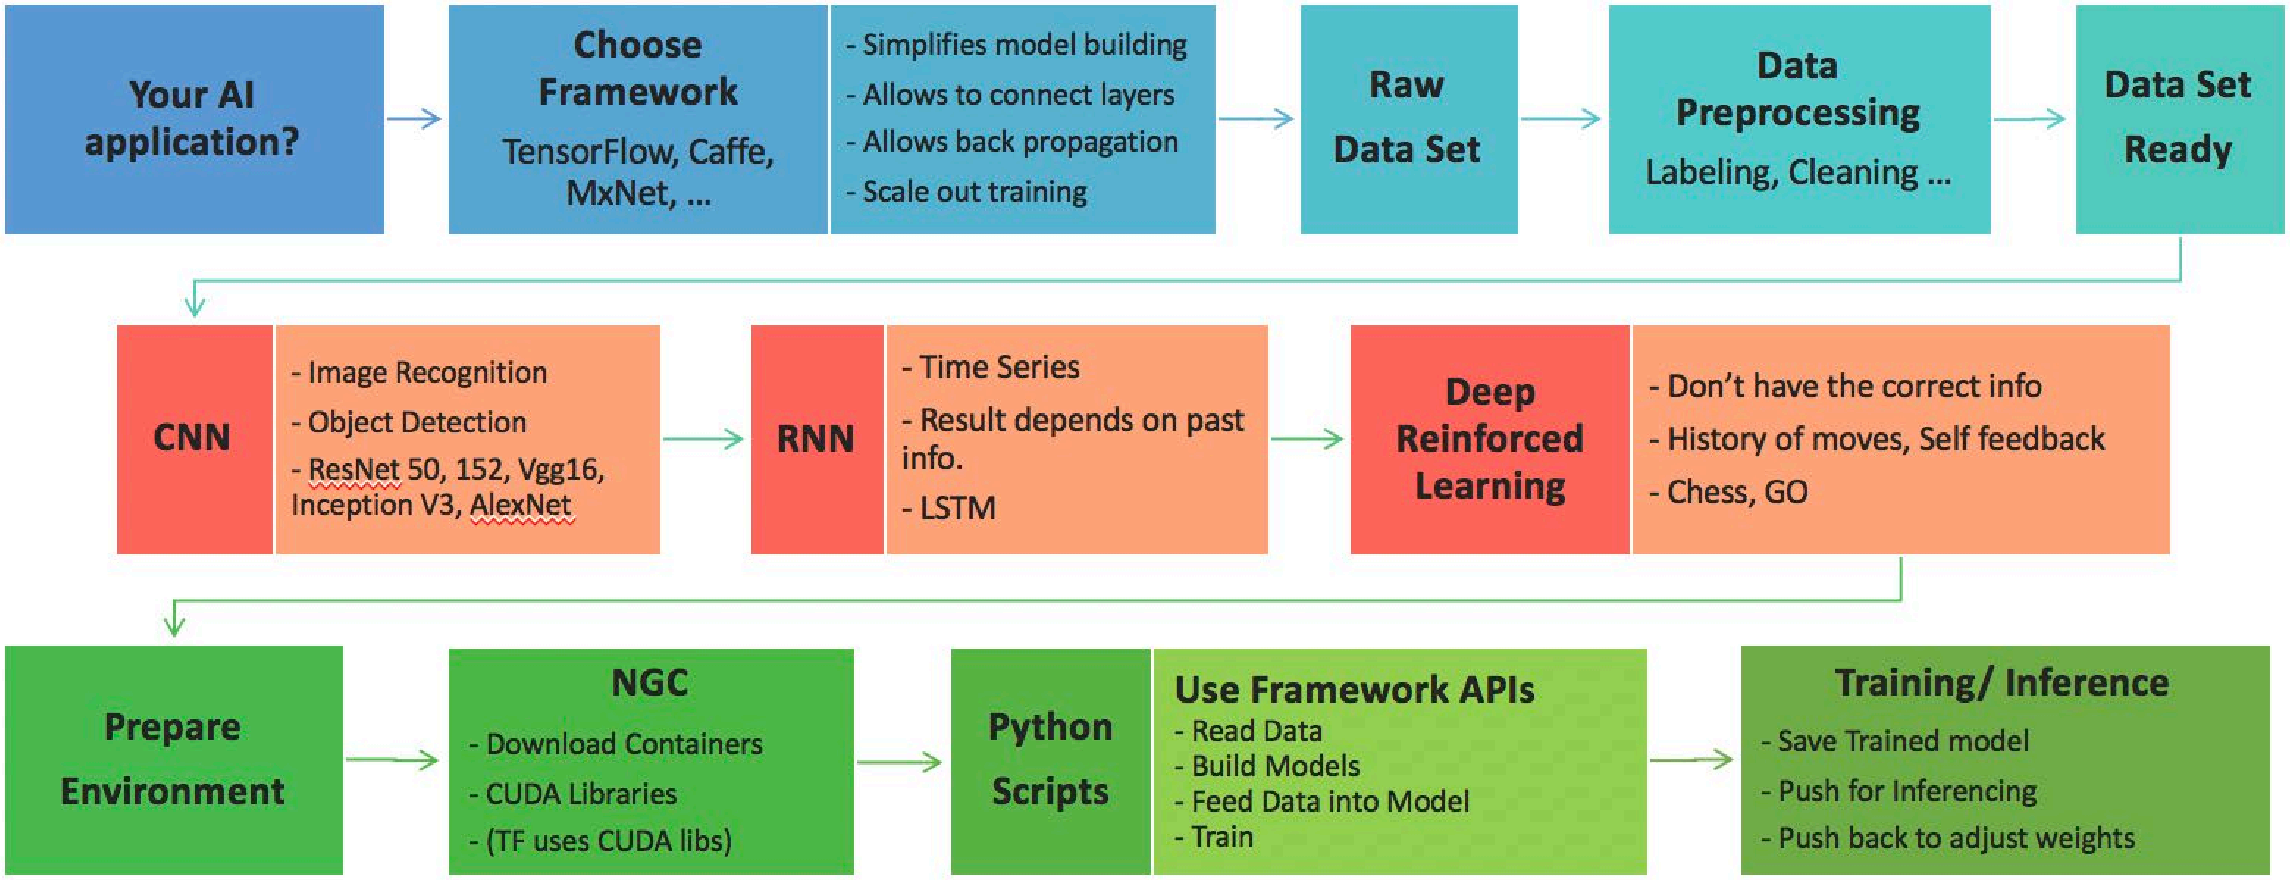

Meet Diverse Needs across Data Science and Infra Functions

Data Scientists

Real-world Data for AI / DL

Need Agile Model DevOps :

- Refreshed access to Production Datasets
- Hybrid Cloud for Model Dev
- Distributed Data Science
- Diverse Data Sources
- Model and Data Parallelism
- Multi-Tenant Model Serving
- From Model to Application

Data Architects

Future Proof Architecture

Seek Extensible Architecture:

- Architecture Scales from PoC to Production
- Future Proof to absorb technology changes
- TCO for Massive Datasets
- Maximize Utilization
- Global Scale

Data/IT Admins

Lowest TCO in face of shrinking budgets

Balance Cost & TTM :

- Leverage vs. Dedicate HW Infra
- Stable Operations & Upgrades
- Supported Components & Ecosystems
- Diverse needs across Big Data, AI/DL and HPC

Balanced Architecture to Deliver for Stakeholders

NetApp

Deployment Options for AI

Where is the source data?

Move Data into AI Platform	Data In-Place	Co-Lo Solution	Source Data in Cold Storage	Cloud Deployment
• 80 - 90% deployments • HDFS, Splunk, NoSql , Lustre, GPFS → ONTAP AI (NFS, GPUs) • Leverage Data Movers to move data into ONTAP AI • FabricPool for data tiering	• All reside and deploy on ONTAP • Concept of Unified Data Lake • Data on ONTAP AI • FabricPool for data tiering	• Greater control of data • NPS solution • Data on NPS, GPUs/ Services on the Cloud	• Data is moved in from cold data tiers for model training • Move data from StorageGrid into ONTAP AI	• Data / GPUs provisioned on the public clouds • Use Cloud Volumes Service for file services • GPUs on Cloud for compute

