Information & Brokerage Session

Horizon Europe Cloud-Edge-IoT Call 2022

2 February 2022 | 9:30 - 12:30 CET, online

Organised by

European Commission

In collaboration with

HORIZON CLOUD

NEXT GENERATION IoT

Information & Brokerage Session

Horizon Europe Cloud-Edge-IoT Call 2022

10:45
12:00

PITCH SESSION

Cognitive Cloud

10:50 - 10:55 **Alberto P. Marti,** *OpenNebula*
10:55 - 11:00 **Cagatay Yilmaz,** *Green Continuum Information Technologies Ltd.*
11:00 - 11:05 **Tuna Tugcu,** *Bogazici University*
11:05 - 11:10 **Tatiana Silva,** *Tree Technology*
11:10 - 11:15 **Karthee Sivalingam,** *Huawei*

Decentralised Intelligence

11:20 - 11:25 **Santiago Macho Gonzalez,** *Tree Technology*
11:25 - 11:30 **Andreas Harth,** *Friedrich-Alexander-University Erlangen-Nuremberg and Fraunhofer IIS Nuremberg*

Open Source

11:30 - 11:35 **Ricardo Vitorino,** *Ubiwhere*

Next Generation Edge Cloud Platform

Build & Manage an On-Demand Edge Cloud
Using Resources from Bare-Metal Providers

ONEedge.io

This work has received funding from the European Union's
Horizon 2020 research and innovation programme
under grant agreement ONEedge 880412

OpenNebula Systems

The company developing **OpenNebula** and leading **ONEedge.io**

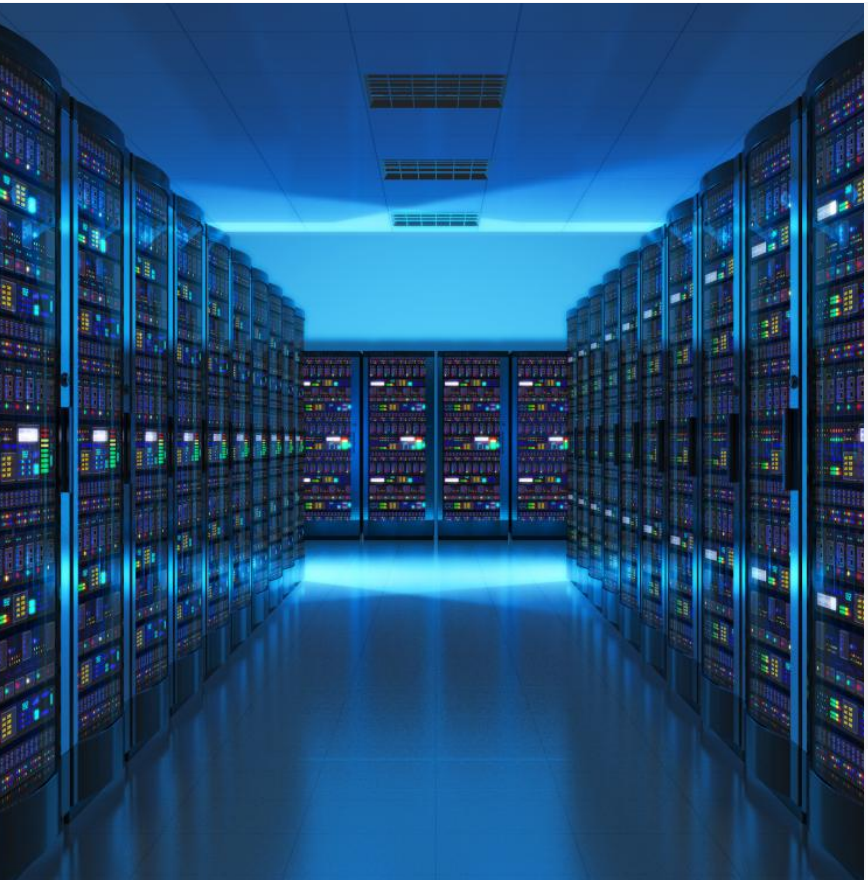

- European open source company.
- The only European open source IaaS solution, born in 2008.
- A success story emerged from EU innovation programs.
- Members of EOSC and GAIA-X.
- Offices in the US (Massachusetts) and Czech Republic.

LINUX FOUNDATION SILVER MEMBER

CLOUD NATIVE COMPUTING FOUNDATION SILVER MEMBER

LF EDGE GENERAL MEMBER

GAIA-X data-infrastructure.eu

Edge Computing Innovation

Towards a 🇪🇺 **meta-orchestrator** for the **Datacenter-Cloud-Edge-IoT** continuum

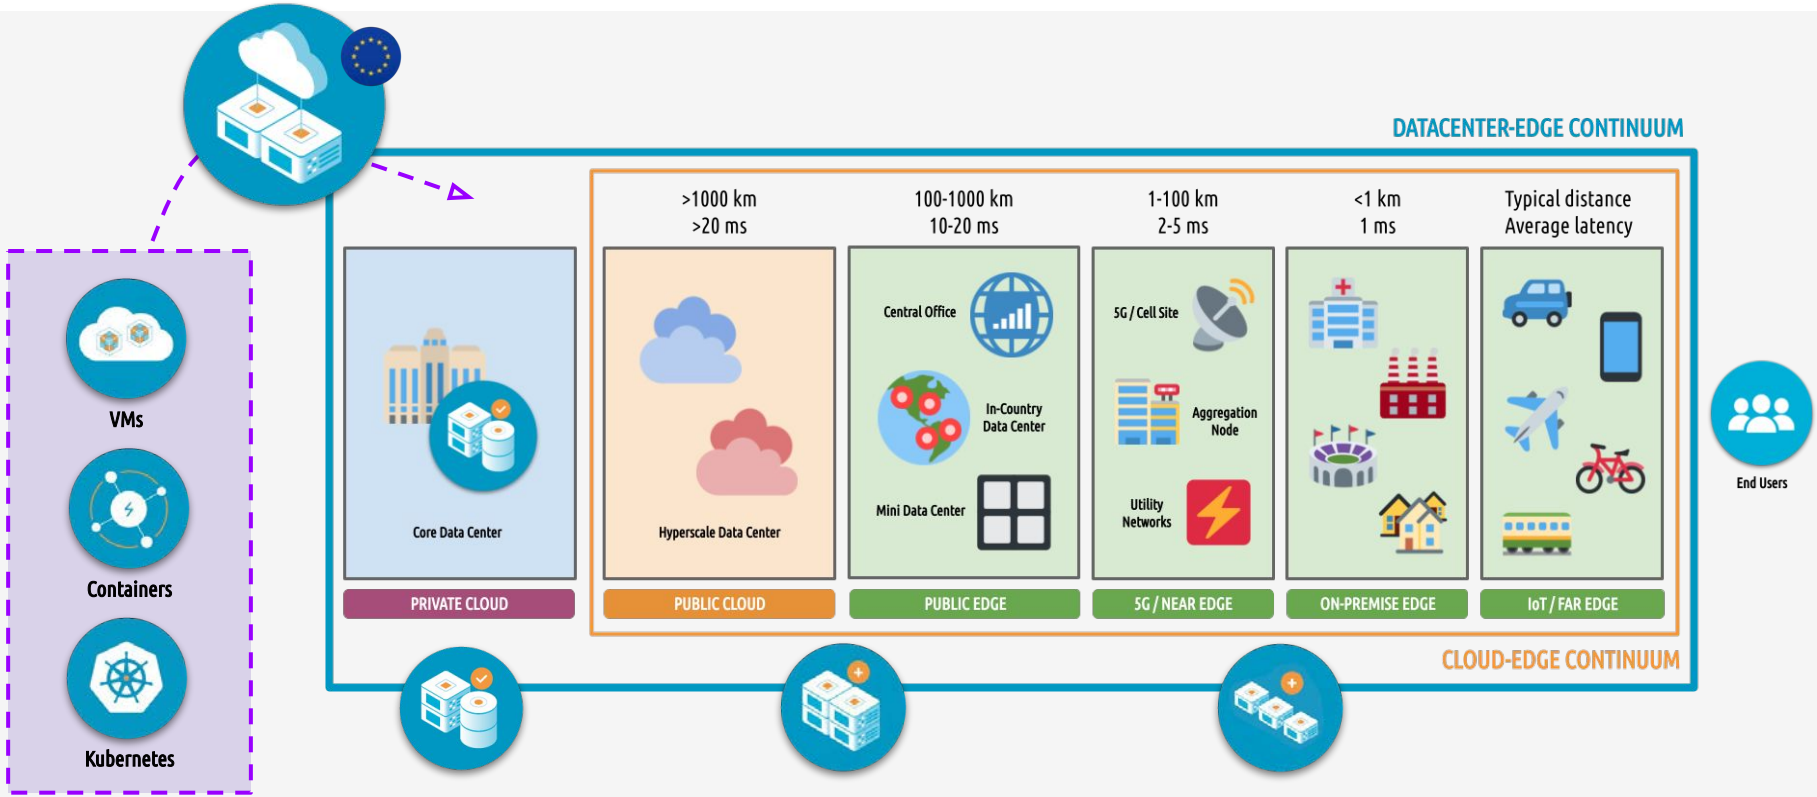

Our Vision on Edge Computing

🚀 Building an **open source, vendor-neutral** platform for Europe

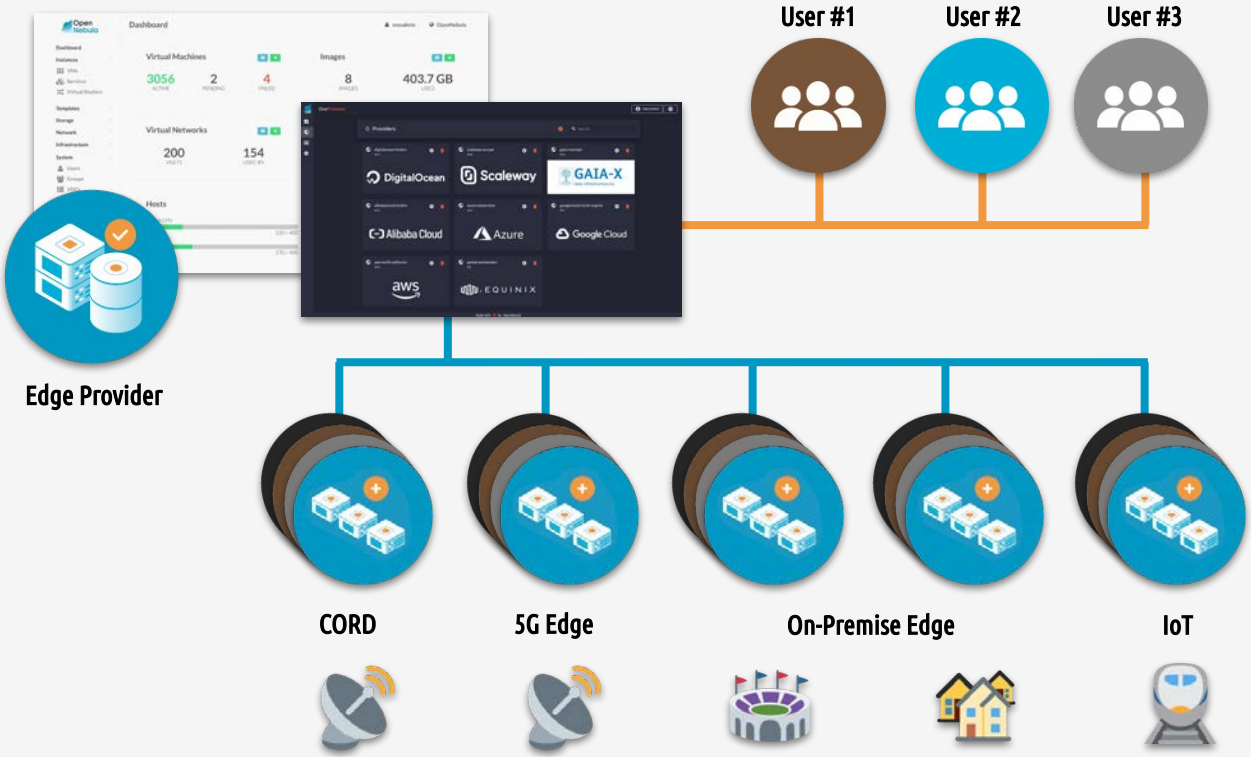

- ✓ Multi-Tenancy
- ✓ Self-Service Portal
- ✓ User Quotas
- ✓ Workload Isolation
- ✓ Multi-Tier Applications
- ✓ Elasticity Rules
- ✓ High Availability
- ✓ Federation
- ✓ Workload Portability
- ✓ Automatic Provisioning
- ✓ Public Marketplace of Virtual Appliances

