Reference architectures for cloud-based platforms:

Convergence vs. Diversification

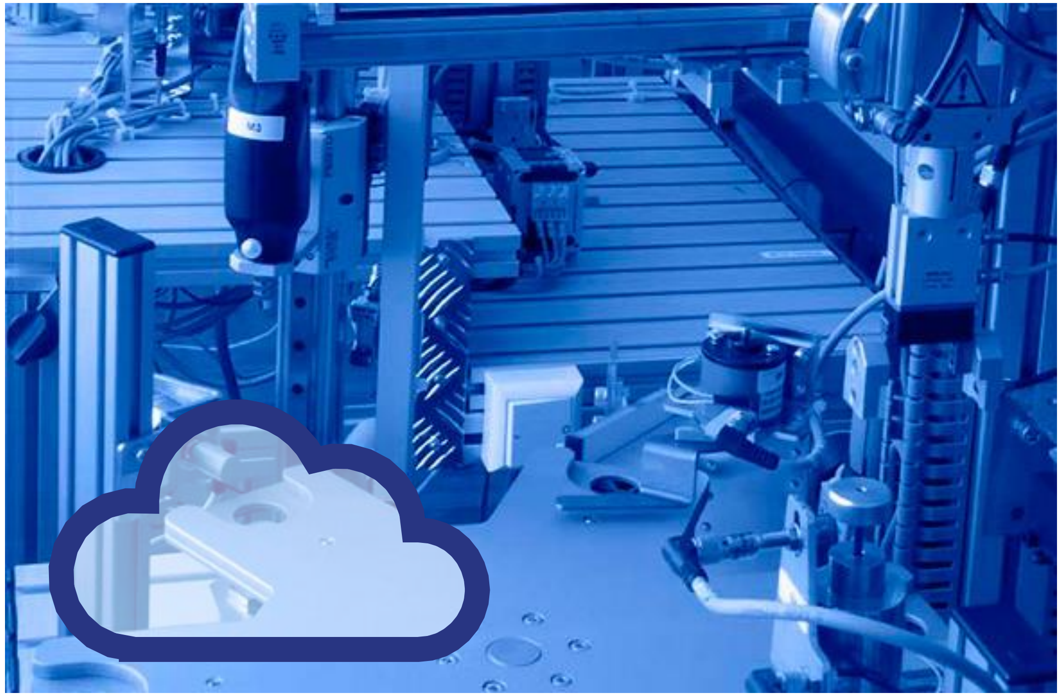

Robert Lovas

Institute for Computer Science and Control

Email: robert.lovas@sztaki.hu

Agenda

1. Steppingstones
   to reference architectures

2. Smart, orchestrated
   reference architectures

3. Applications/development of reference architectures

…focusing on Cyber-Physical Systems

Industry 4.0 in Europe:
Initiatives on Digitizing Industry

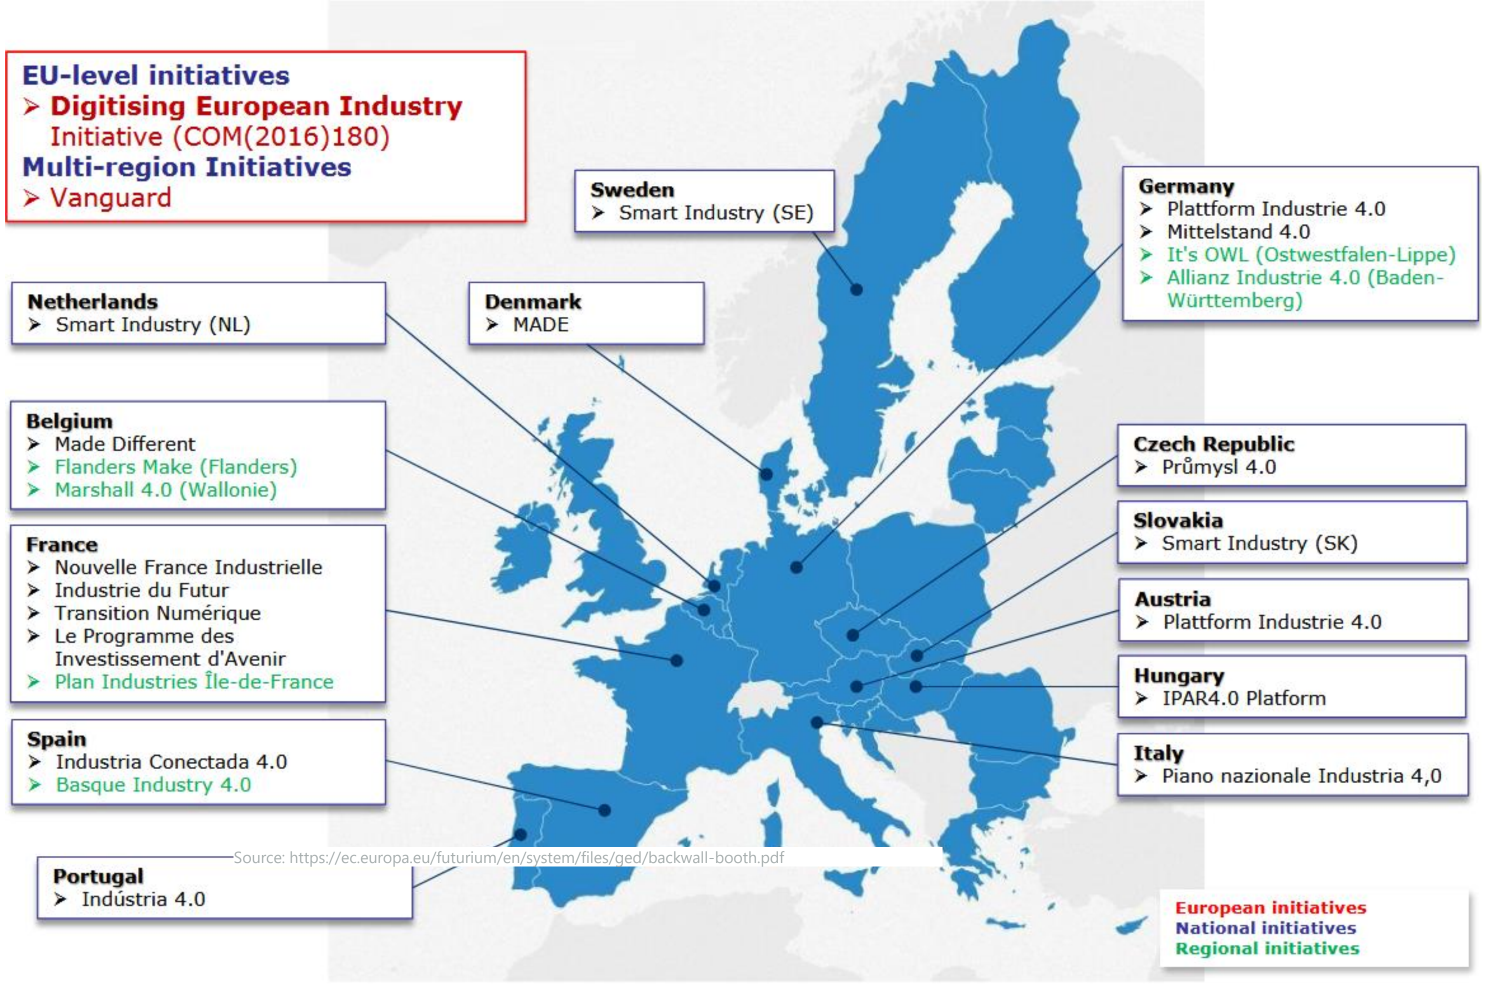

Source: https://ec.europa.eu/futurium/en/system/files/ged/backwall-booth.pdf

IPAR 4.0
NATIONAL TECHNOLOGY PLATFORM

SZTAKI

→ https://www.i40platform.hu/en

Actual trends

- Several solutions are already available from public cloud providers for **Internet of Things** (IoT) and **Big Data** application areas.

HDInsight Microsoft   Amazon Elastic MapReduce   Cloud Dataproc from Google

- **Private** clouds have significant benefits in terms of security and integrability into the enterprise environment but **hybrid** and **multi-clouds** are also widespread.

➔ Growing demand for cloud **ochestrators** and **brokering** tools.

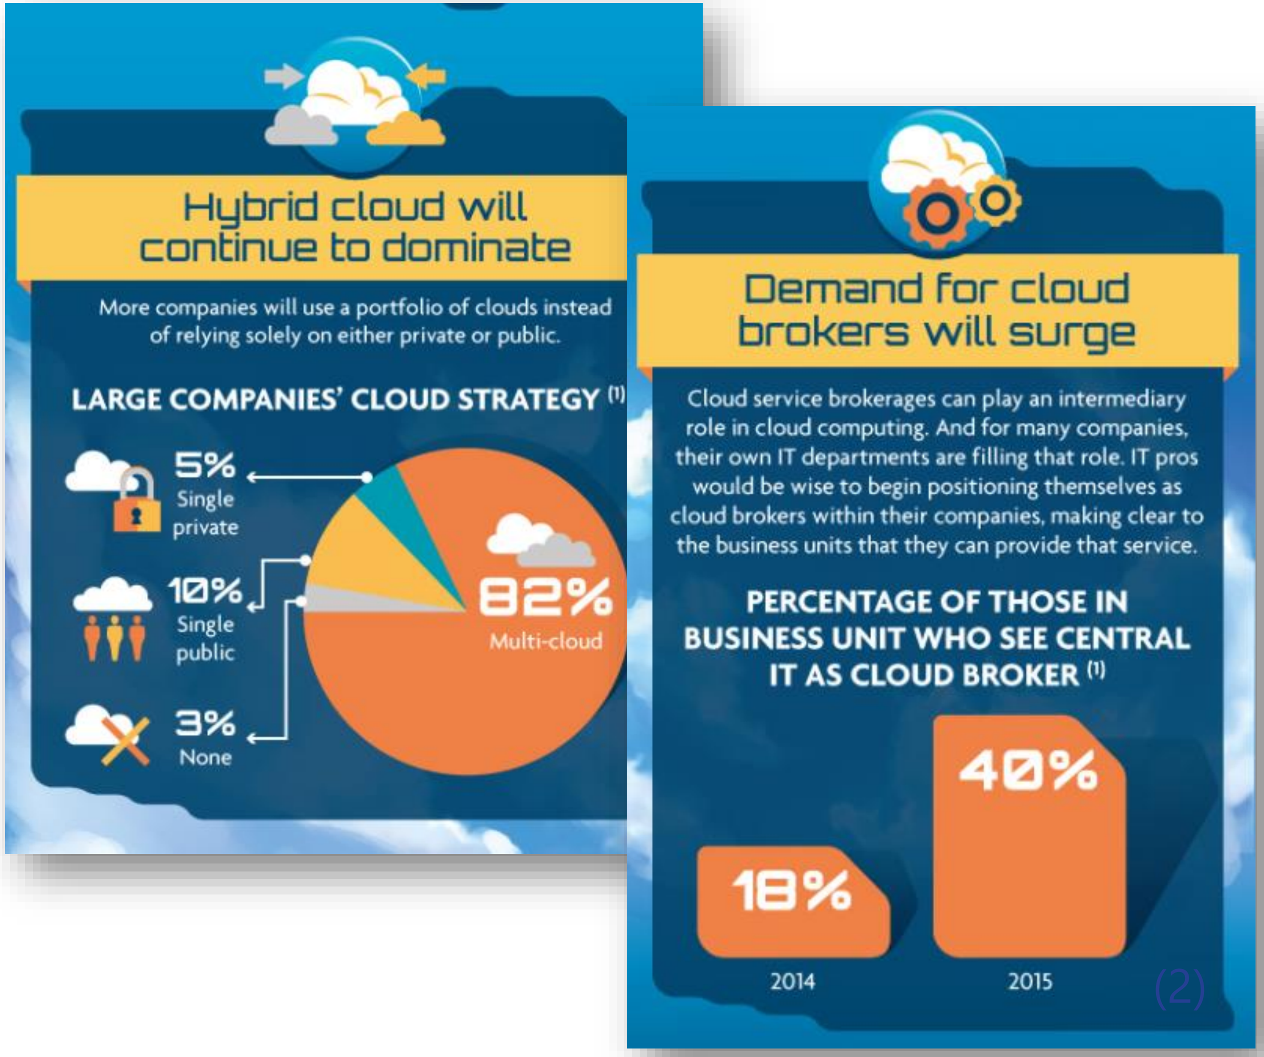

Source: www.computerzone.org

Clouds in production systems: some non-trivial problems

- Application areas:
  - Store and process **sensor data**
  - Historical **analysis**, simulations, predictions, etc.
  - **Visulisation**

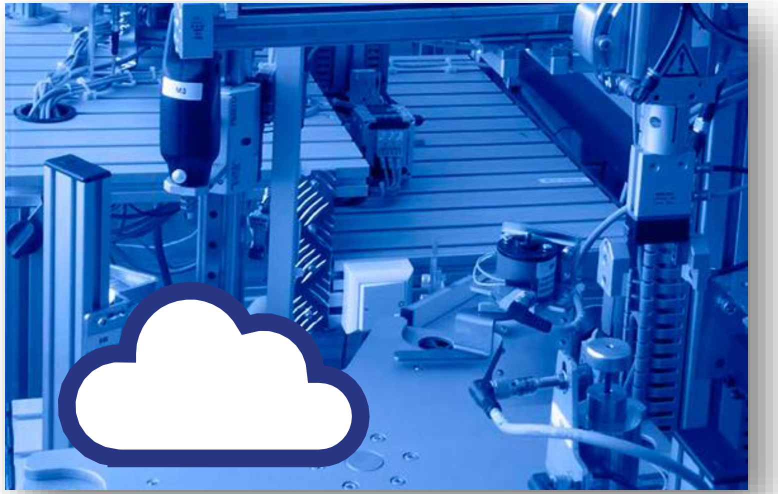

- Industrial users face challenges when they intend to benefit from cloud computing:
  - **migration** of *legacy and new applications* into clouds
  - their **orchestrated** deployment/maintenance,
  - their on-demand **scaling**,
  - **portability**, when a cost-efficient hybrid cloud or cloud agnostic (vendor independent) solution is needed, etc.

➔ emerging new software container and orchestrator solutions

Emerging new software container technologies

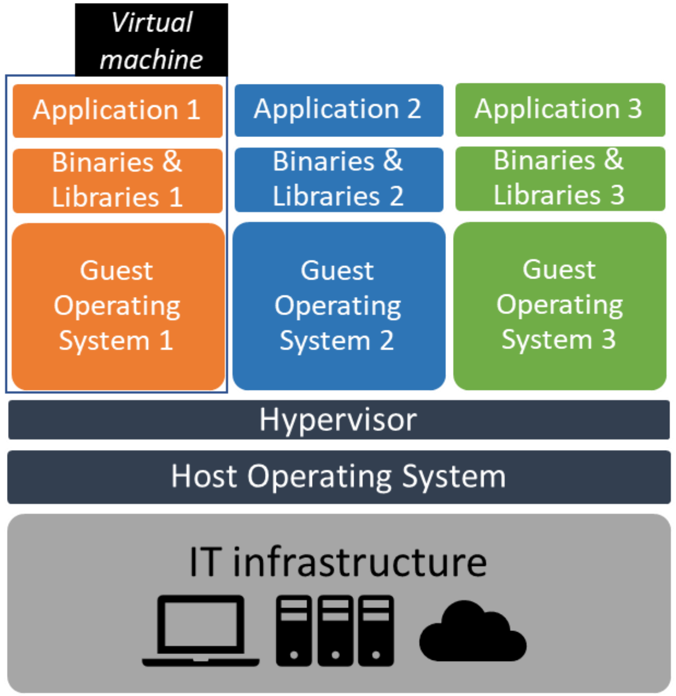

Cloud
Topic
(VM/IaaS)

Docker
Software

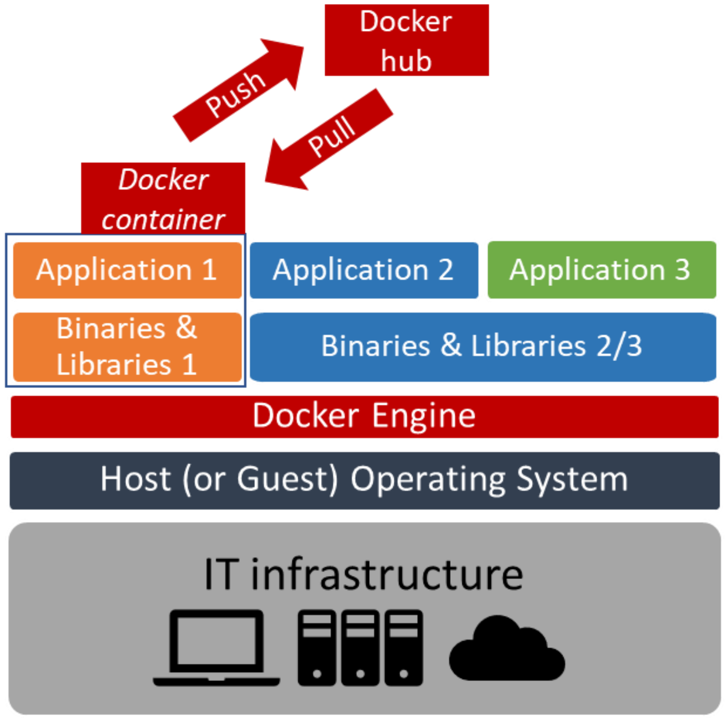

EFFECT OF COVID-19 PANDEMIC

Source: Google trends

Emerging orchestration tools

- "**Orchestration** is the automated configuration, coordination, and management of computer systems and software." (including the deployment and maintenance, i.e. the entire life-cycle)

- A number of tools exist for automation of ***server configuration*** and management
  - Ansible, Puppet, Salt, Terraform, AWS CloudFormation, etc.