NetApp

Move Data to AI Platform

Move Data into ONTAP AI

Data movement from HDFS, MapR-FS, GPFS, Lustre, S3 to AI Platform

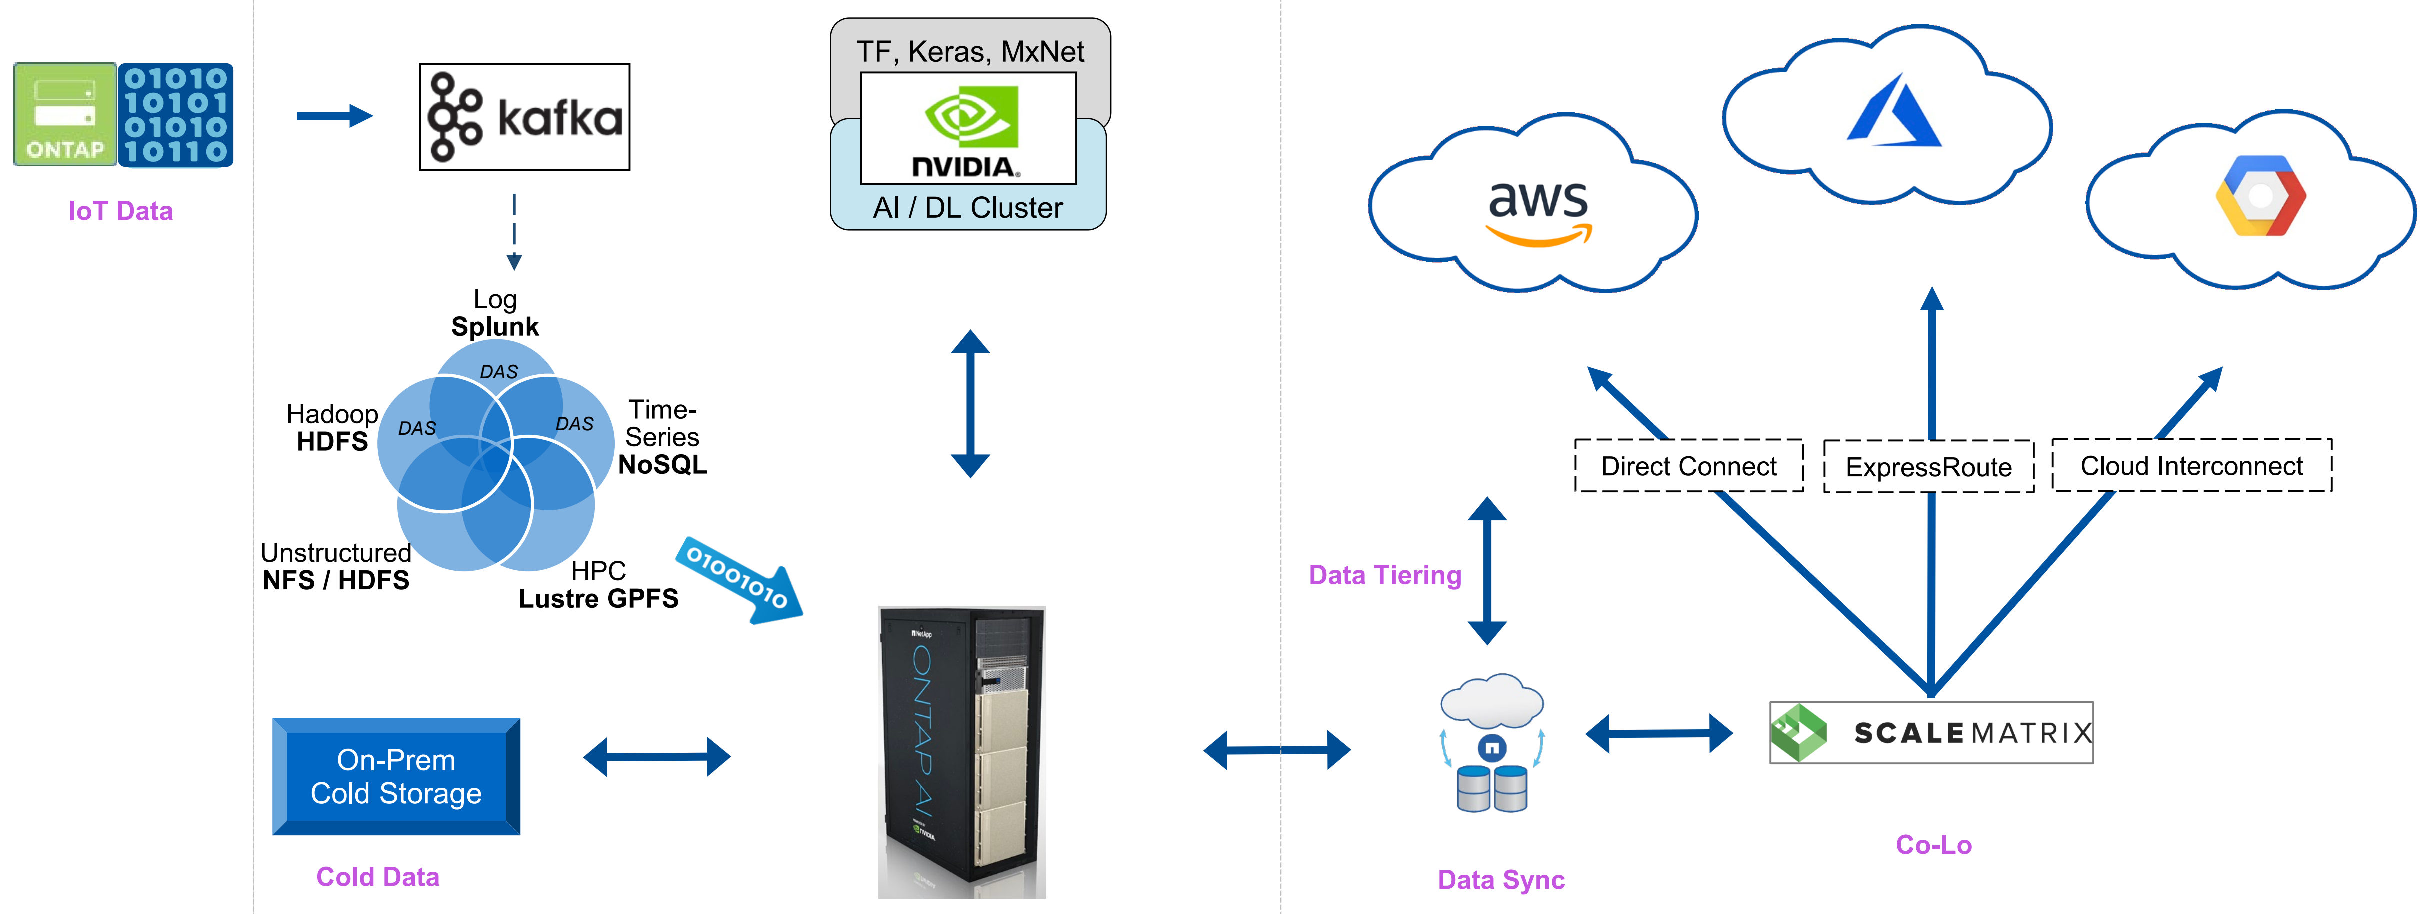

NetApp

In-Place Data with Hybrid Cloud Option

Unified data lake serving CPU and GPU Compute Clusters

Data In-Place

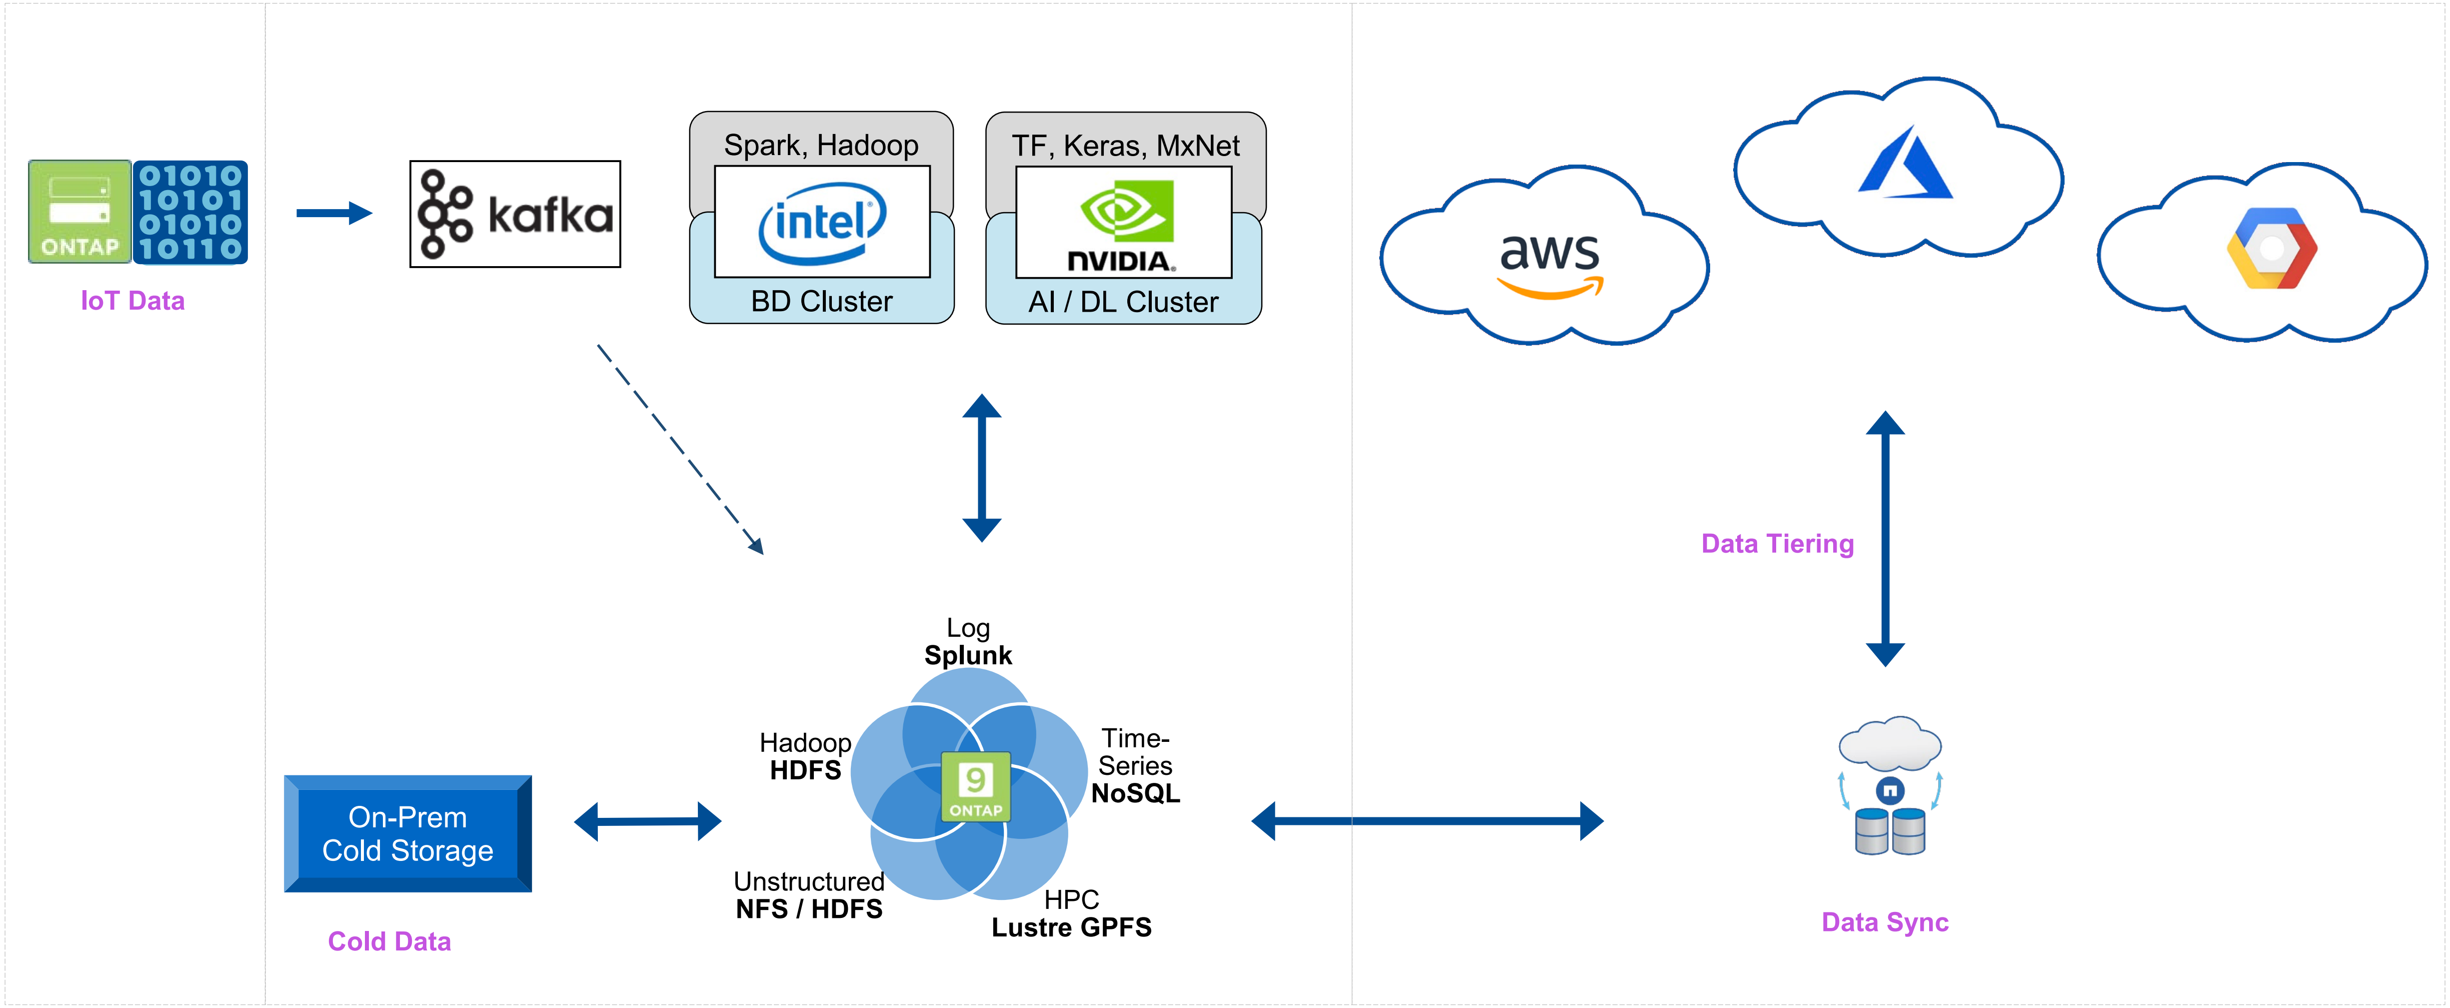

NetApp

AI Platform in Co-Location Data Center

Driven by Power, Cooling, Lease options, As-a-Service based, Consumption based

Co-lo Solution

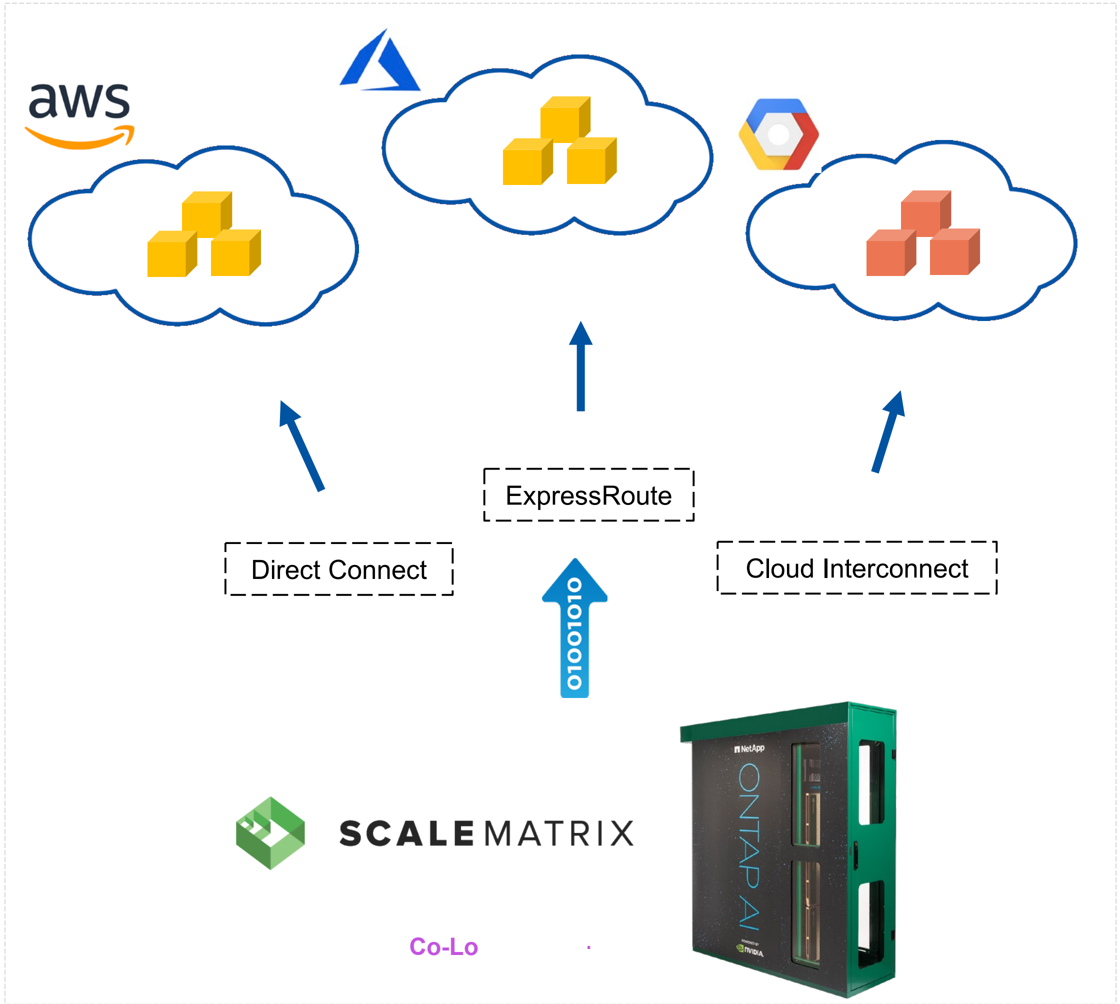

NetApp

Cold Data Tier as Data Source

Data moved from cold into hot

Source Data in Cold Storage

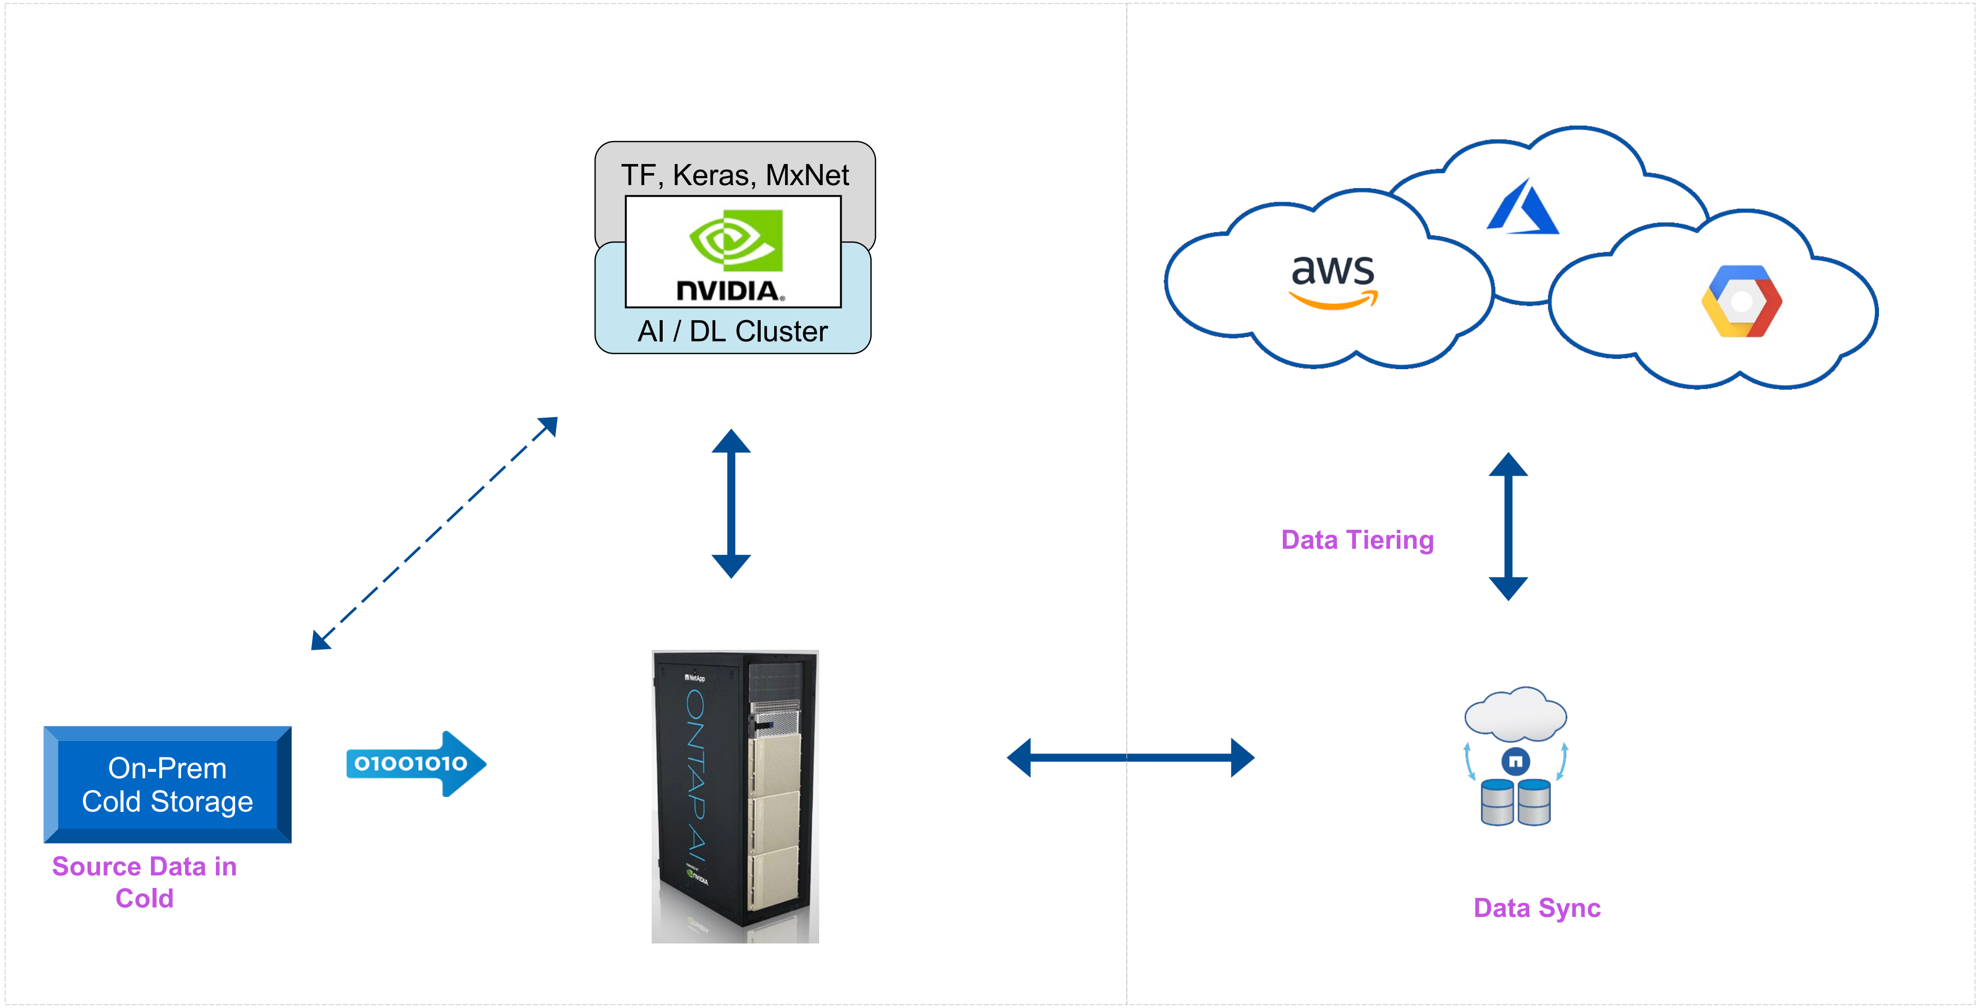

NetApp

Federating ML & DL

Unifying Machine Learning & Deep Learning

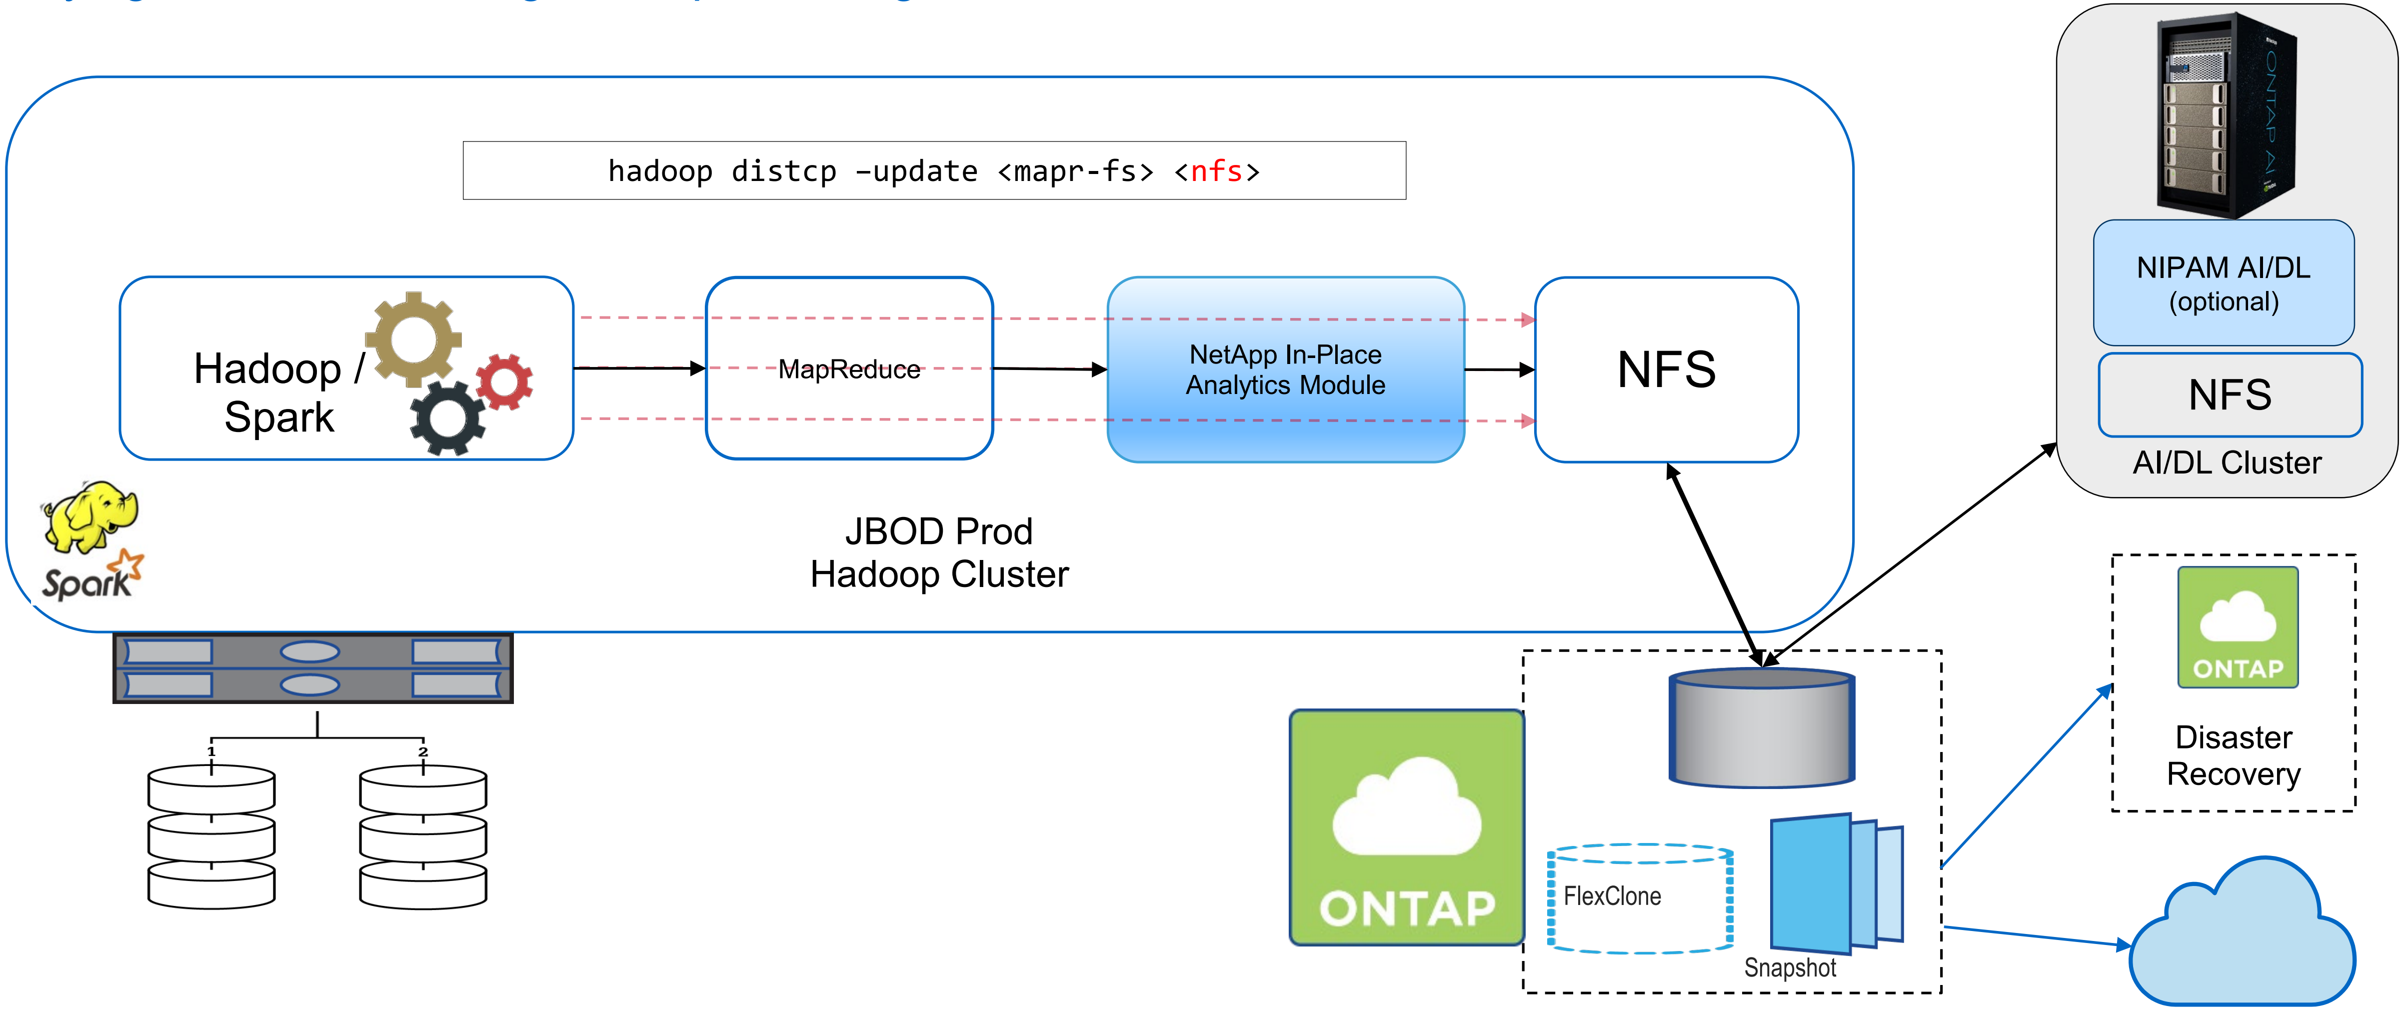

