TeraFlow SDN by ETSI

Introduction to P4

Georgios P. Katsikas, Panagiotis Famelis

ETSI TFS – Hackfest #3, October 16, 2023

Who we are

UBITECH
ubiquitous solutions

Our company

Network Softwarization & IoT (NSIT)

Our unit

Georgios P. Katsikas
NSIT Tech. Lead

gkatsikas@ubitech.eu

Panagiotis Famelis
SW engineer + PM

pfamelis@ubitech.eu

What is a packet processing pipeline?

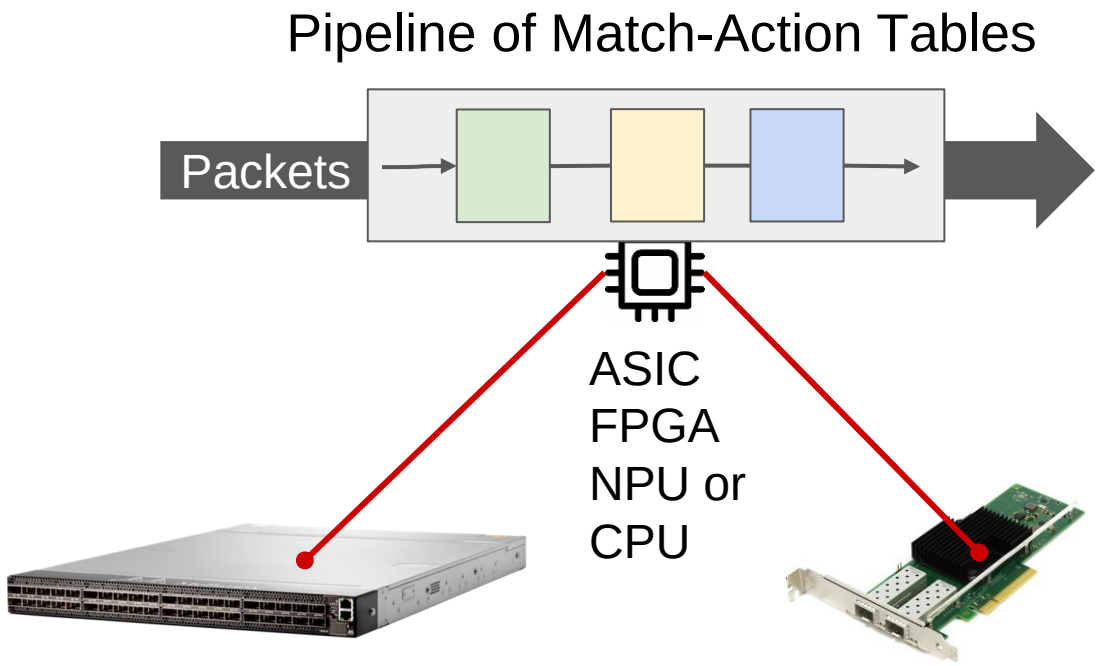

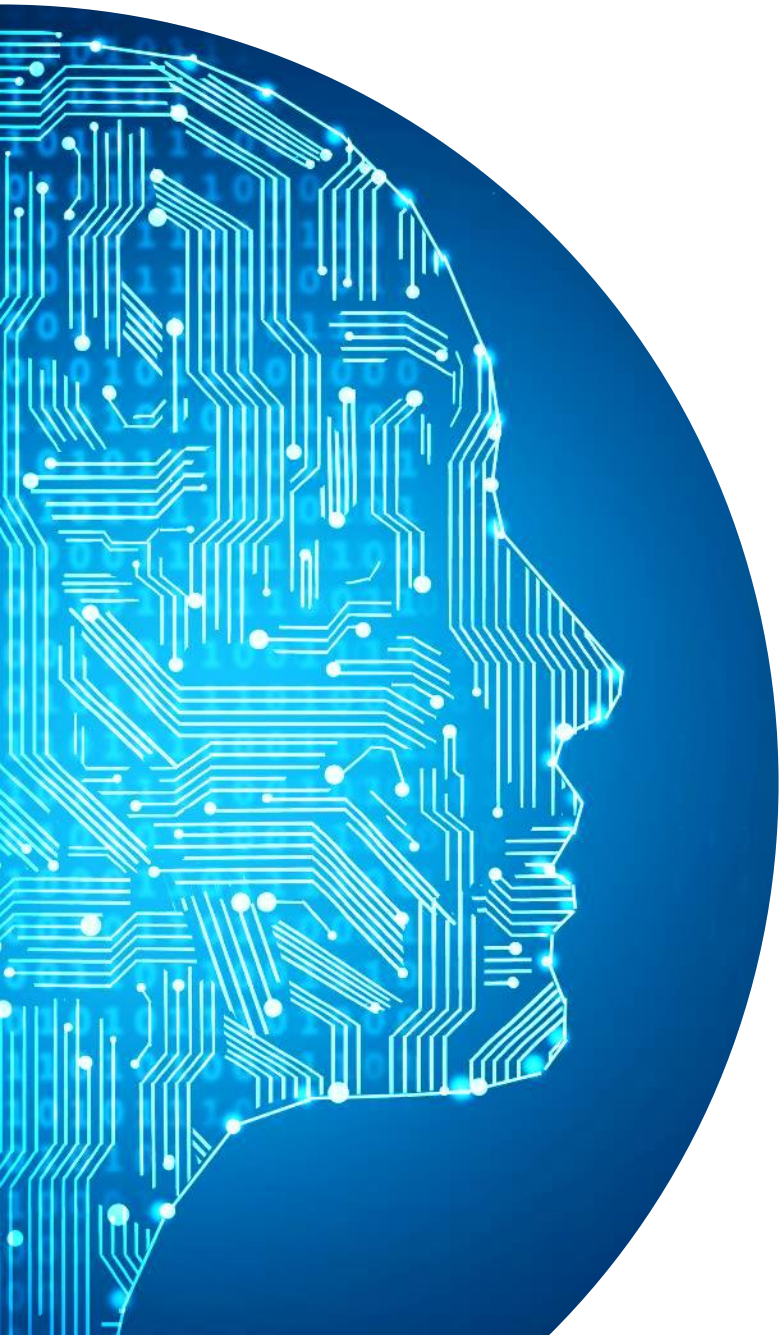

The need for a different SDN

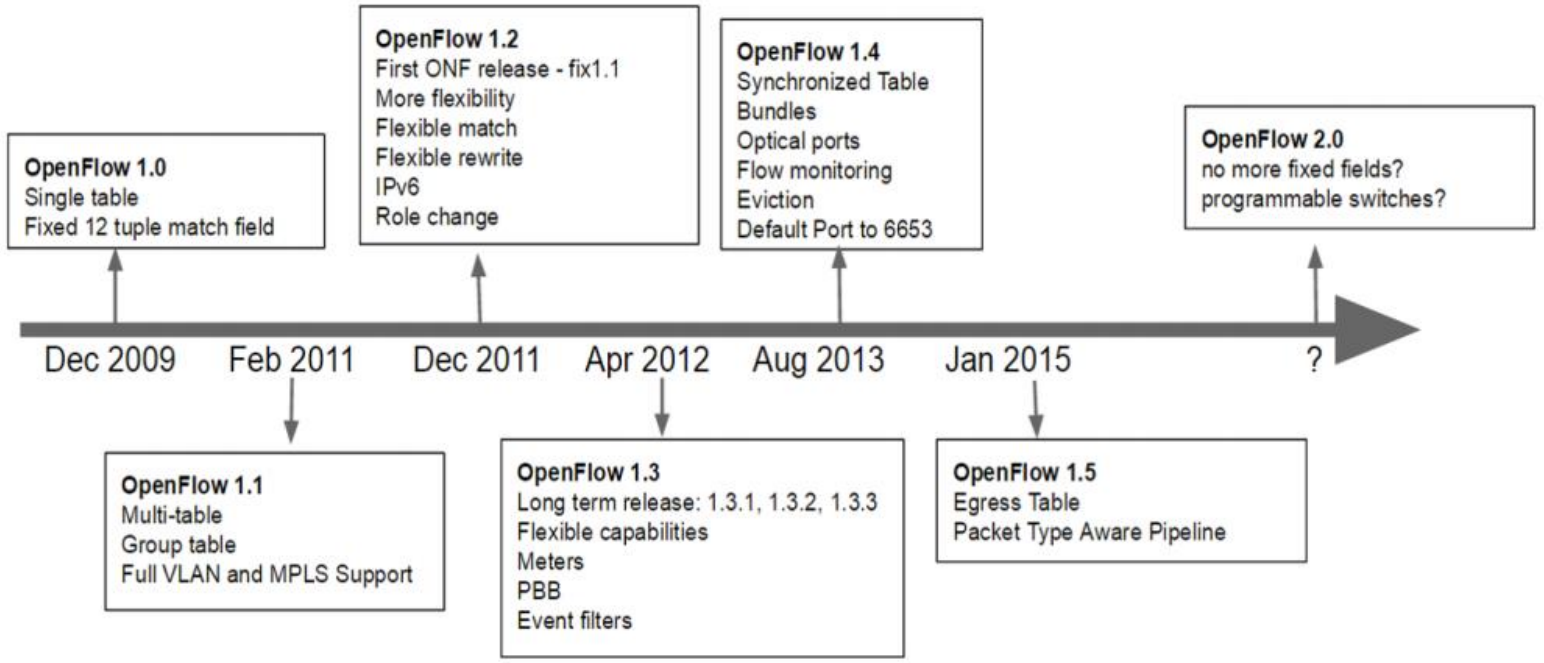

***Figure source:** https://kspviswa.github.io/OpenFlow_Version_Roadmap.html

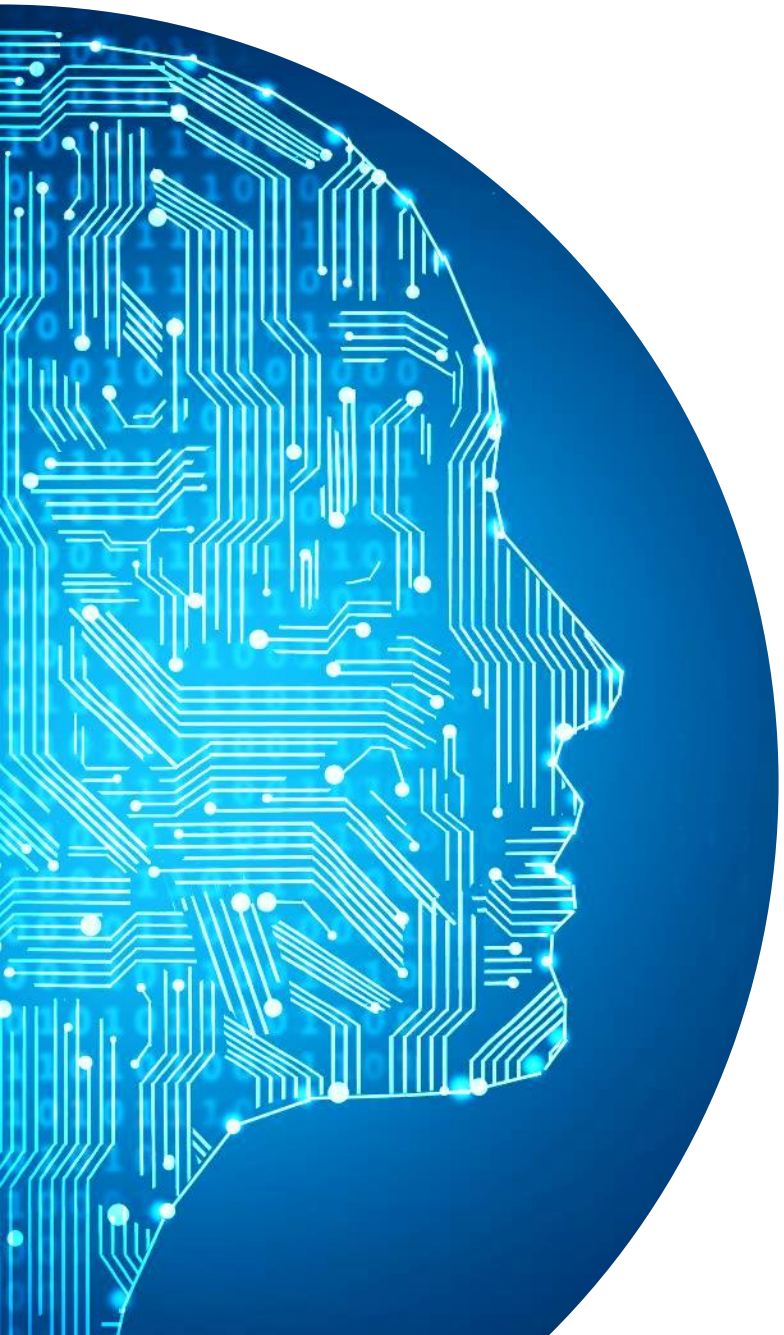

