

# Flightplan: Dataplane Disaggregation and Placement for P4 programs


Nik Sultana   John Sonchack   Hans Giesen   Isaac Pedisich   Zhaoyang Han  
Nishanth Shyamkumar   Shivani Burad   André DeHon   Boon Thau Loo

*University of Pennsylvania*


<https://flightplan.cis.upenn.edu/>

NSDI'21

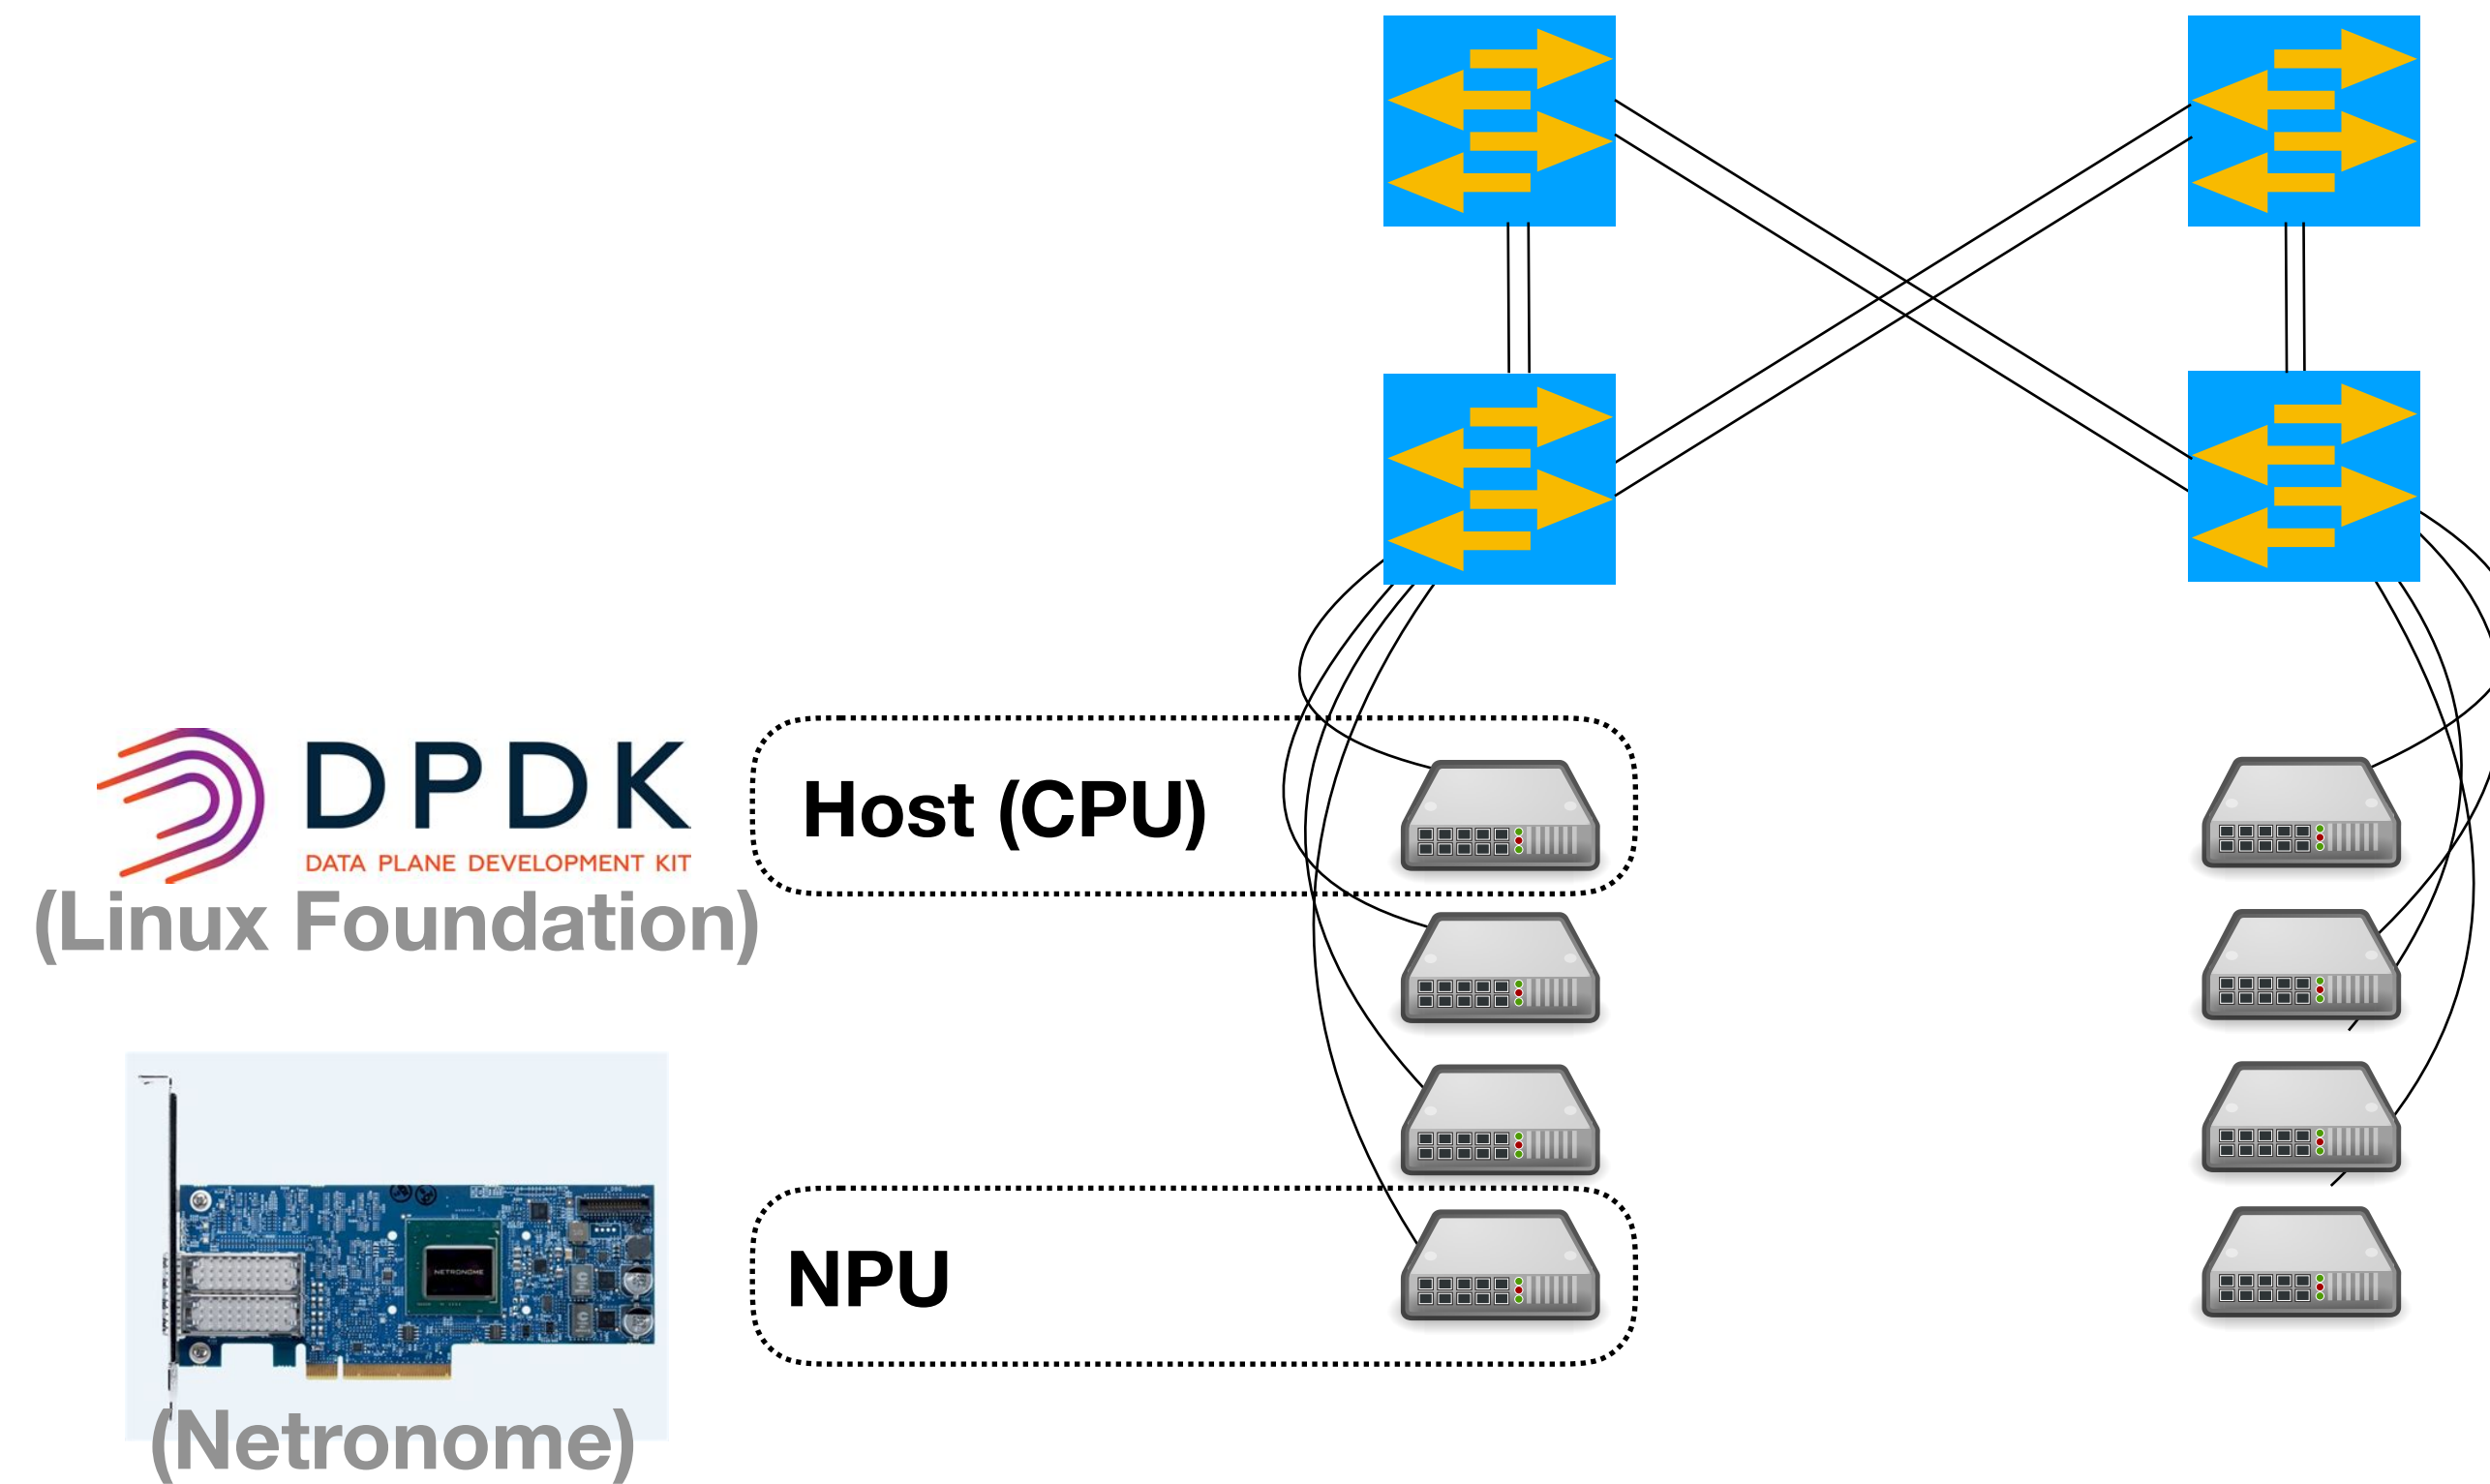
# Dataplane Programmability




# Dataplane Programmability