Discover OpenNebula

Get in contact with us! 📂 research@opennebula.io

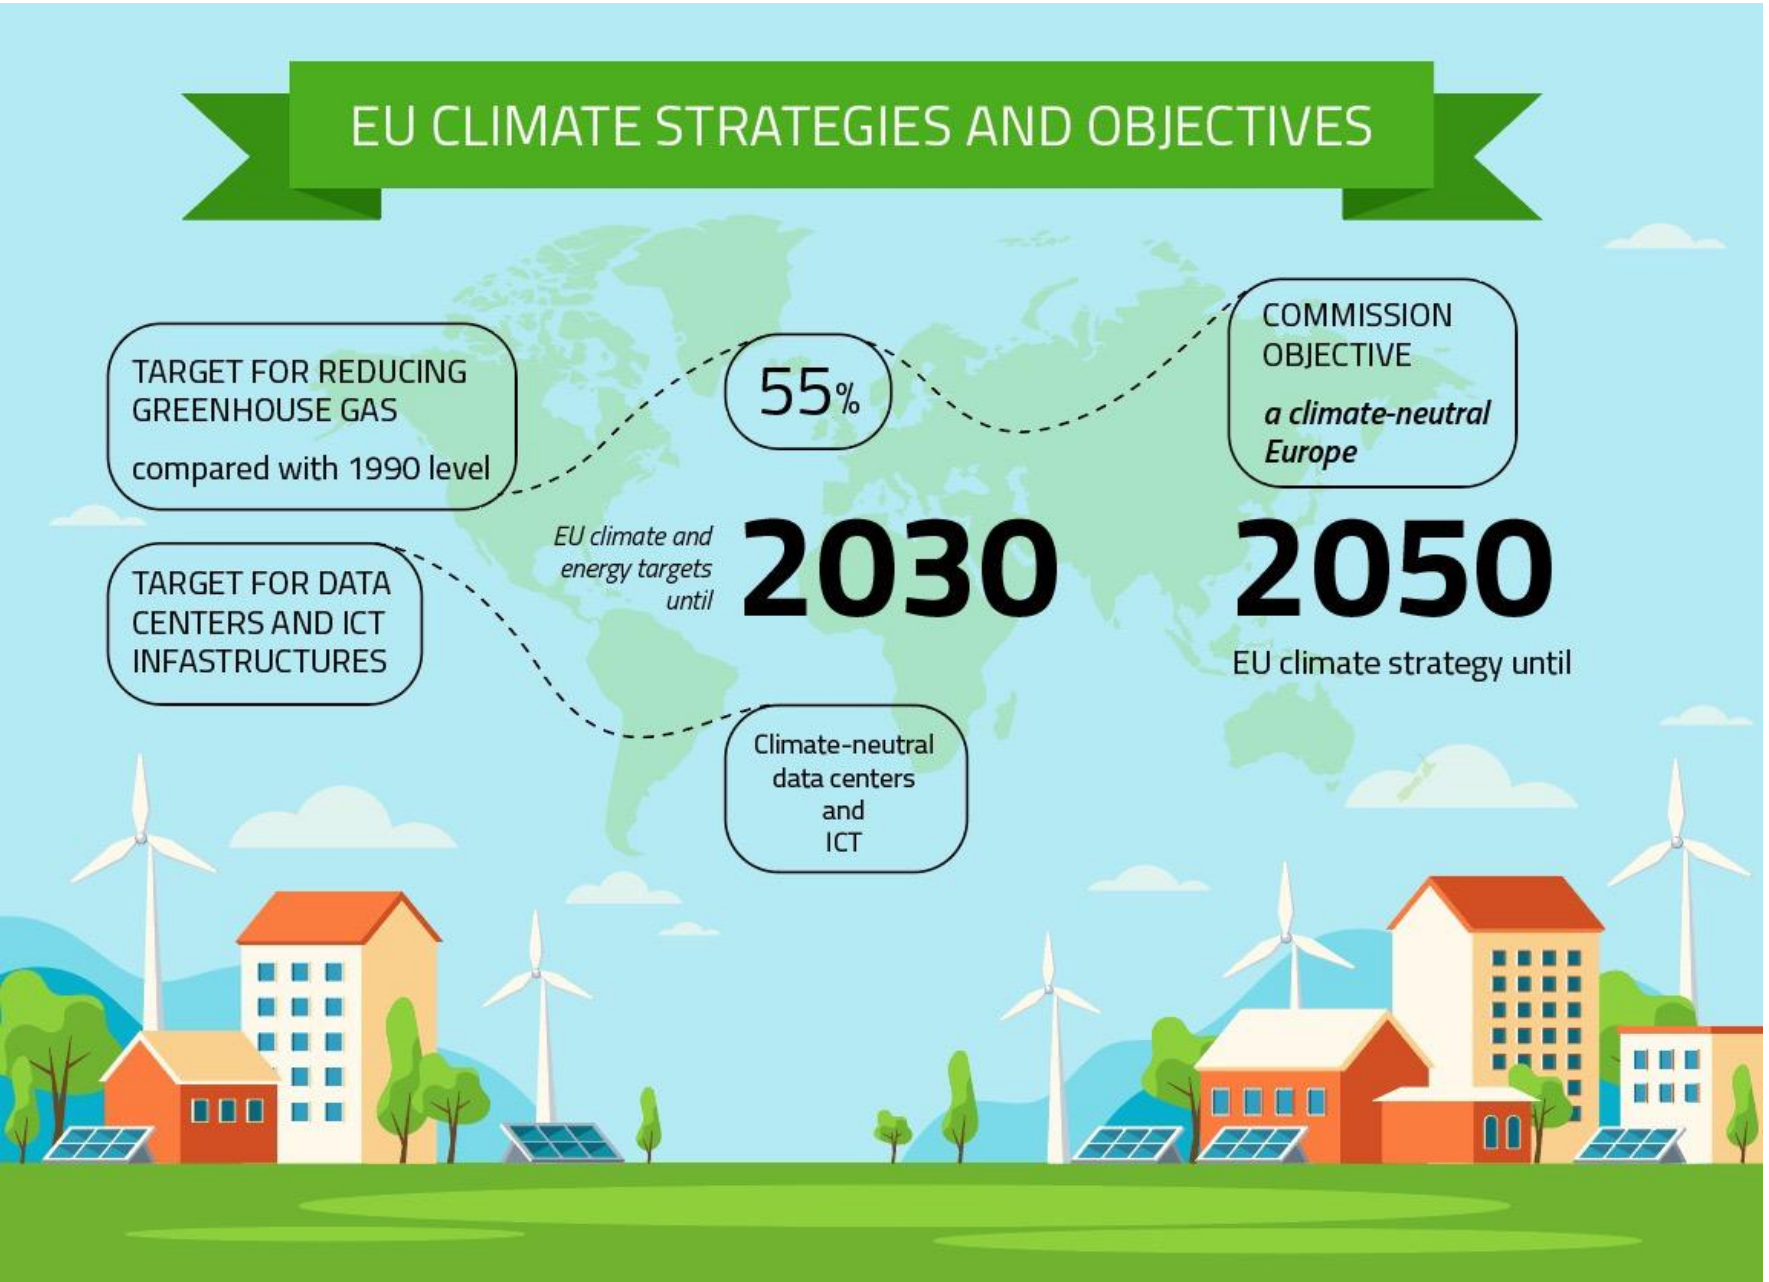
EU CLIMATE STRATEGIES AND OBJECTIVES
TARGET FOR REDUCING GREENHOUSE GAS
compared with 1990 level
55%
COMMISSION OBJECTIVE
a climate-neutral Europe
EU climate and energy targets until
2030
2050
EU climate strategy until
TARGET FOR DATA CENTERS AND ICT INFASTRUCTURES
Climate-neutral data centers and ICT


EUROPEAN
GREEN
DEAL


CODE OF
CONDUCT
DATA CENTRES
Endorser

Uptime Institute®

GREENCONTINUUM

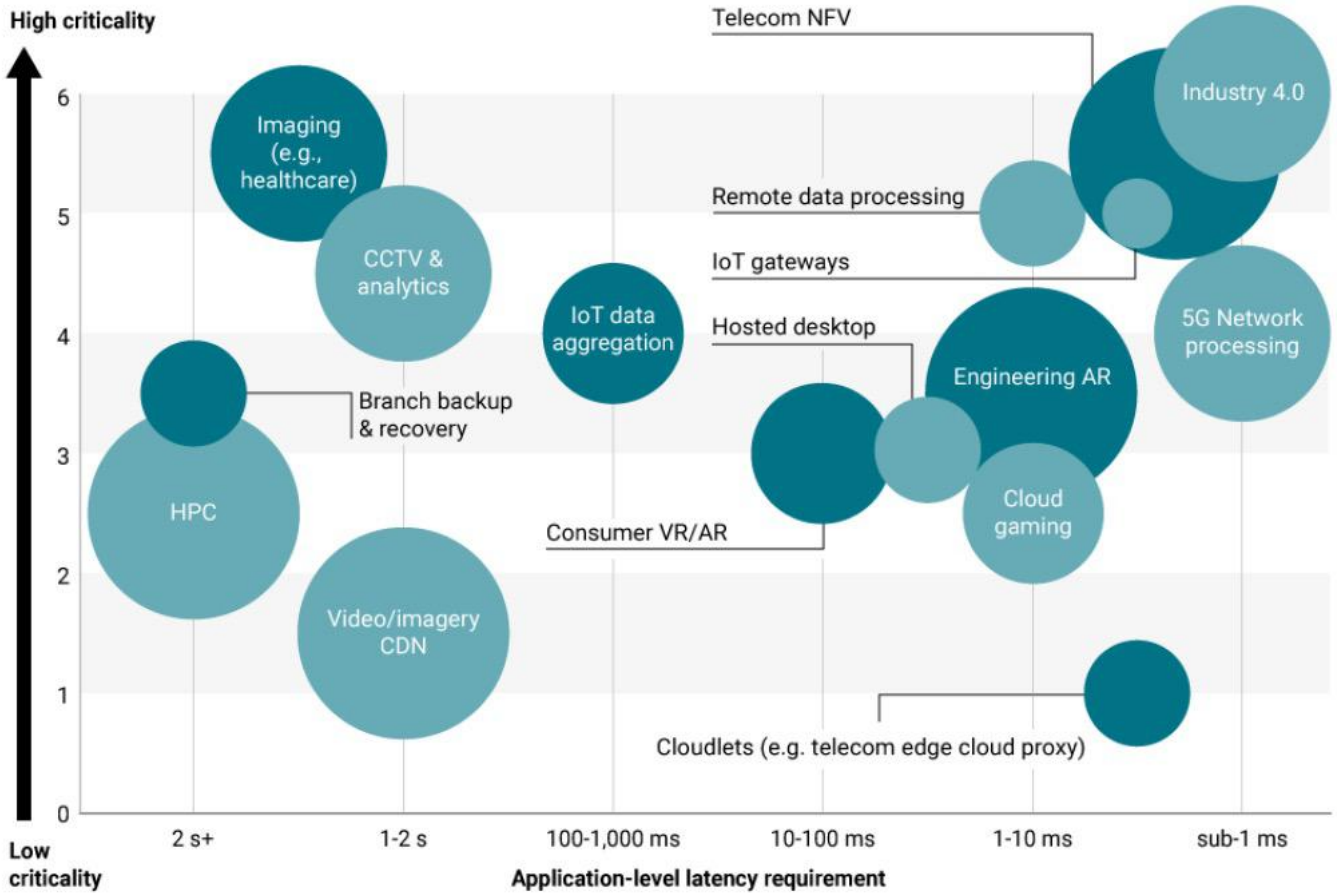

Score value	Volume / day / site	Criticality (availability of data or site)
6	100 TB+ (10 Gbps+)	0 downtime/extreme business risks
5	10-100 TB (1-10 Gbps)	few seconds/ high risks
4	1TB-10 TB (0.1-1Gbps)	few minutes / moderate risks
3	100-1,000 GB (10-100 Mbps)	few hours / low risks
2	10-100 GB (1-10 Mbps)	few days / little risk
1	10 GB (1 Mbps)	irrelevant / no risk

