



Acumos AI

Platform Overview, Releases and Use Cases

LFAI Day, Châtillons, France



Agenda

– Acumos AI

- What is Acumos AI
- Platform Overview
- Value and Benefits
- Community
- Architecture, CI/CD, OSS
- Clio Release

– Key Features & Use Cases

- Network AI
- ONAP Integration

- 5G/ORAN RIC Deployment
- Pluggable ML Workbench
- Collaboration, Federated Training
- Real Time Video Analytics

– LFAI Landscape & Acumos AI

– Machine Learning Stack

– Q&A

What is Acumos AI?



Open Source Platform to Harmonize AI Solutions



Distributed AI Marketplace

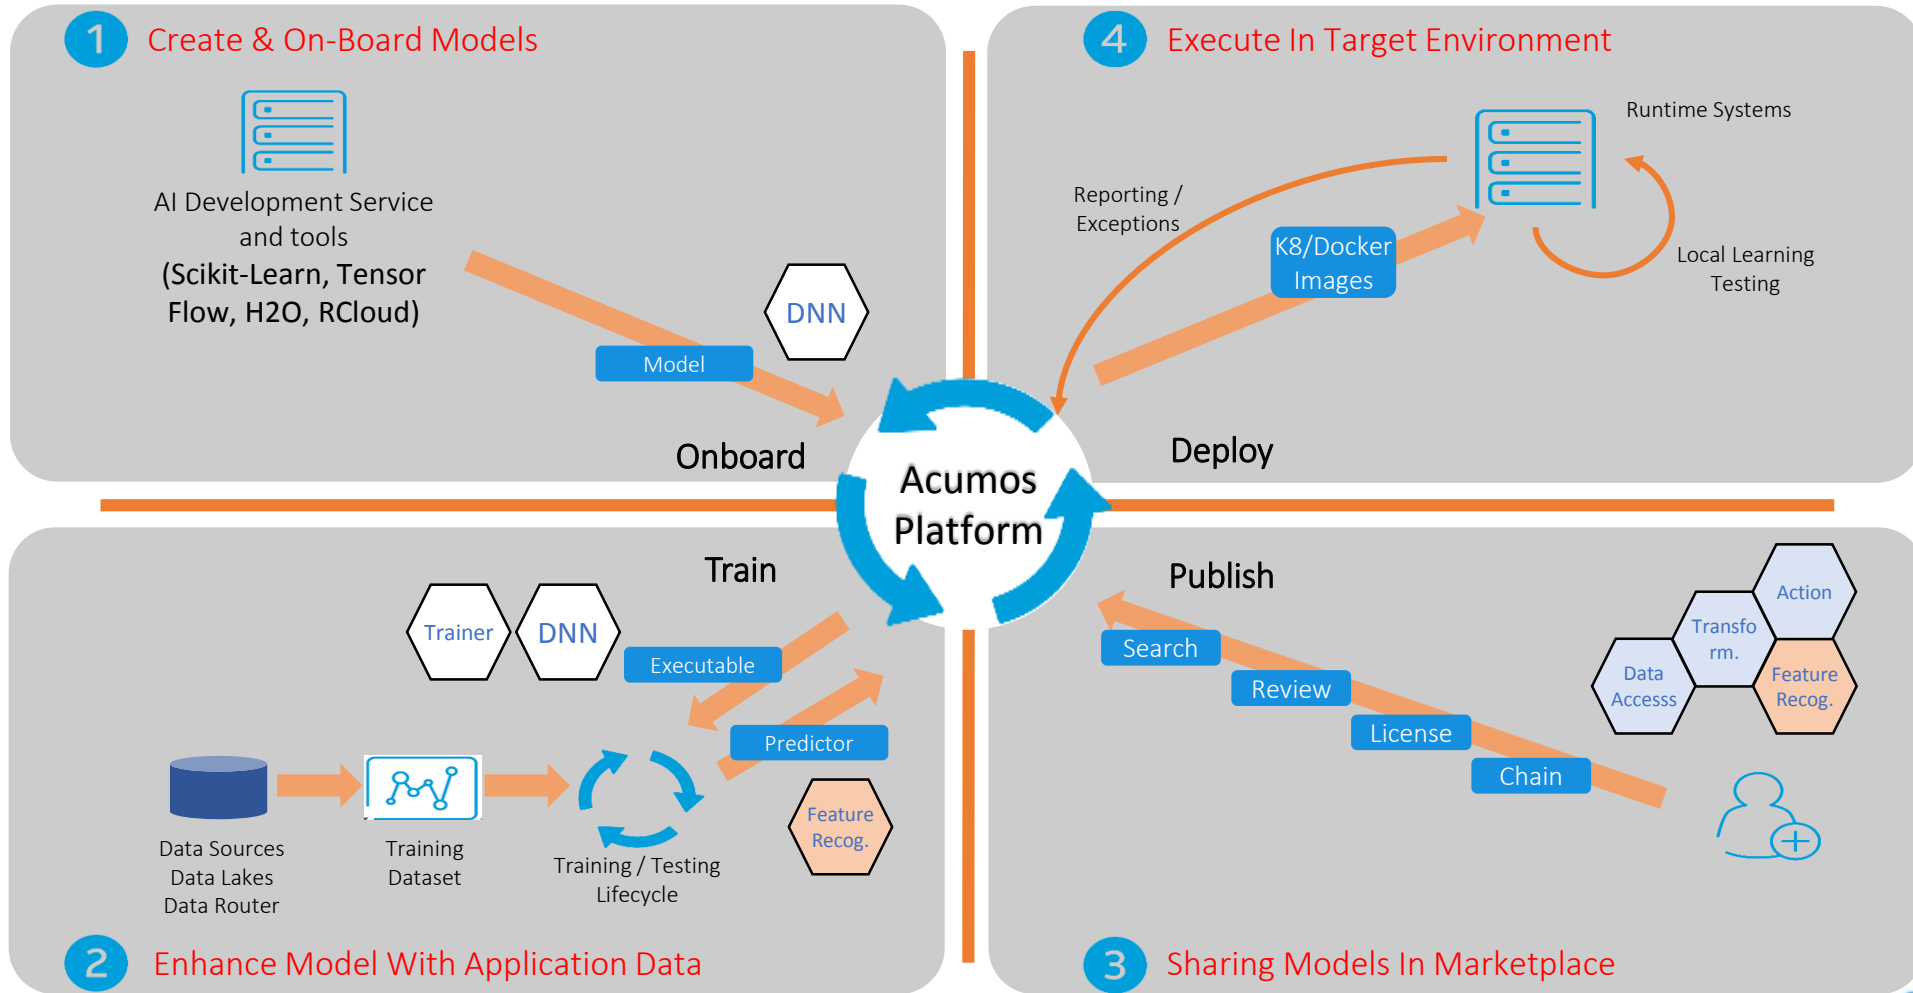


Interoperable Microservices



LFAI Ecosystem to Expedite Collaboration & Innovation

Platform Overview



Platform Value and Benefits





Models are containerized microservices which can be deployed where necessary to any application. Any system can reuse the logic if it supports Docker.

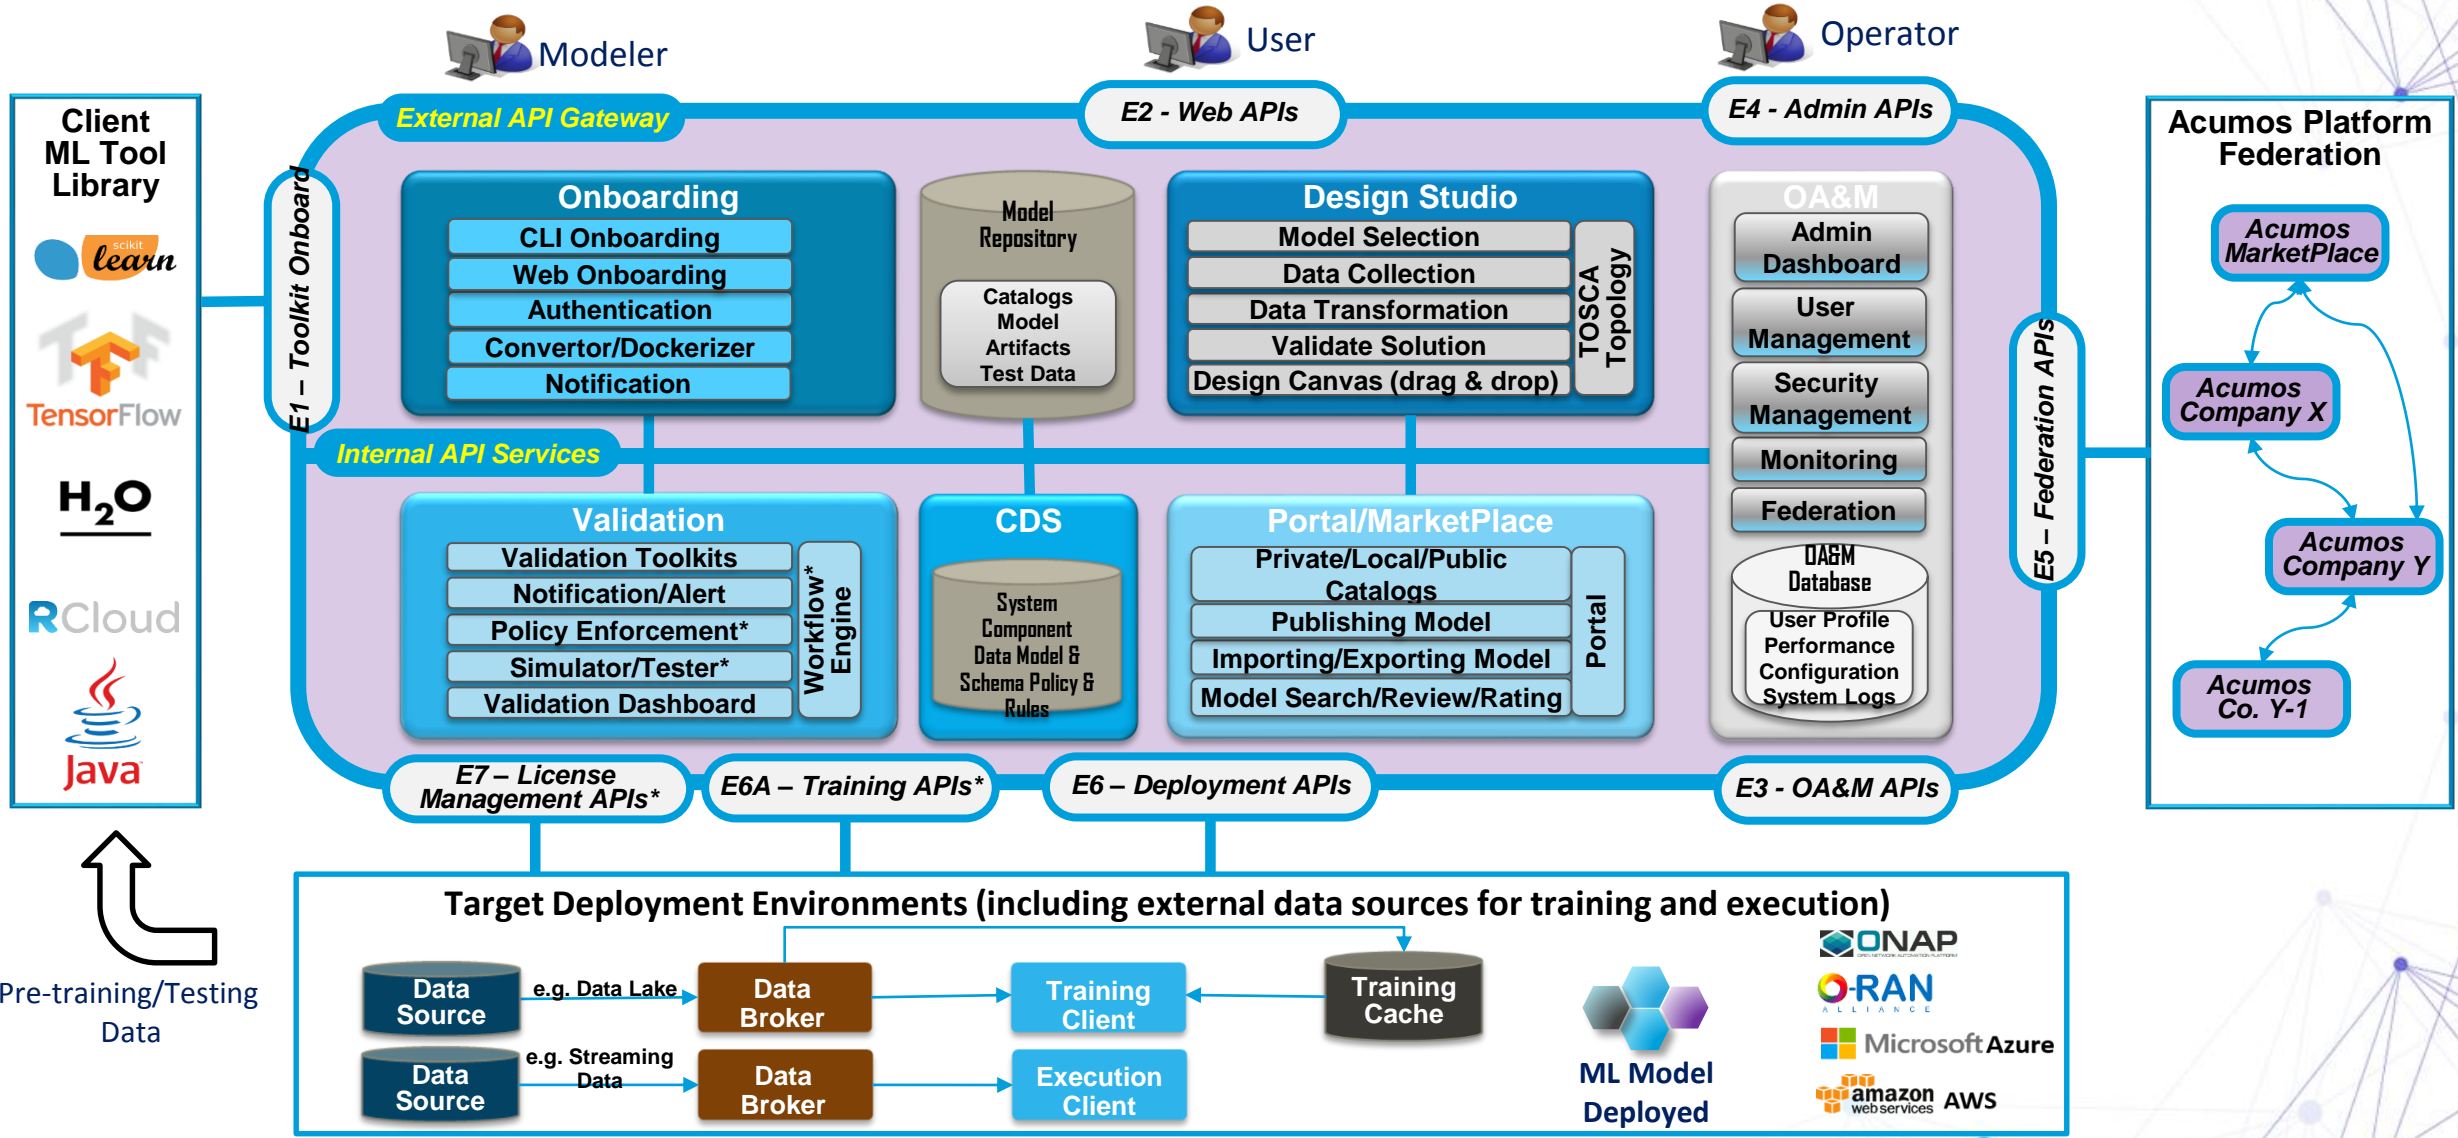
Models are adapted to different problems through training which can be done anywhere in order to comply with security principles (e.g. HIPAA) even if the model is to be executed elsewhere.