- For ***container orchestration*** there are different solutions
  - Kubernetes or
  - managed services such as AWS EKS

Occopus hybrid cloud / container orchestrator

- Multi-cloud solution
- Contextualization
- No vendor lock-in
- Portable descriptor file
- Big Data support
- Enable auto-scaling

MONTHLY FEE: 4 NODE HADOOP CLUSTER

Microsoft Azure

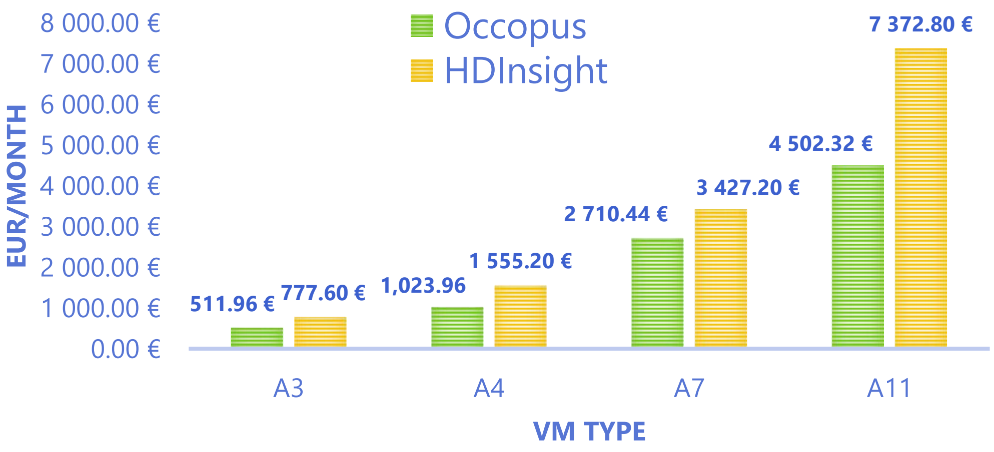

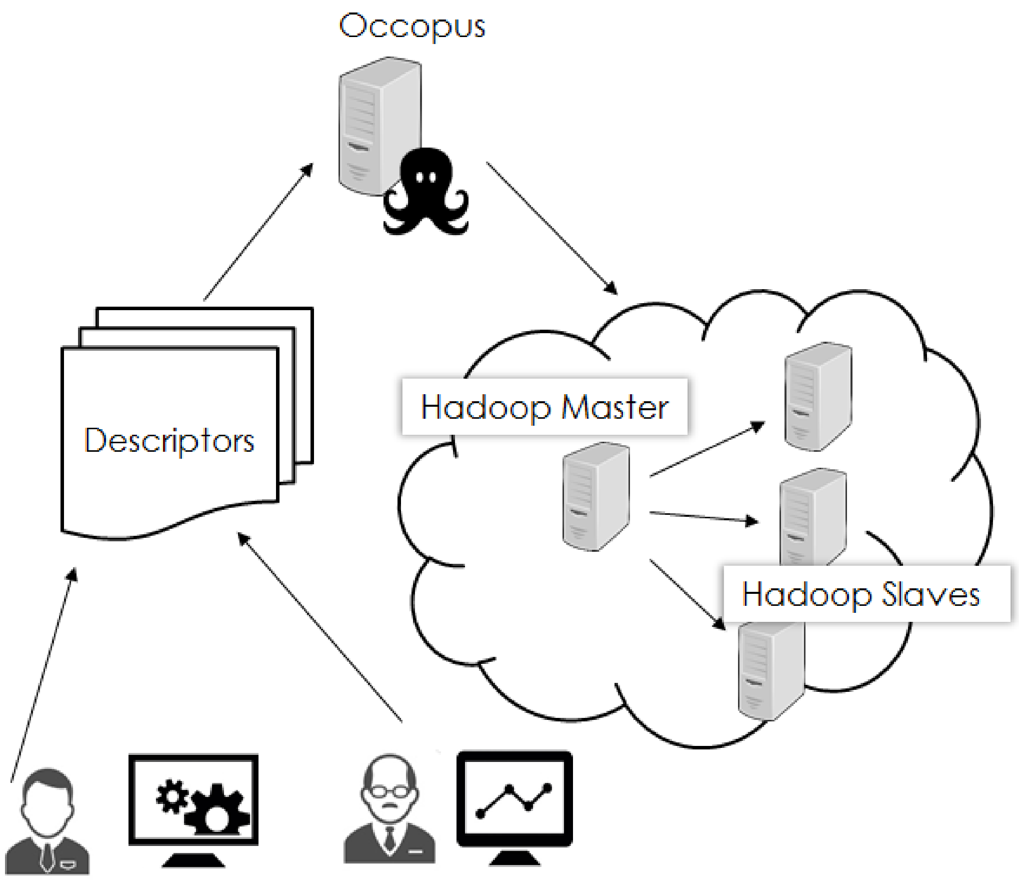

amazon web services

openstack

Microsoft Azure

OpenNebula

CloudSigma

➔ occopus.lpds.sztaki.hu

Flowbster

- quick deployment of the workflow as a pipeline infrastructure in the cloud
- once the pipeline infrastructure is created in the cloud it is activated and data elements of the data set to be processed flow through the pipeline
- as the data set flows through the pipeline its data elements are processed as defined by the Flowbster workflow

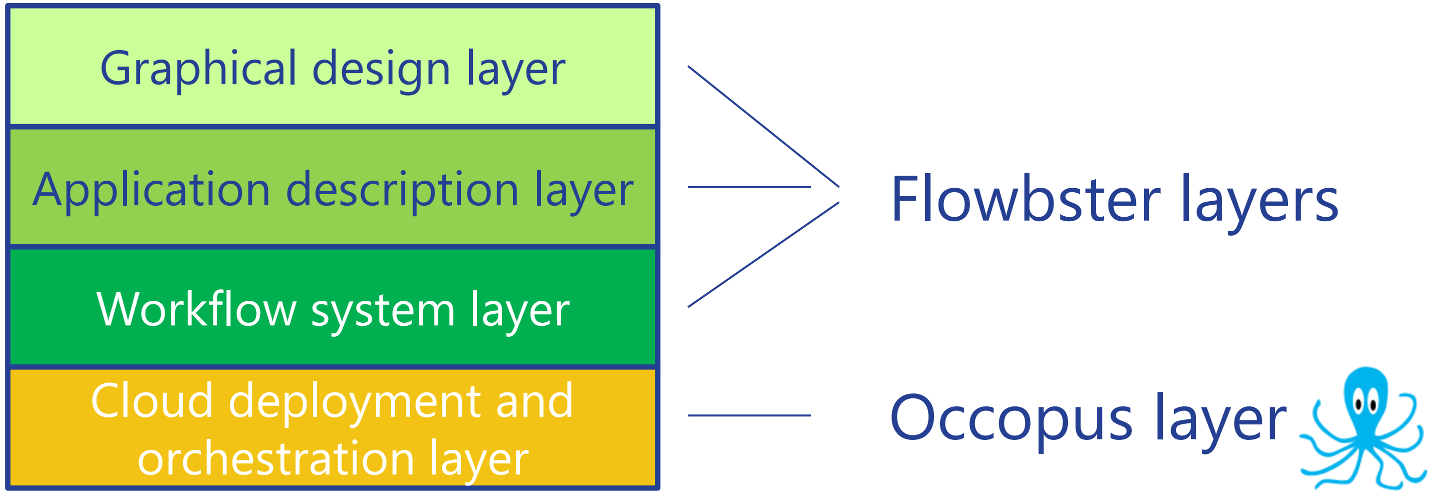

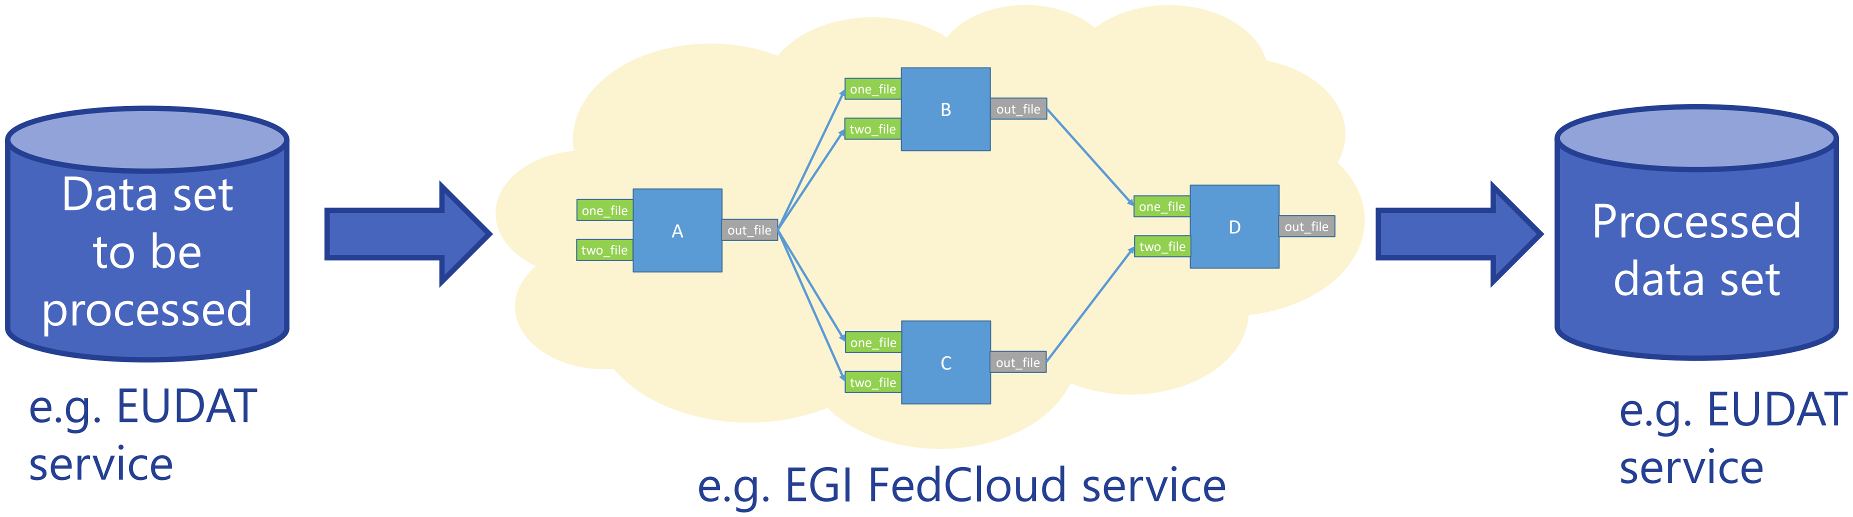

MTA CLOUD

- Joint project with Wigner Datacenter
- Dedicated for academic research communities
- Operation since Q4/16

- Wigner Datacenter:
  - 1040 vCPU
  - 3.25 TB RAM
  - 432 TB HDD
  - 4 V100 GPU
- SZTAKI:
  - 2952 vCPU
  - 2 TB RAM
  - 330 TB HDD
  - 8 K80 GPU
- OpenStack based community cloud
- 100+ research project supported by MTA Cloud

- Upgrade in Q4/2020

➔ www.cloud.mta.hu

Felhasználást segítő szolgáltatások

- DataAvenue
- Cloud alkalmazásokat támogató portál indítása
- Occopus cloud orchestrator indítása
- Apache Hadoop klaszter kiépítése
- Apache Spark klaszter kiépítése
- Apache Spark klaszter RStudio stack-el
- Apache Spark klaszter Python stack-el
- Docker-Swarm klaszter kiépítése
- Flowbster - Autodock Vina
- TensorFlow, Keras, Jupyter Notebook stack
- TensorFlow, Keras, Jupyter Notebook GPU stack

Enzimológia rendezetlen fehérje MD elemzések, Tompa csoport	Enzimológiai Intézet	Horváth Tamás
Fokent MD elemzesek, ill. statisztikai szamolasok rendezetlen feherjekkel kapcsolatosan.		
Eukarióta eredet	Balatoni Limnológiai Intézet	Dr. Zachar István
Az eukarióták kialakulása a prokariótákból az egyik legnehezebbnek tartott nagy evolúciós átmenet. Számos alapvető evolúció újítás jelent meg ekkor, ezek közül is kiemelkedik a energiatermelő prokarióták endoszimbionta fenntartása, azaz a mitokondriumok megjelenése. A szimbionta integrálódása evolúciós és ökológiai kérdések sorozatát veti fel.		
Fast Reacot Monte Carlo simulation	MTA Atommagkutató Intézet	Batki Bálint
Az ALLEGRO és SFR gyorsearktorok zónáját szeretném vizsgálni pálca szintű részletességgel, amihez nagy erőforrást igénylő Monte Carlo számításokra van szükség. Cél csoportállandók generálása a teljes zónára különböző paraméterek függvényében, valamint gyorsreaktorokra jellemző effektusok vizsgálata, kutatása. Az eredményekből várhatóan cikk születik.		
High-tech hunt for secrets - ERC	Szociológiai Intézet	Takács Károly
A kutatásban okoseszközöket használunk szociometriai és hálózatkutatási célokra.		

Agrodat.hu project

AGRØDAT

Main objective: **knowledge centre** and **decision support** system

- based on data gathered by an innovative, ***complex sensor*** system and from international ***open repositories***
- relying on ***big data, cloud***, and ***HPC*** technologies

to support **precision agriculture**.

