DESIRE6G

END-TO-END DATA PLANE ABSTRACTION FOR SUPPORTING DEEP SLICING IN 6G

Sándor Laki, assistant professor

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

Co-funded by the European Union

6GSNS

06/06/2023

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

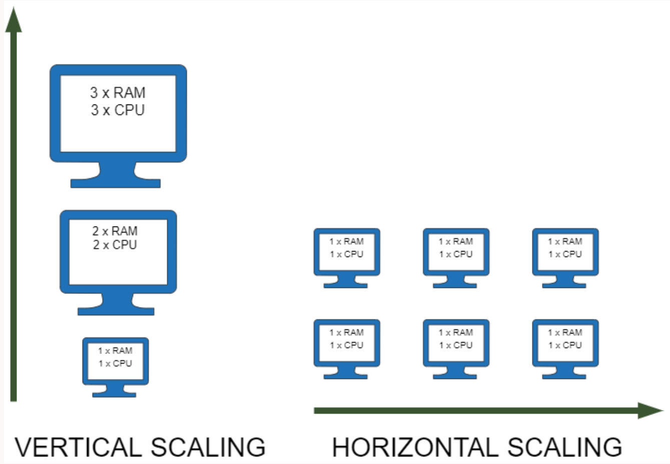

DESIRE6G