Key for volume/day/site

Uptime Institute Intelligence, 2021

Uptime Institute | INTELLIGENCE

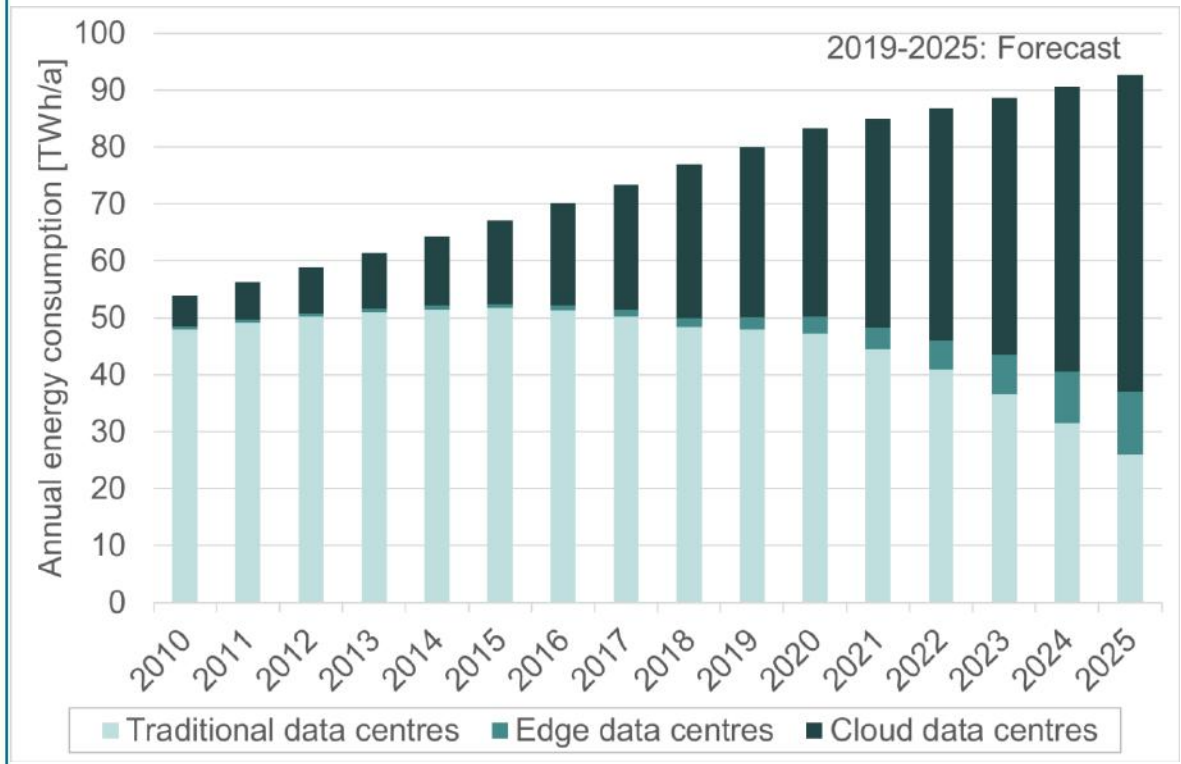

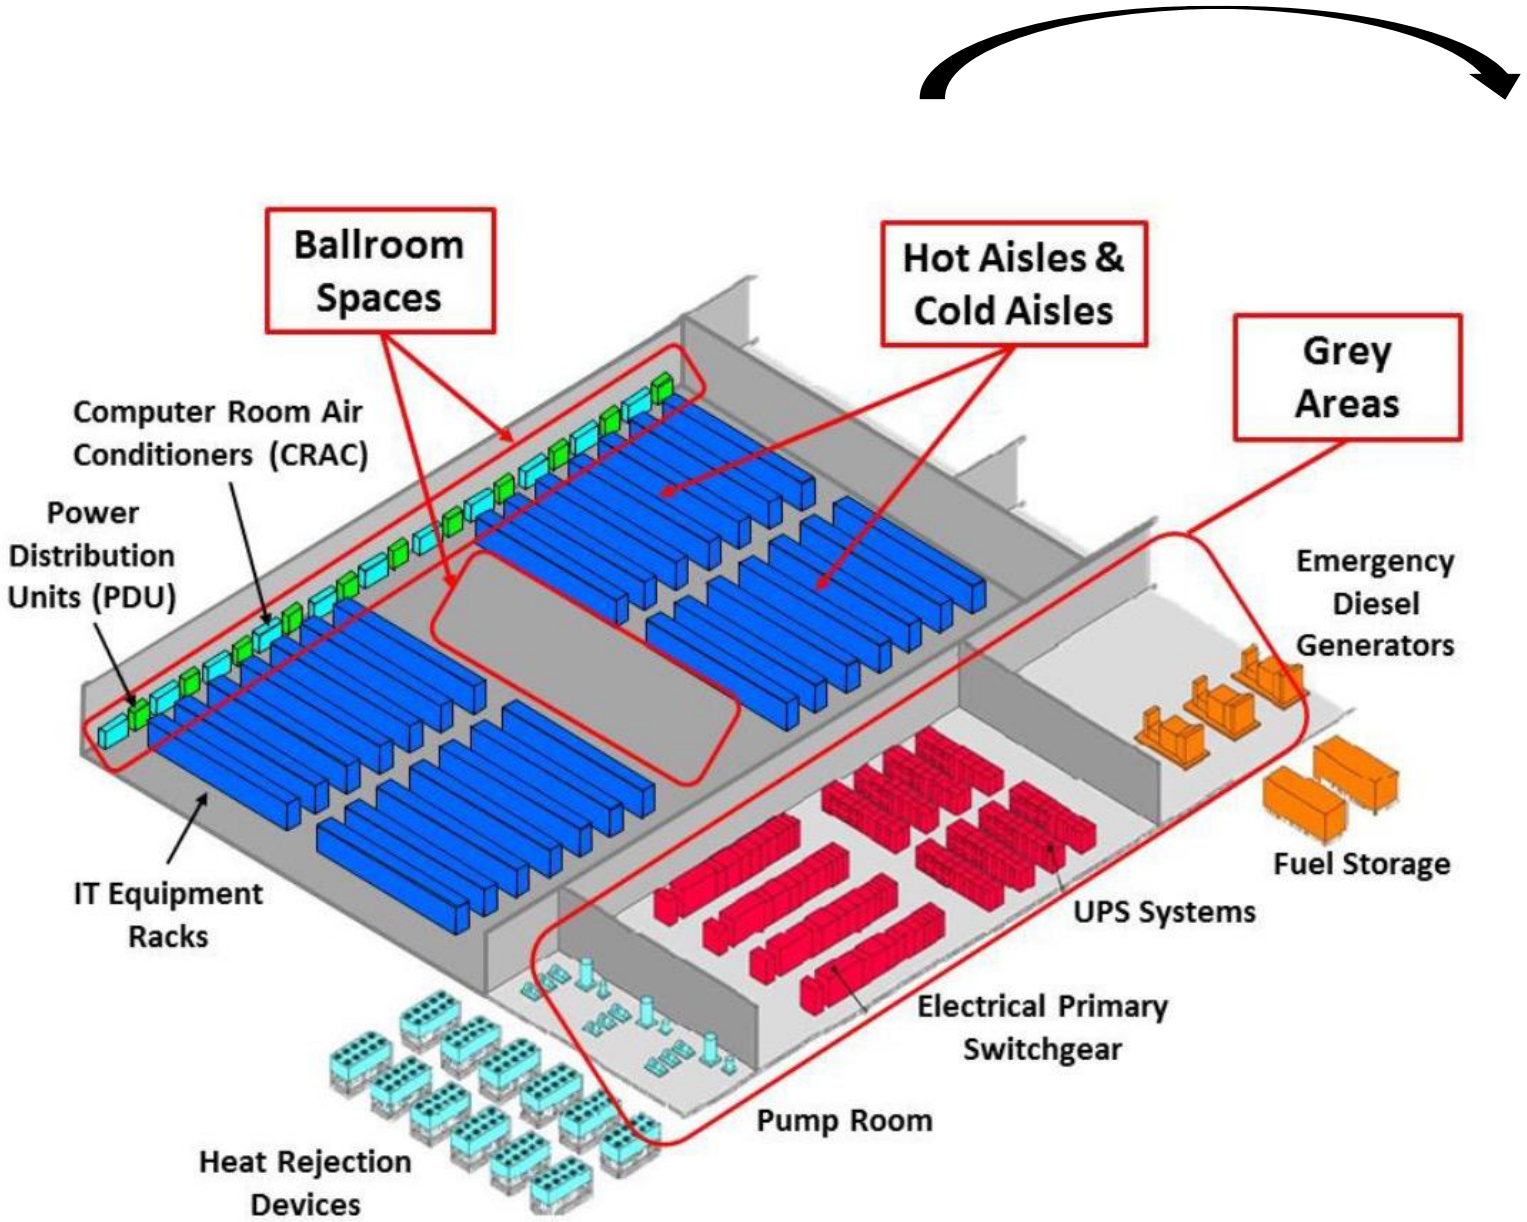
Ballroom Spaces
Hot Aisles & Cold Aisles
Grey Areas
Computer Room Air Conditioners (CRAC)
Power Distribution Units (PDU)
Emergency Diesel Generators
Fuel Storage
IT Equipment Racks
UPS Systems
Electrical Primary Switchgear
Pump Room
Heat Rejection Devices


ecoqube


A
B
C
D
E
F
G

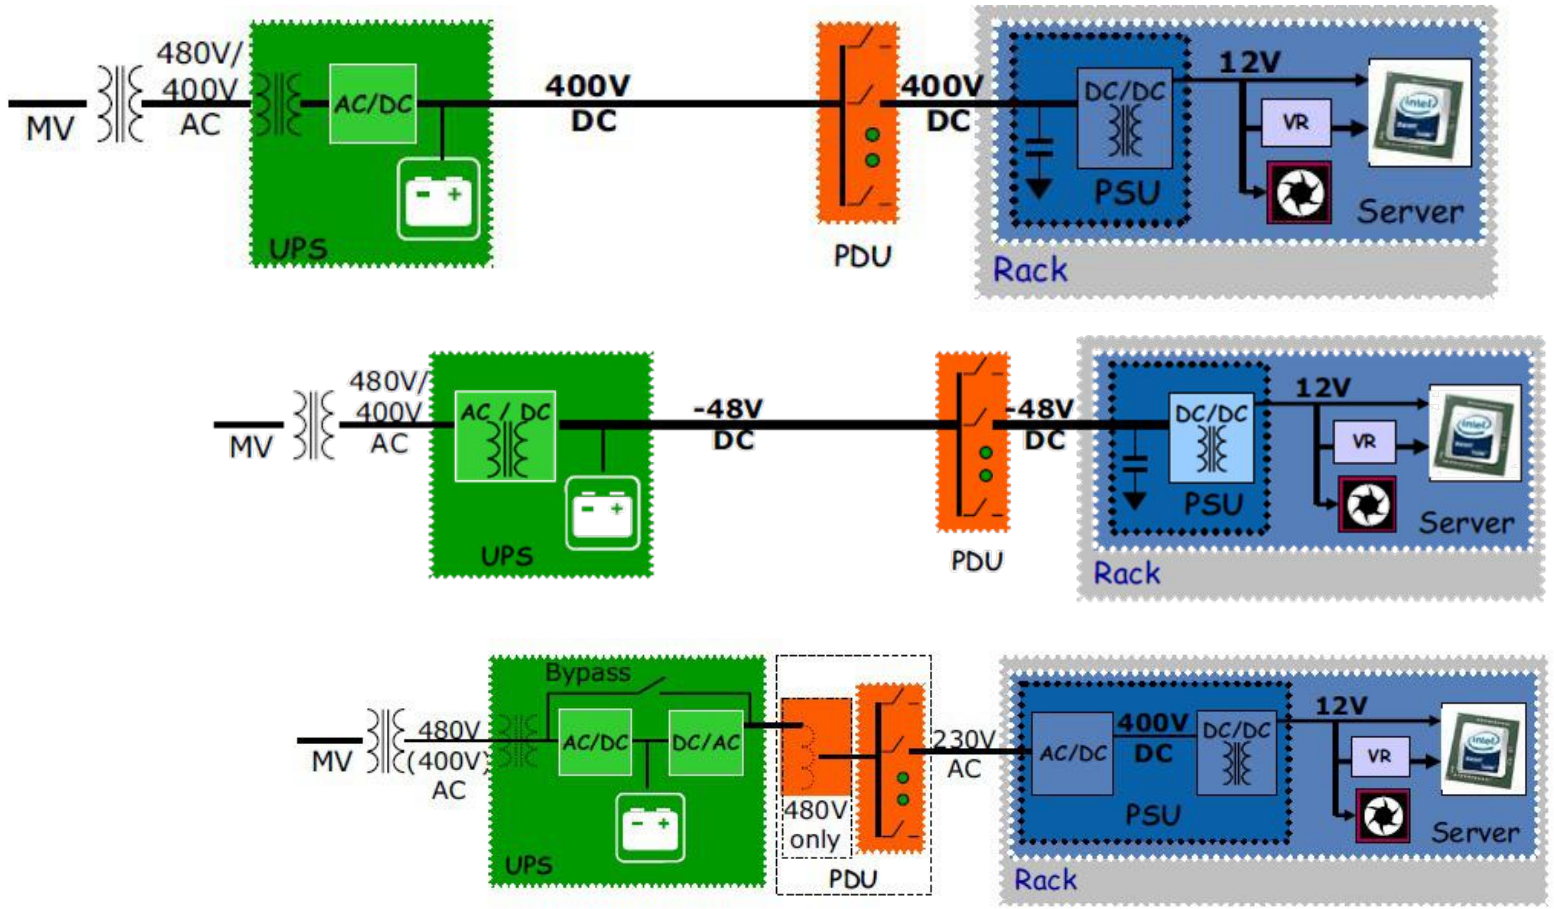
MV
480V/
400V
AC
AC/DC
UPS
400V
DC
PDU
400V
DC
DC/DC
PSU
12V
VR
Server
Rack
MV
480V/
400V
AC
AC / DC
UPS
-48V
DC
PDU
-48V
DC
DC/DC
PSU
12V
VR
Server
Rack
MV
480V
(400V)
AC
Bypass
AC/DC
DC/AC
UPS
480V
only
PDU
230V
AC
AC/DC
400V
DC
DC/DC
PSU
12V
VR
Server
Rack