The need for a different SDN

TeraFlow SDN

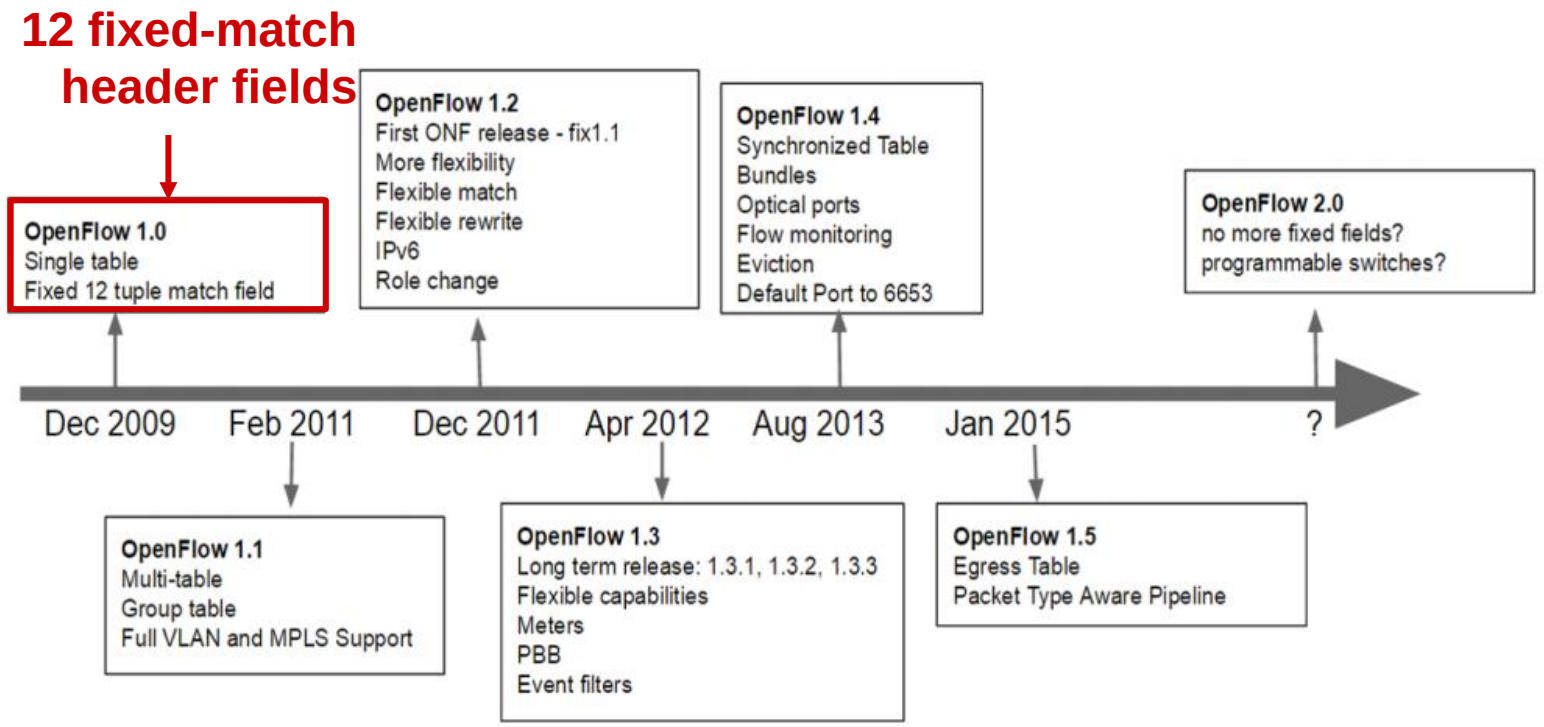

***Figure source:** https://kspviswa.github.io/OpenFlow_Version_Roadmap.html

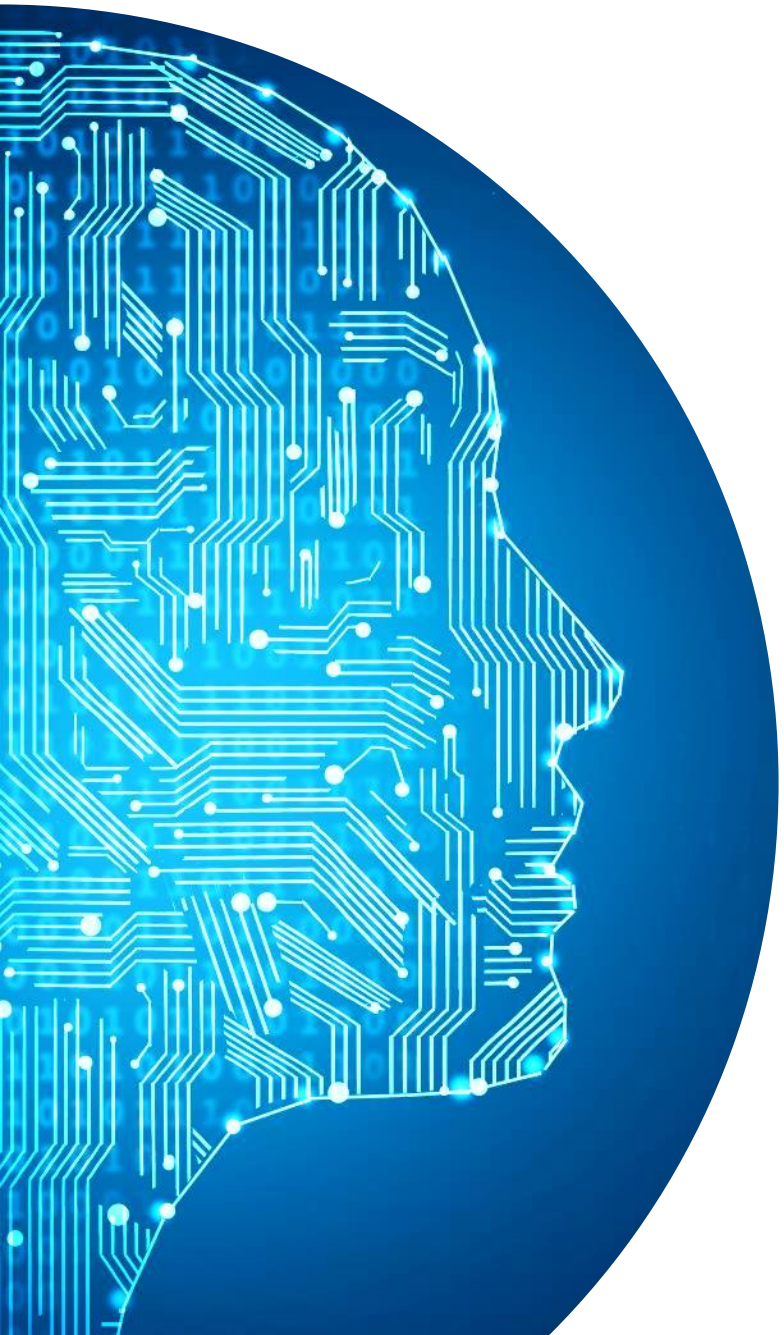

The need for a different SDN

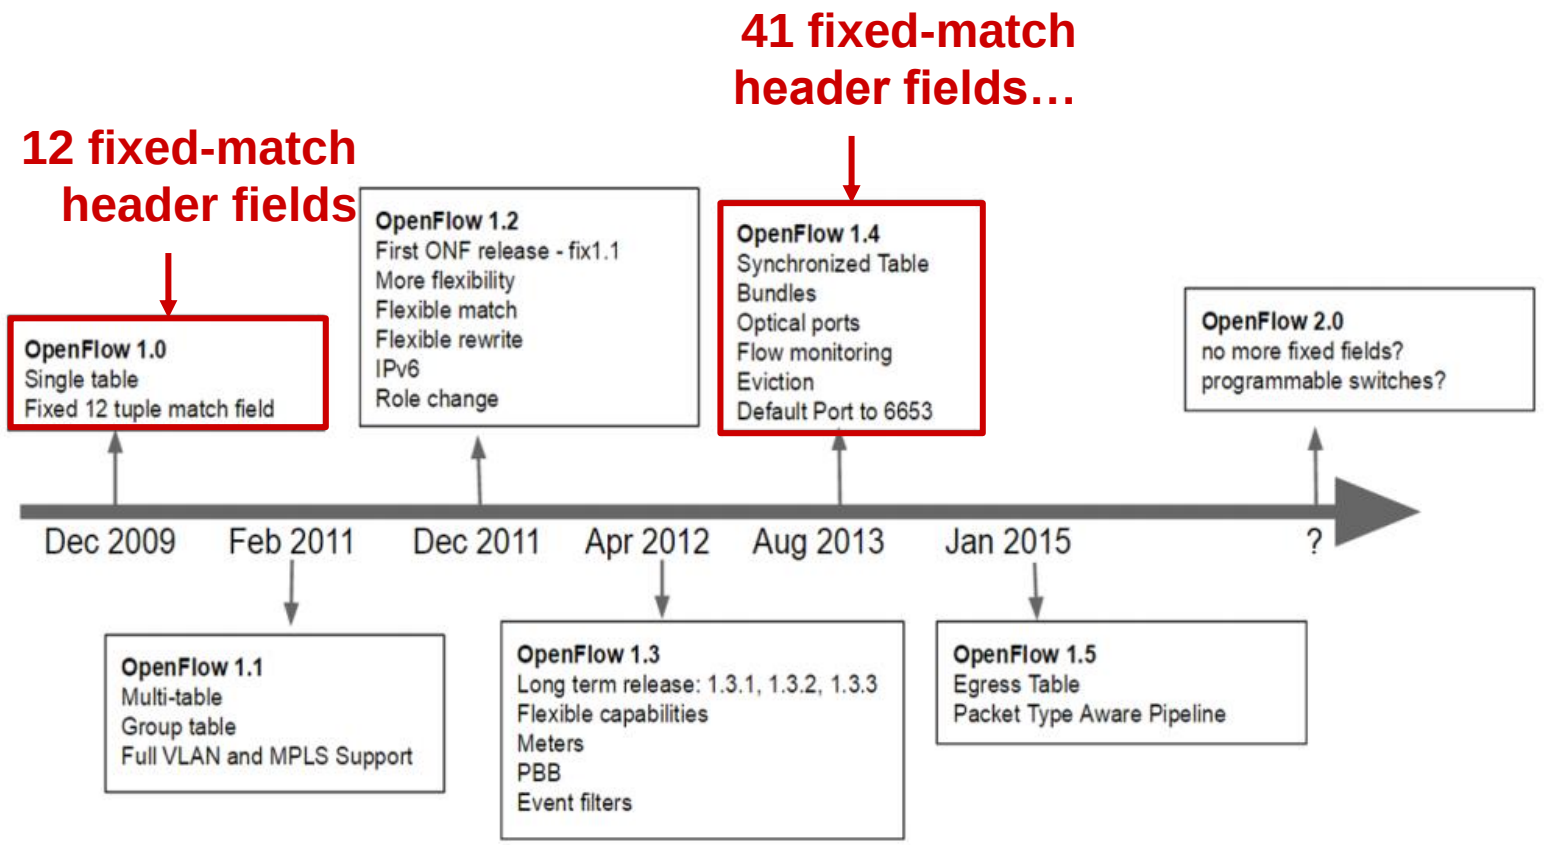

***Figure source:** https://kspviswa.github.io/OpenFlow_Version_Roadmap.html

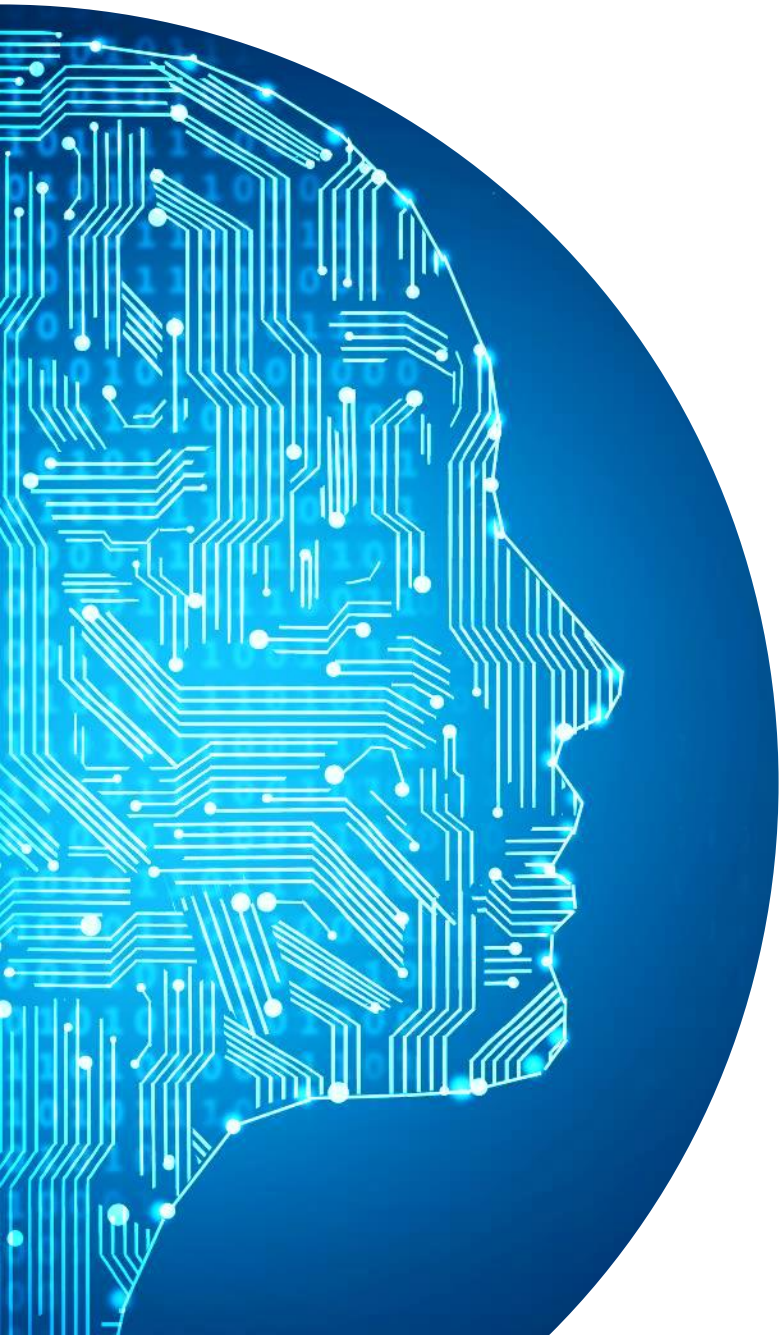

The need for a different SDN

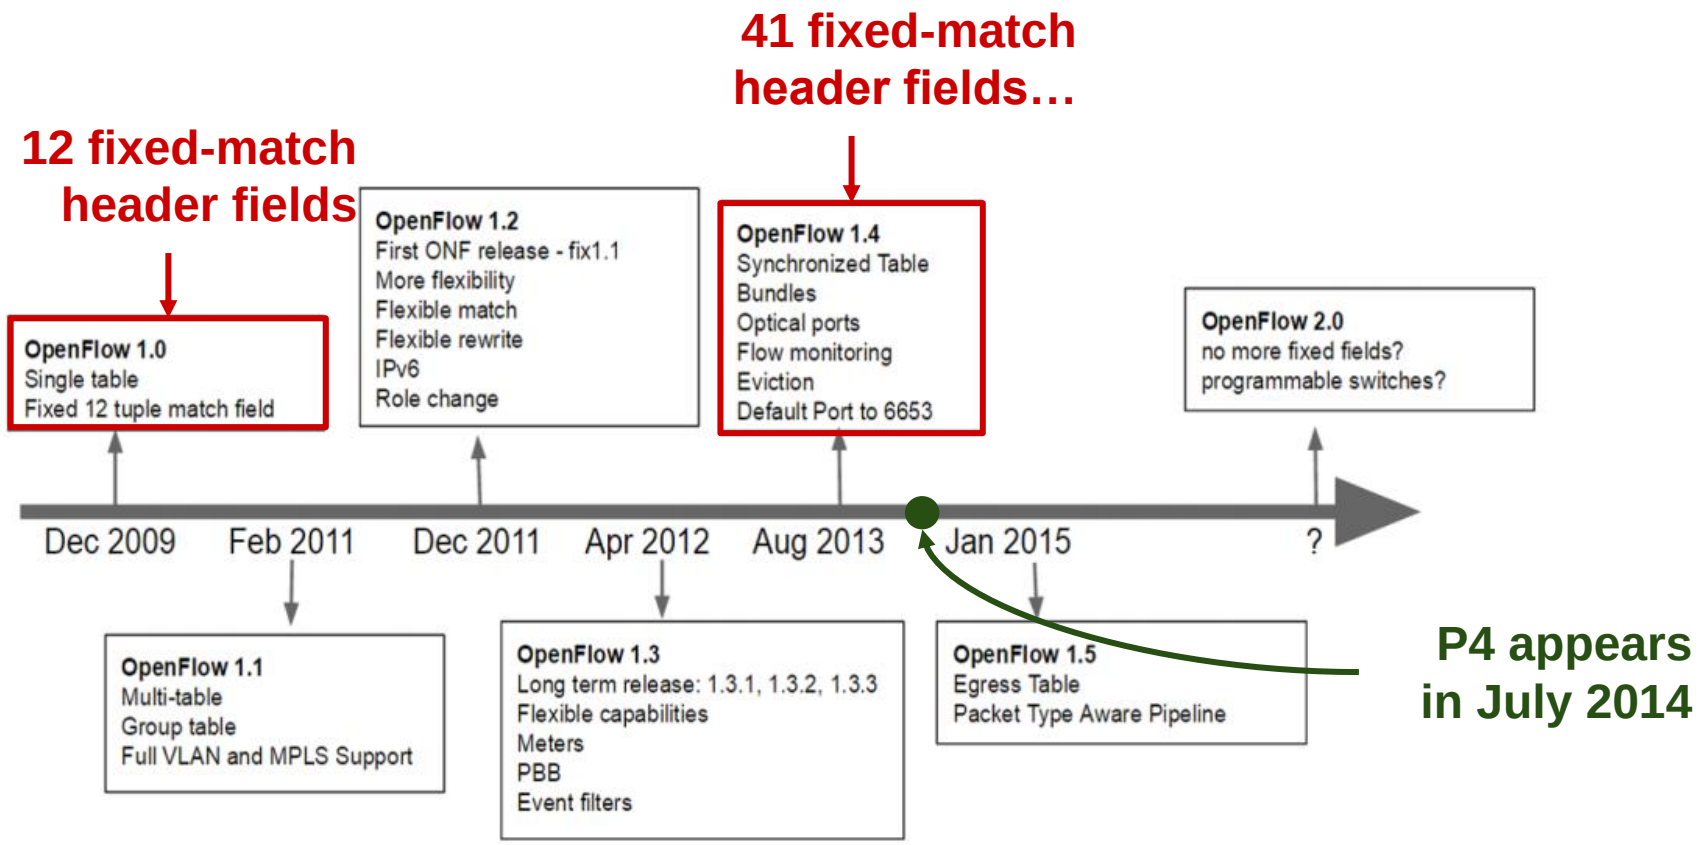

***Figure source:** https://kspviswa.github.io/OpenFlow_Version_Roadmap.html

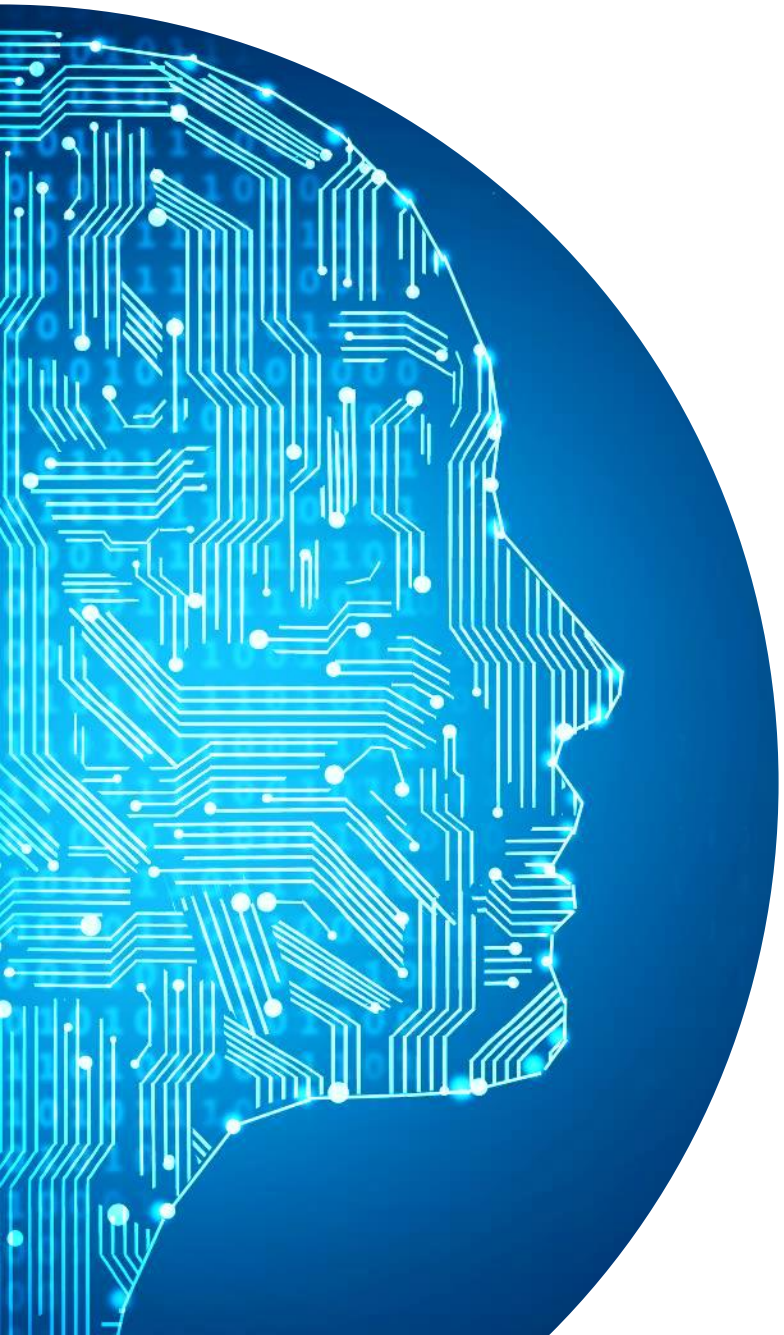