# Dataplane Programmability




# Dataplane Programmability



# Dataplane Programmability

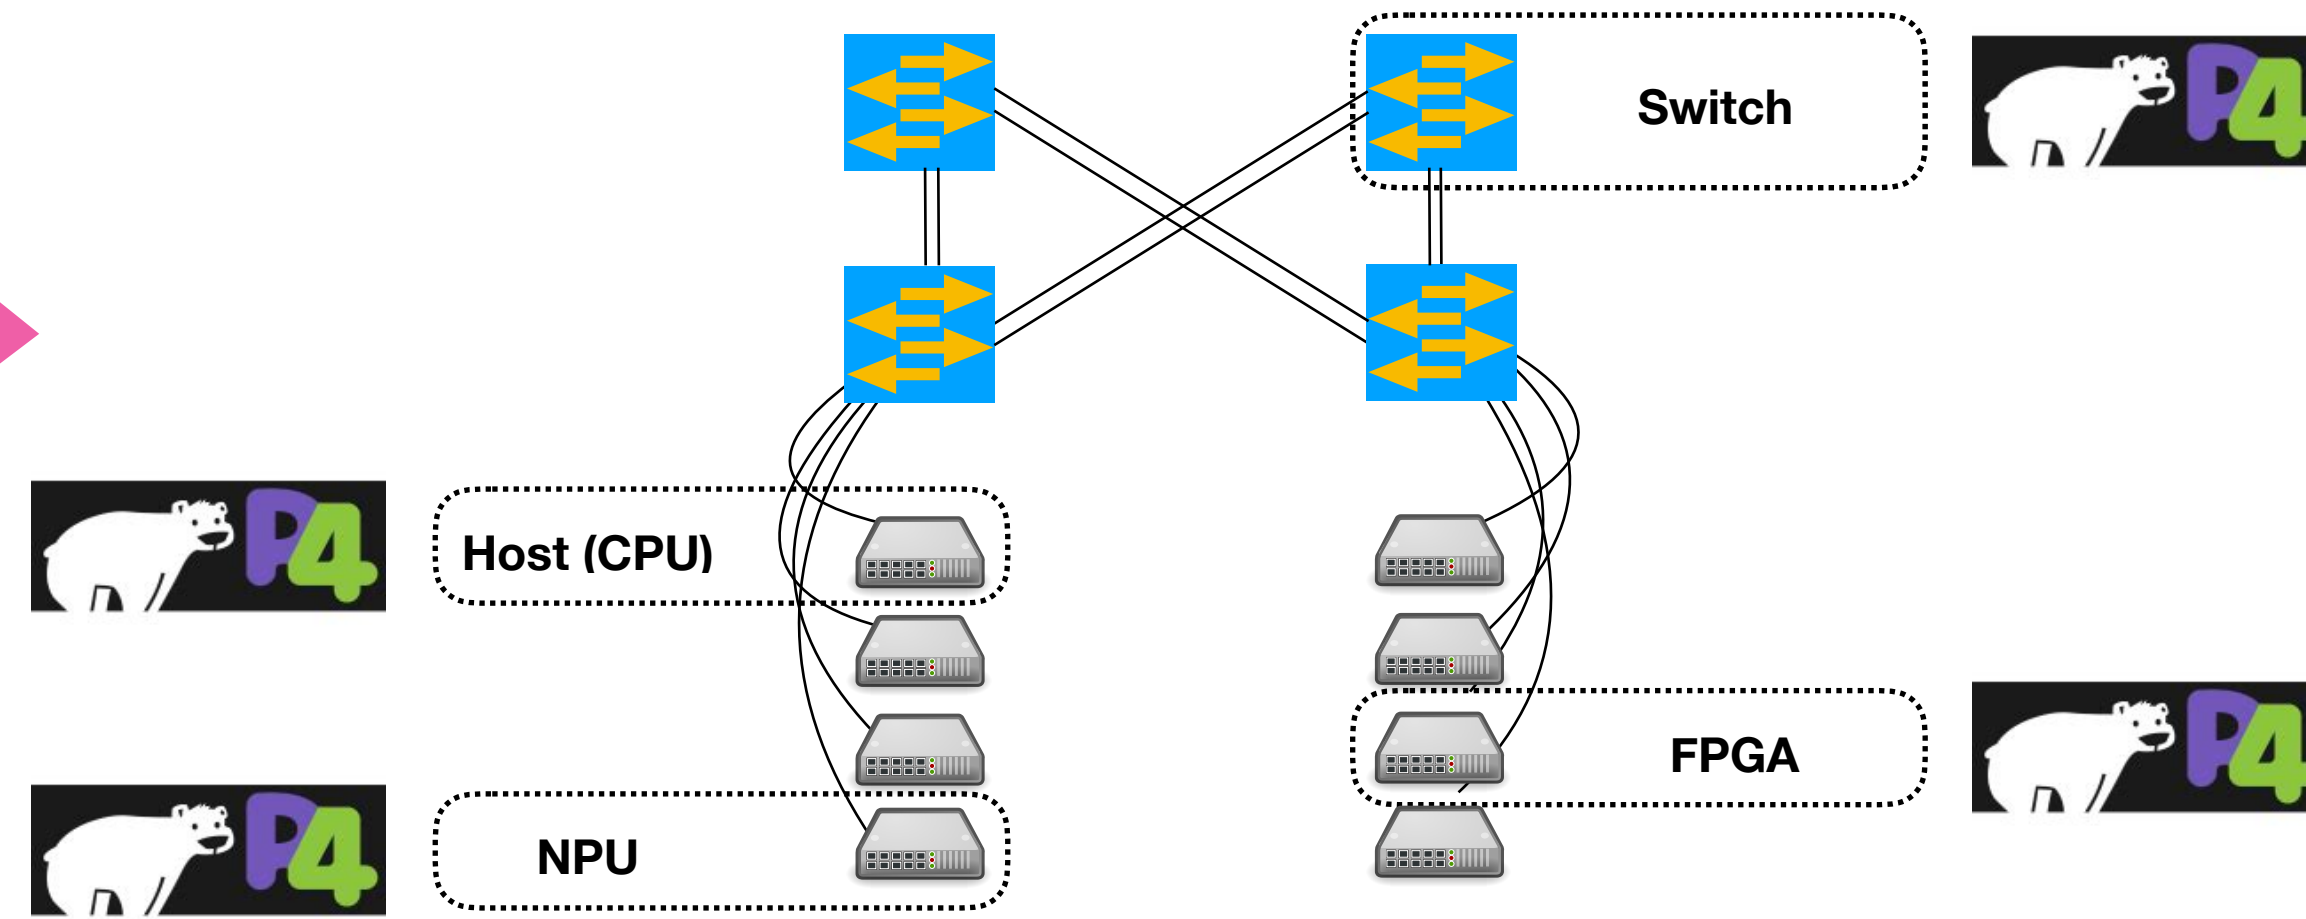


# Dataplane Programmability



# Dataplane Programmability

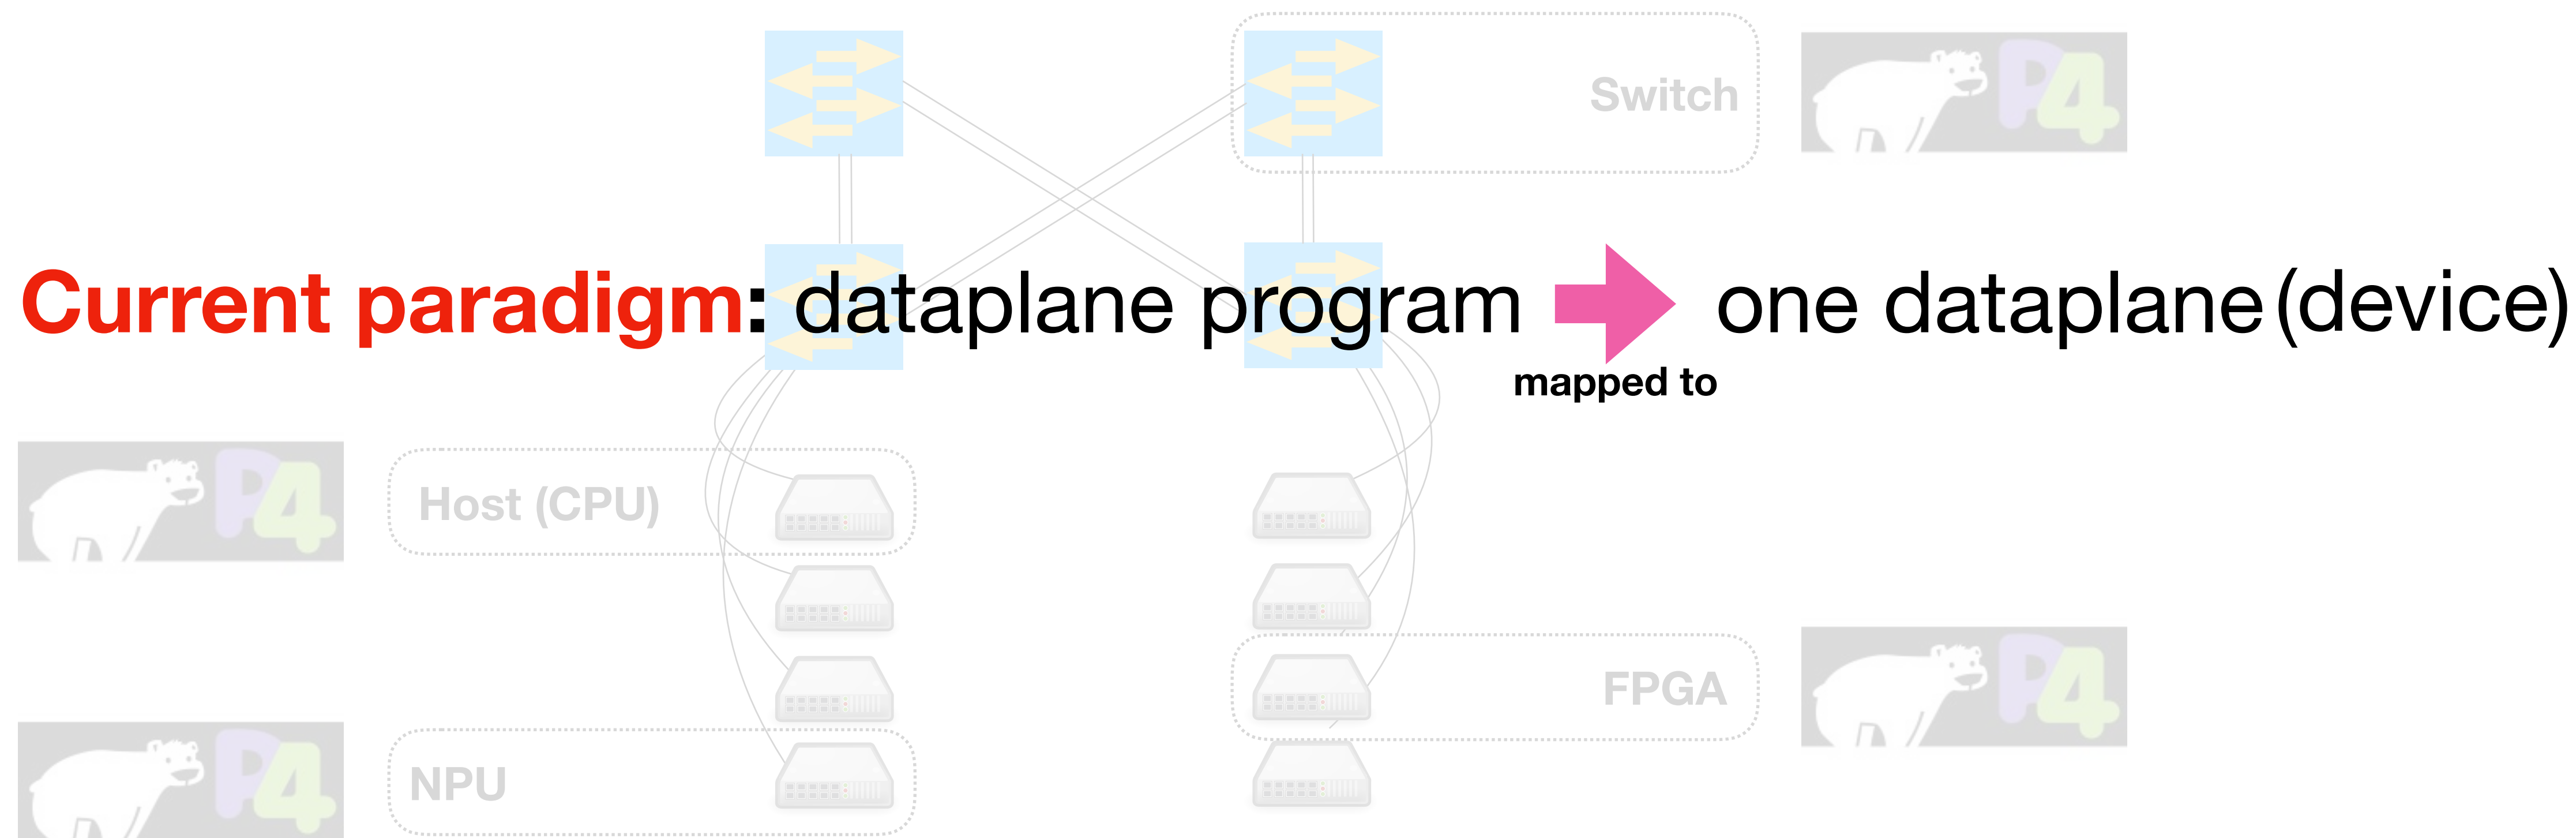
Research



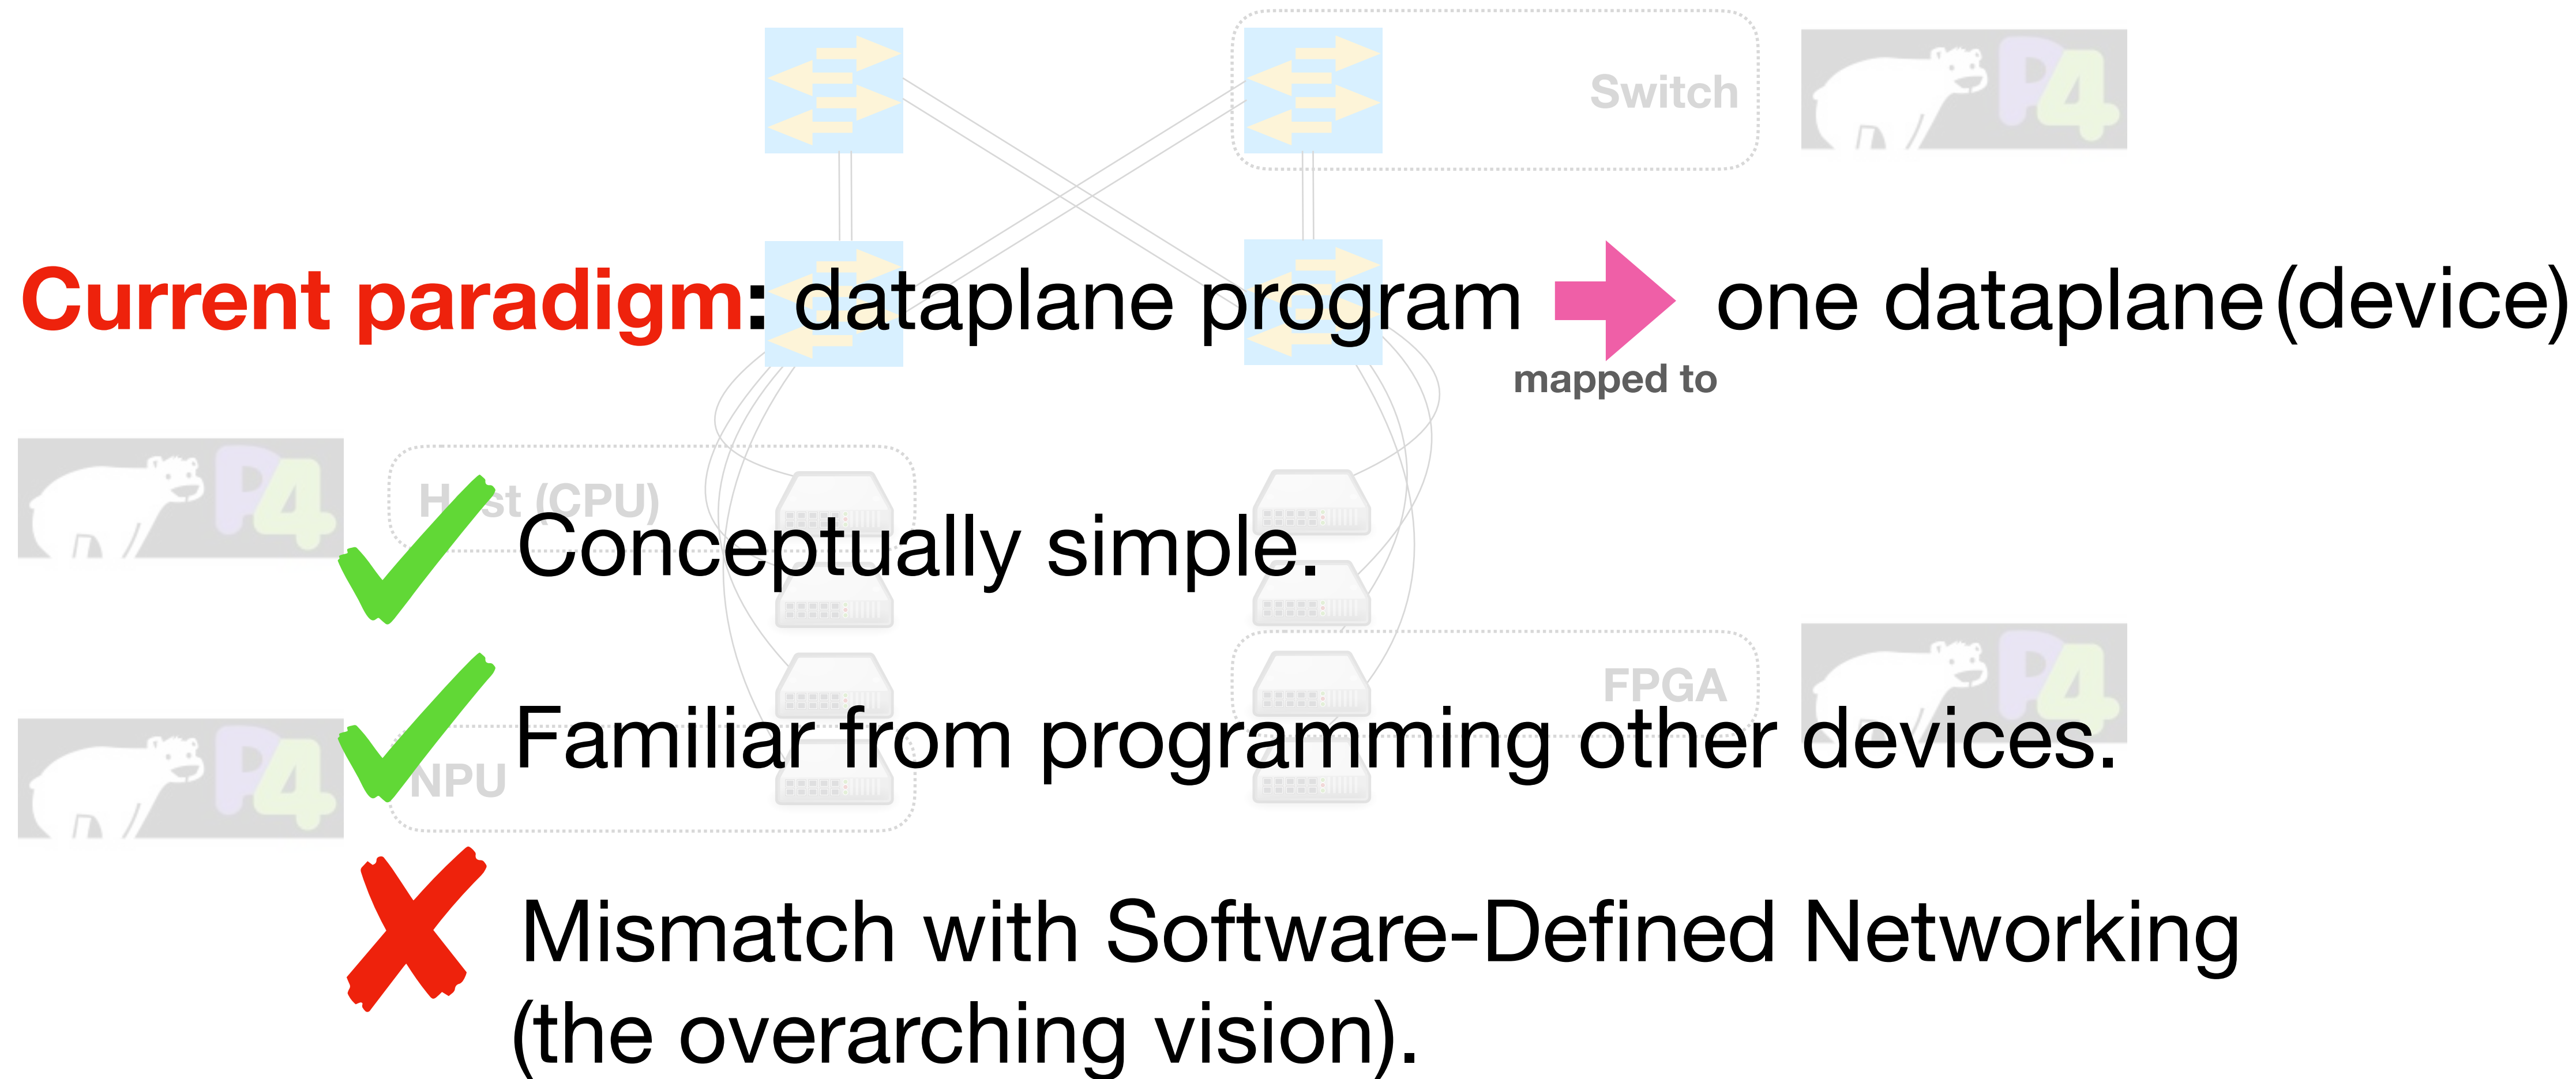
Products




# Dataplane Programmability



# Dataplane Programmability



**Current paradigm:** dataplane program  one dataplane  
mapped to




## Mismatch with Software-Defined Networking

- **Target's resources are dedicated to the program.**  
Inefficient use of individual target resources.
- **Program must entirely execute on a single target.**  
Unnecessary constraints of program functionality.
- **We're meant to be "programming the network".**  
Programmability is scoped too conservatively.