Power DC

GREENCONTINUUM

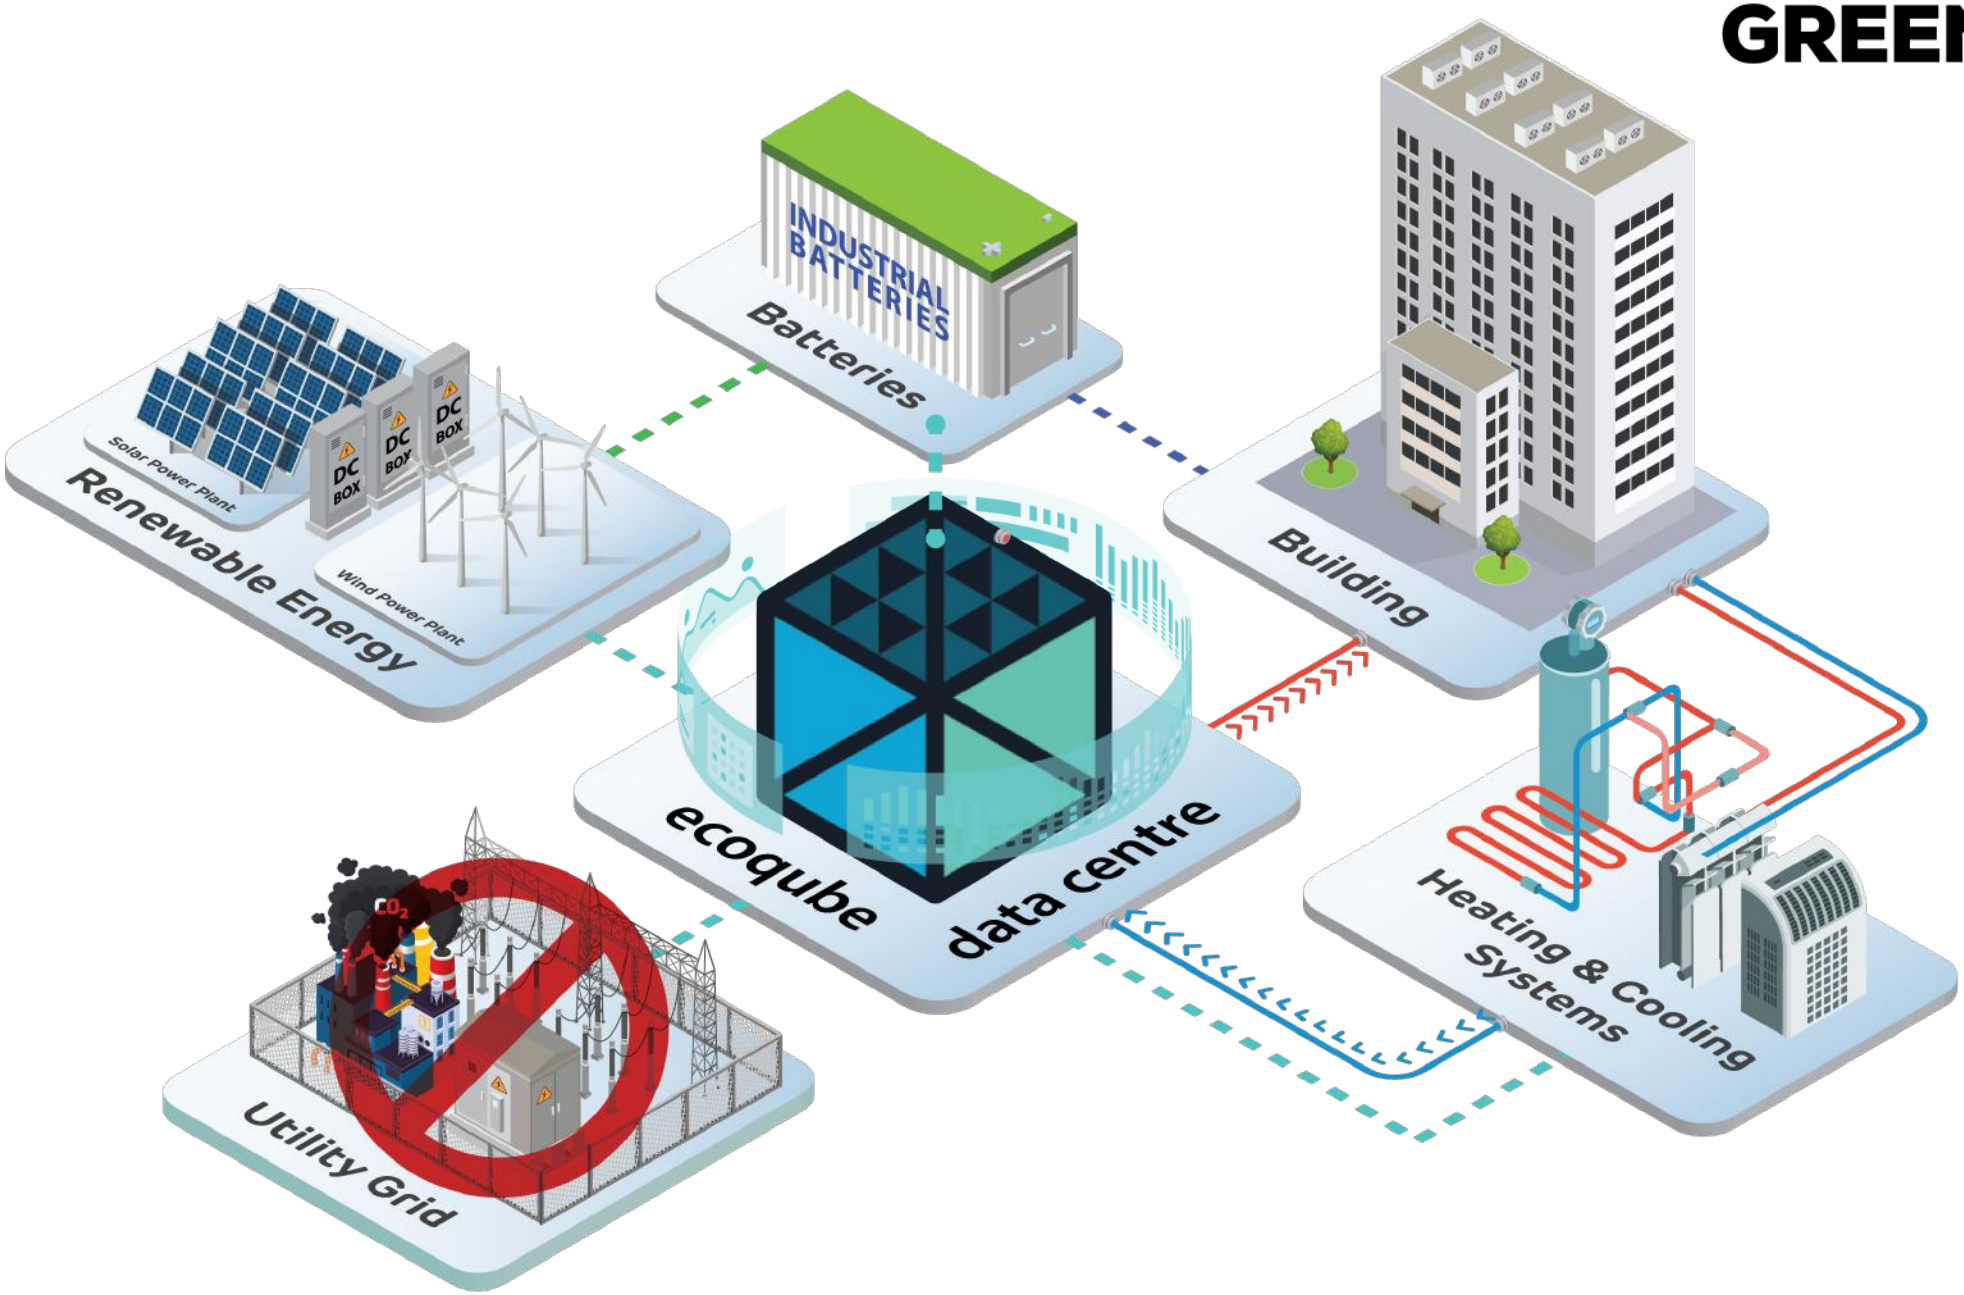

ecoqube

GREENCONTINUUM

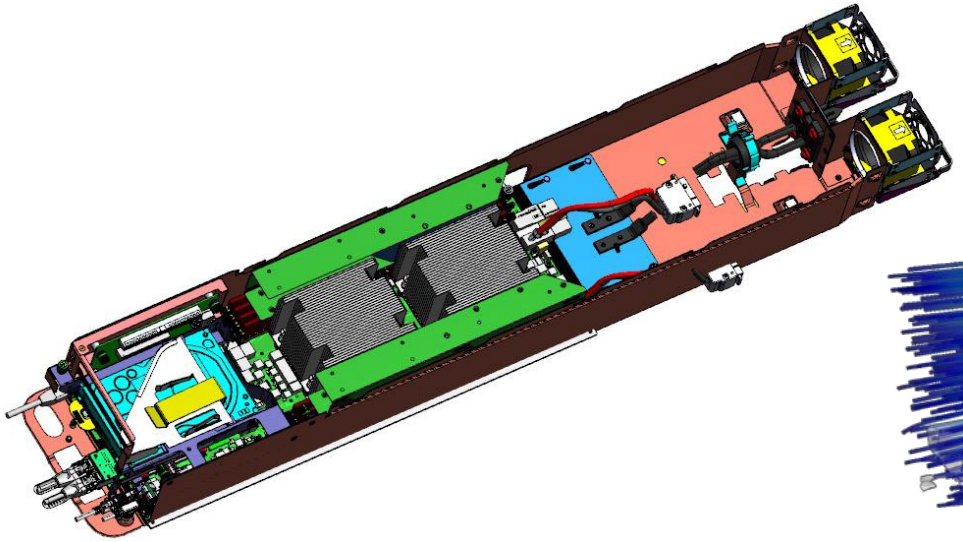

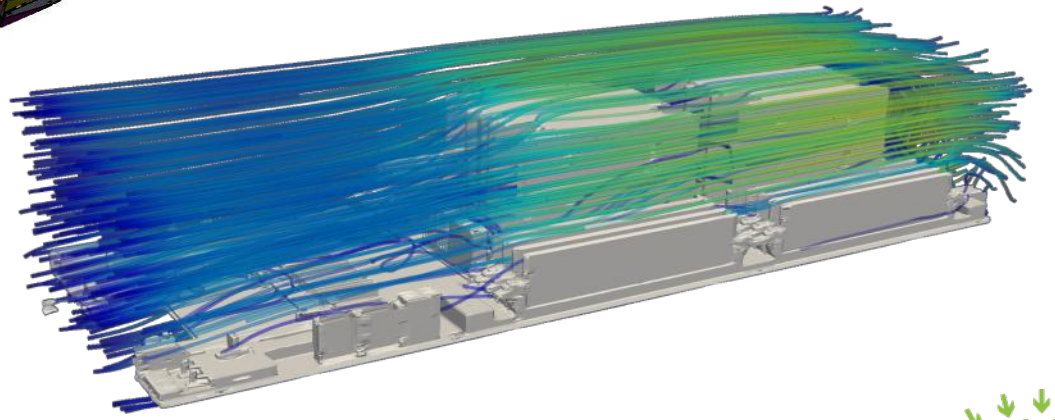

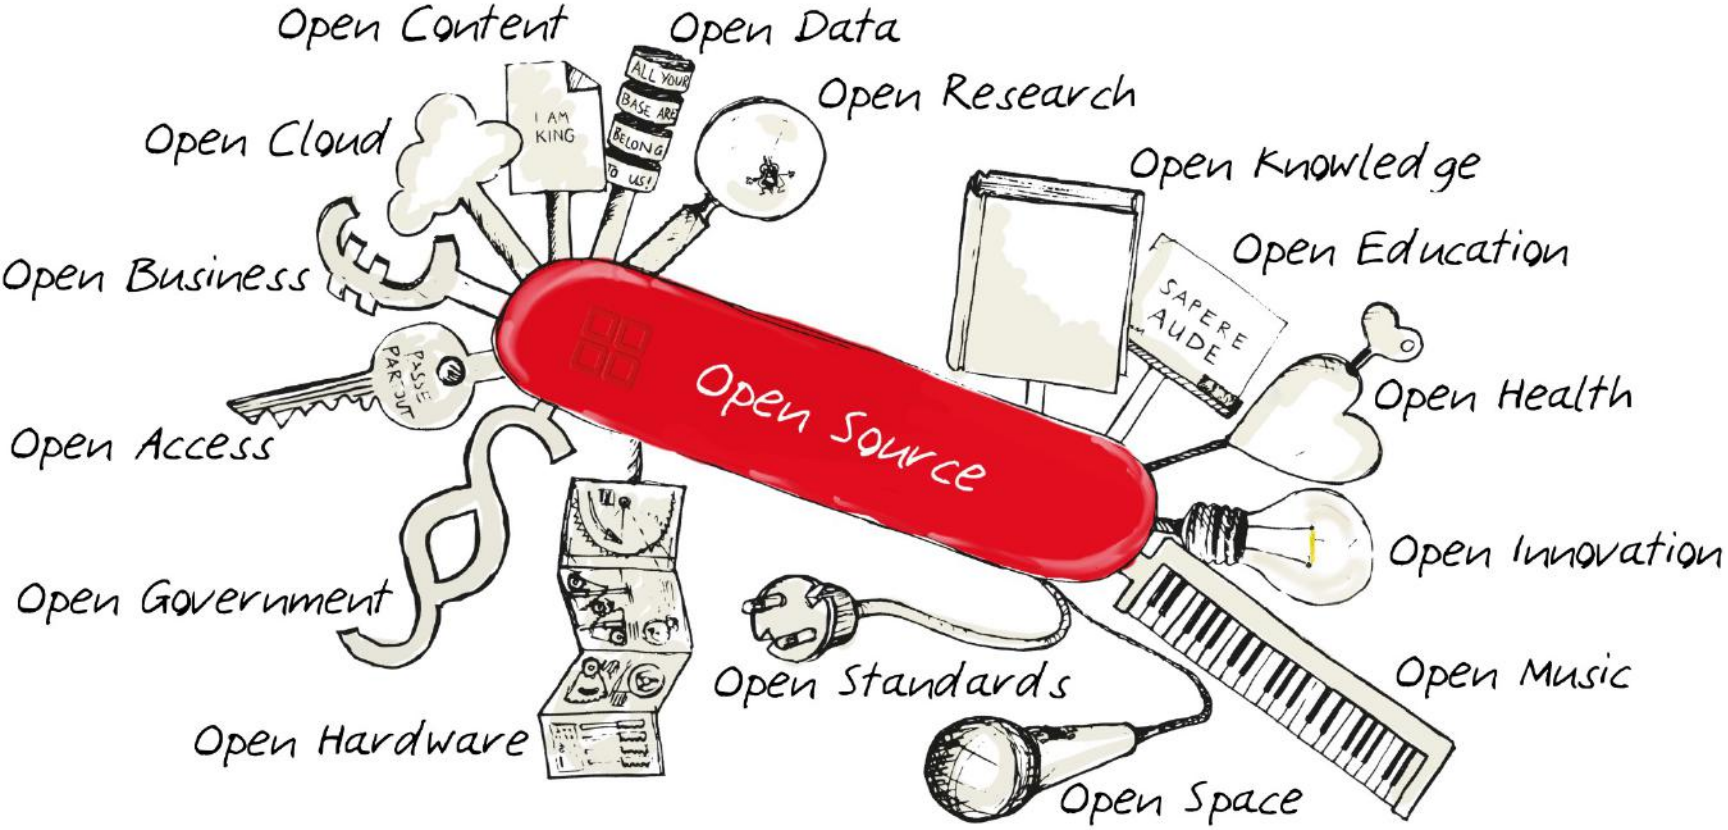

OPEN Compute Project®

OpenTR

Proposal by Bogazici University Team: Distributed PAN-European Cloud+Edge Platform

- **Aim:** Design an open interface cloud+edge framework to enable:
  - Flourishment of cloud & edge provider companies at different scales in all European countries
  - Competitive and efficient, yet constructive, business environment for cloud/edge/IoT/mobile services through data and service brokerage
  - Transnational service mobility across Europe
  - Europe-based data analytics and AI/ML services provided to businesses and services through standardized APIs
  - Transparency and accountability to ensure user anonymity and data protection provided by data analytics companies

- **Sample Use Cases:**

1) X-Company organizes a synchronized event in several e-sports arenas across Europe over edge servers of individual countries unified as a single e-sports challenge.

2) Y-Company wants to migrate to a different cloud server in another country for a better offer.

3) Z-Company offers courses and deploys exams over cloud servers in some countries and edge servers in others.

Stakeholders in the Proposed Architecture

- **End User (EU):** Subscriber of services (human, IoT device, etc.)
- **Application Provider (AP):** Company (at any scale) that provides applications/services to EUs
- **Edge Provider (EP):** Company (at any scale) that provides compute services to APs (via SBs) in a narrow zone (town or city)
- **Cloud Provider (CP):** Company (typically at a larger scale) that provides compute and storage services to APs (via SBs) in a region, country, or whole Europe
- **Data Analytics Provider (DAP):** Company that provides data analytics services to APs by pre-processing anonymized data from DBs (likely to be specific to a market such as mobile operators or online retailers)
- **Service Broker (SB):** Company or government institution that serves as a matchmaker for the requests of APs and available capacities of EPs/CPs
- **Data Broker (DB):** Company that collects data from various data sources, performs AI/ML-based algorithms, and provides anonymized analytics to DAPs
- **Data Source (DS):** Any company/institution that provides data to APs via DBs
- **Regulation Body (RB):** Body authorized by EU Commission or the local governments to oversee proper and lawful functioning of SBs & DBs