Duration: 2014-2017
Budget: appr. 8 MEUR
URL: www.agrodat.hu

Consortium:

eNET
Hewlett Packard Enterprise
MTA SZTAKI Hungarian Academy of Sciences Institute for Computer Science and Control
SZÉCHENYI ISTVÁN EGYETEM SZE · GYŐR

center:
844 CPU Core
5274 GB Memory
564 TB SSD/HD

GPGPU:
21504 CUDA Core
488 Xeon Phi Core

Network:
40 Gb / Infiniband for HPC
10 Gb copper
1 Gb copper for mgm.
8/16 Gb FC for SAN
Connected to HBONE

Cloud middleware
HPE Helion
openstack

SZÉCHENYI 2020
HUNGARIAN GOVERNMENT
NATIONAL RESEARCH, DEVELOPMENT AND INNOVATION FUND
INVESTING IN YOUR FUTURE

Scalable IoT back-end for IaaS clouds

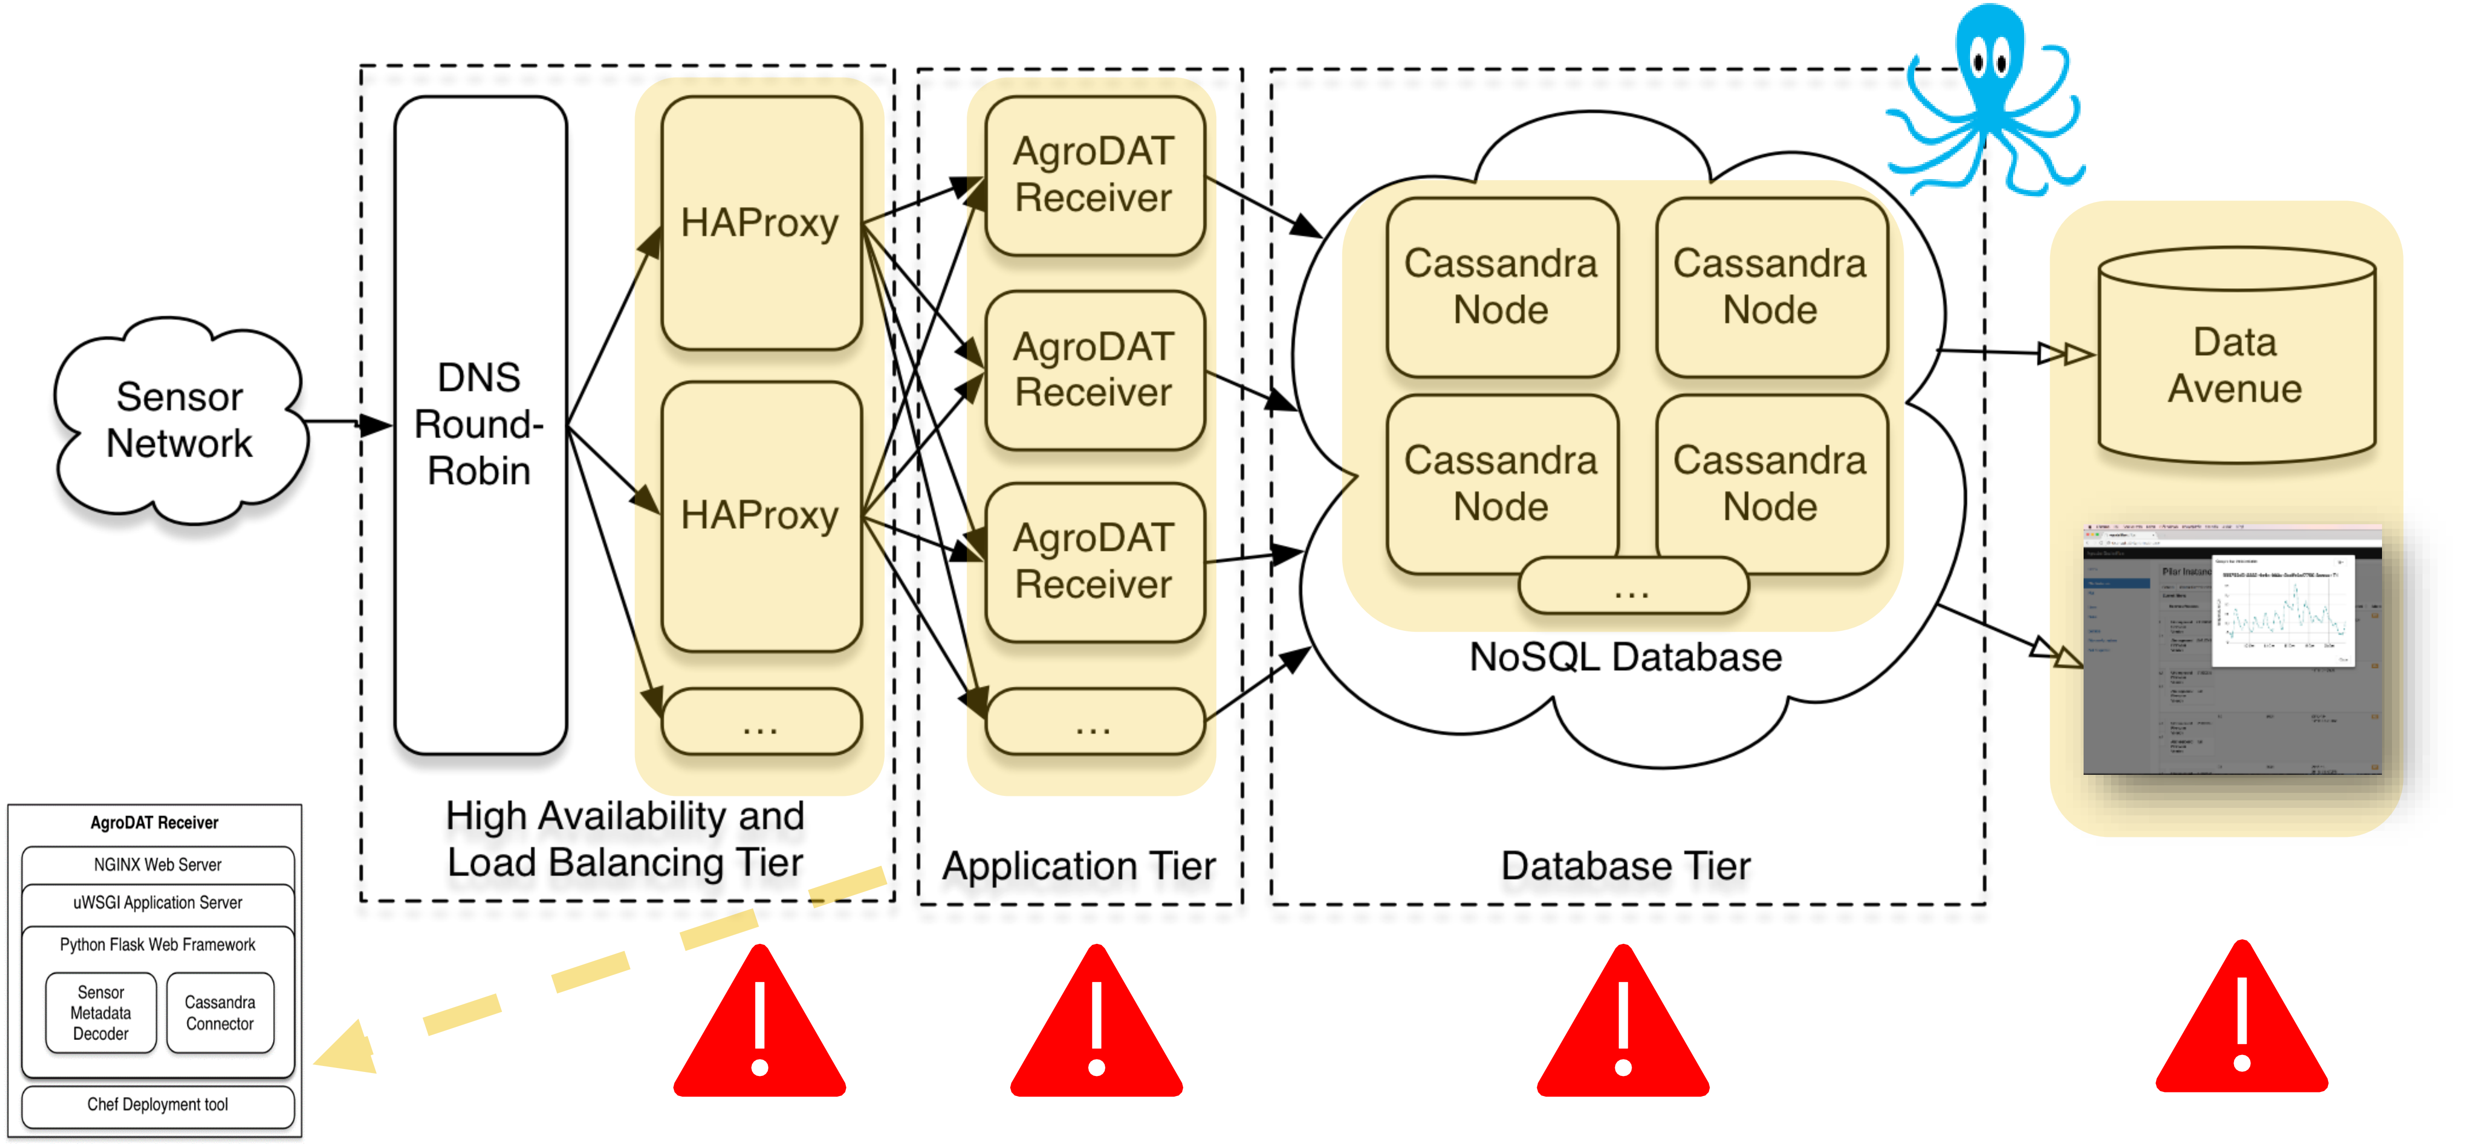

Connected cars project:
Controller Area Network (CAN) Data Collector and its applications (2014-2017)