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

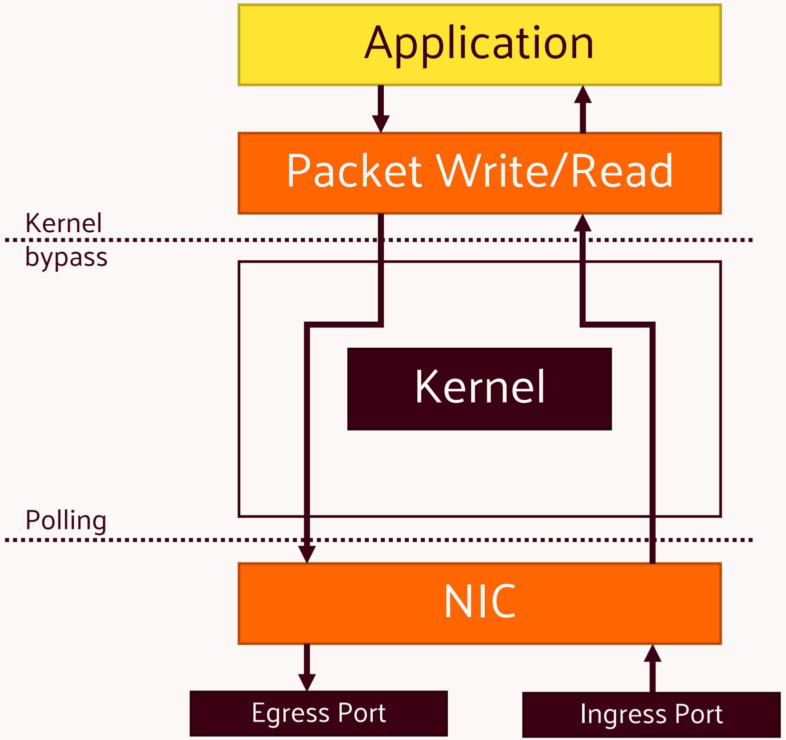

PROGRAMMABLE NETWORK DEVICES AS NF(V) BACKENDS

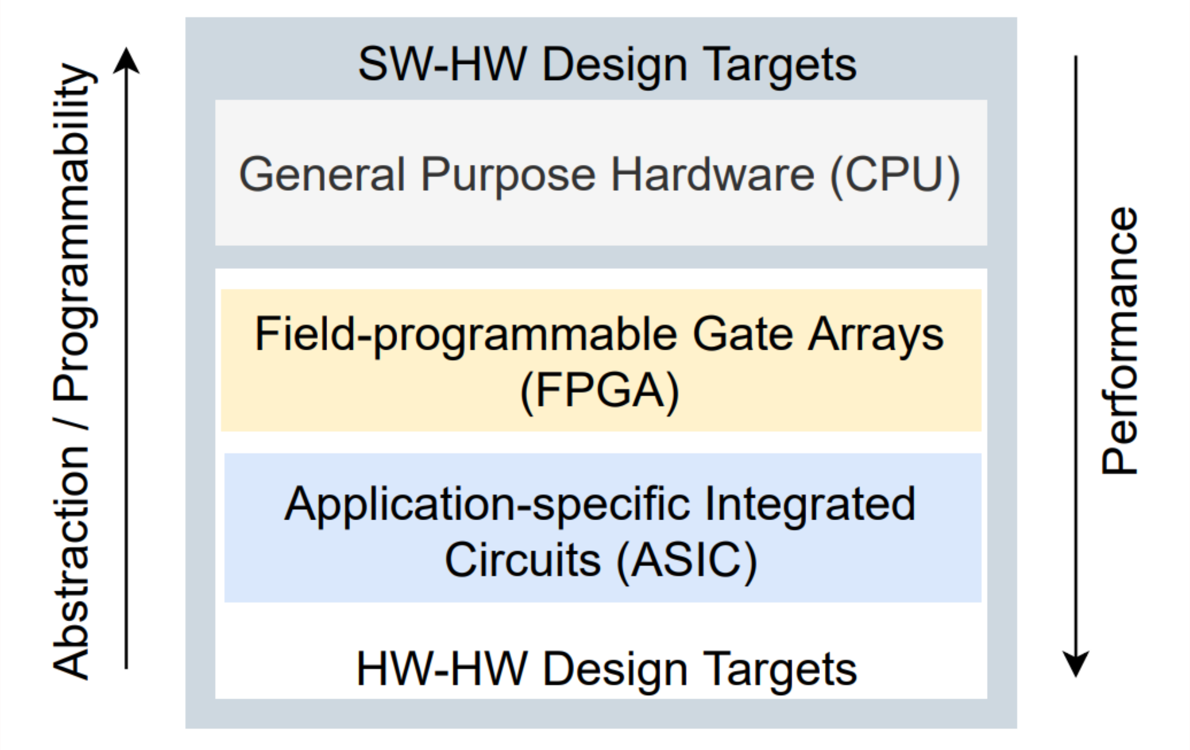

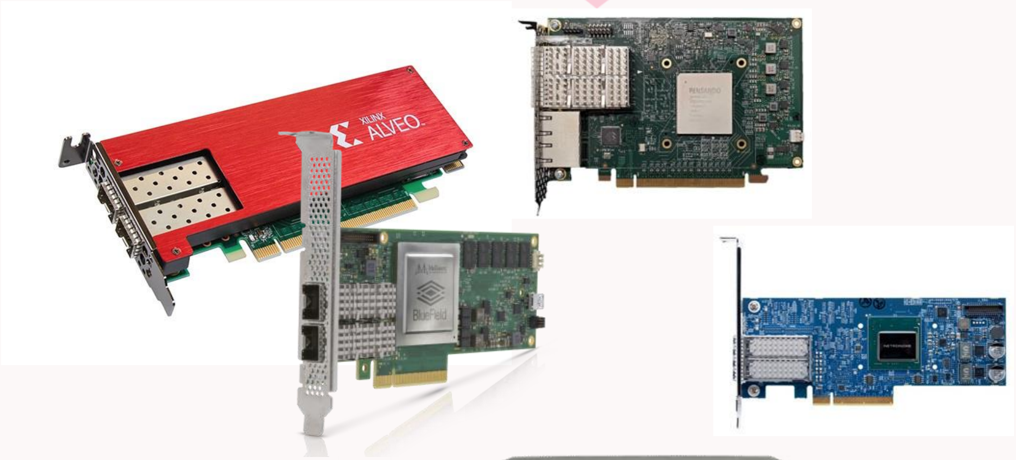

DPDK

eBPF

P4

E2E PROGRAMMABILITY VISION

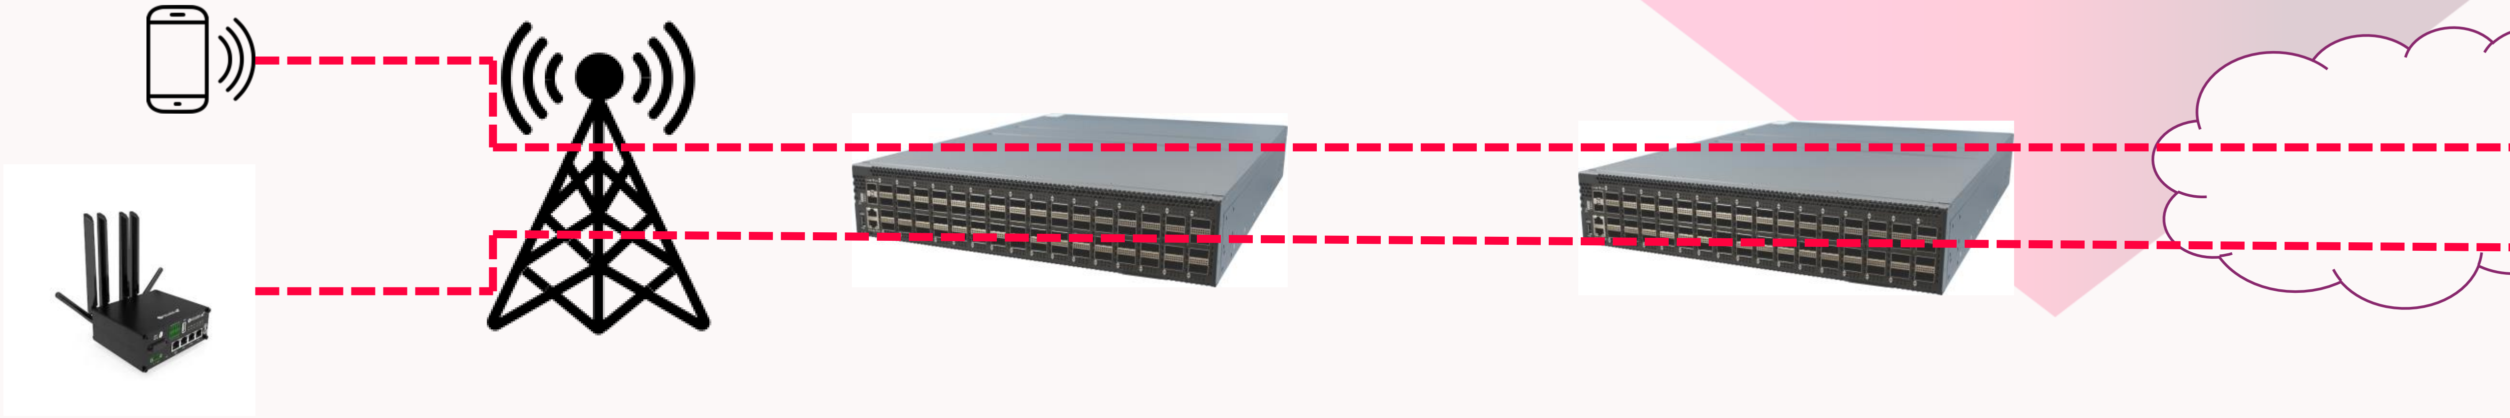

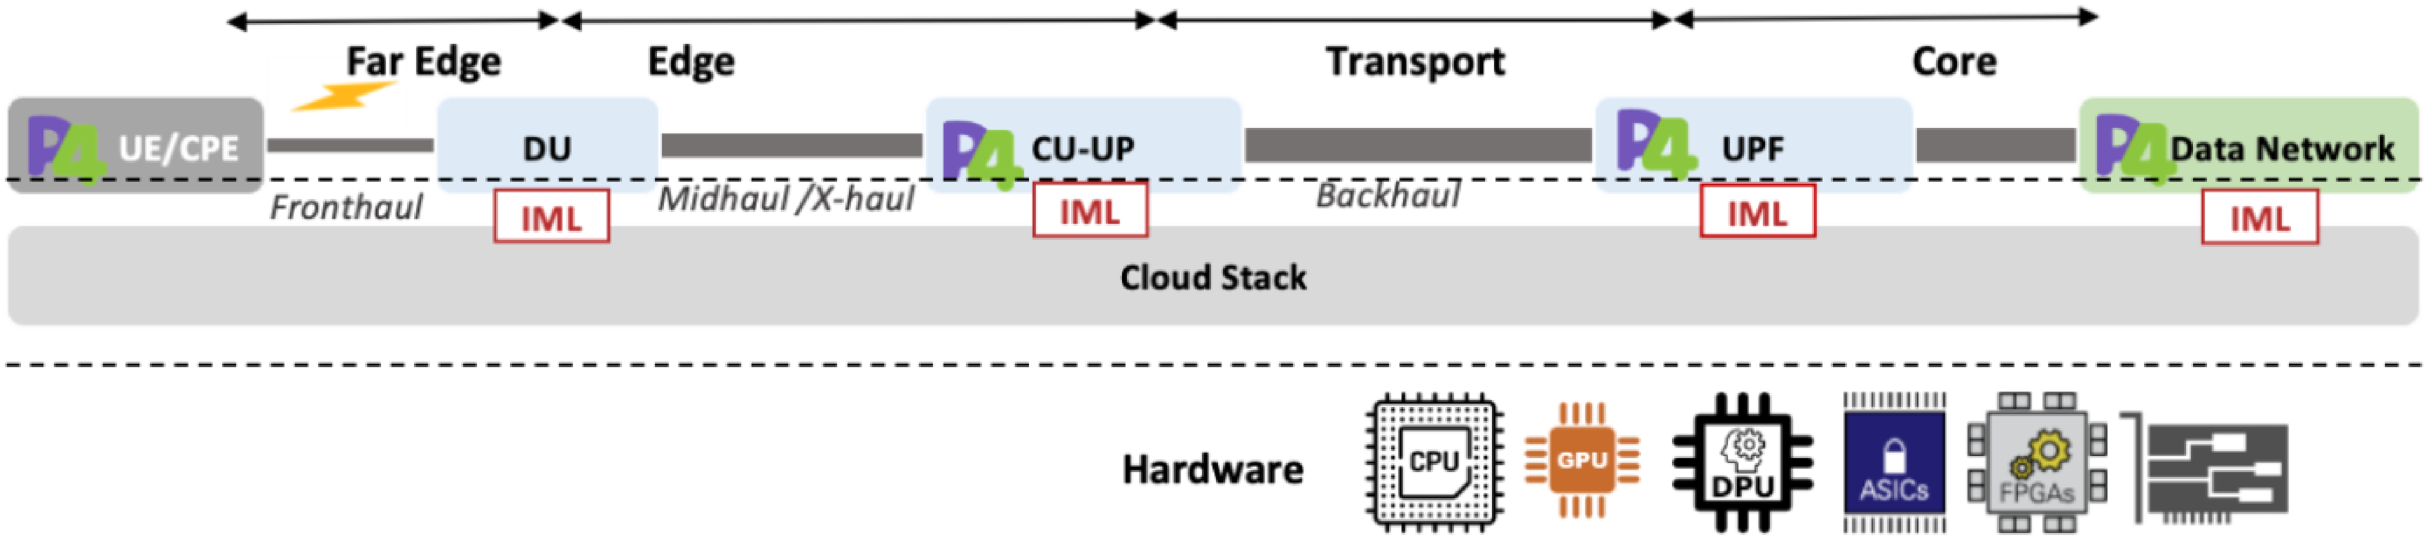

SHARED INFRASTRUCTURE

