DESIRE6G

ON THE WAY TO 6G –
DISTRIBUTED INTELLIGENT CONTROL, DEEP PROGRAMMABILITY AND HARDWARE ACCELERATION

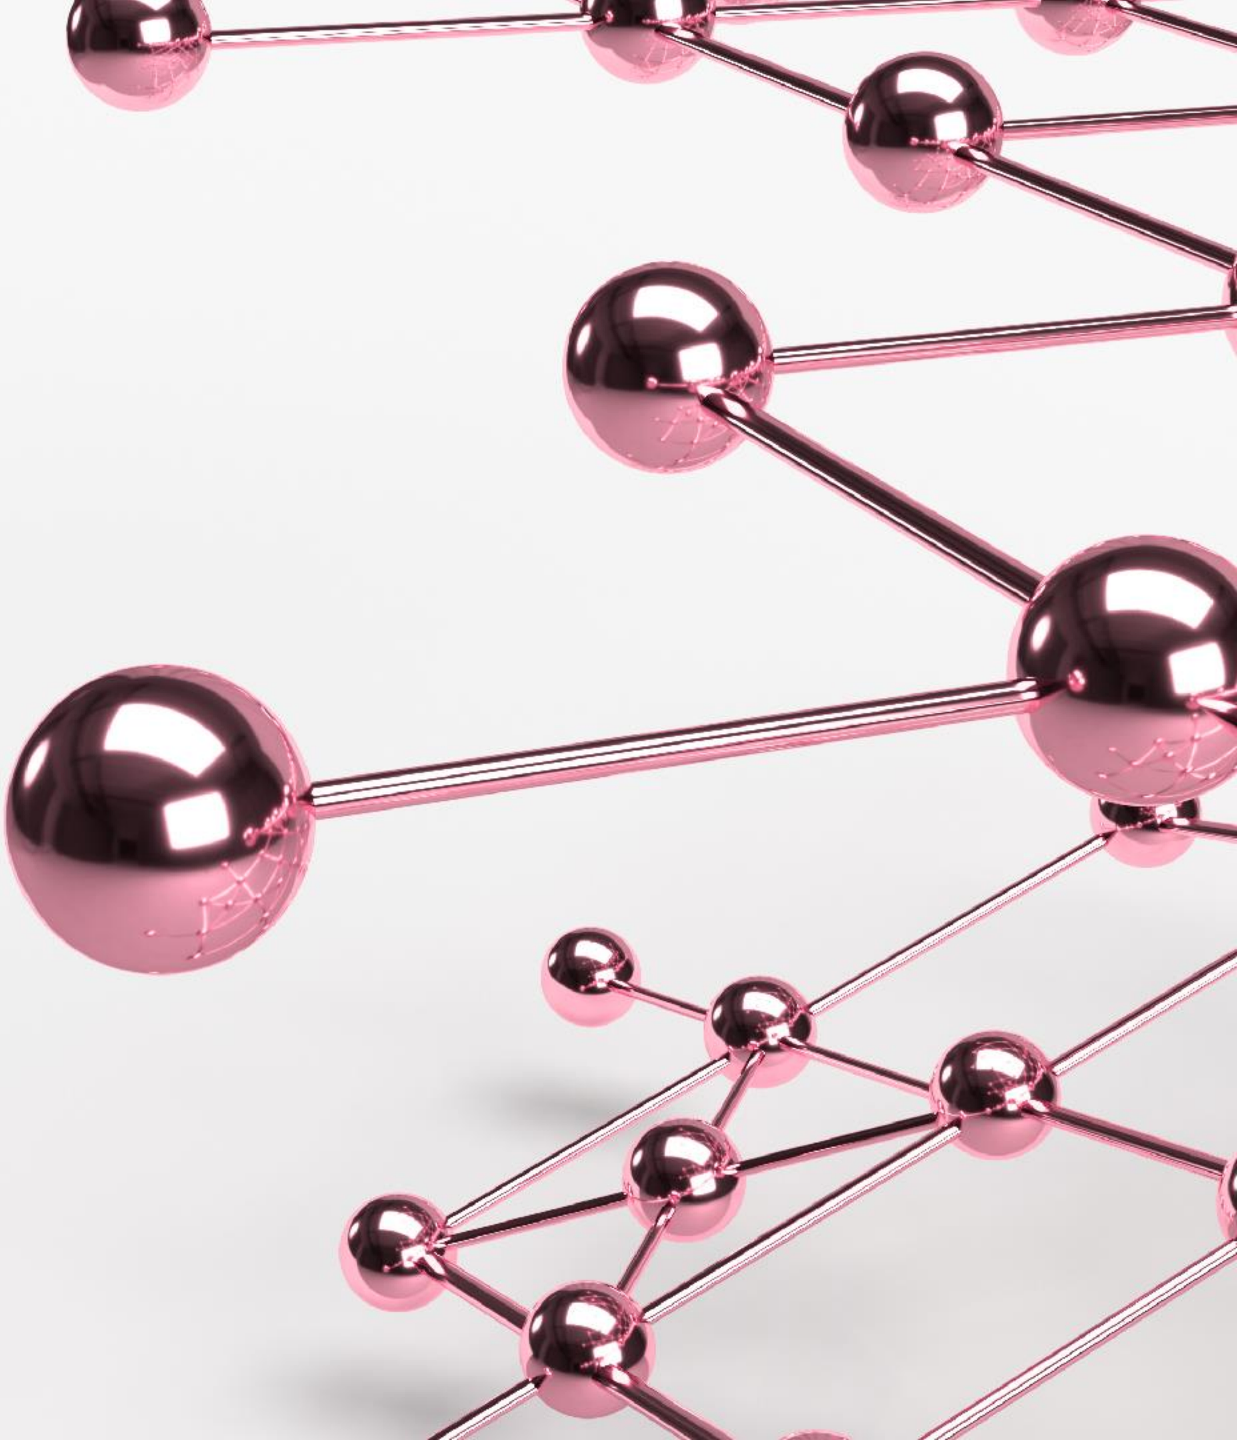

Sándor Laki, assistant professor

Communication Networks Laboratory
Faculty of Informatics
ELTE Eötvös Loránd University
Budapest, Hungary

Co-funded by
the European Union

6GSNS

10/09/2023

WHAT IS 6G?

- No general globally-accepted vision on 6G
- European vision (6G-SNS)
  - Massive digitalization – Phy representation
  - Connected intelligence – Awareness, real-timeness
  - Network as Compute Fabric - Decisions, actions
- Key values
  - Sustainability
  - Inclusion
  - Trustworhiness

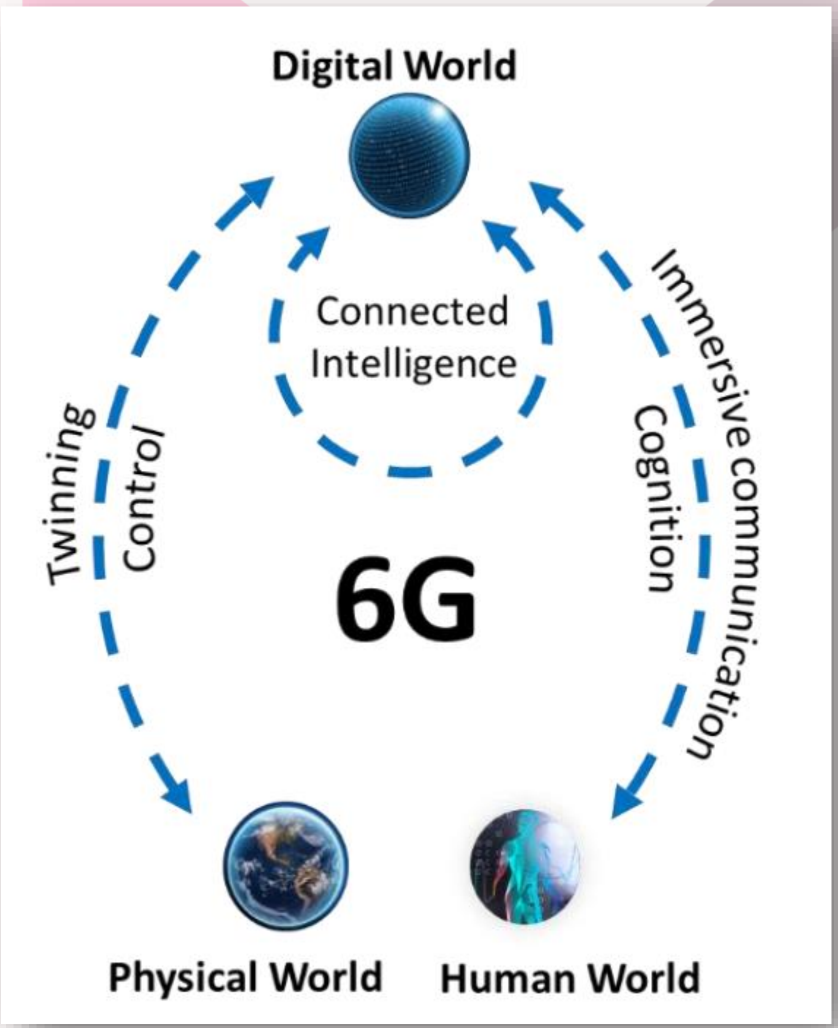

Source: https://5g-ppp.eu/wp-content/uploads/2021/06/WhitePaper-6G-Europe.pdf

6G GOALS

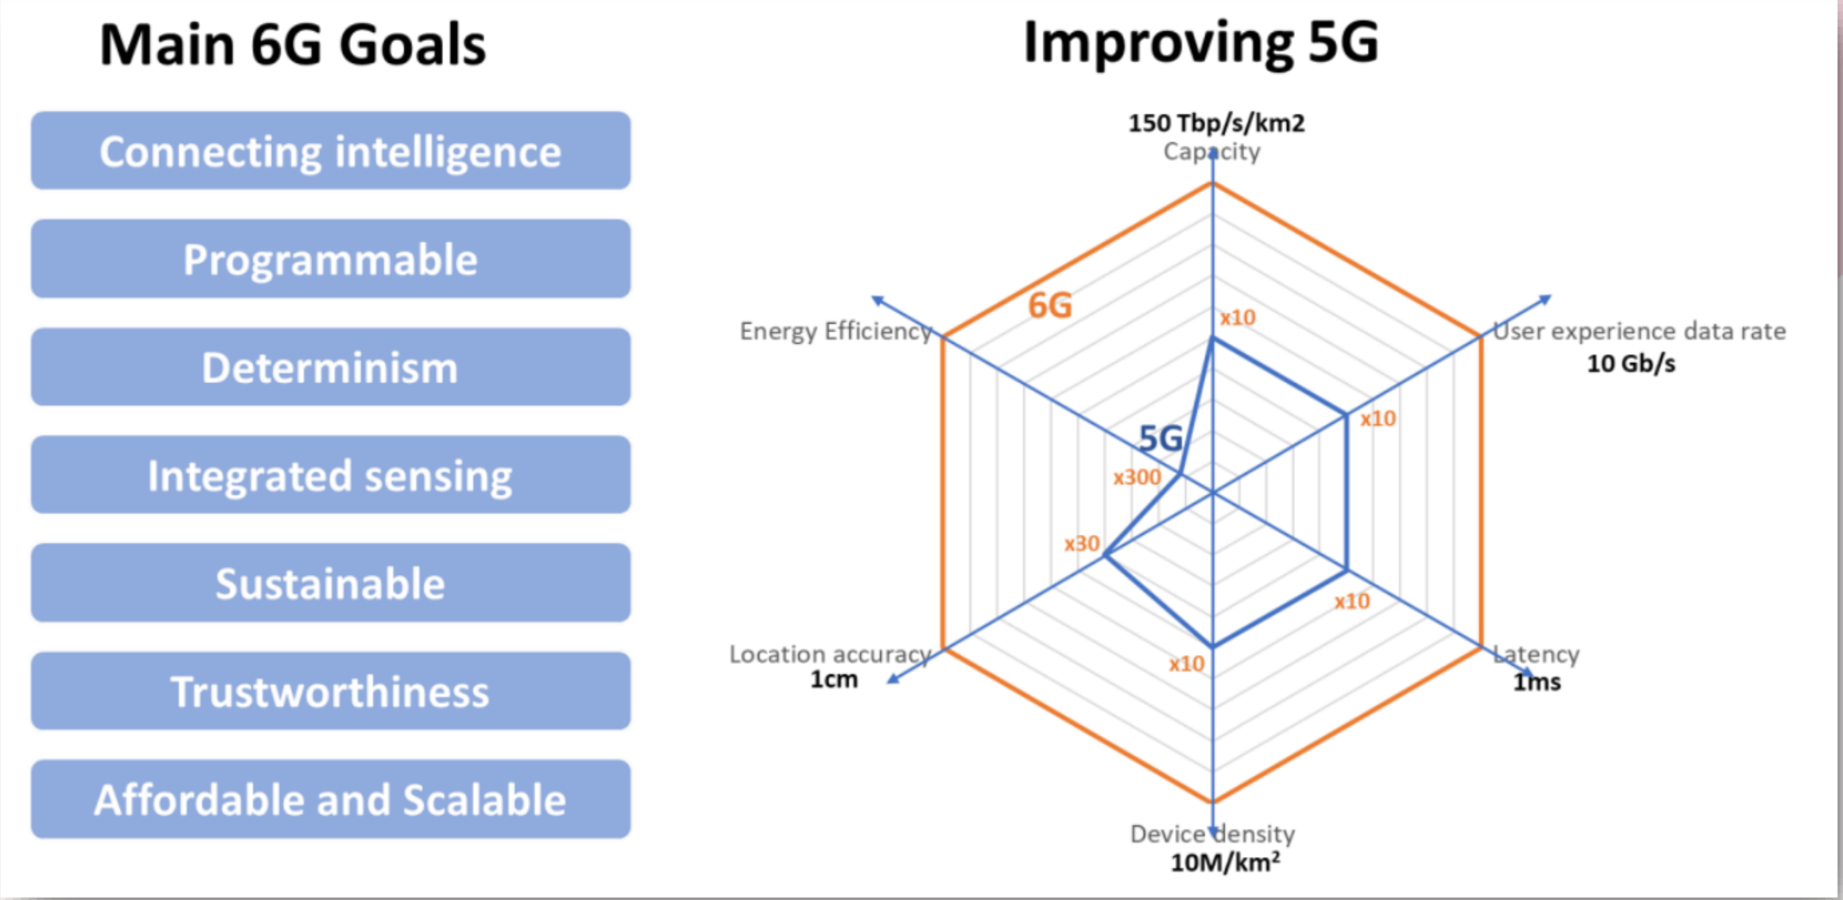

Source: https://5g-ppp.eu/wp-content/uploads/2021/06/WhitePaper-6G-Europe.pdf

USE CASES

- Holographic teleportation
- Extended reality – AR/VR
- Pervasive connectivity – Internet of Everything
- UAV services
- Autonomous services
- Ambient connectivity

RESEARCH CHALLENGES

- Radio HW
  - Achieve 1 Tbps; semiconductors, optics and new materials in THz applications;
- Physical layer
  - Ultra-low-power communication, physical layer security, high spectral efficiency
- Networking infrastructure
  - Embedded trust, attack protection and mitigation, differentiated service quality, high flexibility, network as a computing platform, end-to-end point of view, dynamic service management, distributed control/intelligence, zero-touch operation
- New service enablers
  - Support for a wide range of services, user-specific computations and intelligence at edge cloud, increased sensing and accurate positioning, increased trust and privacy, deterministic networking

DESIRE6G

More complete list: https://dl.acm.org/doi/pdf/10.1145/3571072

WHAT WERE THE MAIN PROMISES OF 5G?

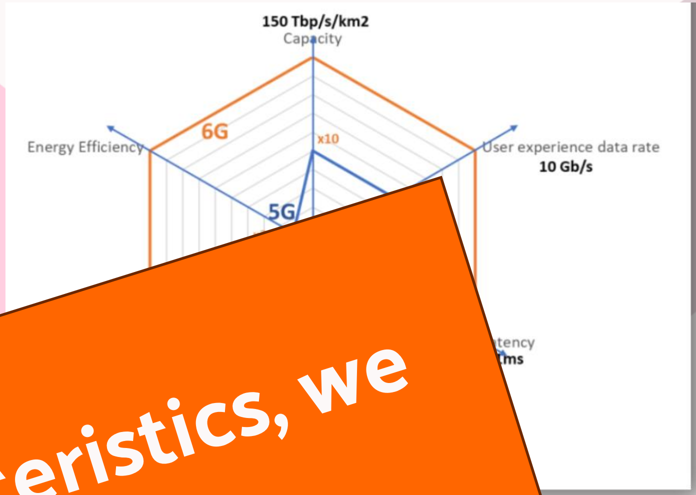

- High performance/High throughput
- A converged infrastructure (radi
- Automation and
- Int
- g and diverse performance

So besides further improving radio characteristics, we need to consider architectural changes too

DESIRE6G

ARCHITECTURAL CHALLENGES FOR 6G

- Main questions of all architecture discussions:
  - How should the functions be grouped / split?
  - How should the interfaces and procedures look like?

- 5G was addressing complexity issues, but only with partial success:
  - “Service Based Architecture” (SBA) became heavier and less cloud-native than expected
  - User plane remained mainly node-based, no “cloud-native” evolution happened there
  - Too detailed standards, less room for vendor innovation
  - The standard does not really count on using IT frameworks/tools to simplify the architecture

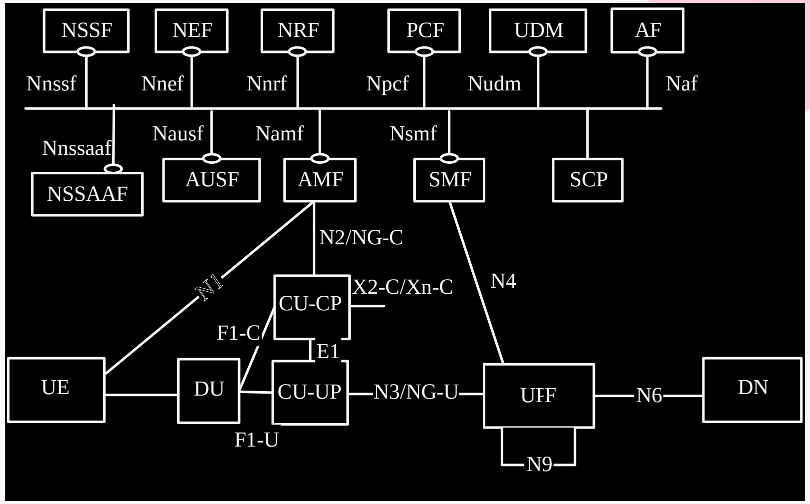

?

DESIRE6G

DEEP PROGRAMMABILITY & SECURE DISTRIBUTED INTELLIGENCE FOR REAL-TIME END-TO-END 6G NETWORKS

Project coordination:
University of Amsterdam

Technical coordination:
Ericsson Hungary

Duration:
01/01/2023 - 31/12/2025

Total Cost:
6.227.919€

Follows us on: desire6g.eu @DESIRE6G_EU @DESIRE6G

DEEP PROGRAMMABILITY & SECURE DISTRIBUTED INTELLIGENCE FO… 6G NETWORKS

Project coor…
University of…

Technical co…
Ericsson Hun…

Duration:
01/01/2023 -…

Total Cost:
6.227.919€

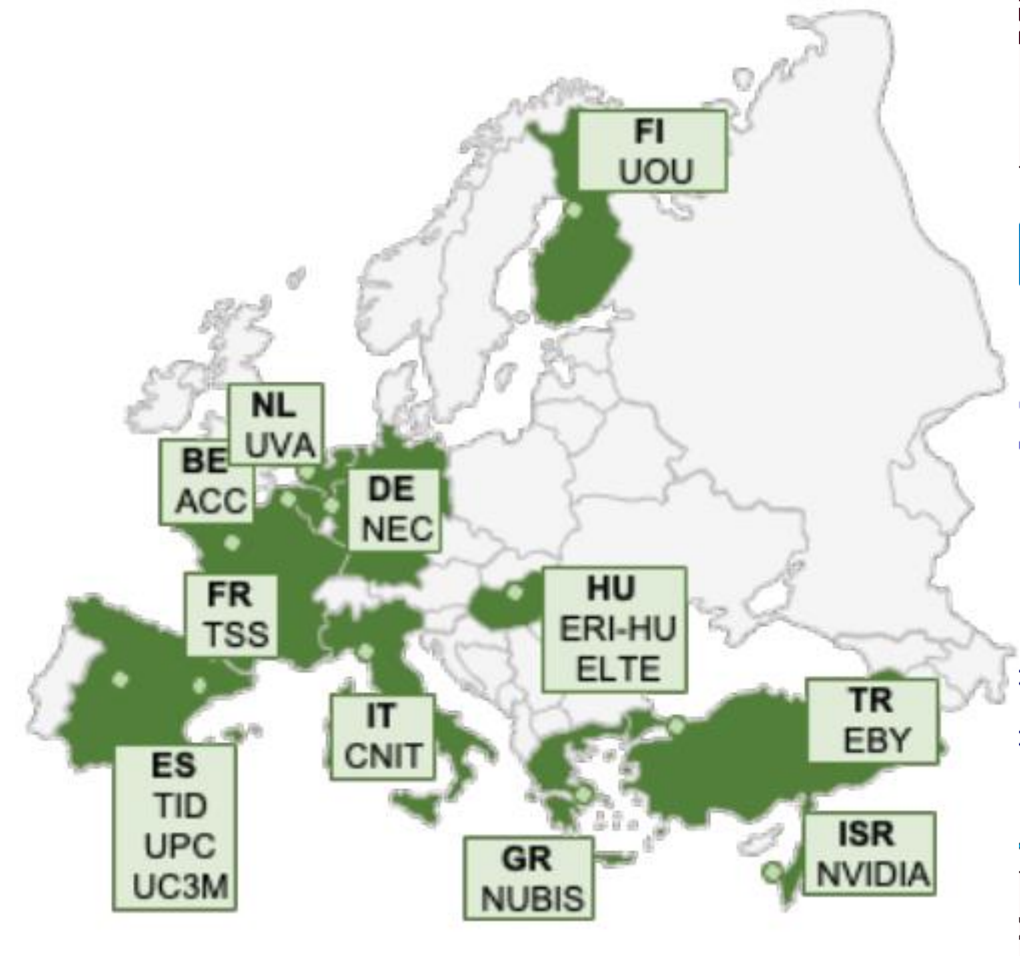

UNIVERSITAT POLITÈCNICA DE CATALUNYA BARCELONATECH

UNIVERSITY OF OULU

ELTE EÖTVÖS LORÁND UNIVERSITY

A((elleran

…ID SHIELD

Follows … …DESIRE6G

DEEP PROGRAMMABILITY & SECURE DISTRIBUTED INTELLIGENCE FO[...] 6G NETWORKS

ESIRE6G

WHY DESIRE6G?

What is the difference between D6G and the other 6G projects?