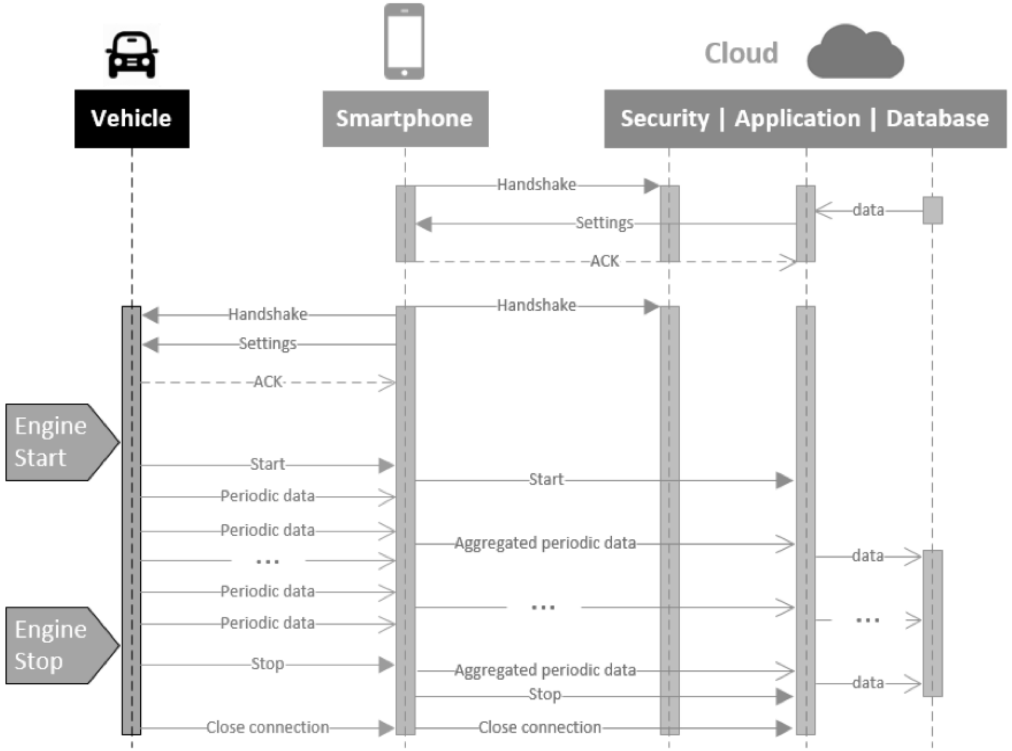

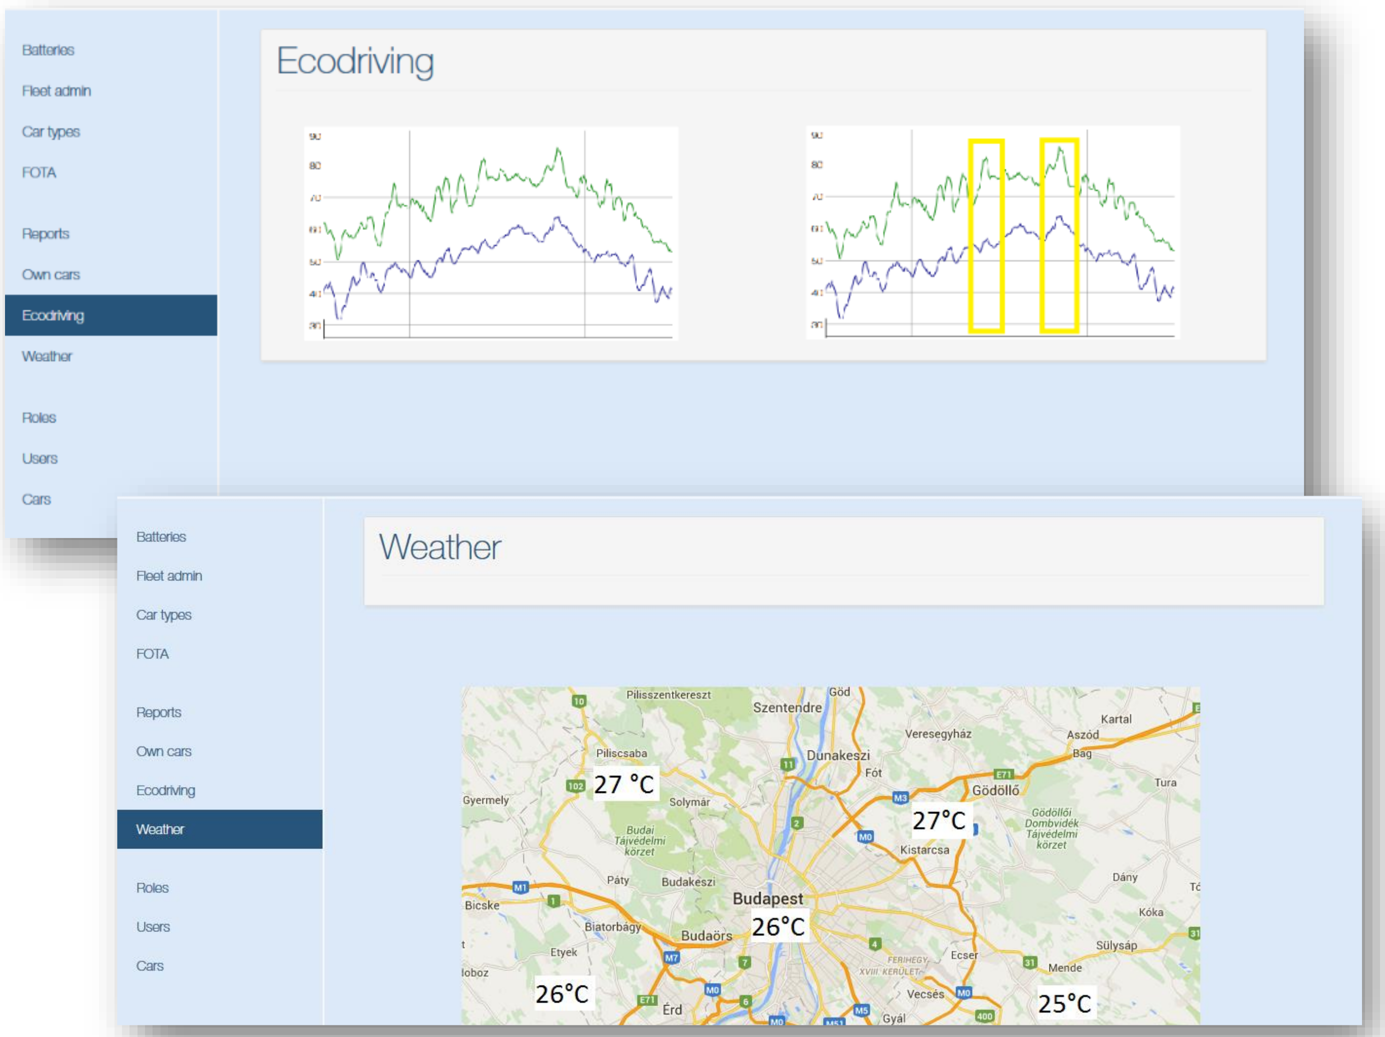

Three tiers of the IoT back-end architecture for Connected Cars

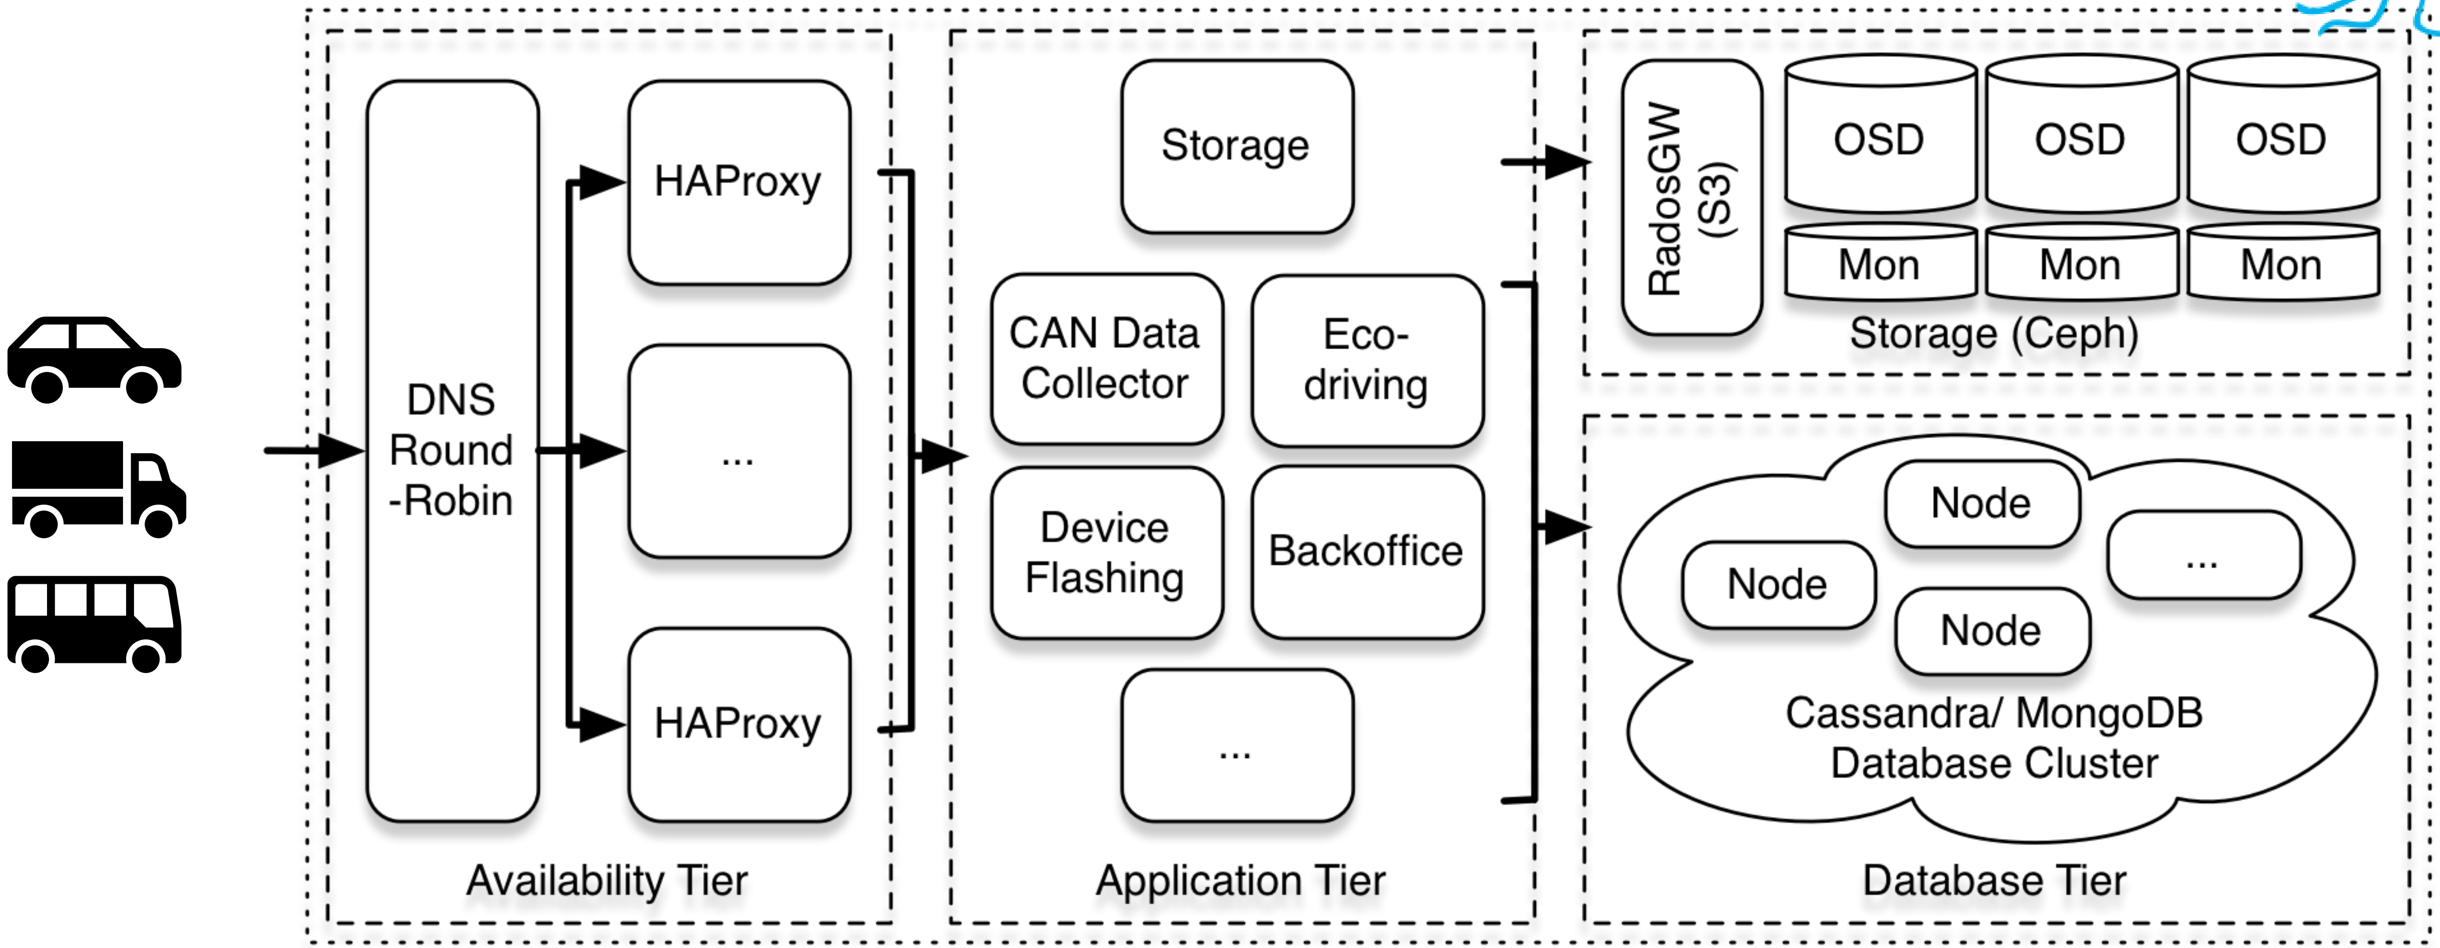

What is the next step?

NEW GENERATION OF

REFERENCE ARCHITECTURES

with advanced, smart orchestration methods

(along with new wave of AI tools and platforms)

Background: Reference architectures for AI/ML

- Machine Learning-as-a-service (MLaaS) on cloud supports various fundamental steps for machine learning (ML), such as **data pre-processing, model training, and model evaluation, with further prediction**.
- Representative **commercial cloud** providers with MLaaS solutions:
  - Amazon SageMaker,
  - Azure Machine Learning Studio, etc.
  - Google Cloud AI hub / platform, etc.

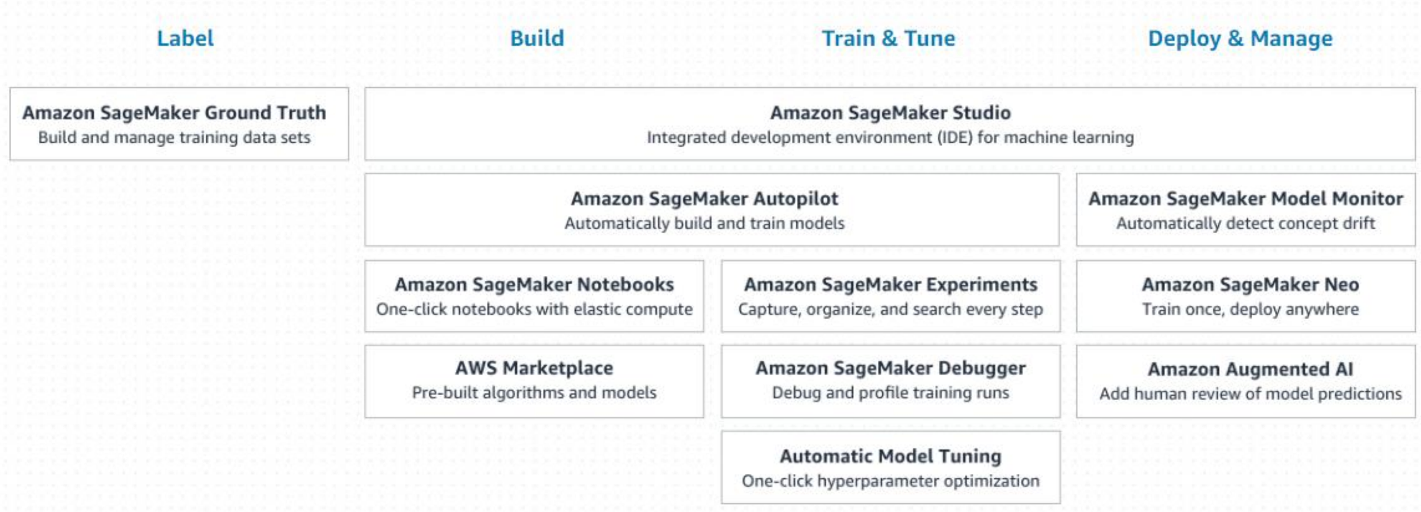

- For the various AI application scenarios, the major cloud providers offer **reference architectures with building blocks/connections** (such as Azure Reference Architectures) including
  - **recommended practices**, along with considerations for
  - **scalability**,
  - **availability**,
  - **manageability**,
  - and **security**.

➔ NOT ONLY "GLUING" TOGETHER SERVICES BUT
STRONG FOCUS ON NON-FUNCTIONAL REQUIREMENTS AND FEAUTURES

Background: Reference architectures for AI/ML (cont.)

- Similar state-of-the-art reference architectures are available from several **HPC vendors**:
  - Hewlett Packard Enterprise elaborated its reference architecture for AI
    - copes not only with open-source software components, but their own *propriety cluster-based hardware* platform as well
  - IBM and other vendors provide similar solutions, etc.
- All these approaches leverage mostly on *open source tools* and frameworks, such as TensorFlow or Apache Spark.

- Concentrating on the **manufacturing sector**, a reference architecture has been recently published by Fraunhofer IOSB (Q3/2019).
  - Designed for scalable data analytics in smart manufacturing systems, and complies with the higher-level *Reference Architecture Model for Industrie 4.0* (RAMI 4.0).
  - Implemented and validated in the Lab Big Data at the SmartFactoryOWL based on various open-source technologies (Spark, Kafka, Grafana).