# Current paradigm: dataplane program one dataplane

mapped to




xor



xor

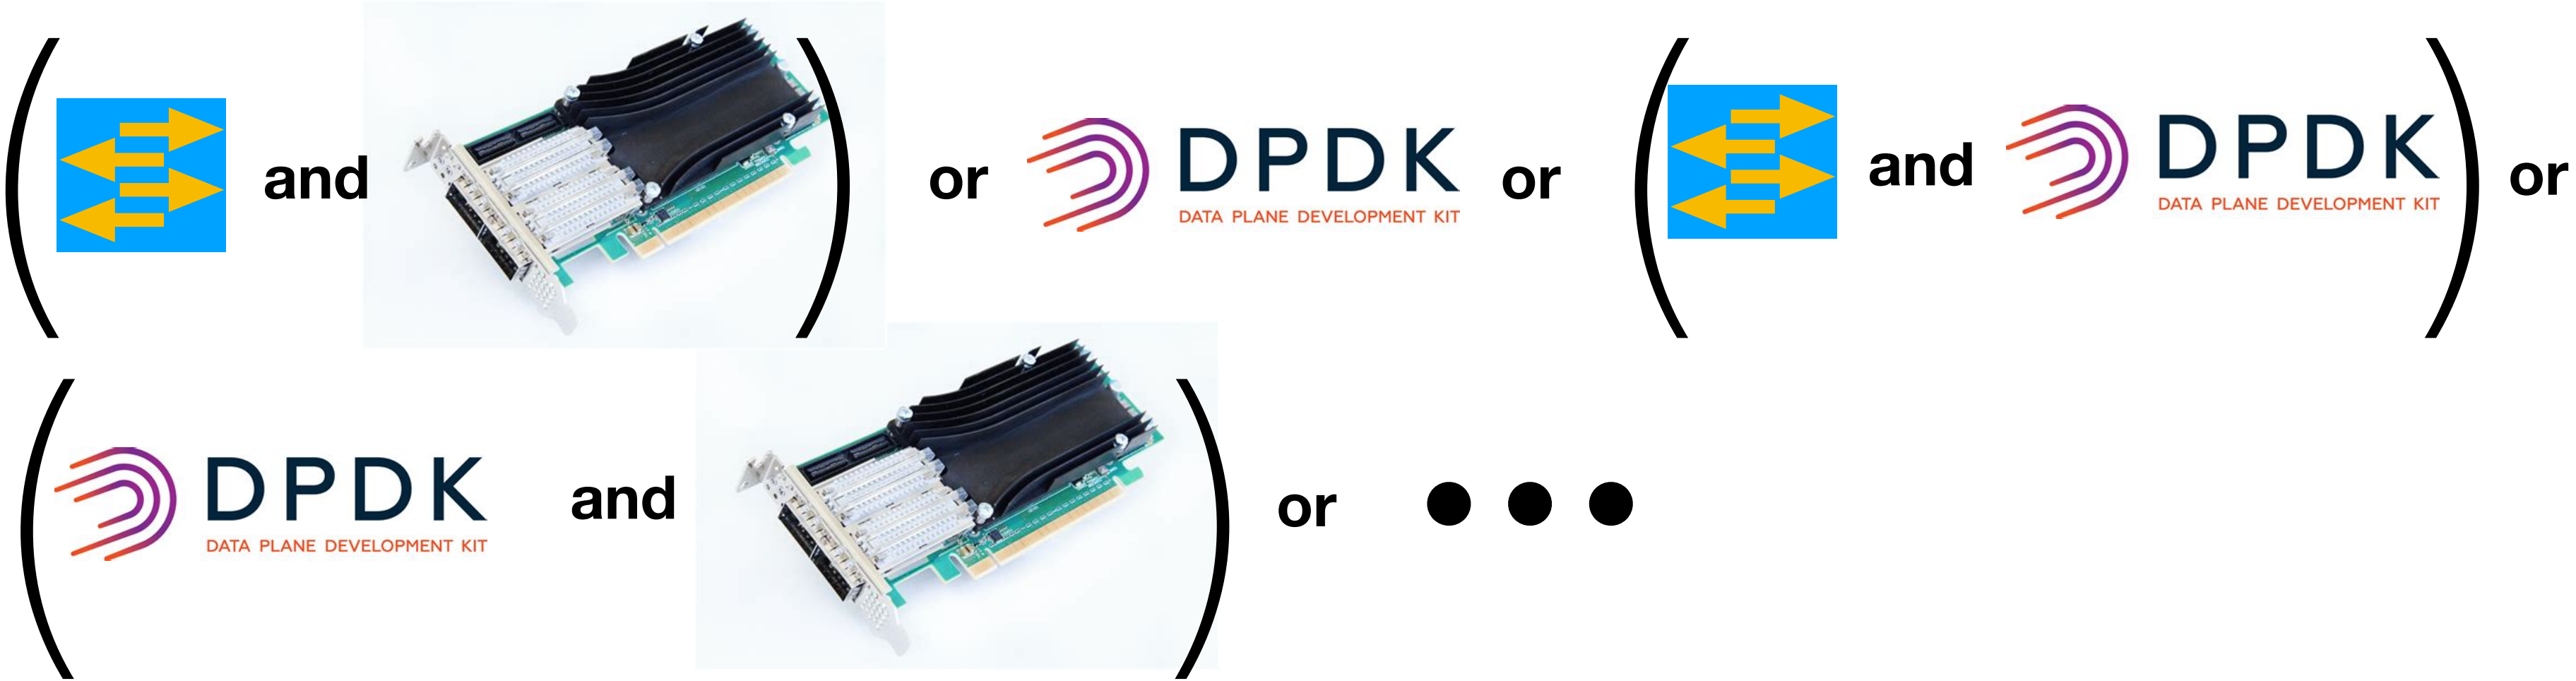


xor




## New paradigm:

dataplane program  suitable mix of dataplanes




**Current paradigm:** dataplane program → one dataplane





xor



xor




xor




**New paradigm:**

dataplane program → suitable mix of dataplanes

- Resource-based program decomposition.
- Split into set of 1-1 dataplane programs.