Conceptual System Architecture

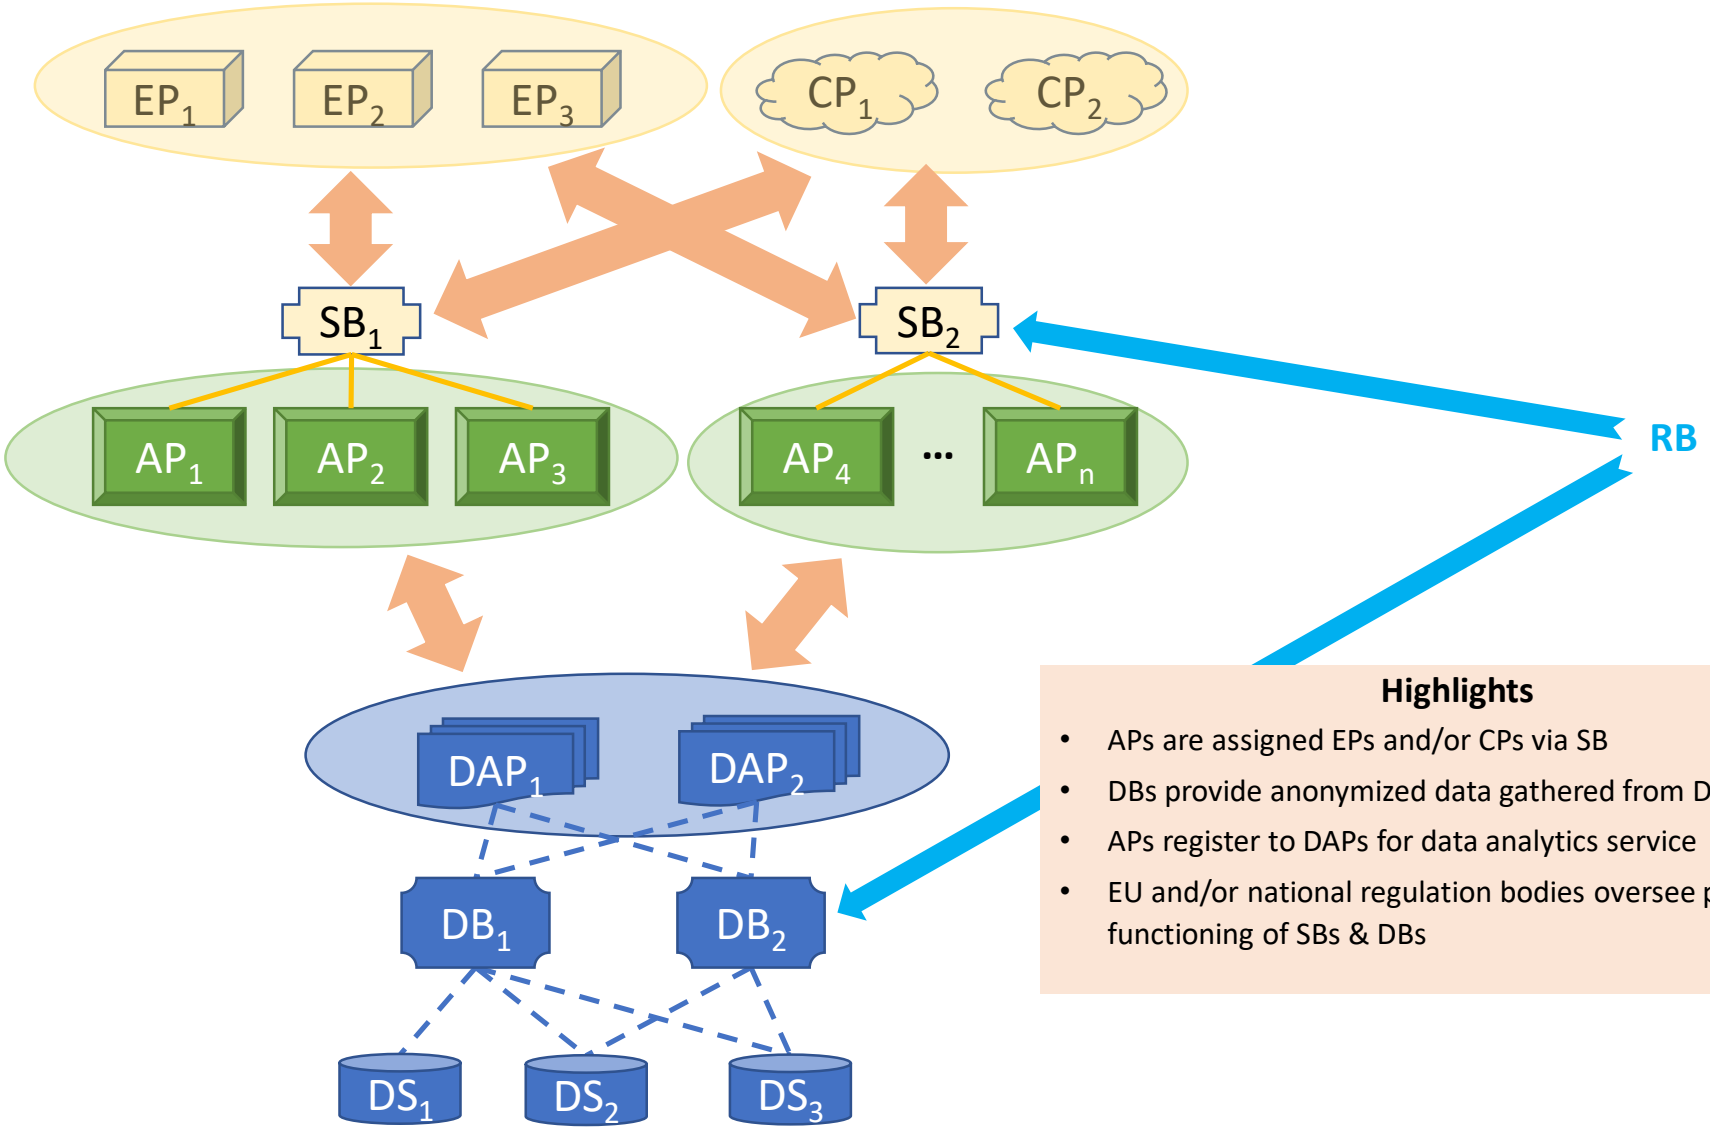

Highlights

- APs are assigned EPs and/or CPs via SB
- DBs provide anonymized data gathered from DSs to DAPs
- APs register to DAPs for data analytics service
- EU and/or national regulation bodies oversee proper functioning of SBs & DBs

** The links in the figure represent conceptual relations, not physical links*

Related Previous Work by Bogazici Team

- Team members: Atay Ozgovde, Cem Ersoy, Birkan Yilmaz, Ali Emre Pusane, Tuna Tugcu
  - E-mails: {ozgovde, ersoy, birkan.yilmaz, ali.pusane, tugcu}*@boun.edu.tr*
  - Bogazici University, Computer Networks Research Laboratory

- **Multiaccess Multi-tier Edge-Cloud [1,2]:**
  - Optimal joint task offloading and resource planning
  - Optimal multi-tier edge/RSU design for intelligent transportation
  - Optimal slicing for service oriented multi-tier edge computing

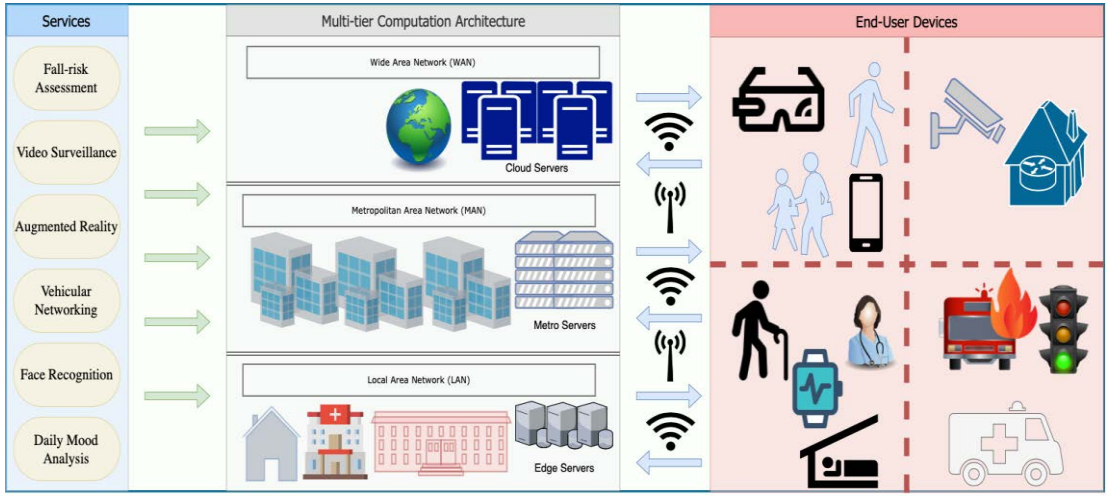

- **EdgeCloudSimSimulator [3]:**
  - Developer of popular EdgeCloudSim simulator for Edge Computing
  - Incorporates 3G/4G/5G
  - https://github.com/CagataySonmez/EdgeCloudSim

[1] C. Baktir, B. Ahat, A. Ozgovde, N. Aras, C. Ersoy, "SLA-Aware Optimal Resource Allocation for Service-Oriented Networks" 2019

[2] C. Sonmez, A. Ozgovde, C. Ersoy "Fuzzy Workload Orchestration for Edge Computing" 2019

[3] C. Sonmez, A. Ozgovde, C. Ersoy "EdgeCloudSim: An Environment for Performance Evaluation of Edge Computing Systems" 2018

Related Previous Work by Bogazici Team

- **C-RAN (CloudRAN) Testbed [4]:**
  - Virtual BBUs as VMs
  - Live migration
  - Low latency, scalability, low OPEX/CAPEX, interoperability, energy efficiency, flexibility

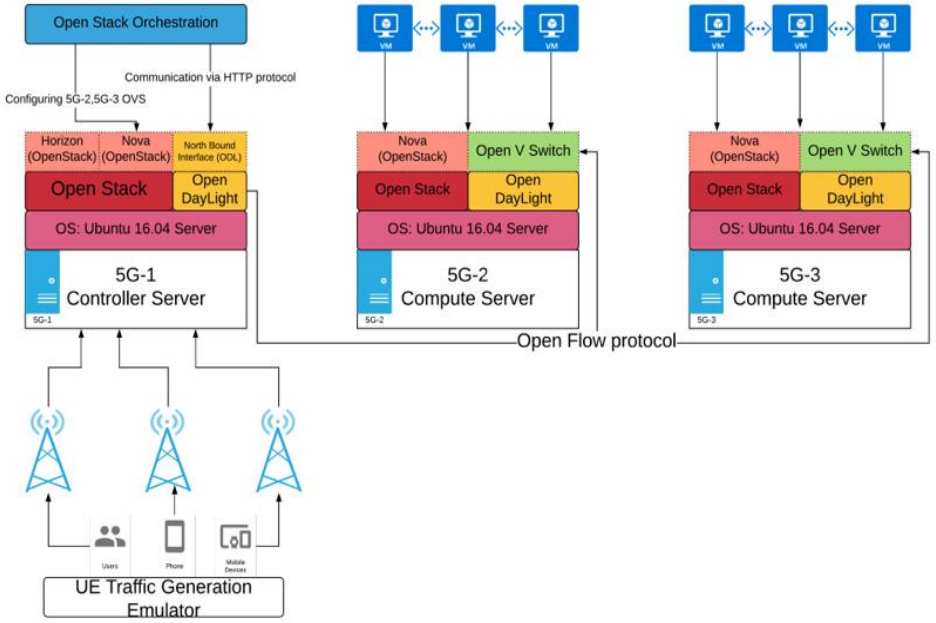

- **NWDAF w/ Machine Learning [5]:**
  - NWDAF gathers network data and provides data analytics for 5G NFs
  - Designed according to 3GPP specs
  - Several machine learning techniques implemented

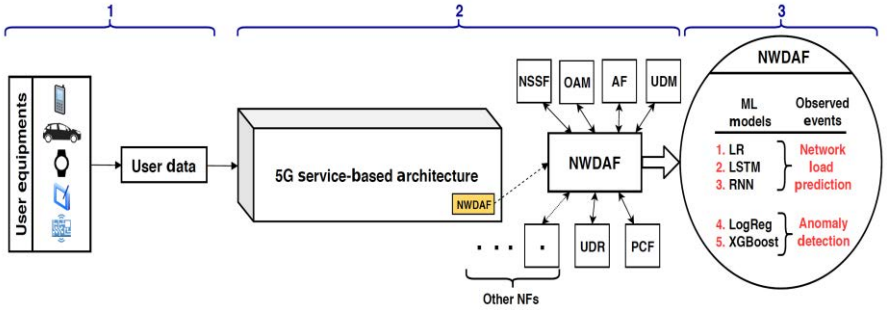

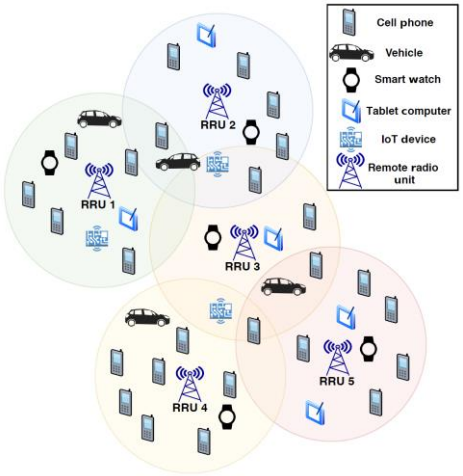

[4] A. Uzumcuoglu, K. Gökarslan, M. Turan, S. Sevgican, T. Tugcu "Coordinated Multipoint-aware Baseband Unit Placement in C-RAN" 2019

[5] S. Sevgican, M. Turan, K. Gökarslan, H. B. Yilmaz, T. Tugcu "Intelligent network data analytics function in 5G cellular networks using machine learning" 2020

TREE TECHNOLOGY

https://www.treetk.com/en

- Expertise in Big Data and Artificial Intelligence
- Spanish SME
- Large experience in EU projects (17 ongoing H2020 projects)
- Role:
  - Coordinator?
  - WP leader
- Contact:
  - Tatiana Silva (tatiana.silva@treetk.com)

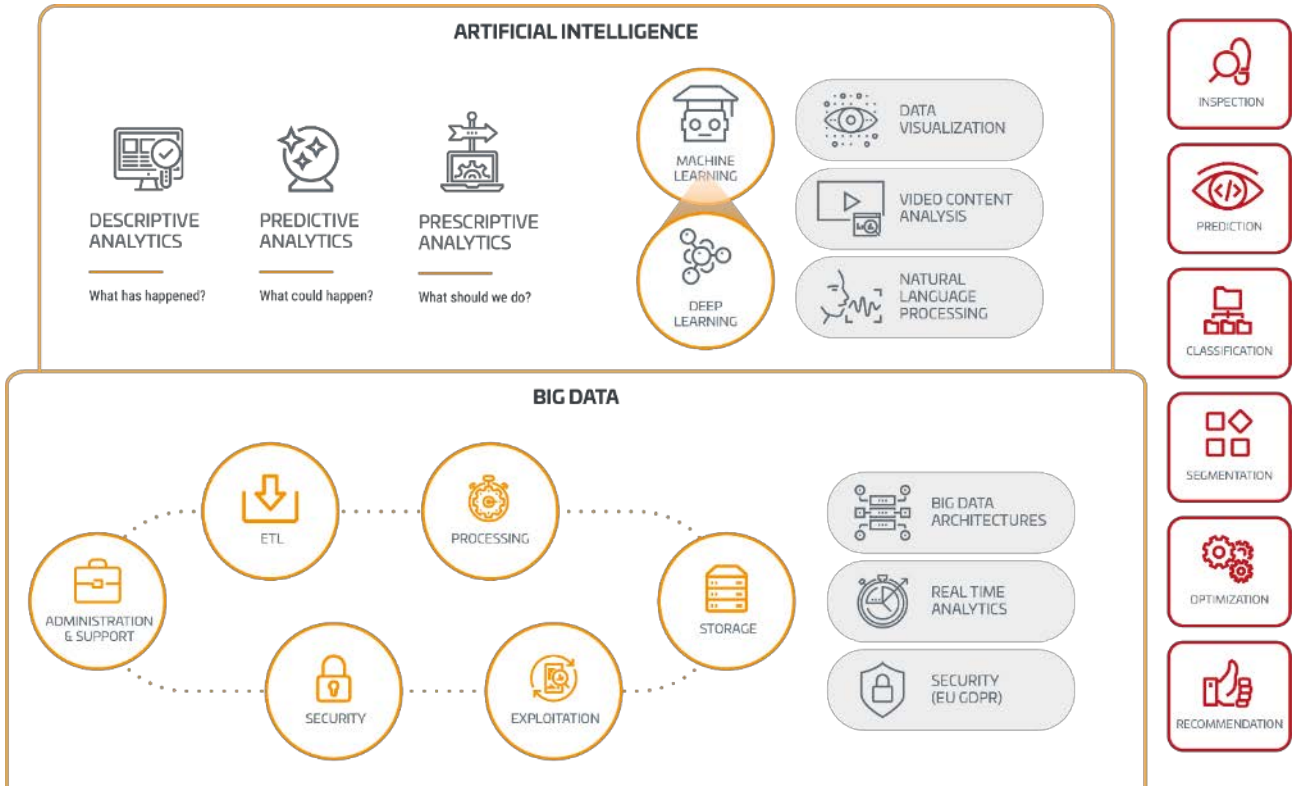

TOPIC: HORIZON-CL4-2022-DATA-01-02 “Cognitive Cloud: AI-enabled computing continuum from Cloud to Edge (RIA)”

Federated Learning

Federated Learning

- Different devices collaborative learn a shared predictive model.
- Data never leave the devices
- Decouples machine learning from the need to store the data in a single place.
- Brings the training to the devices

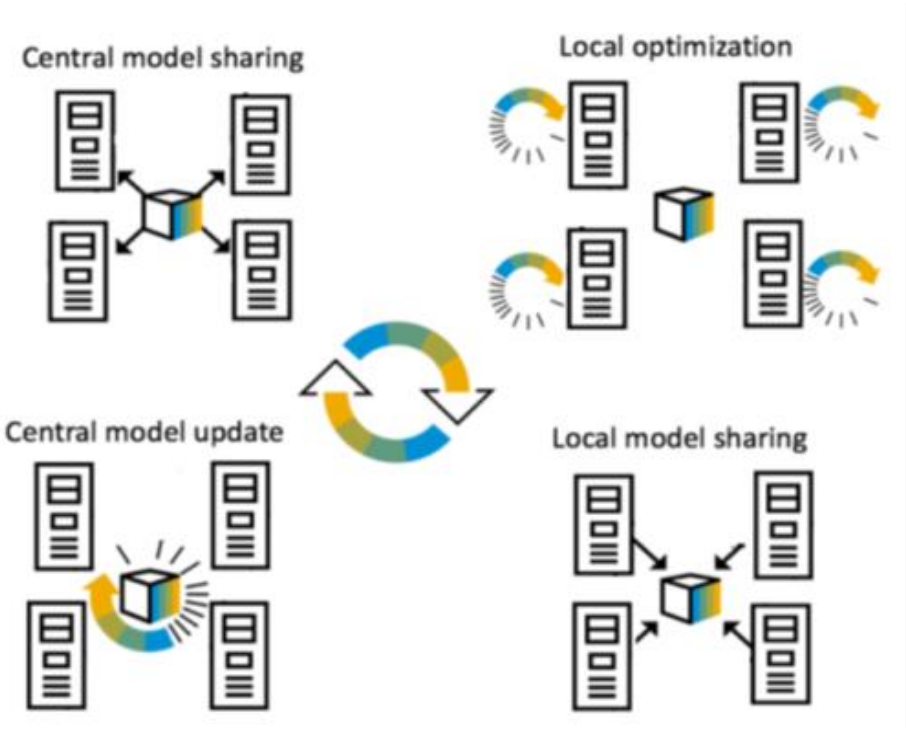

Federated Learning and Edge Computing

Problems in edge computing

- A hierarchical approach can be used.
- In every use case, the cloud, edge servers and IoT devices have completely different computational power, memory or communication capabilities.
- Many decisions may produce bottlenecks in a federated architecture and incurs in computational cost not affordable by the architecture.