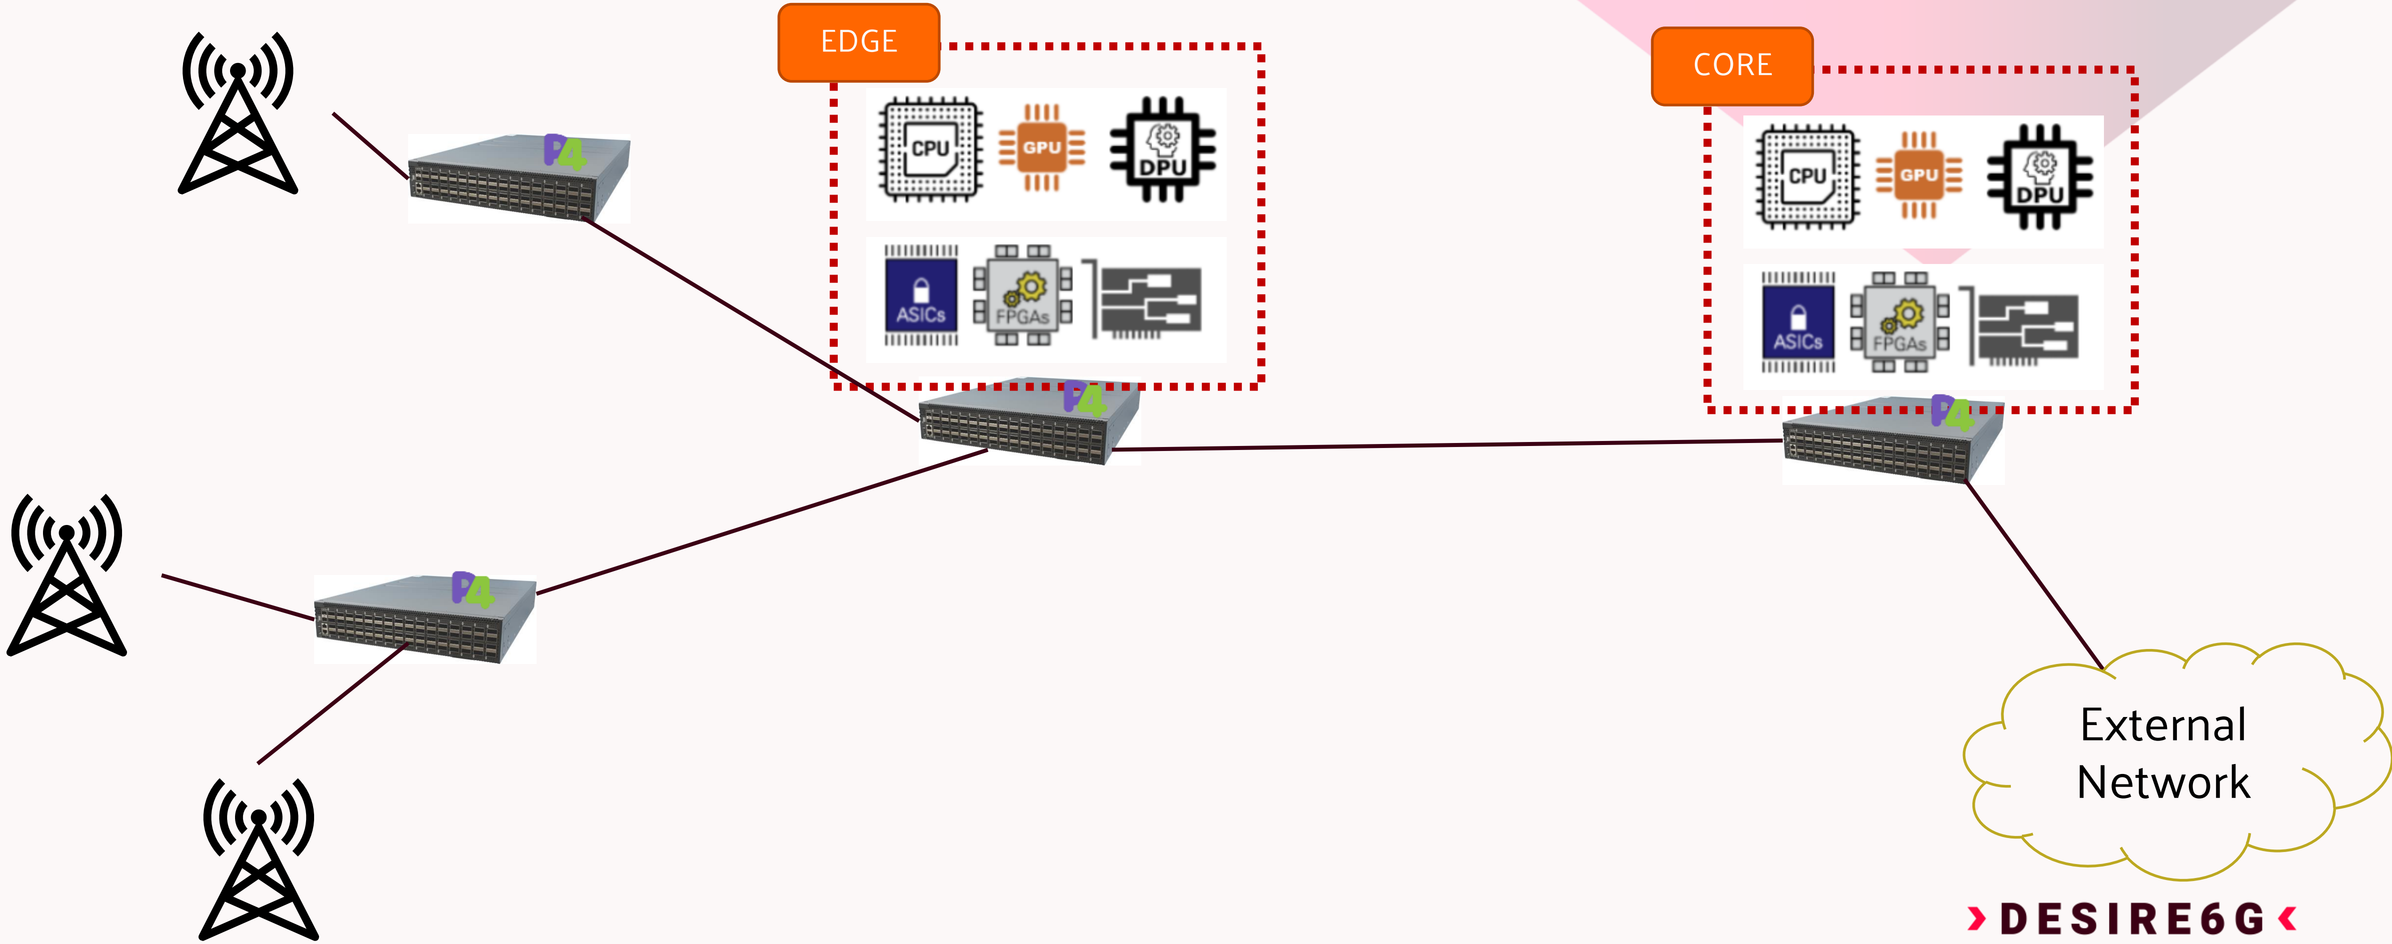

DESIRE6G

SHARED INFRASTRUCTURE

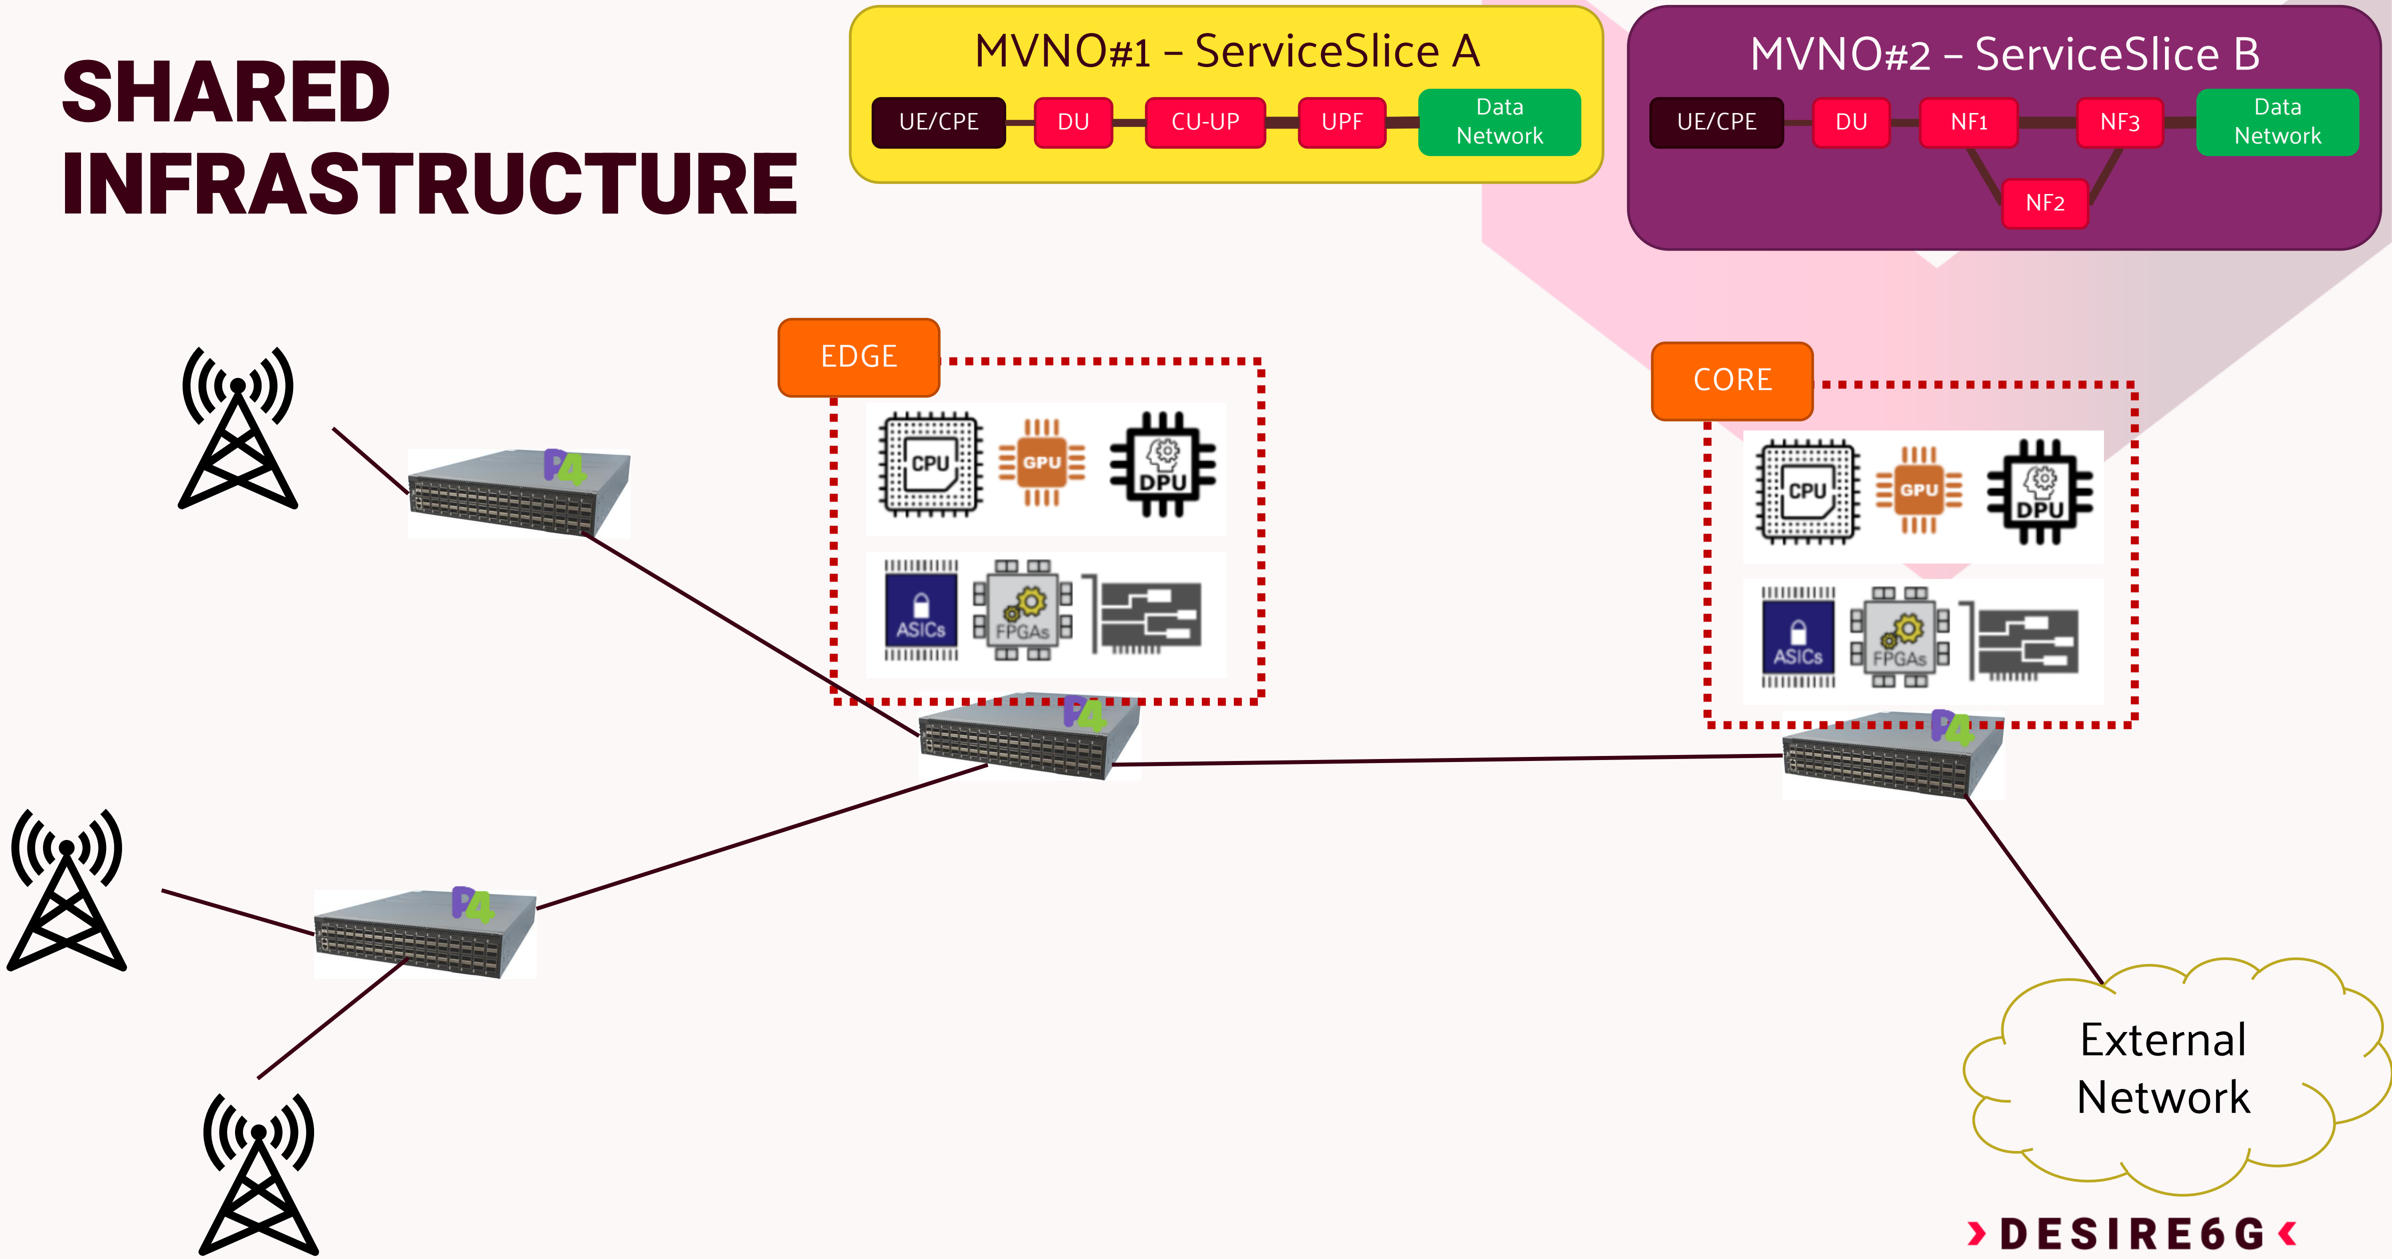

DESIRE6G

SHARED INFRASTRUCTURE

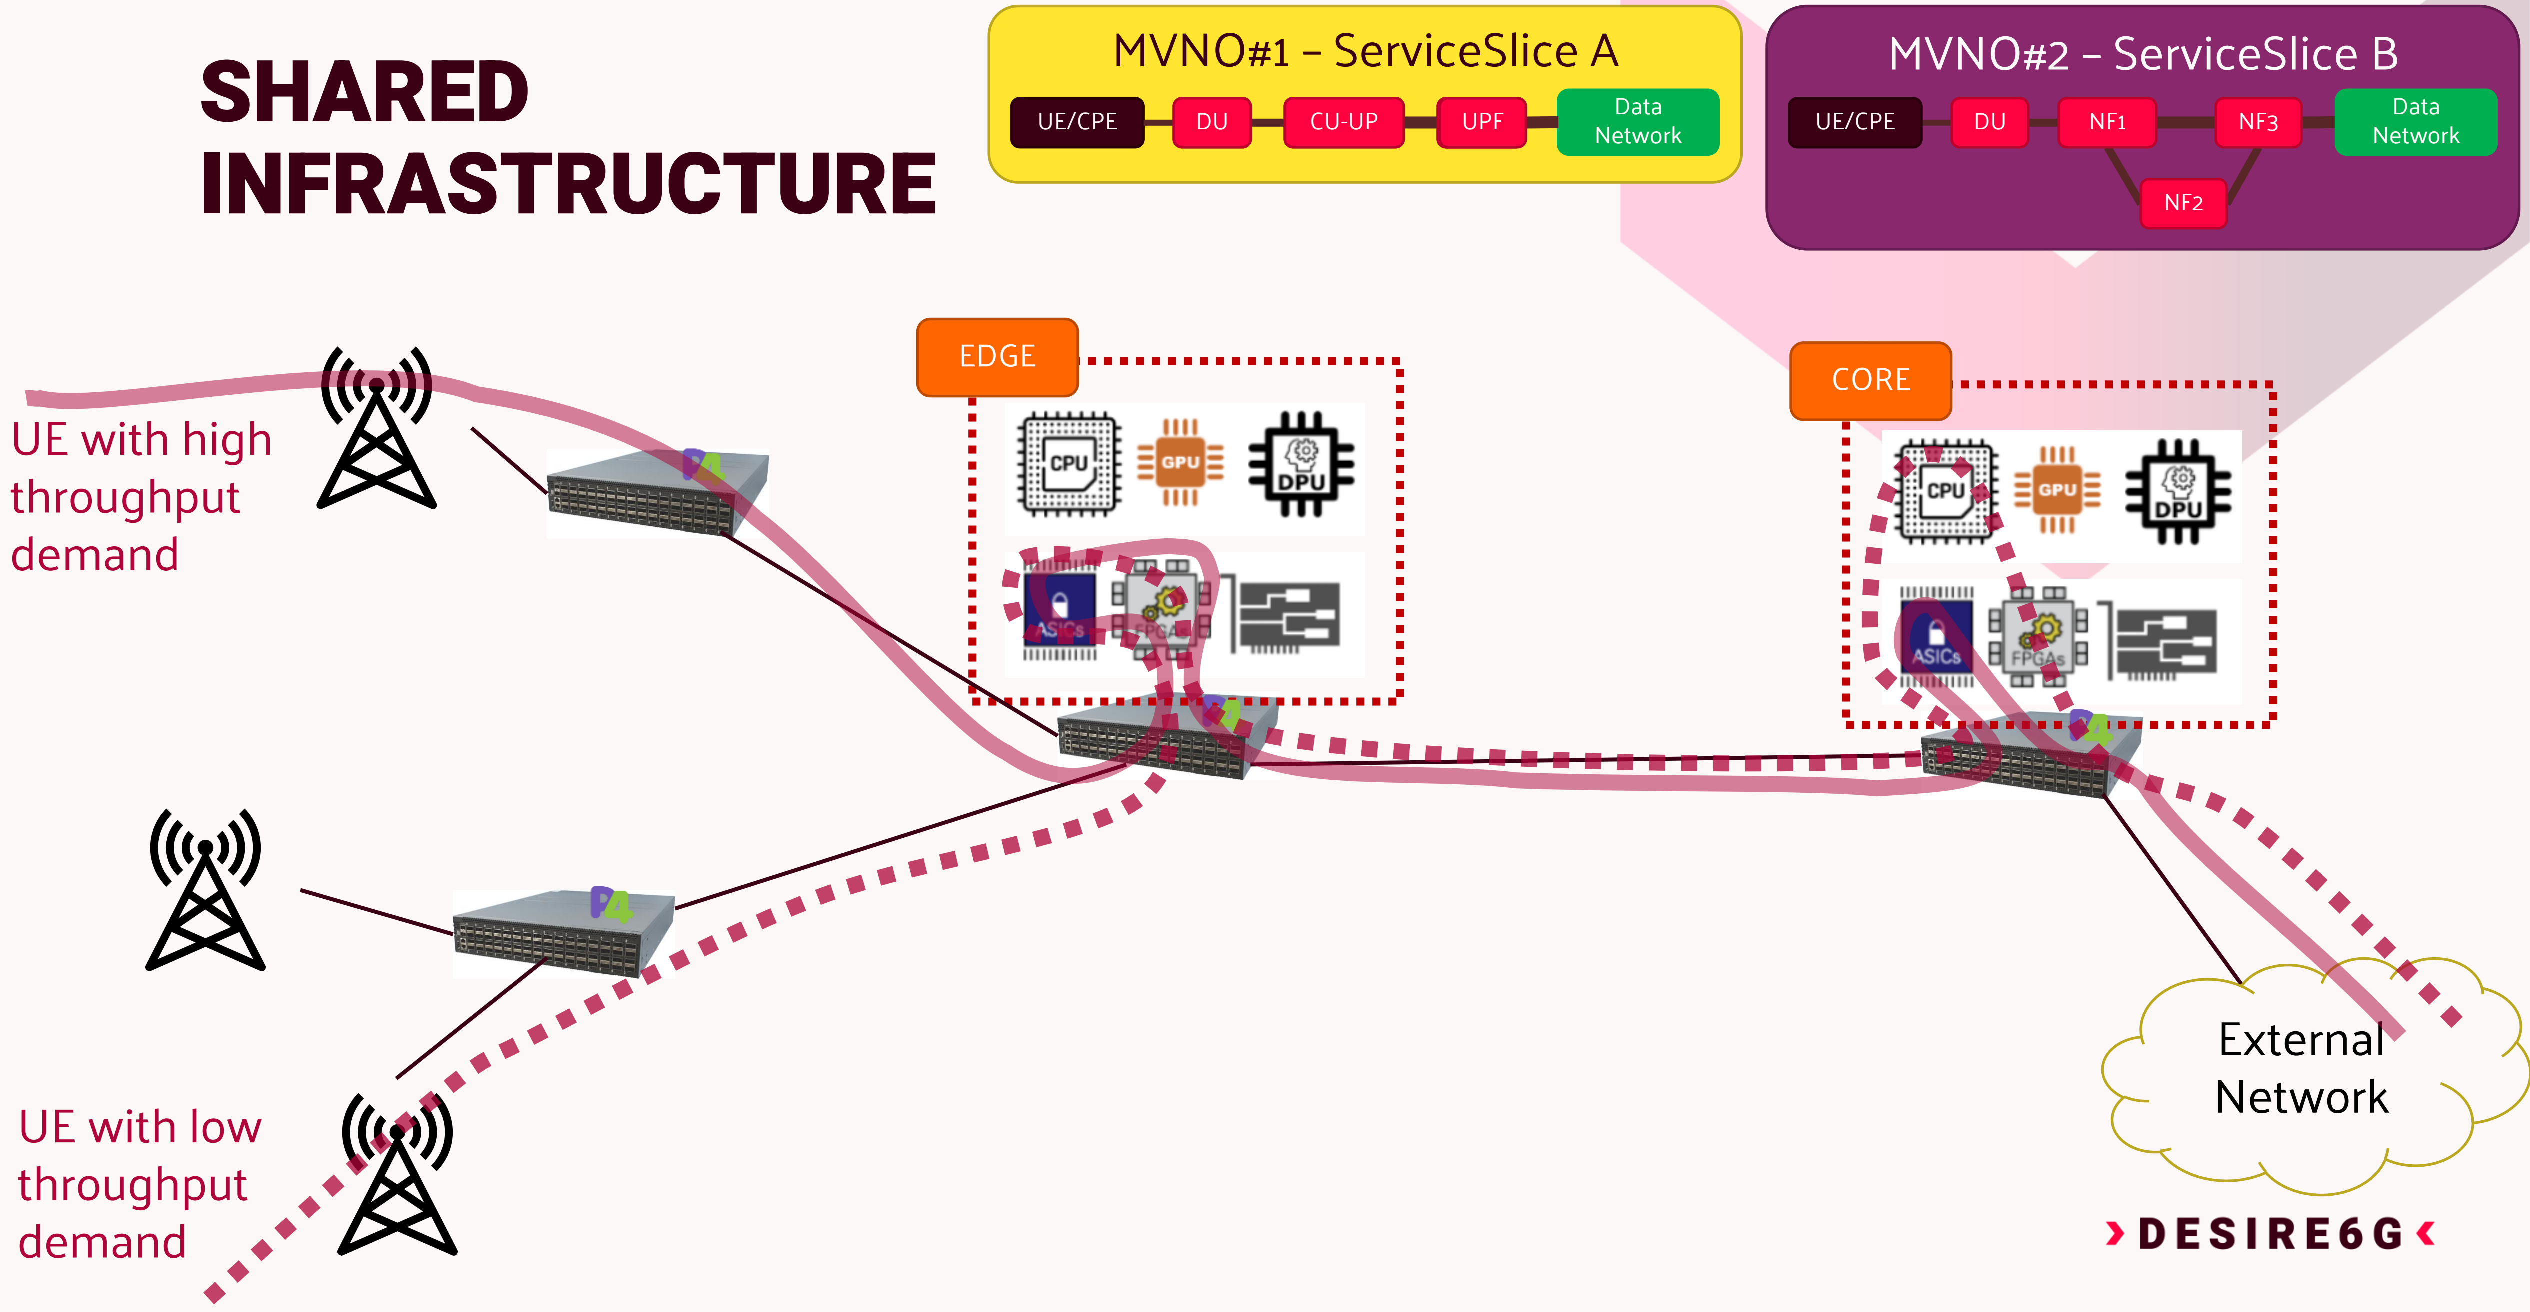

DESIRE6G

SERVICE SLICE: INSTANCE OF A NETWORK FUNCTION GRAPH TEMPLATE

Abstract NF (or service) graph

- Describes the end-to-end packet processing logic of one service
  - e.g., Internet access, robot control @MEC
- A user/application can join a network service (i.e., the slice implementing it) by instantiating the template between the end points
  - Done by mostly configuration, but redeployment of NFs may also be required
- One graph per direction (UL/DL) – the functionality is not always the same