A reference architecture and its implementation from Fraunhofer IOSB

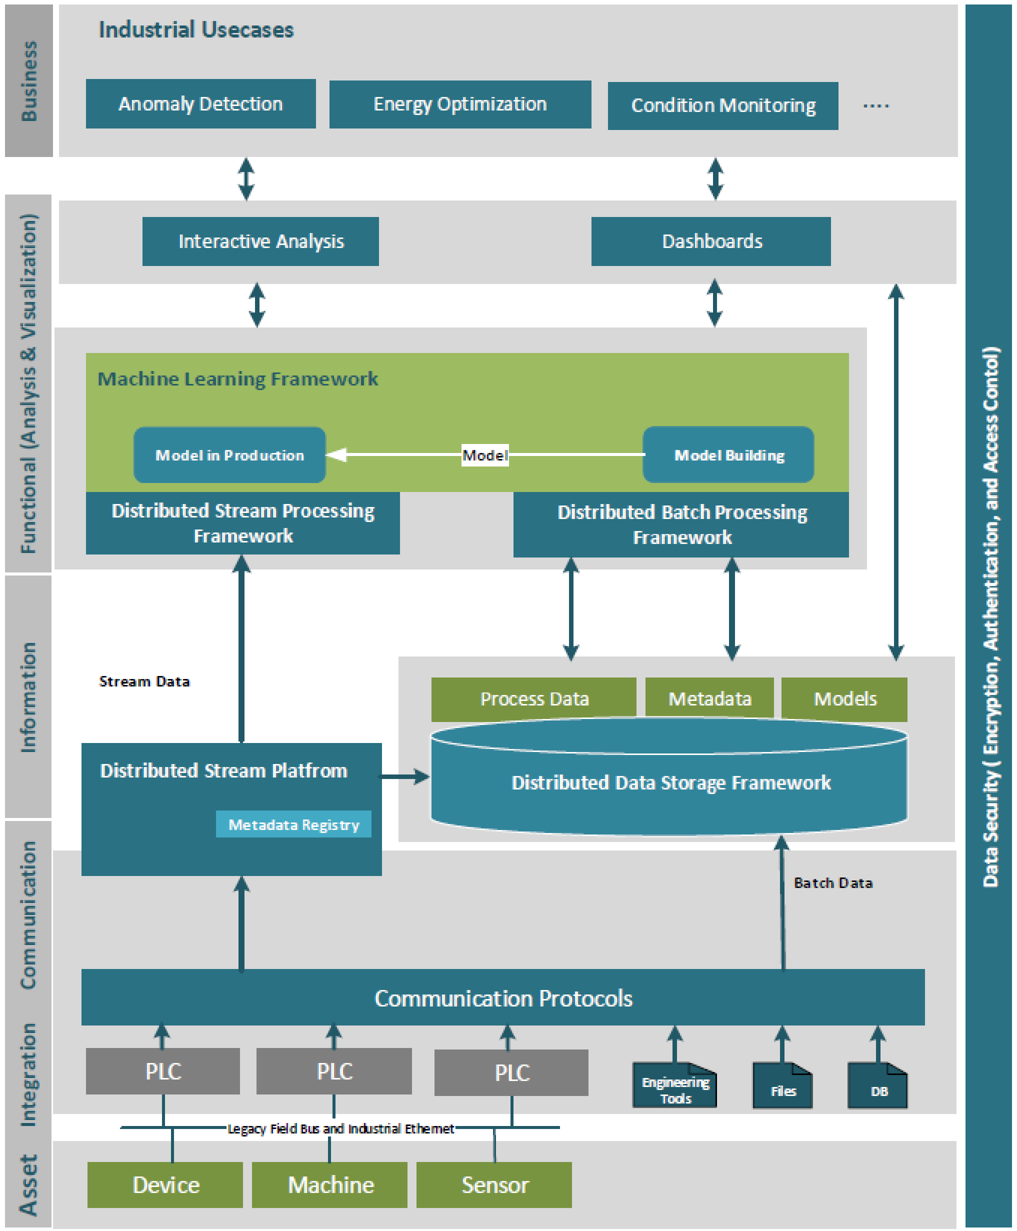

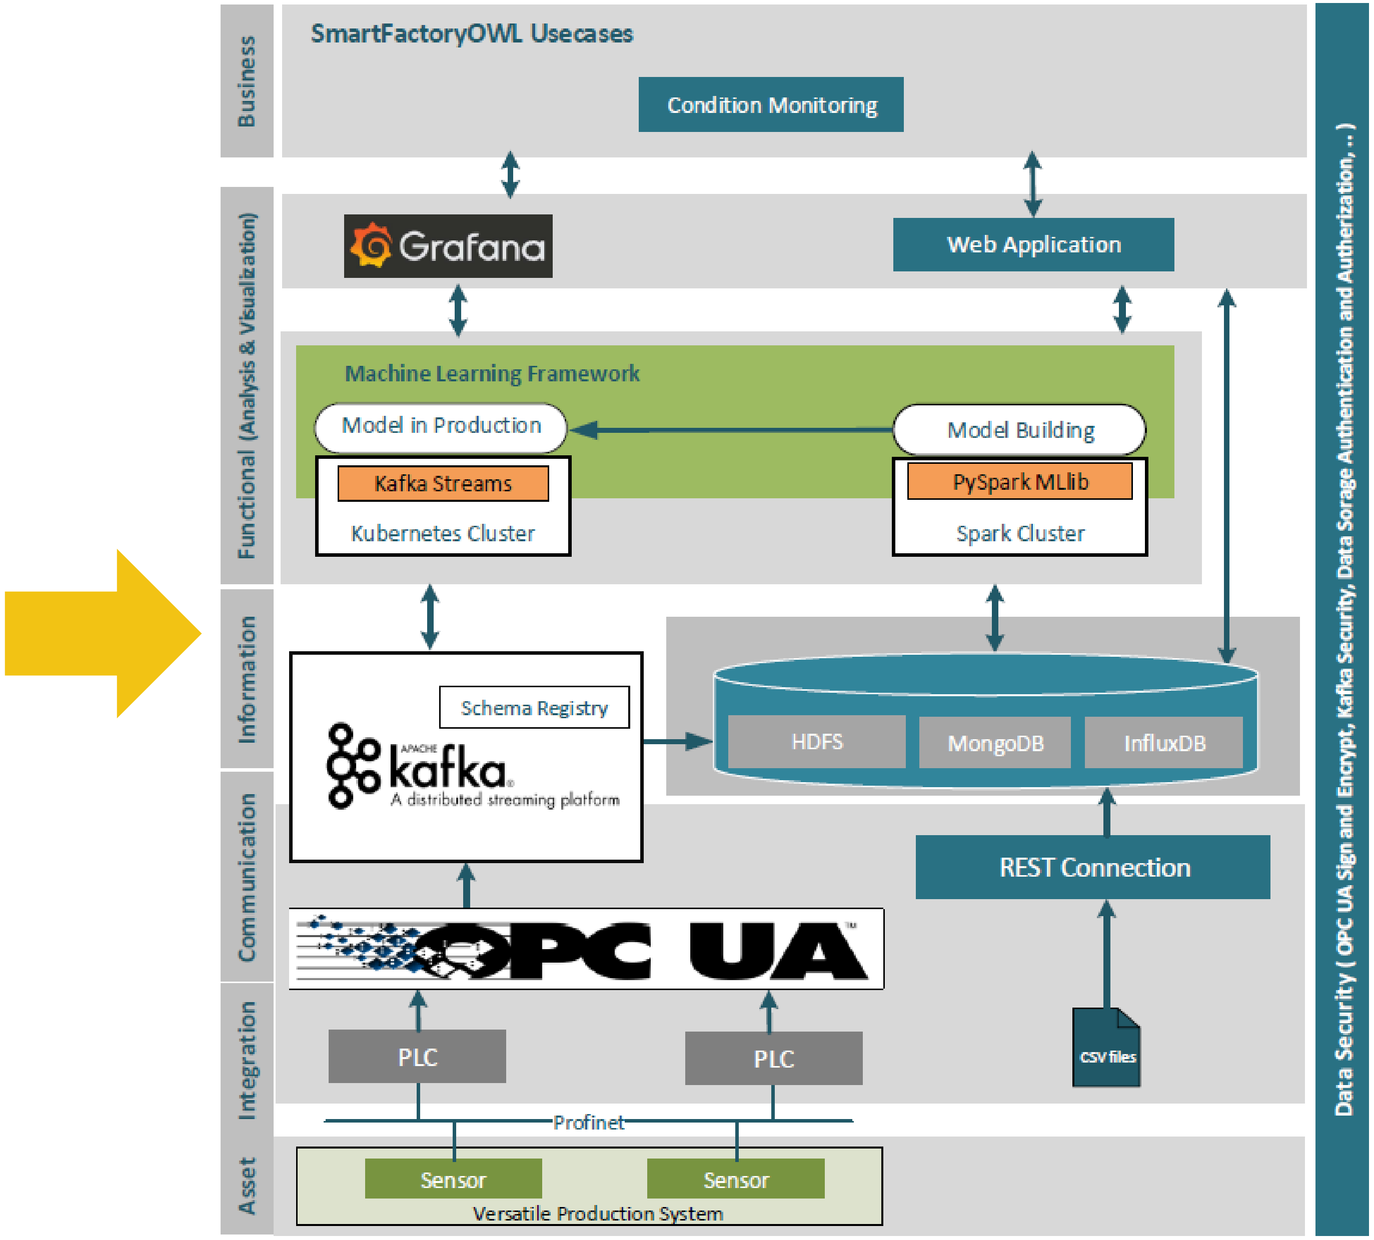

New orchestrated AI/ML reference architectures
from SZTAKI

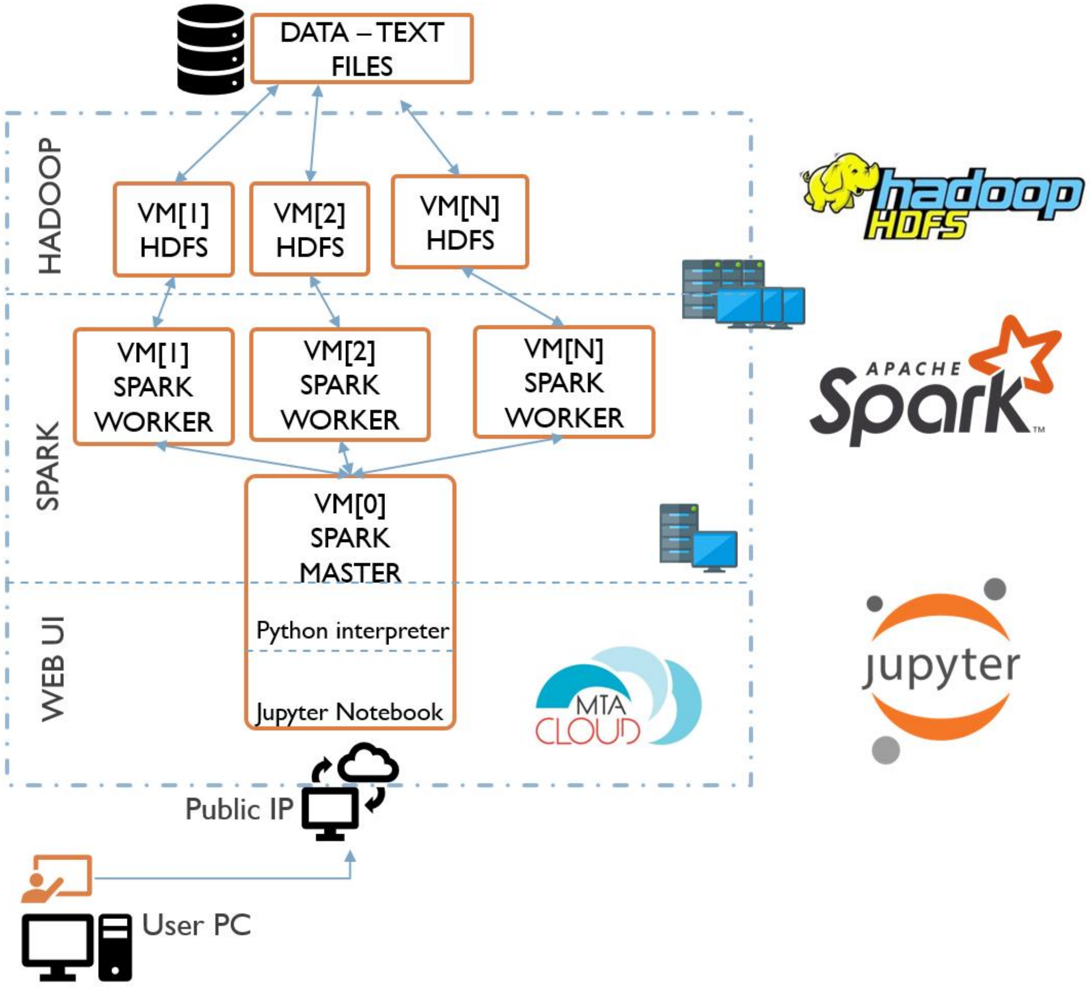

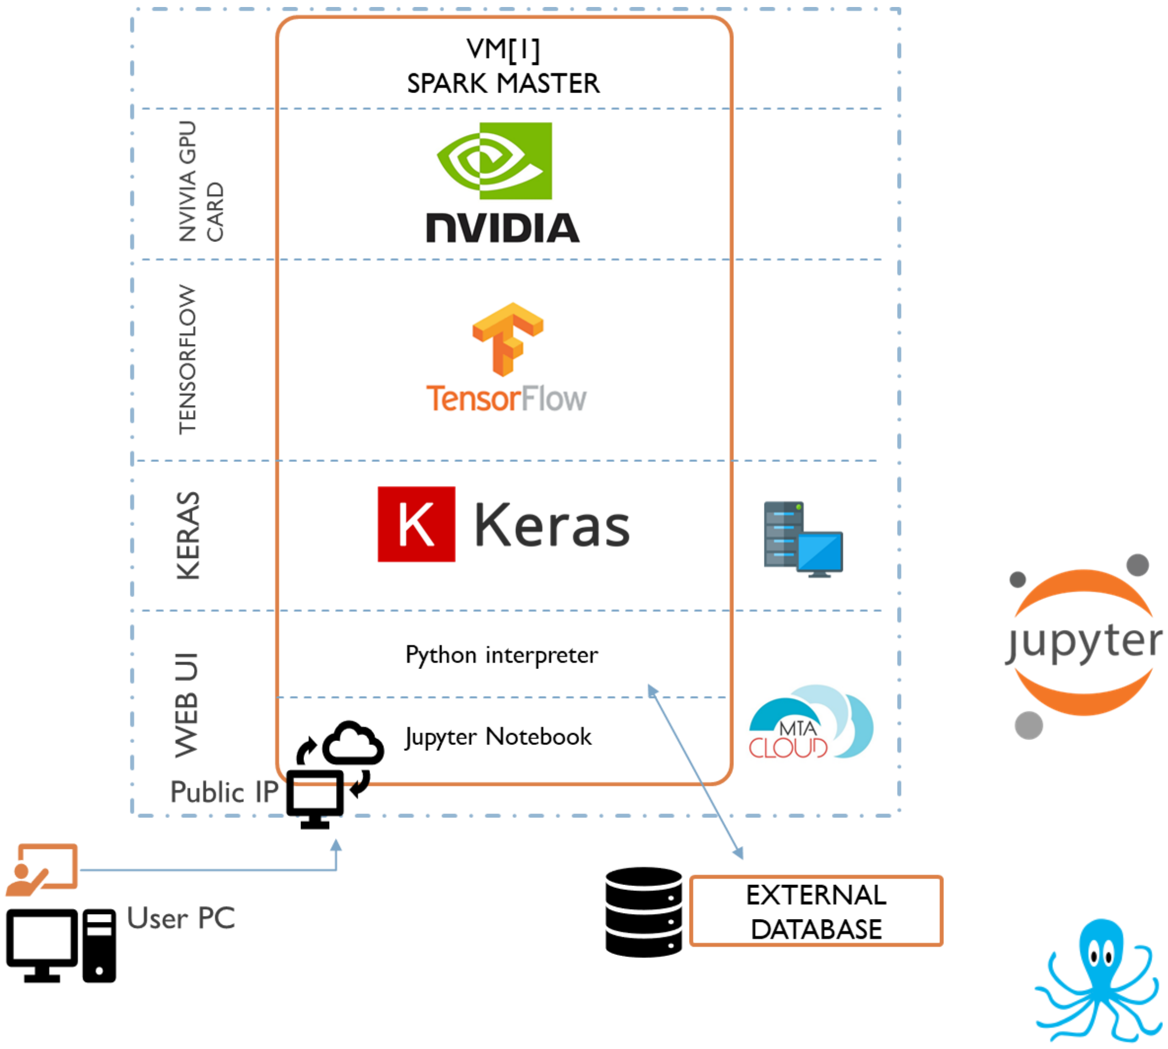

Further orchestrated data-oriented reference architectures
from SZTAKI

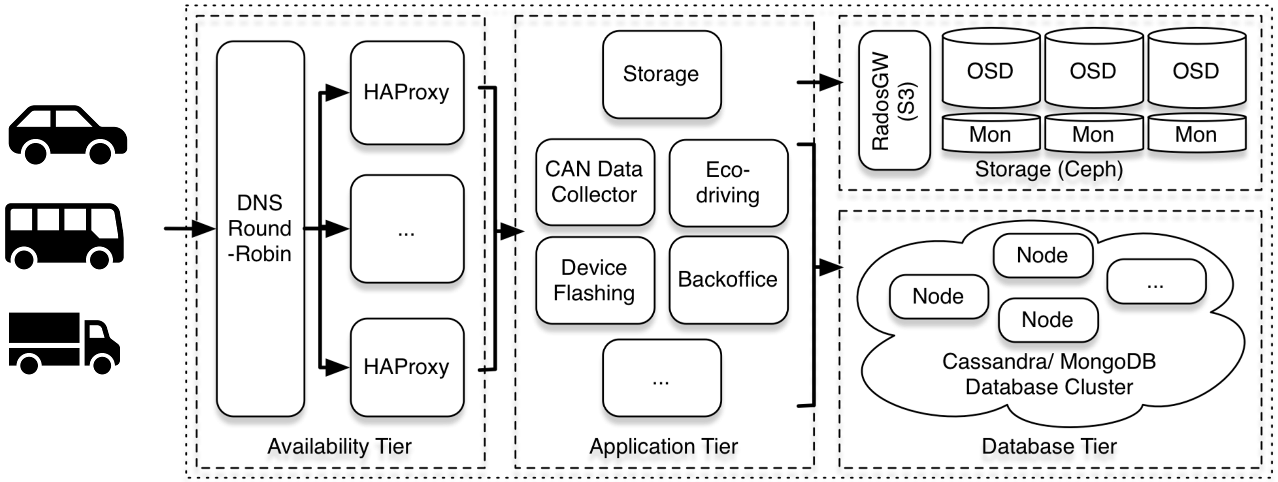

Various application areas (agriculture, connected cars, CPPS, ...)

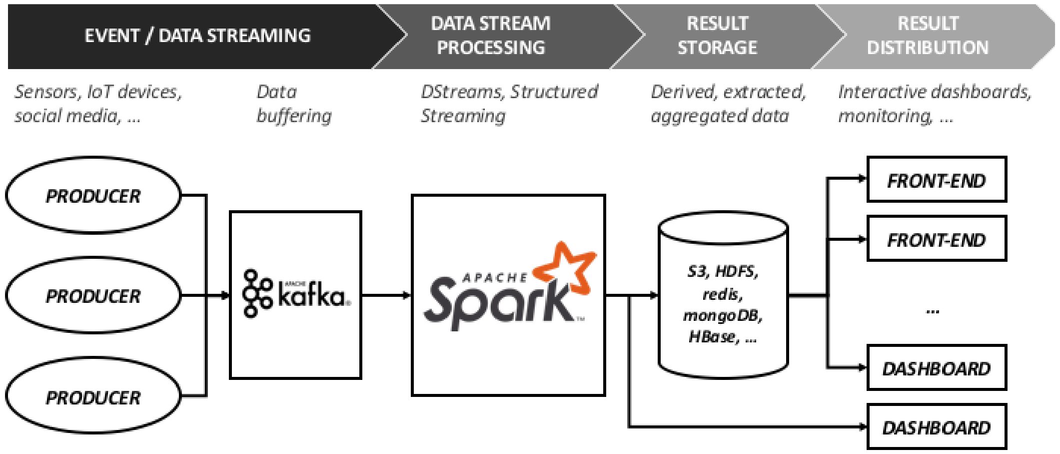

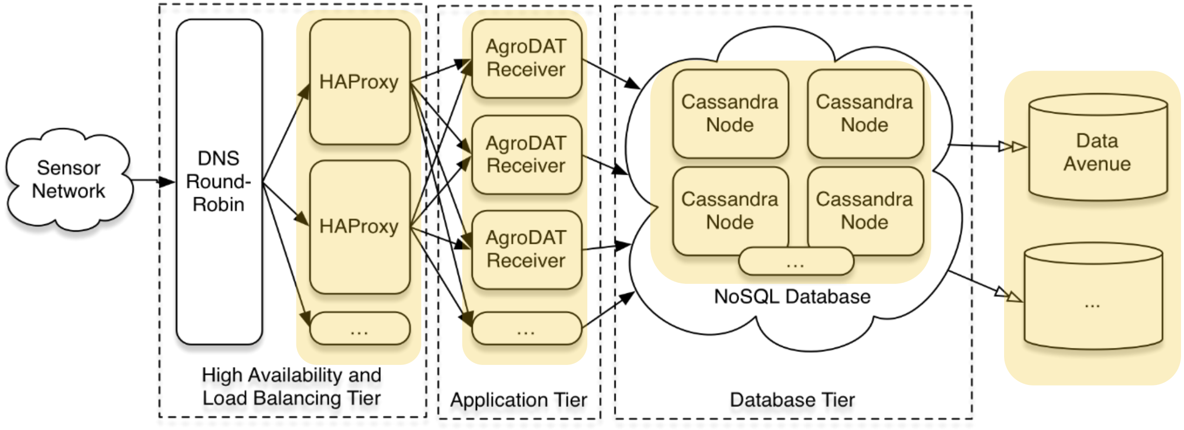

Background: autoscaling

- **Autoscaling** is a method used particularly in cloud computing, where the amount of computational resources in a server farm/cluster scales automatically based on the load on the farm.
- ***Predictive autoscaling*** may combine recent usage trends with historical usage data as well as other kinds of data to predict usage in the future, and autoscale based on these predictions.