We study

- How end-to-end network programmability helps in solving really challenging use cases / KPIs (such as below ms latency)
- How to solve the complexity problem of centralized control and optimization with a distributed agent-based system
- And how can we put this together as simply as possible with other innovative methods, like AI-driven telemetry, blockchain-based federation and a DLT-backed software security framework
- So D6G has a **bottom-up** view and focuses on proof of concept **demos** to validate the value proposition

D6G ARCHITECTURE

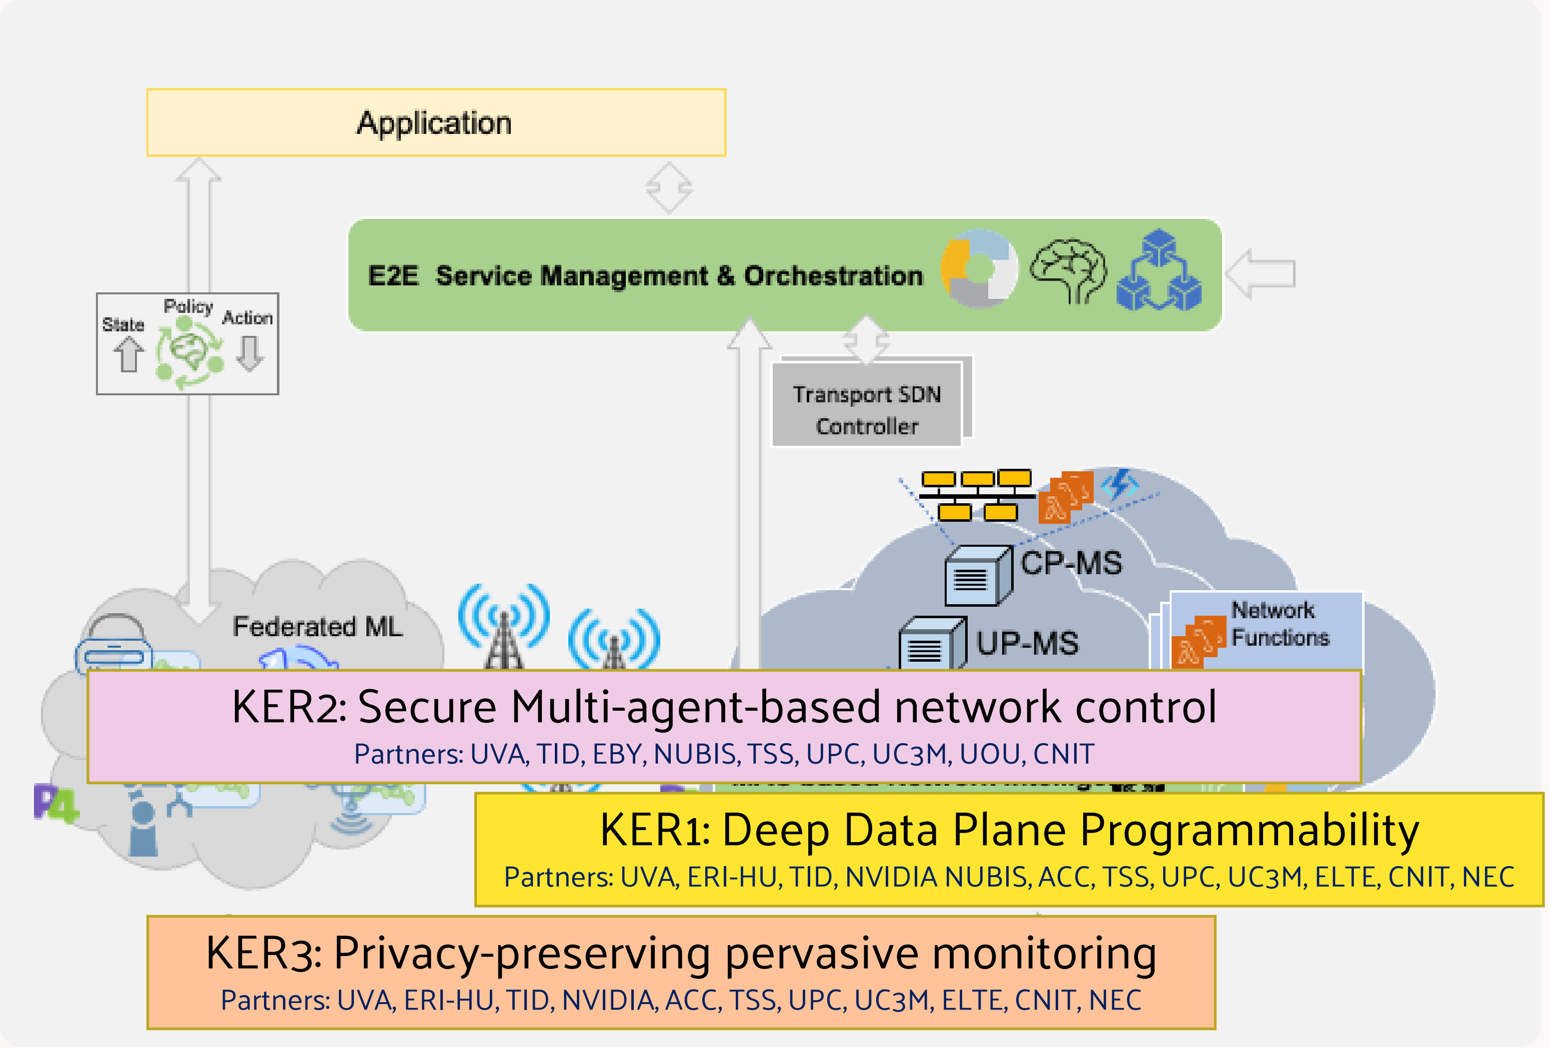

KER2: Secure Multi-agent-based network control
Partners: UVA, TID, EBY, NUBIS, TSS, UPC, UC3M, UOU, CNIT

KER1: Deep Data Plane Programmability
Partners: UVA, ERI-HU, TID, NVIDIA NUBIS, ACC, TSS, UPC, UC3M, ELTE, CNIT, NEC

KER3: Privacy-preserving pervasive monitoring
Partners: UVA, ERI-HU, TID, NVIDIA, ACC, TSS, UPC, UC3M, ELTE, CNIT, NEC

Innovation

- Intelligent Applications
- End-to-end and per-domain intent-based orchestration
- Lightweight blockchain-based federation
- xURLLC services
- Serverless Architecture
- RAN-Core Convergence
- Cloud-to-edge continuum
- Secure Distributed Intelligence
- E2E Data Plane Programmability
- Multi HW acceleration
- E2E Telemetry

DEEP PROGRAMMABILITY

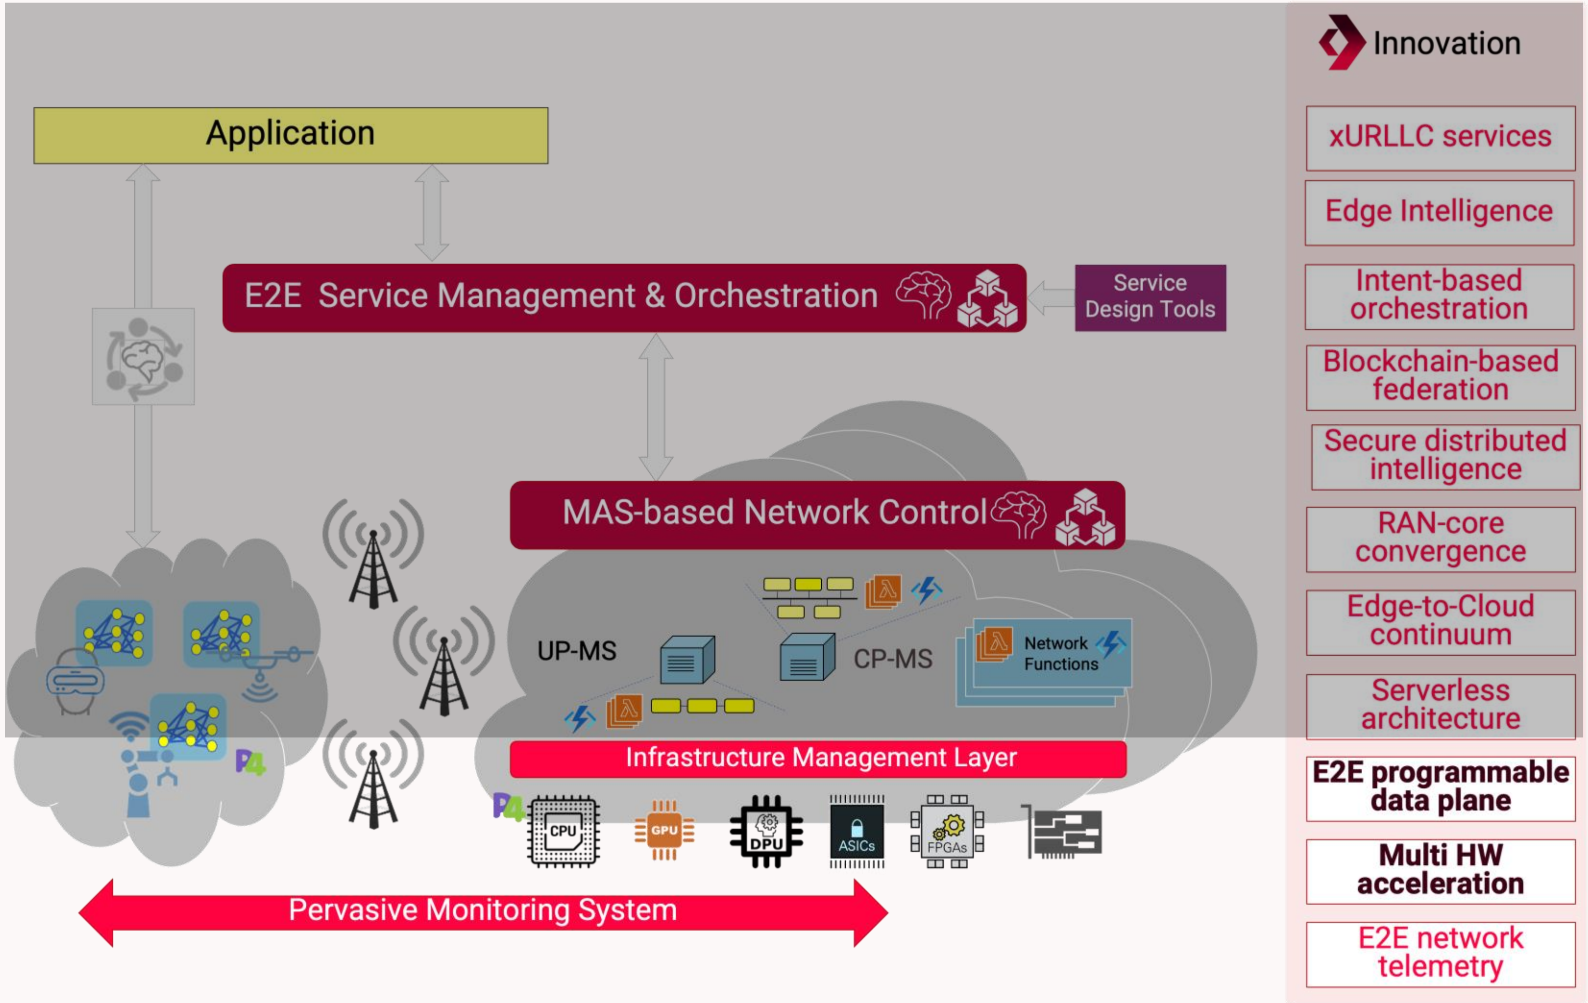

DEEP PROGRAMMABILITY

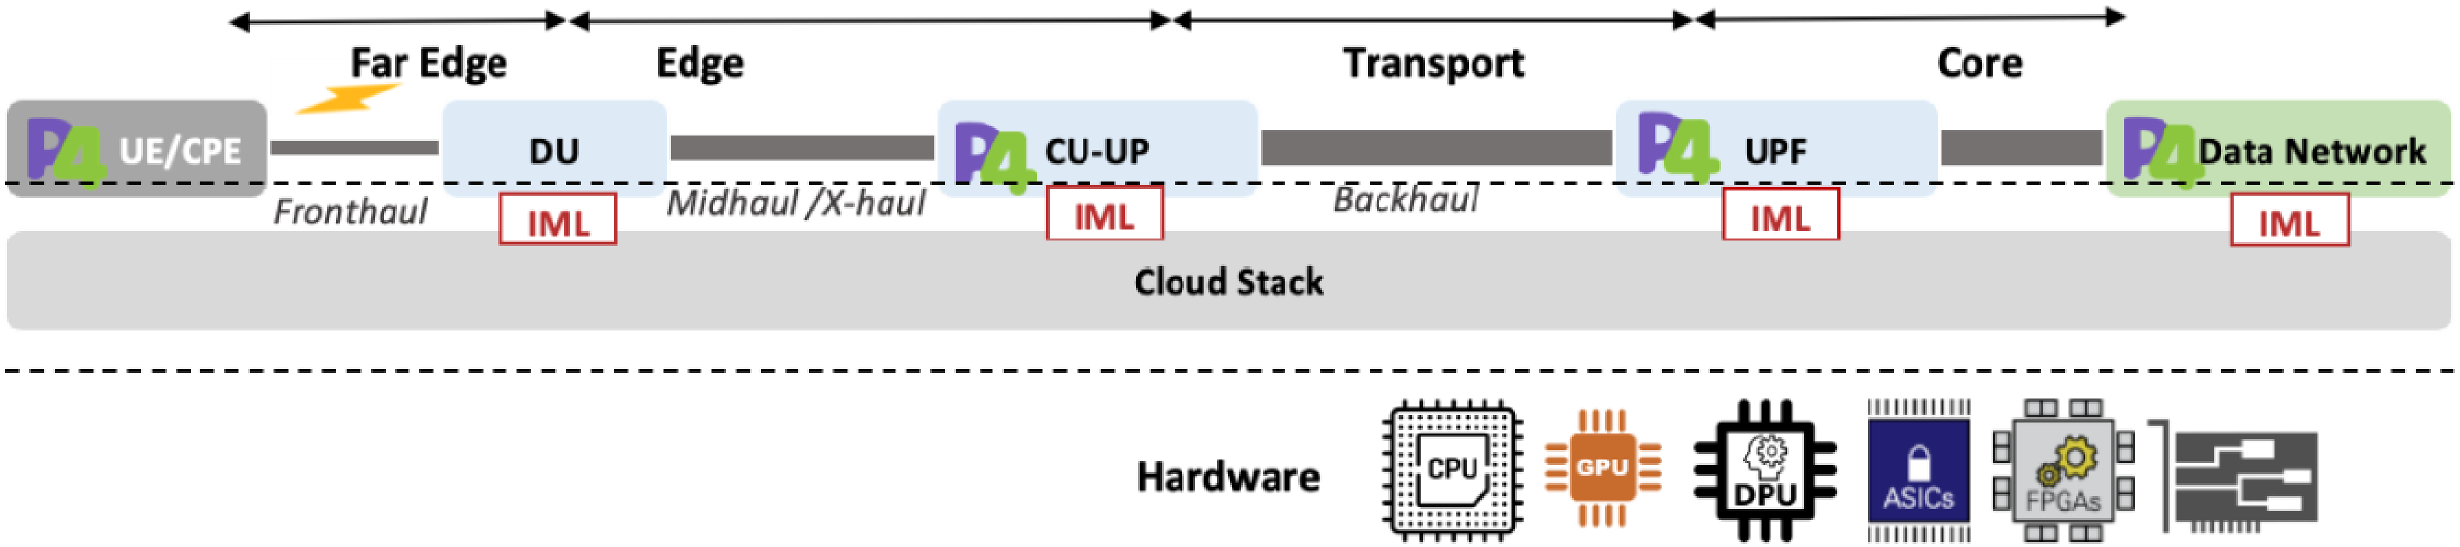

E2E NETWORK VISIBILITY

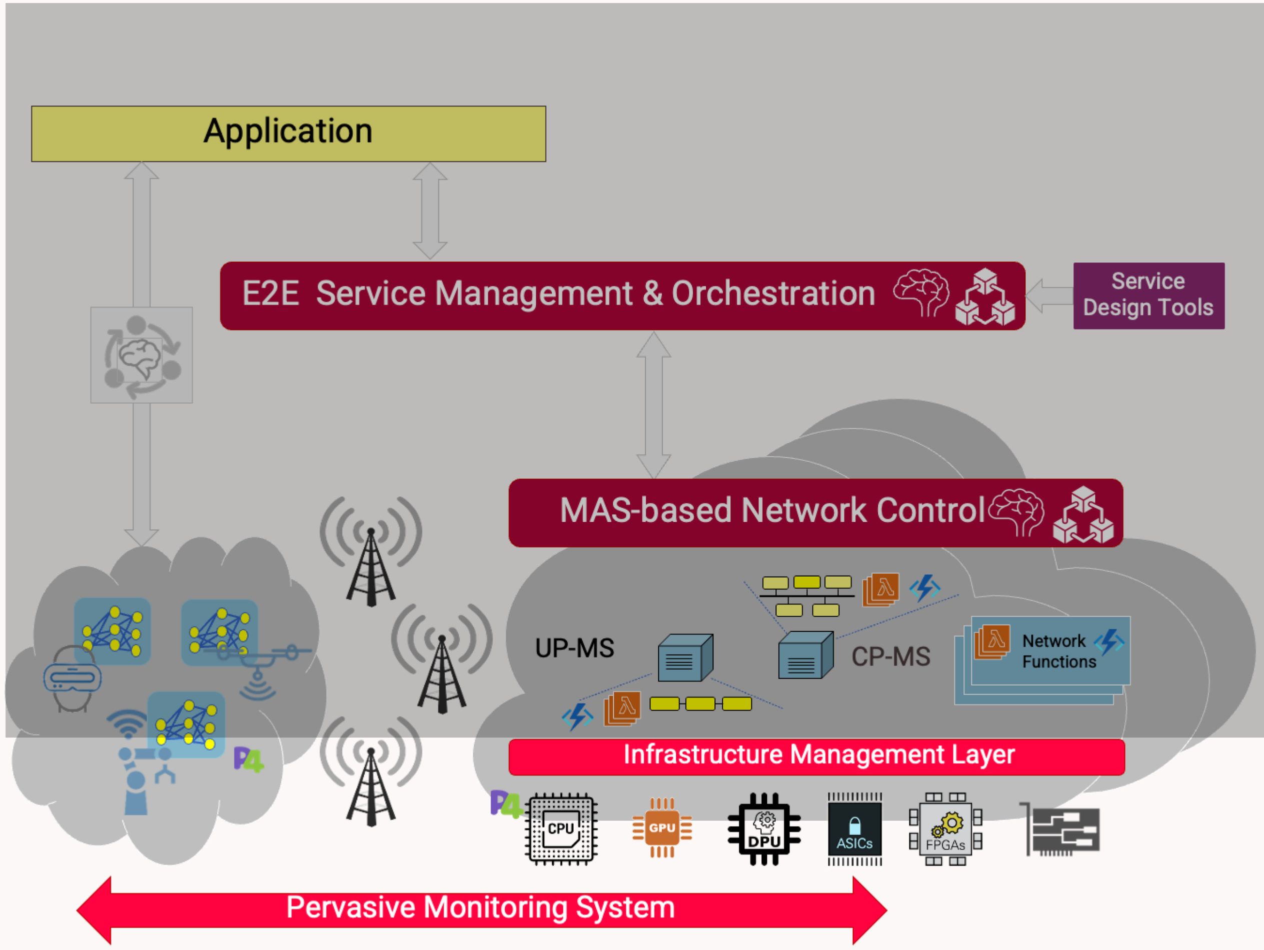

Innovation

- xURLLC services
- Edge Intelligence
- Intent-based orchestration
- Blockchain-based federation
- Secure distributed intelligence
- RAN-core convergence
- Edge-to-Cloud continuum
- Serverless architecture
- **E2E programmable data plane**
- Multi HW acceleration
- **E2E network telemetry**

CLOUD NATIVE

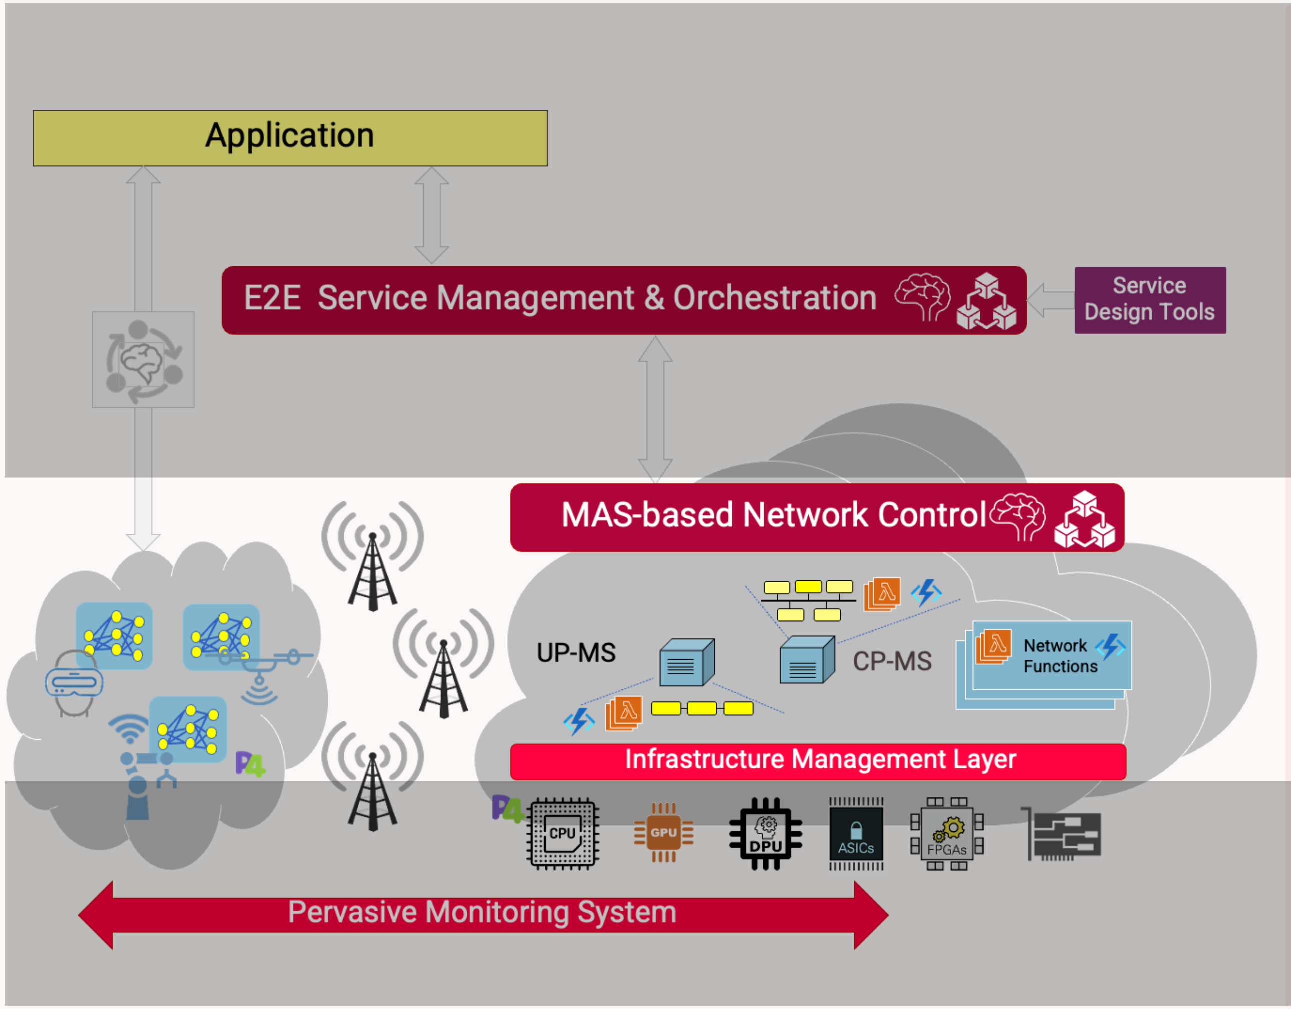

Innovation

- xURLLC services
- Edge Intelligence
- Intent-based orchestration
- Blockchain-based federation
- Secure distributed intelligence
- **RAN-core convergence**
- **Edge-to-Cloud continuum**
- **Serverless architecture**
- E2E programmable data plane
- Multi HW acceleration
- E2E network telemetry

