

Distributed Backends for Machine Learning-based Simulation Solvers

Allen Mao
Kliment Minchev

CTO Office

© Copyright 2023 ANSYS, Inc.



Background

Instant Prediction of High-Fidelity Simulation

Industry landscape:

- Apple, Nvidia, Google, Intel
- DesignAI by Altair, Nvidia Modulus / SimNet + Omniverse

Cloud HPC integration

Cloud-first Machine Learning Platform

1000x or faster simulation

Distributed Computing + Microservices Architecture

- Input simulation data
- Build a deep learning model
- Use existing datasets, models from library

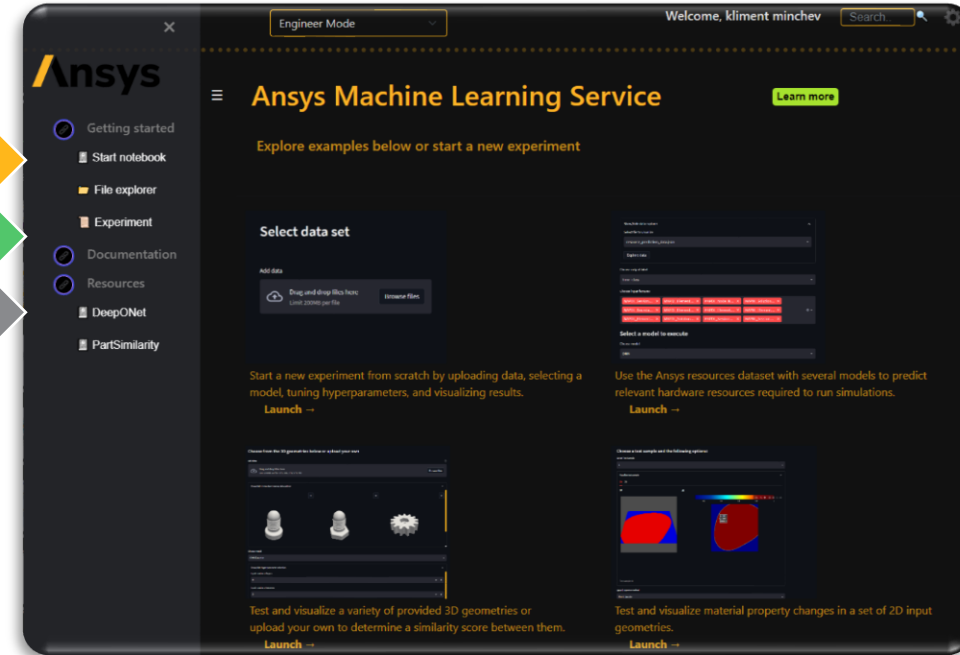
End-to-End Deployment

Focus on:
simulation-based datasets
visualization heavy

Peers:

- AWS SageMaker
- Azure ML

Improve Ansys ML developer productivity



Why Ansys ▾ Products & Services ▾ Learn ▾



Home > Blog > Prepare for the Machine Learning Revolution with Emerging MLaaS Capabilities from Ansys

ANSYS BLOG


Prepare for the Machine Learning Revolution with Emerging MLaaS Capabilities from Ansys

FEBRUARY 2, 2023



Machine learning (ML) is, justifiably, receiving a lot of attention today. ML is helping companies in every industry identify performance

Public datasets



Private datasets



Synthetic datasets



Models + Code



Motivation for a Backend

Composable Machine Learning Simulator (CoMLSim)

Design choices

- Pipeline of Neural Networks
- Built-in optimization
- Generated output and file formats

Intricacies

Solver: BCs, volumetric source terms, surface solution fields, geometry
Trains on and evaluates full-fidelity steady-state PDE solutions
R&D in progress

Containerization & Deployment

Distributed System

- Server + requests
- Functional isolation
 - Data processing
 - Model training
 - Inference

Microservice Architecture

- Containerization
- Interoperability
- Scalability and compute resource requirements (RAM, GPU)

Cloud Resources

- AWS compute instances
- Storage, Container registry, Databases, Queue
- SageMaker, Kubernetes