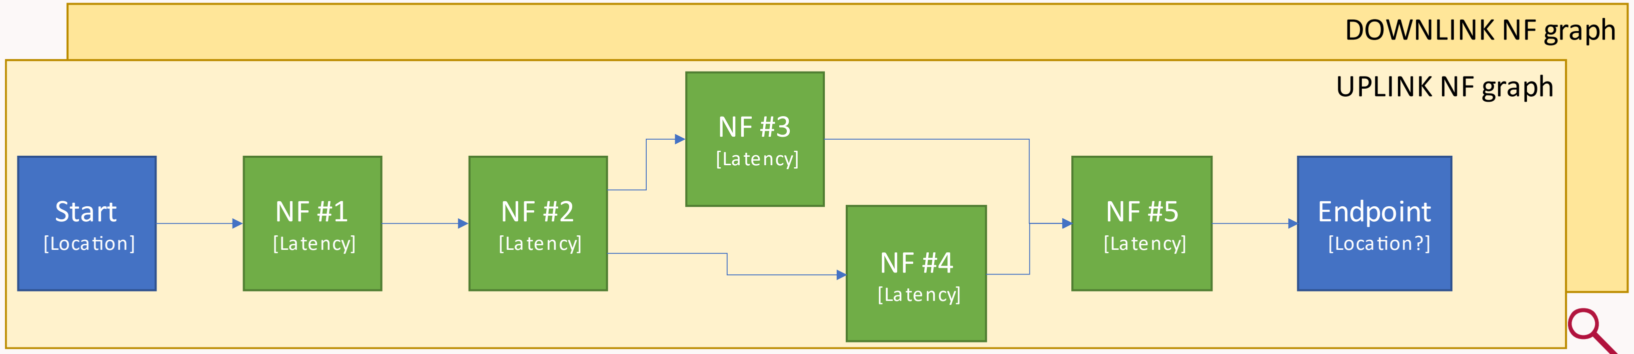

QOS SLICES

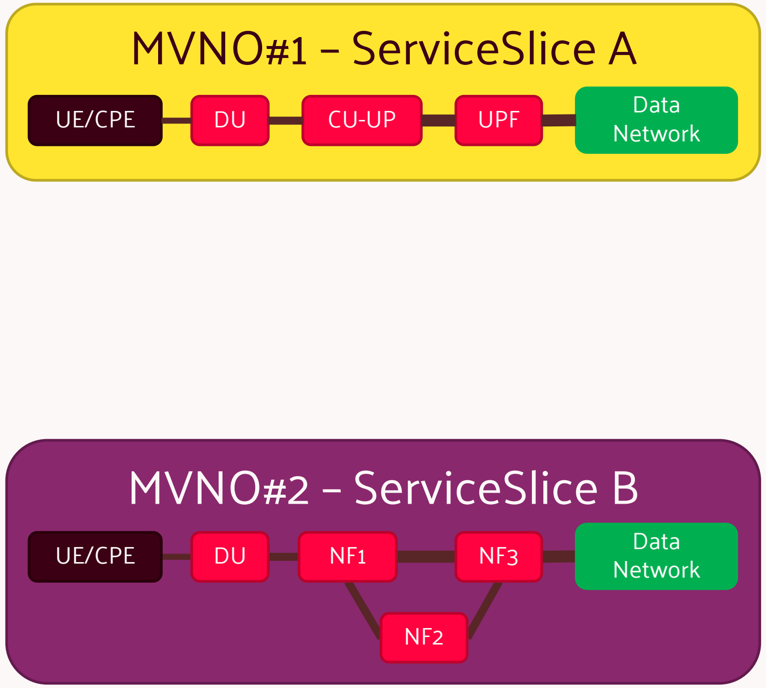

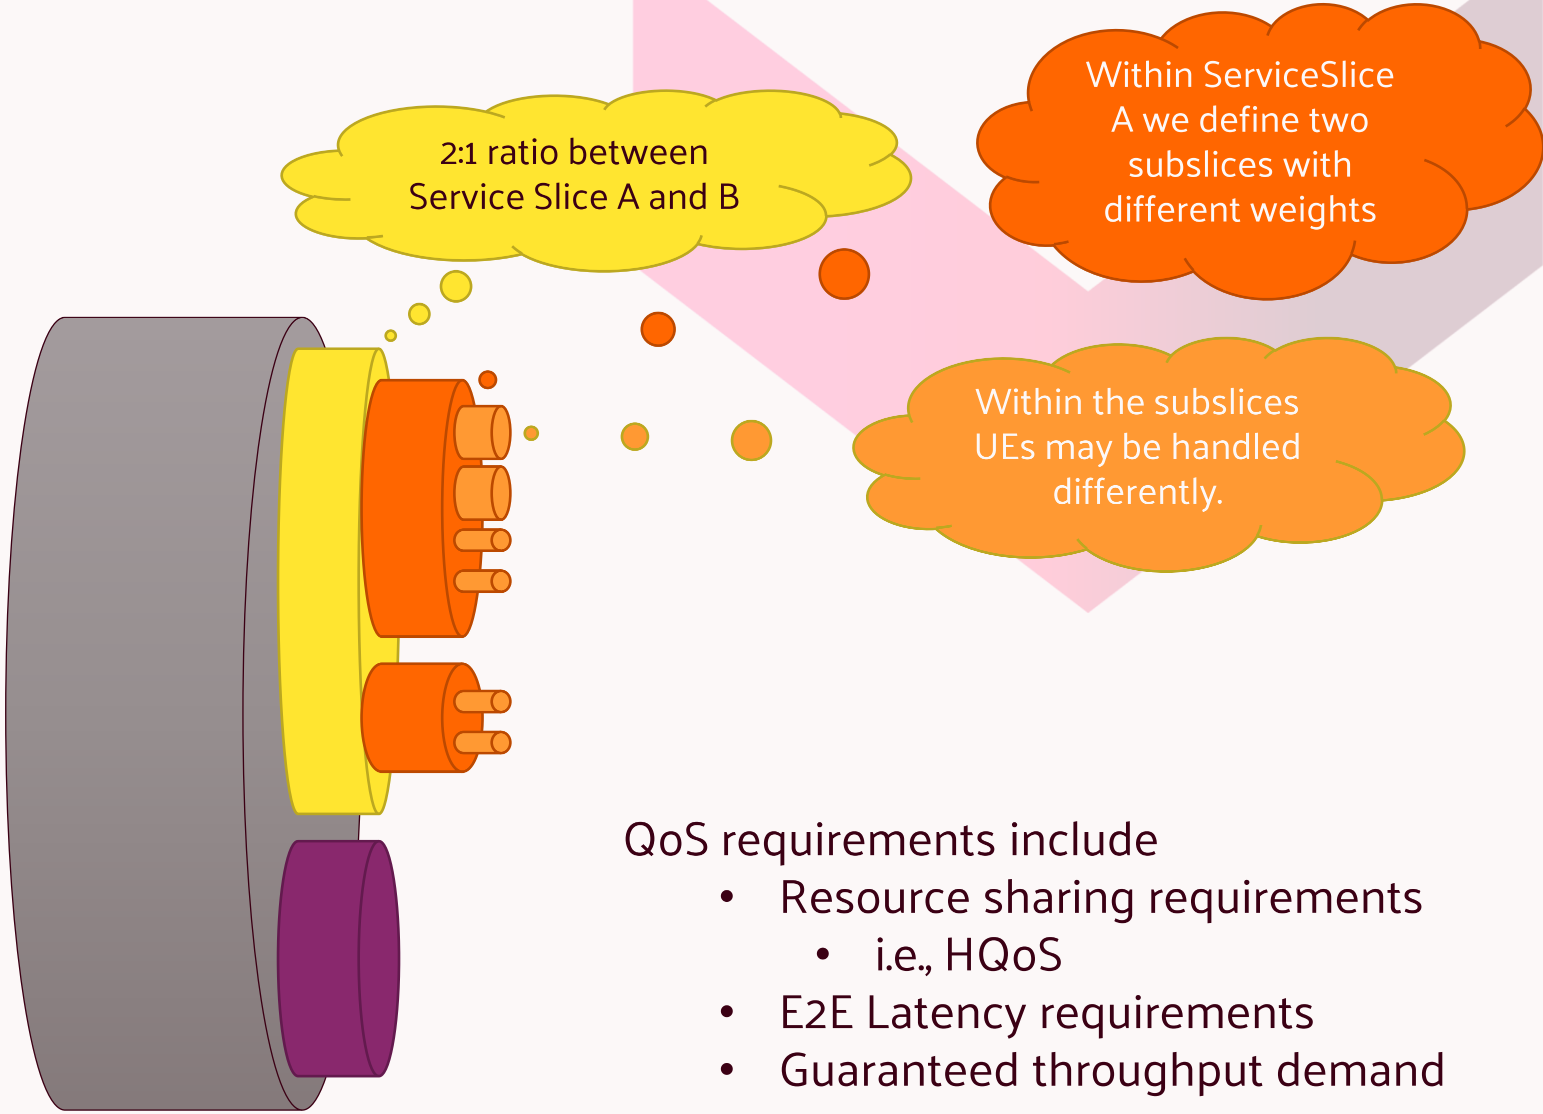

QoS requirements include

- Resource sharing requirements
  - i.e., HQoS
- E2E Latency requirements
- Guaranteed throughput demand

DESIRE6G

DEEP 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
- 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