A data pipeline (acquire, select, combine, transform, etc.) can be used with any model so data access can be done once and used to build many solutions.

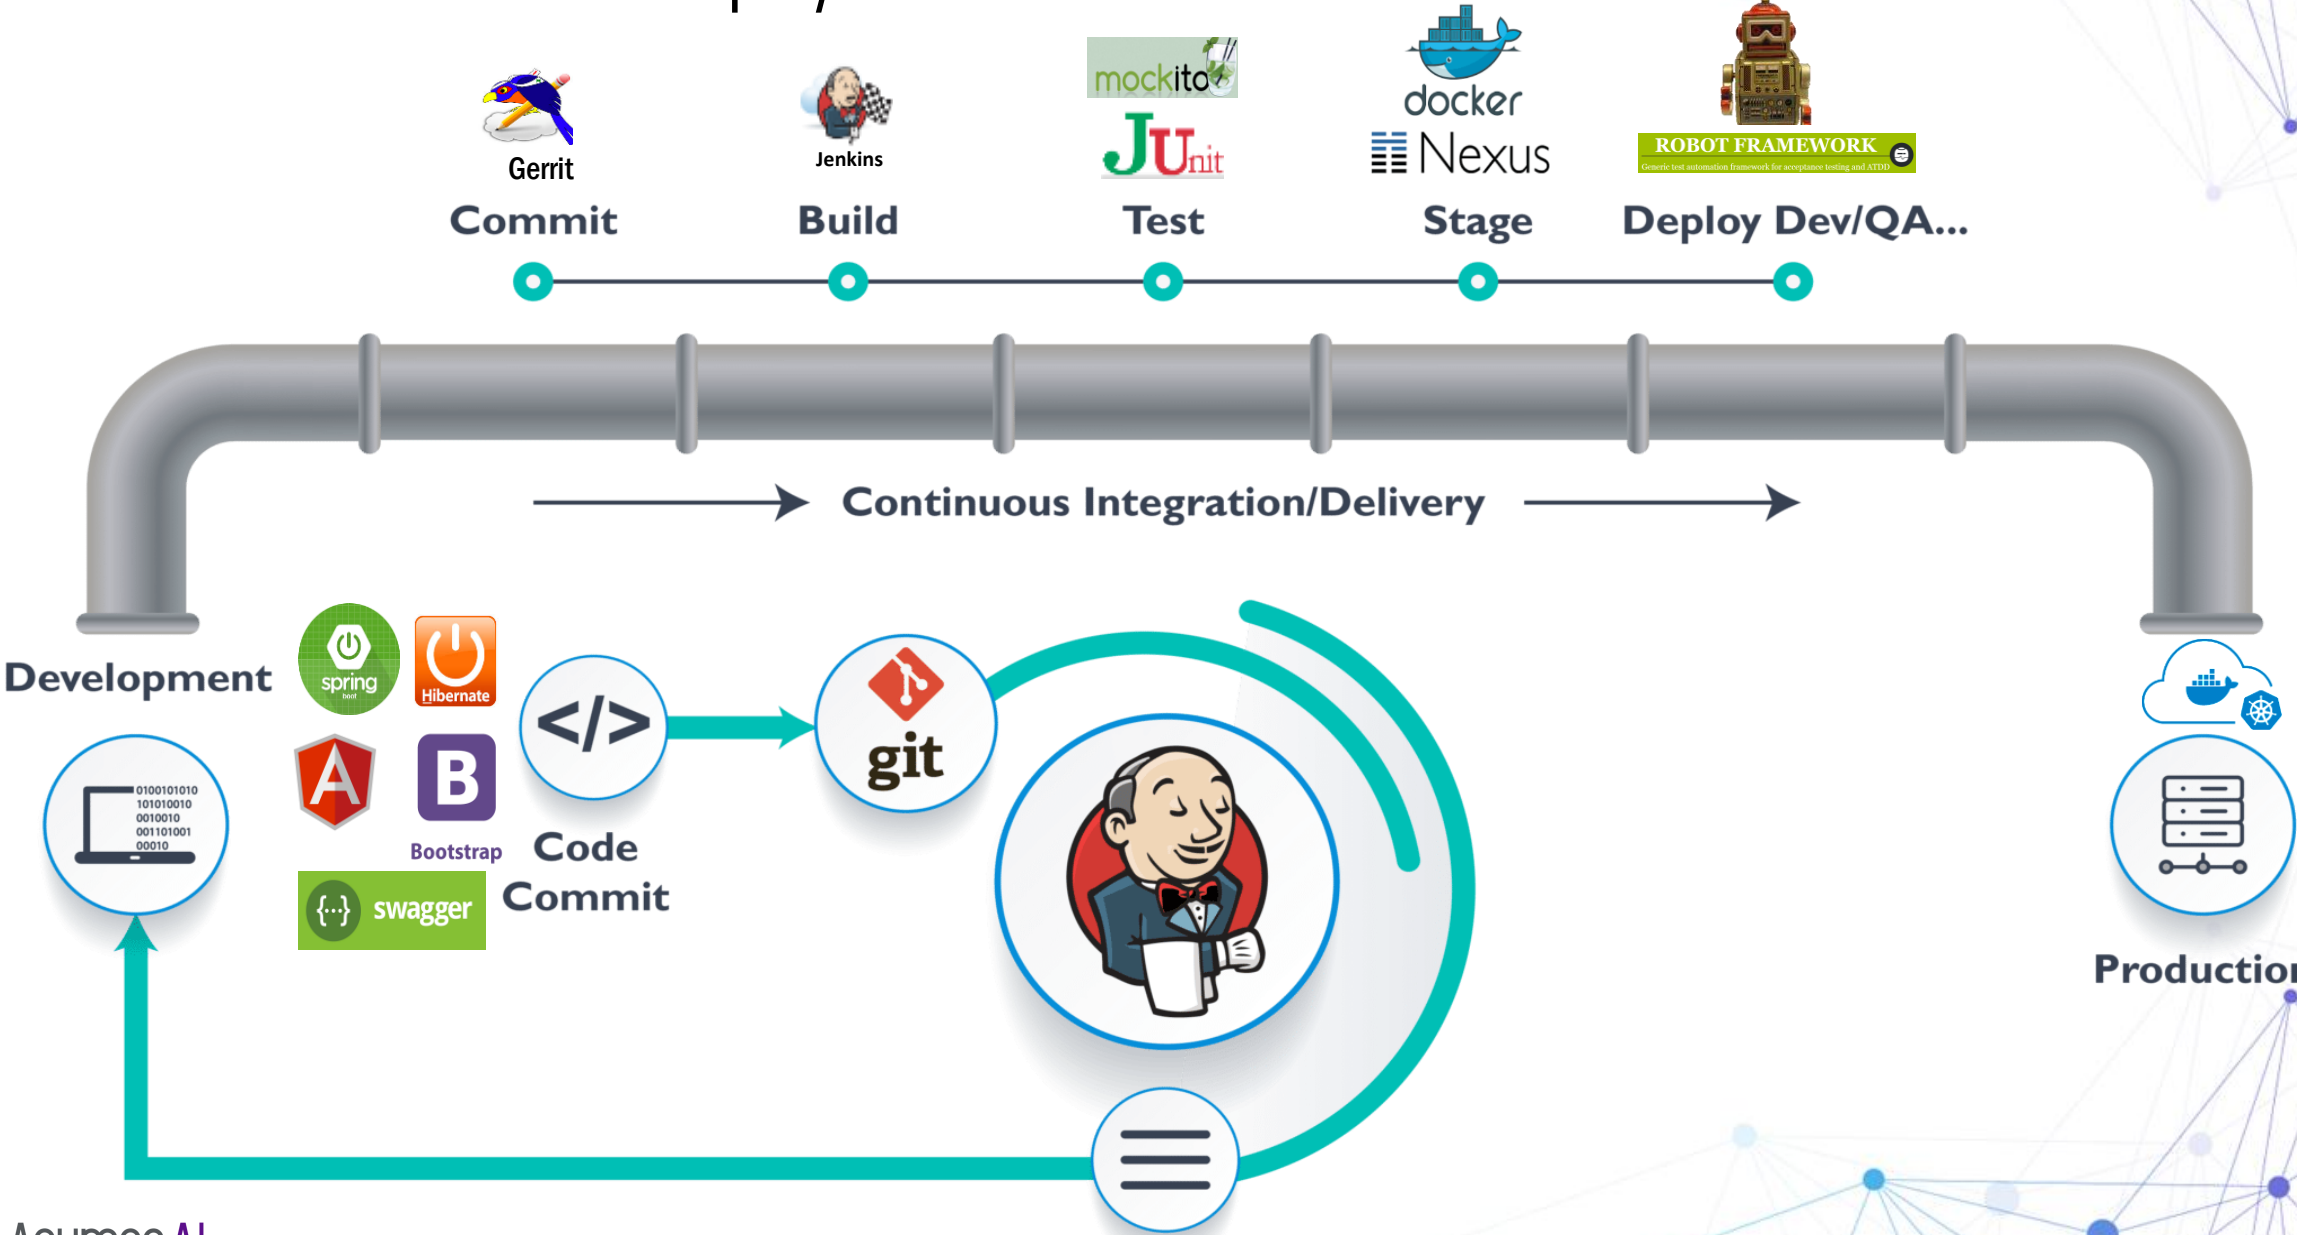
Linux Foundation Acumos AI Members





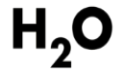





















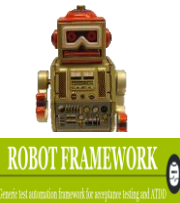
















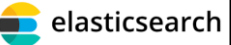


Acumos AI Platform Architecture



CI/CD - Test Automation and Deployment




Acumos AI Open Source Software Technologies and Tools

AI/ML Toolkits	Agile	Code	Data	Test	Code Analysis	Build Automation	Platform Deployment	Model Deployment	Logging Reporting
 TensorFlow  scikit-learn  H₂O  PC  RCloud  python™  java  jupyter  APACHE nifi	 JIRA  Confluence	 spring  hibernate  AB Bootstrap  swagger  ReadTheDocs  git  Gerrit	 MariaDB™  HIPPO  Sonatype Nexus	 JUnit  mockito  Se  ROBOT FRAMEWORK <small>Generate test automation framework for acceptance testing and TDD</small>  SoapUI <small>by SMARTBEAR</small>	 sonarqube  fossology  Sonatype Nexus Q	 Jenkins  maven  Nexus	 docker  kubernetes  kubernetes  White box  openstack  docker	 Microsoft Azure  kubernetes  Stack  elasticsearch  logstash  kibana	