Backend Consumption & UX

Frontend

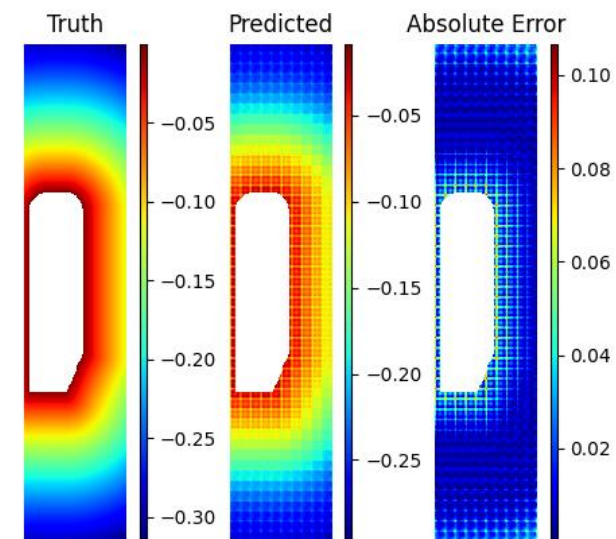
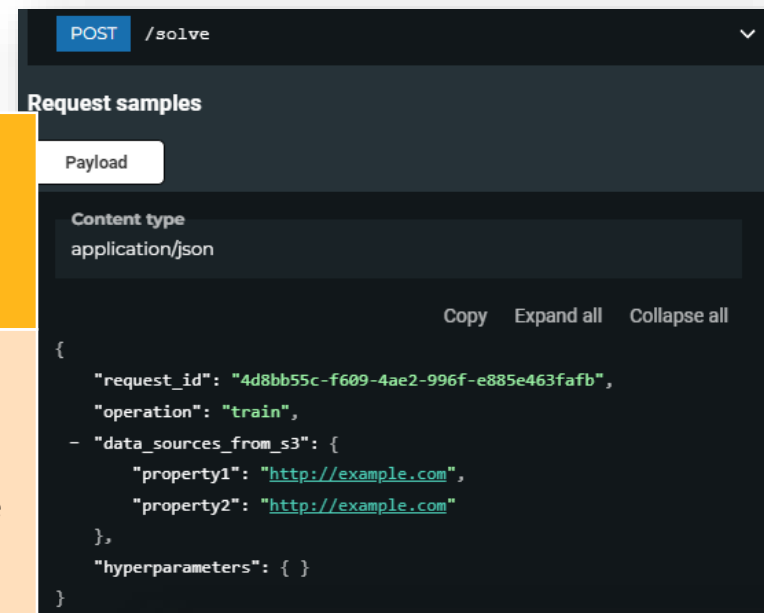
- No-Code Web Platform
- Notebook-style code execution

RESTful API

- Endpoints
- Payload schema
- Authentication

ML-based Simulation Paradigm

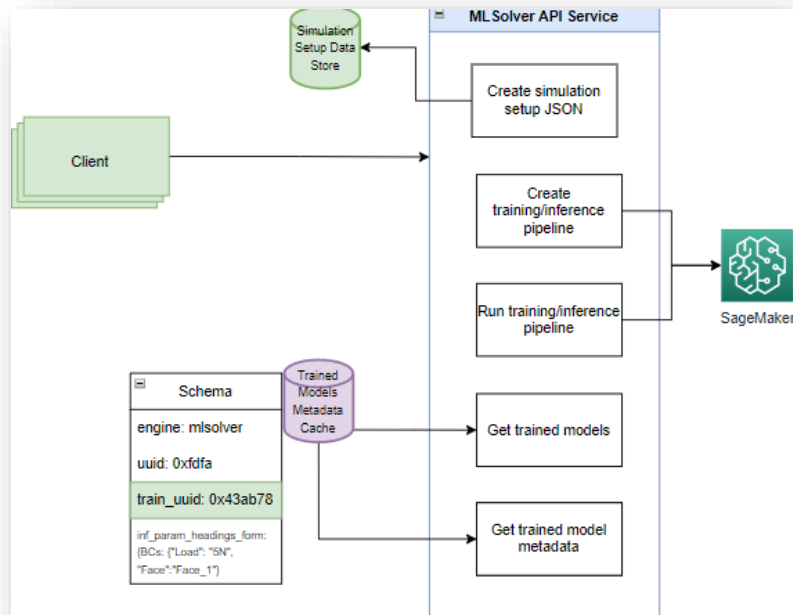
- Producer & Consumer Roles
- Custom workflow creation: training, model tuning, analysis, investigation
- Inference as a Service