and



or



or



and



or




and



or



**Current paradigm:** dataplane program → one dataplane




**New paradigm:**  
dataplane program → suitable mix of dataplanes

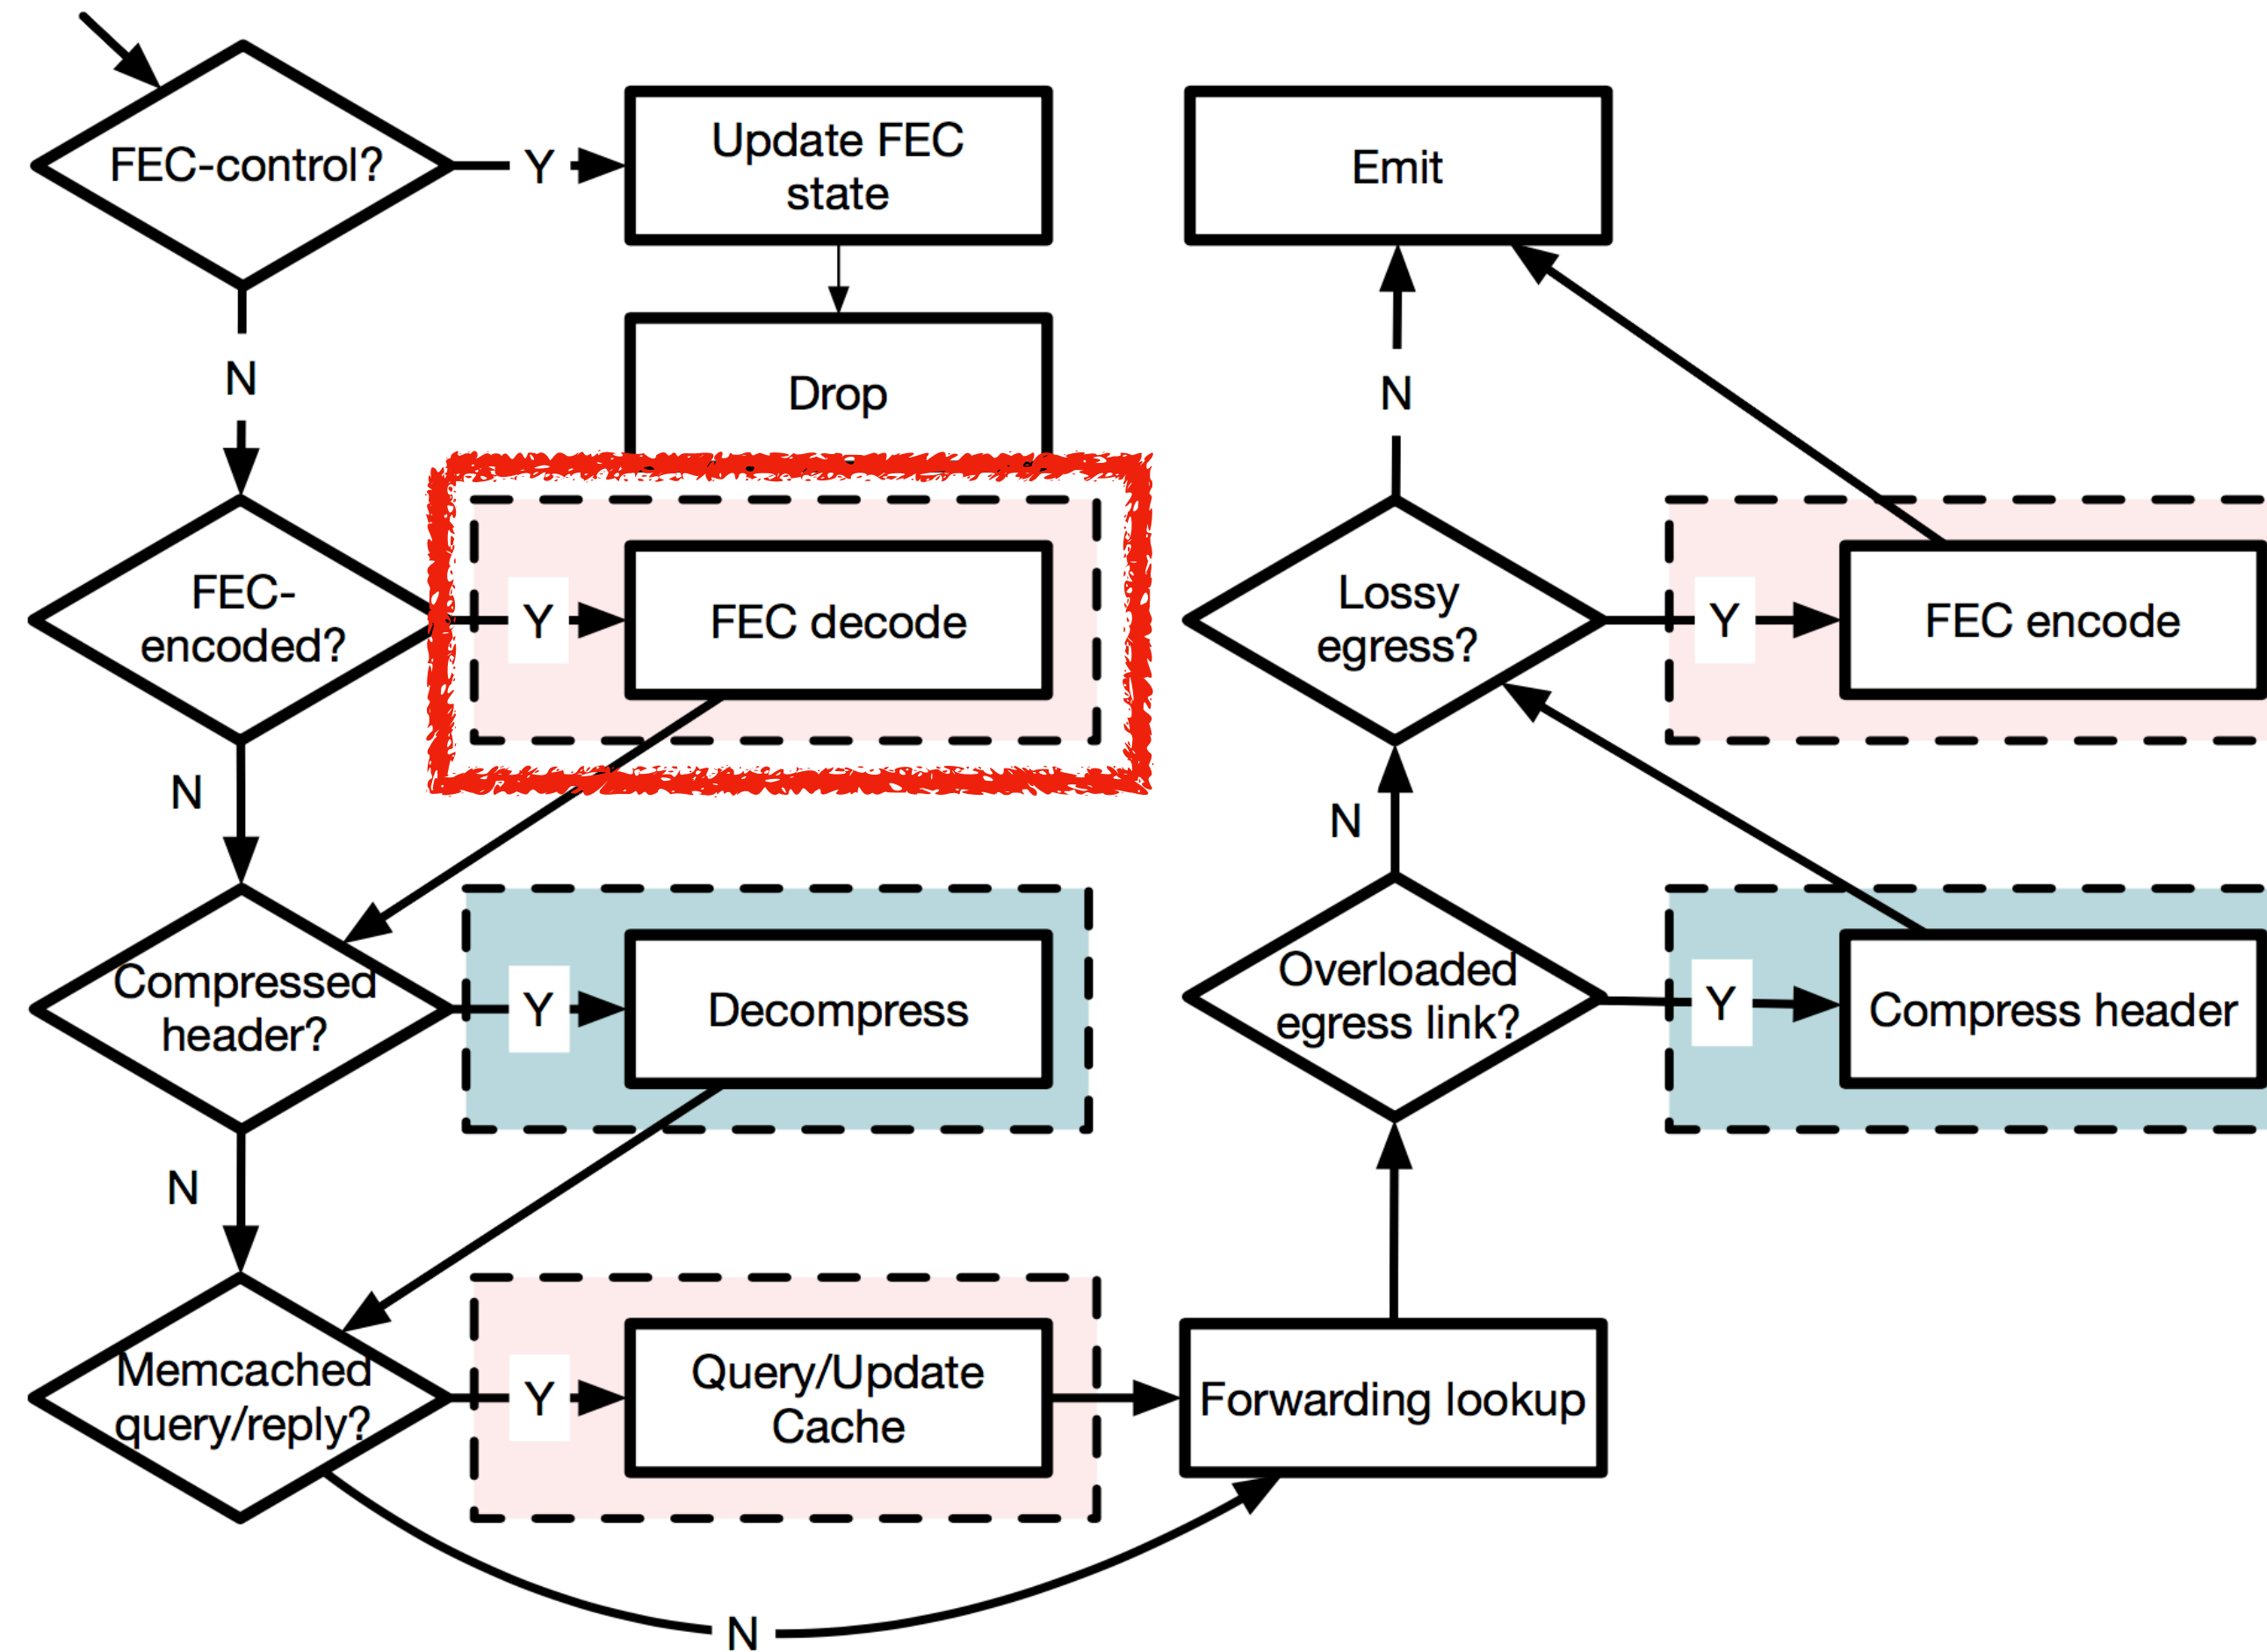
- Resource-based program decomposition.
- Split into set of 1-1 dataplane programs.

✓ Uses existing vendor toolchains, language & hardware.