➔ Joint development with University of Westminster

MiCADO

Microservices-based Cloud Application-level Dynamic Orchestrator

- **Open source automated cloud deployment** and **autoscaling** framework
- Automated scaling based on **highly customisable scaling policies**
- Scaling at both **container and virtual machine levels**
- **Multi-cloud** support
- **Standardised** TOSCA-based application and policy description
- **Modular architecture** based on open source components
- **Easy** Ansible-based **deployment**
- Intuitive **dashboard**
- Policy driven **advanced security settings**

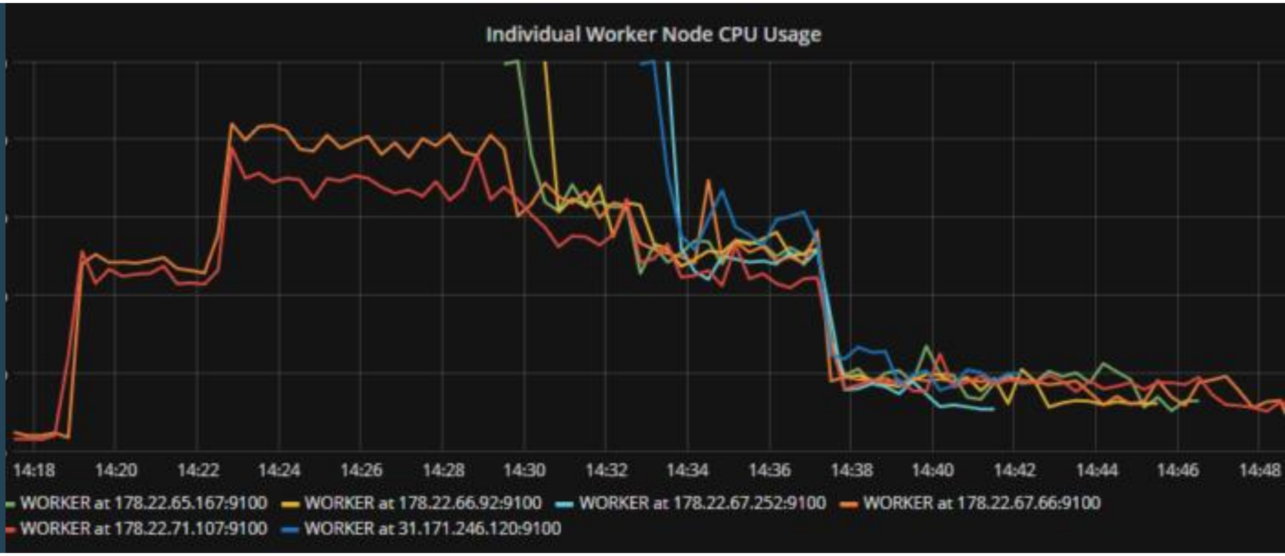

MiCADO: autoscaling framework

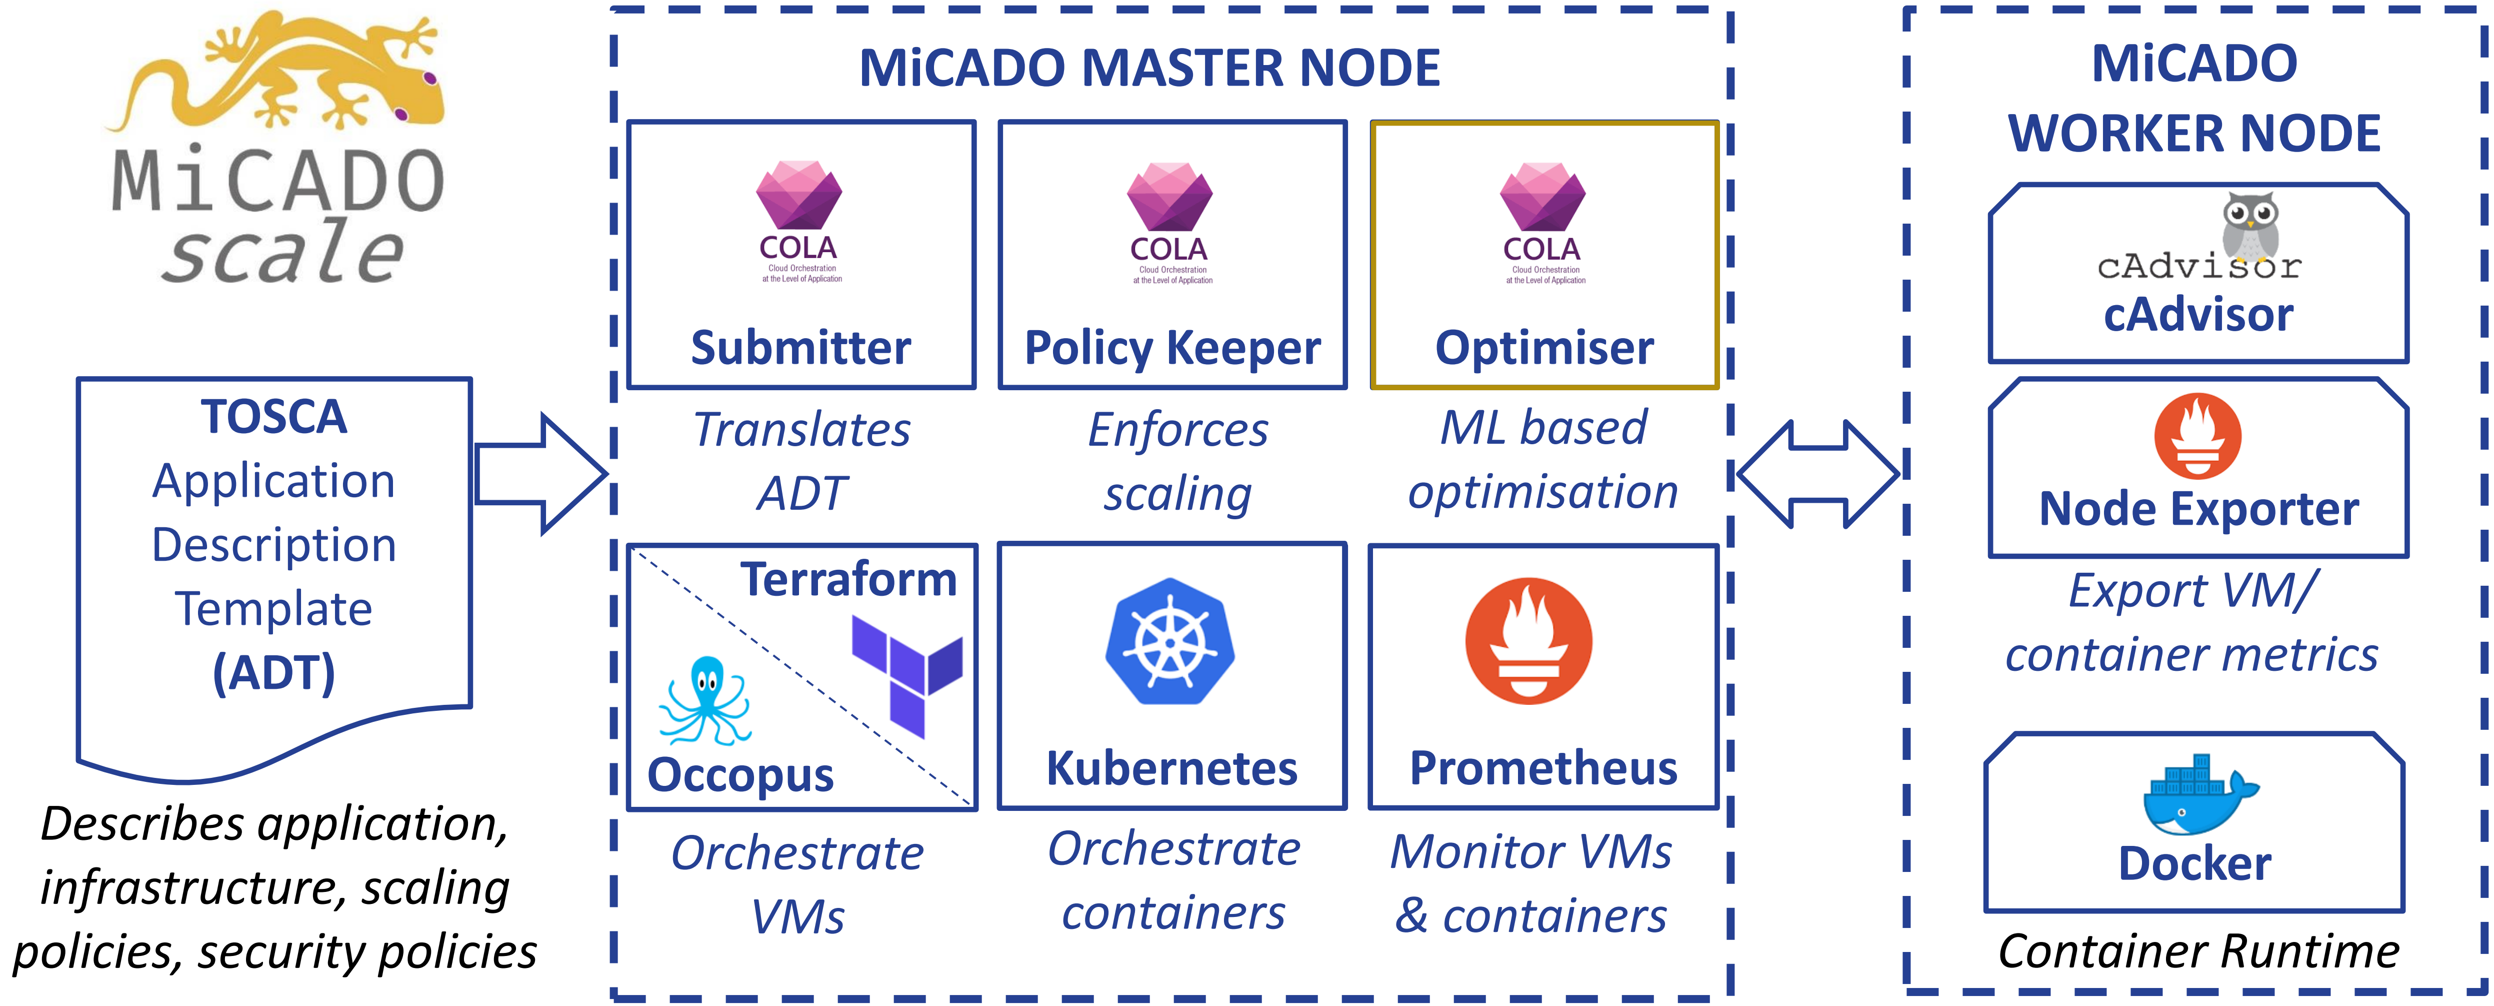

MiCADO: customizable autoscaling and deployment framework (1/2)

- Highly customizable monitoring subsystem
  - Monitored metrics are collected by **dynamically attachable data collectors** (Prometheus exporters)
- Highly customizable scaling logic
  - Scaling logic is **fully programmable** (using Python)
- Many different scaling policies are supported
  - **Application types** (job execution, web applications, …)
  - Different **metrics** (cpu, network, number of jobs, …)
  - Various **strategies** (load-based, deadline-based, event-based, …)
- Scaling containers and virtual machines are supported
  - Scaling at **both levels in parallel**, independently or cooperatively
  - **Container-only** and **VM-only** scaling
- Possible to use **predefined** scaling policy or **own-developed**

```
- scalability:
    type: tosca.policies.Scaling.MiCADO
    targets: [ cqueue-worker ]
    properties:
      min_instances: 1
      max_instances: '{{MAXCONTAINERS}}'
      scaling_rule: |
        required_count = 0
        if ITEMS>0:
          required_count = ceil(AET/(REMAININGTIME/ITEMS)) if REMAININGTIME>0
          else MAXCONTAINERS
          m_container_count = min([required_count, len(m_nodes) * 5])
        else:
          m_container_count = 0
```

```
tosca.policies.Scaling.MiCADO.VirtualMachine.CPU.stressng:
  derived_from: tosca.policies.Scaling.MiCADO
  description: base MiCADO policy defining data sources, constants, queries,
  alerts, limits and rules
  properties:
    alerts:
      type: list
      description: pre-define alerts for VM CPU
      default:
      - alert: node_overloaded
        expr: '(100-(avg(rate(node_cpu_seconds_total{node="{{ NODE_NAME }}",
        mode="idle"}[60s]))*100)) > {{NODE_TH_MAX}}'
        for: 1m
      - alert: node_underloaded
        expr: '(100-(avg(rate(node_cpu_seconds_total{node="{{ NODE_NAME }}",
        mode="idle"}[60s]))*100)) < {{NODE_TH_MIN}}'
        for: 1m
      required: true
    scaling_rule:
      type: string
      description: pre-define scaling rule for VM CPU
      default: |
        if len(m_nodes) <= m_node_count and m_time_since_node_count_changed > 60:
          if node_overloaded:
            m_node_count+=1
          if node_underloaded:
            m_node_count-=1
        else:
          print('Transient phase, skipping update of nodes...')
      required: true
```

MiCADO: customizable autoscaling and deployment framework (2/2)