Backend Architecture Roadmap

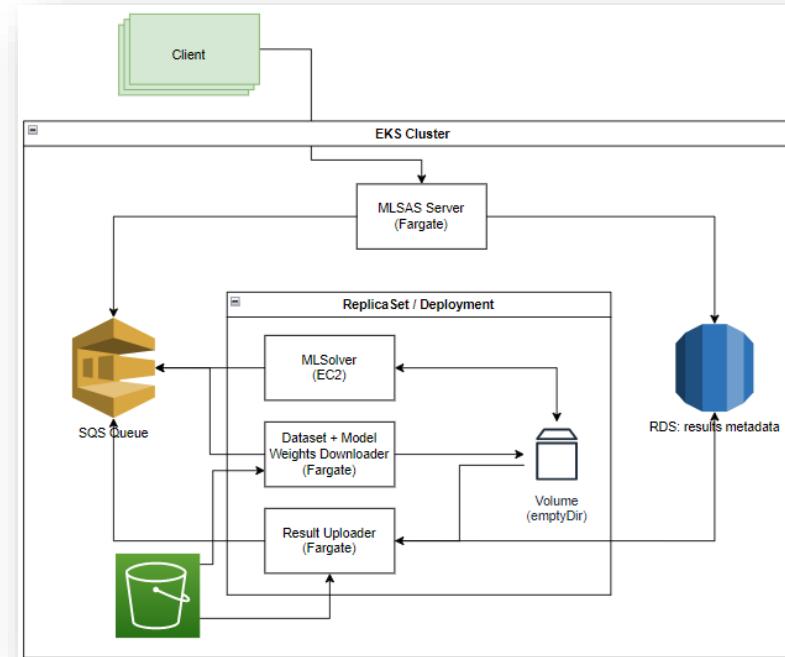
Job Submission and Code Execution

- Server responds to requests
- AWS SageMaker Jobs (compute instance, S3 download/upload): training & inference
- S3 input: case files, S3 output: results, DB: metadata