The need for a different SDN

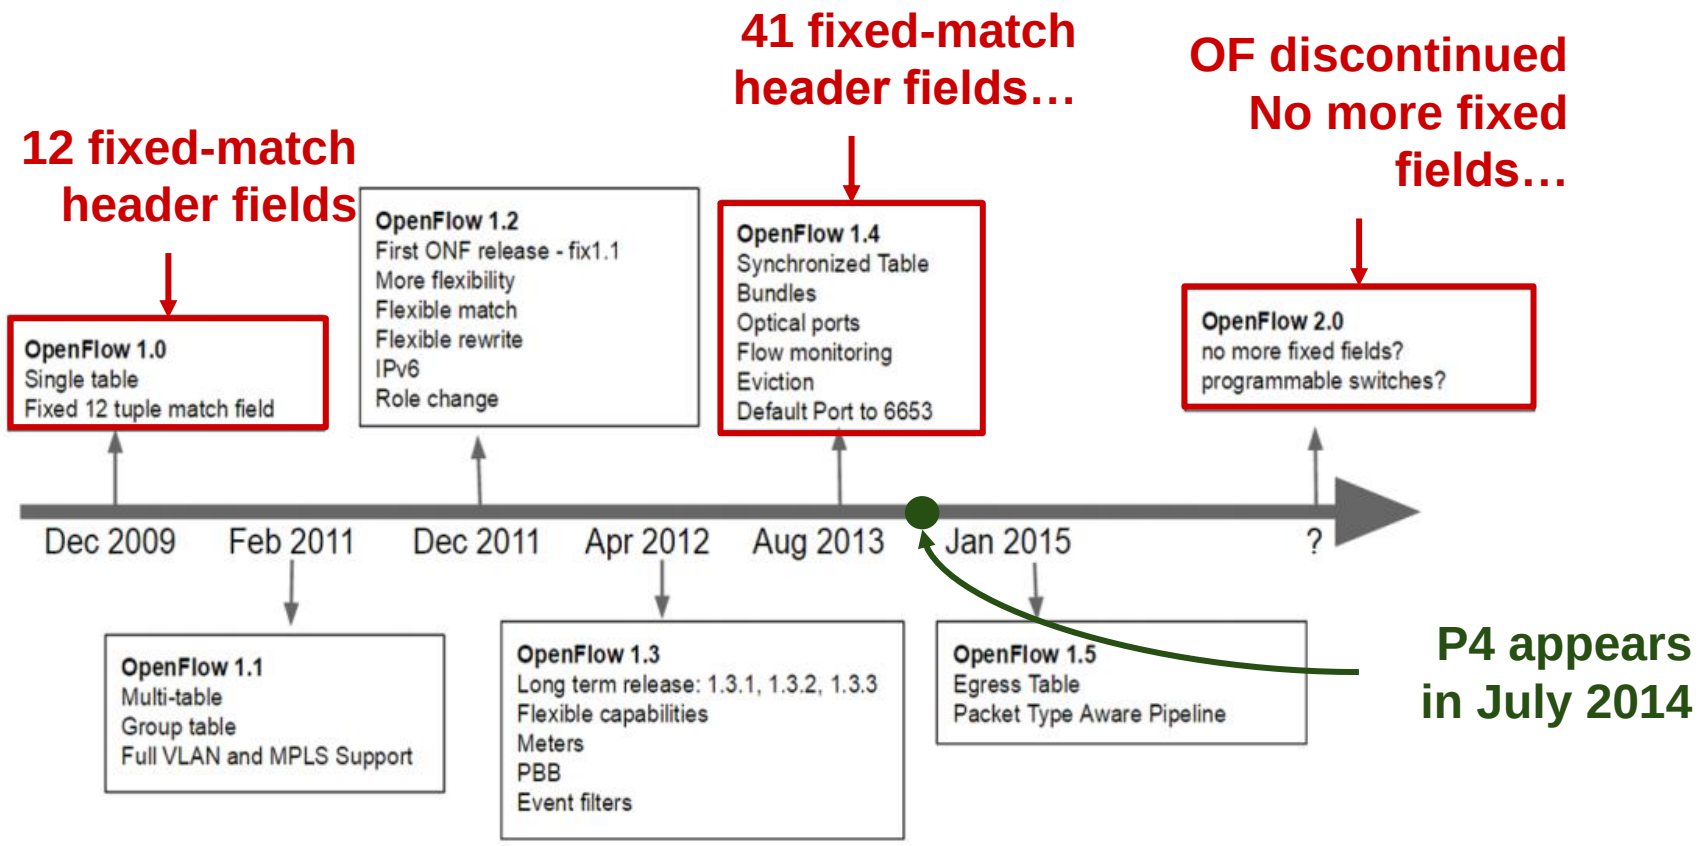

***Figure source:** https://kspviswa.github.io/OpenFlow_Version_Roadmap.html

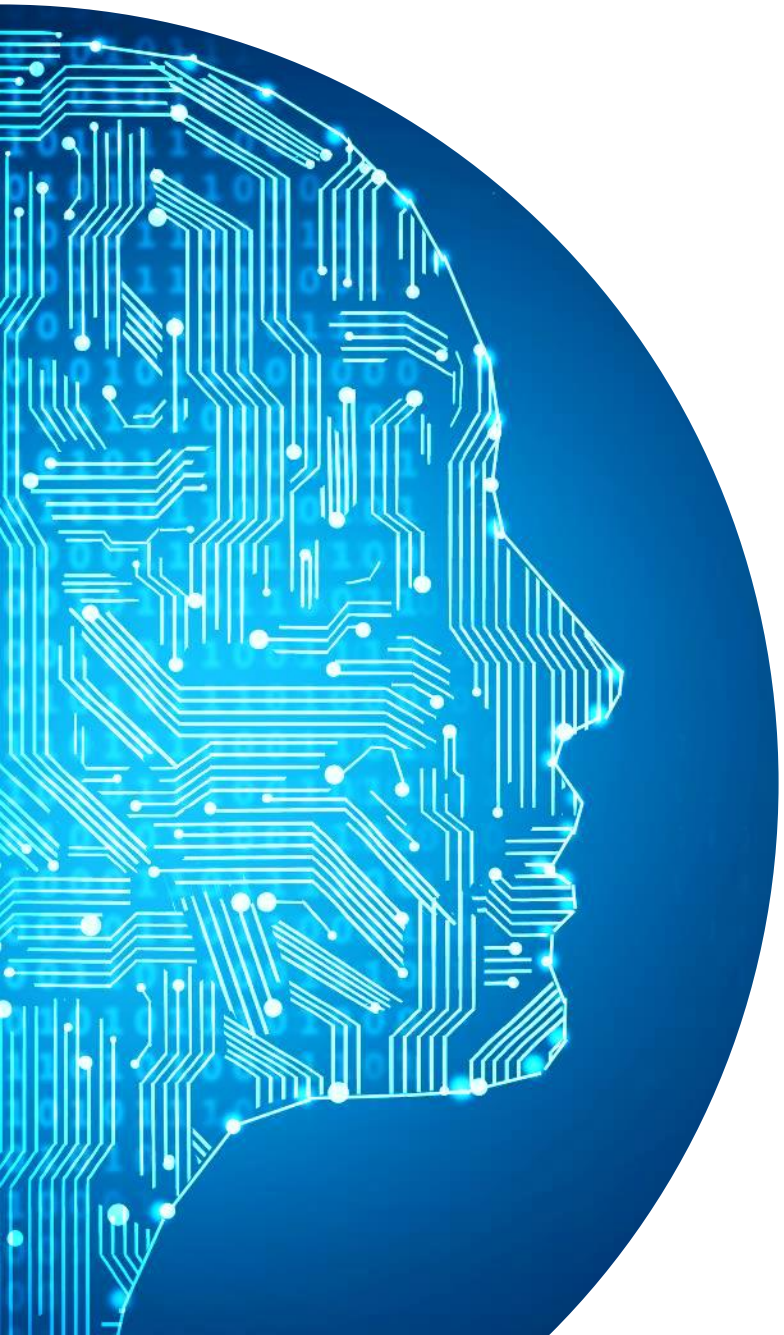

The sad reality about OpenFlow

TeraFlow SDN

Specification	OpenFlow v1.4	41 fixed-match header fields	17 action types

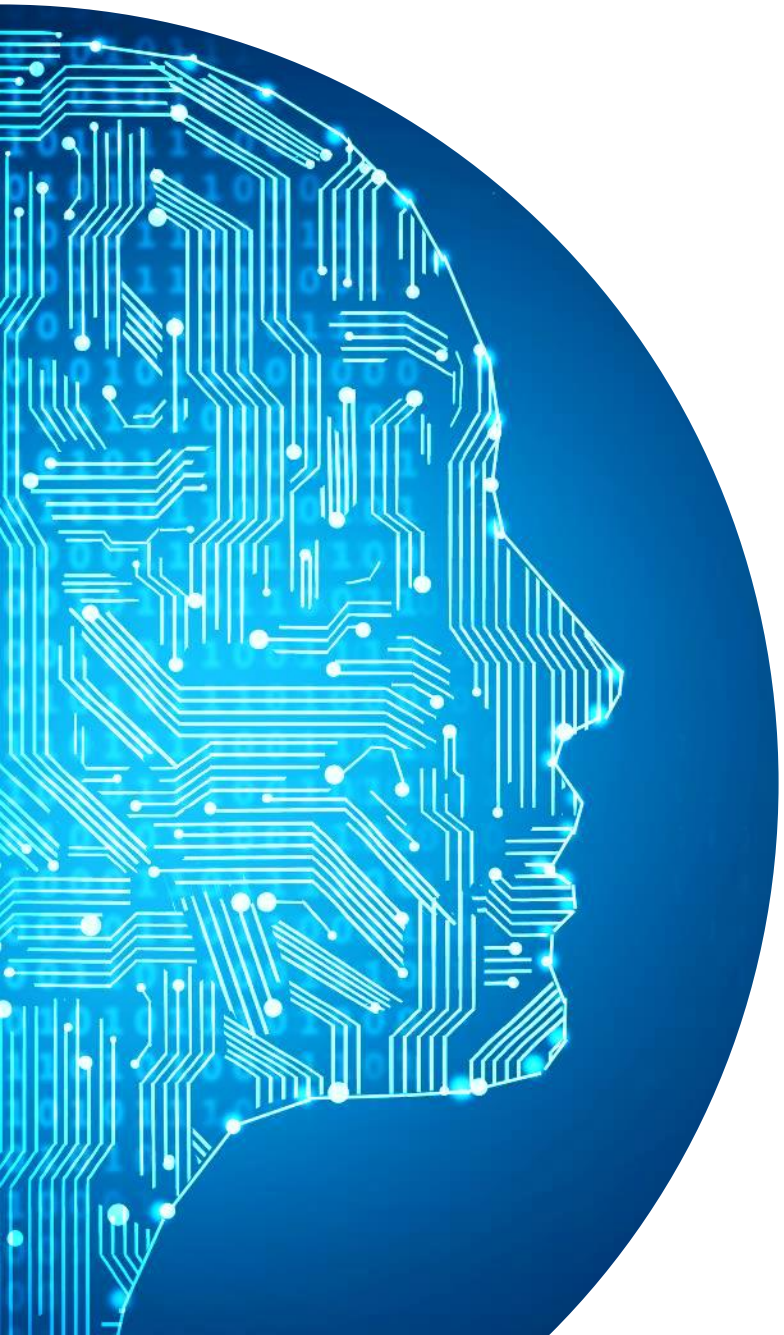

The sad reality about OpenFlow

TeraFlow SDN

Specification	OpenFlow v1.4	41 fixed-match header fields	17 action types

In reality

Most hardware switches only support a limited match/action set due to ASIC limitations