Clio Open Source Release, Nov 2019

Features

- ONAP DCAE Integration
- O-RAN RIC/XApp Deployment
- ML Workbench, Notebooks
- Data/Training Pipelines, NiFi
- Model Ratings
- Predictor Scoring
- CI/CD for Kubernetes
- IPR Licensing Protection
- Cross Company Licensing
- Pluggable MLWB Framework
- Integrated Catalog Search
- Locality/International Design
- Platform Co-branding
- Support for C/C++ Client
- Onboarding Spark/Java



Key Stats (60%)

- 8 Projects
- 62 Repositories
- 1.27M+ lines of code
- 4K+ Commits
- 84 Contributors
- 78 Reviewers
- 33 Epics, 104 US, 28 Tasks
- 178 Jira Items
- 200+ Internal API Functional
- 50+ External API's
- 50%+ Code Coverage

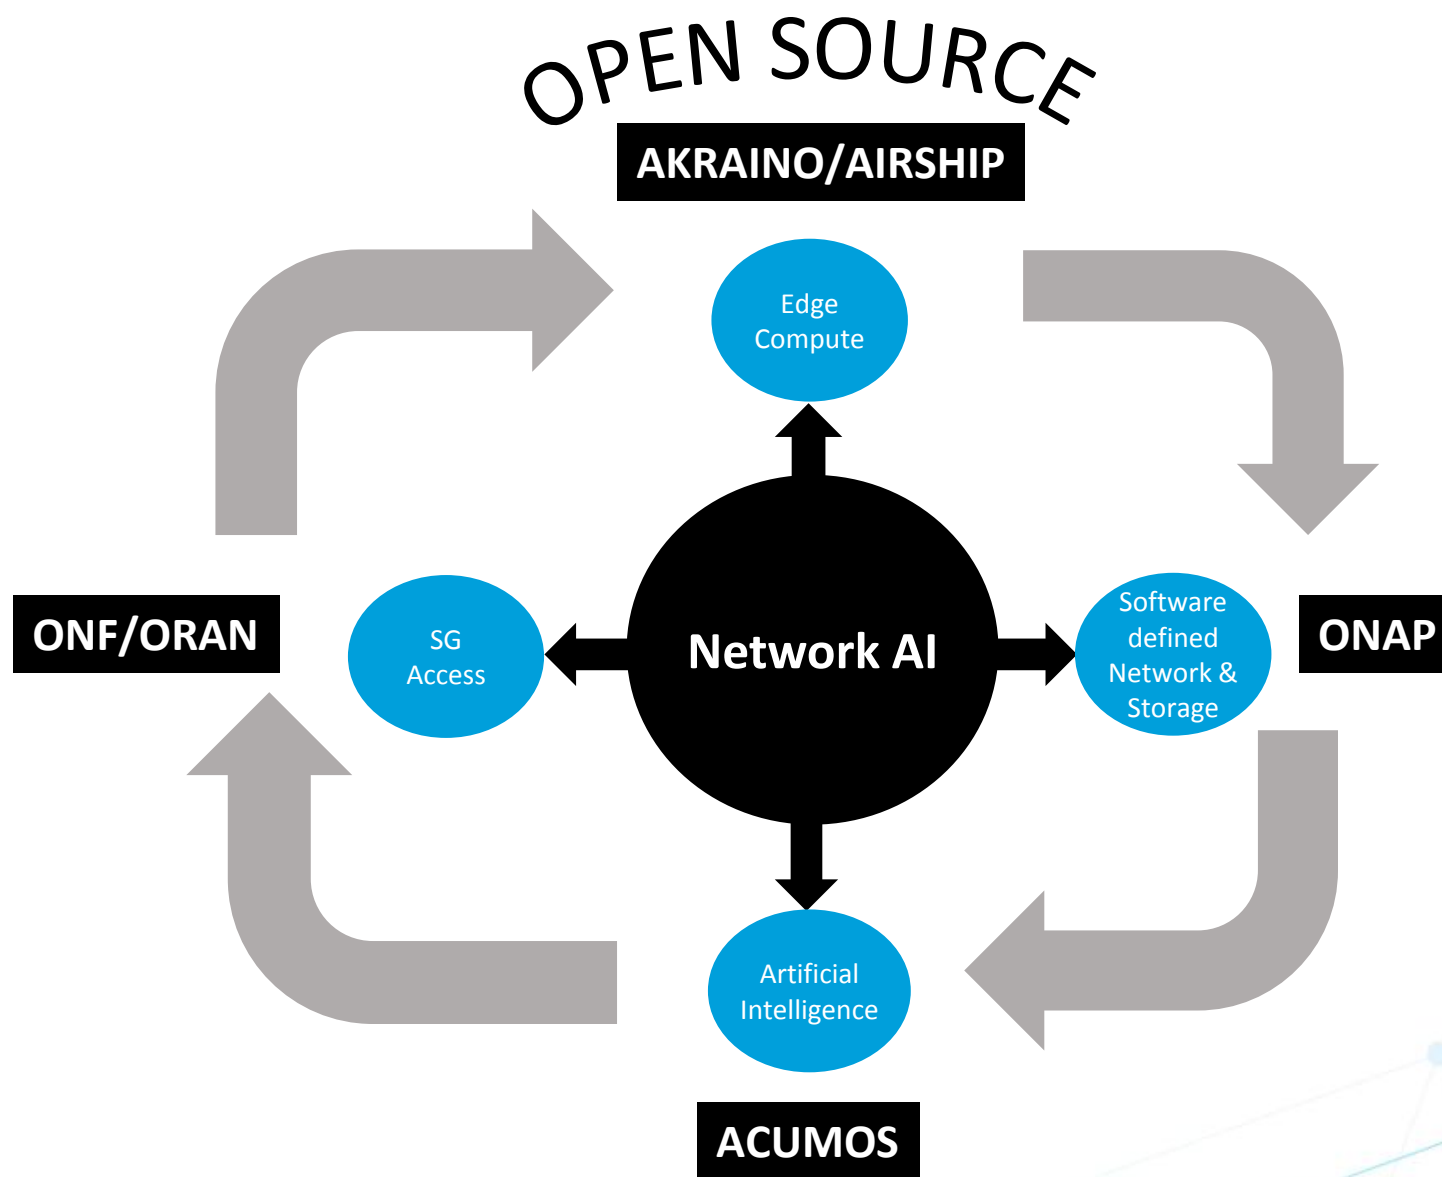
Clio Release Revised - Sprint Calendar 2019

Milestones	M1:05/20							M2:06/11							M3:07/08							M4:08/12							M5:09/30							R0:10/21							R1:11/04							Release C:11/13																																																	
	Week 21	Week 22	Week 23	Week 24	Week 25	Week 26	Week 27	Week 28	Week 29	Week 30	Week 31	Week 32	Week 33	Week 34	Week 35	Week 36	Week 37	Week 38	Week 39	Week 40	Week 41	Week 42	Week 43	Week 44	Week 45	Week 46																																																																									
Mon	20-May	27-May	3-Jun	10-Jun	17-Jun	24-Jun	1-Jul	8-Jul	15-Jul	22-Jul	29-Jul	5-Aug	12-Aug	19-Aug	26-Aug	2-Sep	9-Sep	16-Sep	23-Sep	30-Sep	7-Oct	14-Oct	21-Oct	28-Oct	4-Nov	11-Nov																																																																									
Fri	24-May	31-May	7-Jun	14-Jun	21-Jun	28-Jun	5-Jul	12-Jul	19-Jul	26-Jul	2-Aug	9-Aug	16-Aug	23-Aug	30-Aug	6-Sep	13-Sep	20-Sep	27-Sep	4-Oct	11-Oct	18-Oct	25-Oct	1-Nov	8-Nov	15-Nov																																																																									
Clio Workdep ZEP @ Bedimider, N : 05/09/19 & if required 05/10/19	Sprint Planning and Design														Dev Sprint 1							Dev Sprint 2 - Sprint 1 Demo 7/30							Dev Sprint 3 - Sprint 2 demo 8/20							Dev Sprint 4 - Sprint 3 Demo 9/10							ETE Clio-1, 2, 3 - Sprint 4 Demo 10/1							ETE Clio-4							Clio Release																																										
	Epic Identified							User Stories/Tasks Identified & Review							Map Epic and US/Task							Sub Tasks Identified & Review							Bugfix Sprint1							Bugfix Sprint2							Bugfix Sprint3							Maint R Prep								Test Planning & Preparation of Test Cases							IST Sprint-1							IST Sprint-2							IST Sprint-3							IST Sprint-4							Release Preparation						
	Preliminary Documentation														Finalize Documentation																																																																																				