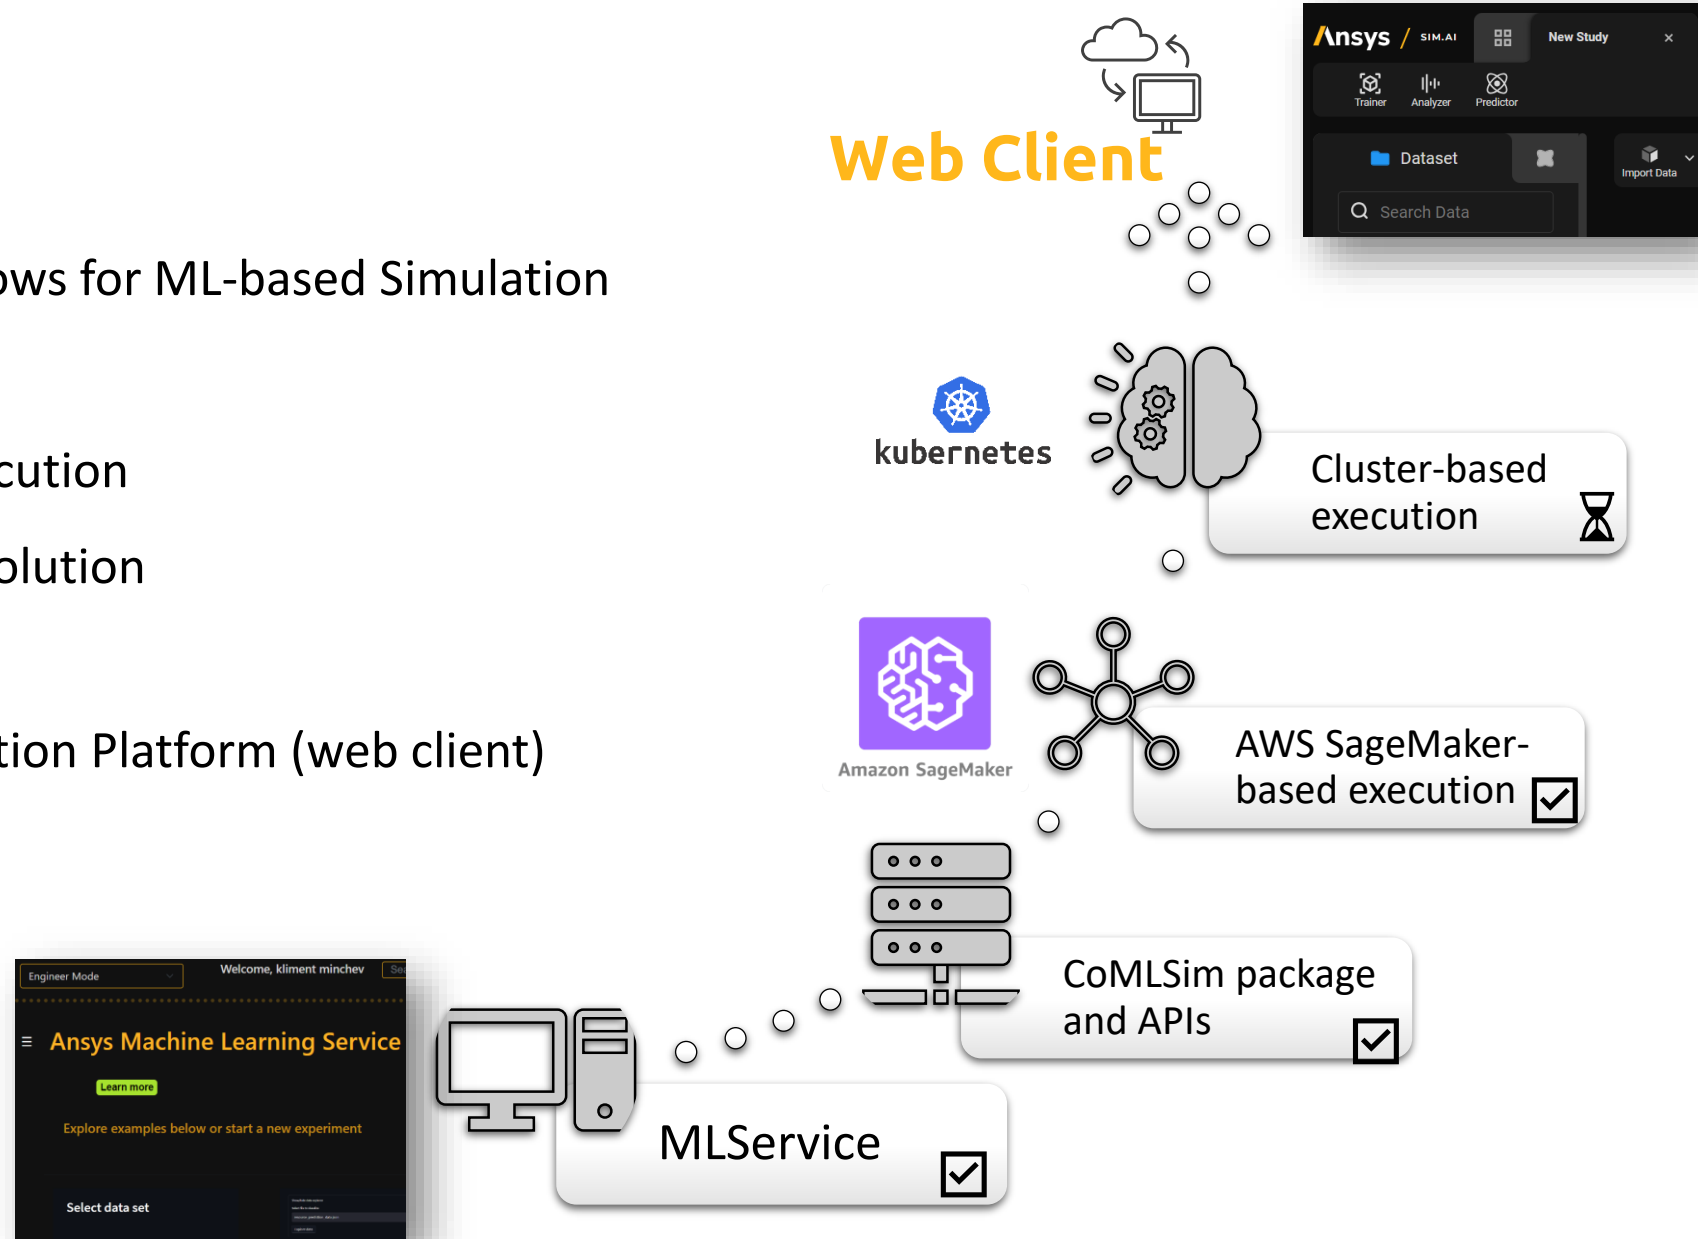
Service-based

- Server responds to requests
- AWS Fargate/ECS/Kubernetes cluster
- Pod in cluster executes compute job (highly available)
- S3 input: case files, S3 output: results, DB: metadata