Hardware re-design requires long development cycles and increases cost

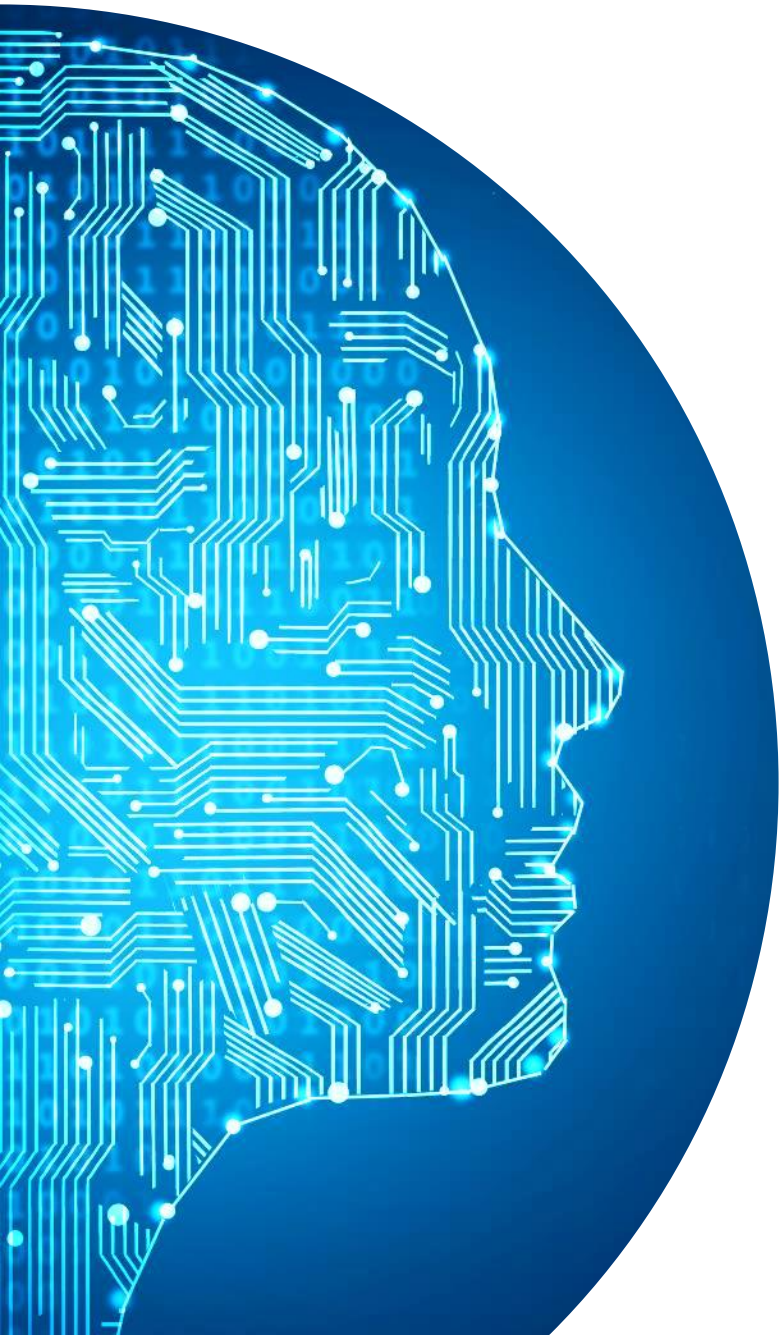

Why P4?

P4 motivation: Instead of repeatedly extending the OpenFlow standards, let's define a whole new abstraction for programming the data plane

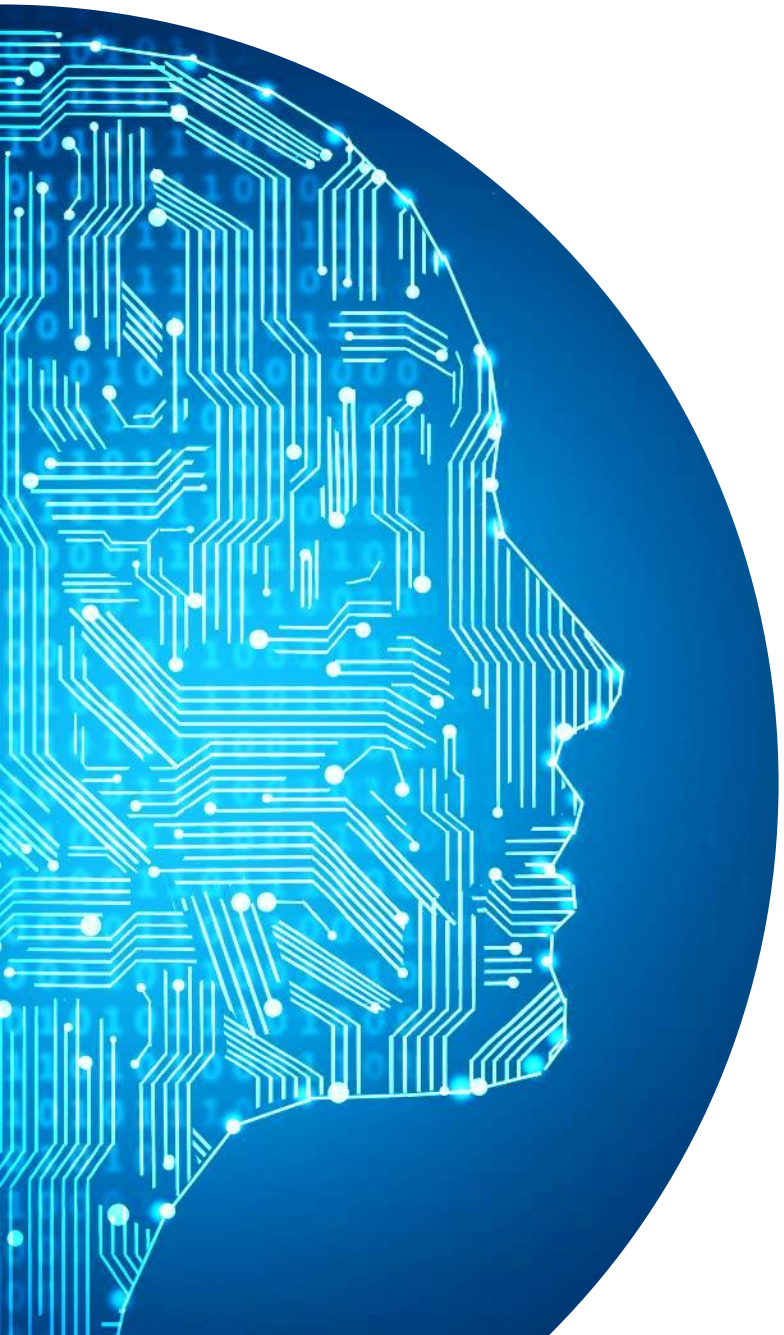

Why P4?

TeraFlow SDN

P4 motivation: Instead of repeatedly extending the OpenFlow standards, let's define a whole new abstraction for programming the data plane

P4 principles:

- a domain-specific language to formally define the data plane pipeline
  - Describes proto headers (existing+new), tables, actions, counters, etc.
  - Describes both fast (ASIC, FPGA) and slow (e.g., soft. switch) pipelines

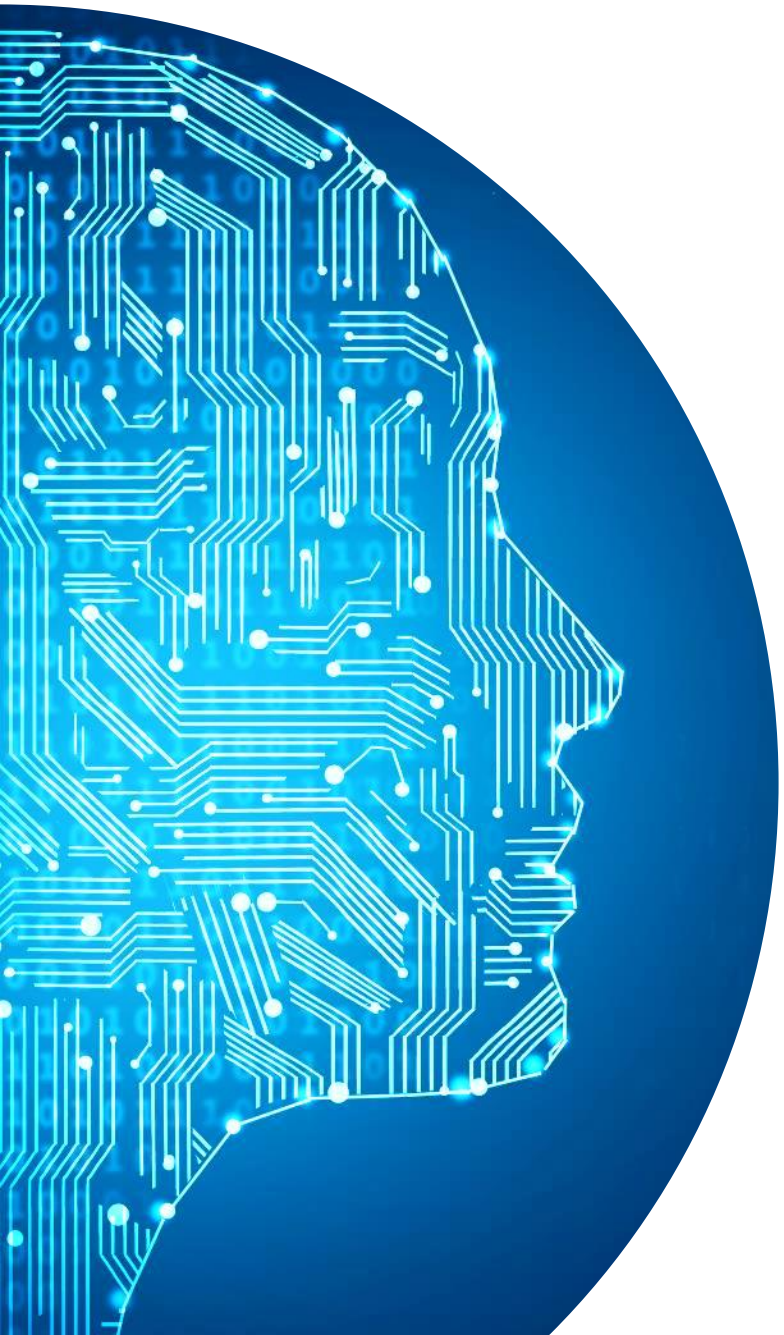

Why P4?

TeraFlow SDN

P4 motivation: Instead of repeatedly extending the OpenFlow standards, let's define a whole new abstraction for programming the data plane

P4 principles:

- a domain-specific language to formally define the data plane pipeline
  - Describes proto headers (existing+new), tables, actions, counters, etc.
  - Describes both fast (ASIC, FPGA) and slow (e.g., soft. switch) pipelines
- a common interface to parse packets and match (arbitrary) header fields
  - Defines a "contract" between the control and data plane