AI-NATIVE

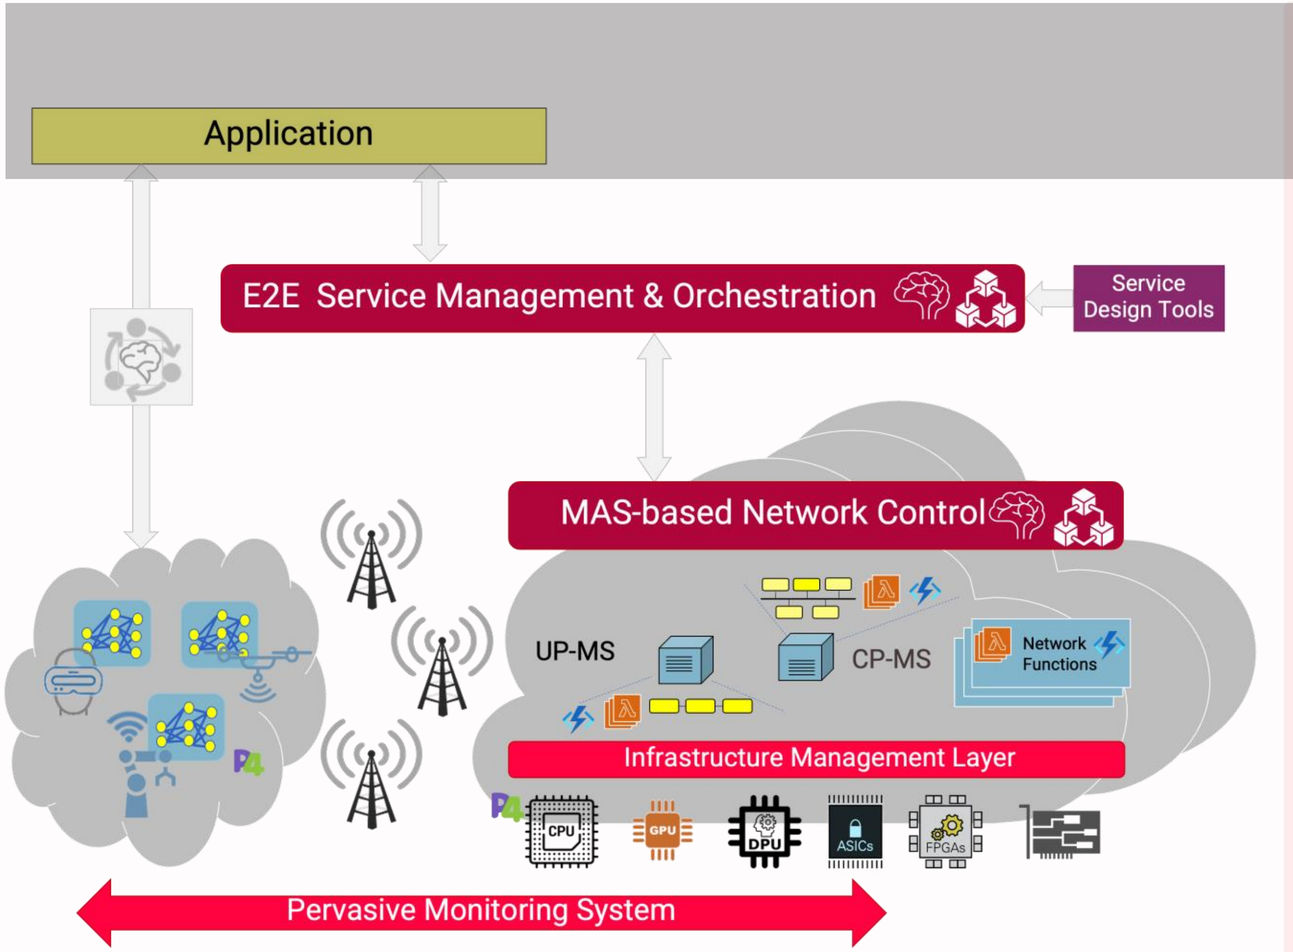

Innovation

- xURLLC services
- Edge Intelligence
- **Intent-based orchestration**
- Blockchain-based federation
- **Secure distributed intelligence**
- RAN-core convergence
- Edge-to-Cloud continuum
- Serverless architecture
- E2E programmable data plane
- Multi HW acceleration
- E2E network telemetry

DLT FOR ZERO-TRUST ARCHITECTURE

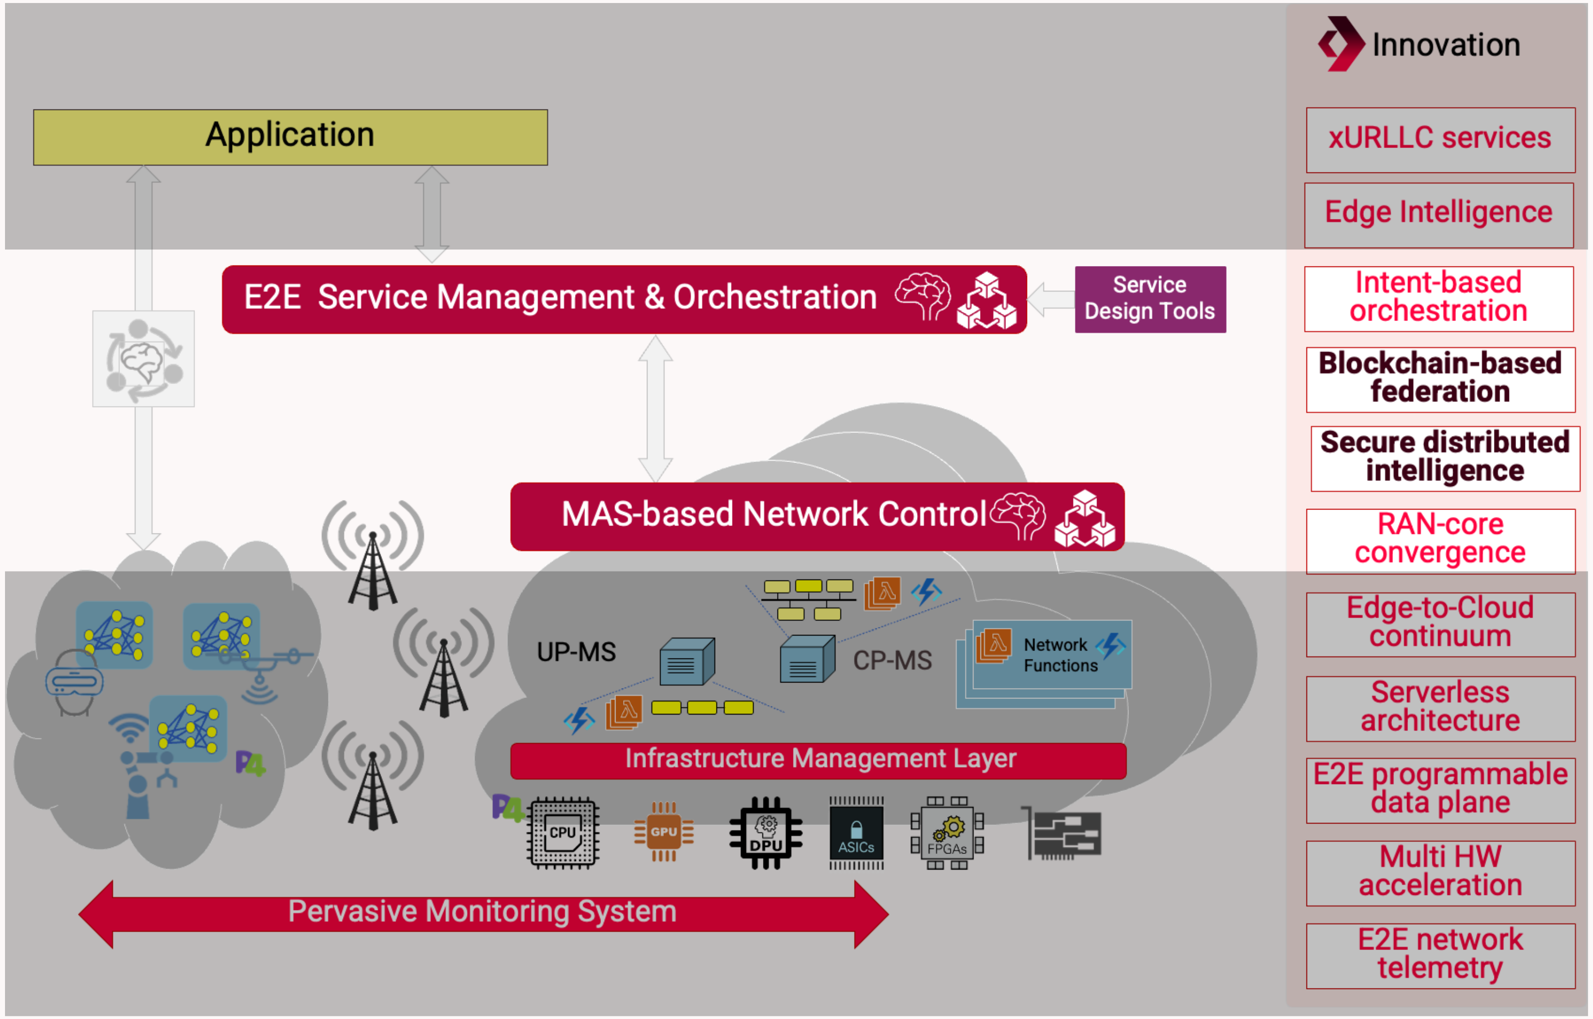

WP STRUCTURE

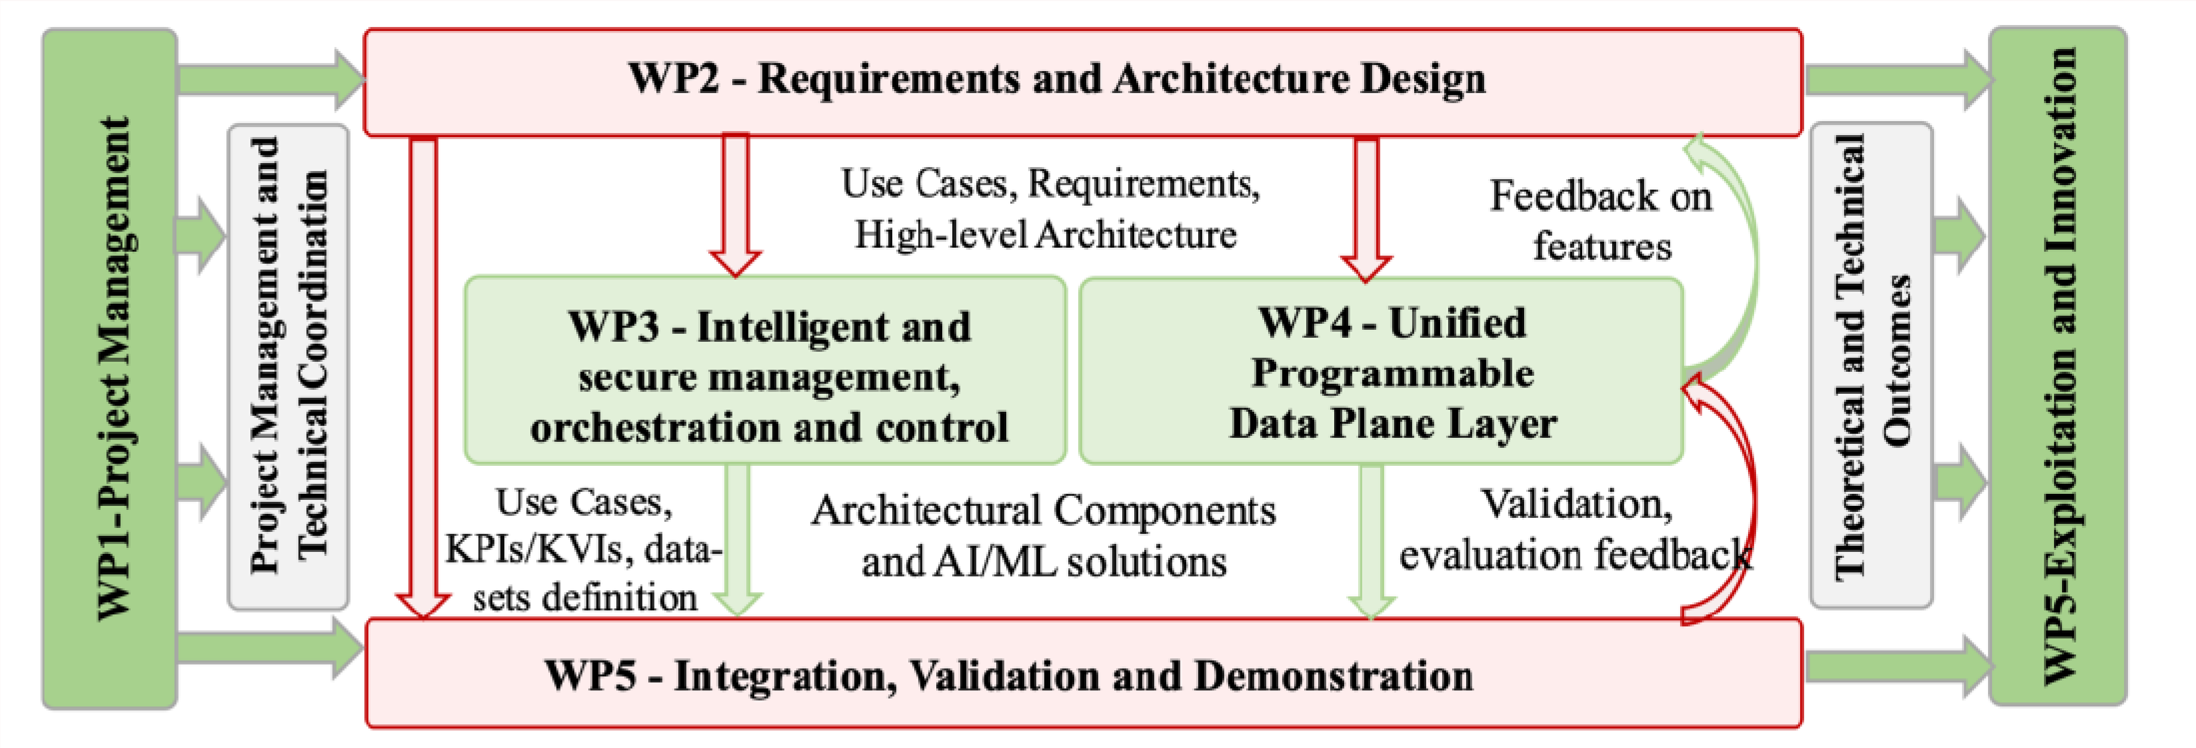

DESIRE6G

THE DATA/USER PLANE ARCHITECURE: TOWARDS A UNIFIED CLOUD-NATIVE DATA PLANE

6GSNS DESIRE6G has received funding from the Smart Networks and Services Joint Undertaking (SNS JU) under the European Union's Horizon Europe research and innovation programme under Grant Agreement No 101096466.

Co-funded by the European Union

SOFTWARIZATION TREND IN PACKET CORE NETWORKS

- Delivering new functionalities
  - Timely and customized way
- Softwarized packet core
  - Packet processing in software
  - Running on commodity servers
- High flexibility and good scalability
  - Software instances can be scaled up or down
  - Network Function Virtualization

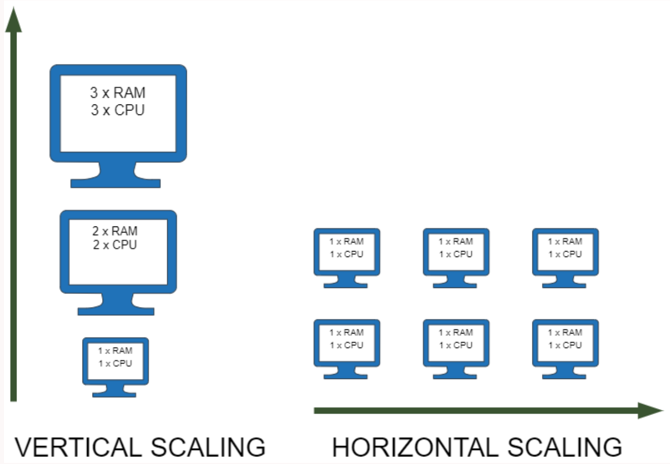

DRAWBACKS

- Unpredictable latency and problems with low latency guarantees
  - Commodity hardware not designed for packet processing
- Throughput limits
  - Several bottlenecks: PCIe speed, cache misses, memory access, etc.

- Kernel-bypass techniques
  - High performance packet processing
  - Needed for good througput
  - Fully utilized CPU cores
  - Constantly polling NICs
- High energy consumption
  - W/pps
  - Increasing OPEX

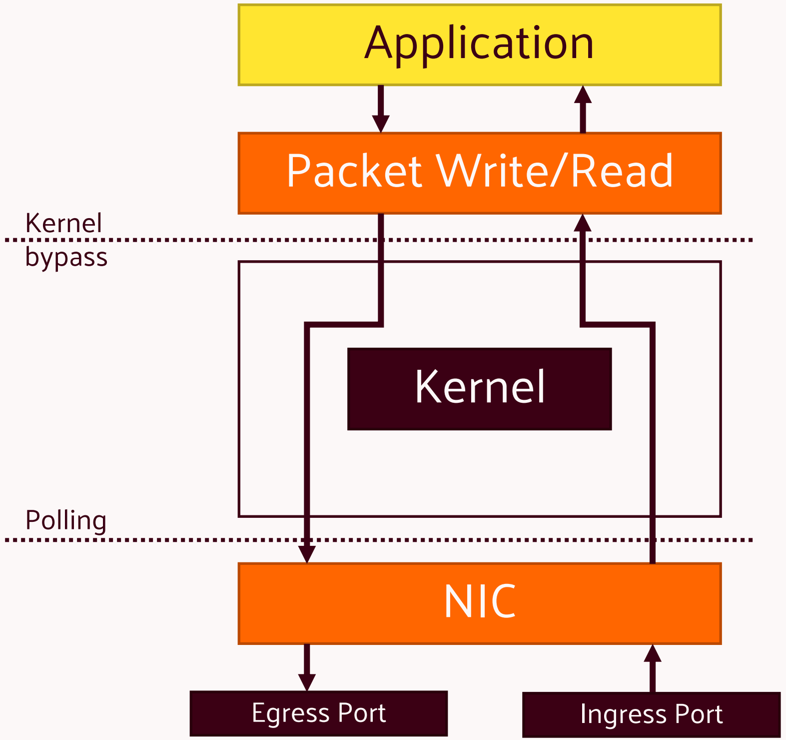

DESIRE6G

PROGRAMMABLE NETWORK DEVICES AS NF(V) BACKENDS

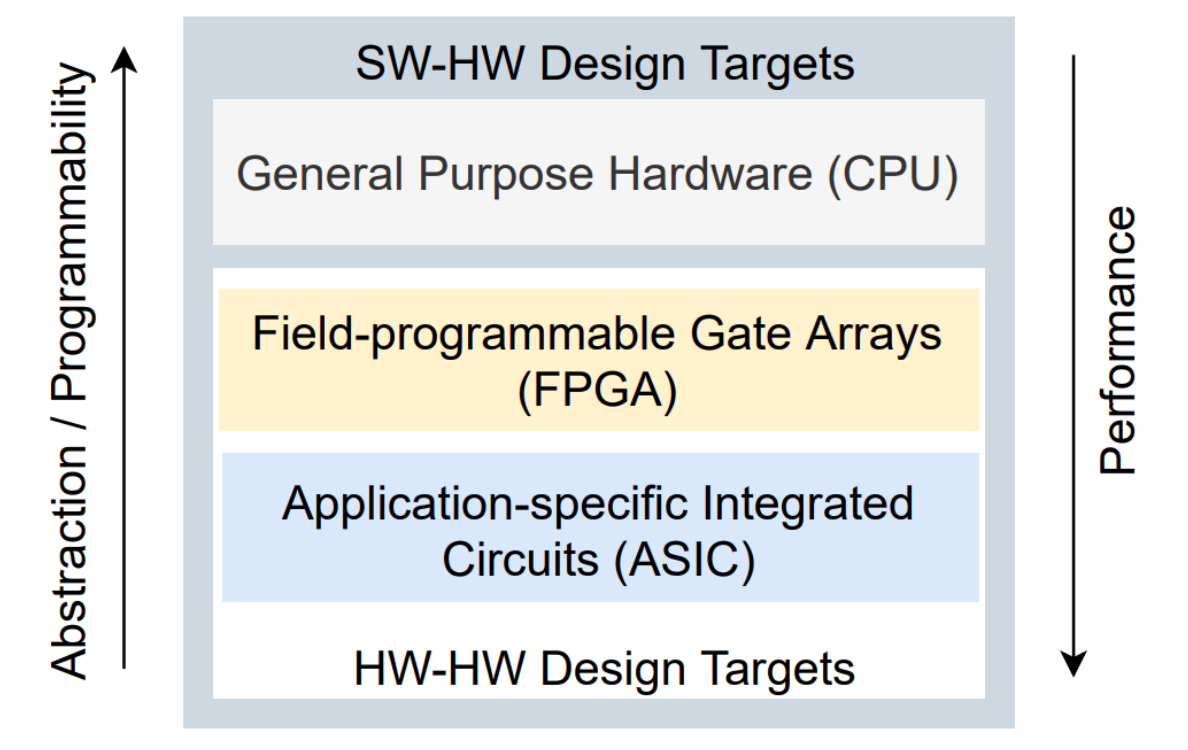

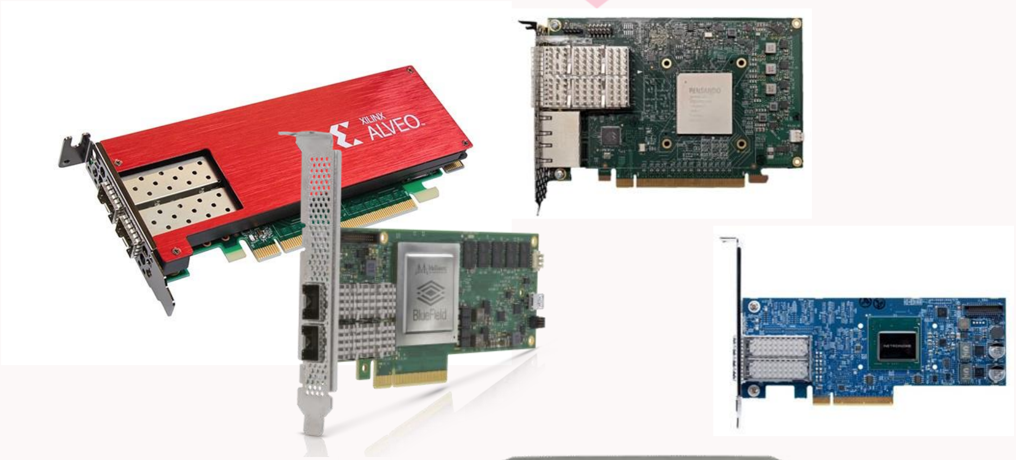

DPDK eBPF P4

PROGRAMMABLE NETWORK DEVICES AS NF(V) BACKENDS

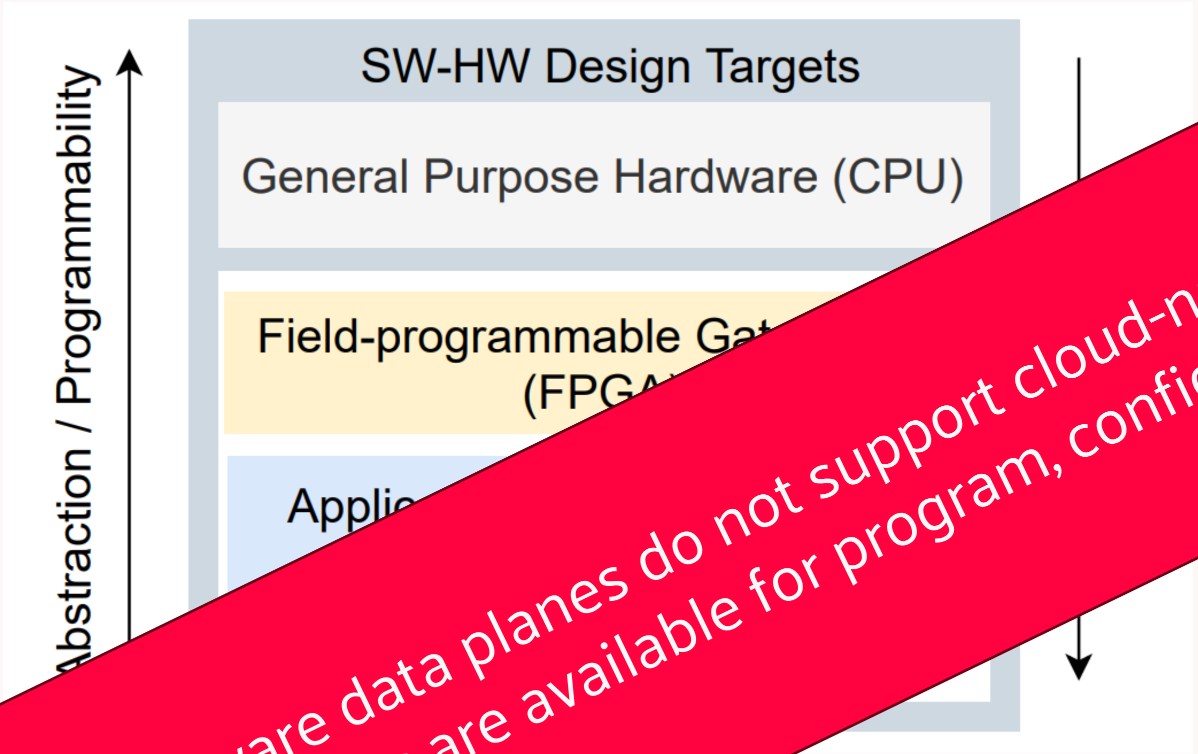

Hardware data planes do not support cloud-native integration. Only low level tools and APIs are available for program, configure and deploy network functions.