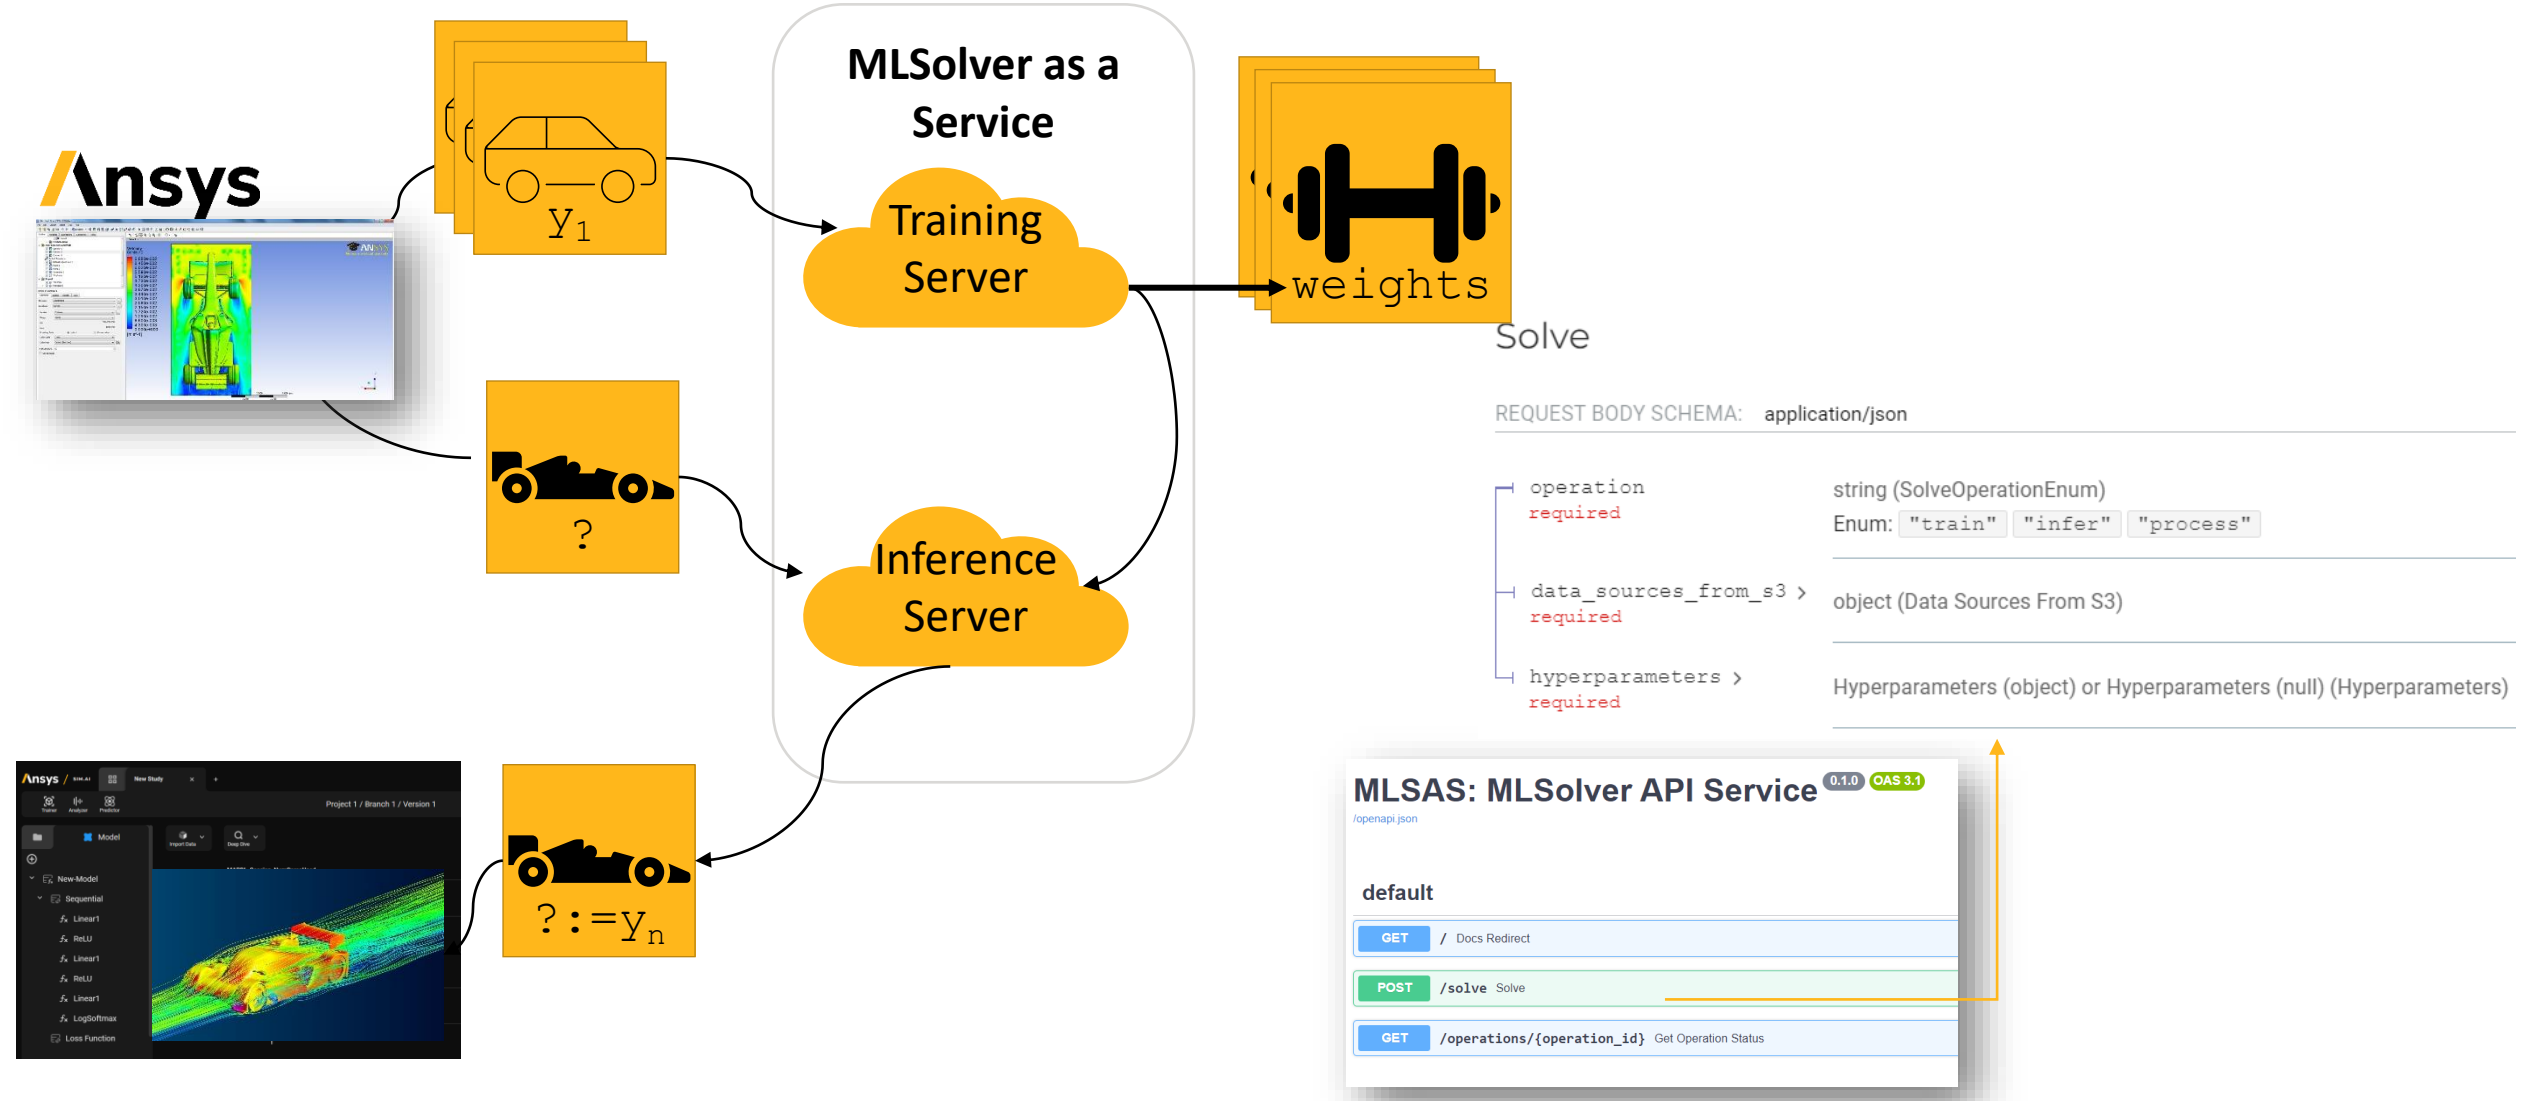


Project Evolution

- MLService: configurable workflows for ML-based Simulation
- CoMLSim python package
- AWS SageMaker: job-based execution
- Cluster-based highly-available solution
 - ECS, EKS
- Ansys Machine Learning Simulation Platform (web client)



Training as a Service, Inference as a Service



The Stack and its Properties

Raw Case Files

Data
Processor

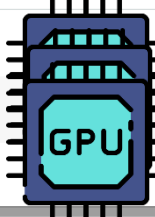
Machine
Learnable Data

Trainer

Machine Learning
Models + Weights

Inference
Solver

Vector Field
Predictions

Accelerated Computing		vCPU	Memory
ml.p3.2xlarge		8	61 GiB
ml.p3.8xlarge		32	244 GiB
ml.p3.16xlarge		64	488 GiB

MLSAS: MLSolver API Service 0.1.0 OAS 3.1

/openapi.json

API Request




Amazon SageMaker



Amazon
EC2



Demo Video



Summary

- Cloud backend and web API for the Composable Machine Learning Simulator (CoMLSim)
 - Next generation ML-based simulation
 - Universally accessible
- Evolving CoMLSim package, maturing cloud backend in parallel
 - Sophisticated, distributed (compute cluster), Infrastructure req's
- Compute jobs-based execution for initial rollout
 - Training as a Service, Inference as a Service