PISA: Protocol Independent Switch Architecture

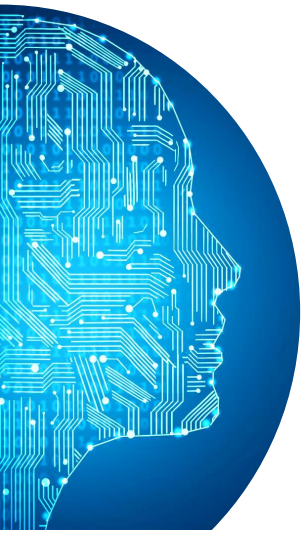

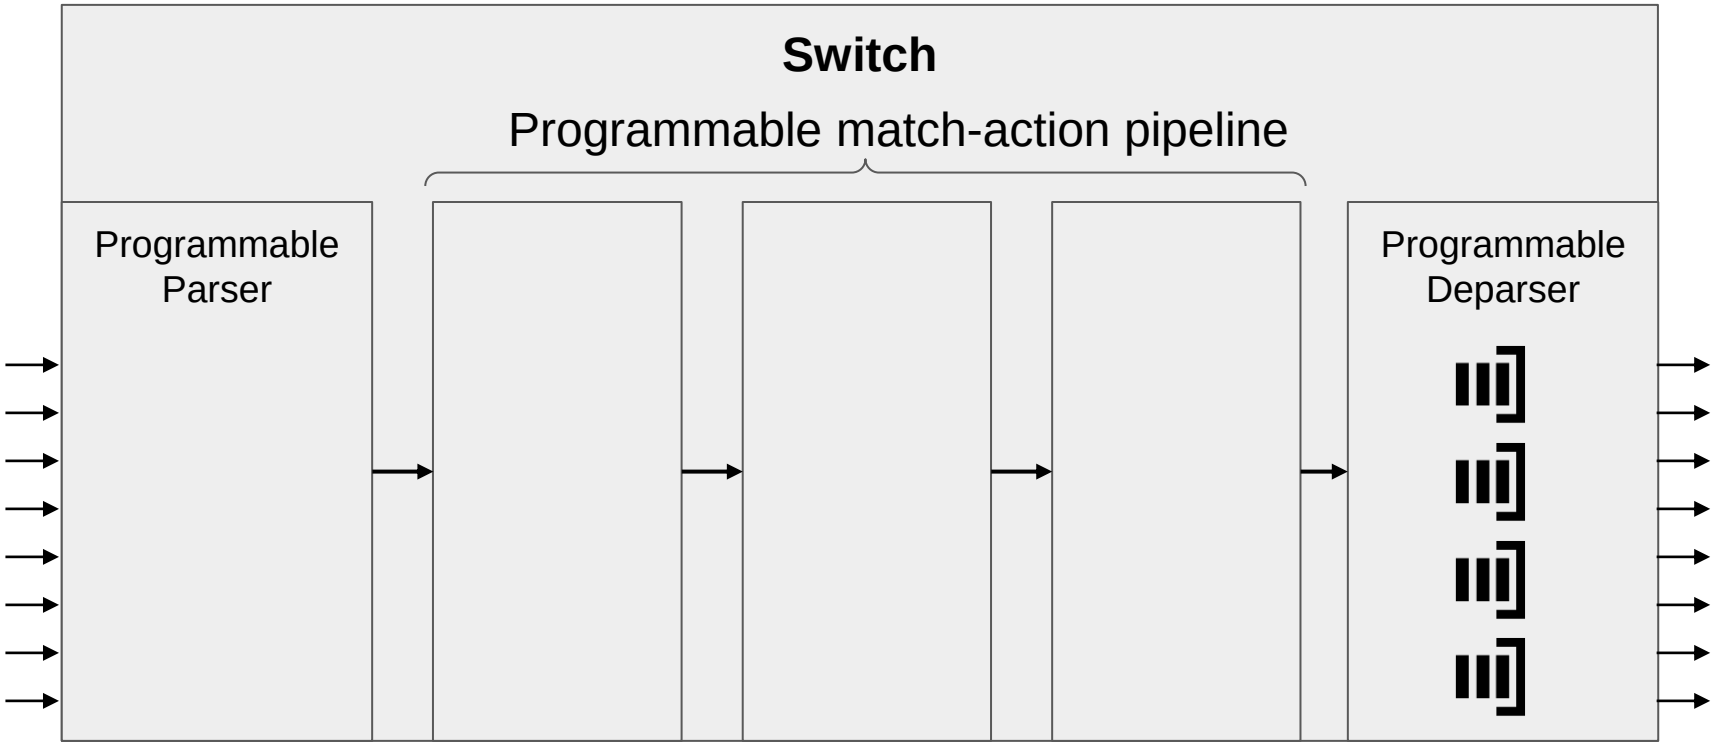

PISA: Protocol Independent Switch Architecture

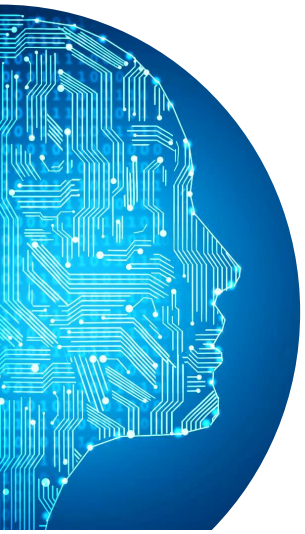

L2-L3.p4 program

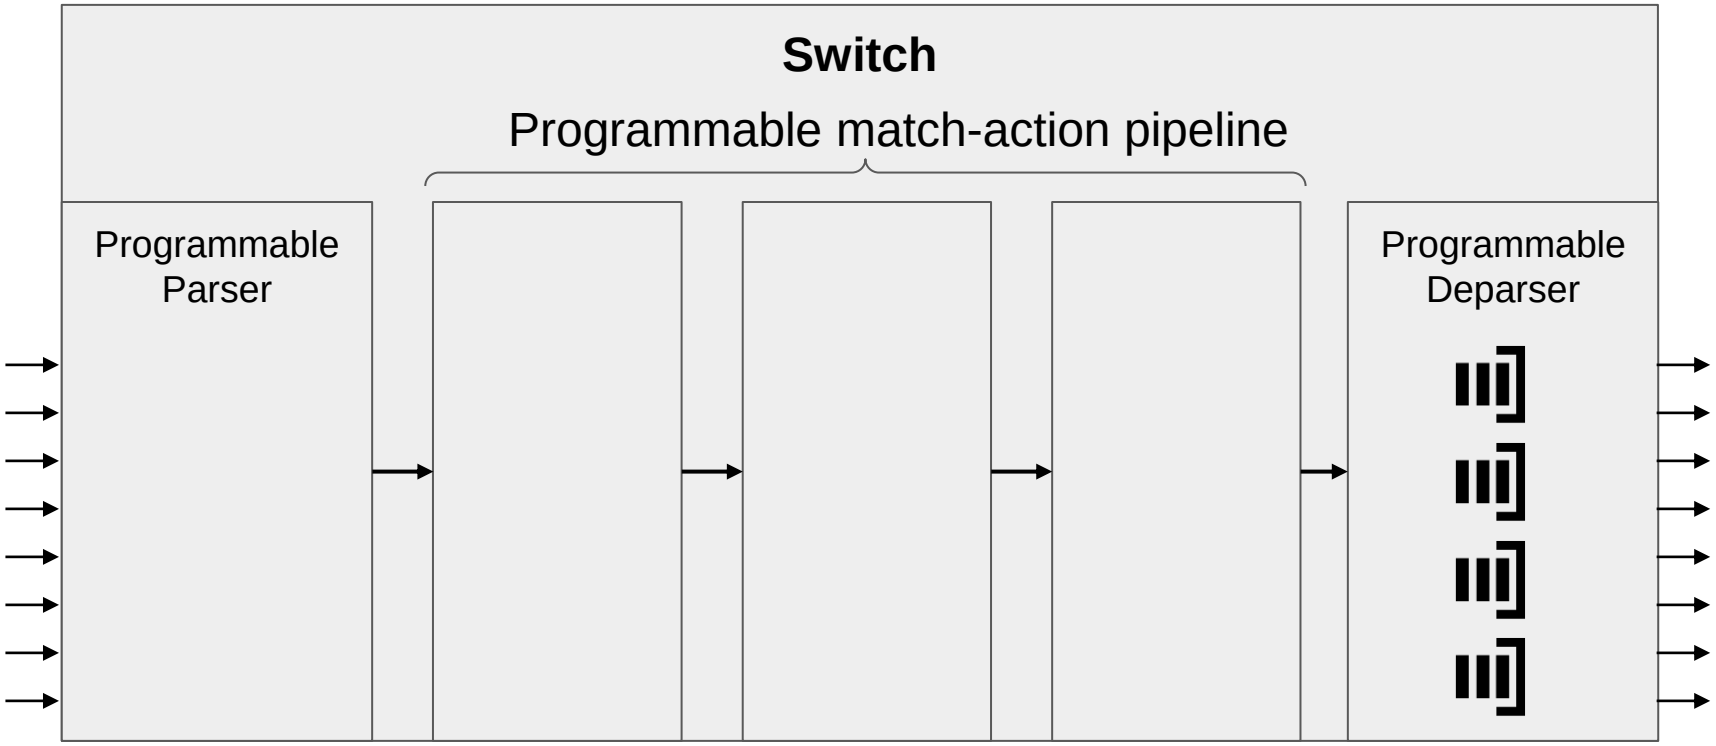

PISA: Protocol Independent Switch Architecture

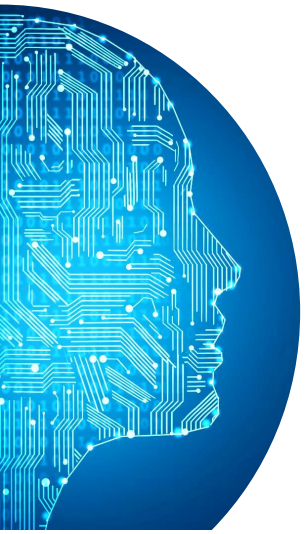

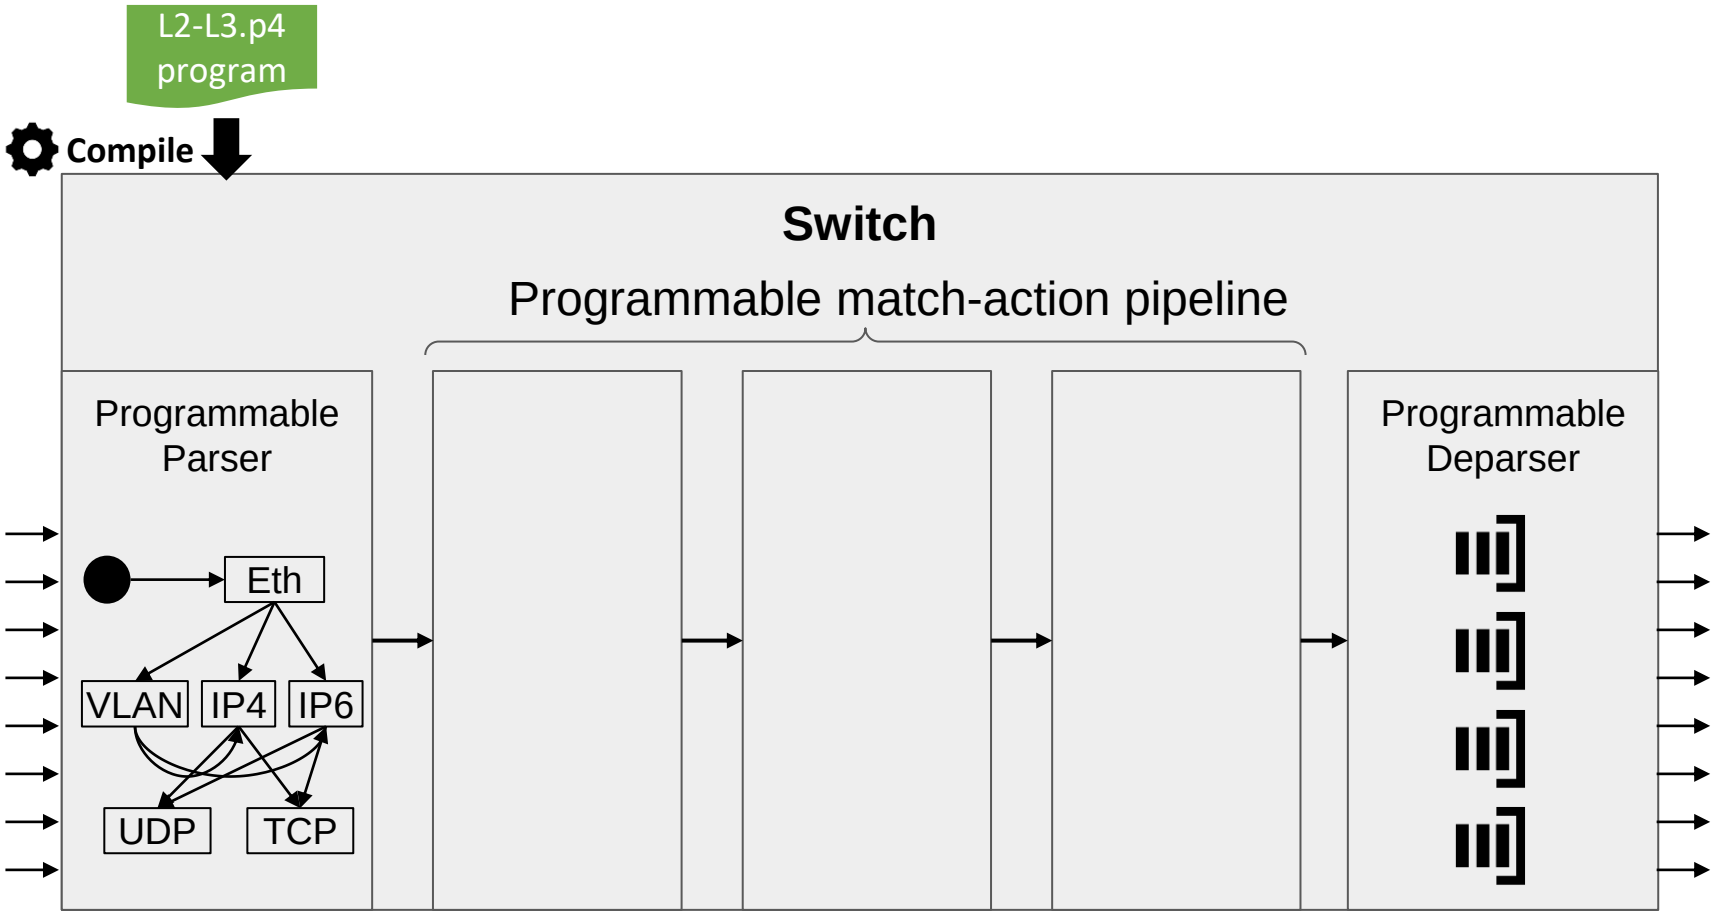

PISA: Protocol Independent Switch Architecture

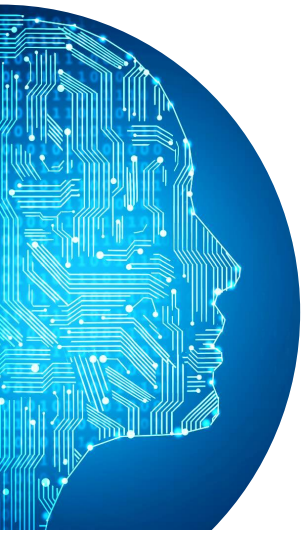

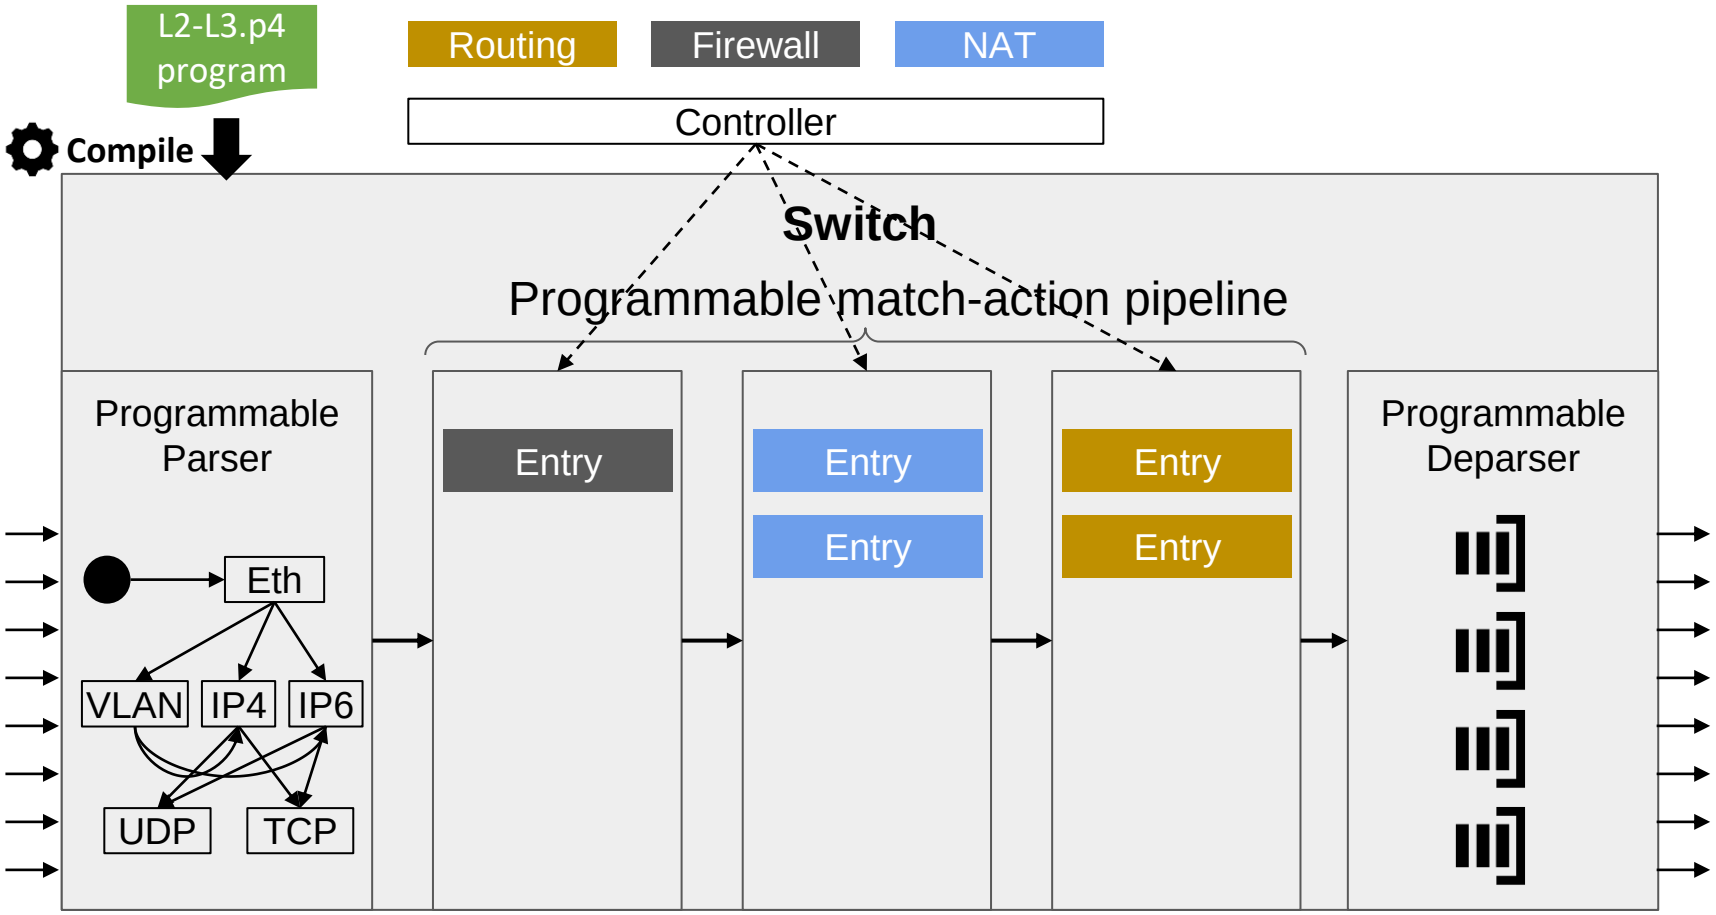

PISA: Protocol Independent Switch Architecture

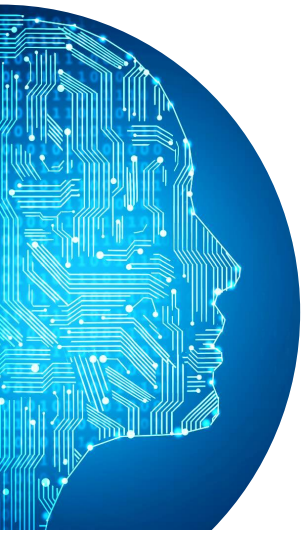