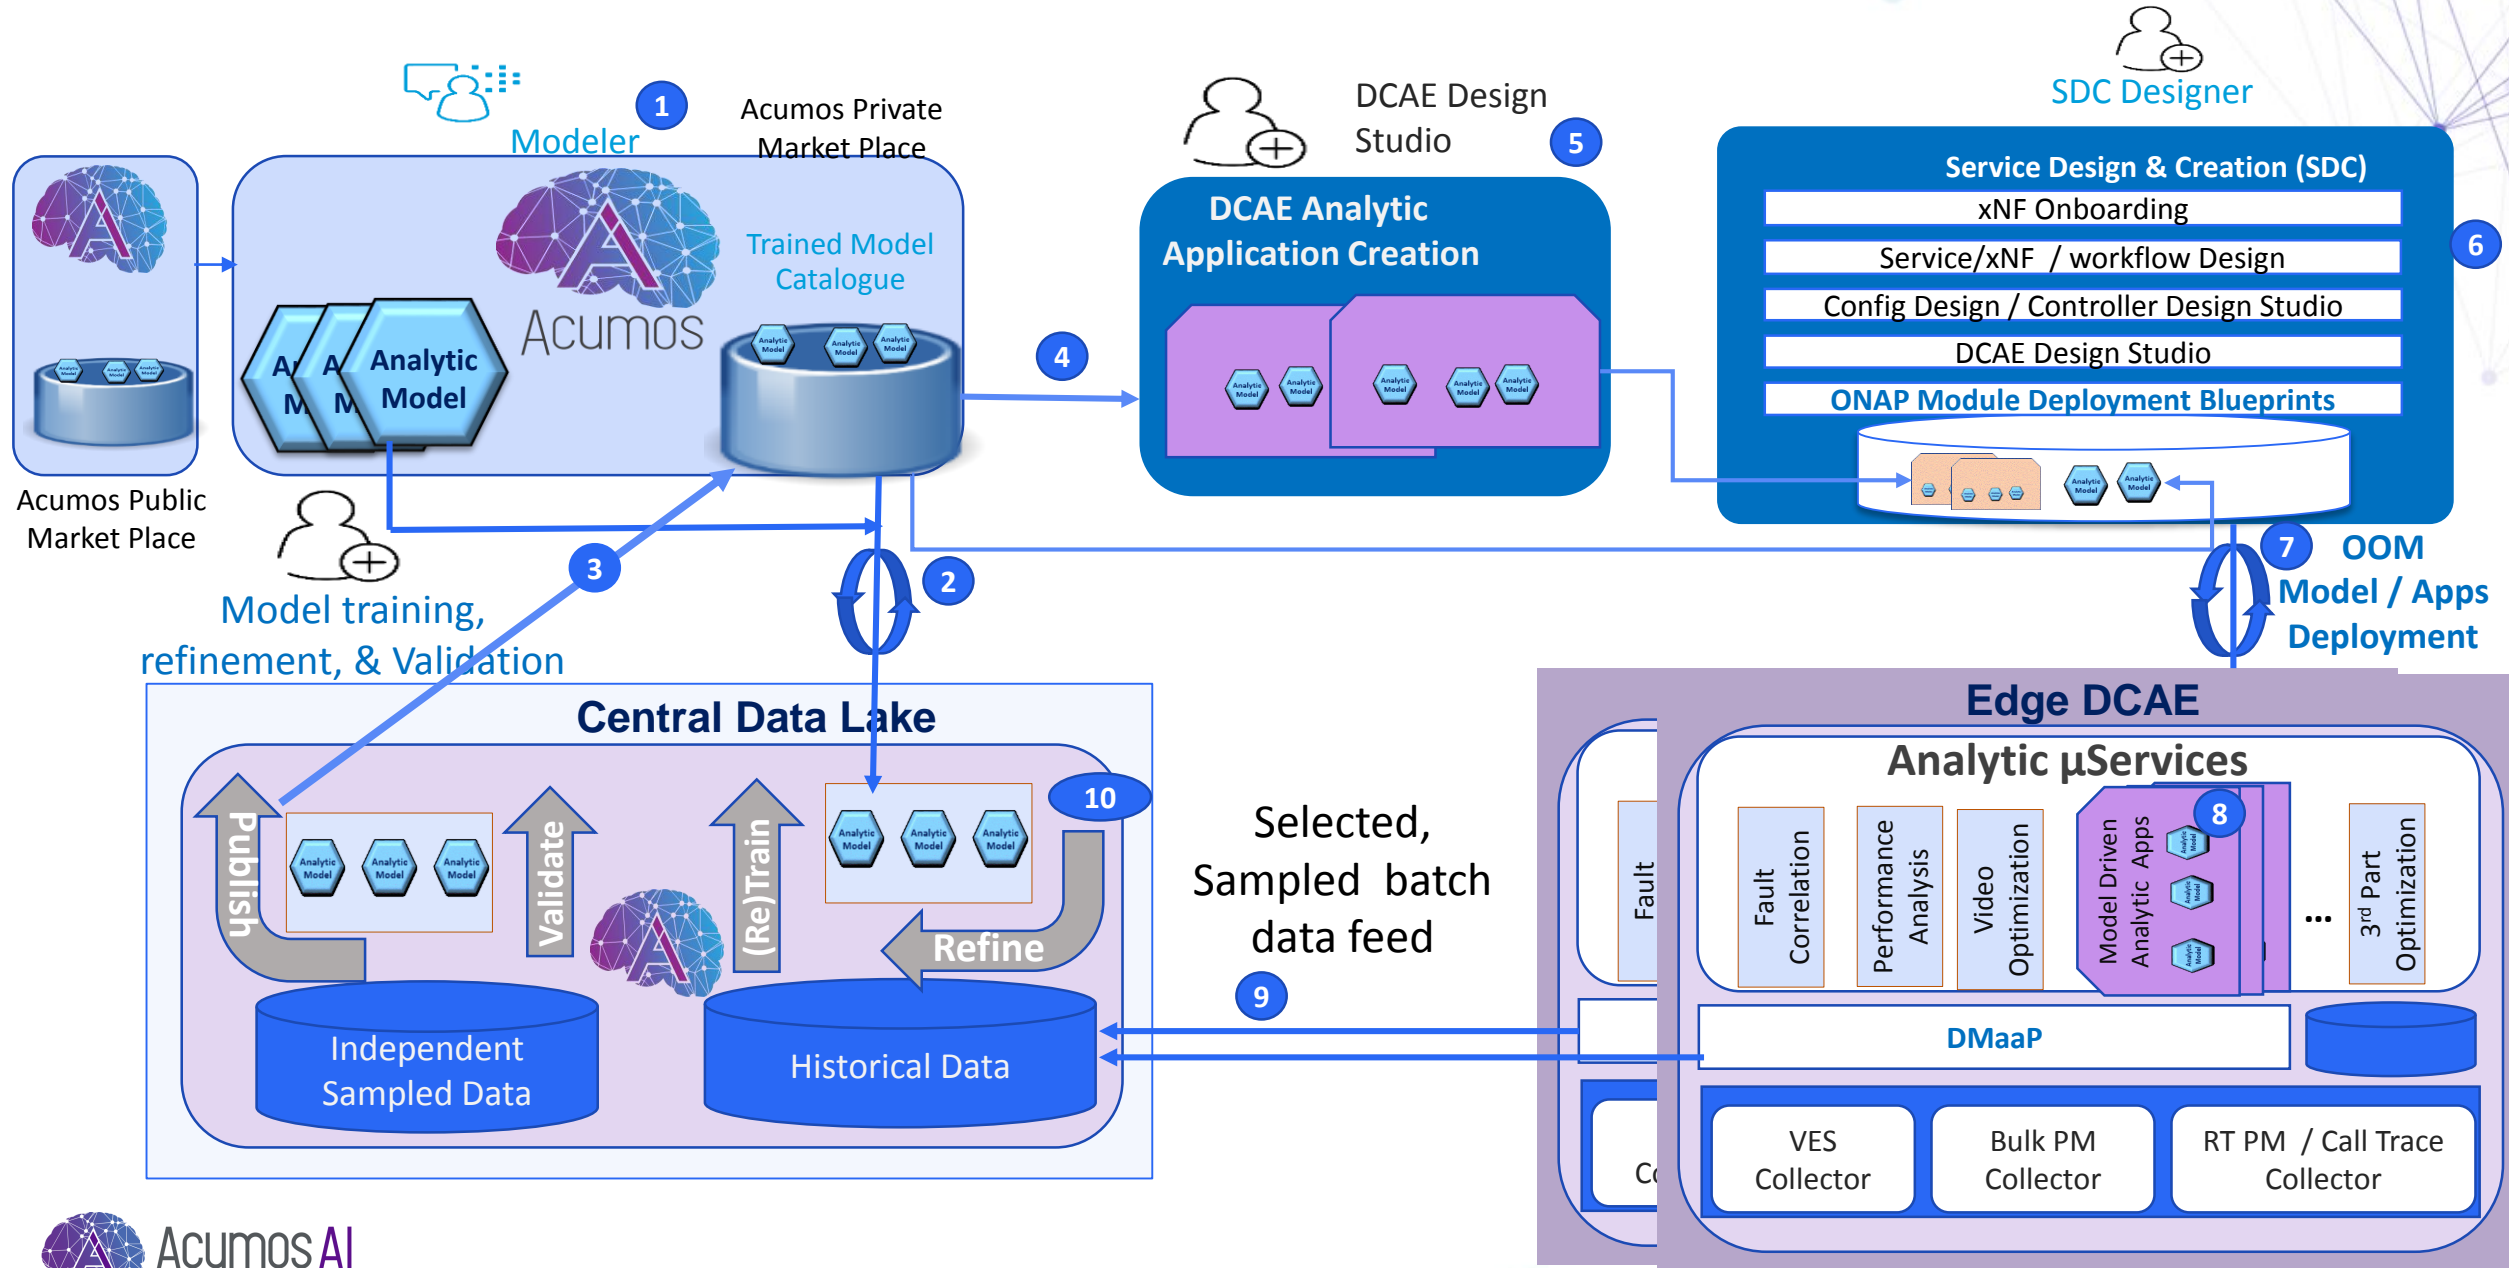
Key Features and Use Cases

AT&T Network AI

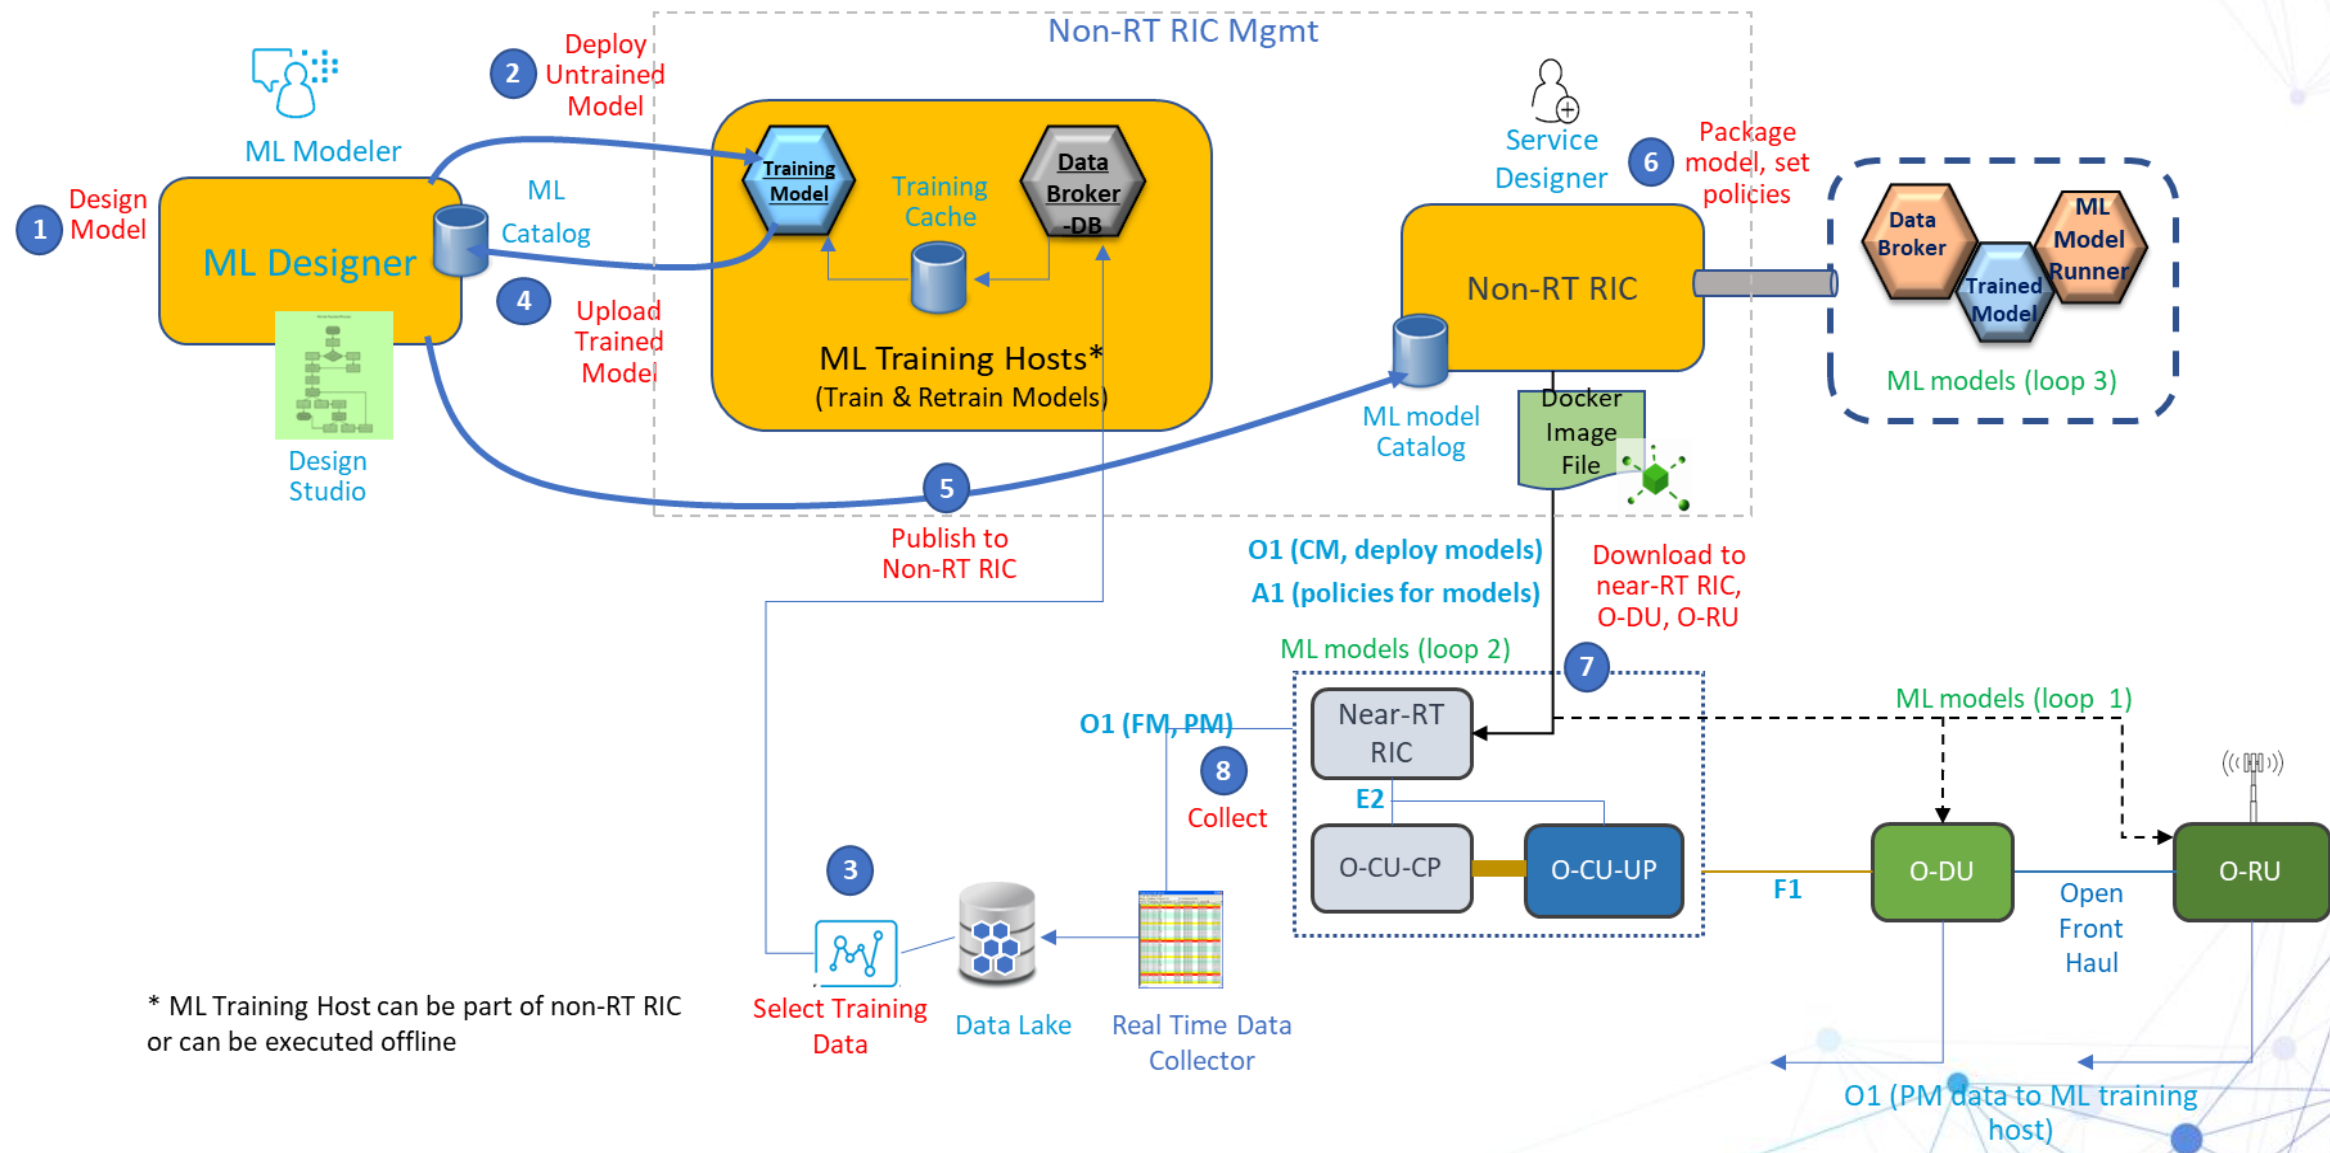
Autonomous, Contextual & Predictive Network to drive new experiences at higher velocity