eBPF

E2E PROGRAMMABILITY VISION

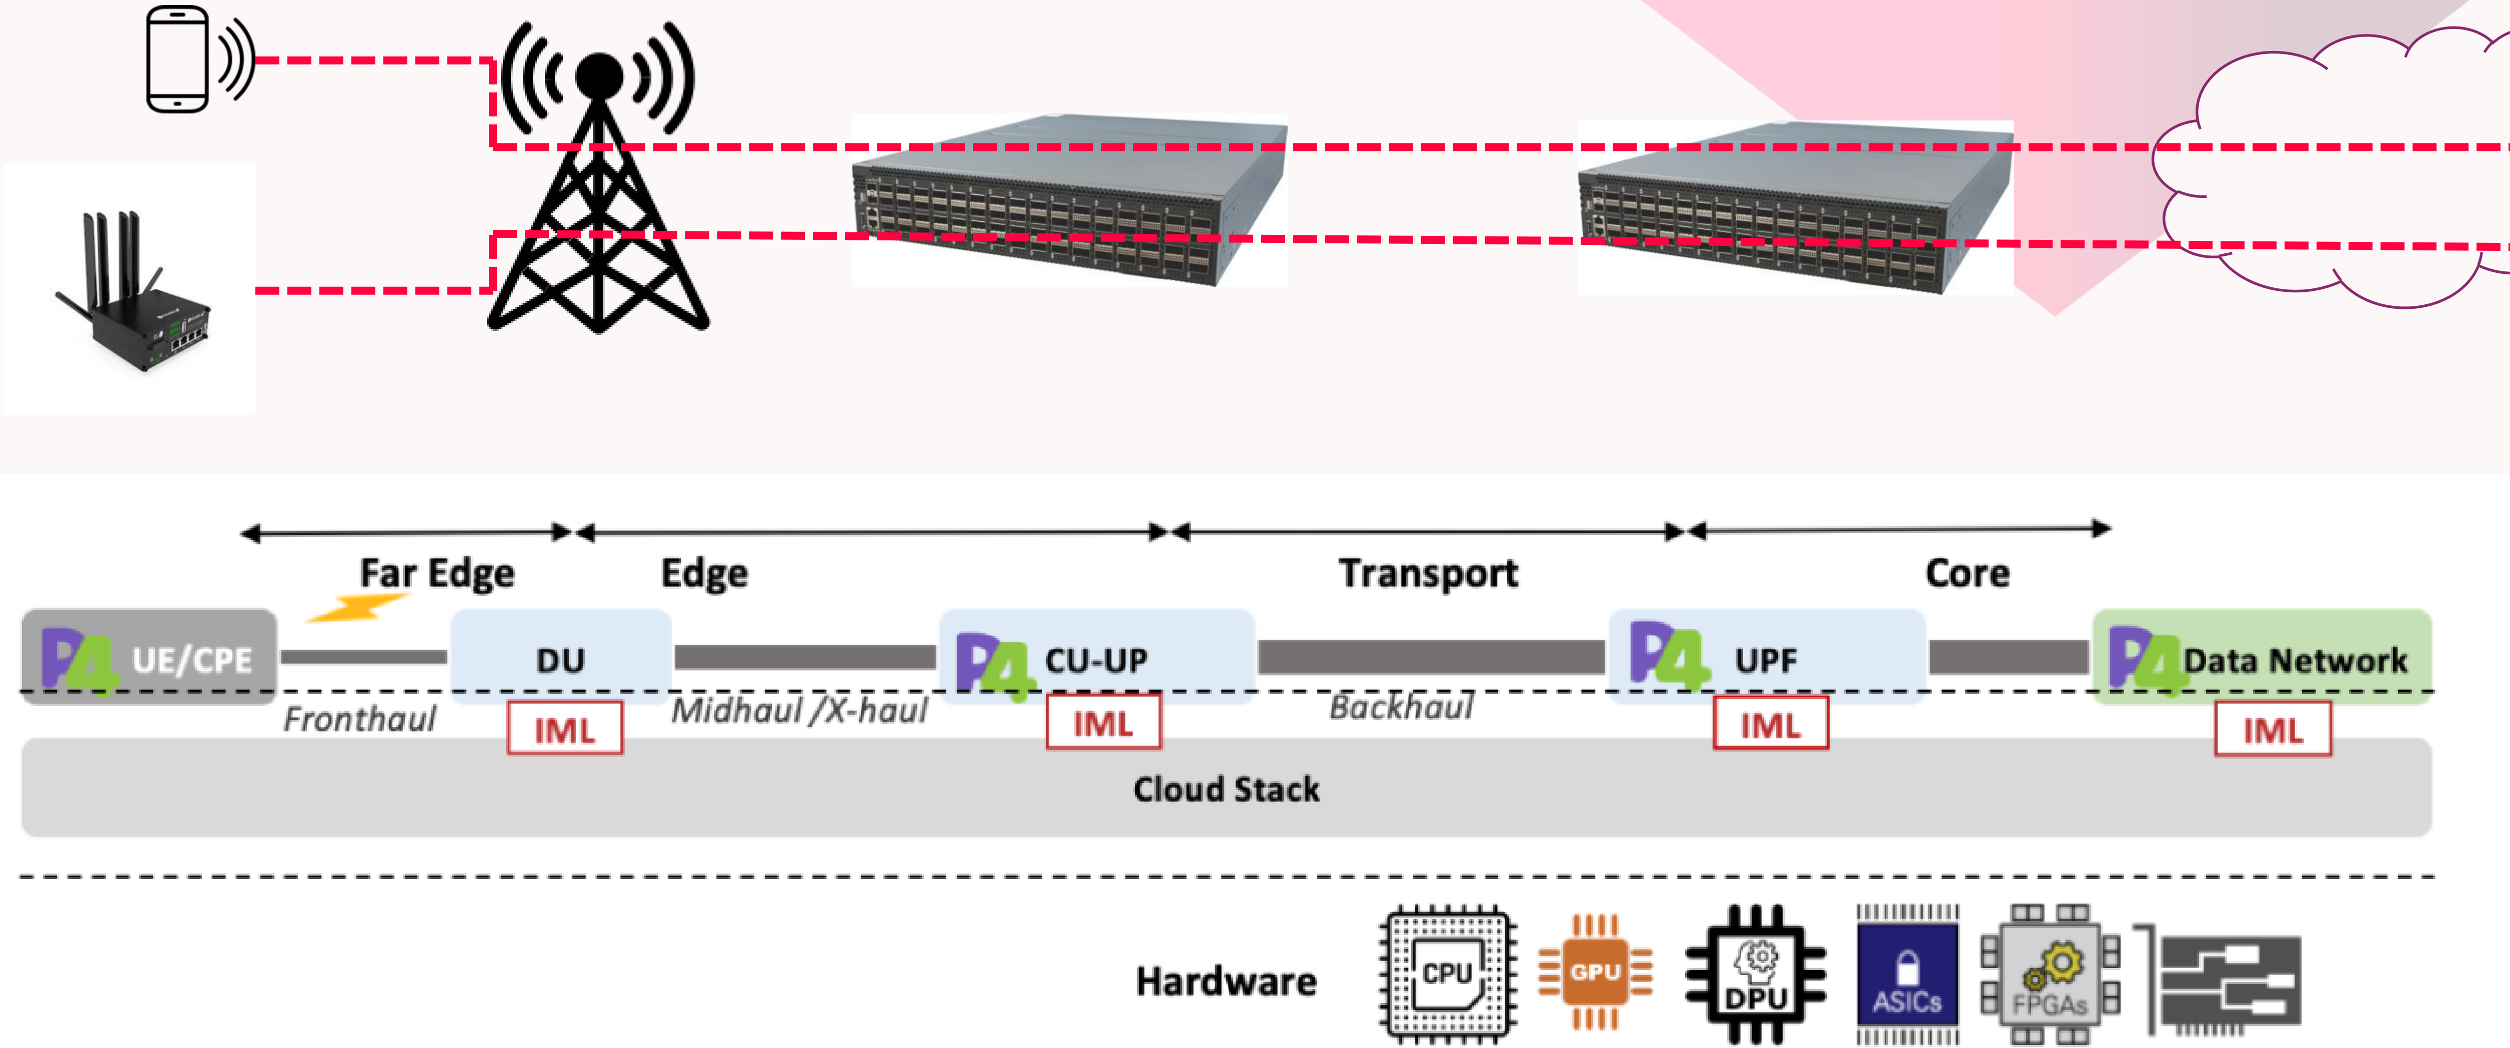

SHARED INFRASTRUCTURE

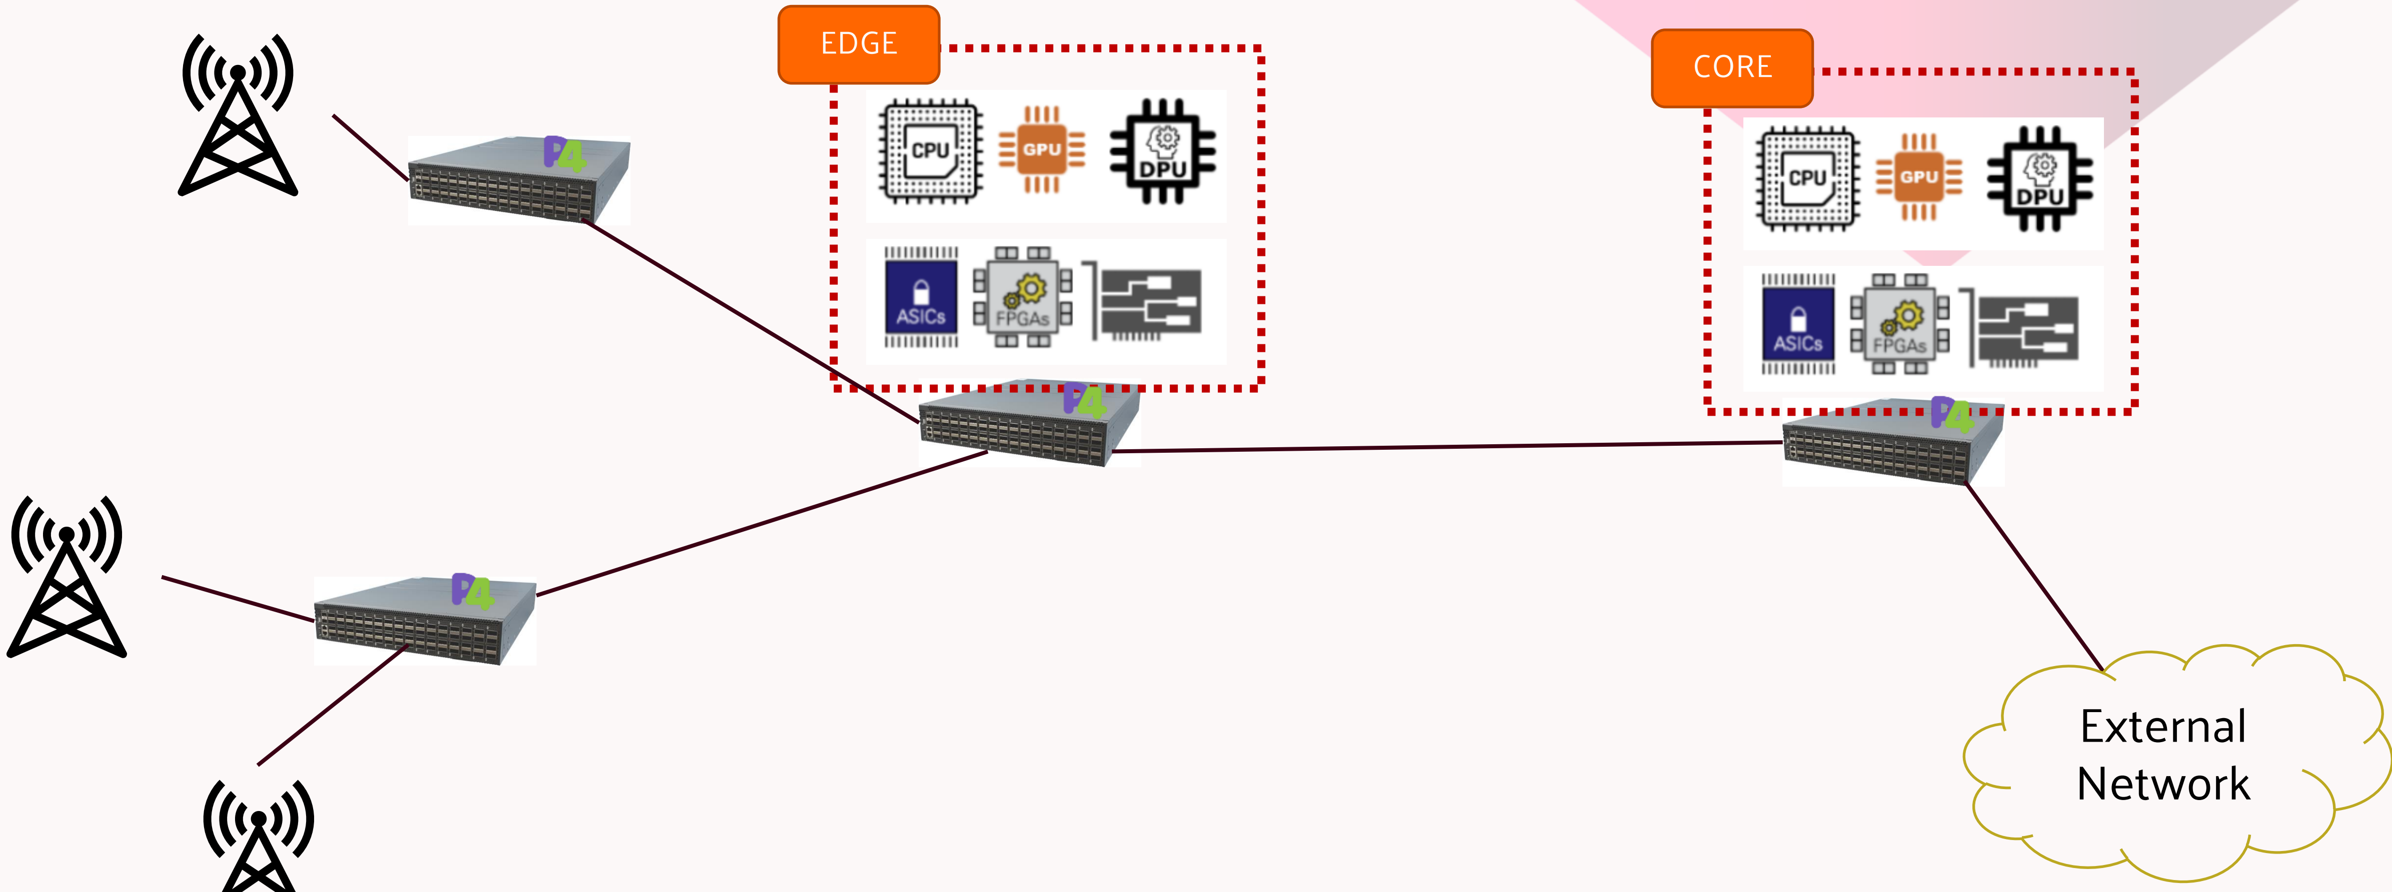

DESIRE6G

SHARED INFRASTRUCTURE

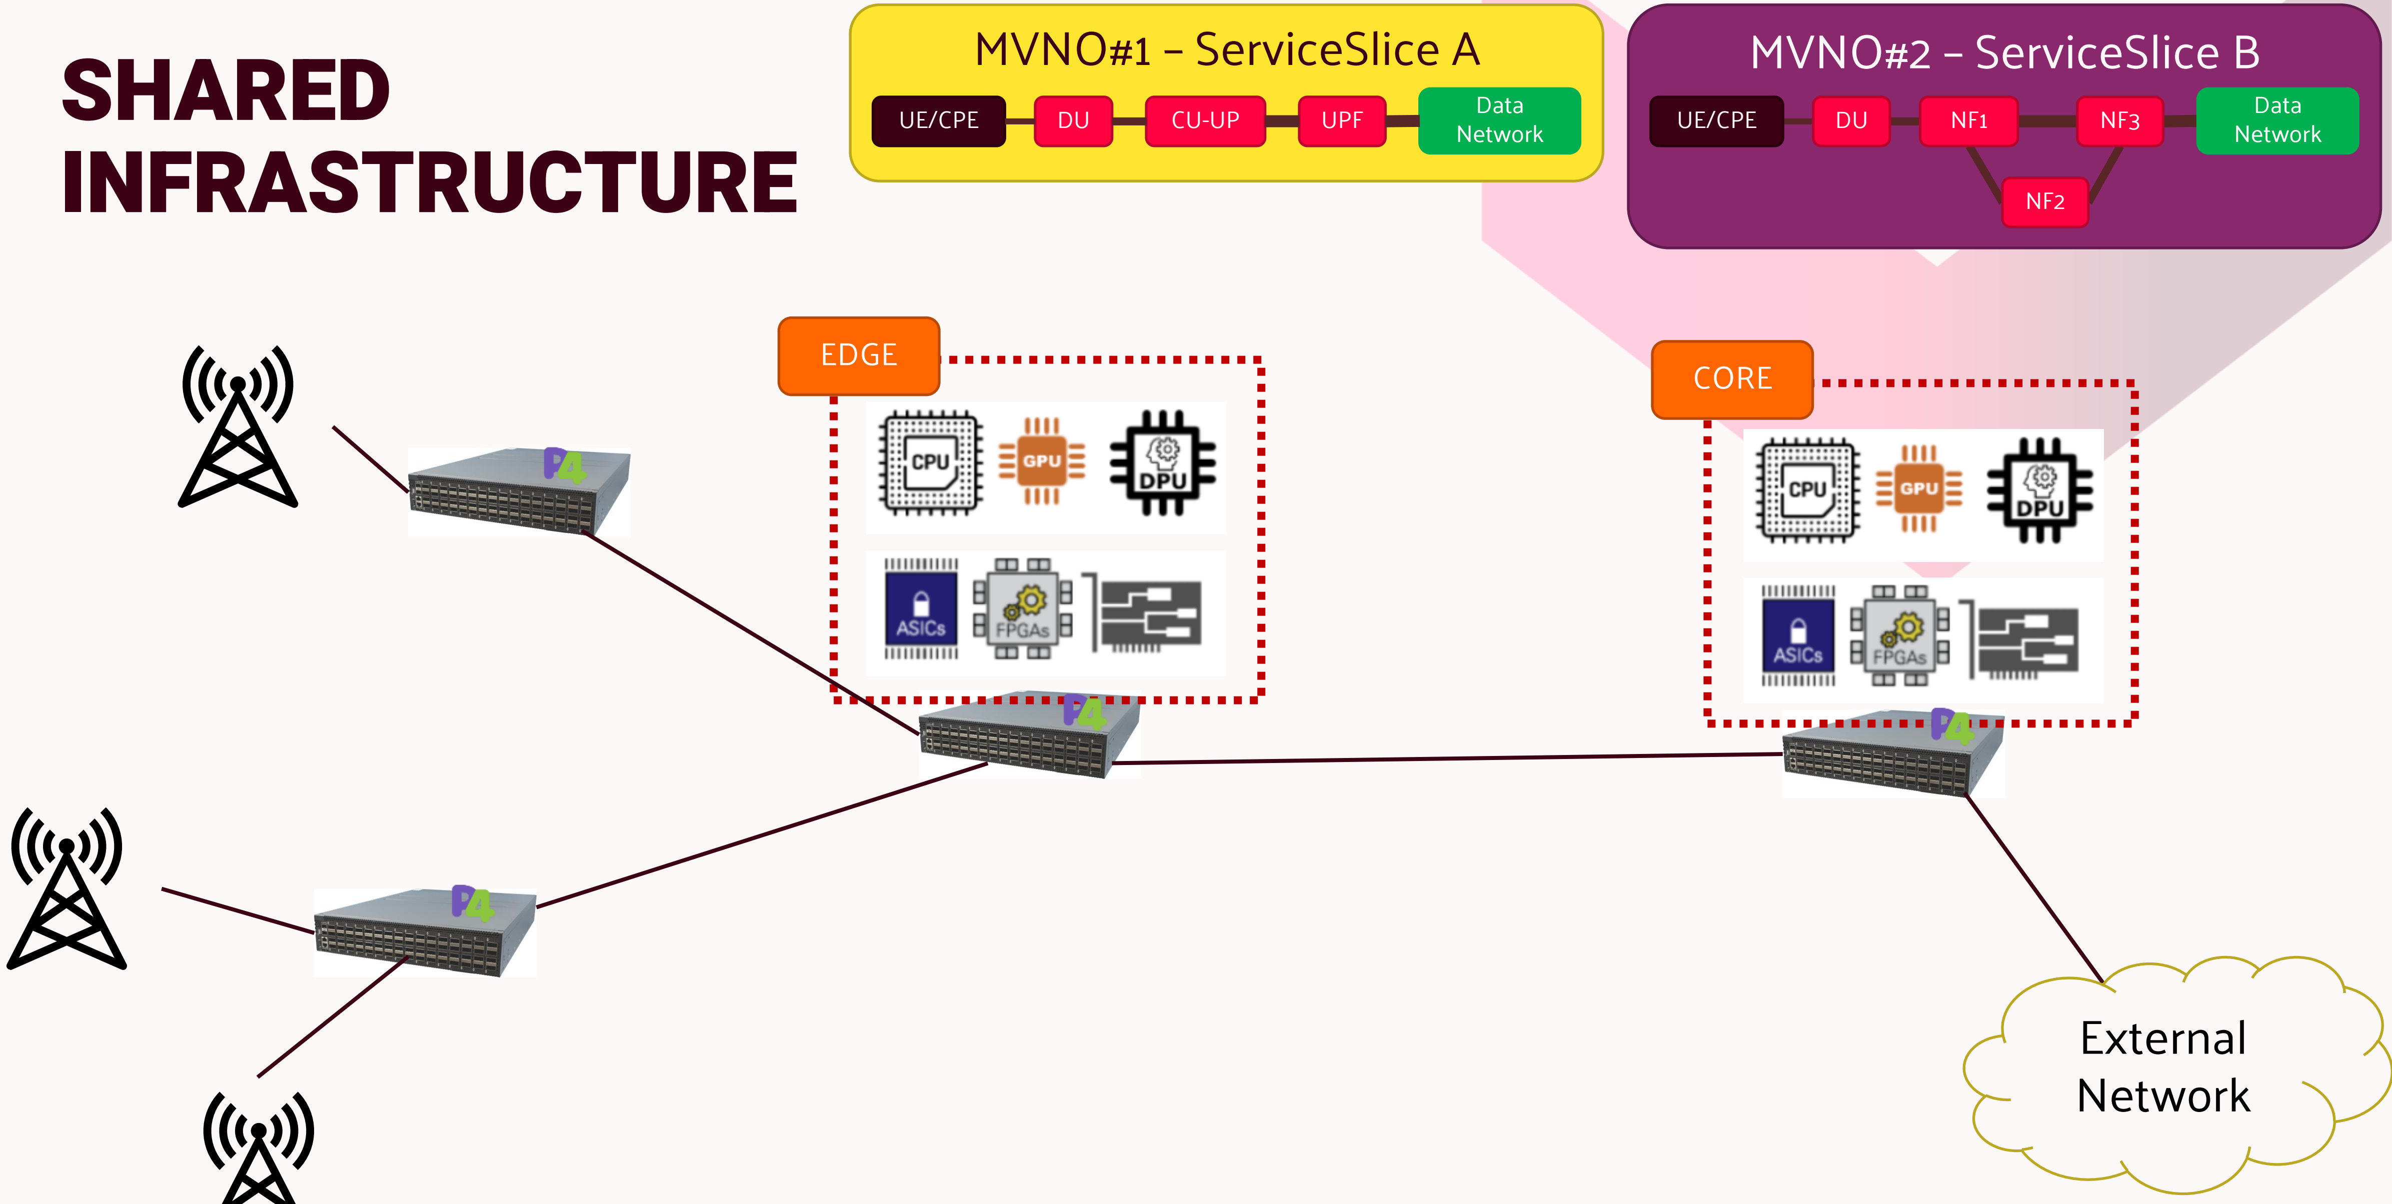

DESIRE6G

SHARED INFRASTRUCTURE

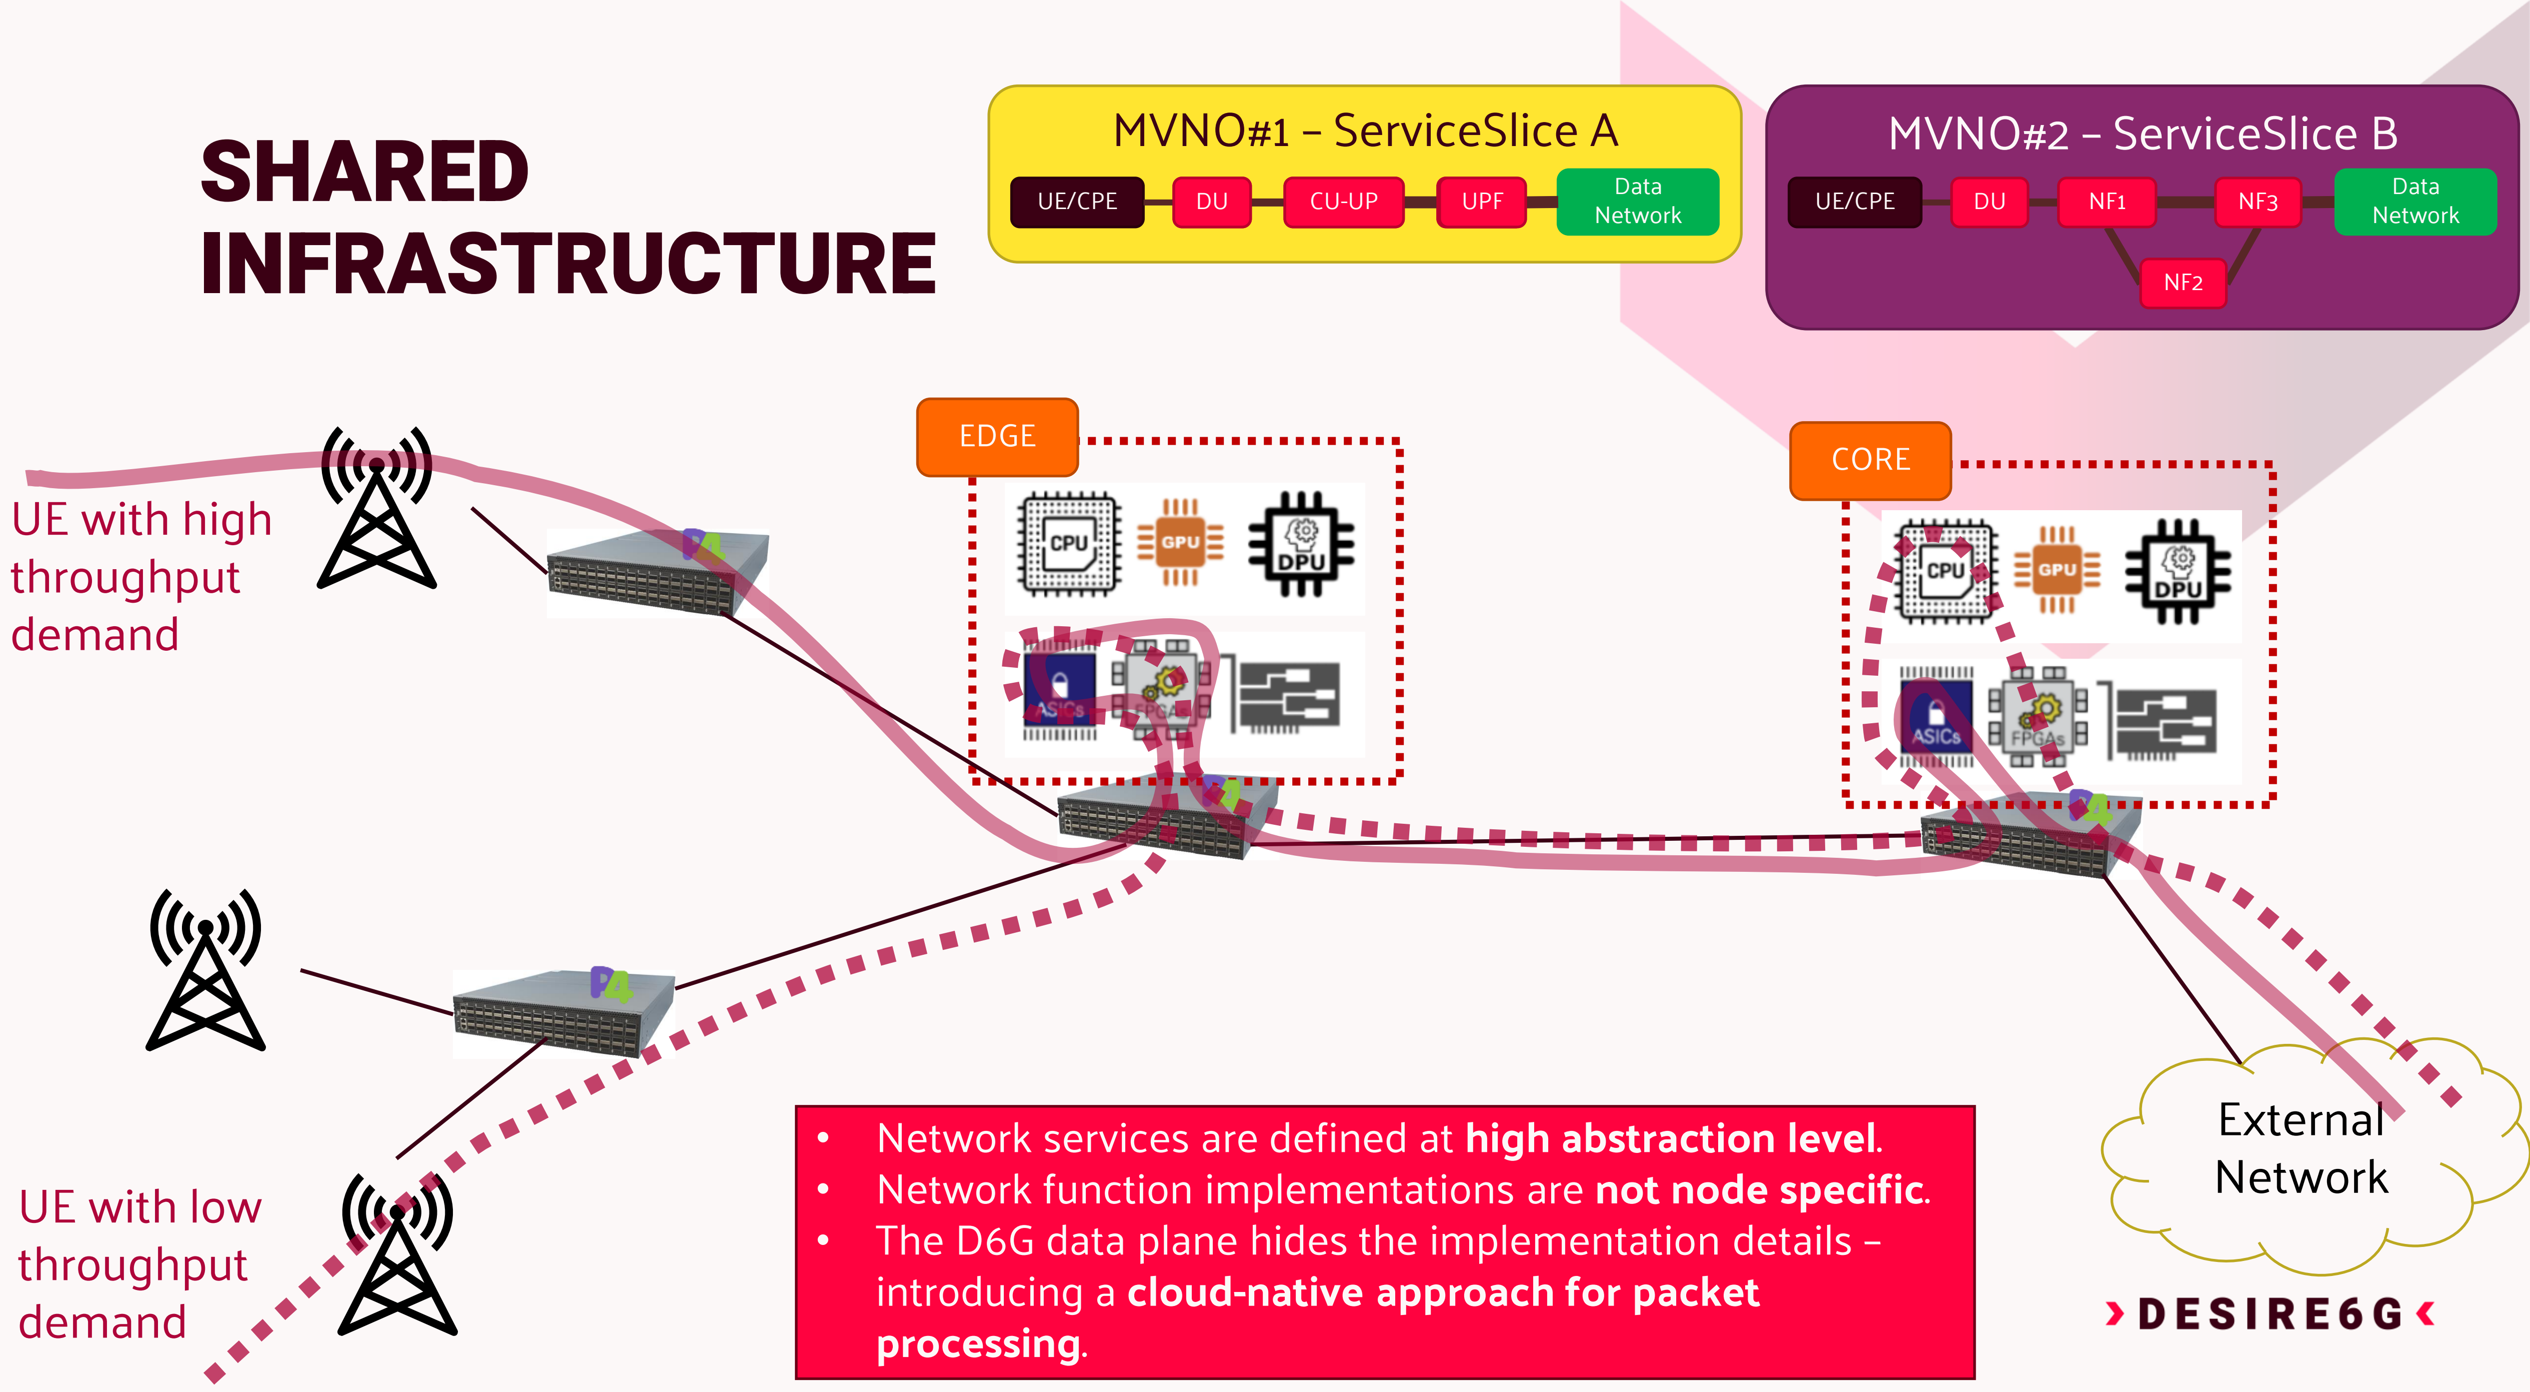

- Network services are defined at **high abstraction level**.
- Network function implementations are **not node specific**.
- The D6G data plane hides the implementation details – introducing a **cloud-native approach for packet processing**.

NETWORK SERVICE (AKA SLICE): INSTANCE OF A NETWORK FUNCTION GRAPH TEMPLATE

Abstract network function (NF) (or service) graph

- End-to-end packet processing logic of one network service/slice
  - e.g., Internet access, real-time voice/video calls, robot control @MEC, XR communication, DetNet over L3
- Dynamic creation and configuration of network services
  - A user/application can join a network service (i.e., the slice implementing it)
  - Instantiation of service template between the end points
  - Mostly configuration, but redeployment of NFs may also be required