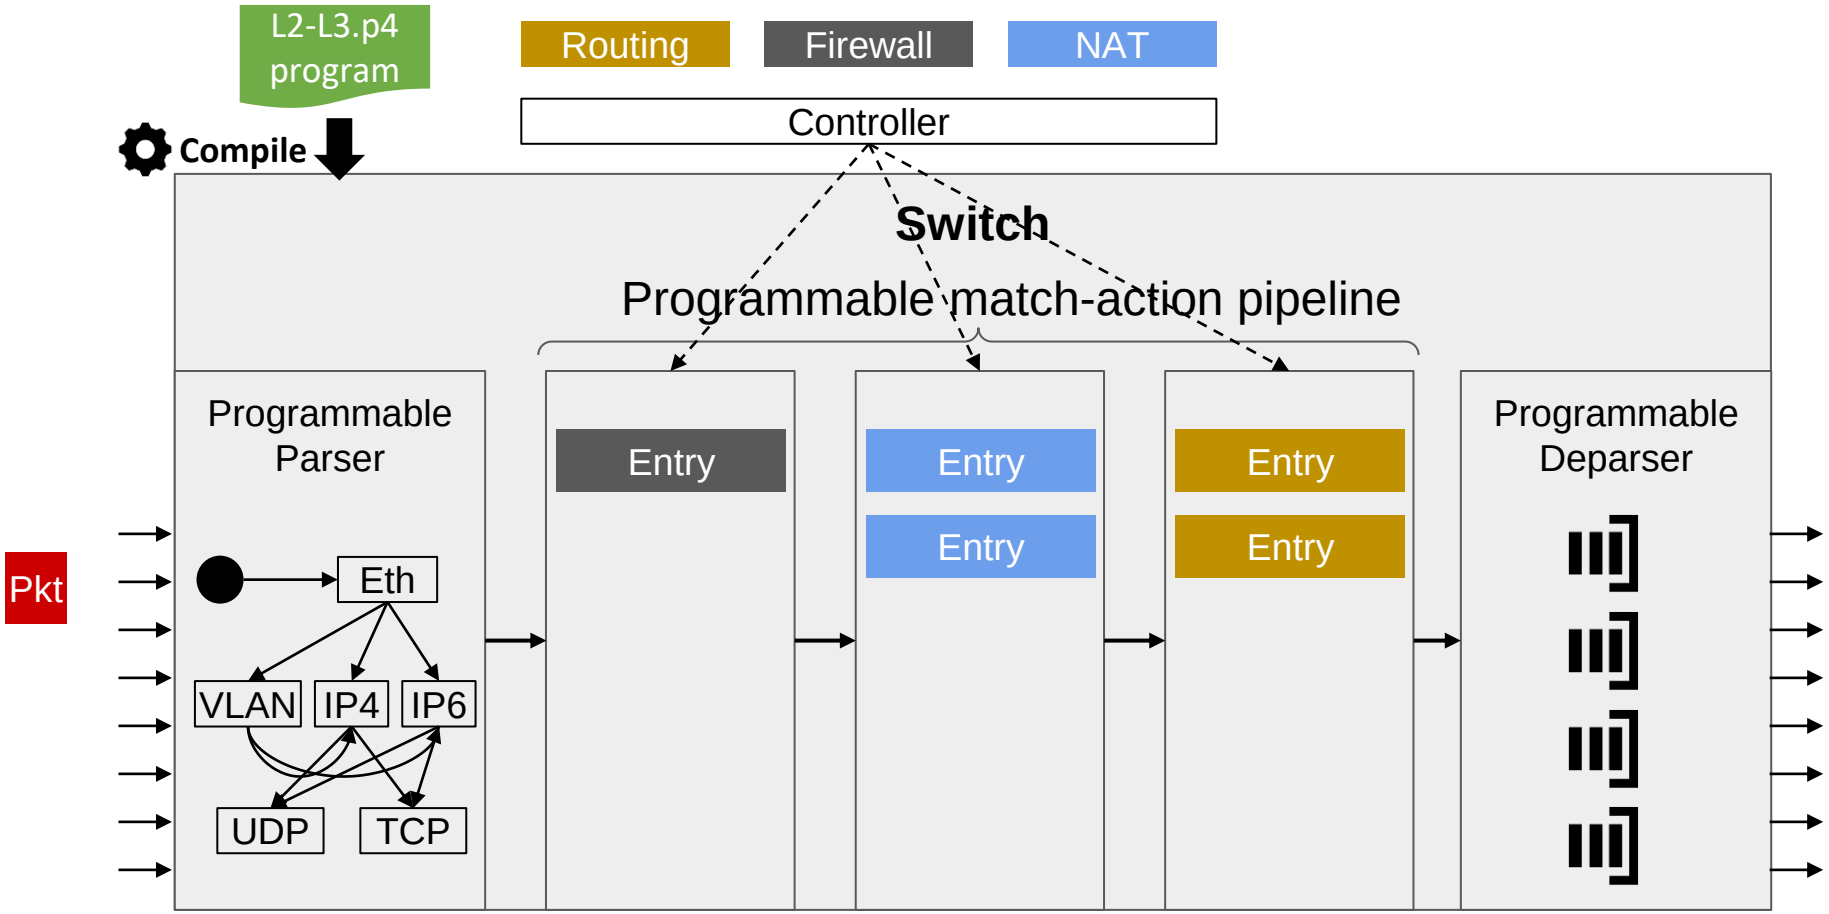

PISA: Protocol Independent Switch Architecture

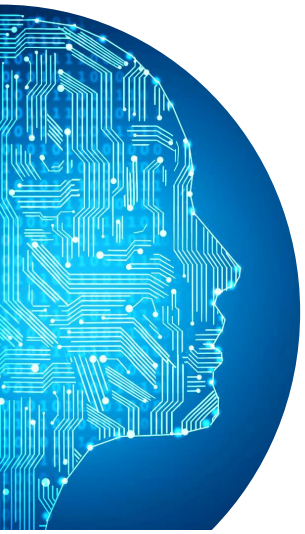

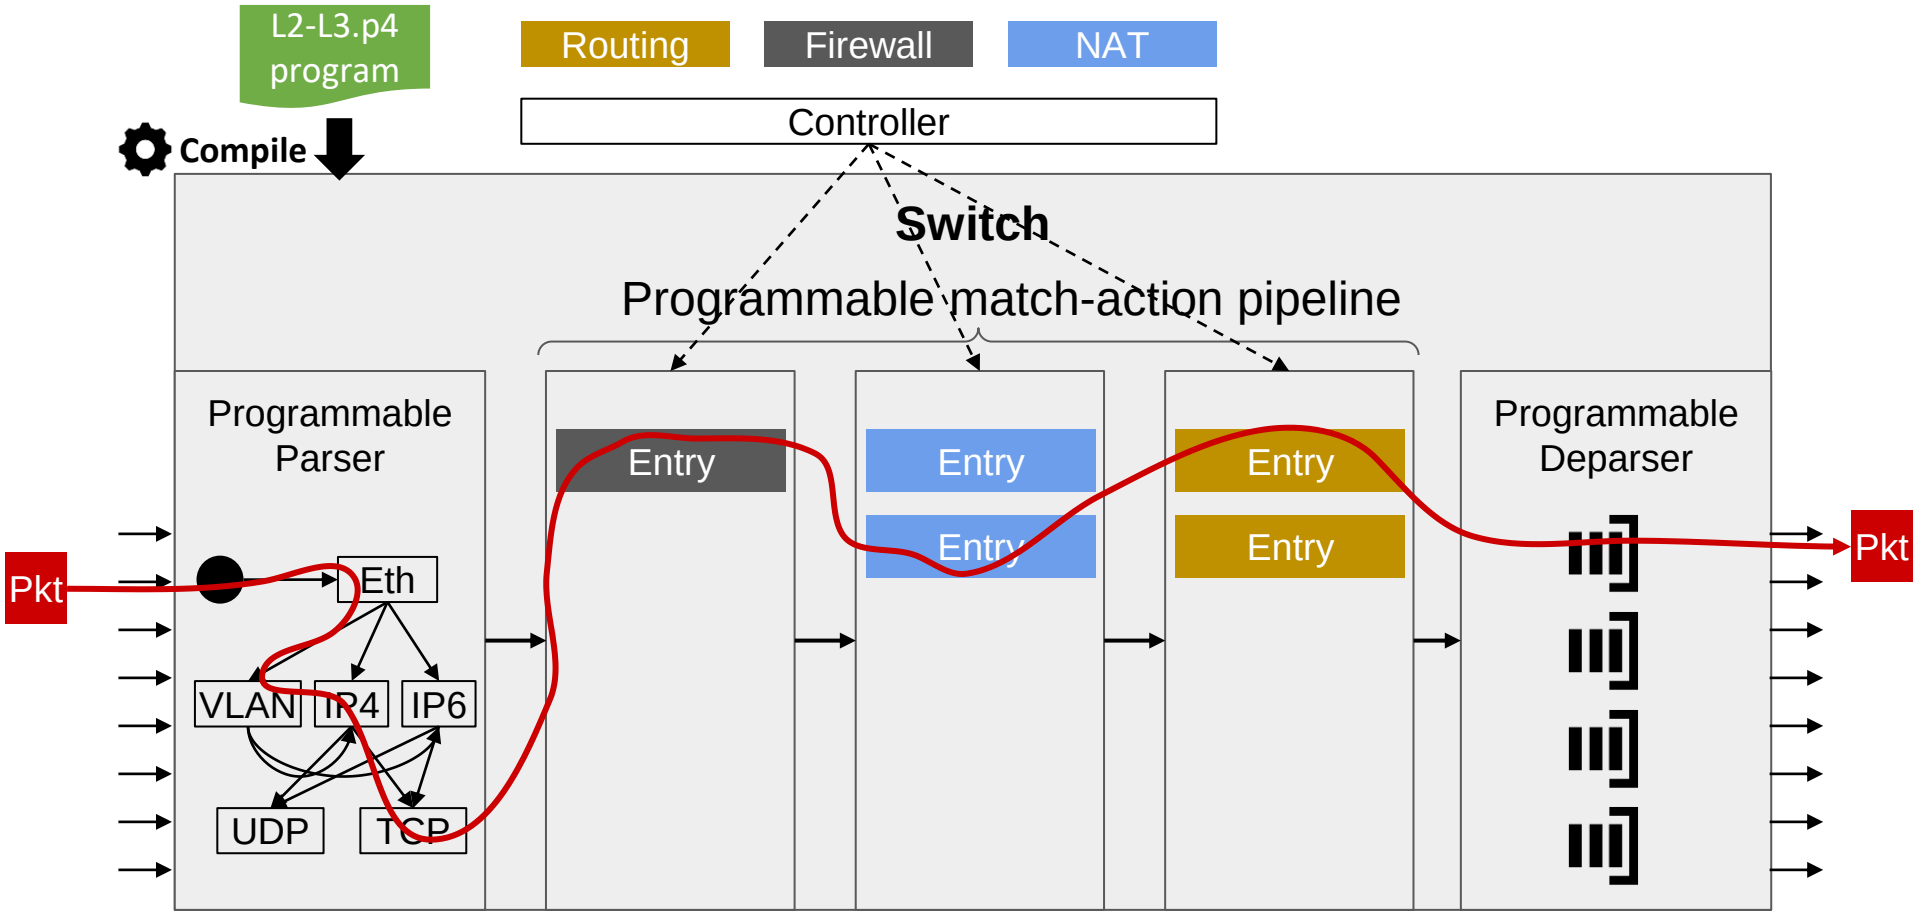

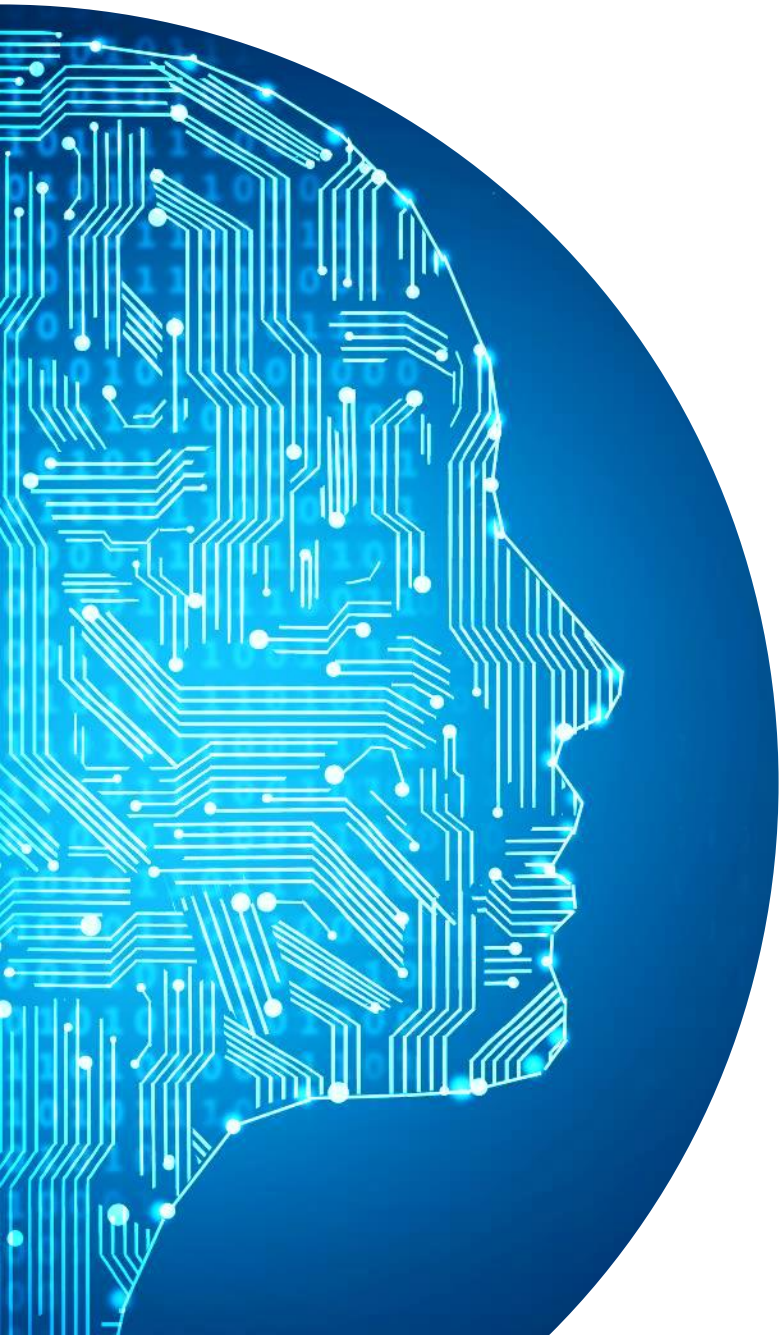

But switches still have ASICS…

TeraFlow SDN

Correct, but…

- New custom ASICs offer decent flexibility even at several Tbits/sec

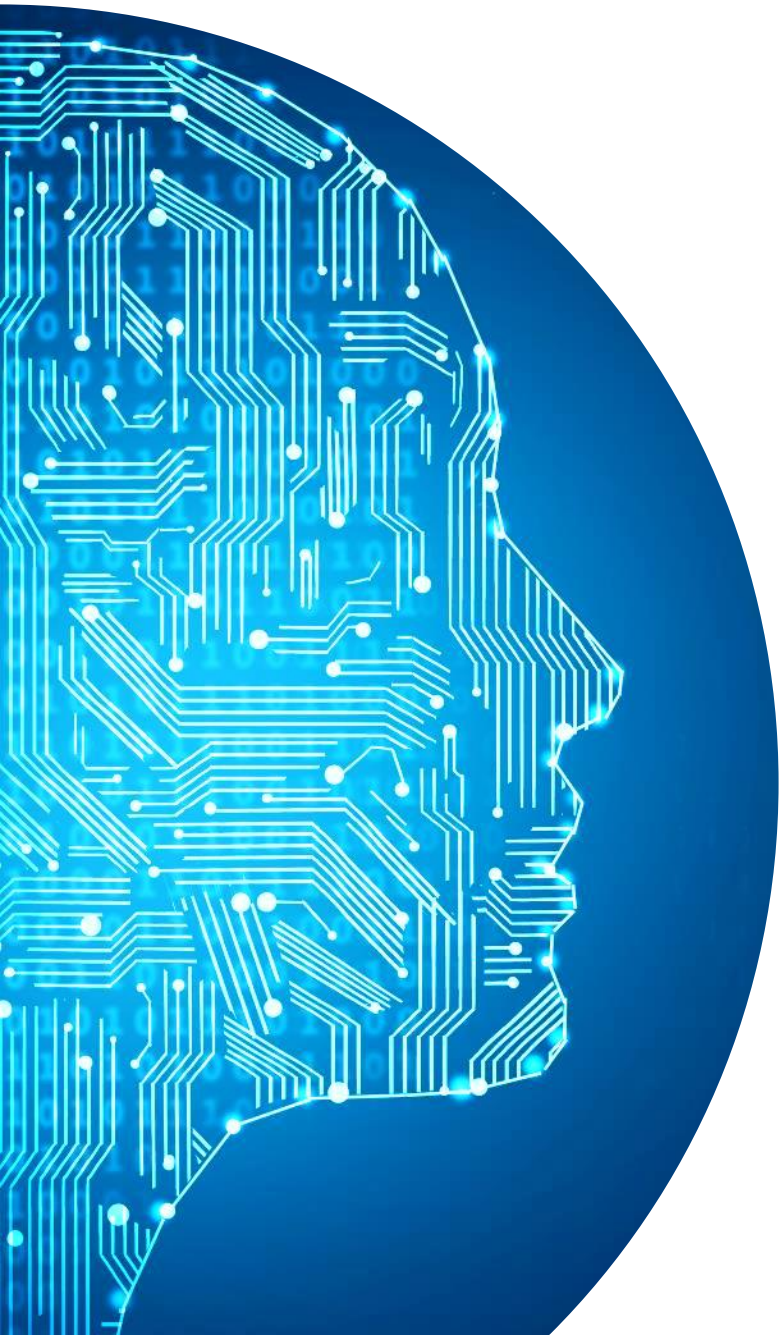

But switches still have ASICS…

TeraFlow SDN

Correct, but…

- New custom ASICs offer decent flexibility even at several Tbits/sec
- Some switches offer higher programmability than others:
  - FPGAs (e.g., Intel, Xilinx)
  - NPUs (e.g., Netronome, EZchip)
  - Software-based (e.g., OVS)