- Scaling **multiple, independent set** of virtual machines

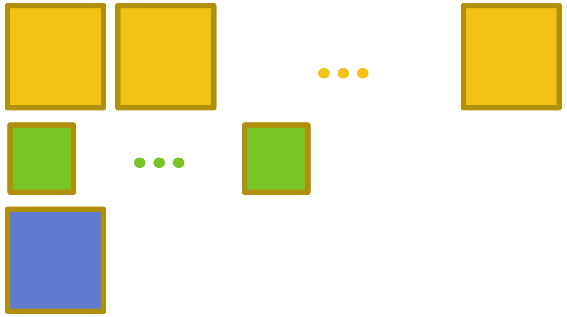

- Supporting **multi-cloud** environment

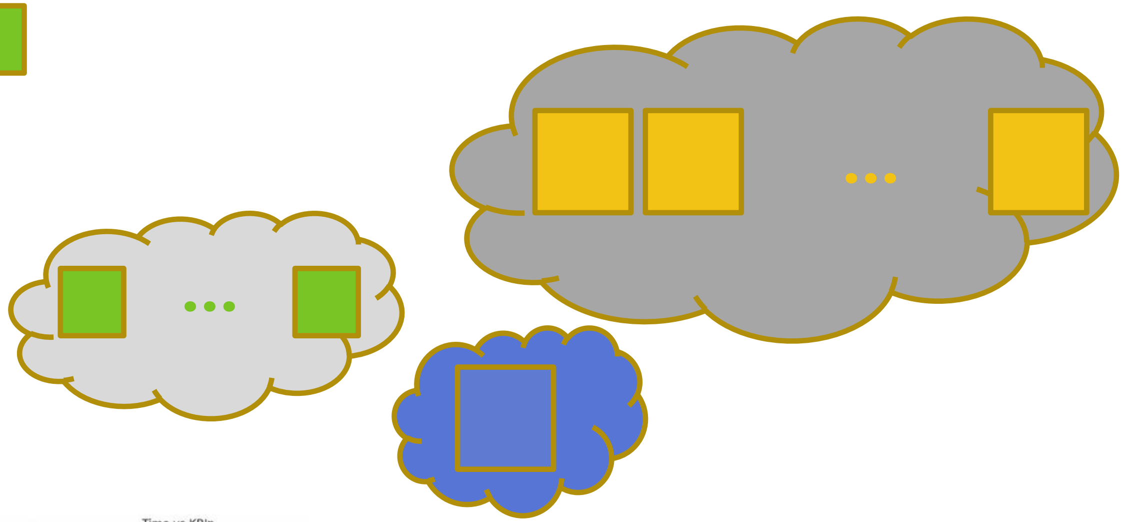

- Supporting **optimizer-based** scaling

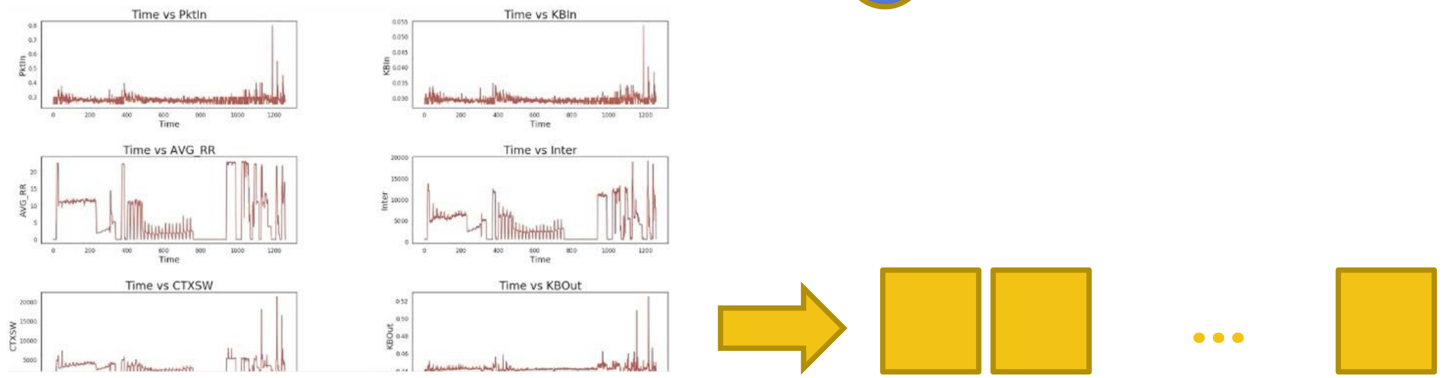

Concept of Optimizer: Machine-Learning

- **Target** variable to be kept in range
- **System** variables in relation
- Continuously **monitoring** the system and target variables
- **Neural network** for learning the relation **between the system and target** variables

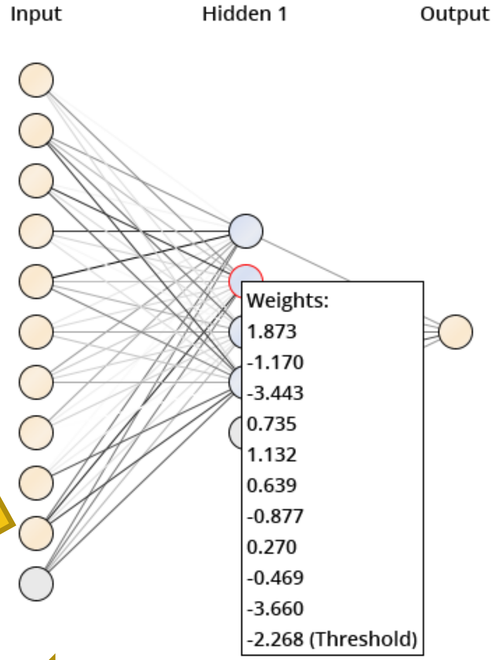

- **Linear regression** to learn the relation between **before and after scaling** event of system variables
- Combining the two aproximators enables to **predict the effect of any scaling action**
- Scaling action with **best predicted result is selected** and suggested

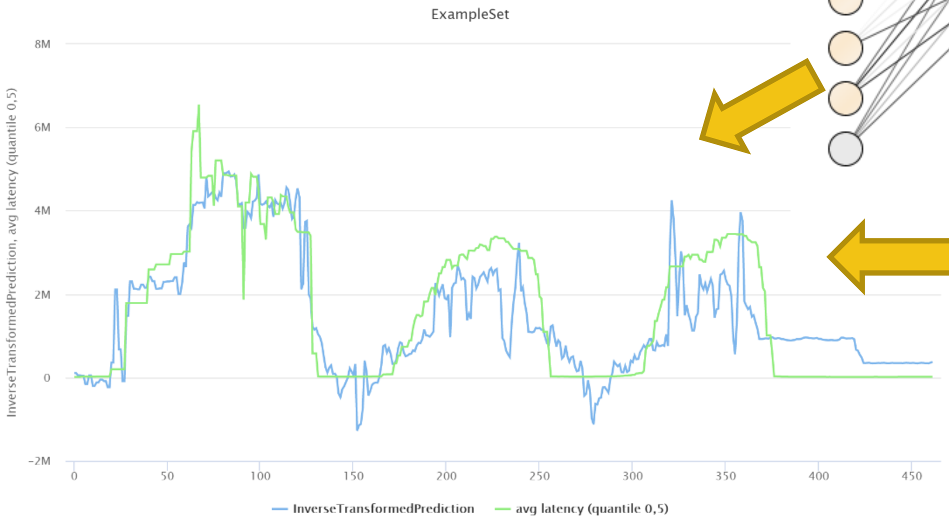

- **More training improves prediction**

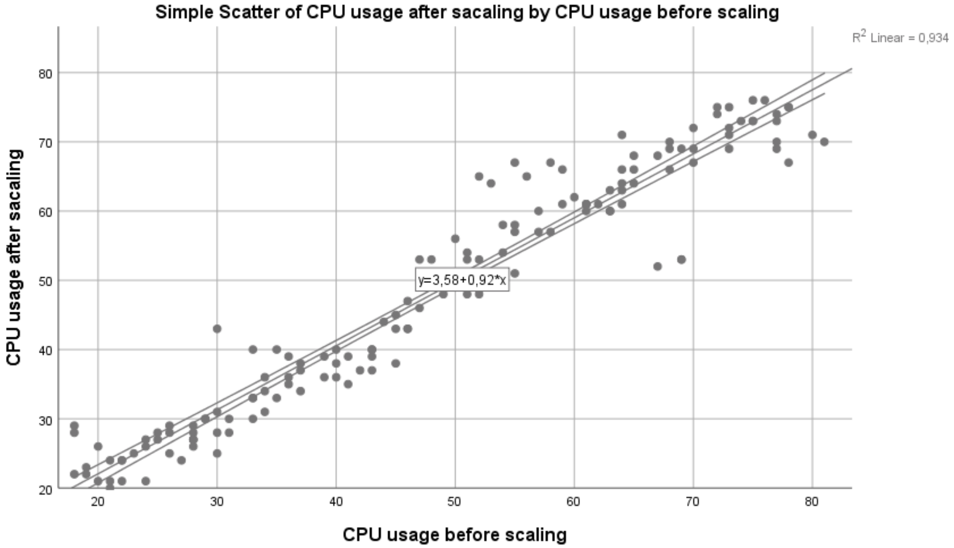

Towards smart orchestration of reference architectures
➔ higher **availability** by detecting/predicting critical failures

"Modelling and enhancement of orchestration methods for virtual research platforms with machine learning" project (2019-2023)

supported by OTKA Hungarian Scientific Research Fund

The research addresses the following phenomena and methods:

1. Large number of hierarchical orchestration steps with dependencies
2. The non-deterministic and dynamic behaviour of cloud (and similar) environments with probe-effect

➔ Steered, automated traversing and verification of consistent global states (based on machine learning)

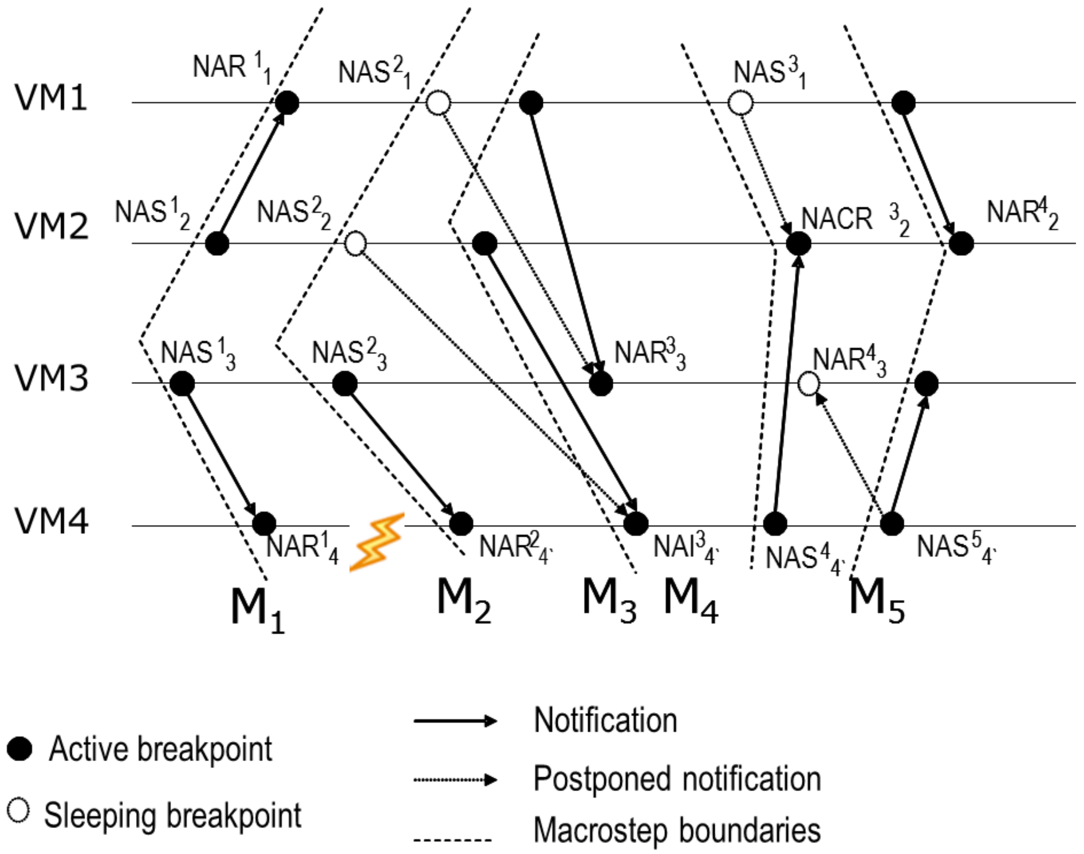

What is the next step?

APPLICATION / DEVELOPMENT OF

REFERENCE ARCHITECTURES

SZTAKI services and on-going international R&D projects

- *Deployment* and configuration of Infrastructure-as-a-service (IaaS) clouds
- Support for *cost efficient* application on cloud and microservice platforms
- *Orchestration of complex processes in clouds and/or on software container based platforms*
- *Design and development of vendor independent (cloud agnostic) prototypes and pilots*

- *EU H2020 CloudiFacturing project: www.cloudifacturing.eu*
- *EU H2020 EOSC-hub project: www.eosc-hub.eu*
- *EU H2020 NEANIAS project: www.neanias.eu*

Cloudi Facturing

Cloudification of Production Engineering for Predictive Digital Manufacturing

Consortium:
Fraunhofer IGD (coordinator), DIHs, ISV, SMEs, etc.
SZTAKI role: WP leader
Duration: 2017-2021
Budget: 9.7 MEUR

➔ www.cloudifacturing.eu

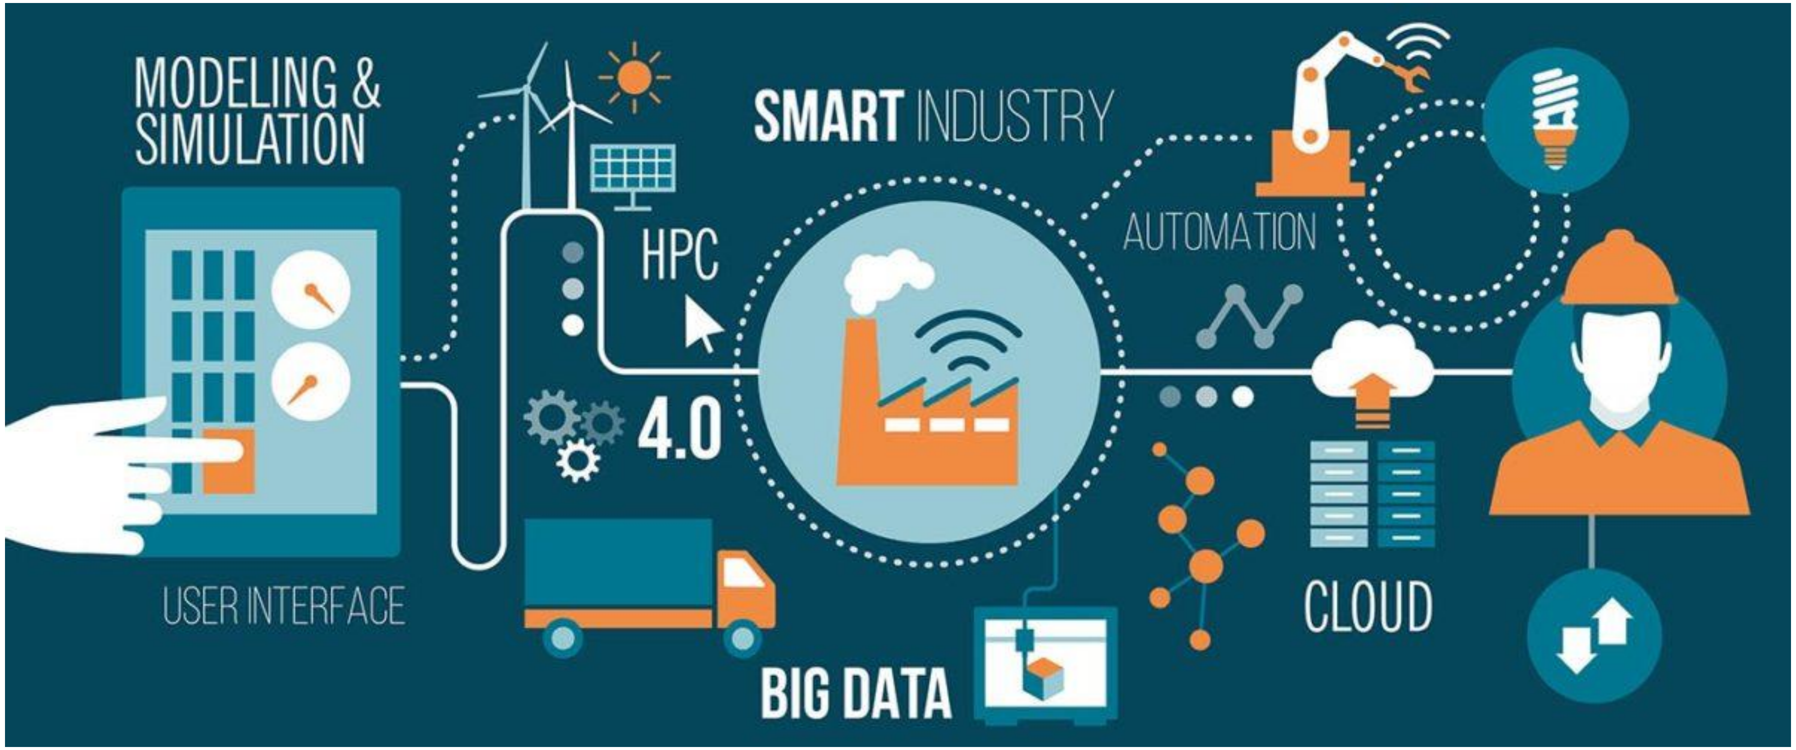

Factories of the Future
Public Private Partnership

I4MS

Funded by the Horizon 2020 Framework Programme of the European Union

7+14 Experiments and 4 Digitial Innovation Hubs

1. Optimizing design and production of electric drives
2. Cloud-based modelling for improving resin infusion process
3. Improving quality control and maintenance at manufacturing SMEs using big data analytics
4. Numerical modelling and simulation of heat treating processes
5. Optimizing solar panel production
6. Optimizing efficiency of truck components manufacturing processes by data analytics
7. Simulating and improving food packaging
8. …

CLOUDIFACTURING: MISSION STATEMENT

•The mission of Cloudi*Facturing* is to **optimize production processes and producibility**

- using **Cloud/HPC-based modelling and simulation**,
- leveraging **online factory data and advanced data analytics**,
- thus contributing to the **competitiveness and resource efficiency** of manufacturing SMEs,
- ultimately fostering the vision of **Factories 4.0 and the circular economy**.