- One graph per direction (UL/DL) – the functionality is not always the same

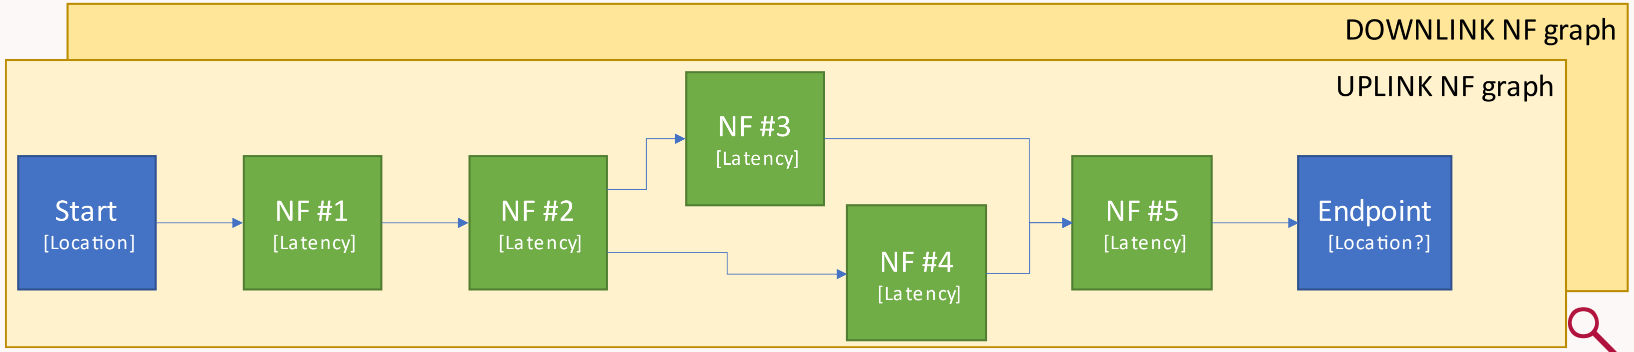

› DESIRE6G ‹

PACKET PROCESSING NETWORK FUNCTIONS: CP+DP

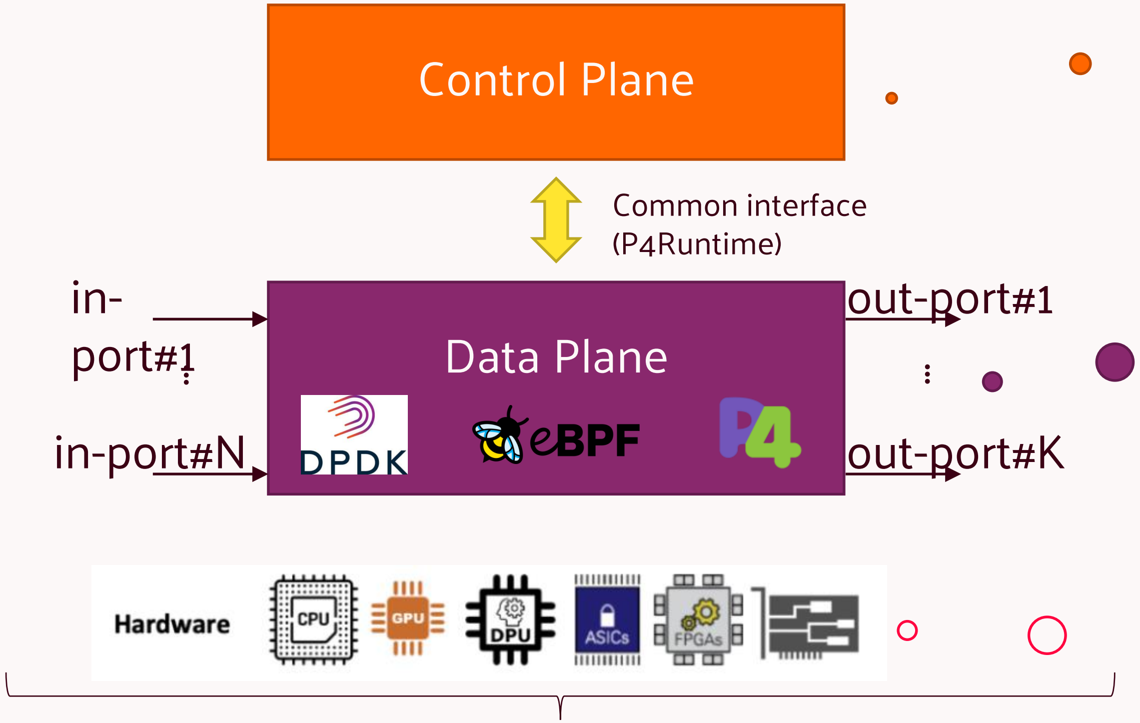

Running on CPU

Different equivalent NF-DP implementations (e.g., P4, eBPF/XDP, DPDK) for the same DP functionality

Different implementations can be executed on various targets

NFs are stored in an NF repository and instantiated by Service Slices
NF-DP implementation properties include

- Execution latency on specific target config
- Max. bitrate/packet rate capacity
- Max. number of Ues to be handled
- …

A CLOUD-NATIVE DATA PLANE MANAGER

- Infrastructure Management Layer (IML)
  - A **cloud-native data plane manager**
- Provides a **simple logical view** of the data plane to control planes
- Ensures **service/slice requirements**
- Hides the underlying **implementation and optimization** details
  - Load balancing
  - Heavy hitter handling
  - Auto-scaling
  - HW offloading/acceleration
  - HW multitenancy