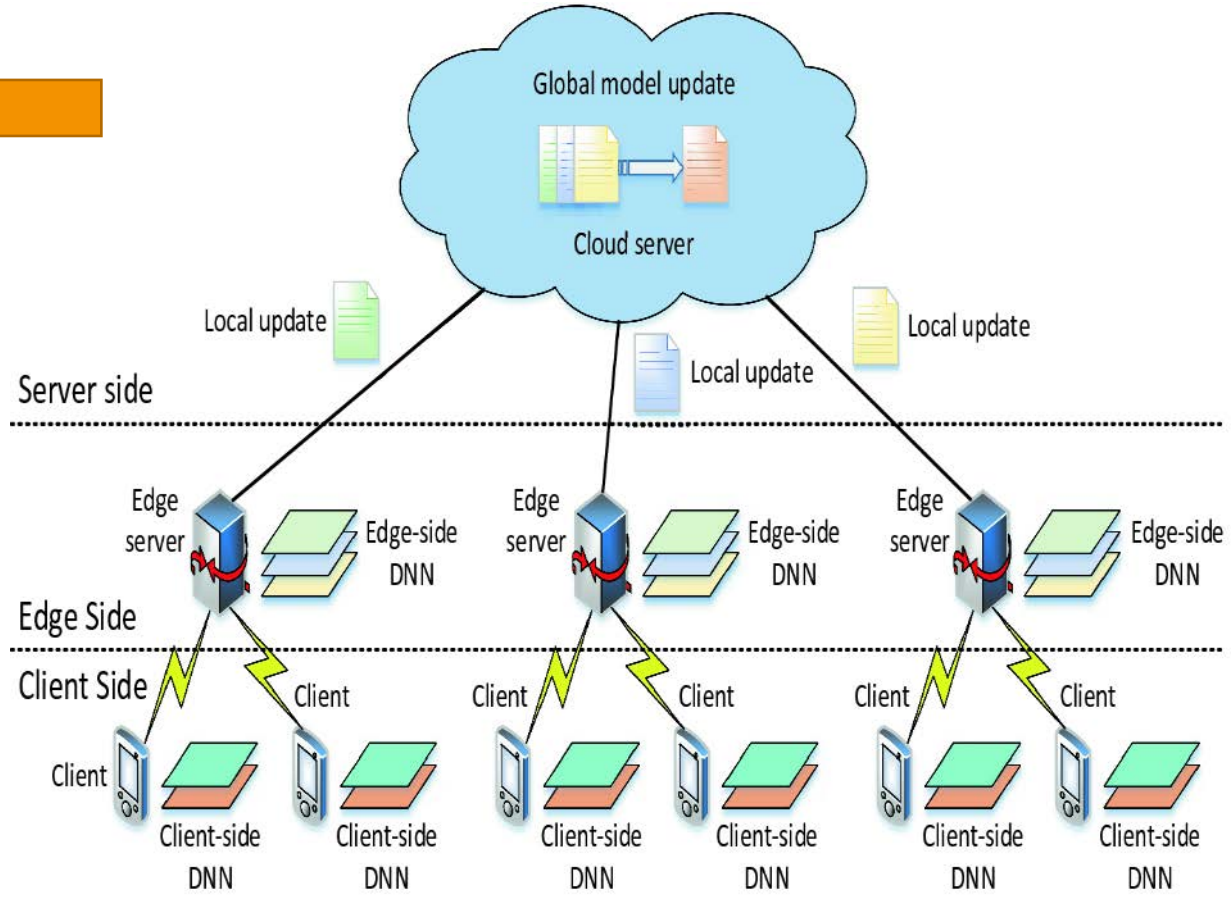

Federated Learning and Edge Computing

Improvements

- Having three different levels (devices, edge servers and cloud servers) we propose a framework to adapt intelligently where data must be allocated, the predictive model must be trained and the global model updated.
- This framework will respond and adapt intelligently to changes in application behavior such as DNN size or data workload by applying AI techniques to automatically adapt connectivity, computing and storage in cloud and edge to optimize where the ML training should be performed.
- Trustworthiness will be provided by including defenses to avoid data poisoning attacks.

Federated Learning and Edge Computing

Looking for

- Potential coordinator
- Experts in edge computing
- Case studies to test and validate the approach

Huawei, a leading global provider of ICT infrastructure and smart devices

Bring digital to every person, home and organization for a fully connected, intelligent world

197,000
Employees

105,000+
R&D employees

170+
Countries and regions

Rank 5th
In Global R&D Investment

49
on Fortune Global 500

2
TOP EPO Applicant

Huawei's research in Europe at a glance

370+ Research projects

230+ Academic institutions are partners

31 Participation in H2020 projects.

Research in Europe

1. **2400** employees, **84%** recruited locally
2. **27** R&D centres in EU, Russia, Ukraine, Belarus and Israel, with ERI HQ based in Leuven/BE
3. Munich Research Center is the largest European Huawei Research Center, established in 2008.

Research Activities (a few examples)

- 5G & Wireless
- Optical Communications
- Data Communications
- Photonics
- Cyber security & Privacy
- Quantum Technologies
- Semiconductors
- Thermal Technologies
- AI/ML/CV
- RF IC
- NFV/SDN
- Cloud/Big Data
- Aesthetics
- User experience

Focus on innovation, continuous dedication and backward compatibility

52 mm x 38 mm x 10 mm

Atlas 200

AI Accelerator Module

- 16 TOPS of INT8
- 16-channel HD video real-time analytics, JPEG decoding
- 4 GB/8 GB memory, PCIe 3.0 x4 interface
- Operating temperature: -25°C to +80°C

Atlas 300

AI Accelerator Card

- 64 TOPS of INT8 @ 67 W
- 32 GB memory
- 64-channel HD video real-time analytics
- Standard half-height half-length PCIe card form factor, applicable to general-purpose servers

reddot award 2019 winner

Atlas 500

AI Edge Stations

- 16 TOPS of INT8
- 25–40 W
- Wi-Fi & LTE
- 16-channel HD video real-time analytics
- Fanless design, -40°C to +70°C environments

reddot award 2019 winner

Atlas 800

Deep Learning System

- Plug-and-play installation
- Ultimate Performance
- Integrated Management

Ascend 310

AI SoC with ultimate efficiency

- Half precision (FP16): 8 TFLOPS
- Integer precision (INT8): 16 TOPS
- 16-channel full-HD video decoder: H.264/265
- 1-channel full-HD video encoder: H.264/265
- Max. power consumption: 8 W
- 12nm

Ascend 910

Highest compute density on a single chip

- Half precision (FP16): 256 TFLOPS
- Integer precision (INT8): 512 TOPS
- 128-channel full-HD video decoder: H.264/265
- Max. power consumption: 350 W
- 7nm

Atlas 900 AI Cluster

The pinnacle of computing power

- Thousands of Ascend 910 AI processors
- High-speed interconnection
- Delivers up to 256 to 1024 PetaFLOPS at FP16
- Can complete model training based on ResNet-50 within 59.8 seconds
- 15% faster than the second-ranking product
- Faster AI model training with images and speech

Atlas 200 DK

Quickly build development environments in 30 minutes

- 16 TOPS of INT8 @ 24 W
- 1 USB type-C, 2 camera interfaces, 1 GE port, 1 SD card slot
- 4 GB/8 GB memory

Kunpeng 920

The industry's highest-performance ARM-based server CPU

- ARM v8.2-architecture
- up to 64 cores, 2.6 GHz
- 8 DDR4 memory channels
- PCIe 4.0 and CCIX
- Integrated 100GE LOM and encryption and compression engines
- Supports 2- or 4-socket interconnects

Storage-intensive ——— **Computing-intensive**

5290	5280	2280	1280	2480	X6000
4U 72-drive storage model	4U 40-drive storage model	2U 2S balanced model	1U 2S high-density model	2U 4S high-performance model	2U 4-node high-density model






Huawei pilot cluster

Hosted at Goethe University of Frankfurt (FIAS)
Supported by Open Edge and HPC Initiative (OEHI)

- Cluster Configuration

28	Standard Compute Node - TaiShan 200 (Model 2280)
1	Development Compute Node - TaiShan 200 (Model 2280)
2	IO Node - TaiShan 200 (Model 5280)
1	AI Training Node - Atlas 800 (Model 9000)
1	AI Interference Node - Atlas 800 (Model 3000)

Connected by 6 Mellanox Infiniband – Each switch provides 36 x 100 Gb/s InfiniBand EDR ports.

Future addition

- *Edge intelligence*
- *OceanStor Pacific storage*
- *Cloud capabilities*

Taishan 200 server

- support of two Kunpeng 920 processors as well as DDR4-DIMM-Slots
- 2 x Kunpeng 920 processor (ARMv8 64bit):
  - 64 core; 2.60 GHz base frequency; 180 Watt TDP;
  - 2933 MHz max. Memory speed incl. heatsinks
- 128 GB main memory

IO nodes

2 x 960 GB SSD local storage + 36 x 1.2 TB HDD data storage

AI Training Node

8 x Huawei Ascend 910 with 32 AI cores
1024 GB main memory + 2 x 960 GB SSD + 4 x 3,2 TB SSD NVMe

AI Interference Node

5 x Atlas 300 AI Inference Card; 32 GB; PCIe3.0 x16

Huawei cluster pilot to streamline HPC software stack

European HPC-AI applications

SPECFEM 3D · NEMO · GROMACS · WRF · PyTorch · MindSpore · QUANTUMESPRESSO · CP2K · GPAW! · OpenFOAM · TensorFlow · ONNX

Cluster Management / Scheduler

slurm workload manager · HUAWEI · HUAWEI CLOUD · TORQUE

Bisheng Compiler · Open MPI · MVAPICH · GCC · LLVM Compiler Infrastructure · arm COMPILER

Middleware / Libraries

OpenBLAS · SYCL · PETSc · openGauss · Kunpeng DevKit · CANN 3.0 · docker

Filesystem & OS

OpenEuler · CentOS · ubuntu

BeeGFS · lustre · ceph · OceanStor FS

HW Infrastructure

Device	Edge	Server	Storage	Cloud
 Smartphones · Atlas 200 · Smart devices	 Atlas 500 · Atlas 300I · Atlas 500 Pro	 Atlas 800 inference server · Taishan server · Atlas 800 training server	 OceanStor Pacific	HUAWEI

TREE TECHNOLOGY

- Expertise in Big Data and Artificial Intelligence
- Spanish SME
- Large experience in EU projects (17 ongoing H2020 projects)
- Role:
  - WP leader
- Contact:
  - Santiago Macho (santiago.macho@treetk.com)

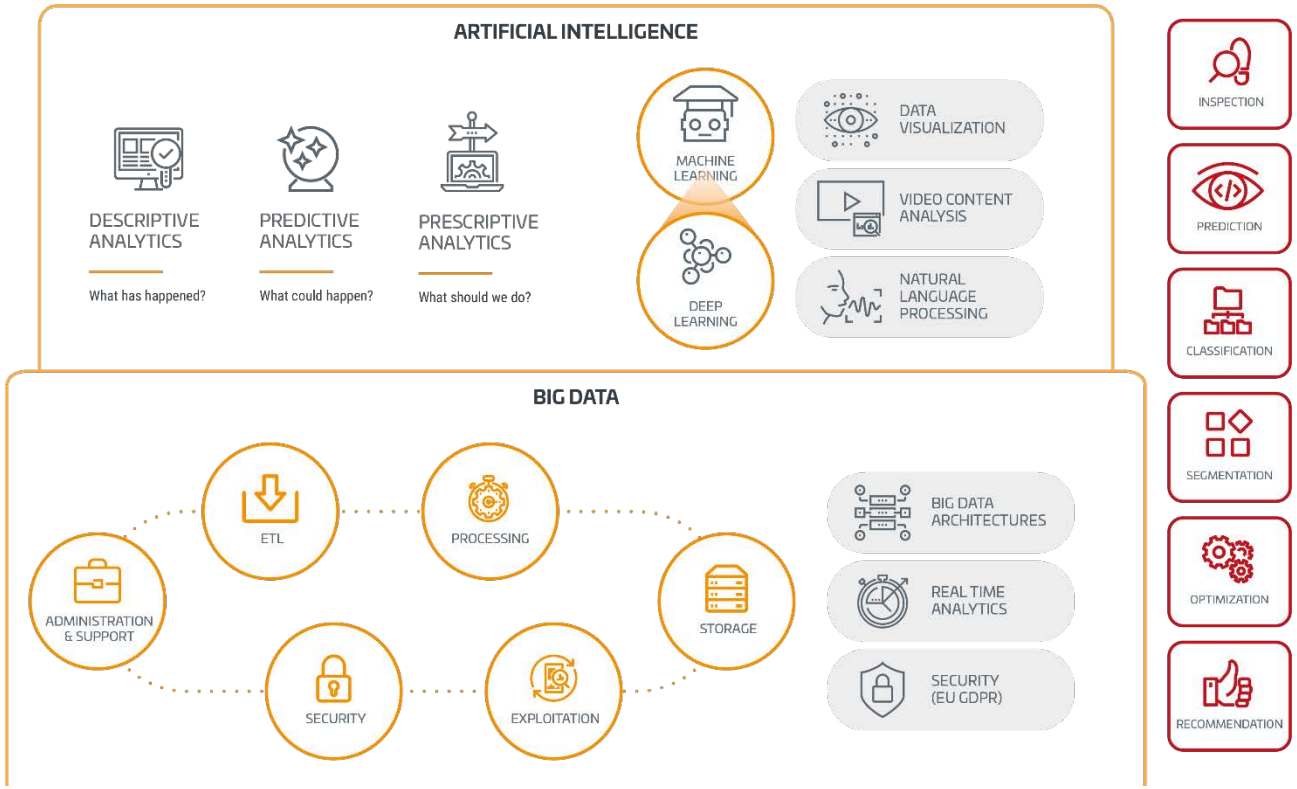

TOPIC: HORIZON-CL4-2022-DATA-01-03 “Programming tools for decentralised intelligence and swarms (RIA)”

Capabilities: Federated Learning

Federated Learning

- Different devices collaborative learn a shared predictive model.
- Data never leave the devices
- Decouples machine learning from the need to store the data in a single place.
- Brings the training to the devices