•**Vision**: closing the loop from factory data back to simulation and forward to influencing factory processes (with support in real-time).

LIVE DIGITAL TWIN MODEL

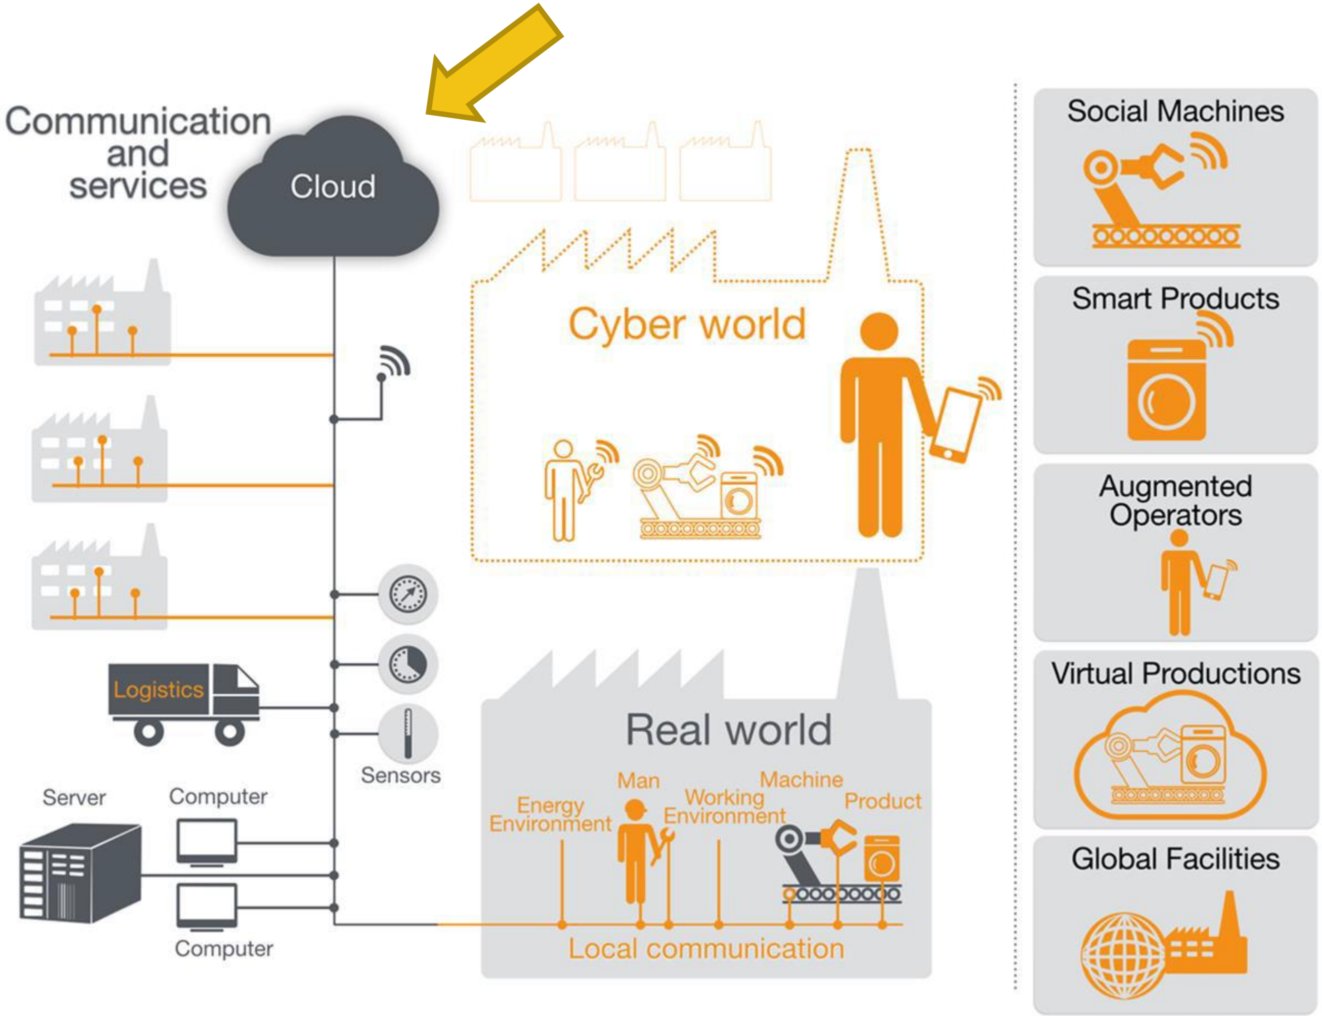

Cloudi*Facturing*: Integrated Platform

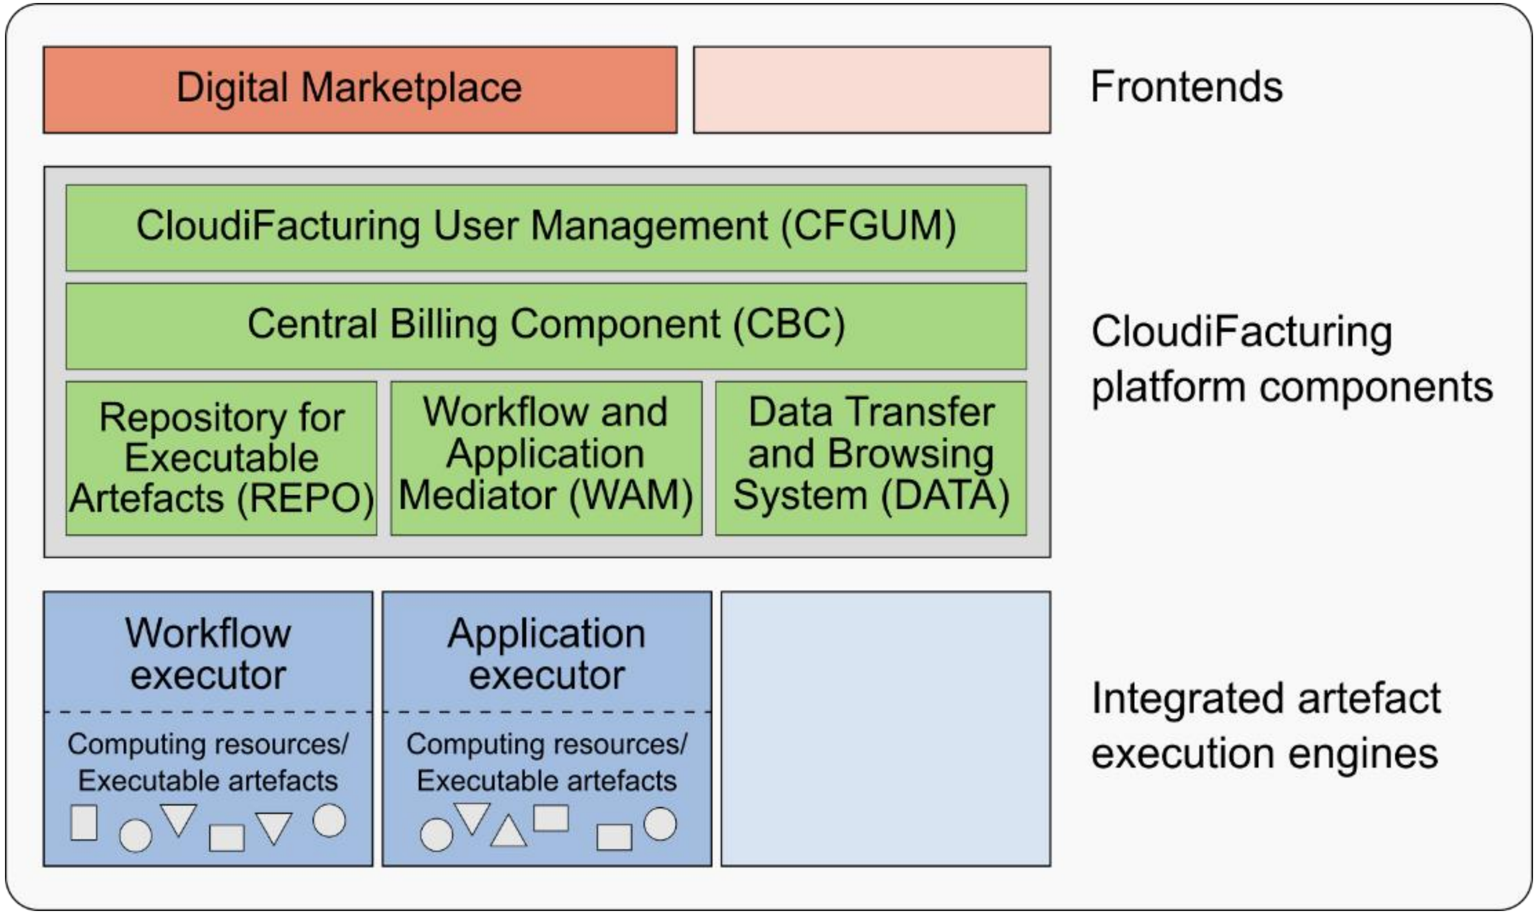

NEW SUPPORTED EXPERIMENT: ERGONOCLOUD (2020-2021)

European Open Science Cloud

- European Commission Horizon 2020 Programme
- 100 Partners, 76 beneficiaries (75 funded)
- 3874 PMs, 108 FTEs, more than 200 technical and scientific staff involved
  - Budget: €33,331,18, contributed by:
- 36 months: Jan 2018 – Dec 2020

- EOSC-hub brings together multiple service providers to create the Hub:
  - a single contact point for European researchers and innovators
  - to discover, access, use and reuse a broad spectrum of resources
  - for advanced data-driven research.
- For researchers, this will mean a broader access to services supporting their scientific discovery and collaboration across disciplinary and geographical boundaries.

Source: EOSC-hub

The project creates **EOSC Hub**:

a federated integration and management system for EOSC

From the consortium AND from **external contributors**

- Data
- Applications & tools
- Baseline services (storage, compute, connectivity)…
- Training, consultants

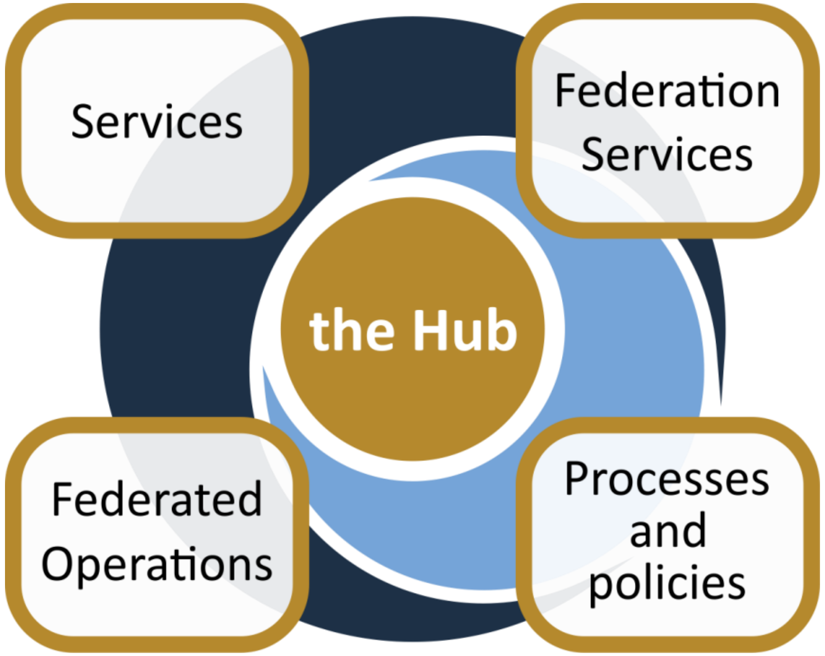

- Marketplace
- AAI
- Accounting
- Monitoring
- …

Usage according to **Rules of Participation**

- Lightweight certification of providers
- SLA negotiation
- Customer Relationship Management
- …

Based on **FitSM**

- Security regulations,
- Compliance to standards,
- Terms of use,
- FAIR implementation guidelines
- …

NEANIAS – Novel EOSC Services for Emerging Atmosphere, Underwater & Space Challenges

Call: INFRAEOSC-02-2019
Project ID: 863448
Duration: 36 months (started from 2019/Q4)

NEANIAS

NEANIAS: Services – Infrastructures – Communities & Business

Complex services where reference architectures are to be used

Technology Stack

Underwater	Atmospheric	Space
U1 – Bathymetry Mapping from Acoustic Data U2 – Seafloor Mosaicking from Optical Data U3 – Seabed Classification from Multispectral, Multibeam Data	A1 – Greenhouse Gases Flux Density Monitoring A2 – Monitor atmospheric Perturbations and Components in Active Tectonic Regions A3 – Air Quality Estimation, Monitoring and Forecasting	S1 – FAIR Data Management and Visualization for Complex Data and Metadata S2 – Map making and mosaicking for multidimensional images S3 – Structure detection on large map images with machine learning techniques

Core Services

C1 – Open-Science lifecycle support services Cataloguing, Discovery, Validation, Identification	C2 – EOSC hub, RIs and cloud integration enabling services Compute, Store, Share, AuthN/AuthZ	C3 – AI services Cataloguing and discovery, Machine Learning	C4 – Visualisation services 3D, 2D rendering, Virtual Reality, Augmented Reality

EOSC / Research Infrastructures' Services

EOSC Hub Marketplace, EOSC Hub Portal, B2Store, B2Share, B2Access, D4Science VRE services,

Conclusion

- **Divergence**: cloud providers, open-source / proprietary code, software stack
- **Convergence**: standards, containers, orchestration, **reference architectures**

1. Potential advantages of orchestrated reference architectures:
   - faster development/delivery
   - lower costs at each stage
   - higher quality (user satisfaction)
2. ... still several challenges: how to make them really **“smart”** by
   1. addressing *all* the typical non-functional requirements, and
   2. covering *every* complex application area/sector

Thank you for your attention!

SZTAKI

Robert Lovas
robert.lovas@sztaki.hu

SZTAKI

About the institute

- Established in 1964
- EU Centre of Excellence in IT (since 2001)
  - Computer Science and Control
- Basic and applied research
- Contract-based R&Đ&I activity mainly on complex systems, turnkey realizations:
  - GE, Hitachi, Audi, Hungarian Telekom, MOL, Knorr-Bremse, Bosch, Opel, ESA, etc.
- Transferring up-to-date results to industry and universities

- Basic research
  - Computer science
  - **Systems- and control theory**
  - **Engineering and business intelligence**
  - Machine perception and human-computer interaction
- Applied research and innovation
  - Vehicles and transportation systems
  - **Production informatics and logistics**
  - **Distributed, cloud / Big Data / AI computing**
  - Energy and sustainable development
  - Security and surveillance
  - Networking systems and services

- Budget
  - 12-13 MEUR/year
  - ~30% basic funding
- Staff
  - 280
  - 67% scientific
- **Fraunhofer Project Center**
  - Production Management and Informatics from 2010
- W3C member
- ERCIM member
- Founder of MTA Cloud

- **Leader of Hungarian National Technology Platform on Industry 4.0**
- EPIC Centre of Excellence
- AI National Lab (from 2020)