Available resources to deploy on

- Initial resource view for our example

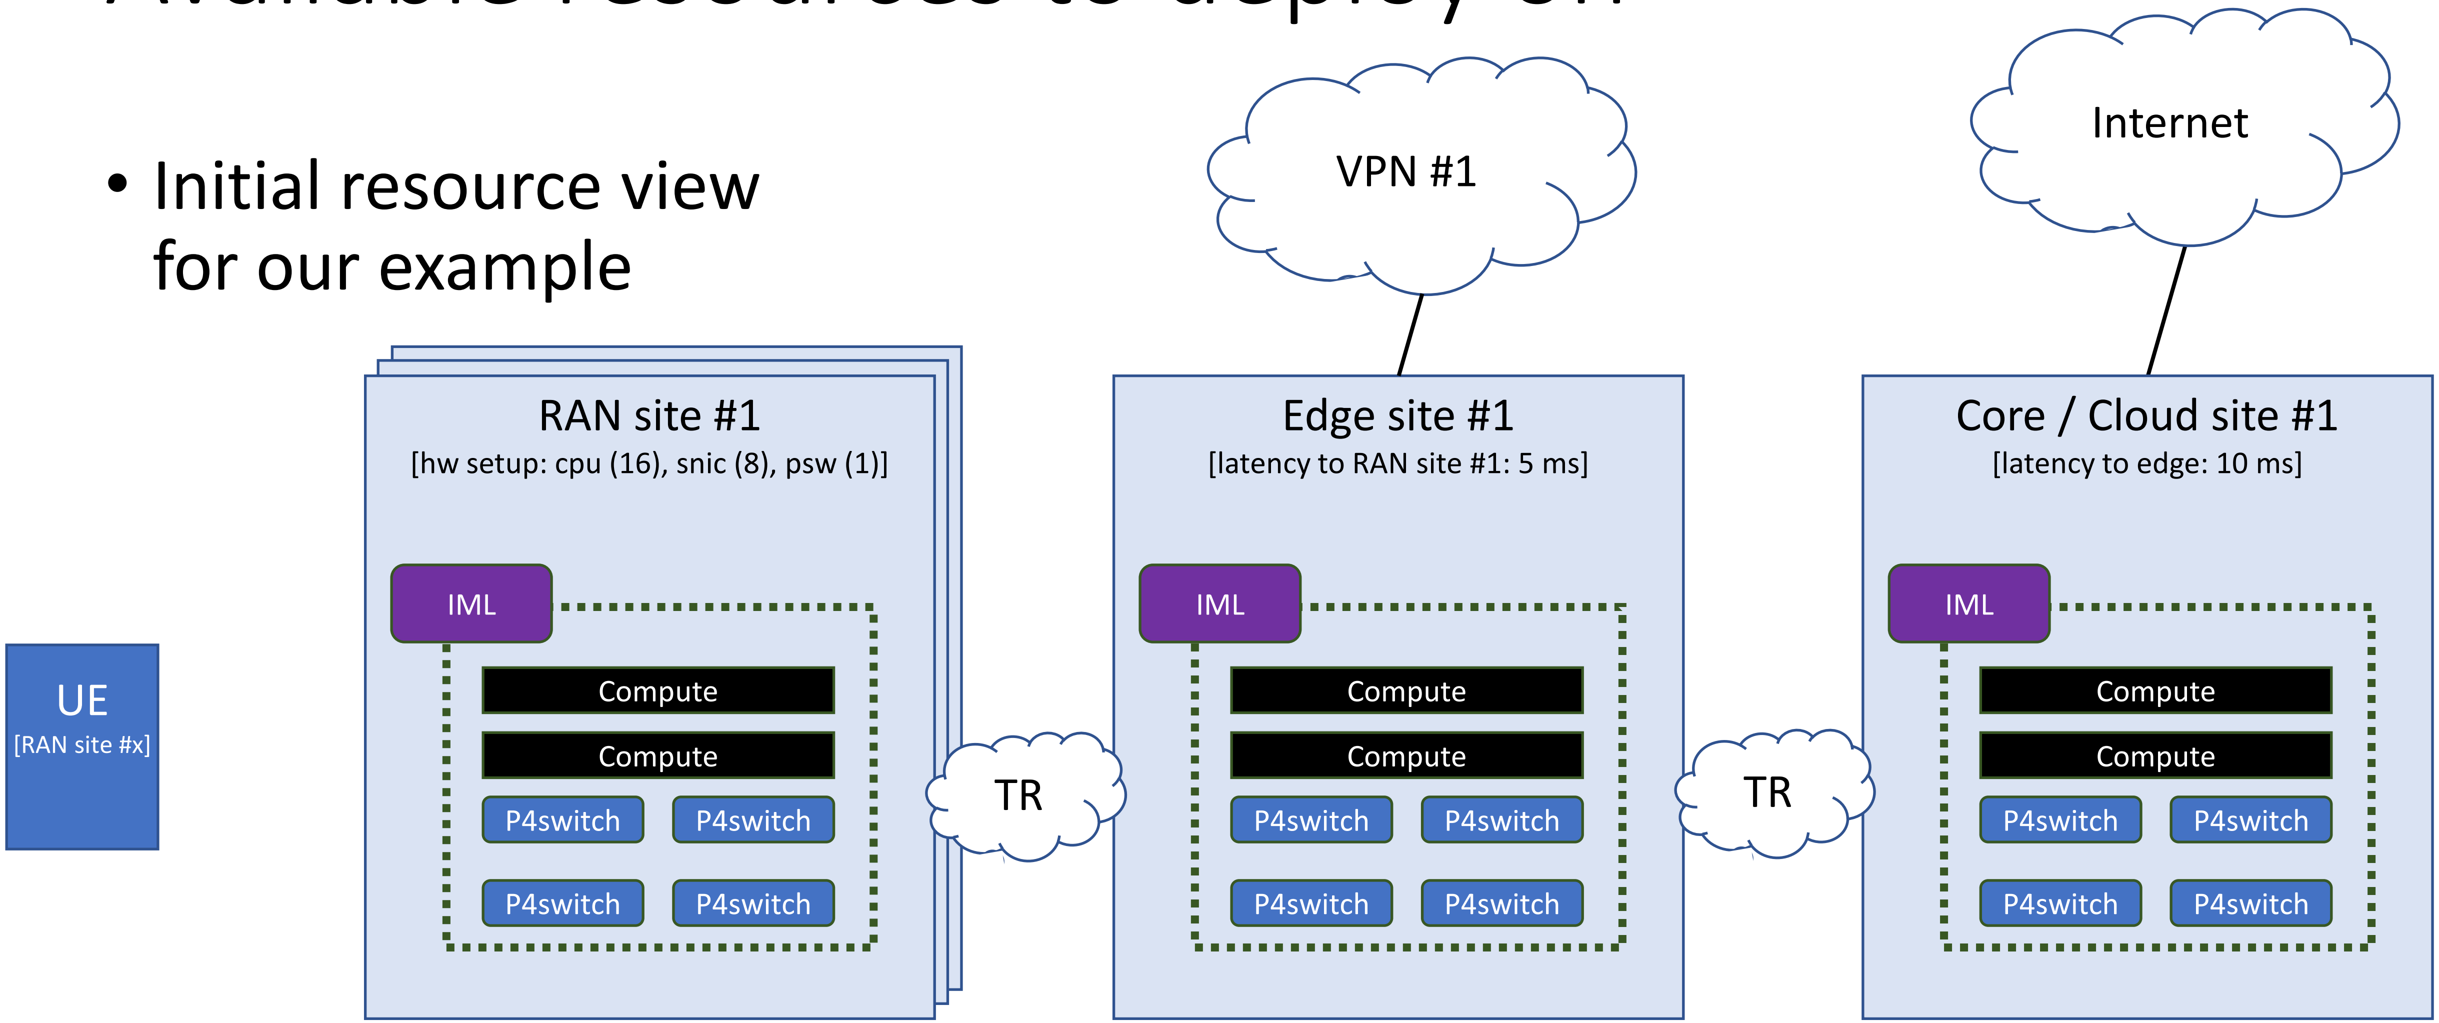

IML and service deployment

- During service deployment IML can (sometimes must) add further NFs to the graph
  - Transport adapters: adapt to transport between two sites
    - Probably via non-programmable devices – we need to connect these domains
  - Network slicing (both separation and QoS)
  - Load balancing for a given NF or graph fragment
  - Heavy-hitter pattern (kind of load balancing)

} Transparent optimizations for NFs

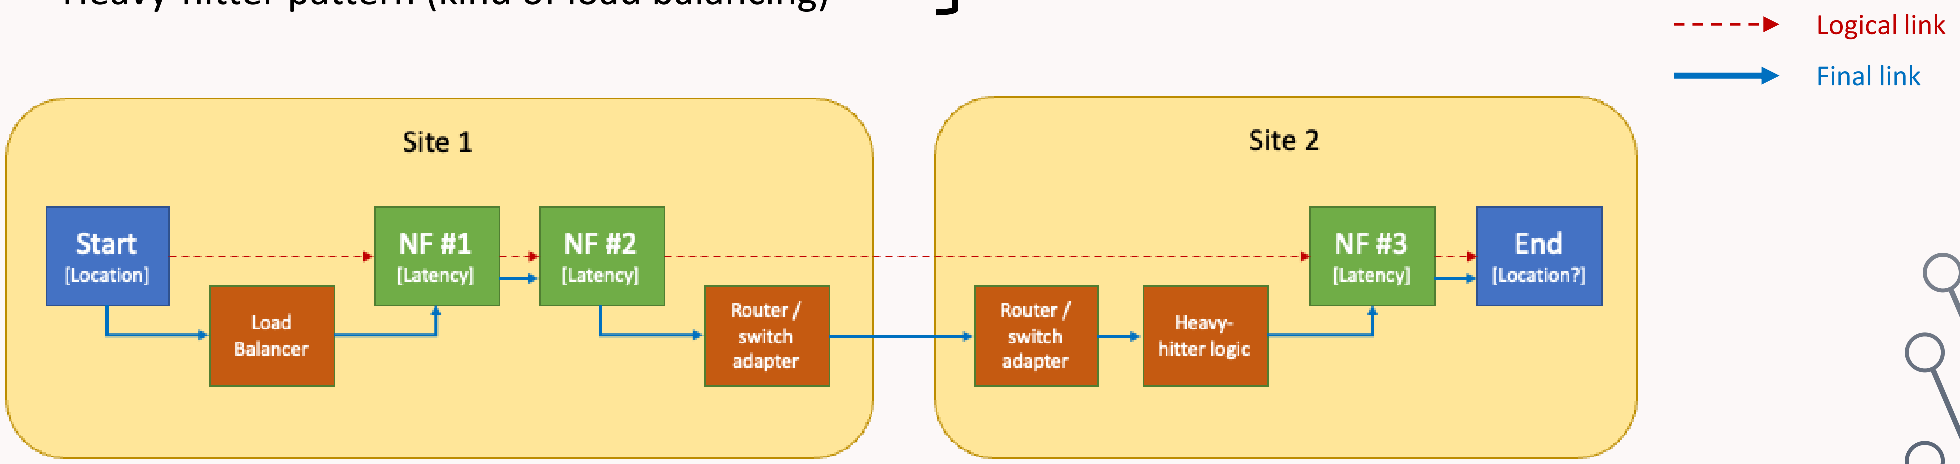

IML: the physical deployment view (site 1)

color codes: service NF, infra NF, static function

- Map logical functions to physical (workers, hardware)
- Map virtual ports to physical ports

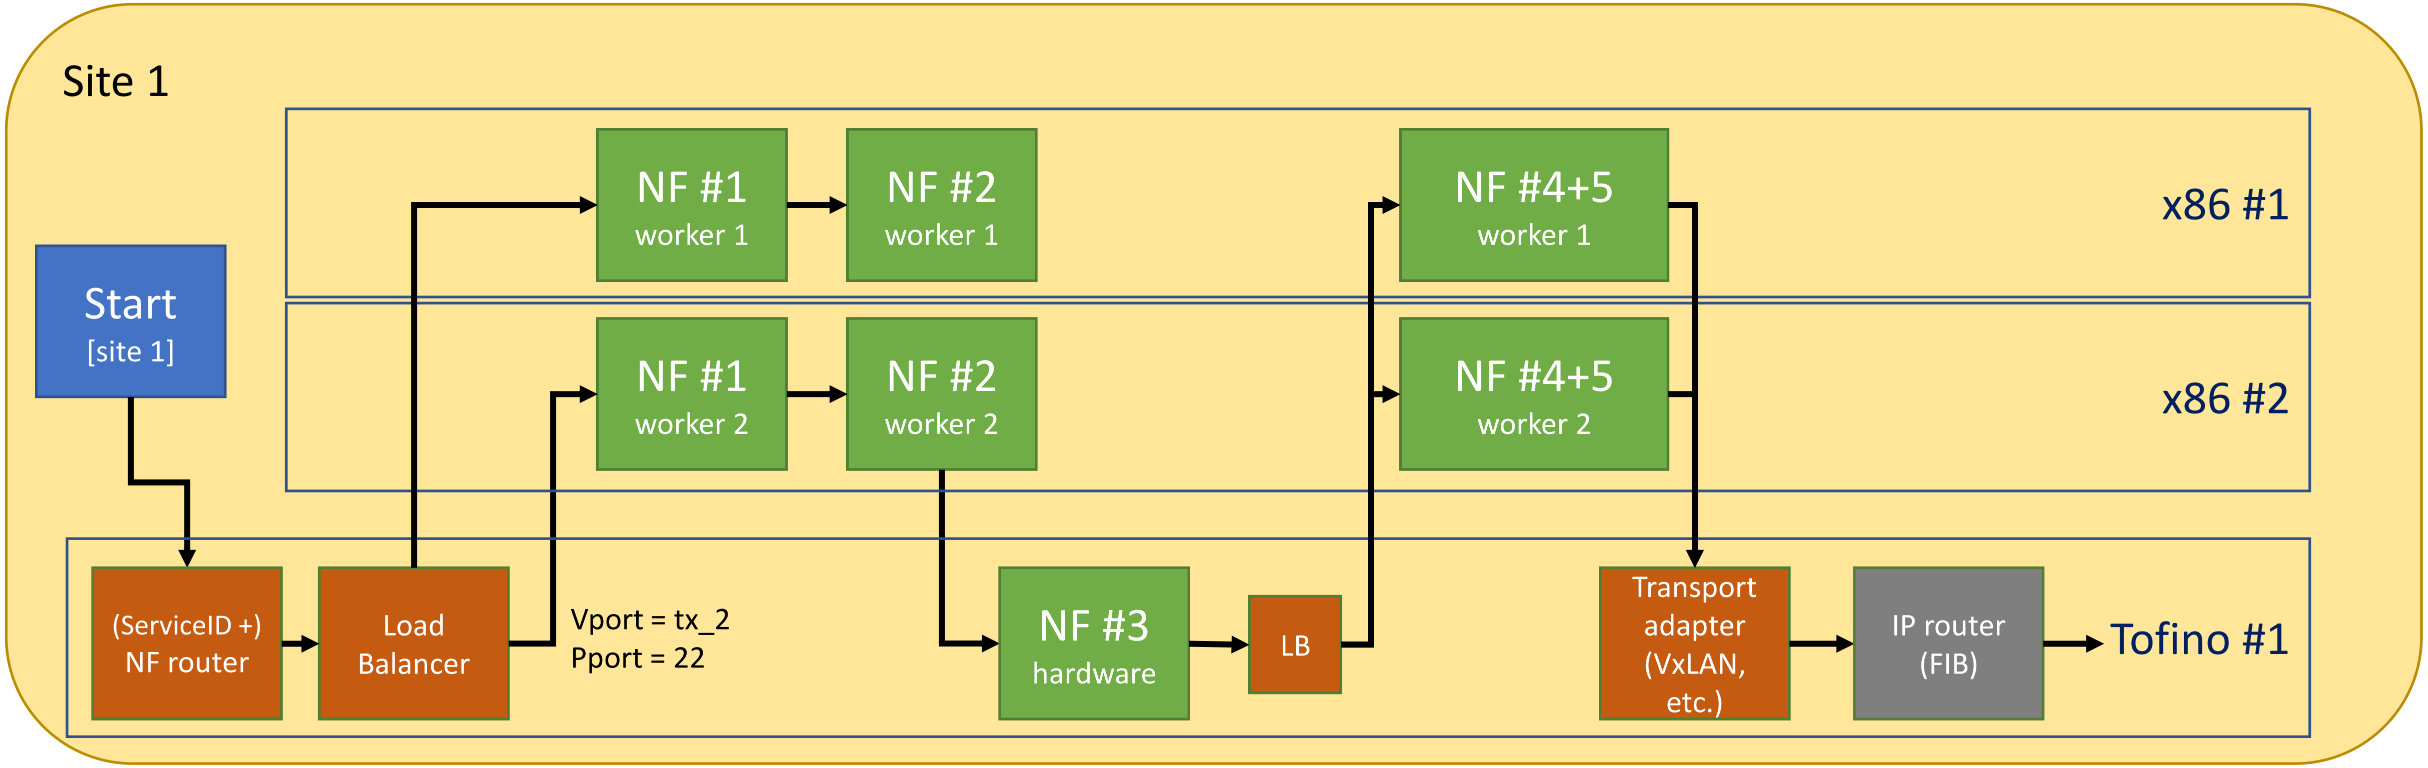

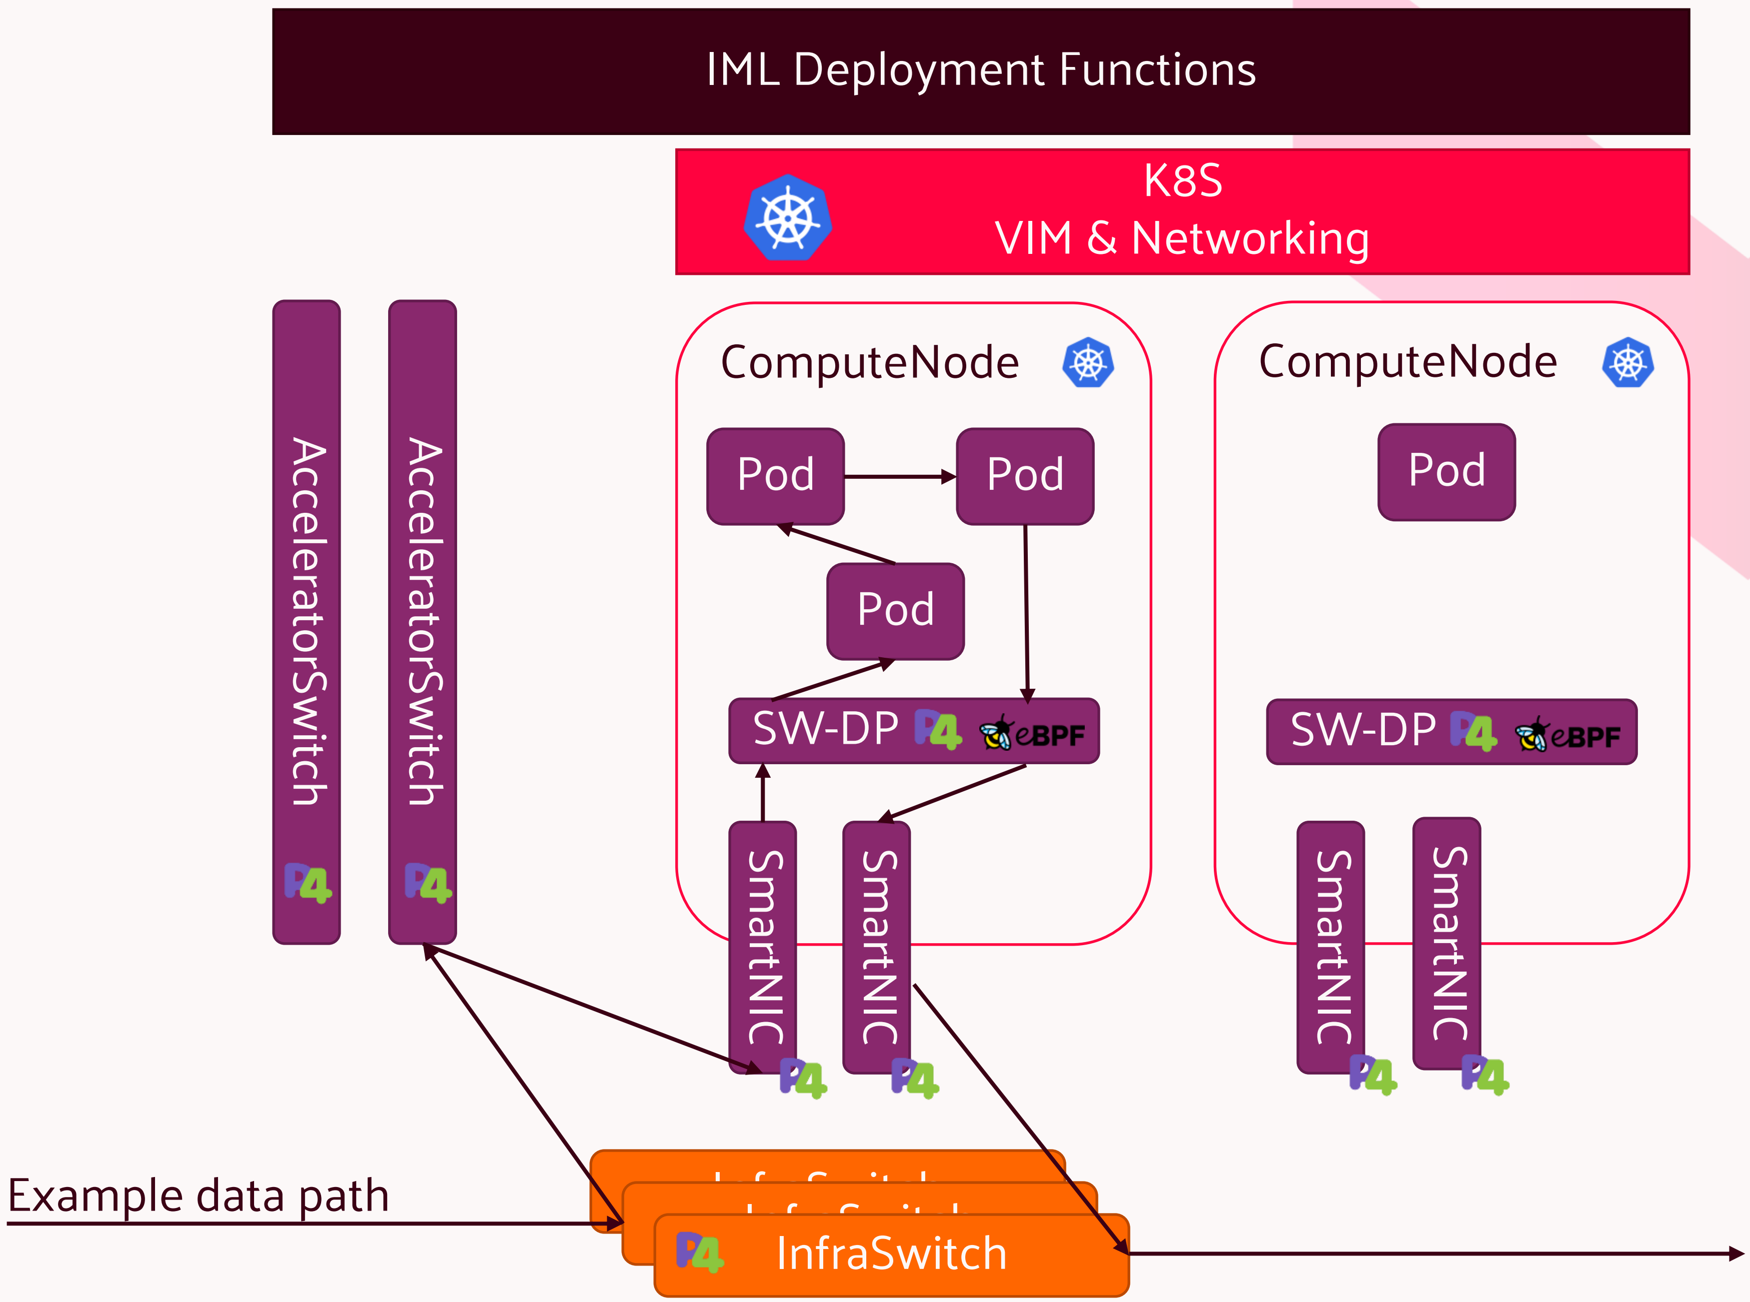

IML Deployment Functions
K8S
VIM & Networking
ComputeNode
Pod
Pod
Pod
SW-DP
eBPF
ComputeNode
Pod
SW-DP
eBPF
AcceleratorSwitch
AcceleratorSwitch
SmartNIC
SmartNIC
SmartNIC
SmartNIC
InfraSwitch
Example data path
DESIRE6G

SLICING REQUIREMENTS AGAINST DATA PLANE

- Resource isolation between service slices
  - Requires multi-tenant support for NF deployment on dedicated PDP HW
- Security isolation
  - Access control between data plane objects and control plane components
- QoS/Performance isolation between slices and subslices
  - Includes routing, traffic management and load balancing implemented by PDP
  - Fine grained and on demand settings
  - SLA enforcement with runtime optimization
- Pervasive monitoring for SLA assurance
  - Fast reaction to failures and performance degradation

QOS SLICES IN D6G (~ETSI ZSM)