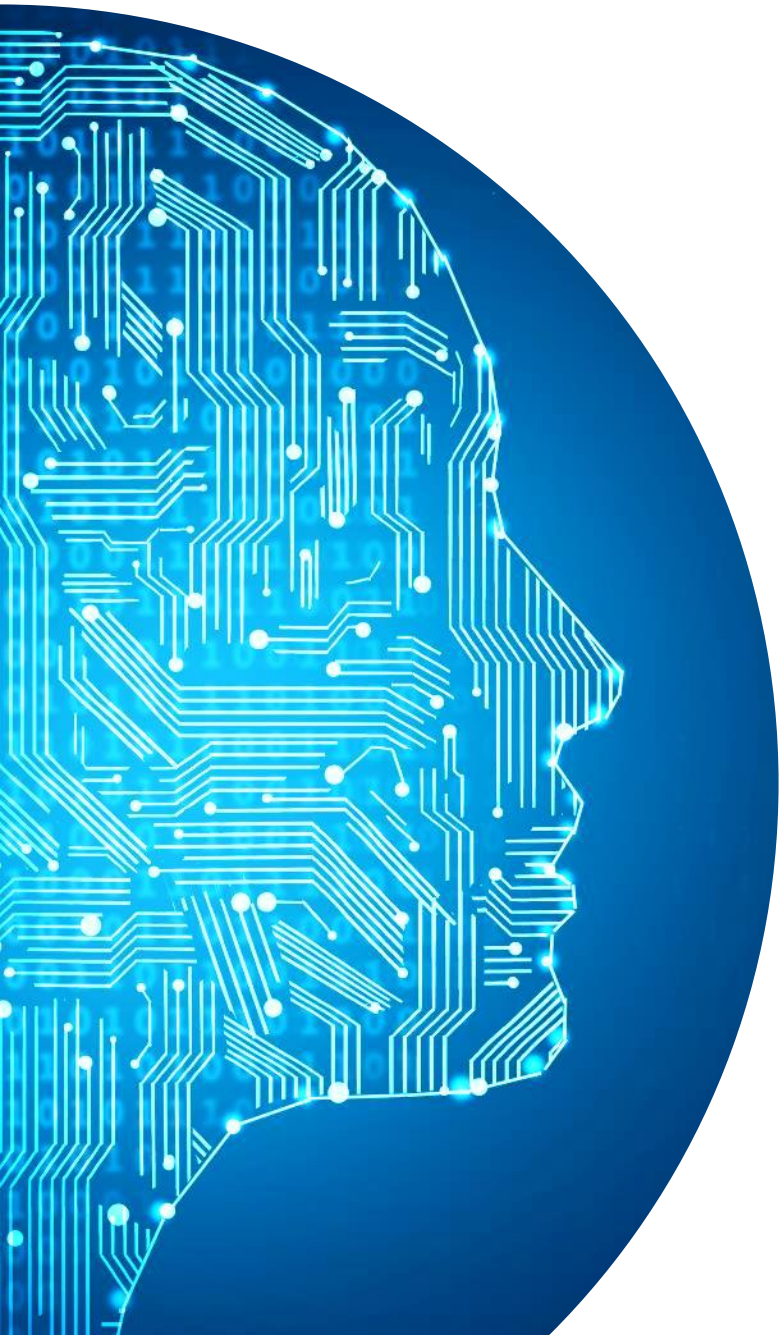

P4Runtime: Not just yet another P4 API?

TeraFlow SDN

P4Runtime is a runtime control API for P4-defined data planes

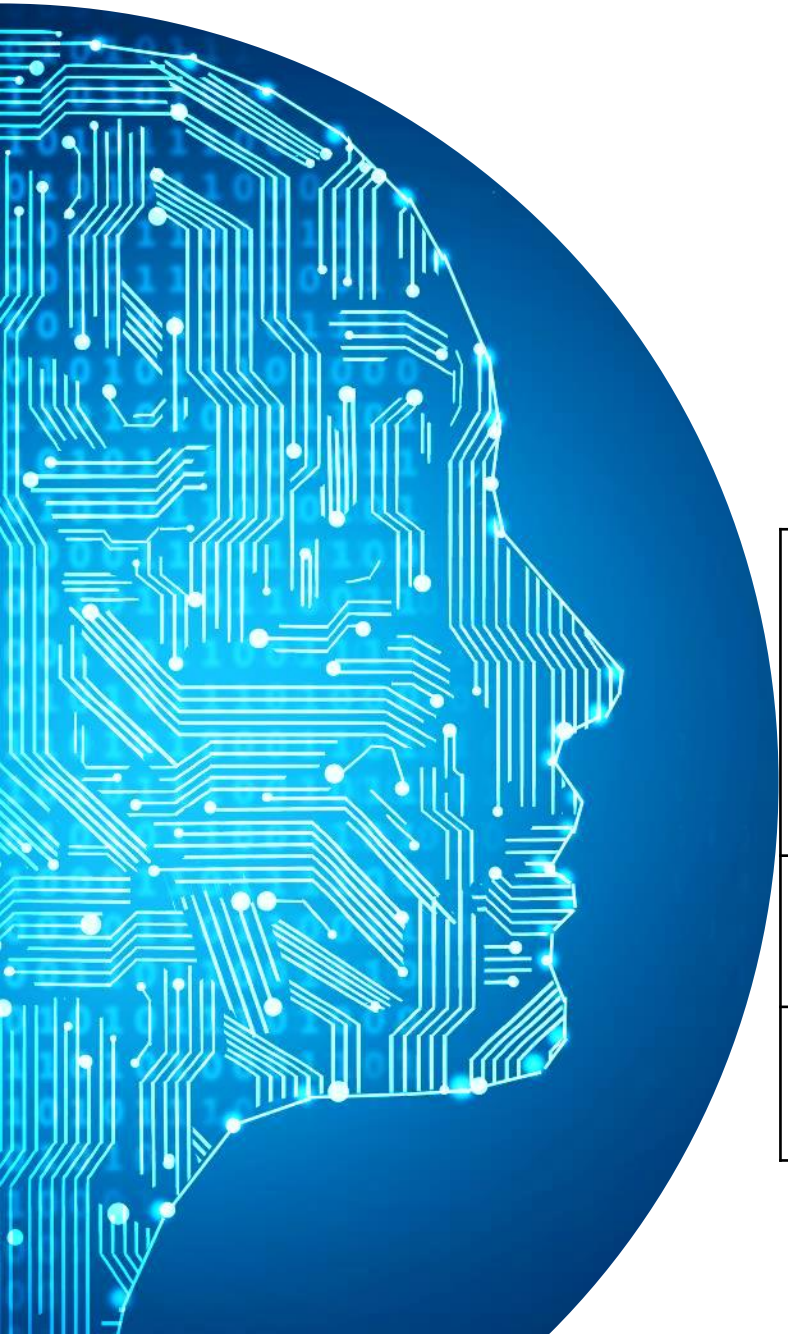

P4Runtime: Not just yet another P4 API?

TeraFlow SDN

P4Runtime is a runtime control API for P4-defined data planes

API	Target independent Same API works with different switches from different vendors	Pipeline independent Same API allows control of many arbitrary pipelines	Protocol independent Same API allows control of any data plane proto (standard/ custom)
OpenFlow	✓	✓ With table type patterns (TTP)	✗ Proto headers and actions hardcoded in the spec
P4Runtime			

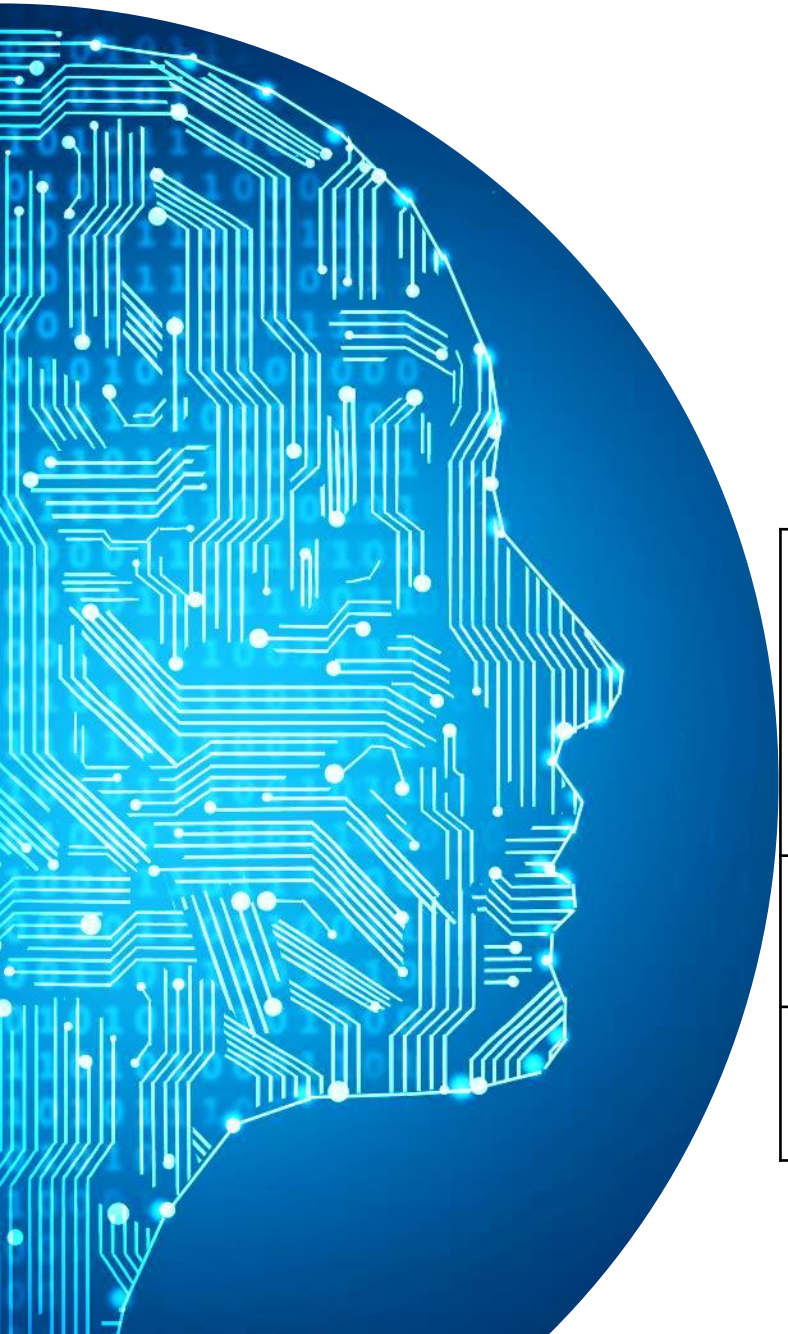

P4Runtime: Not just yet another P4 API?

TeraFlow SDN

P4Runtime is a runtime control API for P4-defined data planes

API	Target independent Same API works with different switches from different vendors	Pipeline independent Same API allows control of many arbitrary pipelines	Protocol independent Same API allows control of any data plane proto (standard/ custom)
OpenFlow	✓	✓ With table type patterns (TTP)	✗ Proto headers and actions hardcoded in the spec
P4Runtime	✓	✓ With P4	✓