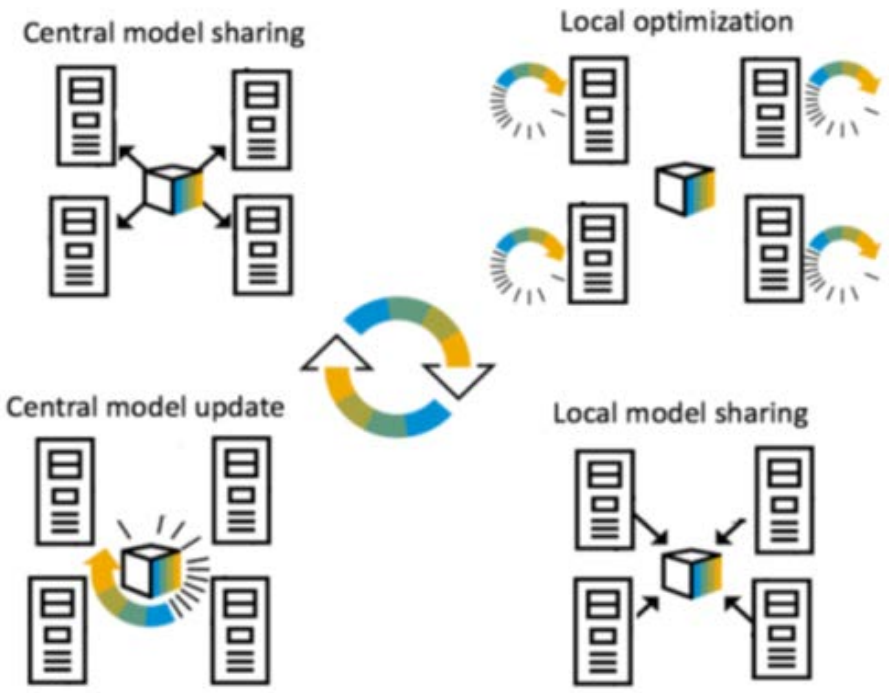

Challenges: Decentralized Federated Learning

Challenges

- In decentralized federated learning, the nodes should be able to coordinate themselves to obtain the global model.
- The specific network and device topology highly affect the performance of the learning process.
- The design and development of this kind of training architectures for real environments is extremely hard.
- In addition, the security threads increase due the lack of a central node to control the training.
- It is one of the open problems in the federated learning field.

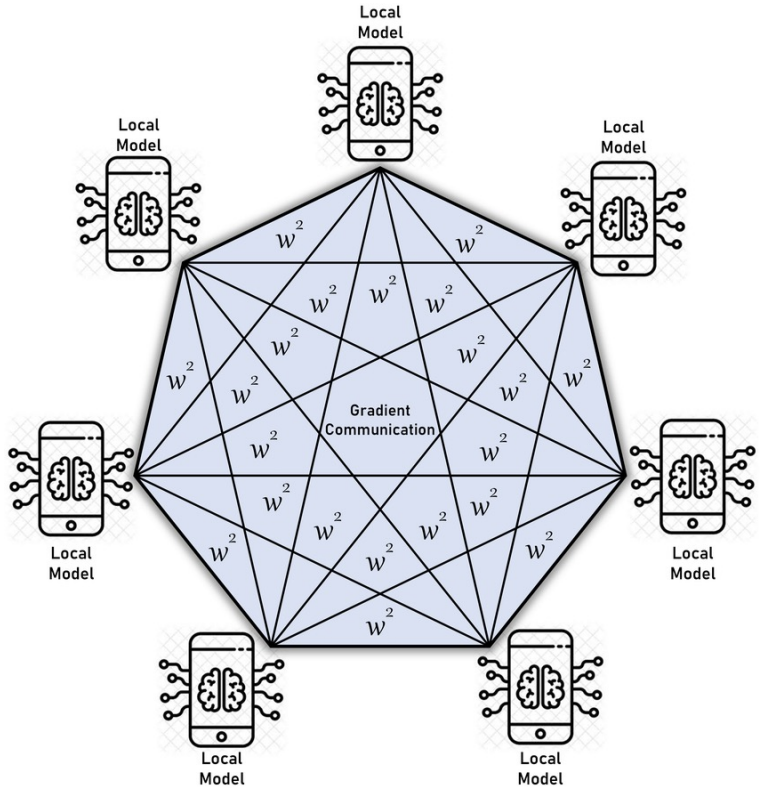

Challenges: Decentralized Federated Learning

IDEA

- Creation of a new integrated environment for developing decentralized federated learning solutions

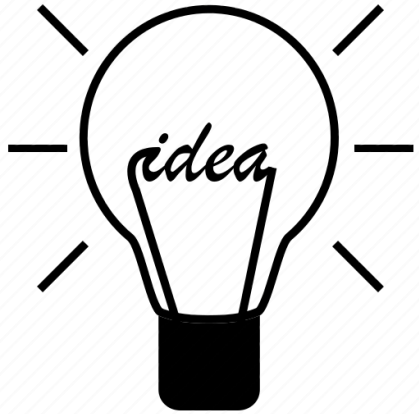

Looking for

- Potential Coordinator
- A consortium to join
- Experts in programming environments
- Case studies to test and validate the approach

From Internet of Things Platforms to Web of Things User Agents

Prof. Dr. Andreas Harth
Department Data Spaces and IoT Solutions, Fraunhofer IIS Nuremberg
Chair of Technical Information Systems, Friedrich-Alexander University Erlangen-Nuremberg

Fraunhofer IIS

FRIEDRICH-ALEXANDER UNIVERSITÄT ERLANGEN-NÜRNBERG

Traditional System Architecture

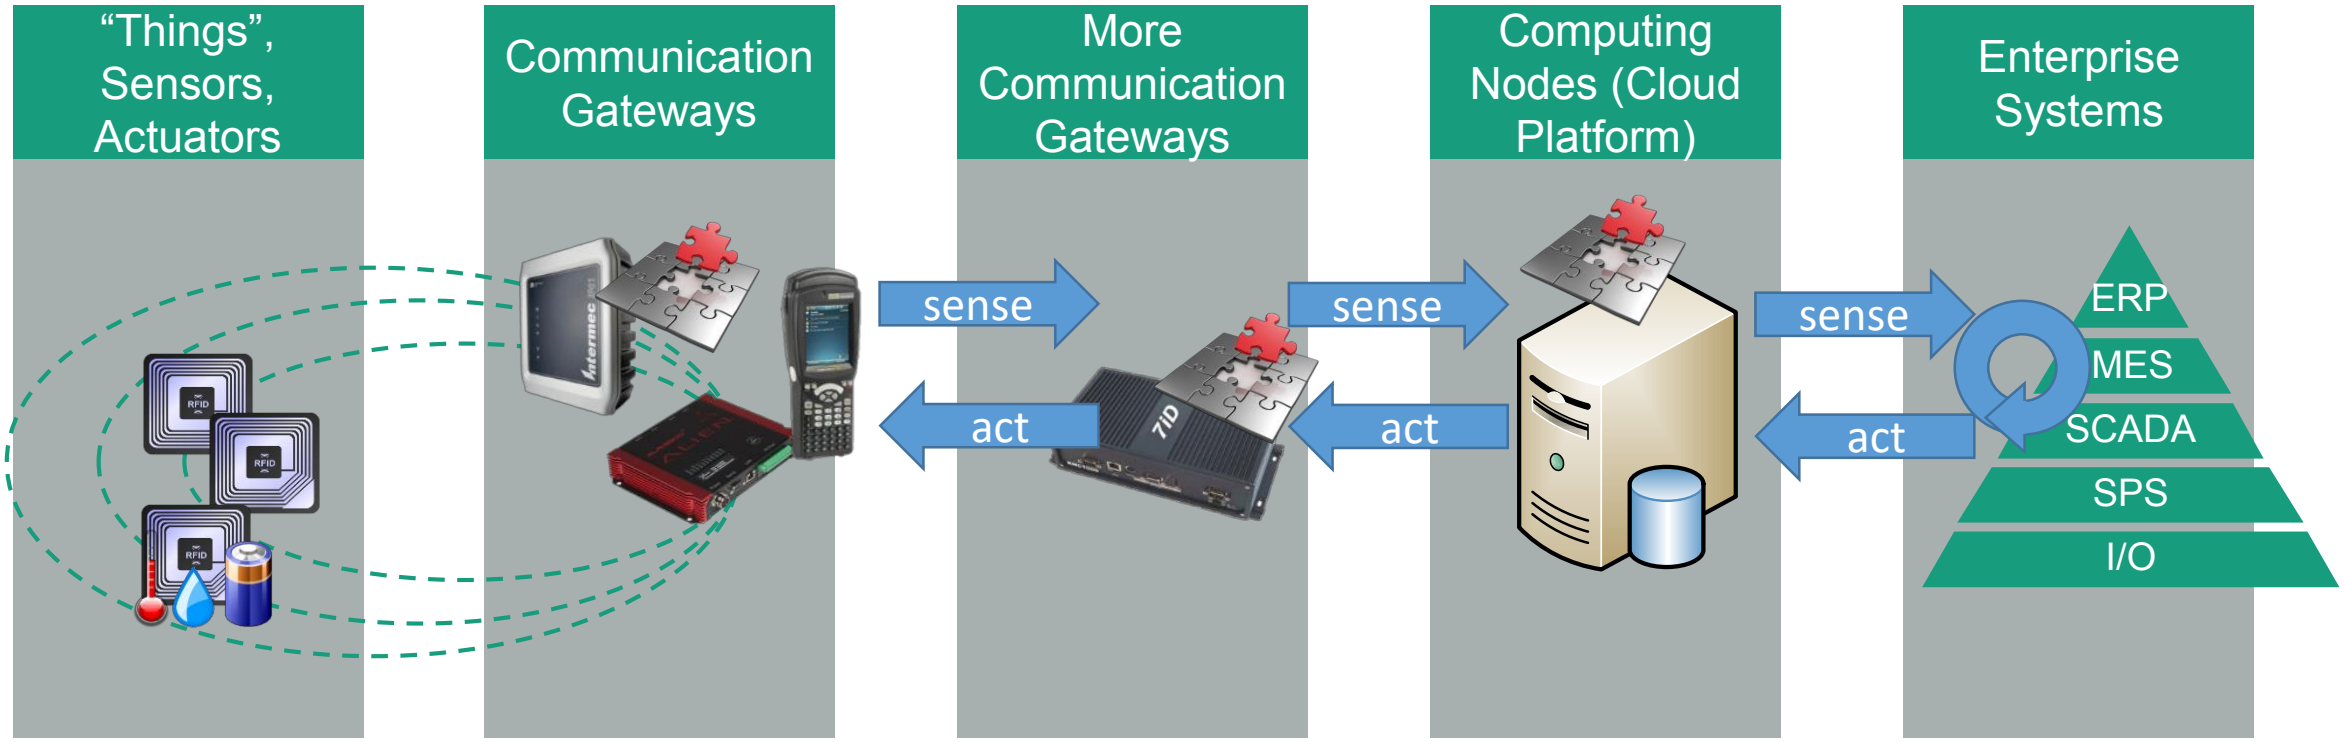

Low-level communication protocols

Communication via TCP/IP ("narrow waist")

Sense-act cycle for control

Selected Research Challenges

- Interoperability: different protocols to access to “things”
  - Standardise data representation and communication protocols
  - Use formal languages for data to enable data portability, data integration and querying
  - Use formal languages to describe network-accessible operations
- Complexity: writing programs that access “things” is hard
  - Provide higher-level programming abstractions
  - Use formal languages to support programming, enable automated composition of processes and program synthesis
- Communication: latency for network connectivity to the cloud
  - Bring computation closer to the devices (“edge computing”)
  - Use formal languages for programs to enable flexible execution at different nodes
- Safety: systems with actuators might harm equipment or people
  - Use formal languages to reason about programs and verify their properties
- Resilience: intermittent communication links, reliance on business decisions and risk of censorship
  - Implement decentralised architectures that reduce or remove centralised components
  - Use formal languages to be able to split up monolithic systems into pieces under diverse ownership

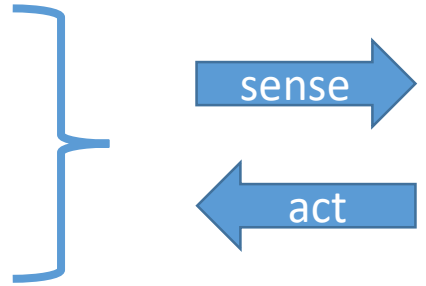

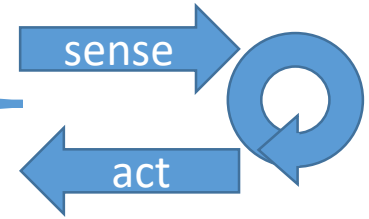

From Representing Knowledge to Representing Behaviour

- We assume a uniform interface to read data from sensors and manipulate actuators
- Access data from sensors
  - Linked Data: GET (read)
  - WoT: read Property, receive Events
- Effect change in actuators
  - Linked Data: PUT (write)
  - WoT: write Property, invoke Actions
- How to represent behaviour?

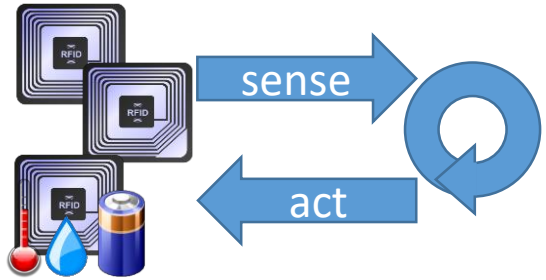

Options for Representing and Executing Behaviour

- Imperative programming languages (e.g., JavaScript)

- Declarative programming approaches
  - (Event-)Condition-Action rules (c.f. Simple Reflex Agents)
  - Process languages, workflow languages

- Declarative programming approaches allow for
  - programming at a higher level of abstraction
  - automated composition of processes and program synthesis
  - flexible execution at different computing nodes
  - reasoning about programs and verifying their properties
  - splitting up monolithic systems under central control into pieces under diverse ownership

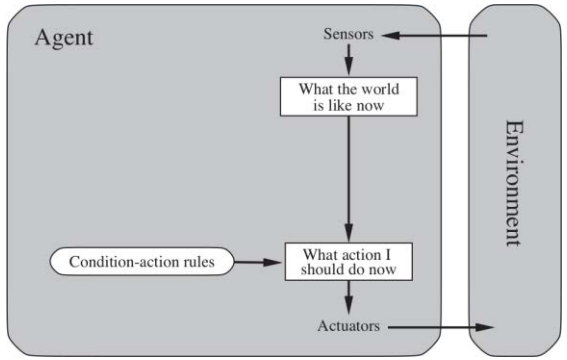

The Economist, Illustration by David Simonds

ubiwhere

Suiting the future of technology

hello@ubiwhere.com

ubiwhere

From Portugal to the world **with love.**

ubiwhere

Can you imagine your city as a single, integrated system?