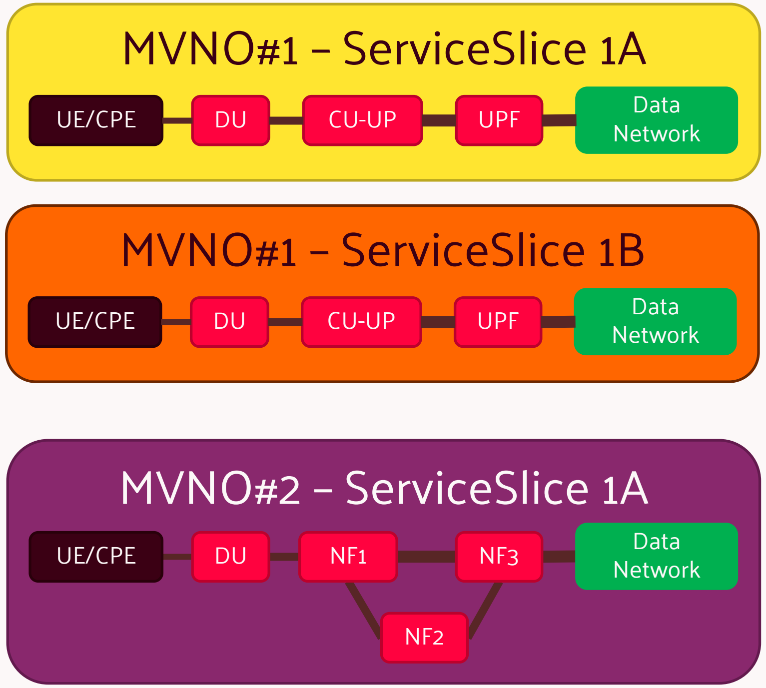

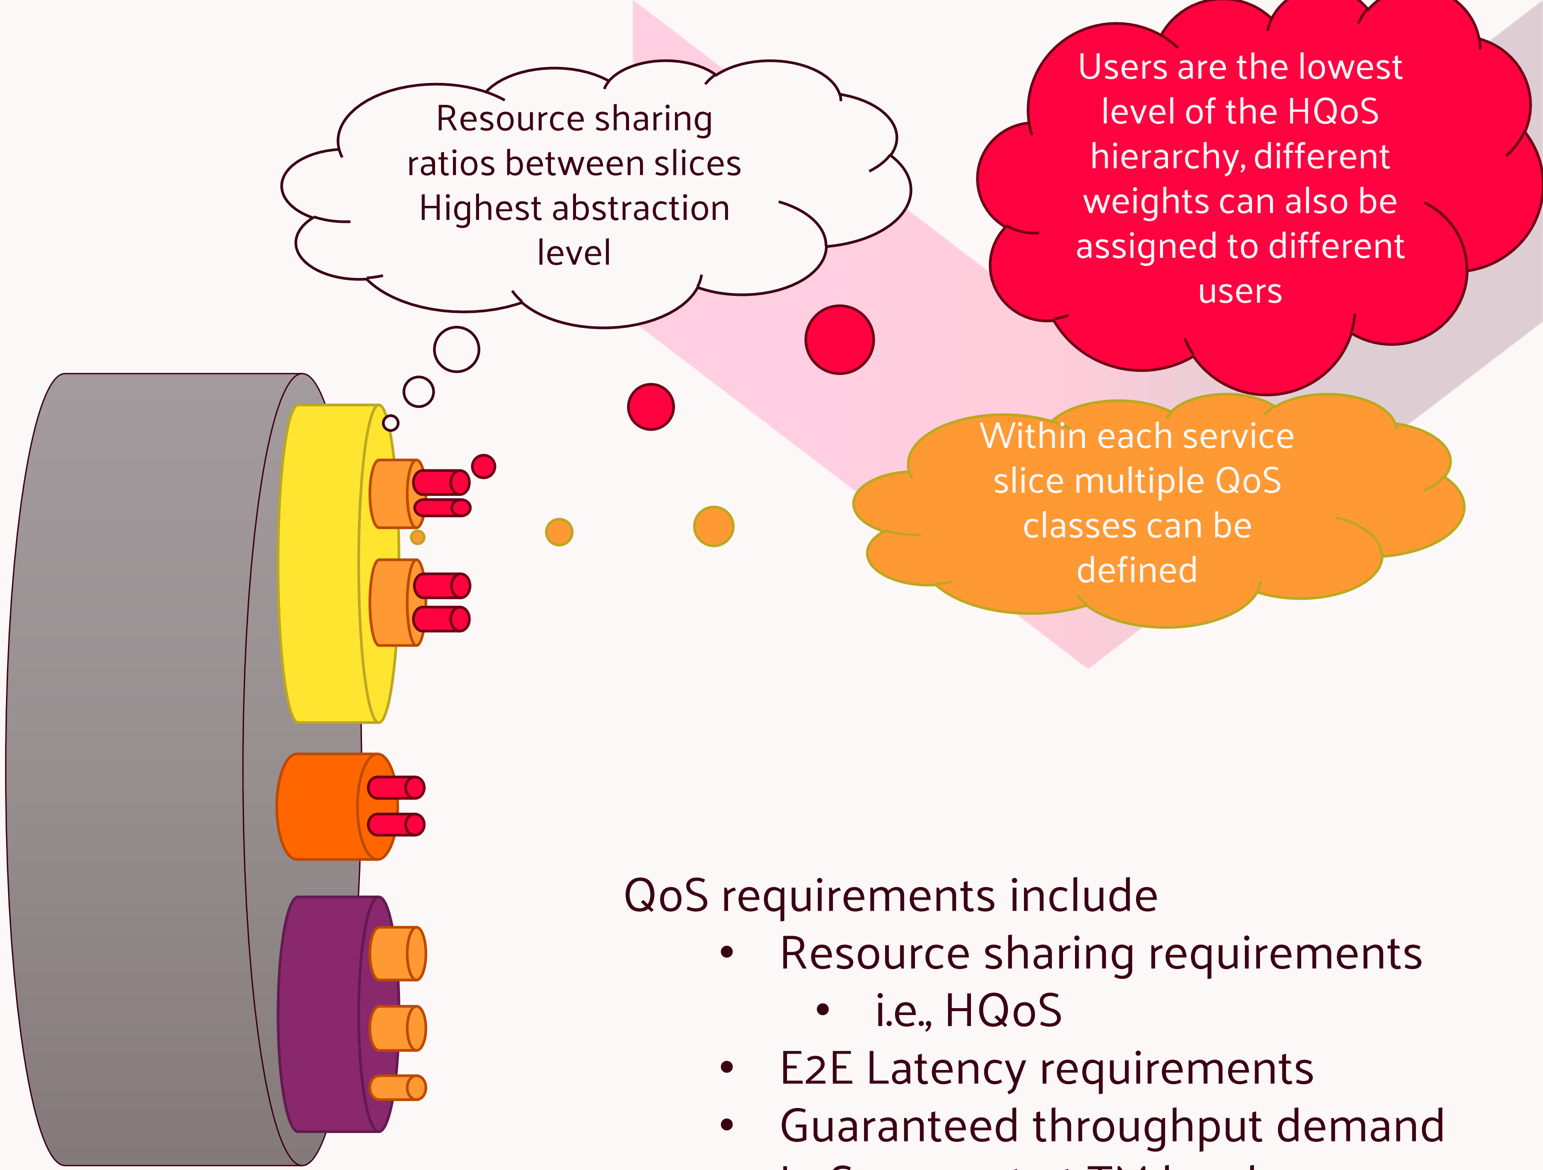

QoS requirements include

- Resource sharing requirements
  - i.e., HQoS
- E2E Latency requirements
- Guaranteed throughput demand
- L4S support at TM level

DEEP SLICING REQUIREMENTS AGAINST DATA PLANE

- Resource isolation between service slices
  - Implemented by P4 program aggregation and slice-based traffic classification in PDP
- Security isolation

EXAMPLE#1

  - Implemented by a Proxy between the Aggregated Data Plane and Control Plane instances
- Performance isolation between slices and subslices

EXAMPLE#2

  - Implemented by so called InfraNFs: routing, traffic management and load balancing
  - Reconfigurable traffic management and load balancing, self-driving pure data plane solutions
- Pervasive monitoring for SLA assurance
  - Implemented as an in-band network telemetry solution, can notify higher layers if needed
  - QoS/SLA measurement techniques for continous monitoring of the provided services

EXAMPLE1: MULTITENANCY SUPPORT ON P4 TARGETS

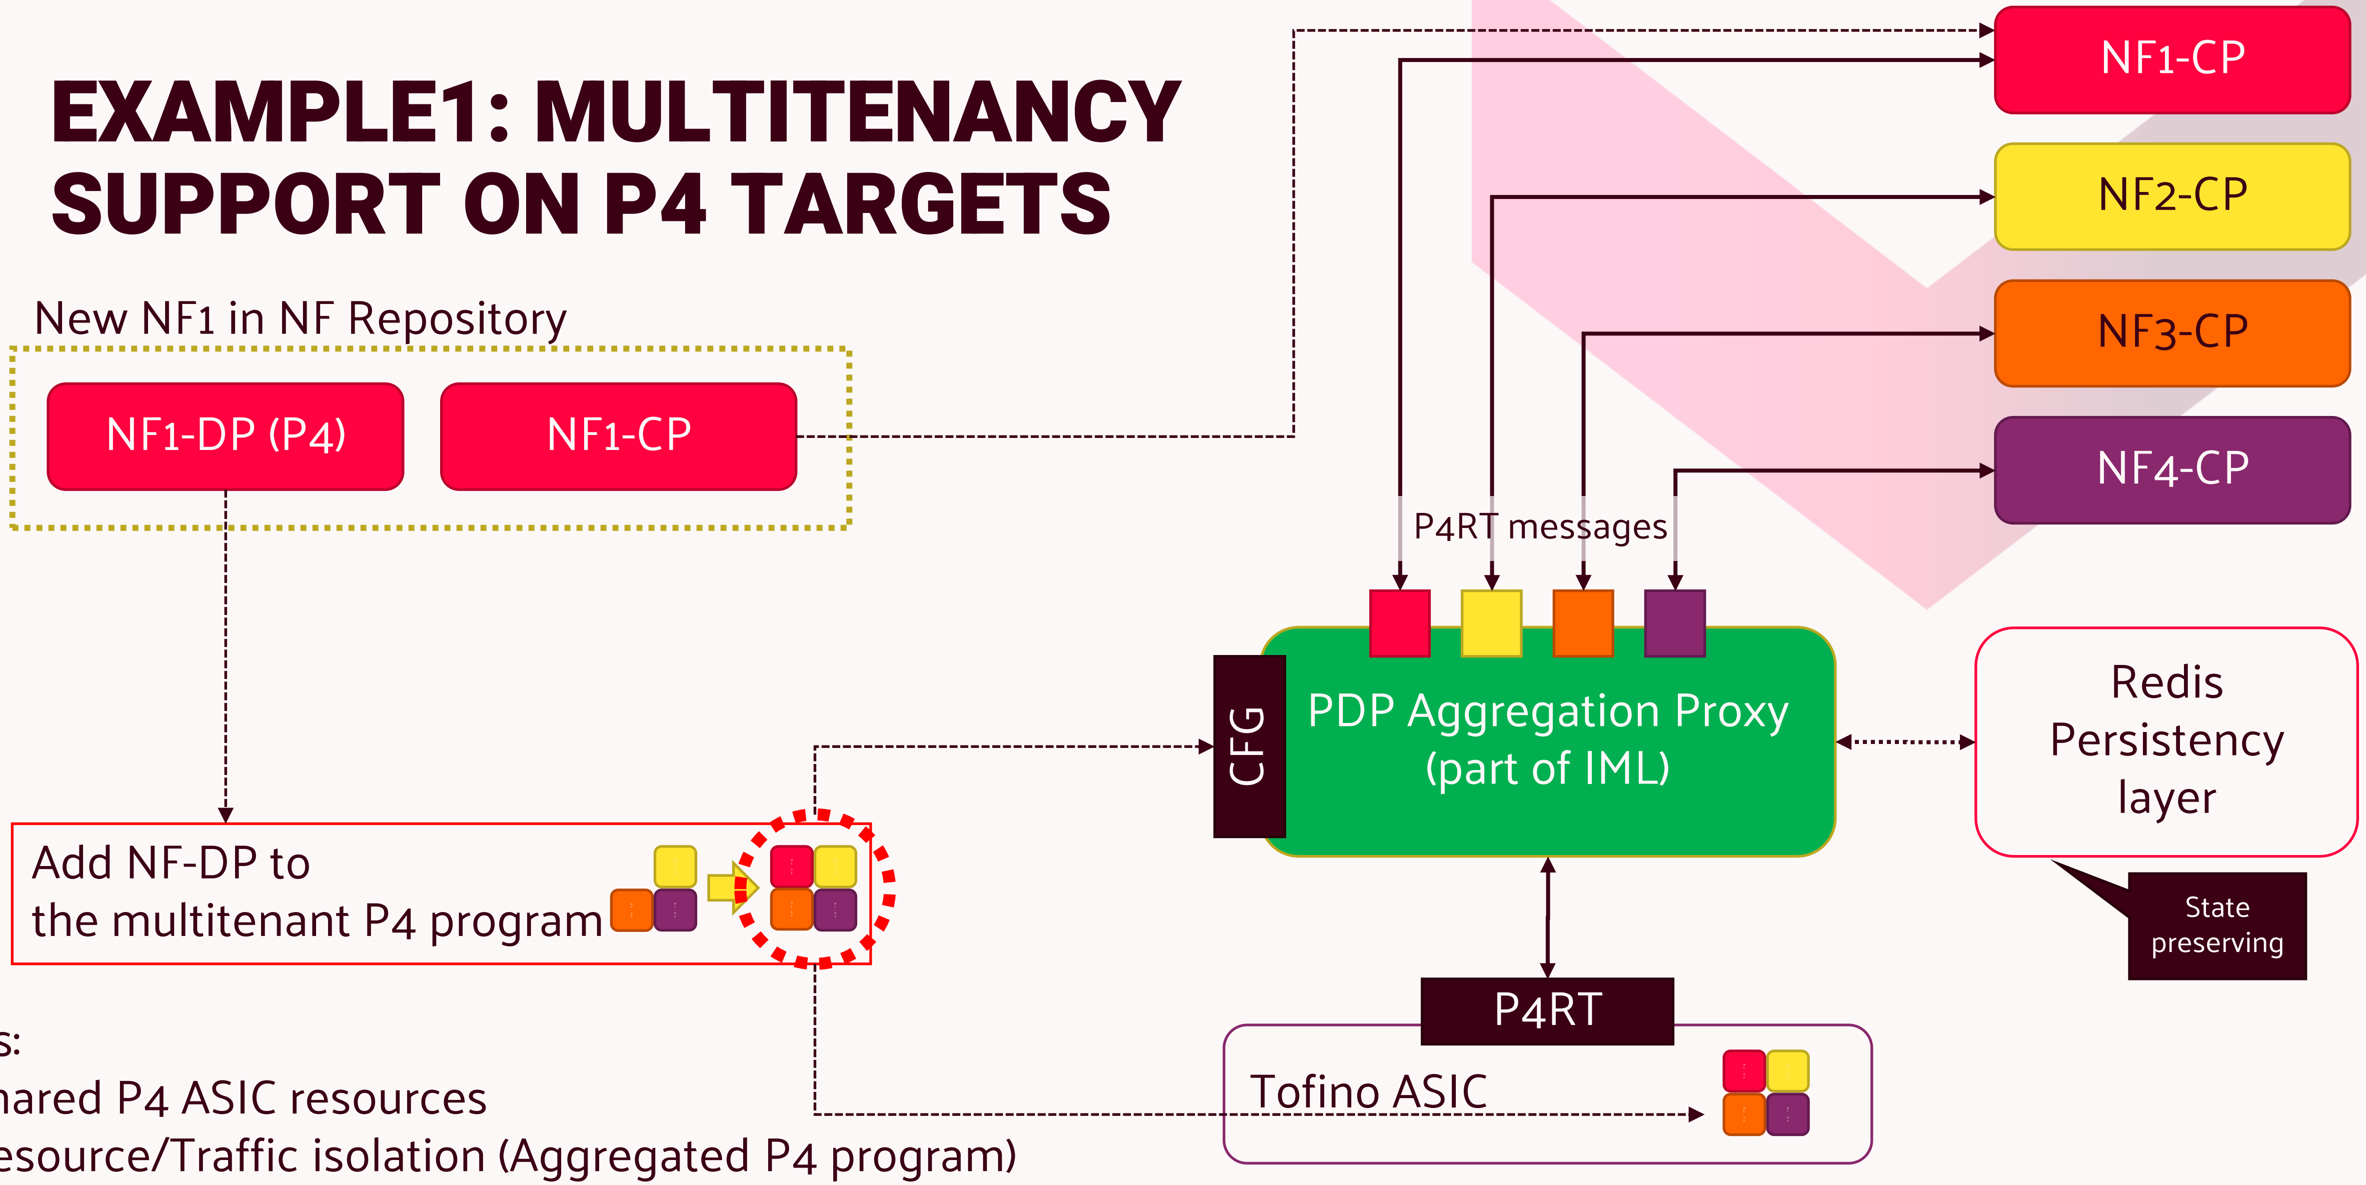

Goals:

- Shared P4 ASIC resources
- Resource/Traffic isolation (Aggregated P4 program)
- Security isolation (PDP Aggregation Proxy)

EXAMPLE2: SEAMLESS LOAD BALANCING/OPTIMIZATION

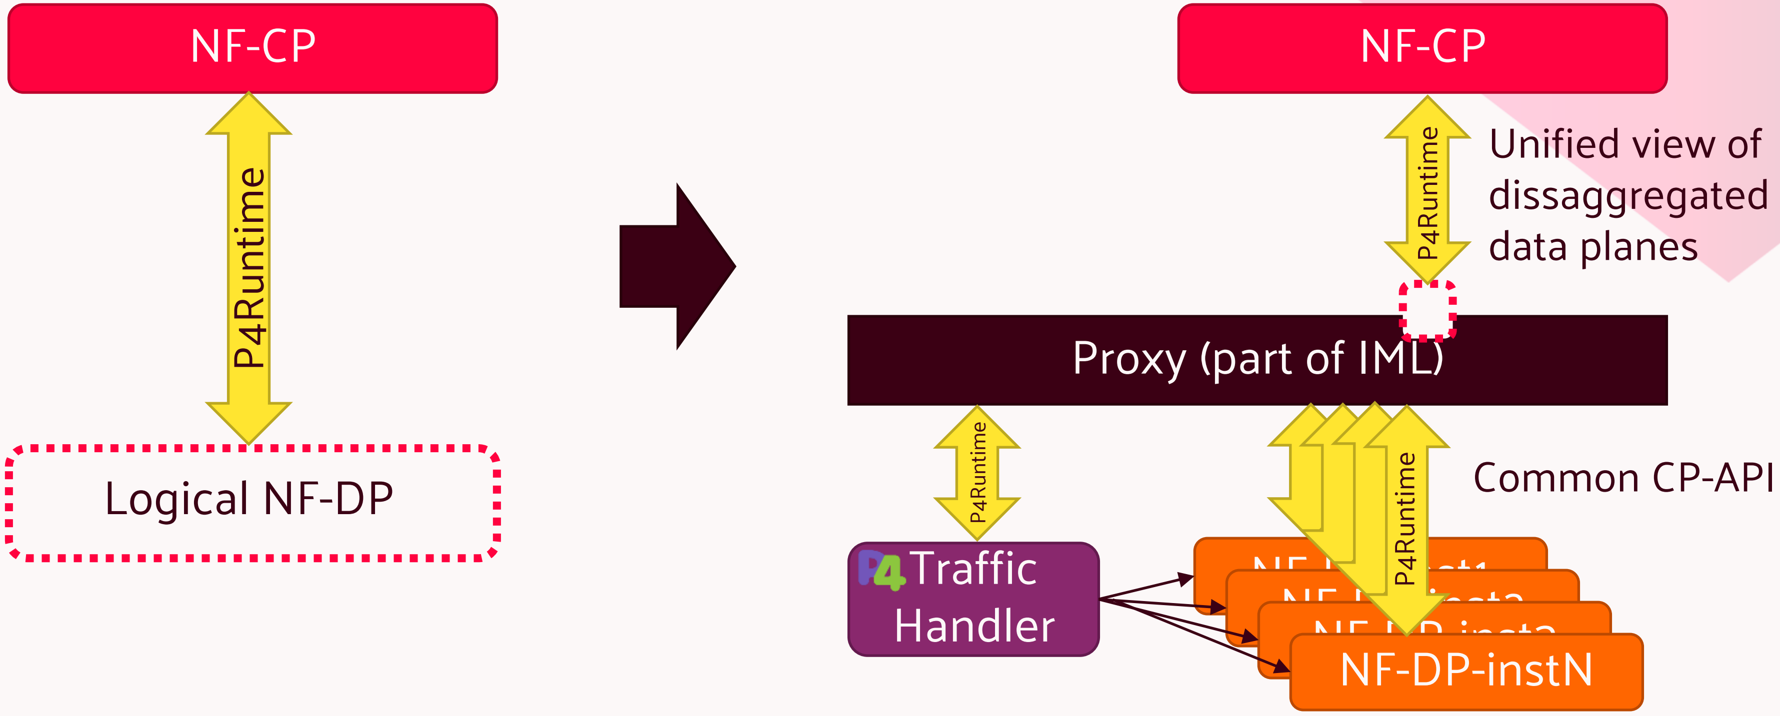

InfraNF on P4 switch, e.g., load balancer, heavy hitter offloader

NF-DP instances on different targets (not only P4)