NetApp

In-Place Data Pipeline Federating ML & DL

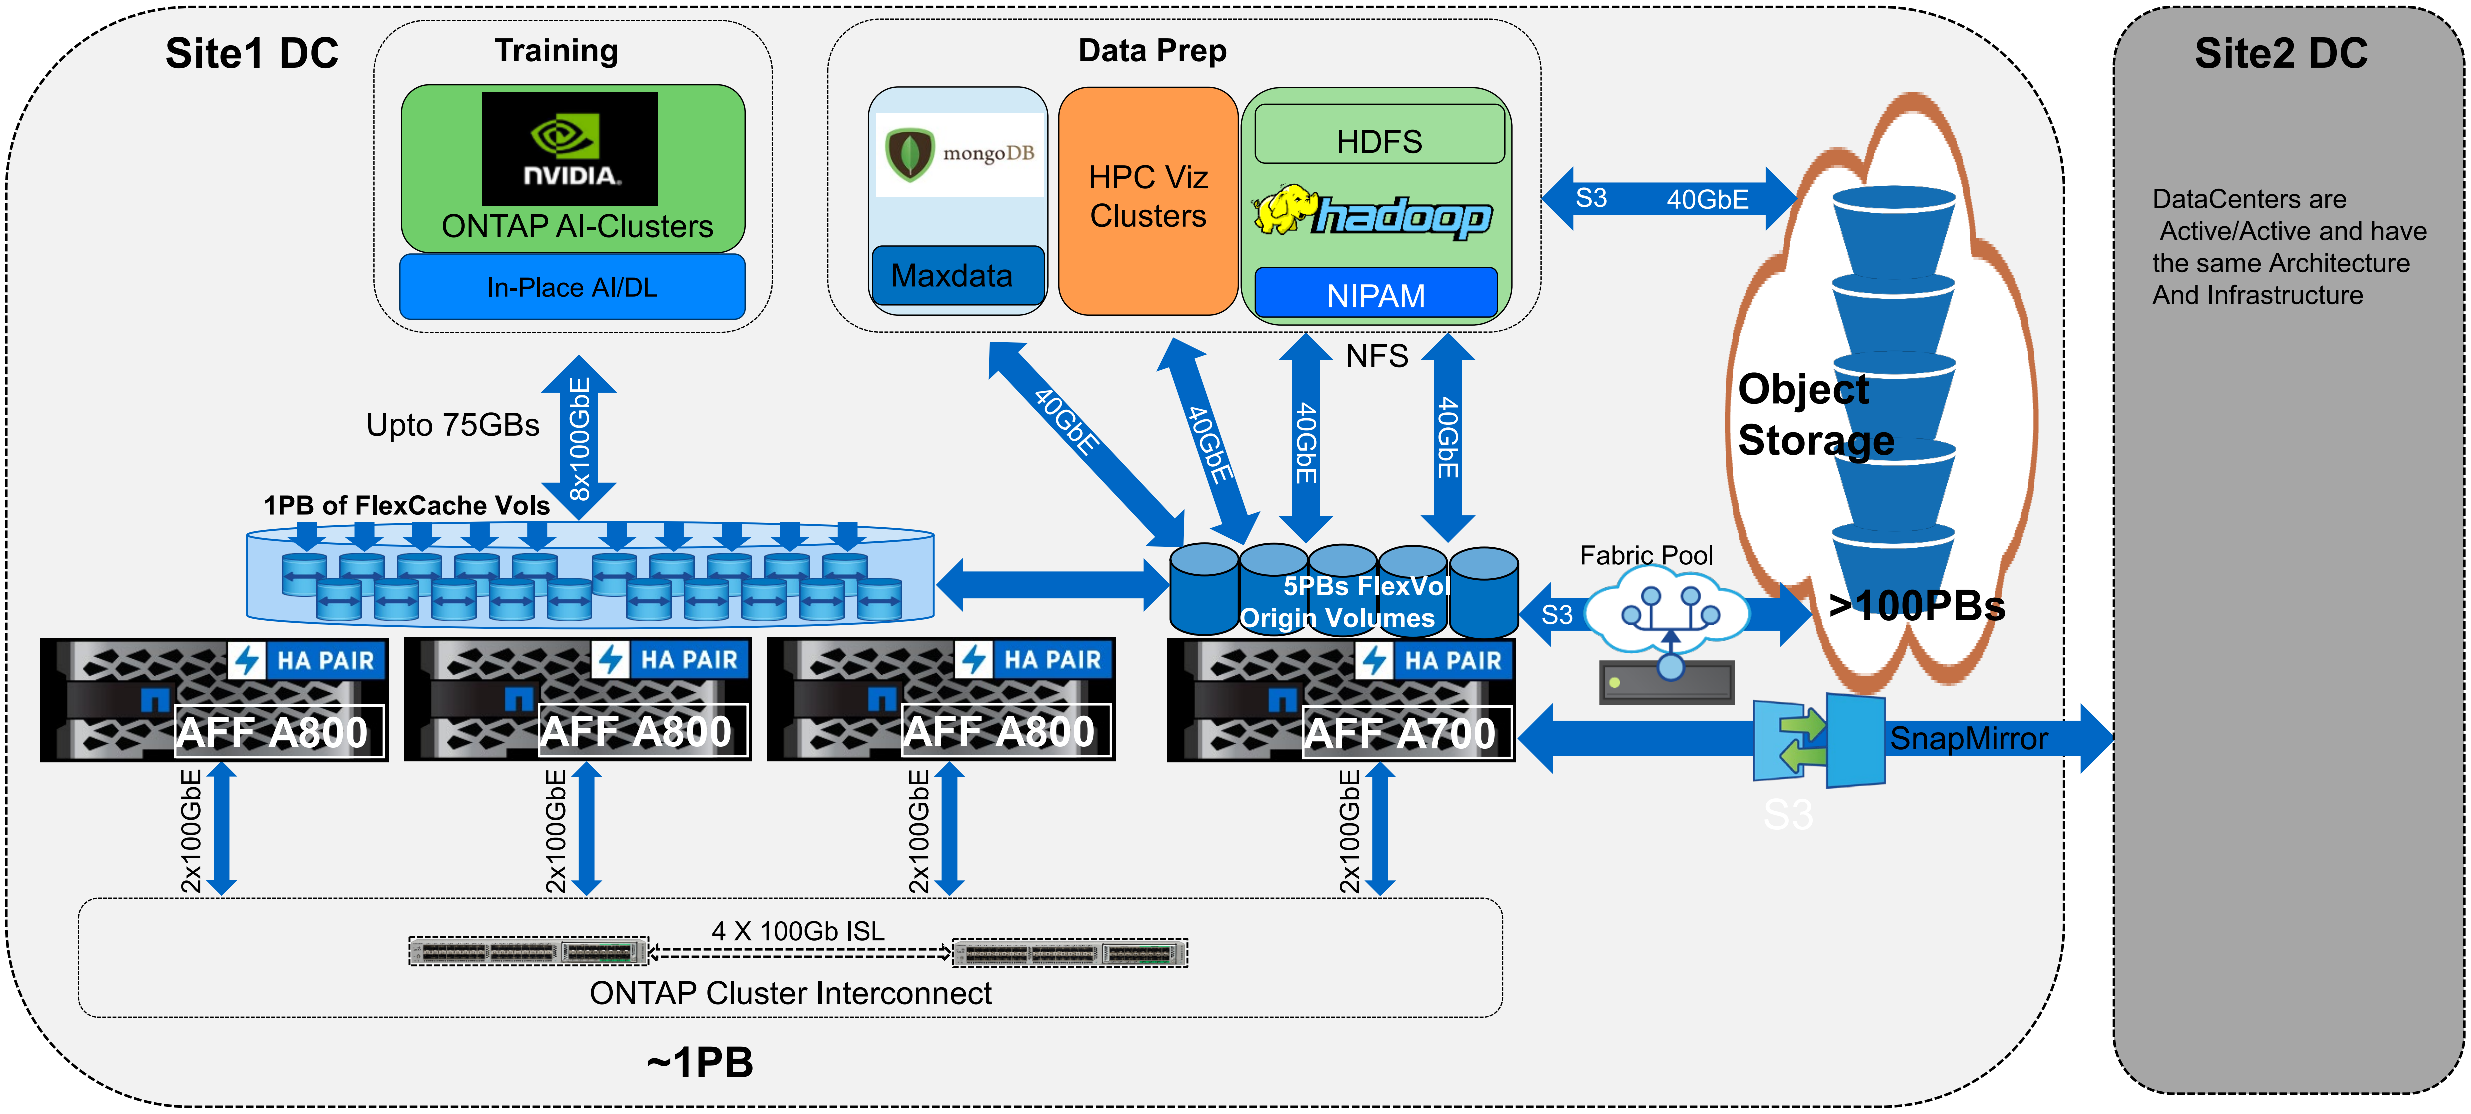

NetApp

Software Platforms for AI

H2O.ai with ONTAP AI

NetApp partner for ML platform

H2O.ai is transforming the use of AI with software with its category-creating visionary open source machine learning platform.

Key Features:

- Predictive analytics with visualization
- Build More Models in Less time
- Machine Learning: supervised and unsupervised algorithm

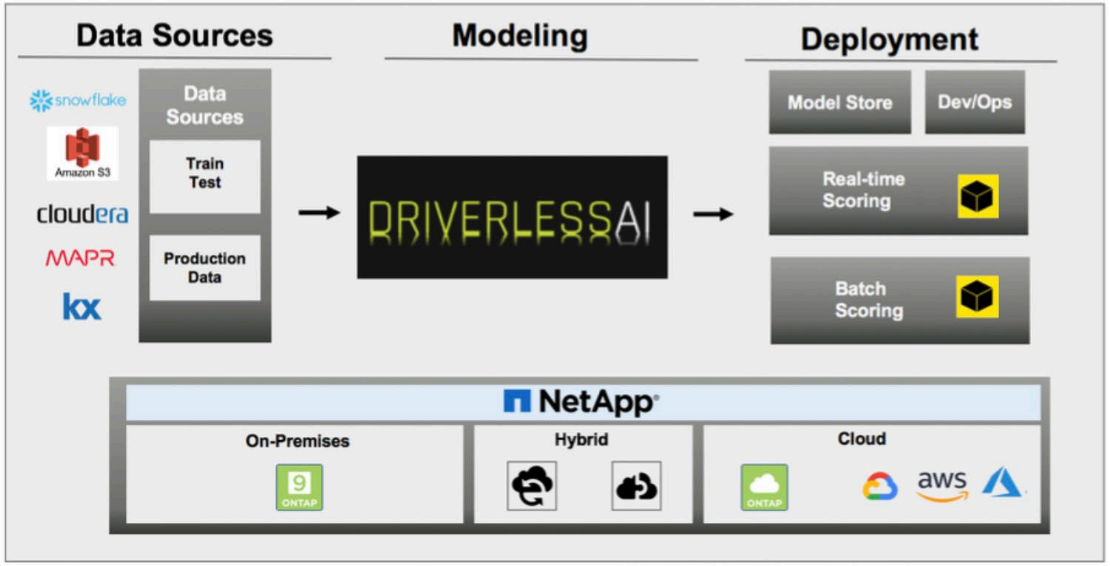

Driverless AI for the enterprise using H2O on ONTAP AI

H2O.ai and ONTAP AI Joint Solution

- Tested and validated the H2O software running on ONTAP AI
- Eliminates design complexities and guesswork with a solution brief.

Why ONTAP AI for H2O.ai ?