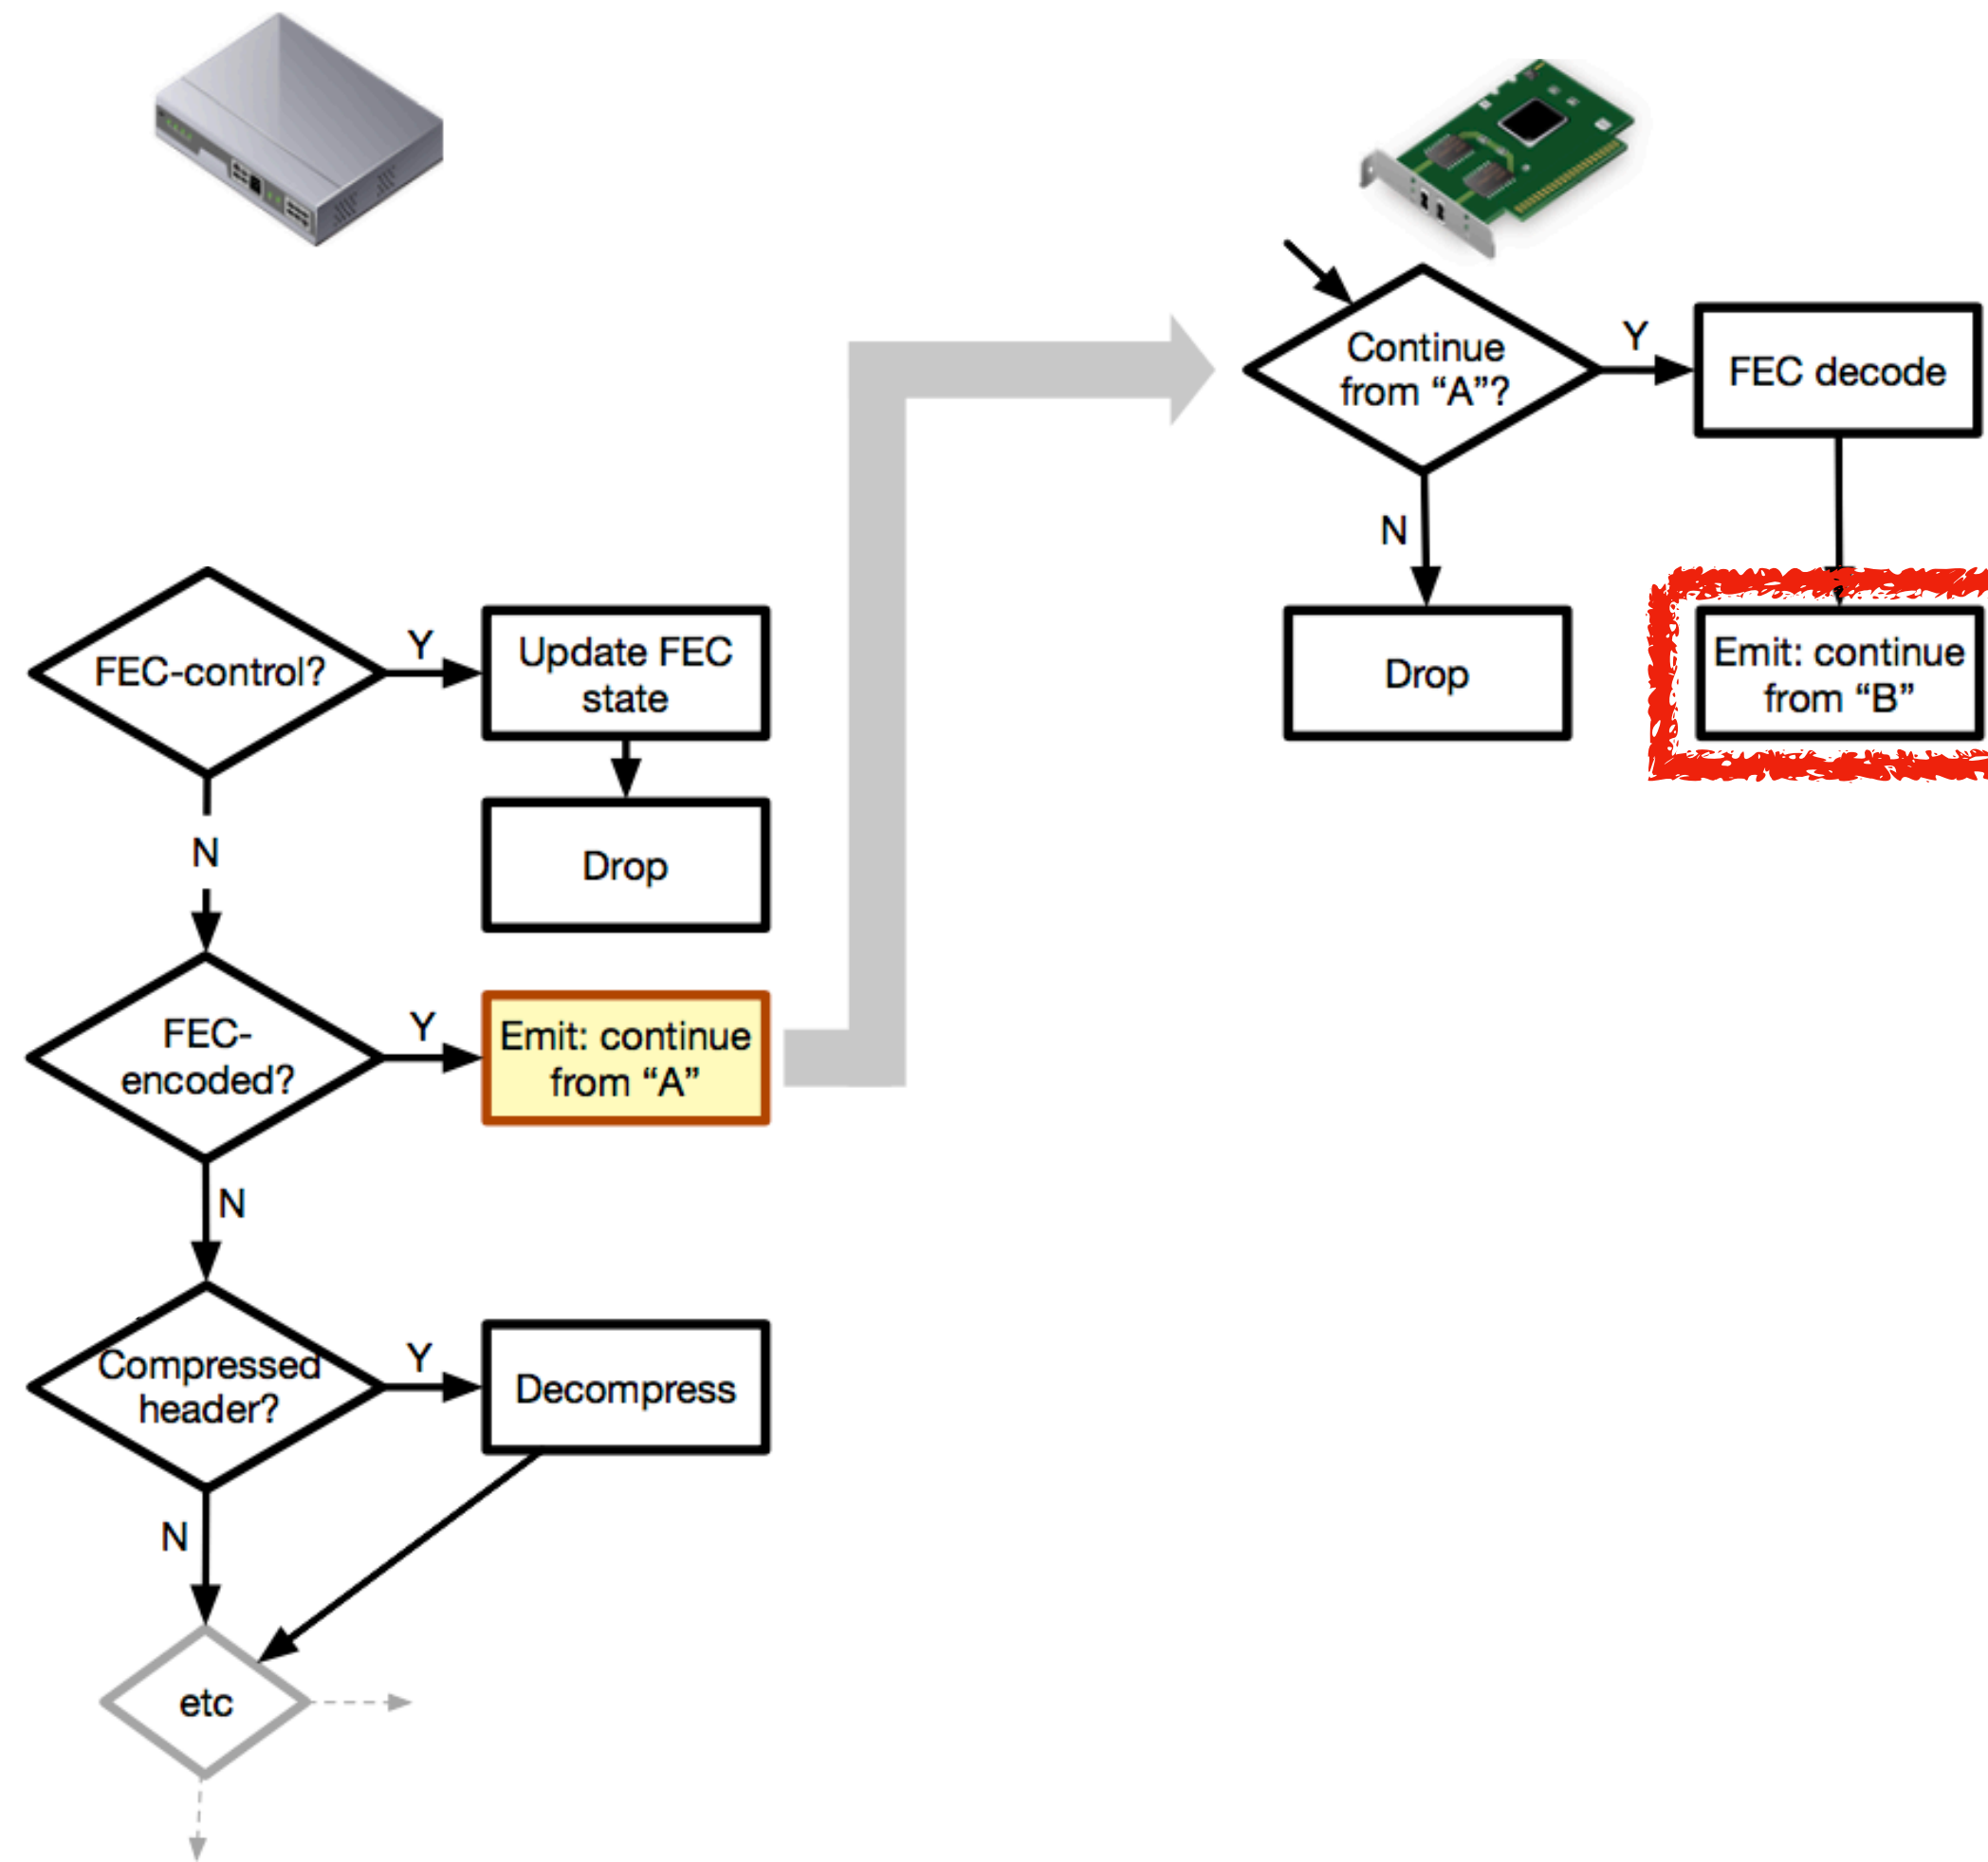
# Example: “Crosspod” Program




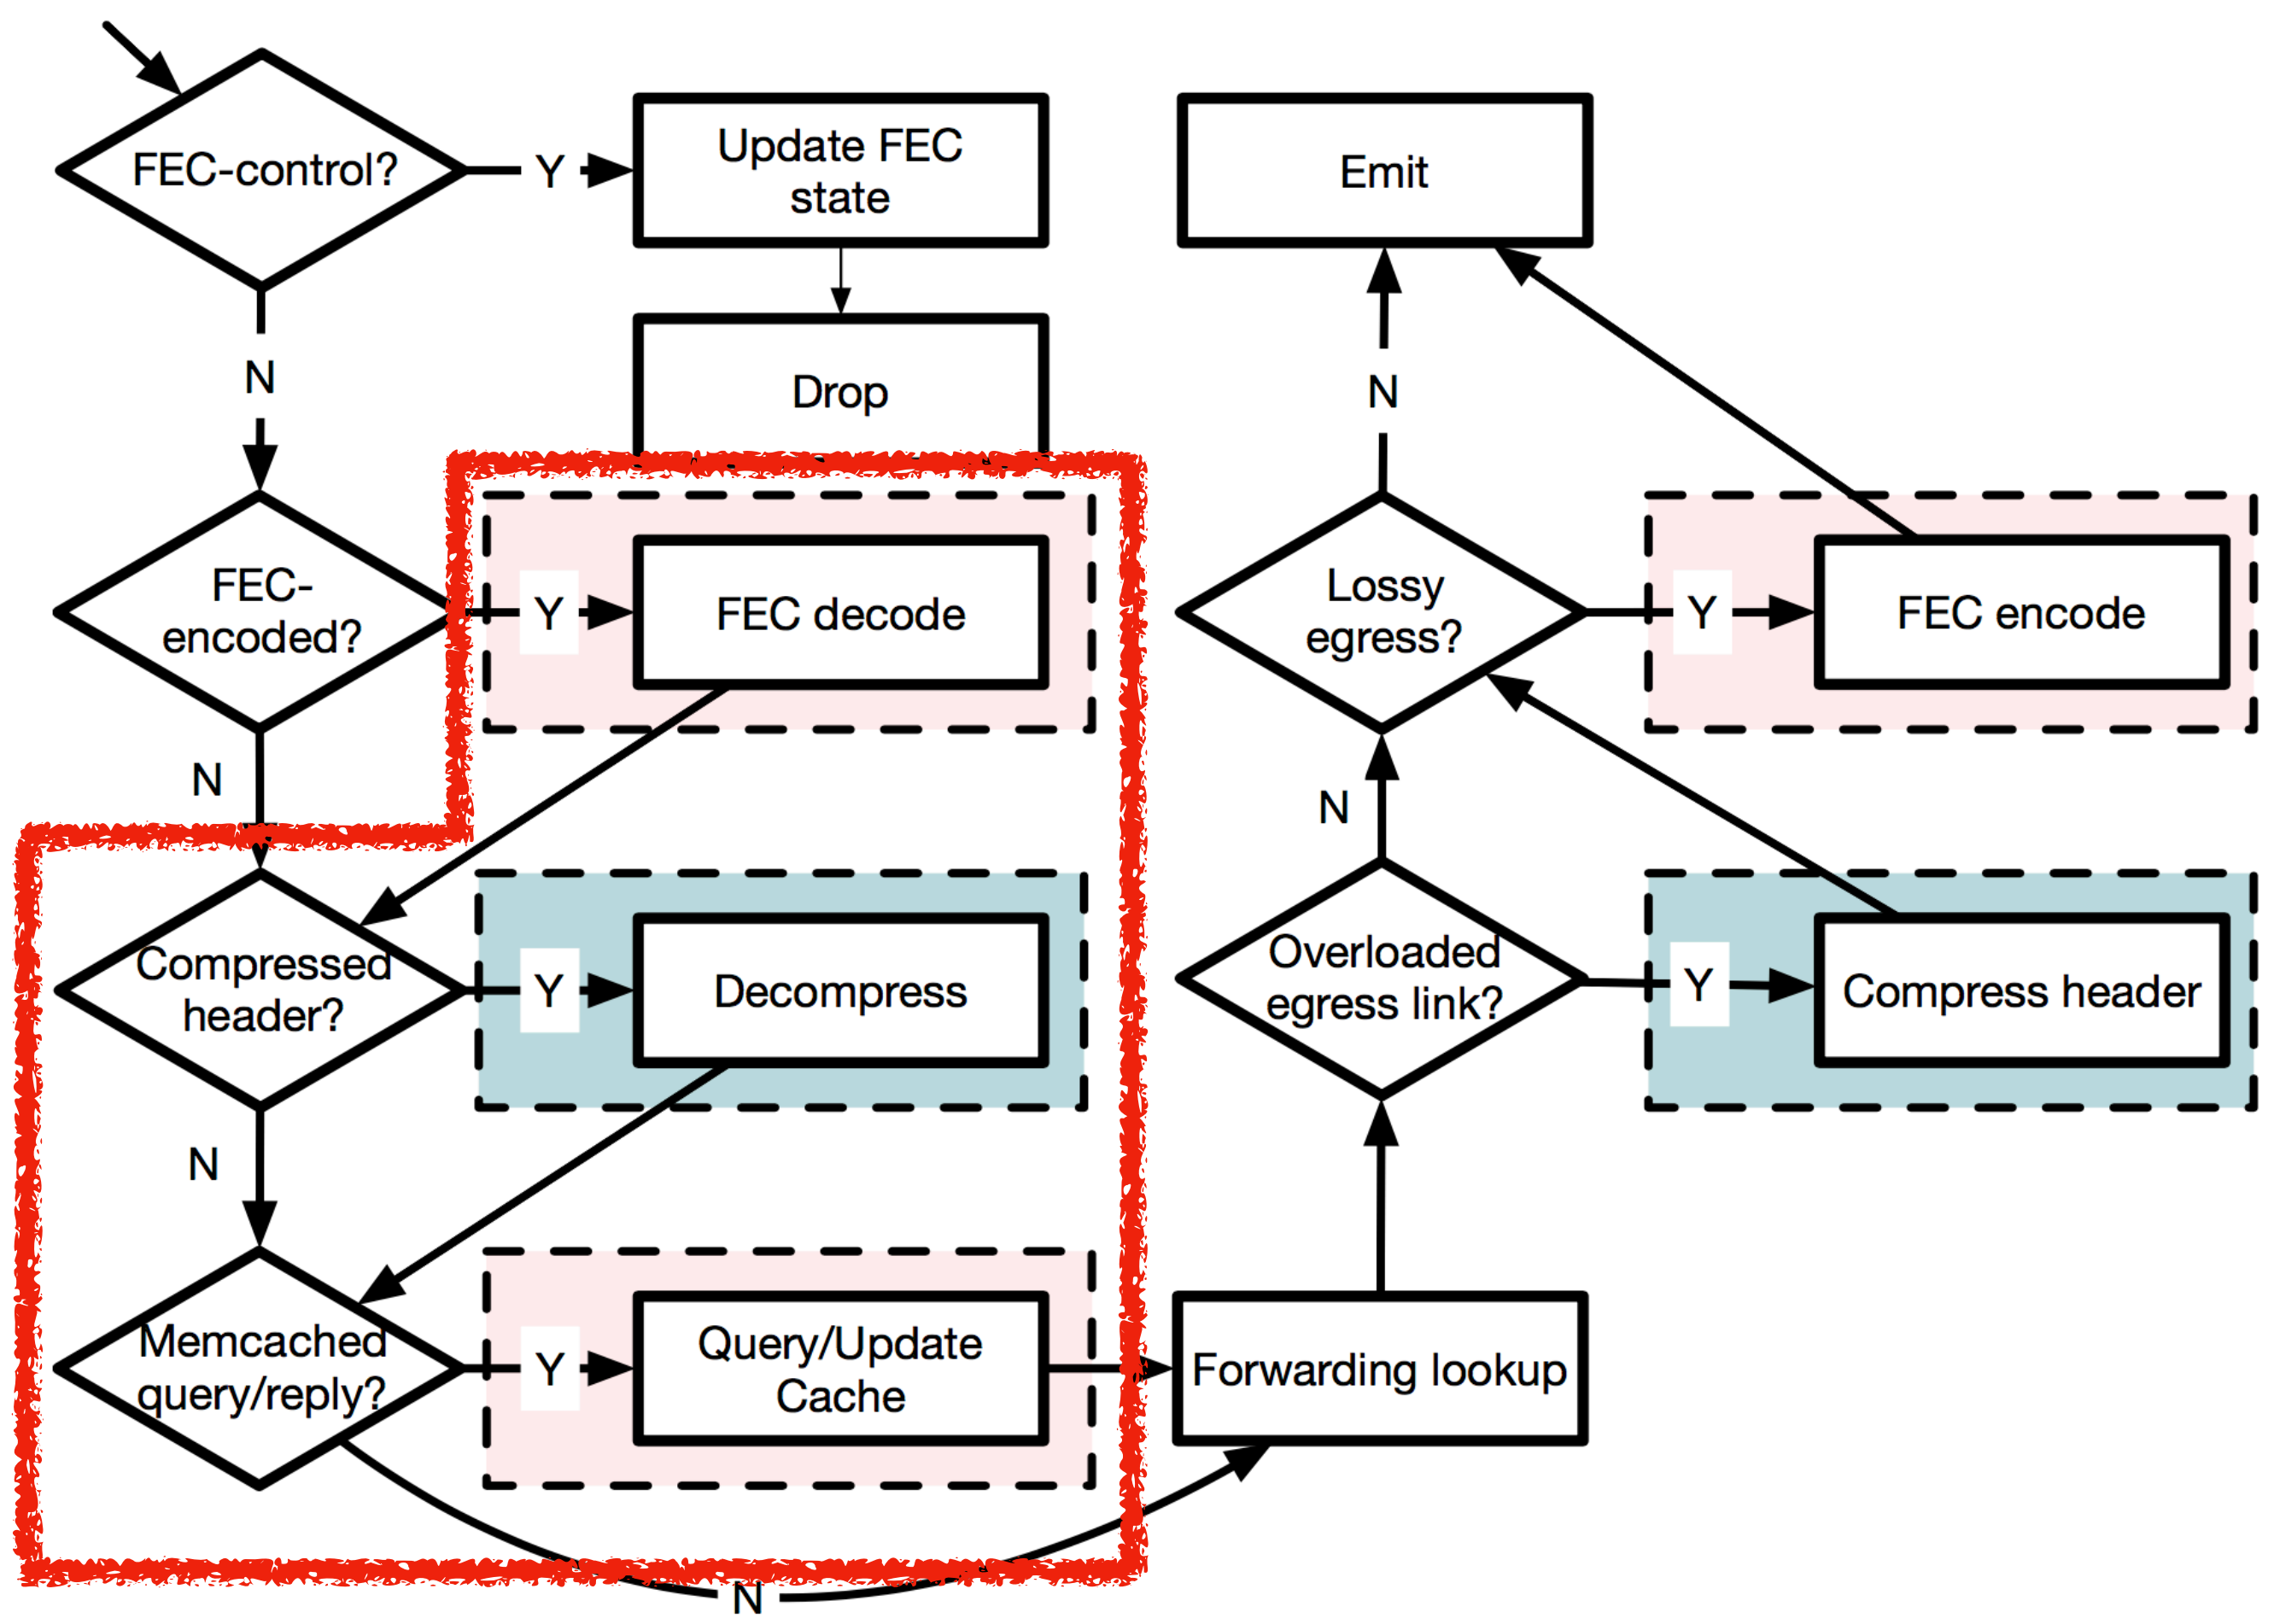
# Example: "Crosspod" Program













**Current paradigm:** dataplane program → one dataplane



**New paradigm: Dataplane Disaggregation**  
dataplane program → suitable mix of dataplanes

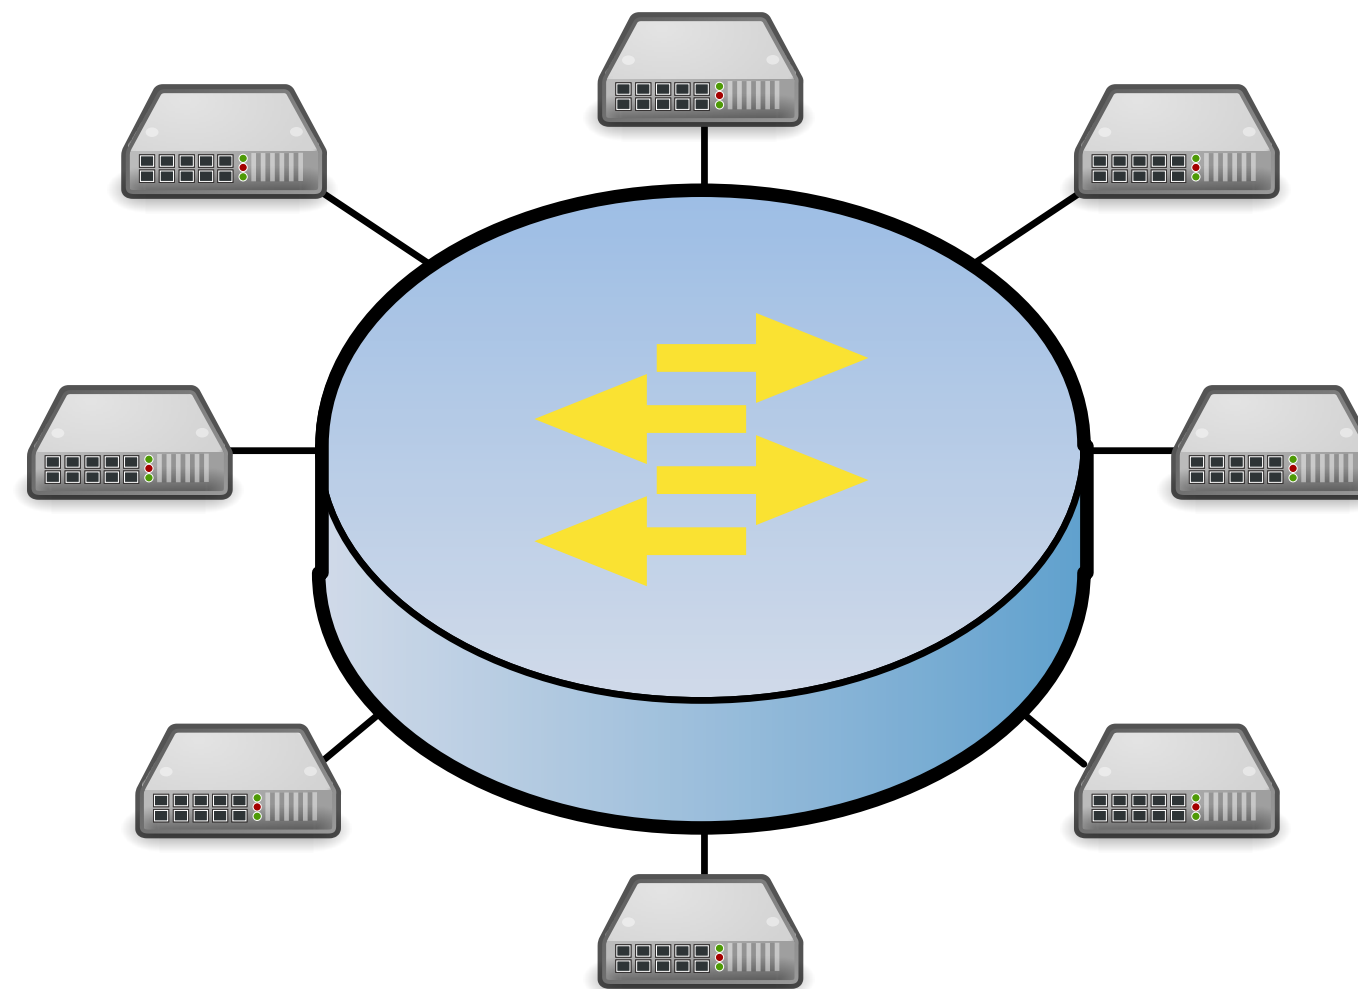
- Resource-based program decomposition.
- Split into set of 1-1 dataplane programs.



# Dataplane Disaggregation


( ≠ Server Disaggregation)

Virtual Dataplane → Set of Physical Dataplanes



“One big switch” → “One big programmable switch”

Current paradigm: dataplane program → one dataplane



xor



xor

**How to implement?**

**Automated support needed**

**Dataplane Disaggregation:**

dataplane program → suitable mix of dataplanes

- Resource-based program decomposition.
- Split into set of 1-1 dataplane programs.




Uses existing vendor toolchains, language & hardware.



If manual: laborious, error-prone, hard to change.