EXAMPLE2: HEAVY HITTER OFFLOADING WITH HYBRID HARDWARE

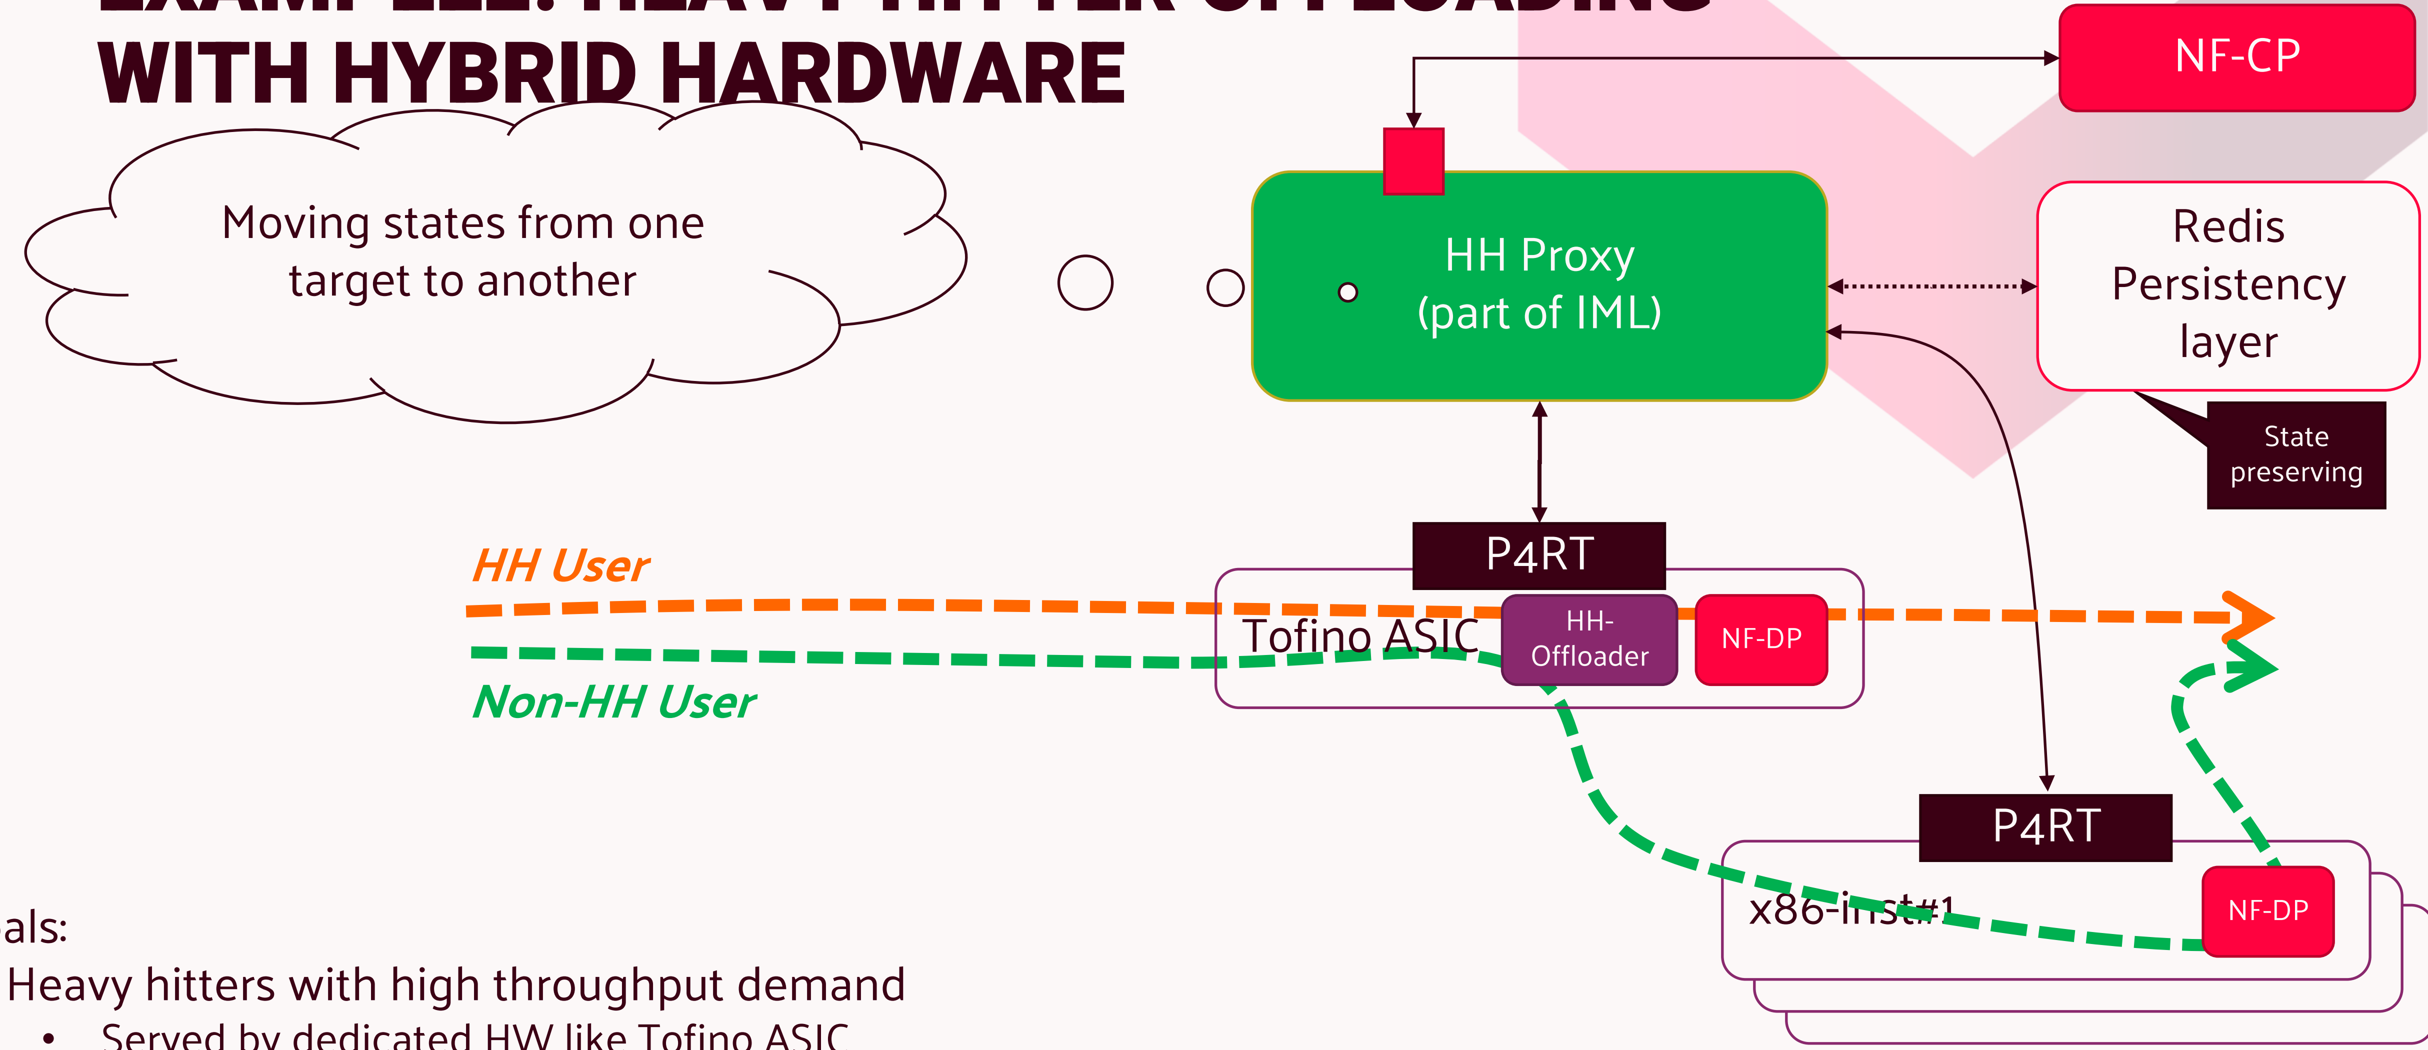

Goals:

- Heavy hitters with high throughput demand
  - Served by dedicated HW like Tofino ASIC
- Non-heavy hitters not requiring dedicated high-speed HW
- Run-time optimization needed

CASE STUDY [1]:
UPF ON HYBRID HW

- Key functions
  - L2 switching/virtualization
  - QoS support
  - Firewall
  - GTP decap/encap
  - L3 routing
- Disaggregation of the pipeline
  - Horizontal split
    - Identical logic, but the traffic is split
  - Vertical split
    - Chain of basic functional blocks

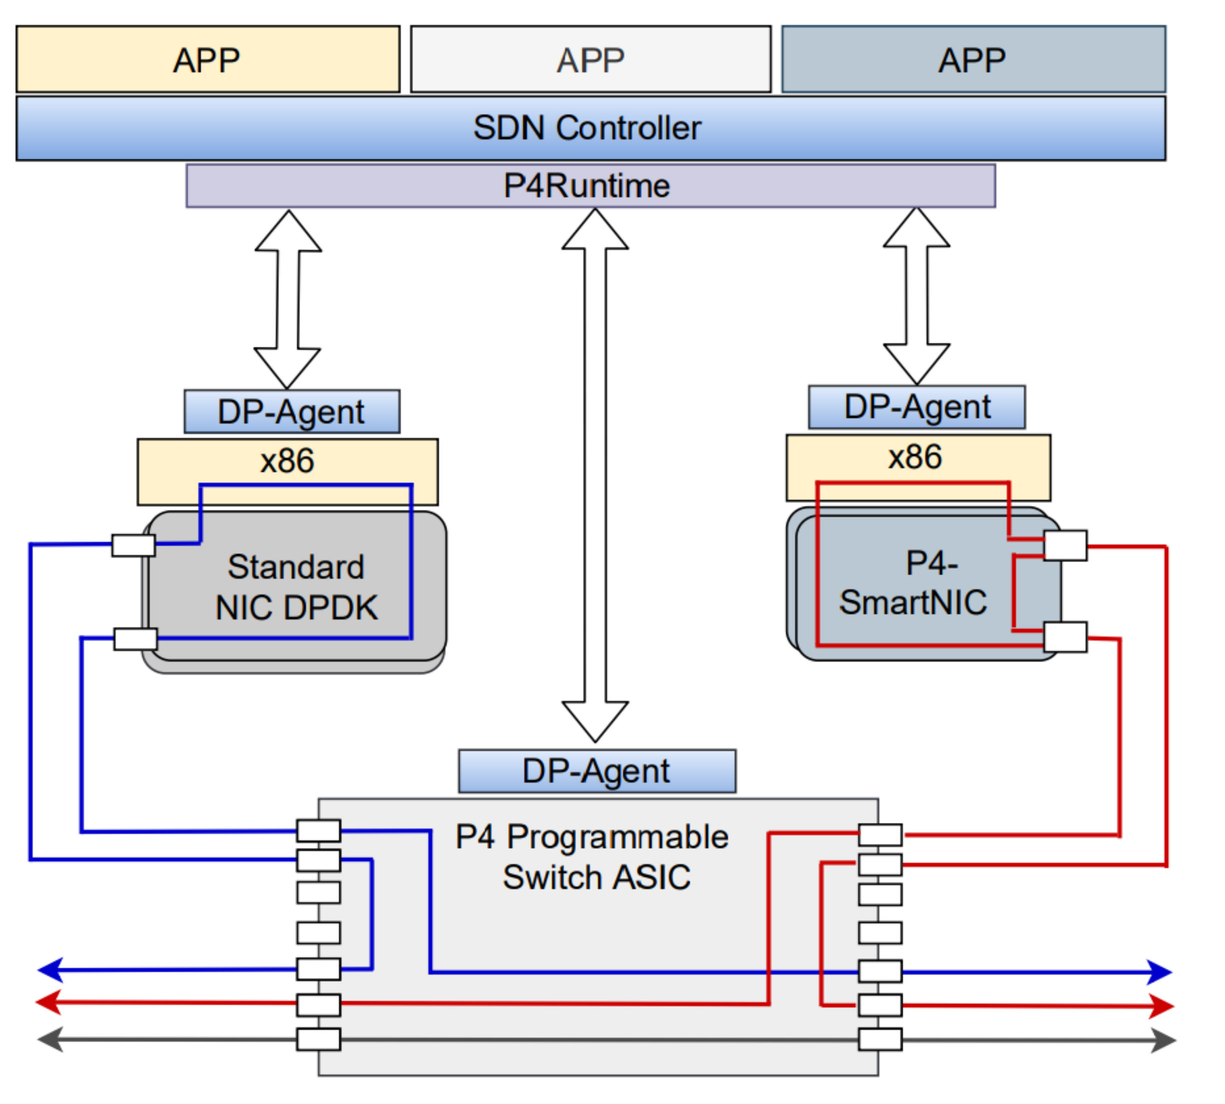

DESIRE6G

[1] S. Kumar Singh et al., "Hybrid P4 Programmable Pipelines for 5G gNodeB and User Plane Functions," in IEEE Transactions on Mobile Computing, 2022, doi: 10.1109/TMC.2022.3201512.

CASE STUDY [1]: UPF ON HYBRID HW

- Tofino ASIC
  - Guaranteed low and bounded per packet delay
  - >6.5 Tbit/sec forwarding capacity
  - Limited SRAM resources - 10000s of UE matches only
  - Good target for crucial control functions like ACL
- Solutions
  - Option 1 – Scaling out to multiple switches
  - Option 2 – Differentiate between UEs
    - 90-95% of UEs are inactive or non-heavy-hitters
    - Only 5-10% have high throughput demand (heavy-hitters (HH))
    - E.g., 5M UEs: 5-10% smart phones (HH), 10-20% wideband IoT (HH), 70-85% narrowband IoT (non-HH)
    - Deploying HHs on Tofino, while non-HHs on x86

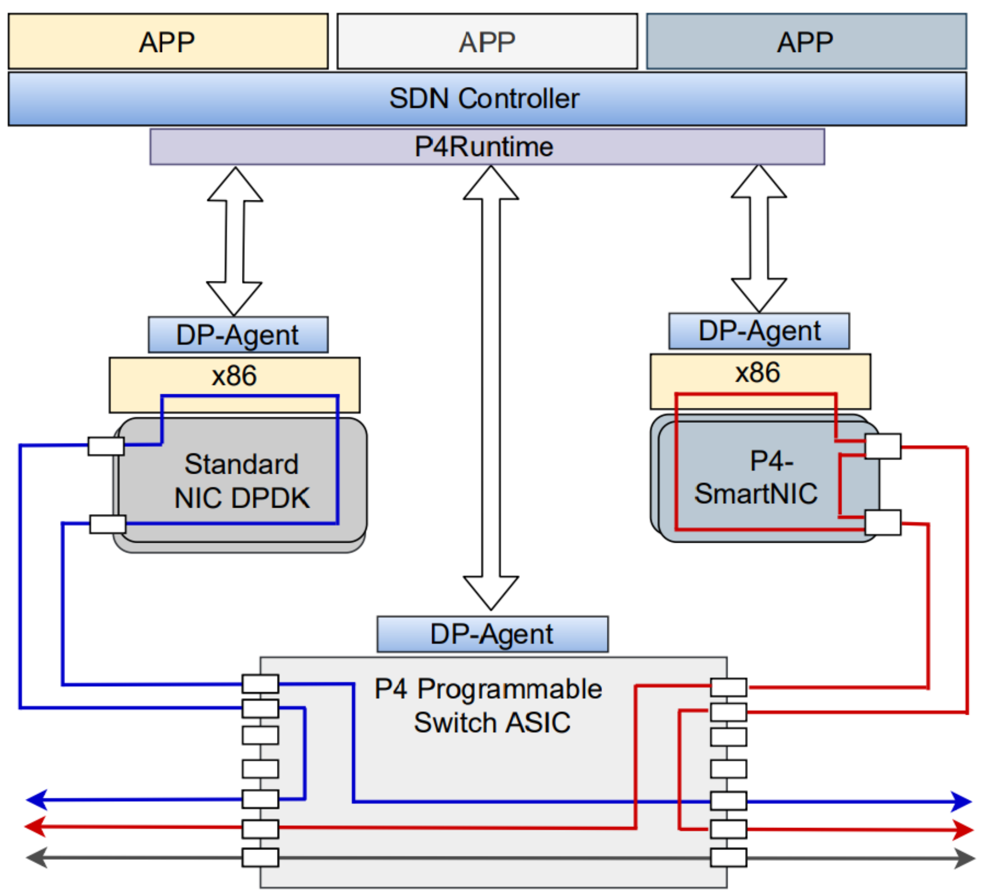

DESIRE6G

[1] S. Kumar Singh et al., "Hybrid P4 Programmable Pipelines for 5G gNodeB and User Plane Functions," in IEEE Transactions on Mobile Computing, 2022, doi: 10.1109/TMC.2022.3201512.

CASE STUDY [1]:
UPF ON HYBRID HW

- Upstream on Tofino only
- Downstream on both
- Heavy hitter detection-based switching
  - Inter Packet Gap-based HH detection
  - High detection accuracy
  - Notification to the control plane
  - Autonomous operation
- Exceptions can be added
  - Low latency flows
  - Slices with low latency requirements

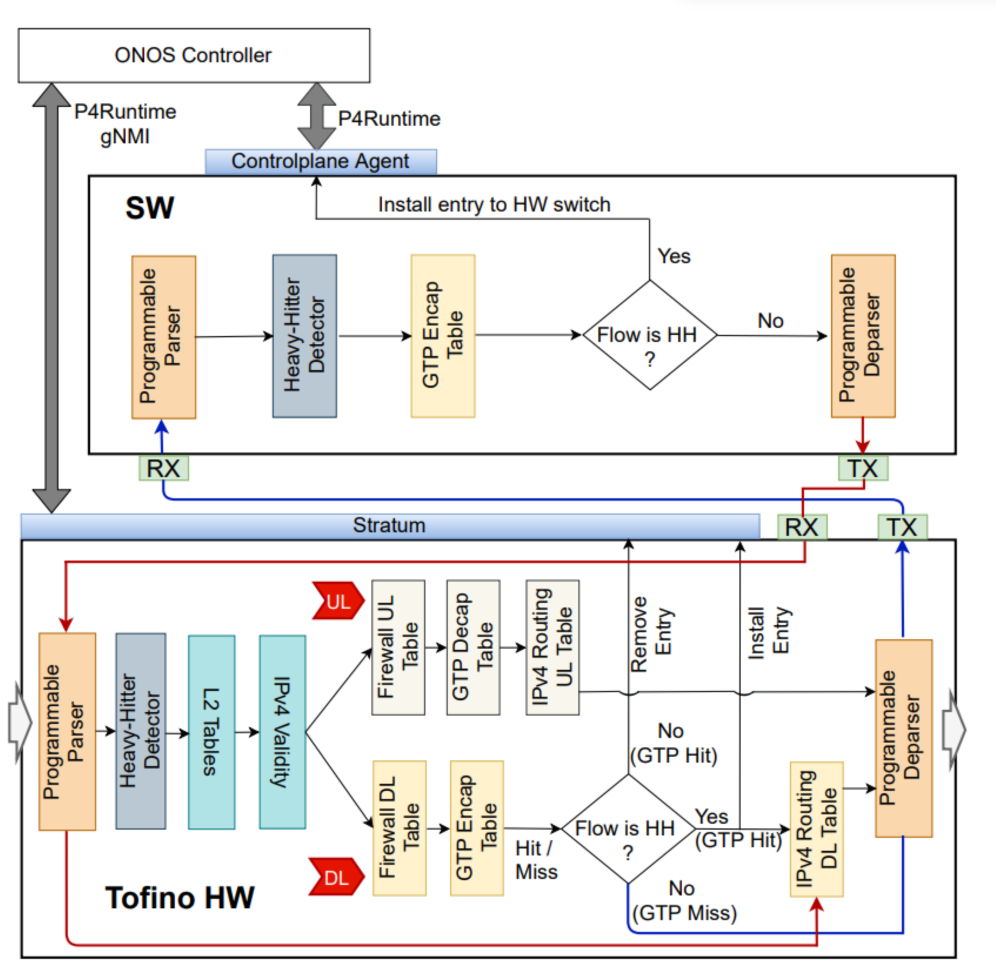

DESIRE6G

[1] S. Kumar Singh et al., "Hybrid P4 Programmable Pipelines for 5G gNodeB and User Plane Functions," in IEEE Transactions on Mobile Computing, 2022, doi: 10.1109/TMC.2022.3201512.

CASE STUDY [1]: UPF ON HYBRID HW

Testbed settings:

- Tofino switch: Edgecore Wedge 100BF-32X
- X86 server: Intel Xeon D-1518 (4C, 2.2GHz) 10G SFP+ ports
- Traffic generator: NetFPGA SUME 10G
- Traffic: CAIDA 2016 ISP traces

	UPF-Tofino	H-UPF(Tofino+x86)
SRAM	60%	34.4%(Tofino)
UEs	850K	430K(Tofino)+15M(x86)

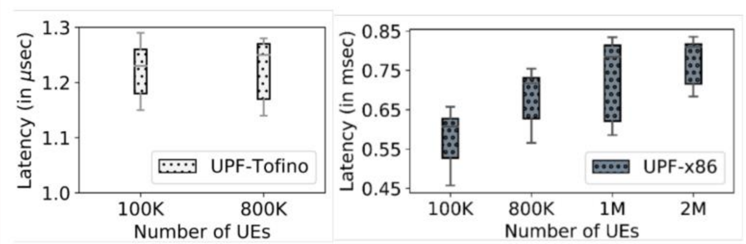

(a) UPF-Tofino Latency (b) UPF-x86 Latency

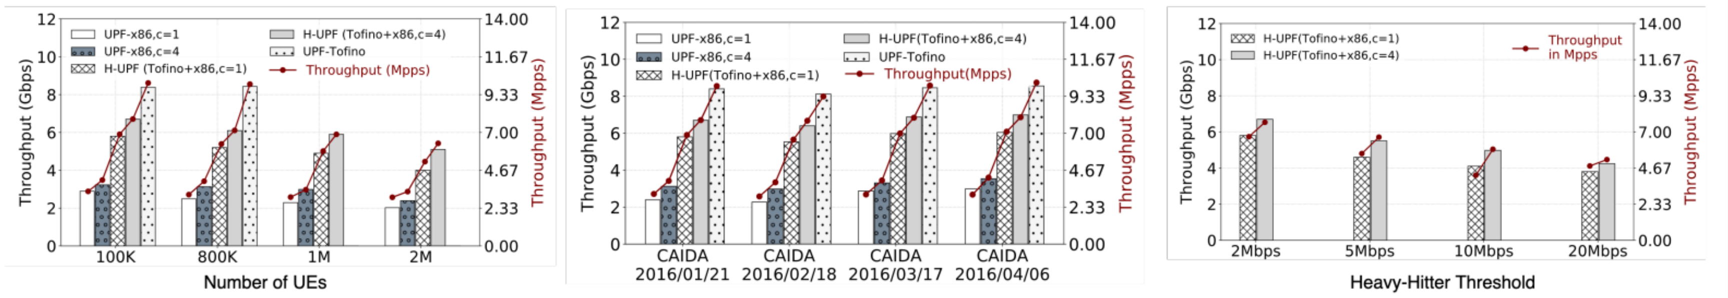

(a) TP for the number of UEs (b) TP for 4 days of CAIDA trace data (c) TP based on different HH thresholds

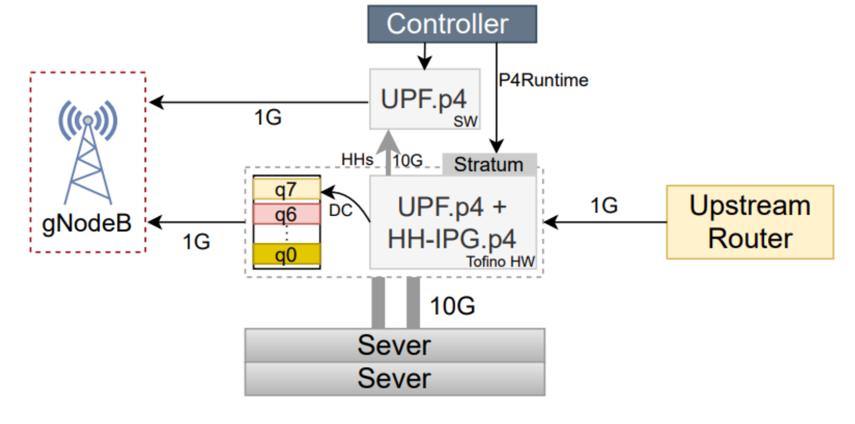

(a) QoS-HH Use-case Scenario

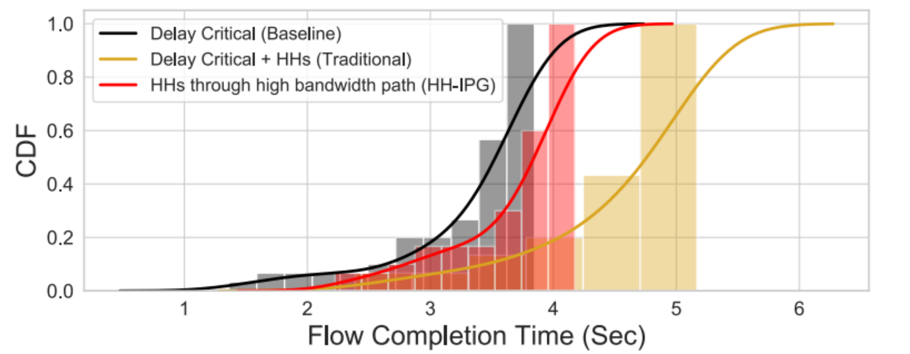

(b) Flow completion time with and without HHs offloading

[1] S. Kumar Singh et al., "Hybrid P4 Programmable Pipelines for 5G gNodeB and User Plane Functions," in IEEE Transactions on Mobile Computing, 2022, doi: 10.1109/TMC.2022.3201512.

TAKE-AWAY

- **Programmable data planes as Technology enablers**
  - Accelerating customized packet processing
  - Quasi deterministic, ultra-low packet processing latency
  - InfraNFs can do runtime optimization at packet processing time-scale
    - Non-traditional traffic management – fine-grained resource sharing
    - Routing/Fast Rerouting
    - Load balancing including heavy hitter offloading can improve scalability – less load on CPU resources
  - Pervasive monitoring via in-band network telemetry
    - Fast notification and reaction to unexpected situations, failures and performance issues
- **Challenges**
  - HW PDPs are not shared resources by default, application of the cloud-native approach is challenging
  - HW PDPs have numerous limitations and many restrictions
  - Migration of stateful NF-DPs
  - Seamless data plane optimization (acc. cloud-native approach)
  - Dealing with non-programmable node in the transport

THANKS!

Sandor Laki

lakis@inf.elte.hu

Co-funded by the European Union

DESIRE6G has received funding from the Smart Networks and Services Joint Undertaking (SNS JU) under the European Union's Horizon Europe research and innovation programme under Grant Agreement No 101096466.
Funded by the European Union. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or European Commission. Neither the European Union nor the granting authority can be held responsible for them.