ONAP / Acumos AI Integration



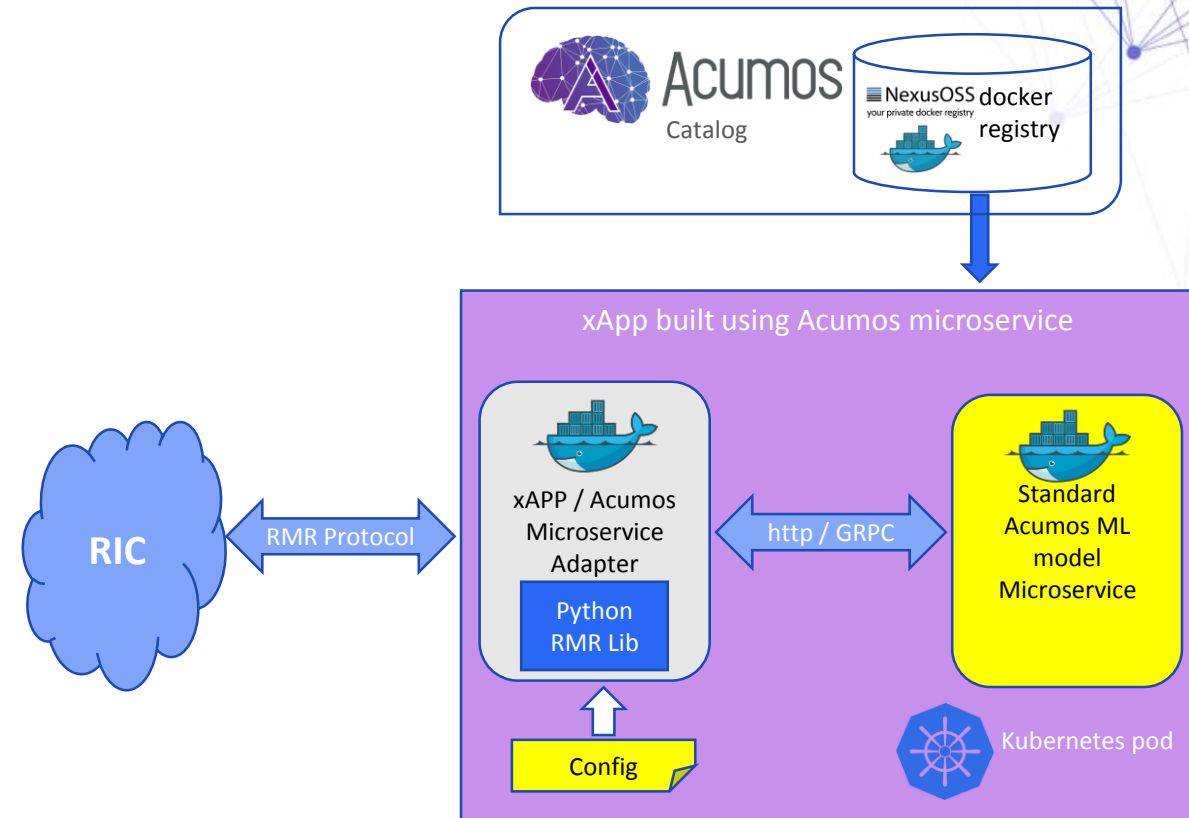
5G/4G ORAN RIC - ML model lifecycle



* ML Training Host can be part of non-RT RIC or can be executed offline

5G/4G/LTE ORAN RIC - Acumos Integration

- Acumos microservices deployment with GRPC/Protobuf and self-contained request/response usage pattern
- New Acumos microservices adapter connecting ML model to RIC via Python interface for messages on the bus using RMR.
- Controlled by JSON Config file for mapping incoming messages of interest onto calls to prediction/classification functions offered by the Acumos MS
- Two Docker images deployed together to serve as a single xApp via Helm chart
- Performance measurement results are acceptable for non-real-time and many near-real-time xApps.
- XApps use cases: 5G traffic steering, QoE optimization, MIMO beam forming optimization, etc



- HOME
- MARKETPLACE
- MY MODELS
- ON-BOARDING MODEL
- DESIGN STUDIO BETA
- SITE ADMIN
- PUBLISH REQUEST
- Q AND A
- REPORT ISSUE
- ML LEARNING PATH

Design Studio BETA
Home / Design Studio / ML Workbench 2.0

- ML WORKBENCH
 - Dashboard
 - Projects
 - Pipelines
 - Models
 - CMLP Studio
 - AcuCompose
 - Notebooks
 - FAQs

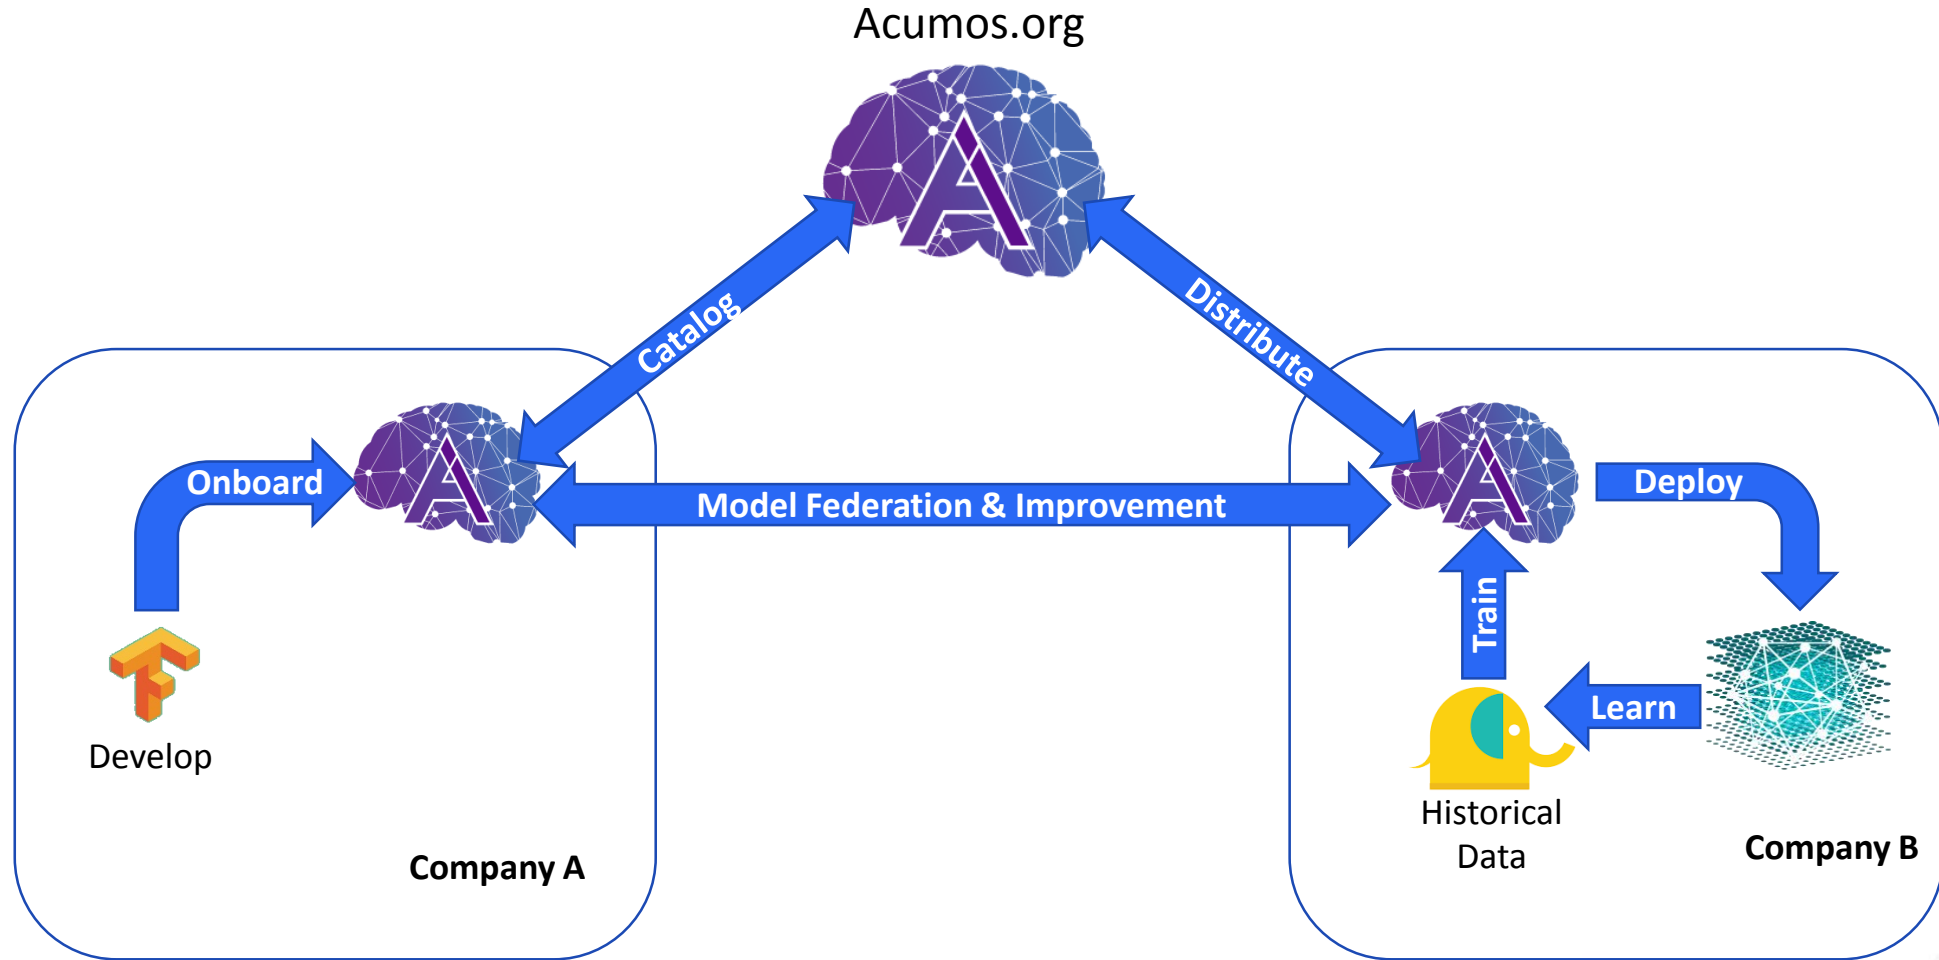
 Projects 67 Catalog Create	 Users 120 Catalog Create	 Pipelines 32 Catalog Create	 Models 231 Catalog Create	 Predictors 67 Catalog Create	 DataSets 102 Catalog Create
--------------------------------------	------------------------------------	---------------------------------------	-------------------------------------	--	---------------------------------------