- **Deliver AI at scale:** Scale from zero to 100TB deployments in a matter of seconds with a simple, easy-to-use cloud interface.
- **Reduce Risk:** Instantaneous snapshots & cloning allow data scientists to experiment with datasets without risking data loss.
- **Dynamically change service levels:** Enhance performance or reduce OpEx with the ability to dynamically change storage service levels.
- **Build more models in less time:** Reduce the time it takes to develop accurate, production-ready models by automating time-consuming data science tasks.
- **Train across hybrid cloud :** Seamlessly move H20 Driverless AI workloads between on-premises ONTAP AI infrastructure and Cloud Volumes Service.

Allegro.ai with ONTAP AI

NetApp partner for DL platform

Allegro.ai offers the first true end-to-end AI product life-cycle management solution with a focus on DL applied to computer vision.

Key Features:

- Automated annotation
- Experiment management + data management
- Continuous learning

Semantic segmentation experiment using Allegro on ONTAP AI

Allegro and ONTAP AI Joint Solution

- Tested and validated the Allegro software on ONTAP AI.
- Eliminates design complexities and guesswork with a validated reference architecture.

Why ONATP AI for Allegro?

- **Accelerate time to completion:** ONTAP AI's high-throughput AFF and NVIDIA GPUs enable you to train more models in less time while extracting maximum performance from your hardware.
- **Increased GPU utilization:** ONTAP AI helps cut down on GPU idle time by keeping the GPUs engaged more often.
- **Simplified deployment:** ONTAP AI's prescriptive architecture eliminates design complexities and enables independent scaling of compute and storage.
- **Fast and big cache size:** ONTAP AI provides large pools of flash storage that can be used for caching the data further reducing the overall job completion times.

Allegro + NetApp

Optimizing Deep Learning Pipelines At Scale

Allegro.ai

Deep Learning computer vision & sensor fusion platform

allegro.ai helps data scientists / engineers manage & operationalize the full lifecycle of deep learning - from data-set creation to model post deployment self-learning.

Allegro.ai

- Fully integrated end2end platform
- Data Management / Experiment Management / Resource Management