P4 Workflow

P4 program

1

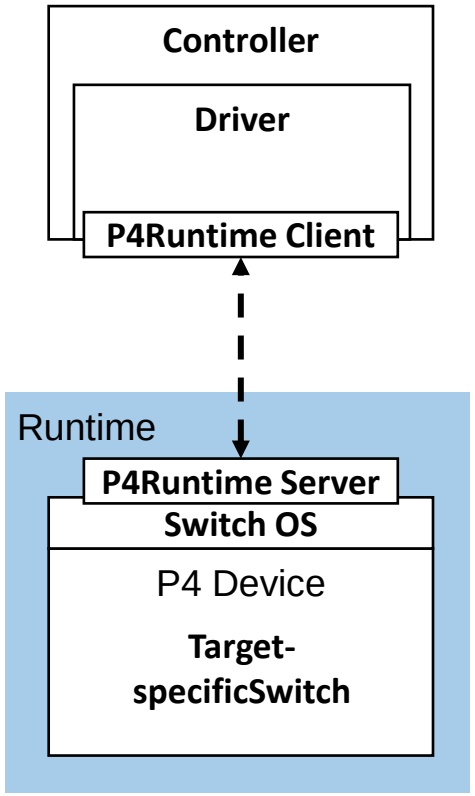

P4 Workflow

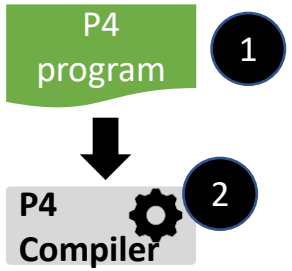

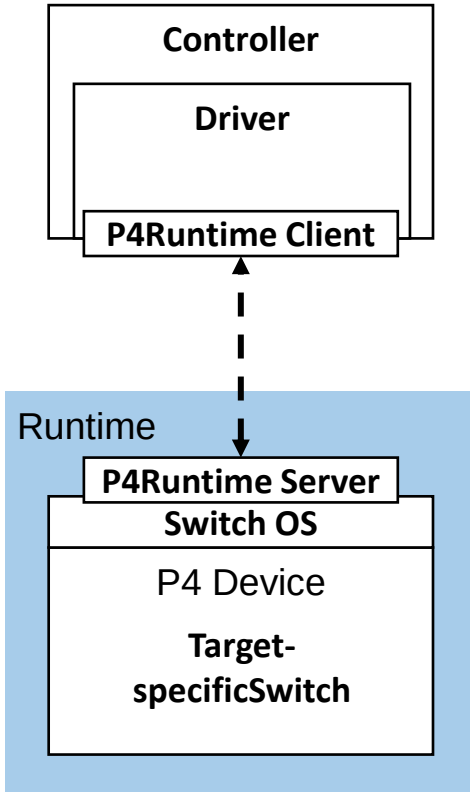

P4 Workflow

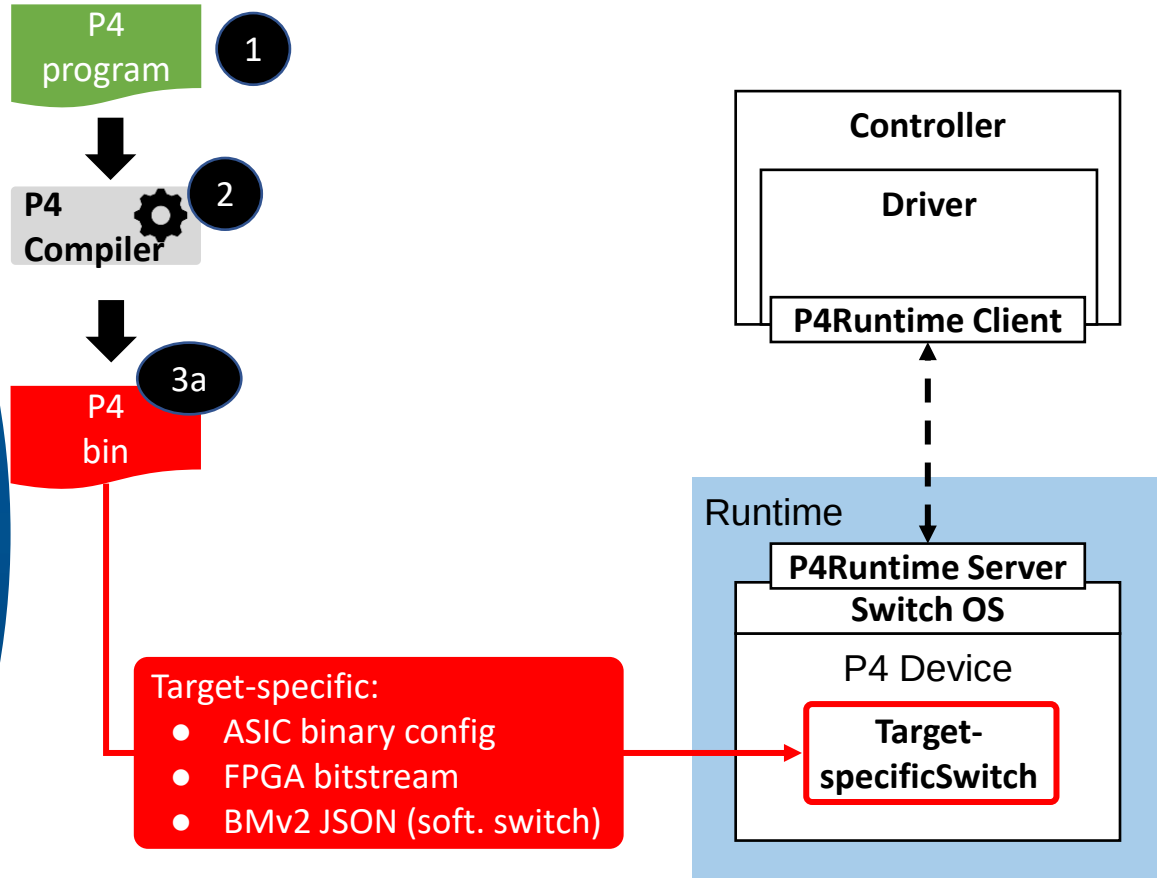

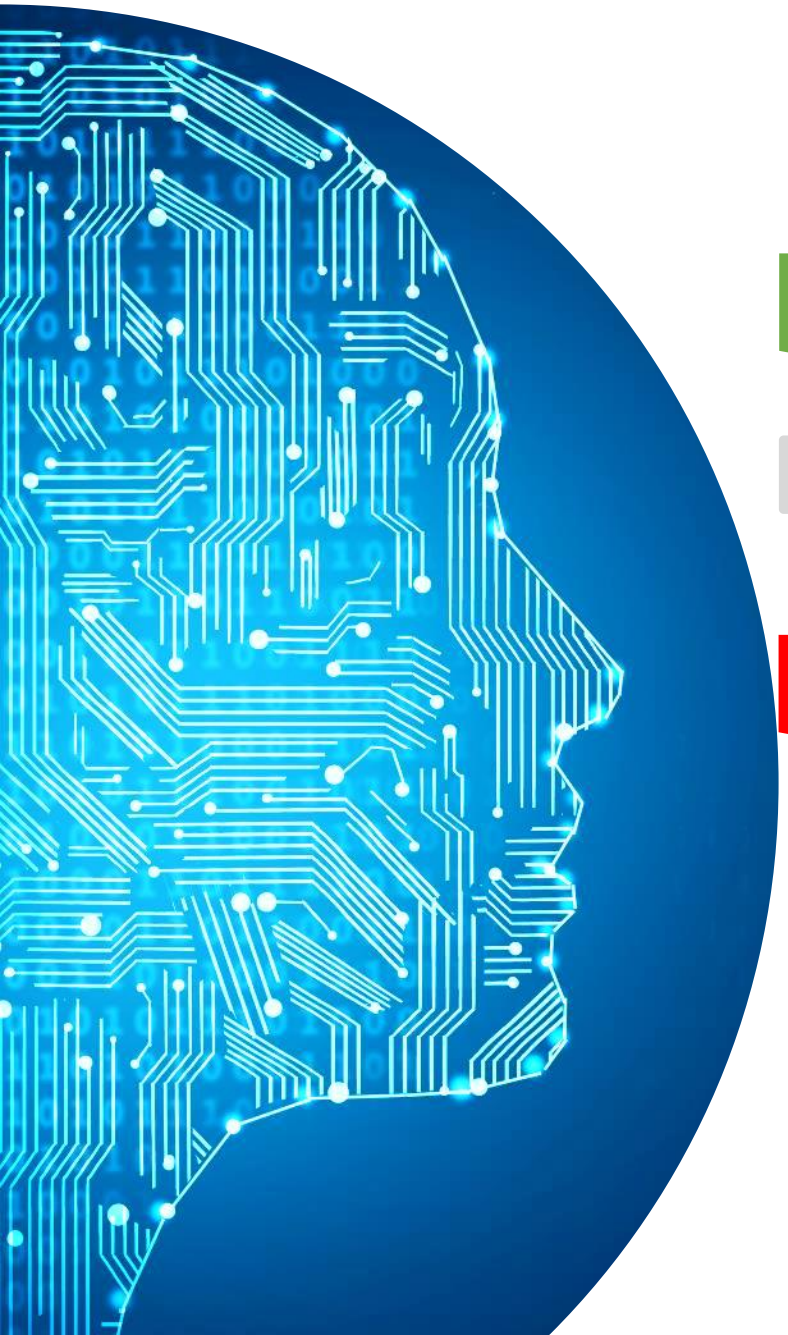

P4 Workflow

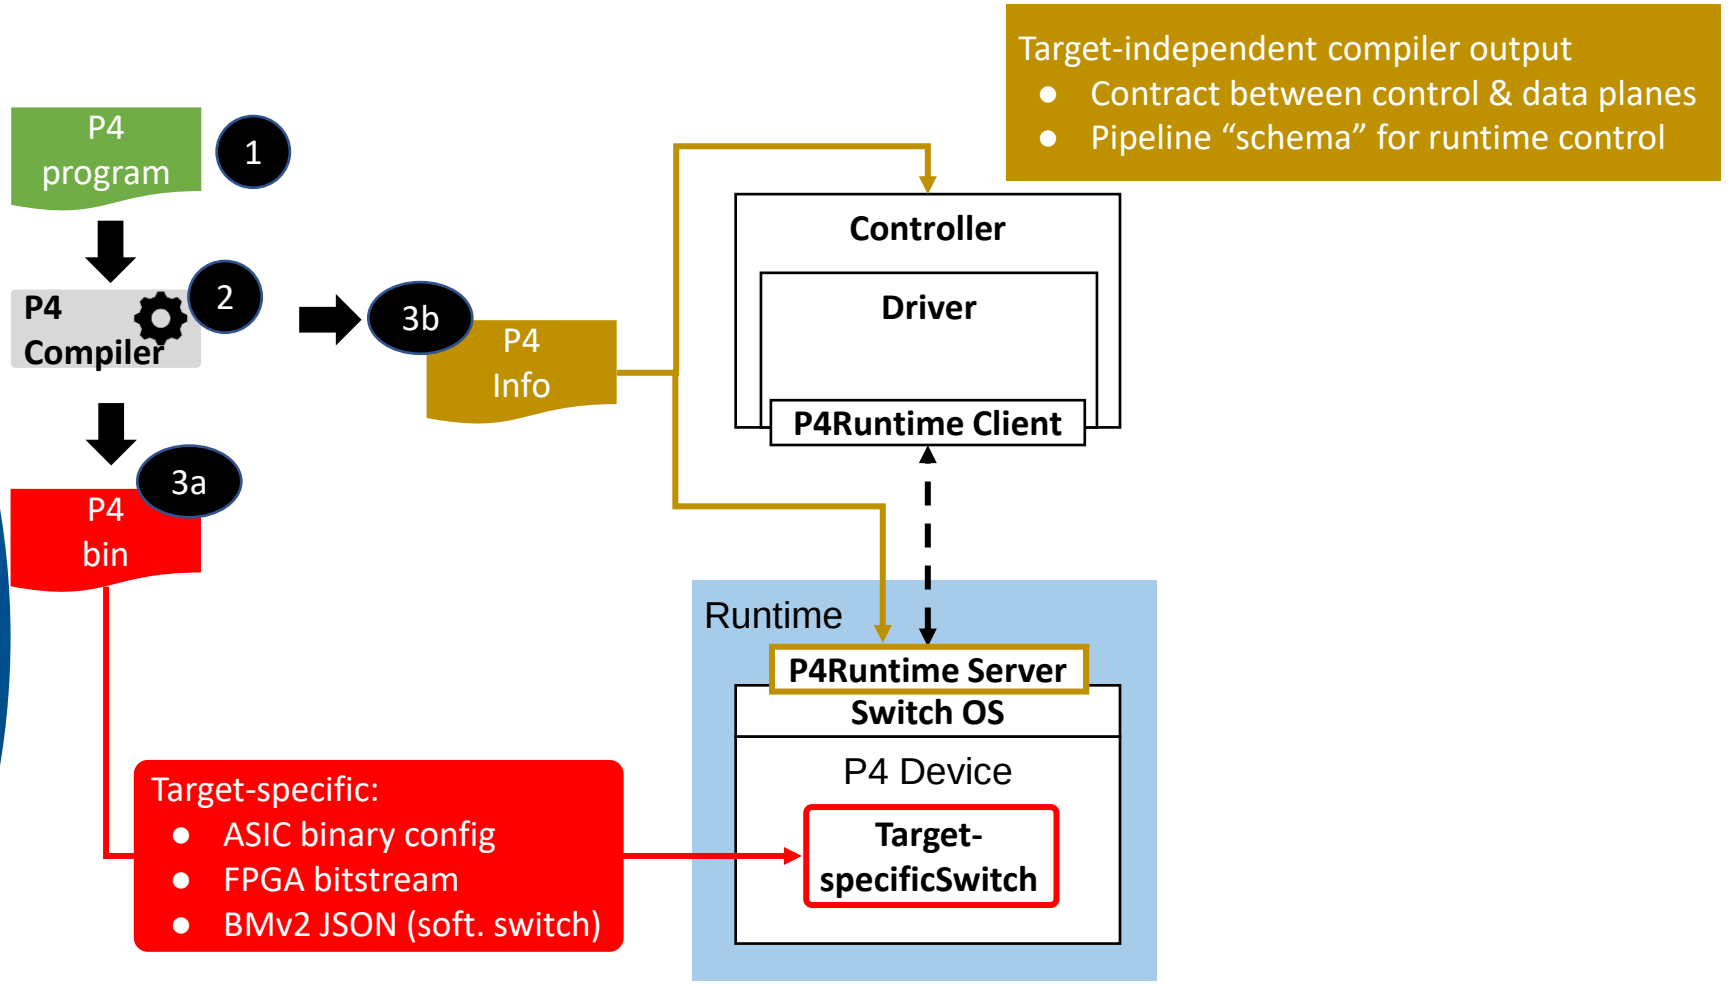

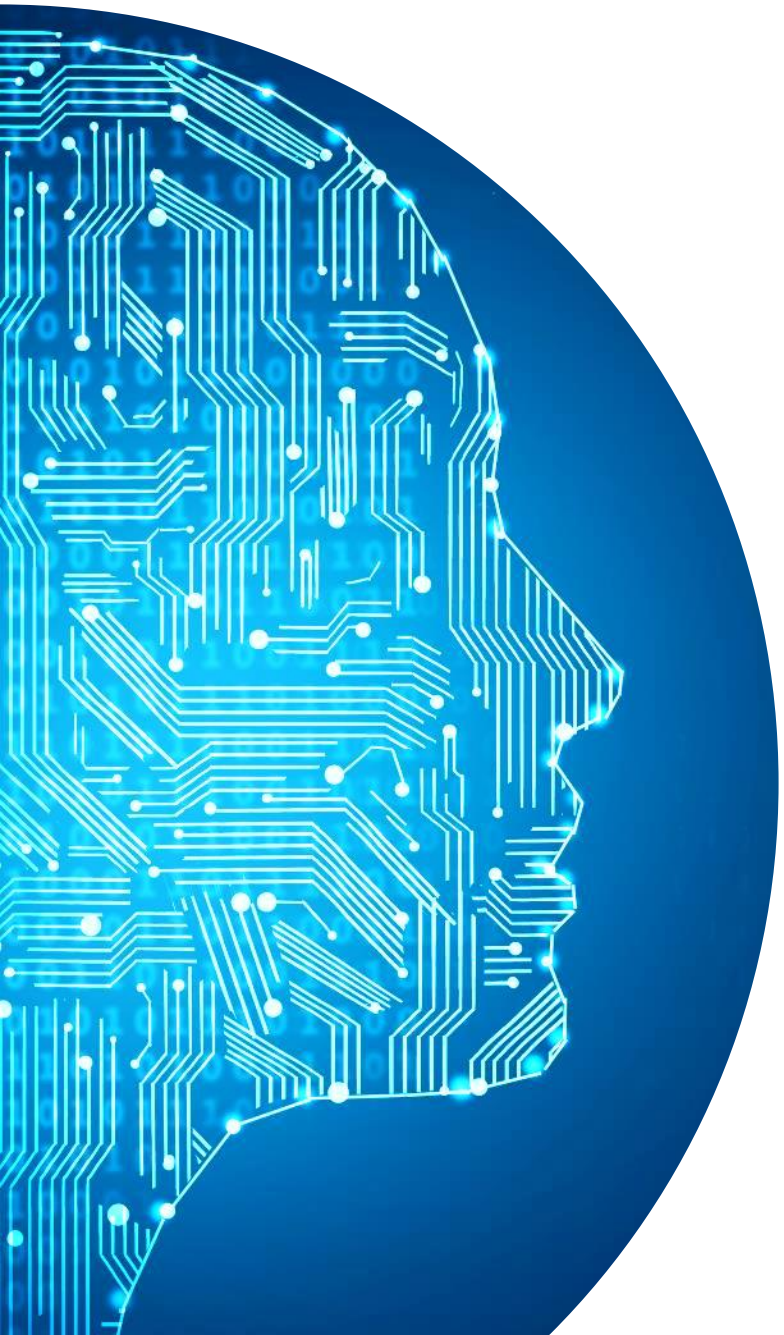

References

ONF's P4 Language Tutorial:
https://opennetworking.org/wp-content/uploads/2020/12/P4_D2_East_2018_01_basics.pdf

ONF's Next generation SDN tutorial:
https://github.com/opennetworkinglab/ngsdn-tutorial

P4 presentation:
https://olivermichel.github.io/doc/p4.pdf

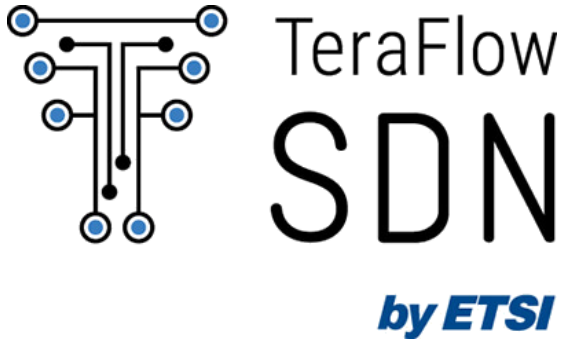

Thank you!

TFSsupport@etsi.org