 CMLP Studio Lorem ipsum is simply dummy text of the printing and typesetting industry. Lorem Ipsum has been the industry's standard dummy text ever since the 1500s, when an unknown printer took a galley of type & scrambled it to make a type specimen. Launch	 AcuCompose Lorem ipsum is simply dummy text of the printing and typesetting industry. Lorem Ipsum has been the industry's standard dummy text ever since the 1500s, when an unknown printer took a galley of type & scrambled it to make a type specimen. Launch	 Zeppelin Notebooks Lorem ipsum is simply dummy text of the printing and typesetting industry. Lorem Ipsum has been the industry's standard dummy text ever since the 1500s, when an unknown printer took a galley of type & scrambled it to make a type specimen. Launch	 Jupyter Notebooks Lorem ipsum is simply dummy text of the printing and typesetting industry. Lorem Ipsum has been the industry's standard dummy text ever since the 1500s, when an unknown printer took a galley of type & scrambled it to make a type specimen. Launch
---	--	--	---

ML Workbench Plug and Play

Acumos Collaboration

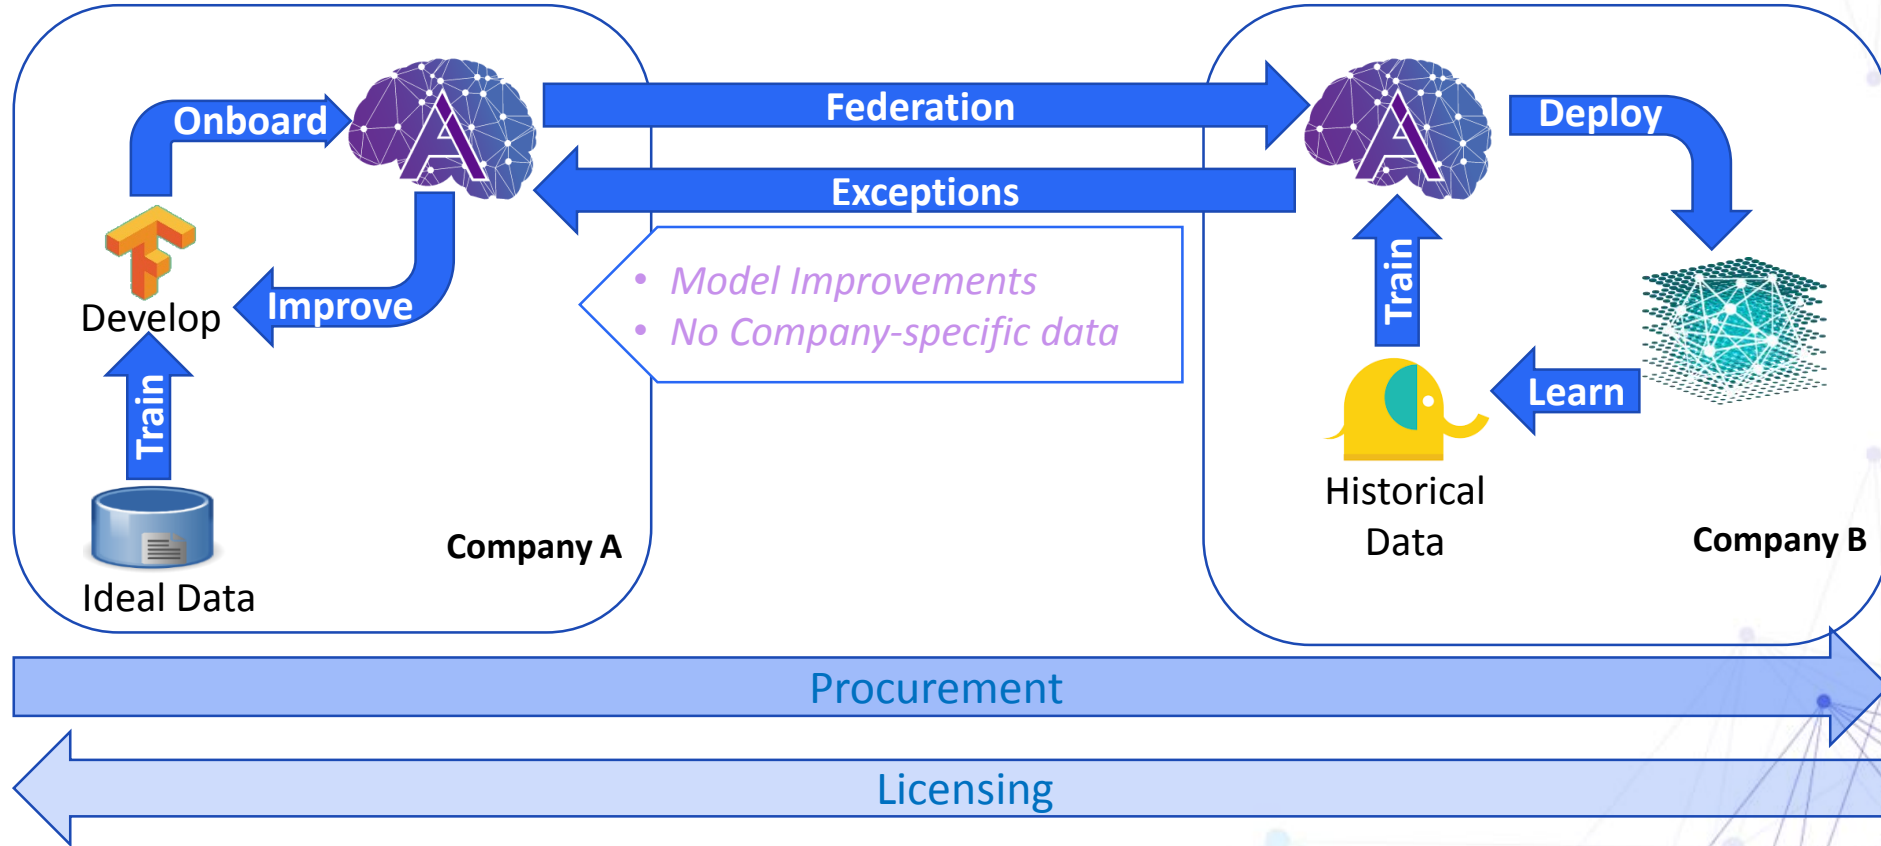
Acumos Marketplace facilitates iterative ML model development, training and improvement among collaborators



Acumos Vendor Collaboration and Federated Training

Acumos Marketplace facilitates iterative ML model development, training and improvement among collaborators

- Vendor Develops Model
- Trains On Ideal Data
- Federates Via Acumos
- Model Retrained On Actual Data
- Model Deployed
- Refinements Shared With Collaborators



Video Analytics for Real Time Streaming Data

Build using Acumos AI open source platform

Uses AT&T Research Labs Acumos AI instance which provides various video analytics machine learning models.

Streaming Video Analytics powered with following Acumos AI Models :

- Object – to classify the images in the video for ad insertion
- Sentiment – to detect the mood of the scene for ad insertion
- Face Detection – to detect human face
- Face Recognition – to recognize well known celebrities
- Privacy – to pixelate and provide privacy of the detected faces.
- Logo Recognition – to detect various company logos for advertising purposes
- Frame Ranking - compute the best frame in the video

Fast, Easy, Drag and Drop building of complex AI Applications using Acumos AI Design Studio by chaining multiple ML models

Video Analytics Models Chaining in AccuCompose


The screenshot displays the AccuCompose Design Studio interface. At the top, the header includes the Acumos AI logo, AT&T Labs Research branding, and user information for Manoop. The main workspace shows a workflow diagram for a video analytics pipeline. The process starts with an 'ingest_video_data1' node, which feeds into a 'Splitter1' node. This splitter branches into two paths: one leading to an 'image_classifier1' node and another to a 'face_privacy_detect_AI1' node. The 'image_classifier1' node's output goes to a 'Splitter2' node, which branches into 'xfer_classification_data1' and 'image_mood_classifier1'. The 'face_privacy_detect_AI1' node's output goes to a 'Splitter3' node, which branches into 'face_privacy_pixelate_AI1' and 'xfer_face_data1'. The 'xfer_classification_data1', 'image_mood_classifier1', 'face_privacy_pixelate_AI1', and 'xfer_face_data1' nodes all feed into a 'Collator1' node. The 'Collator1' node's output goes to a 'return_video_analytics1' node. The 'att_coco1' node also feeds into the 'Collator1' node. The right sidebar shows the 'Properties' panel for the selected 'return_video_analytics1' node, with fields for Node Name (empty), Node Type (return_video_analytics), Author (SANDEEP KUMAR REDDY), Model Provider (SANDEEP KUMAR REDDY), and Toolkit Type (TensorFlow). The left sidebar contains a 'Marketplace' section with 'Solutions' and 'Models' tabs, and a 'Data Transform Tools' section with a list of tools including Collator, Datamodel_MS, DirectTV_Program_FeatureExtractor, pole_fill_attributes, and Splitter.

Video Analytics System Powered by Acumos AI

Pipeline Application Demo (PAD) for Acumos

Samples 60 Latency 306.45 ms Dropped 1.92

all Configuration Options

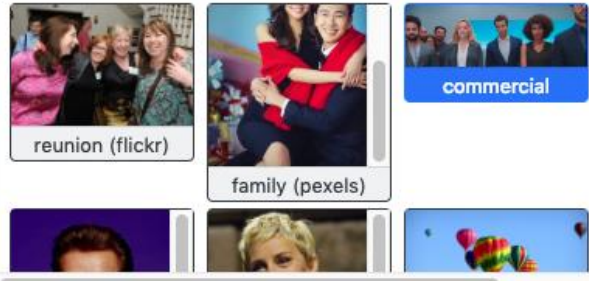



Original Review 59

Pipeline Activation


on	model	proto	input	output	latency
✓	Face Detection				217.8 ms
✓	Face Blur				124 ms
✓	Face Detection (alt)				550.4 ms
✗	Face Recognition				367 ms

Asset Selection Upload





region	x	y	w	h
3	225	67	53	53
2	120	102	45	45
1	2	132	37	37
0	442	90	43	43

Face Detection 59 1




Face Blur 59 0




region	x	y	w	h
5	137	118	24	27
4	264	94	27	36
3	55	133	22	27
2	112	115	21	26
1	194	56	41	56

Face Detection (alt) 58 6



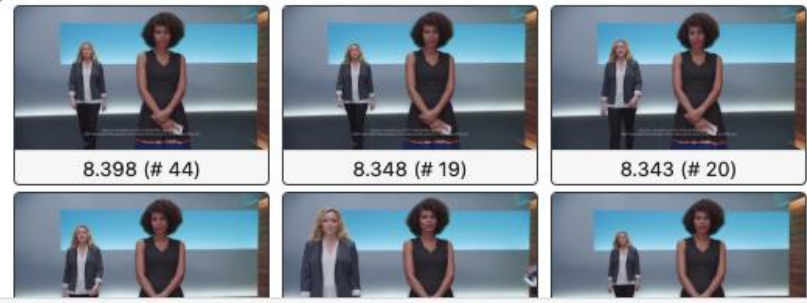
tag	score
suit, suit of clothes	0.816
jean, blue jean, denim	0.037
Windsor tie	0.033
bow tie, bow-tie, bowtie	0.019
gown	0.009

Image Classification 59 3



tag	score
sad	0.223
excitement	0.16
amusement	0.14
fear	0.127
awe	0.12

Image Mood 59 0



Frame Ranking 59 2

LF AI Landscape and Acumos AI

Acumos AI Integration within LFAI Landscape


Greyed logos are not open source

Machine Learning	Framework	Platform	Library	Framework	Platform	Library	Tool	Notebook Environment	Reinforcement Learning
	Accord.NET, FedAI, Apache SystemML, AI, H2O, Cartoost, Spark MLlib, mlpack, Orange3, Microsoft LightGBM, MAHOUT, Kubeflow, mlflow, H2O Driverless AI, QOSMOS, Learn, 將軍 Shogun, Sonnet, ML.NET, RAY, SELDOON, VOPAL WABBIT, Salesforce TransmogriAI, dink XGBoost, xLearn	Deep Learning	TensorFlow, PyTorch	SINGA, Chainer, CNTK, dy/net, DeepDetect, FIDL, DL4J, Borech, Intel Distiller, Dnnnet, Intel Neon, facebook Pythia, PaddlePaddle, Polyaxon, Keras, OpenMind PySytt, plaidML, PyTorch Ecosystem, PYTORCH, TensorFlow	Notebook Environment	Apache Zeppelin, BeakerX, colab, IPython, jupyter, IBM	Reinforcement Learning	OpenAI Gym, Coach, DeepMind Lab, Dopamine, Horizon, OpenAI, Google PlanNet, Google TRFL	