PACKET PROCESSING NETWORK FUNCTIONS: CP+DP

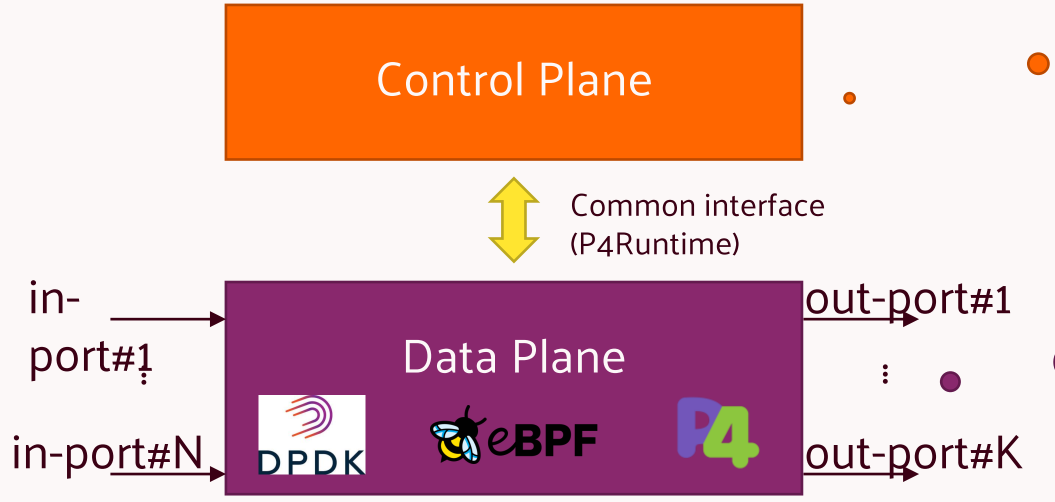

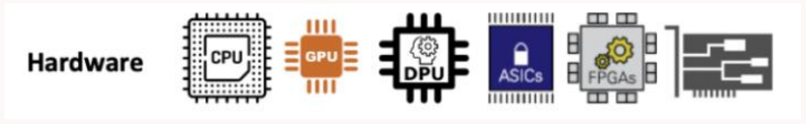

Running on CPU

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
- …

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

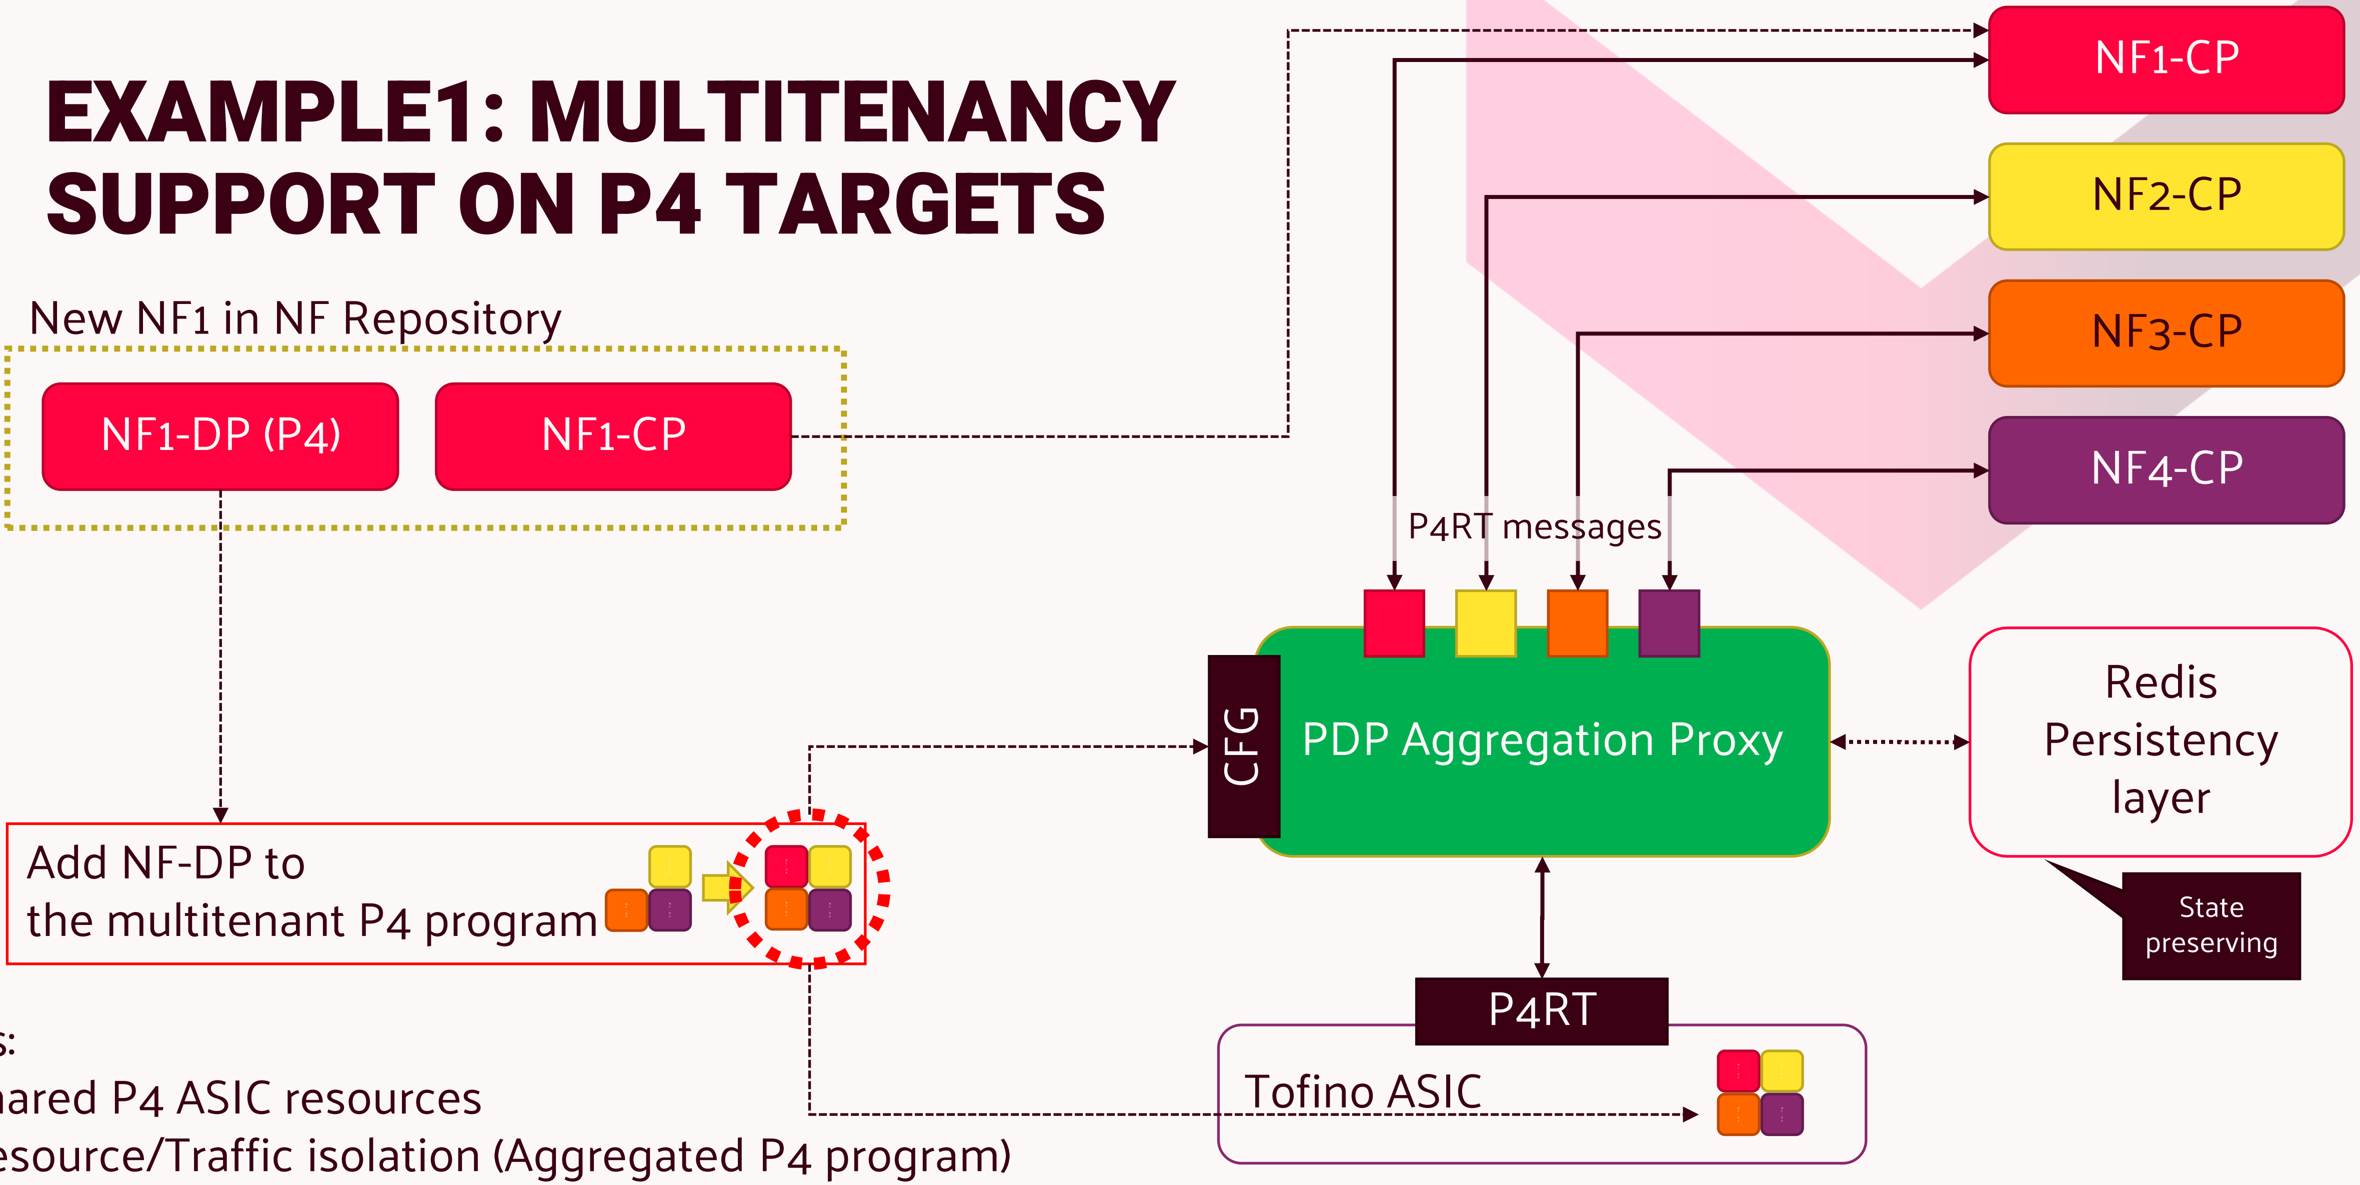

Goals:

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

EXAMPLE2: HEAVY HITTER OFFLOADING WITH HYBRID HARDWARE

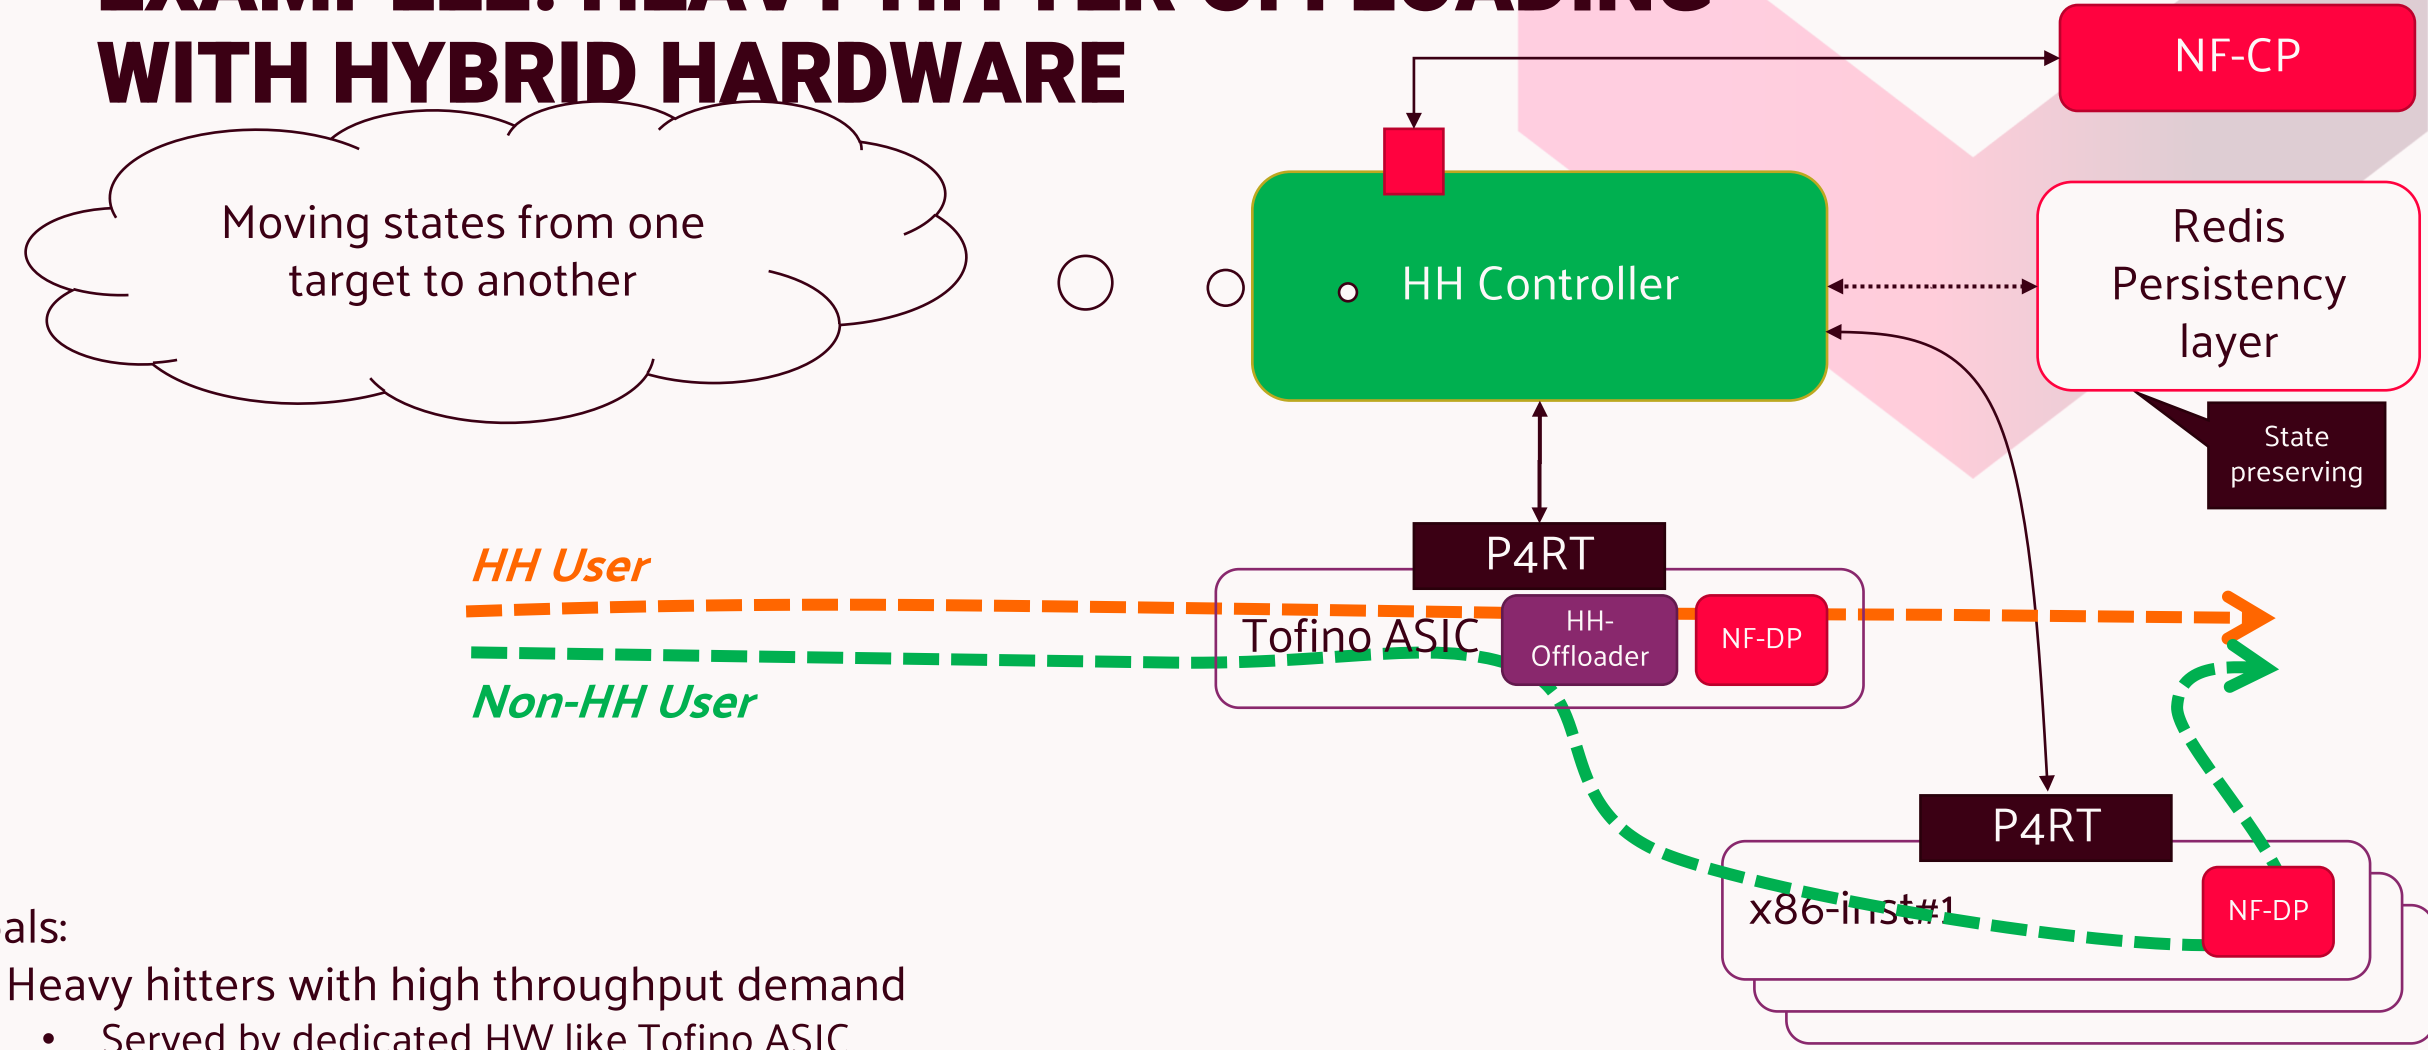

Goals:

- 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

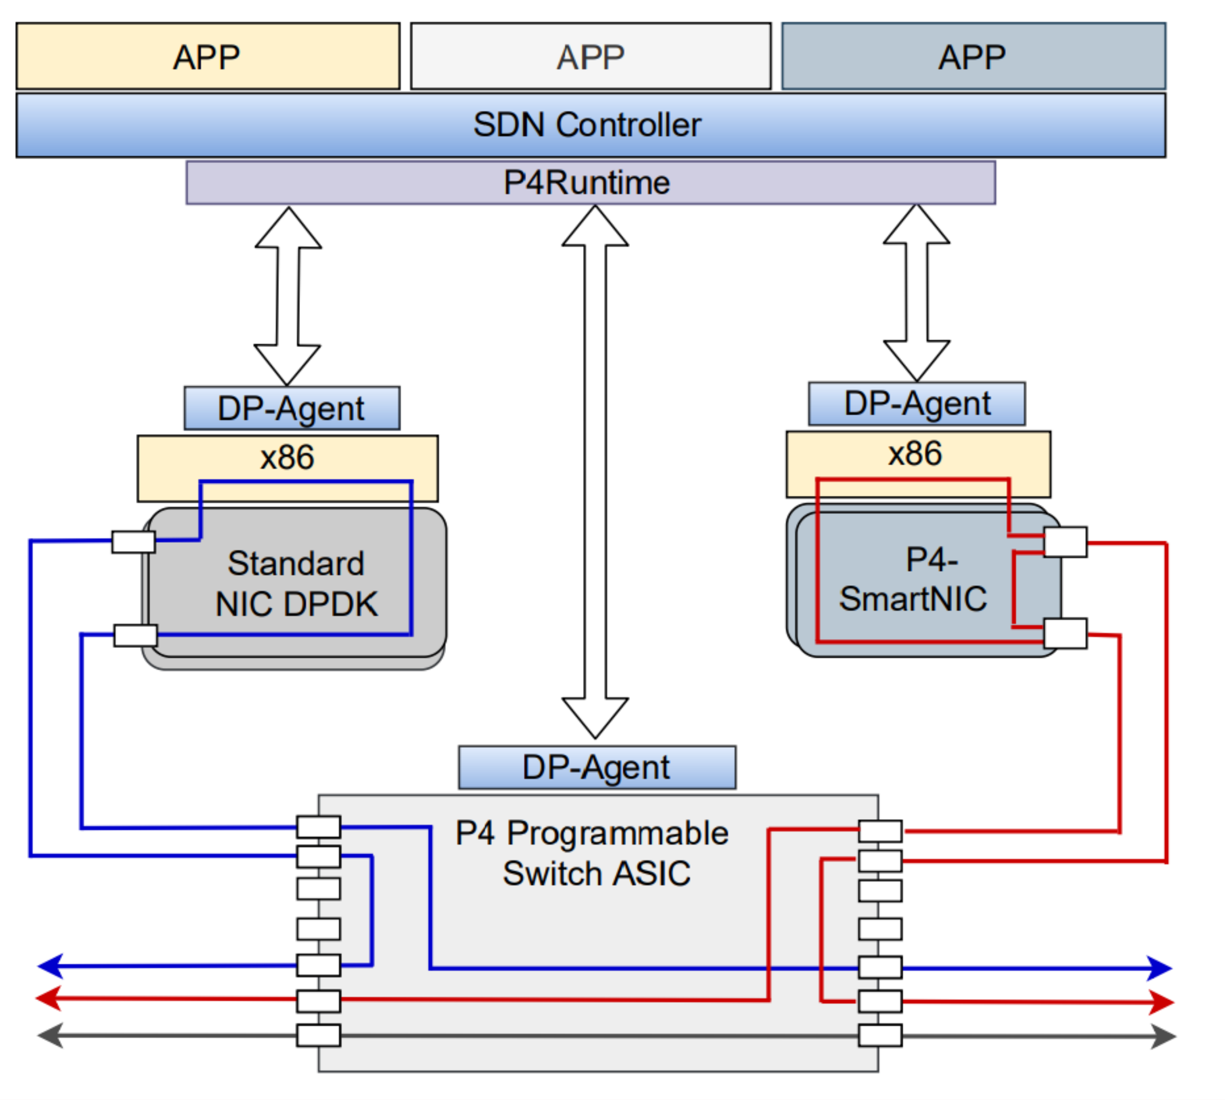

DESIRE6G

[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

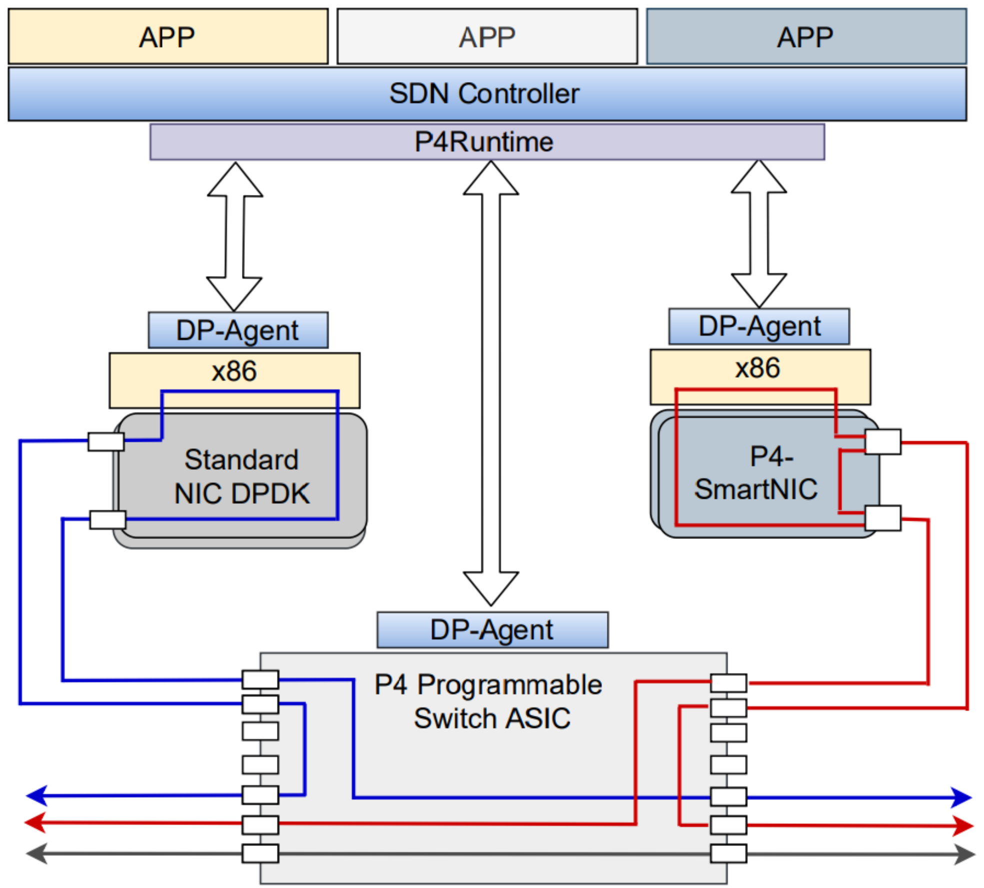

DESIRE6G

[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

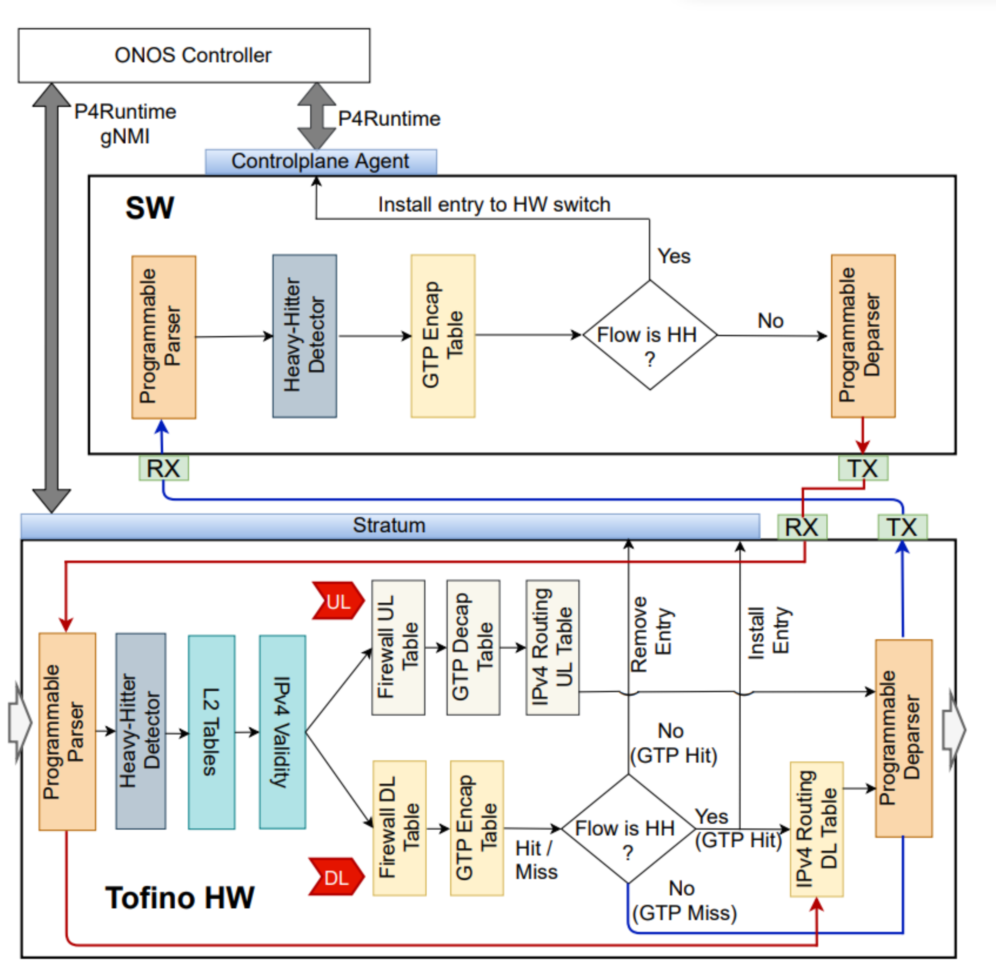

DESIRE6G

[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)

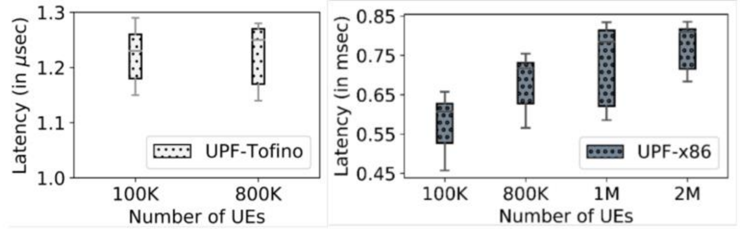

(a) UPF-Tofino Latency (b) UPF-x86 Latency

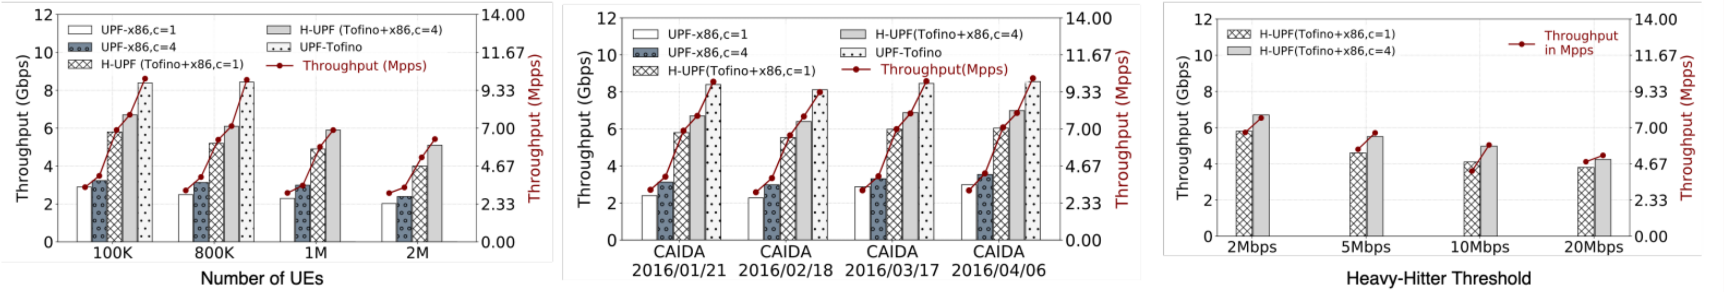

(a) TP for the number of UEs (b) TP for 4 days of CAIDA trace data (c) TP based on different HH thresholds

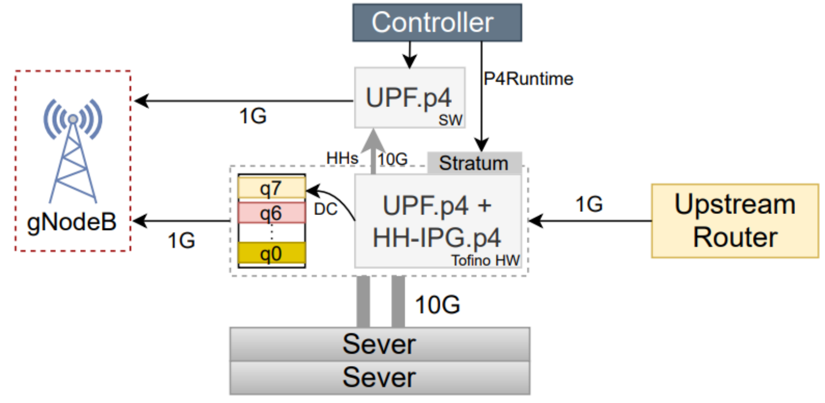

(a) QoS-HH Use-case Scenario

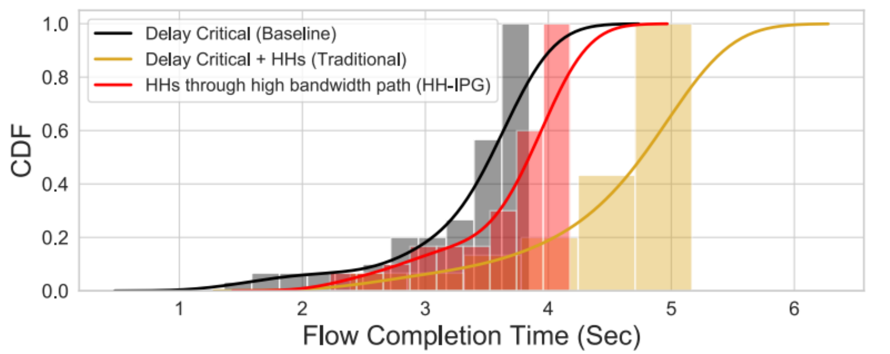

(b) Flow completion time with and without HHs offloading

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.

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
  - HW PDPs have numerous limitations and many restrictions
  - Migration of stateful NF-DPs
  - Dealing with non-programmable node in the transport

DESIRE6G

THANKS!

Sandor Laki

lakis@inf.elte.hu

Co-funded by the European Union

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.
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.