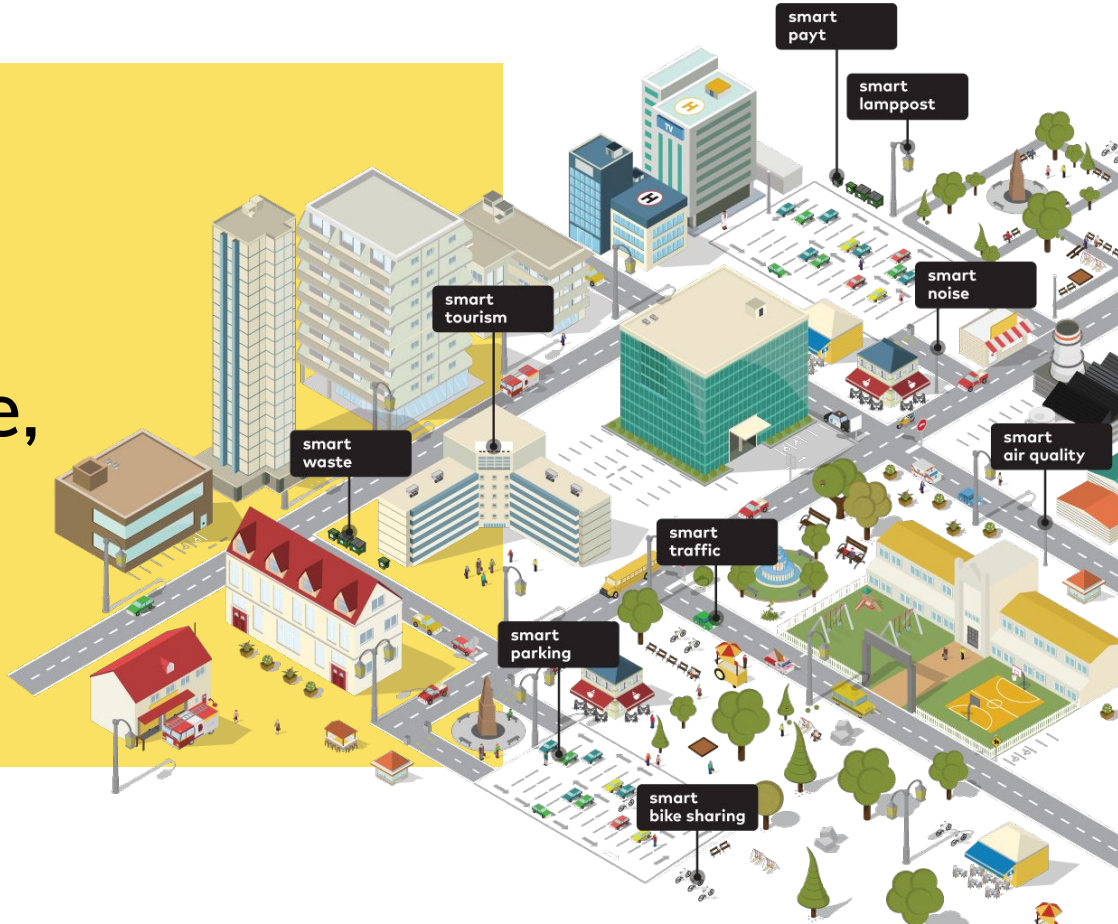

ubiwhere

urbsense

The City Nervous System

An organic approach to

Smart City Management

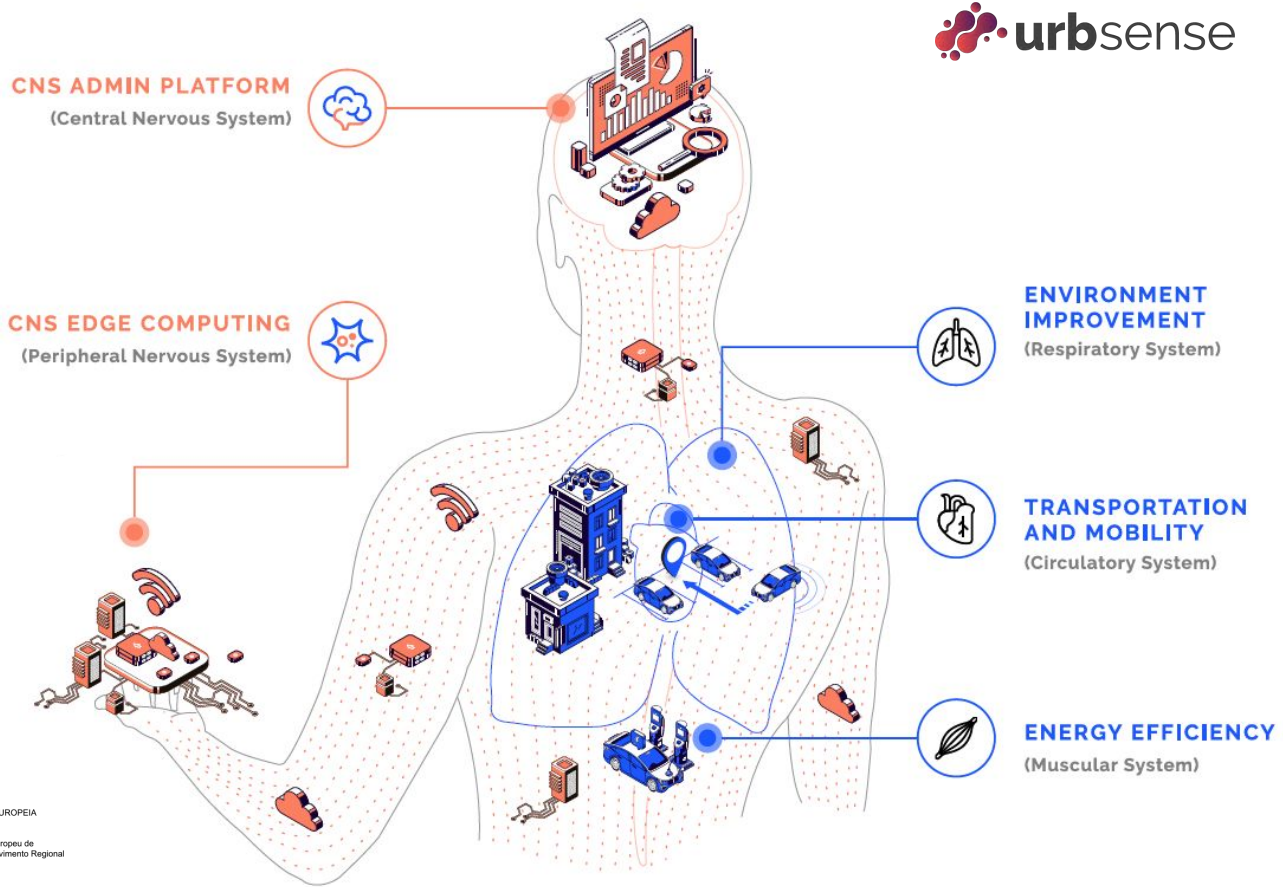

→ urbsense.com

COMPETE 2020 PORTUGAL 2020 UNIÃO EUROPEIA Fundo Europeu de Desenvolvimento Regional

Clusters
SSH Keys
Members
Service Accounts

Delete

Add Machine Deployment | Open Dashboard

cluster-neutral-hosting1

Master Version 1.19.9

Cluster ID: vhd9b2b2gd

Region: NL (Amsterdam)

DigitalOcean

User SSH Key Agent

OPA Integration

cjorge

Created: 3 days ago

Nodes CPU Usage 189/1600 millicores

Nodes Memory Usage 1759/3444 MiB

API Server | etcd | Machine Controller | Scheduler | Controller | User Controller Manager

Machine Deployments

Name	Labels	Replicas	kubelet Version	Operating System	Created
sfdsfd-worker-7kvpb4	system/cluster: vhd9b2b2gd; system/project: qrtlxl6f5d	1 / 1	v1.19.9	Ubuntu	3 days ago

Events | RBAC | Addons

Message	Resource ID	Type	Count	Occurred

No events available.

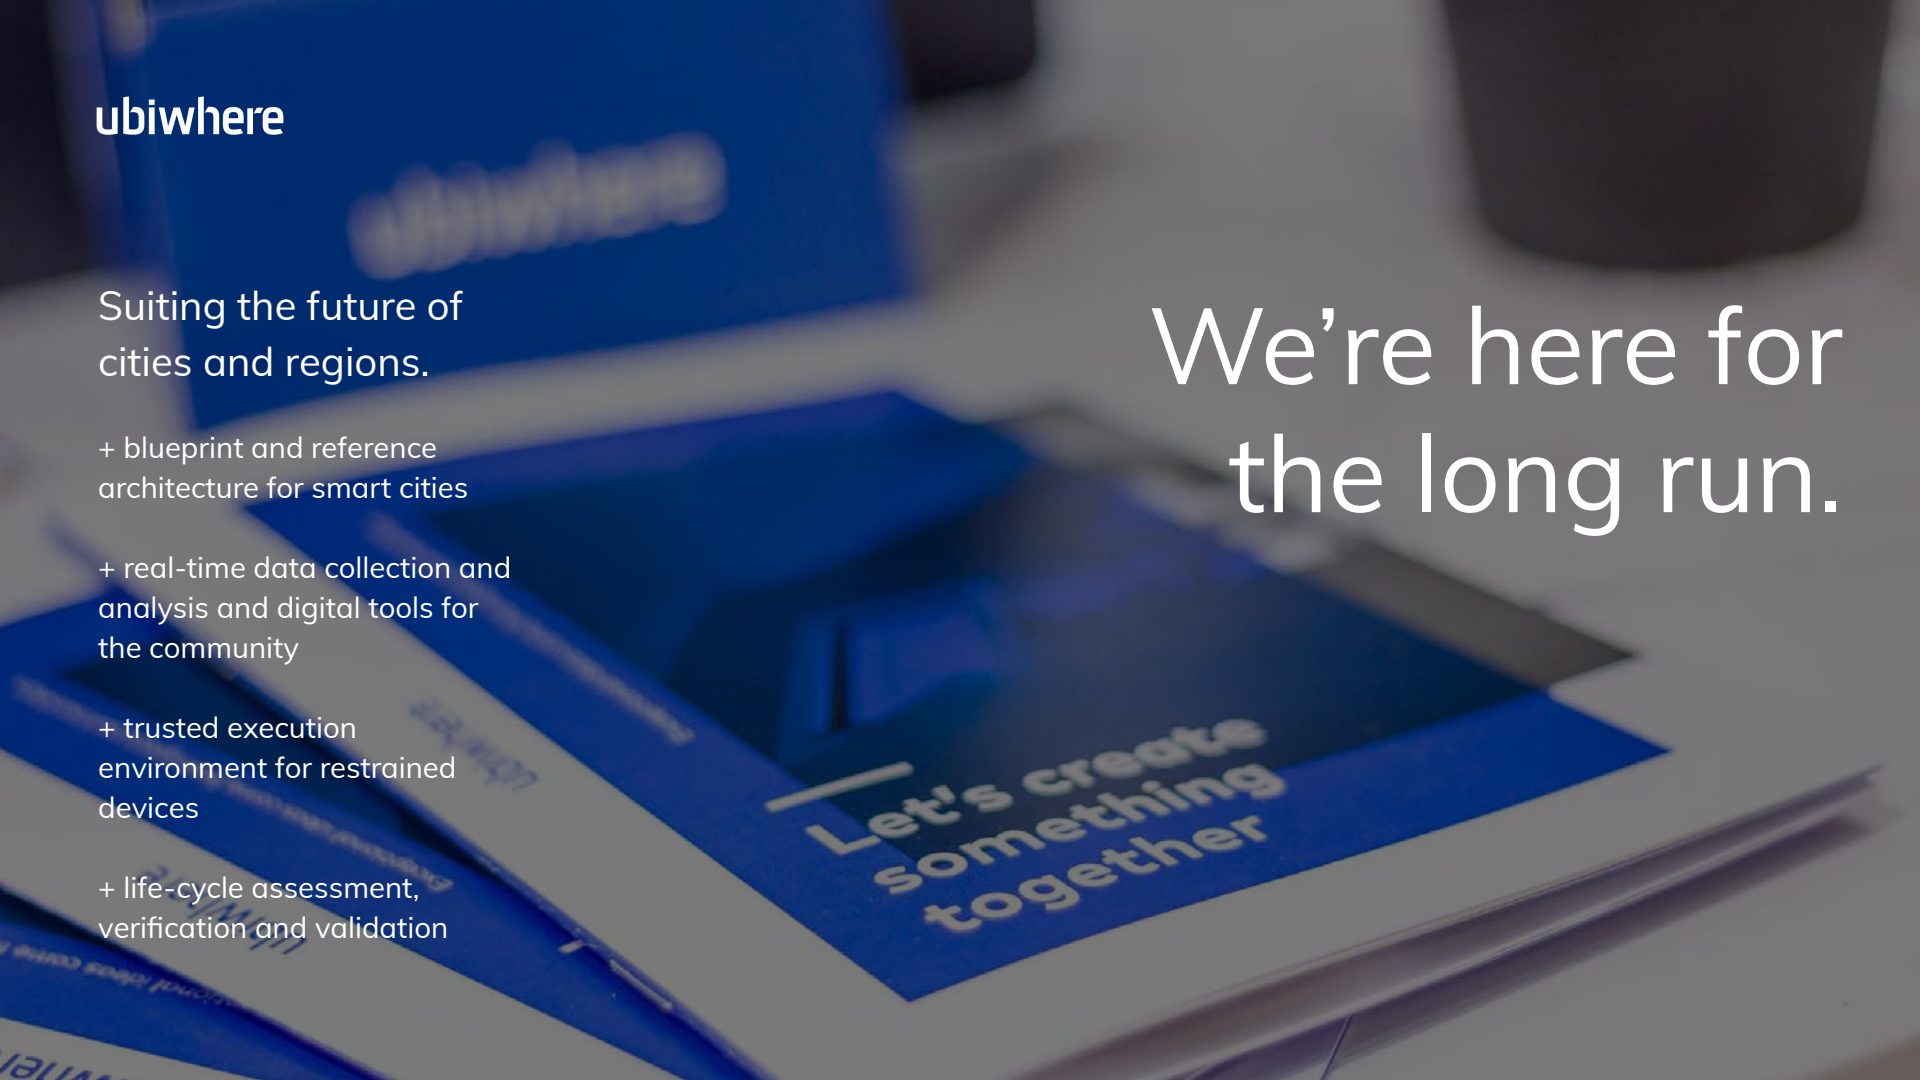
ubiwhere
Suiting the future of cities and regions.
+ blueprint and reference architecture for smart cities
+ real-time data collection and analysis and digital tools for the community
+ trusted execution environment for restrained devices
+ life-cycle assessment, verification and validation
We're here for the long run.
Let's create something together

ubiwhere

Let’s go for it?

Ricardo Vitorino

Head of Innovation (Smart Cities)

innovation@ubiwhere.com

ubiwhere

Information & Brokerage Session

Horizon Europe Cloud-Edge-IoT Call 2022

12:00
12:30

WRAP-UP AND CLOSING SESSION

Rolf Riemenschneider
Head of Sector IoT
European Commission

Maria Tsakali
Programme Officer, Cloud and Software Unit, European Commission

Luis Busquets Pérez
Programme Officer, Cloud and Software Unit, European Commission

Moderator

Monique Calisti
CEO
Martel Innovate

Jan Komarek
Topic Coordinator, IoT Unit, European Commission

THANK YOU !