Distributed Computing	Computing & Management	Interface	Versioning	Store & Format	Pipeline Management	Stream Processing	SQL Engine	Feature Engineering	Visualization	Governance	Labeling
	EDL LFAI Incubating, Apache Storm, Kubernetes, Intel Nauta, Singularity	Apache Ambari, Bahir, Apache TORRE, LOVY, MESOS, Apache Ranger, U	datacensus, DVC, Orc, AVEVOB, ARROW, intel Analytics Zoo, bear, Flink, kafka, ALPACAS DRILL, presto, Google Cloud Feast, Featuretools, Uber deck.gl, Google Facets, NYU, COMMUNITY DATA SCIENCE HUB, EGERIA, intel CVAT	ouilt, Pachyderm, AIRC, DELTA LAKE, druid, marmaray, TROJAN, samza, uber uReplicator, hudi, Parquet, pilosa	Feature Engineering	Visualization	Governance	Labeling			

AI Ethics & Fairness	Natural Language Processing	Benchmarking	Training	Parameter	Format	Marketplace	Workflow	Tool	Programming
	Audit AI, IBM Aequitas, IBM Fairlearn, Microsoft Fairlearn, University of Washington Lime, Google Lucid	fastText, facebook LASER, intel NLP Architect, IBM Watson NLP, Google Bert, Facebook, Cisco, Intel NLP Architect, ParIAI, facebook Pytest, IBM Watson NLP, spaCy	DAWNbench, MLPerf	Horovod LFAI Incubating, Amazon DSSTNE, Microsoft Fairlearn, Waterloo Fairlearn, LEDWT, Petastorm	Angel LFAI Incubating, ONNX	Acumos LFAI Graduated, Kaggle, IBM	Airflow, nifi, Amazon Neo-AI, PipelineAI, turi, Amazon SageMaker, Cadence, Pinterest, Cyclone	amazon Neo-AI, PipelineAI, turi, Amazon SageMaker, Cadence, Pinterest, Cyclone	Pyro LFAI Incubating, PyMC3, SciPy, Stan, PyTorch, Julia, NumPy, NYOKA, pomegranate

LFAI Member Company	Premium	General
	amdocs, AT&T, Ericsson, Huawei, Tech Mahindra, Tencent 腾讯	Baidu 百度, Nokia, ZTE, B-YOND, ciena, DiDi, geminiopencloud, orange, redhat, Syllabs.io



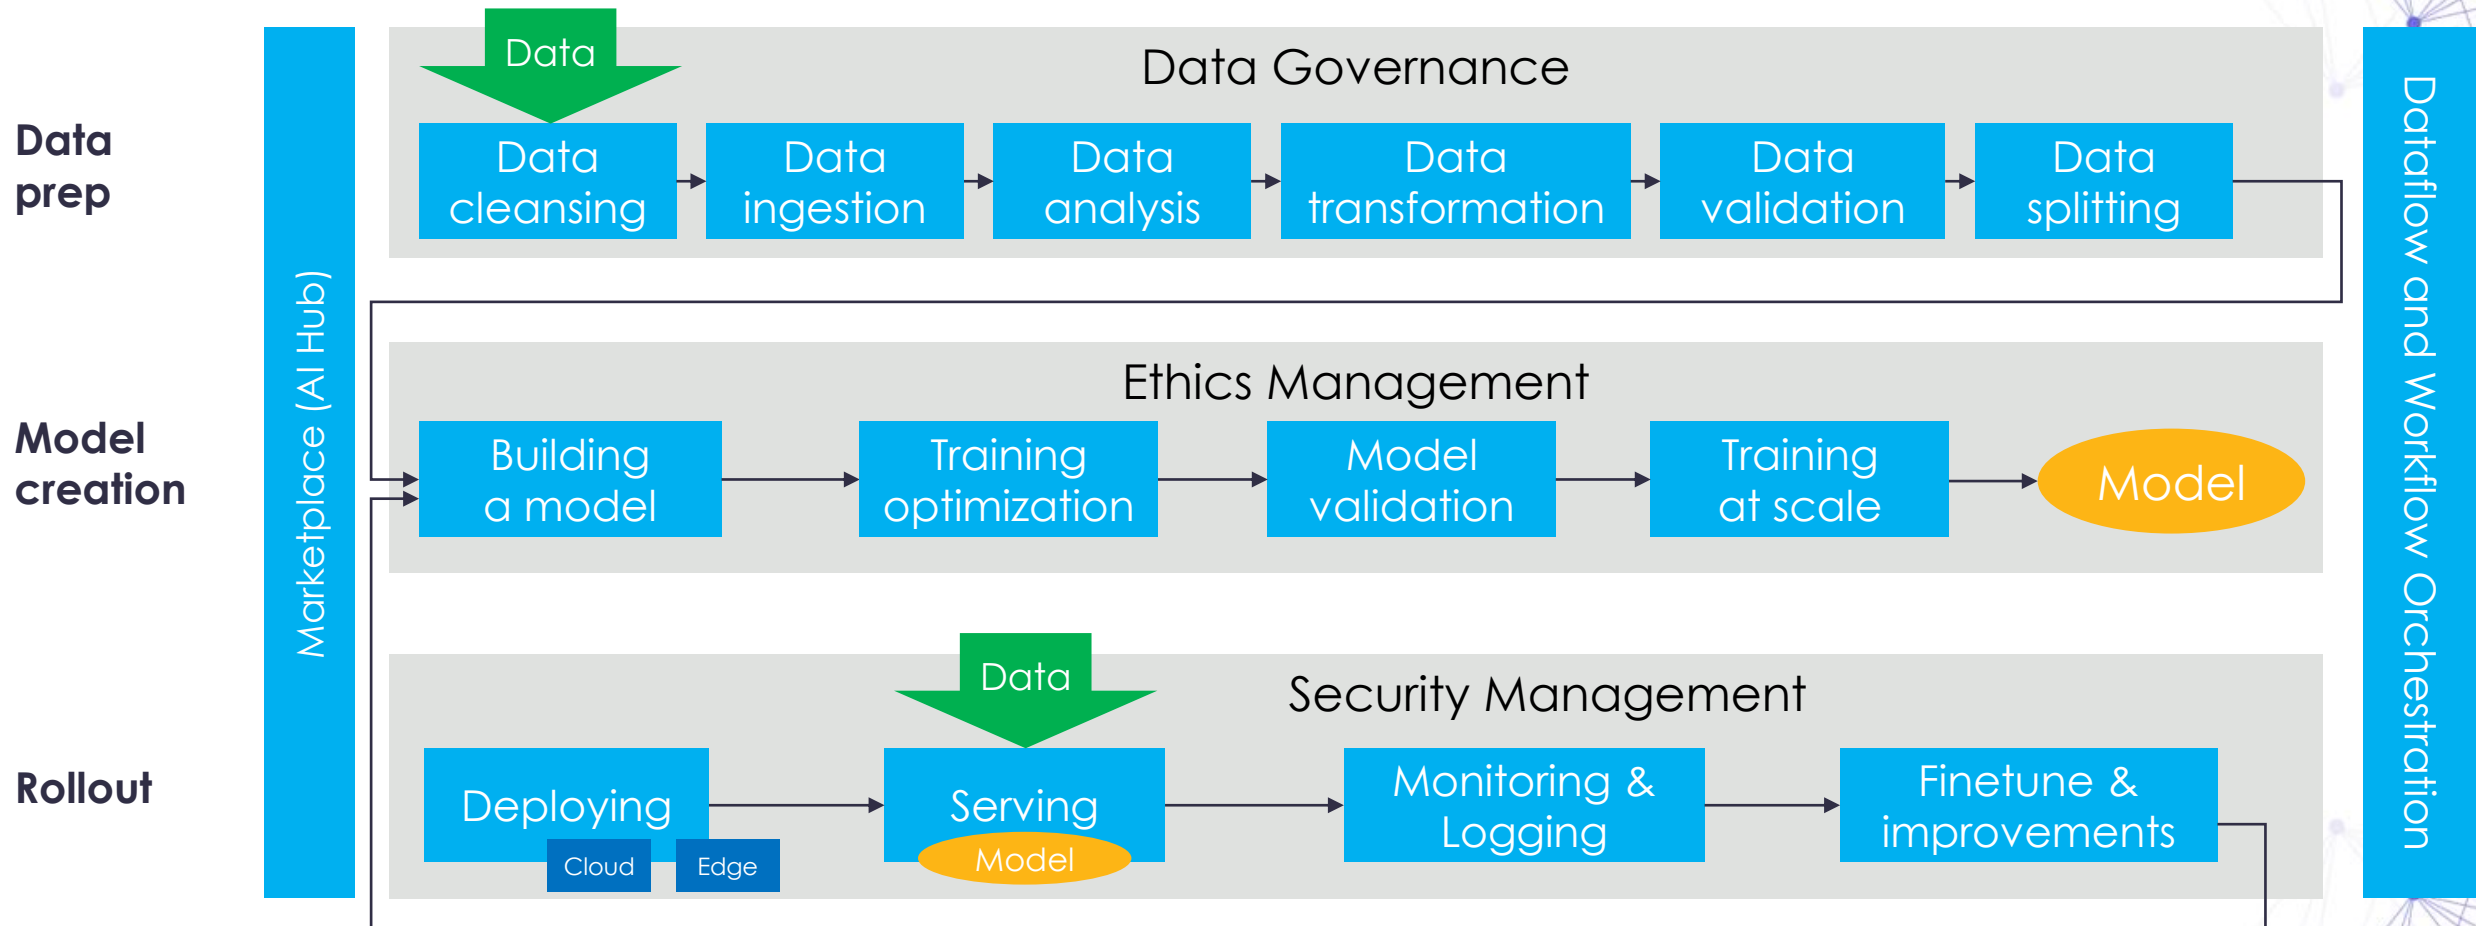
This landscape is intended as a map to explore open source artificial intelligence (AI), machine learning (ML), and deep learning (DL) projects, and also shows the member companies of the LFAI Foundation.