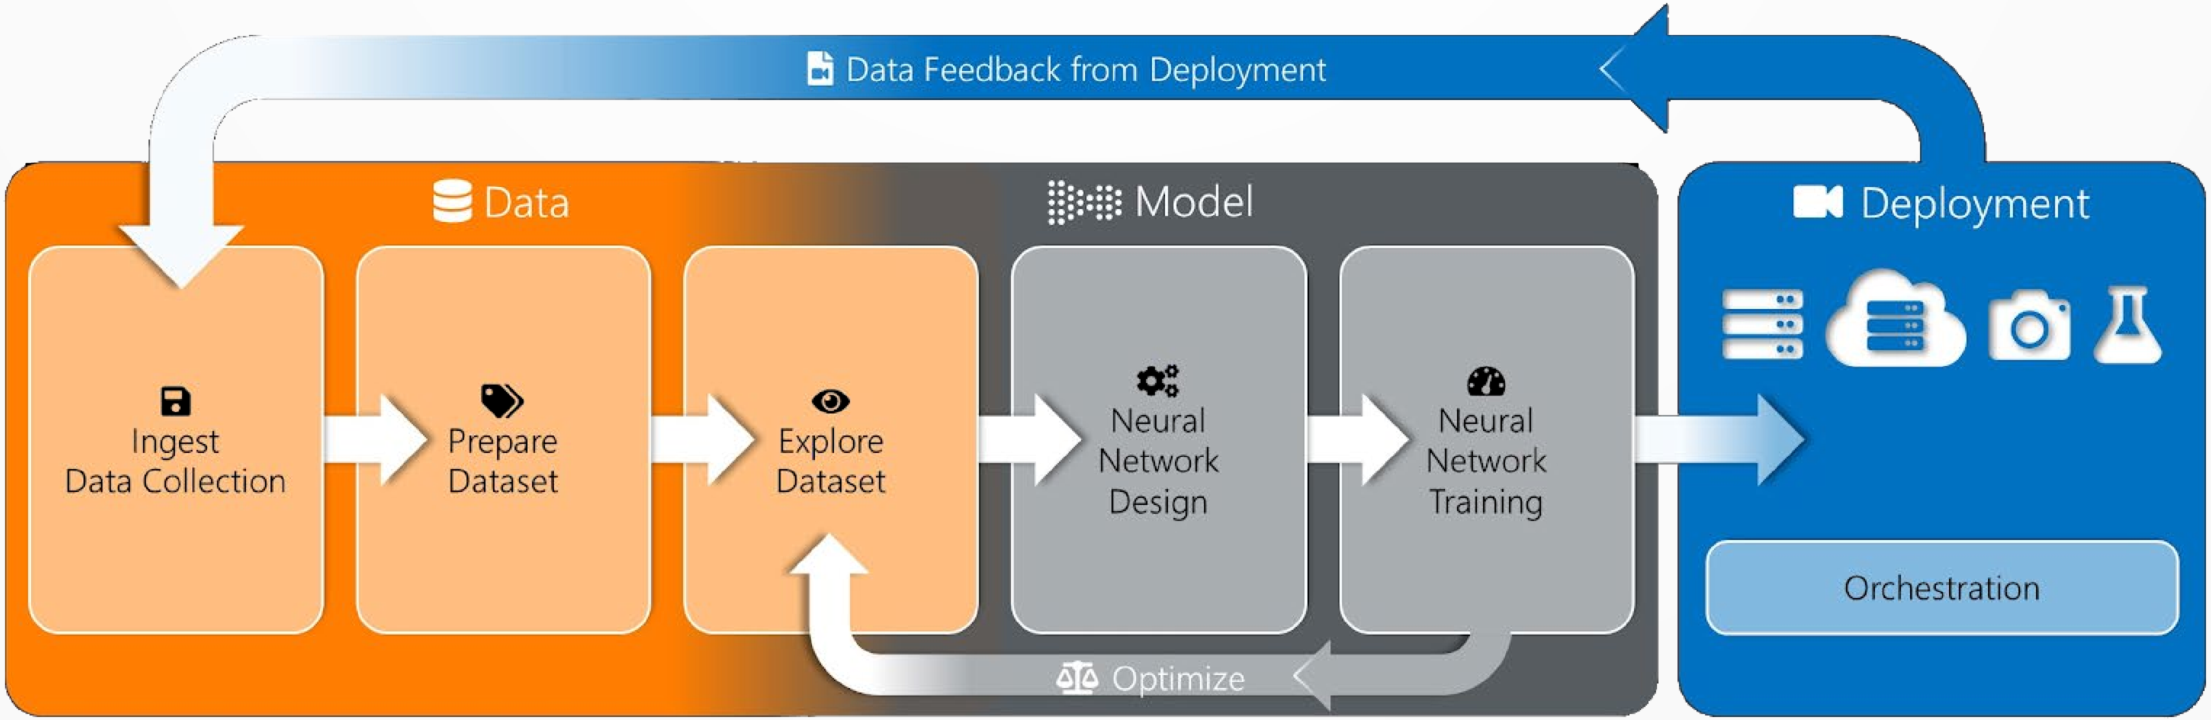

Allegro Data Abstraction

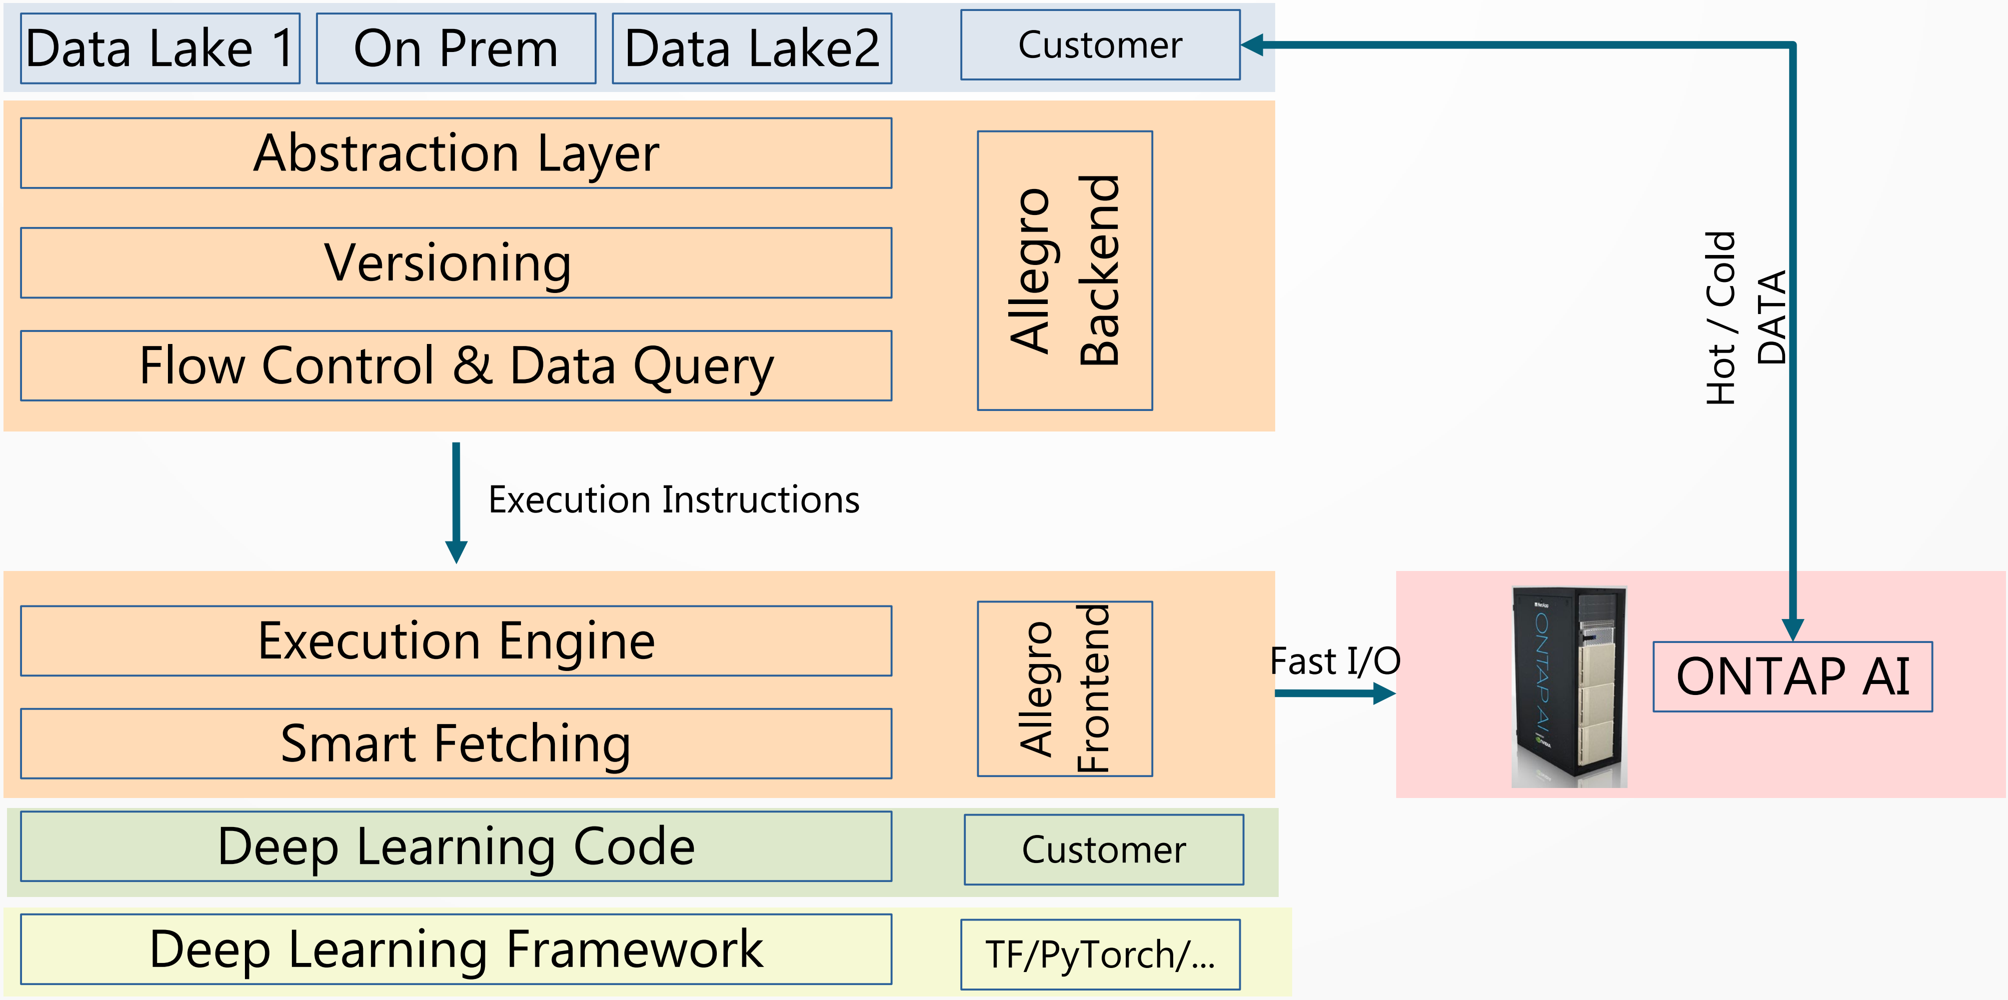

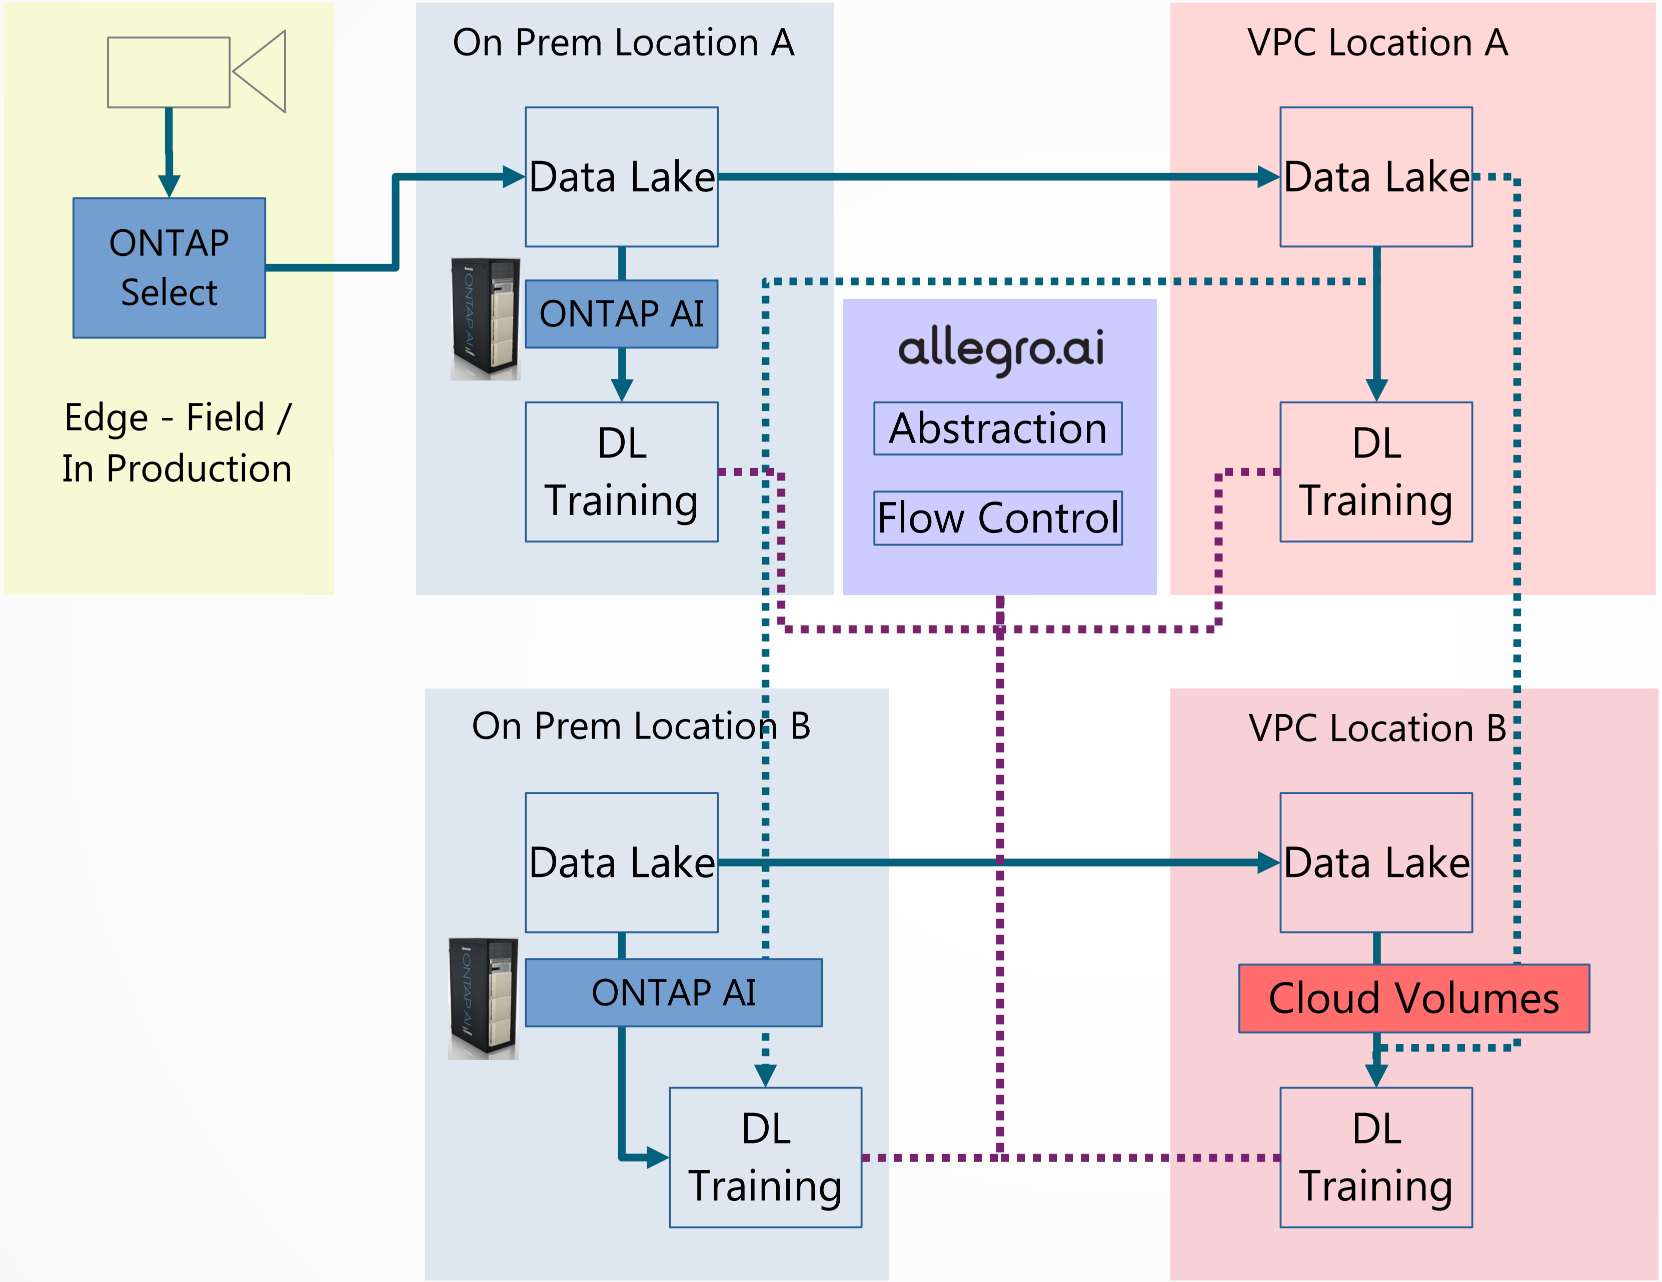

Multi-Site
Real World
Optimized DL
Pipeline
Allegro + NetApp

Edge to Core to Cloud

Seamless data management

Edge | Core | Cloud

Ingest | **Data prep.** | **Training cluster** | **Deployment** | **Analysis / Tiering**

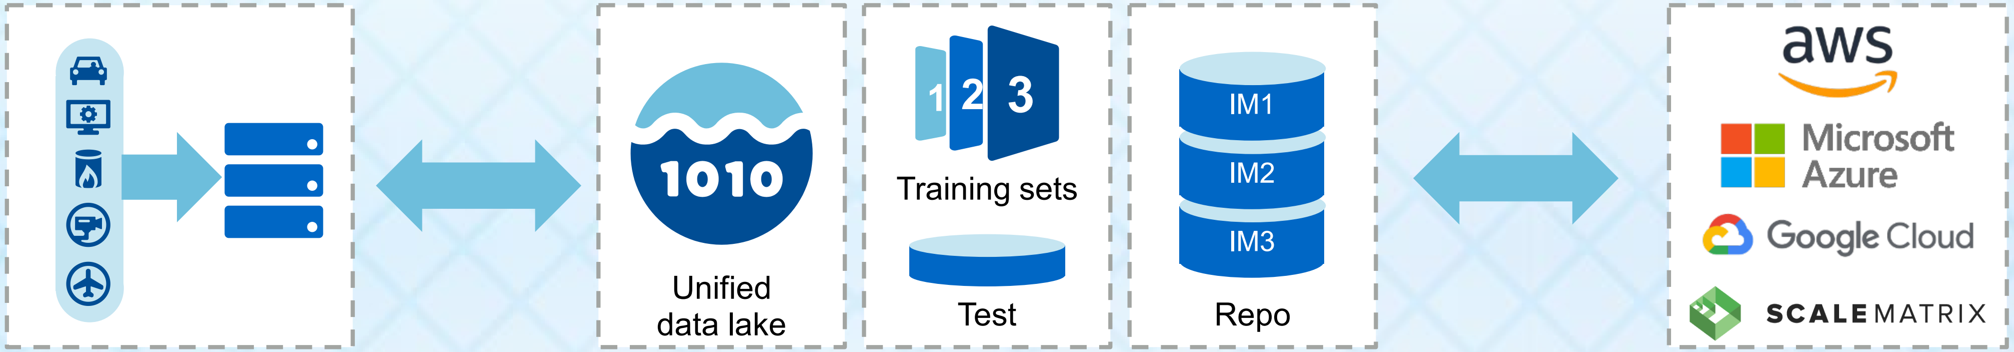

Edge – Ingest
- Data collection
- Edge-level AI

Core – Data prep.
- Aggregation
- Normalization

Core – Training cluster
- Exploration
- Training

Core – Deployment
- Deployment
- Model serving

Cloud – Analysis / Tiering
- Cloud AI (GPU instances)
- Data tiering

ONTAP Select
Vector Data

CloudSync
SnapMirror
Kafka, NiFi

ONTAP AI
FlexPod AI
FlexPod Inference

FabricPool
CloudSync
NDAS

Cloud Volumes
ScaleMatrix ONTAP AI

NetApp

AI Public Resources

netapp.com/ai

#NetAppAI

Collateral

- ONTAP AI Reference architecture NVA-1121-design
- ONTAP AI Deployment guide NVA-1121-deploy
- Building a Data Pipeline for Deep Learning WP-7299
- Edge to Core to Cloud white paper WP-7271
- AI with GPUs on AWS & Cloud Volumes Service TR-4718
- Scalable AI Infrastructure WP-7267
- Designing a Data Pipeline for Your AI Workflows WP-7264
- Solution brief SB-3939
- IDC Technology Spotlight paper
- Cambridge Consultants success story

Recent Blogs

- Your Guide to Everything NetApp at GTC 2019
- AI Across Industries: Manufacturing, Telecom & Healthcare
- How to Configure ONTAP AI in 20 Minutes with Ansible
- Unifying Machine Learning and Deep Learning
- Bridging the CPU and GPU Universes
- Is Your Infrastructure Ready for AI Workflows in Production?
- Accelerate I/O for Your Deep Learning Pipeline
- Addressing AI Data Lifecycle Challenges with Data Fabric
- Choosing an Optimal Filesystem for the AI Pipeline
- Five Advantages of ONTAP AI for AI and Deep Learning
- Deep Dive into ONTAP AI Performance and Sizing

Thank You!

NetApp