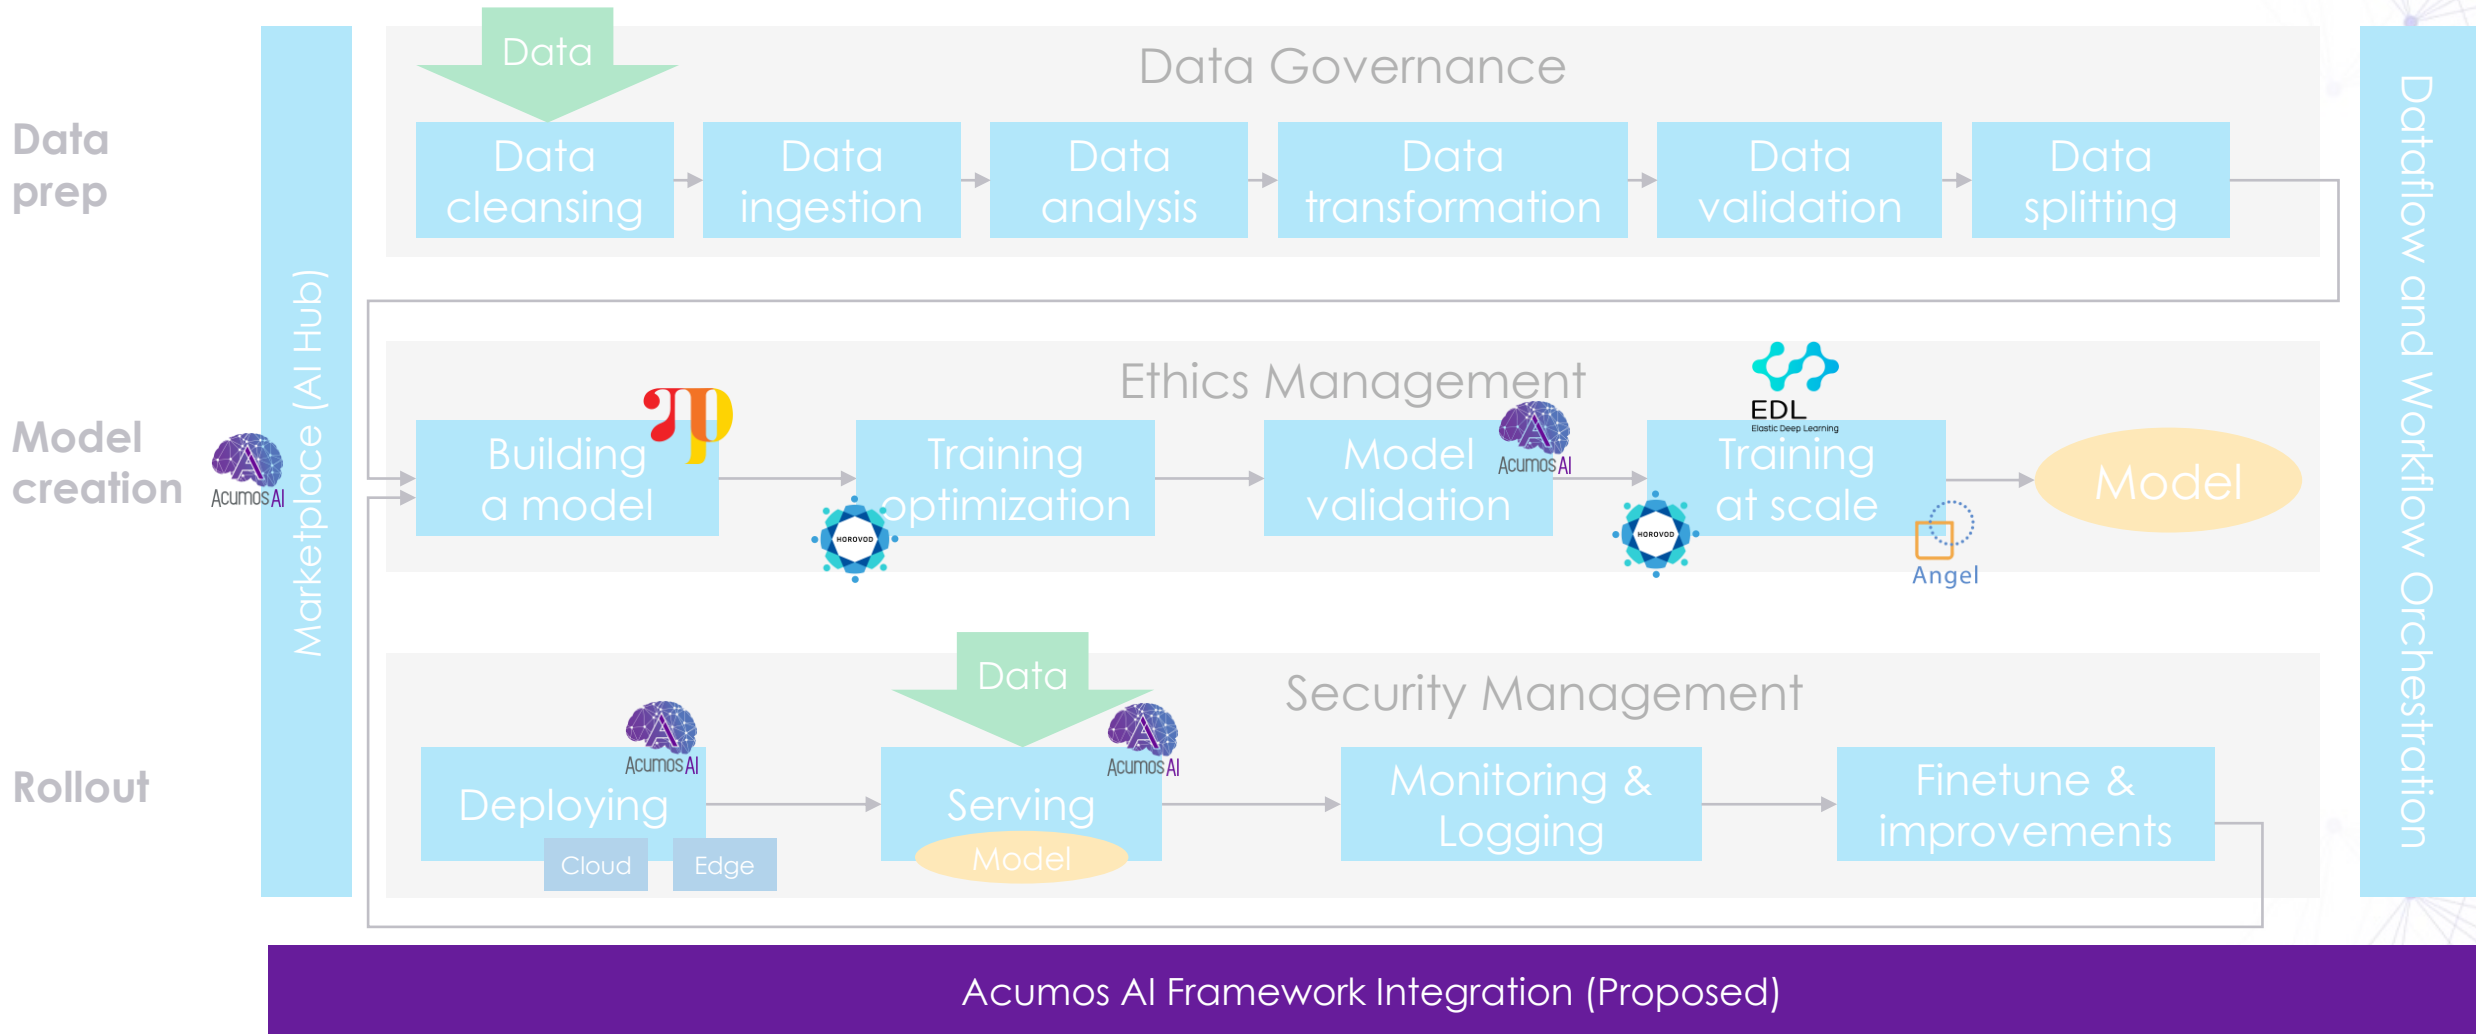
LFAI Landscape

l.lfai.foundation

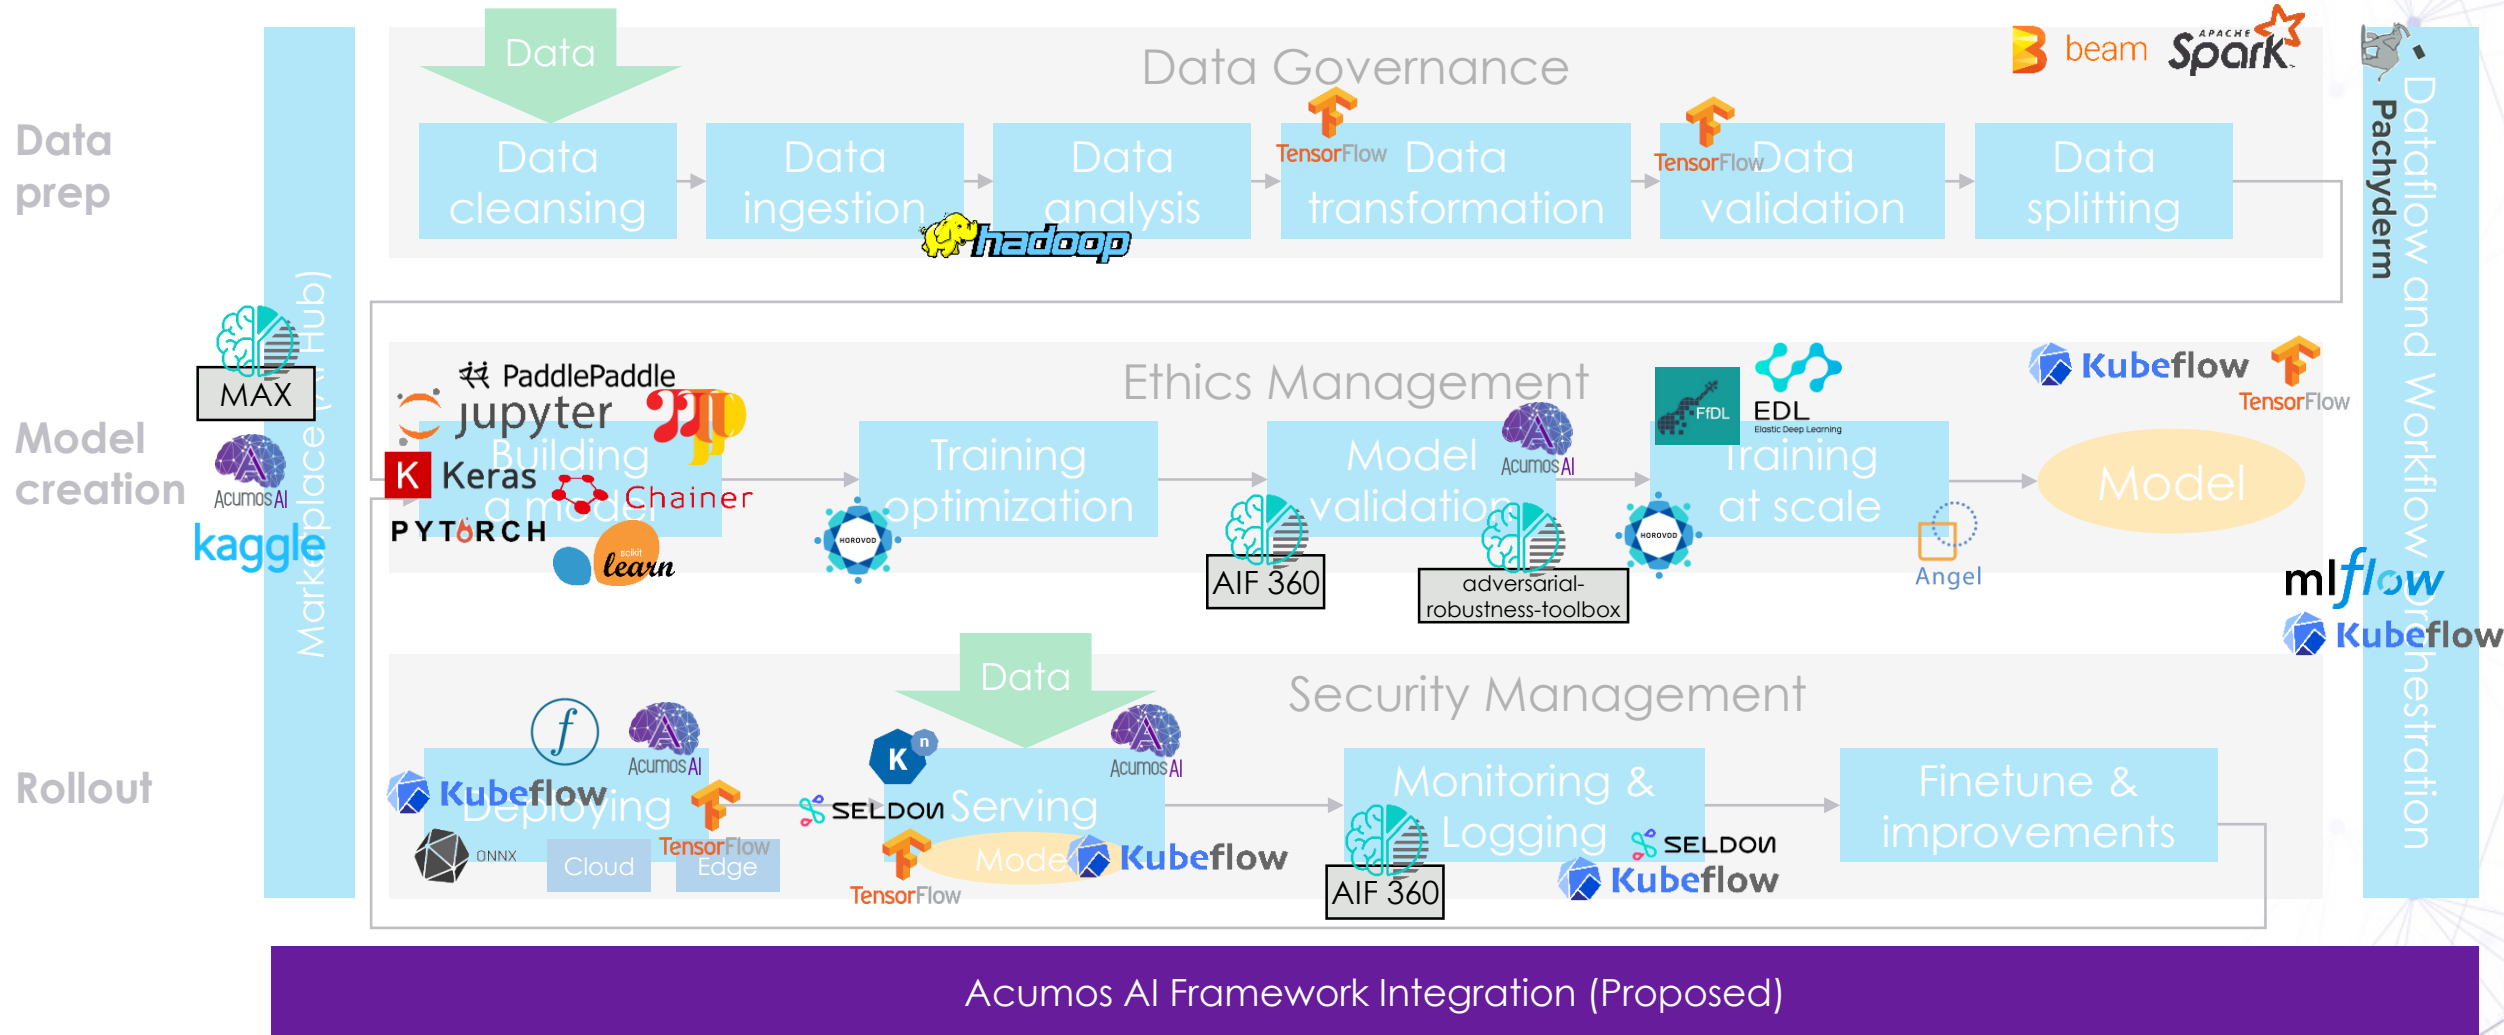
Machine Learning Stack



LFAI Machine Learning Stack



Open Source Machine Learning Stack



Q&A?

Thank You!