Current paradigm: dataplane program → one dataplane




xor



xor



xor




### Dataplane Disaggregation:

dataplane program → suitable mix of dataplanes



- Analyze program's use of resources.
- Split into set of 1-1 dataplane programs.




and



or



or



and



or



and



or




Current paradigm: dataplane program → one dataplane



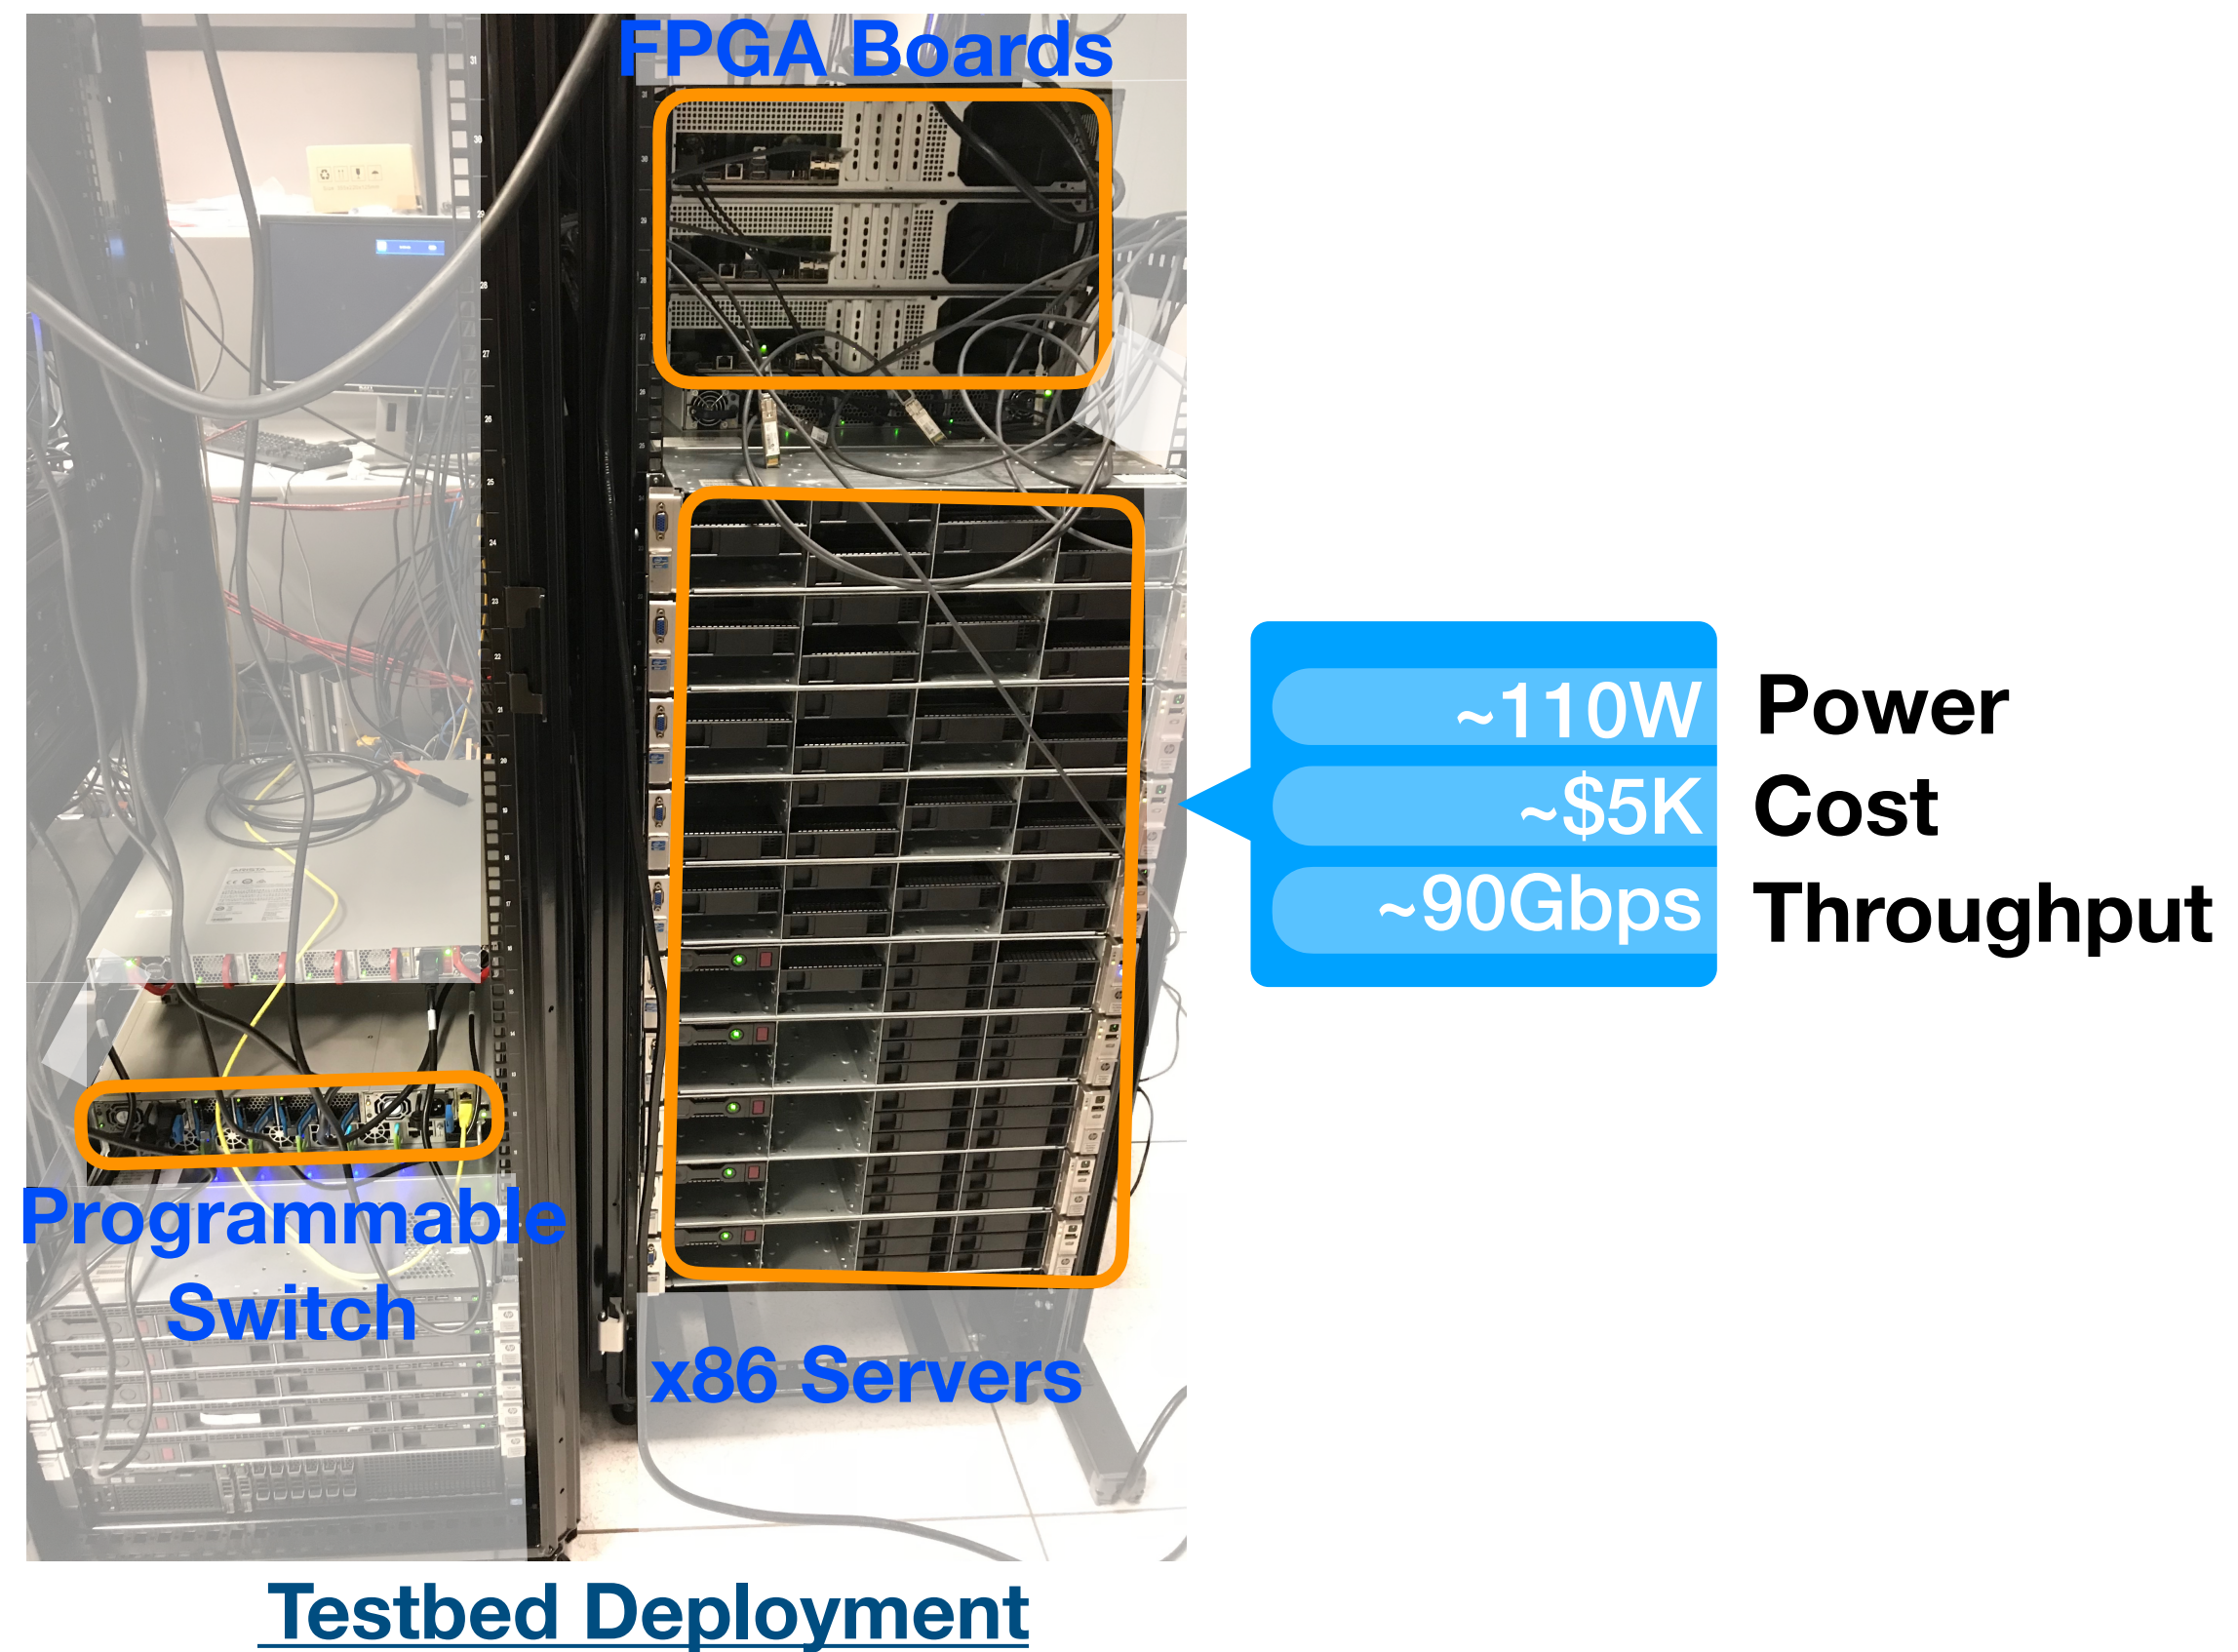


# Heterogeneity

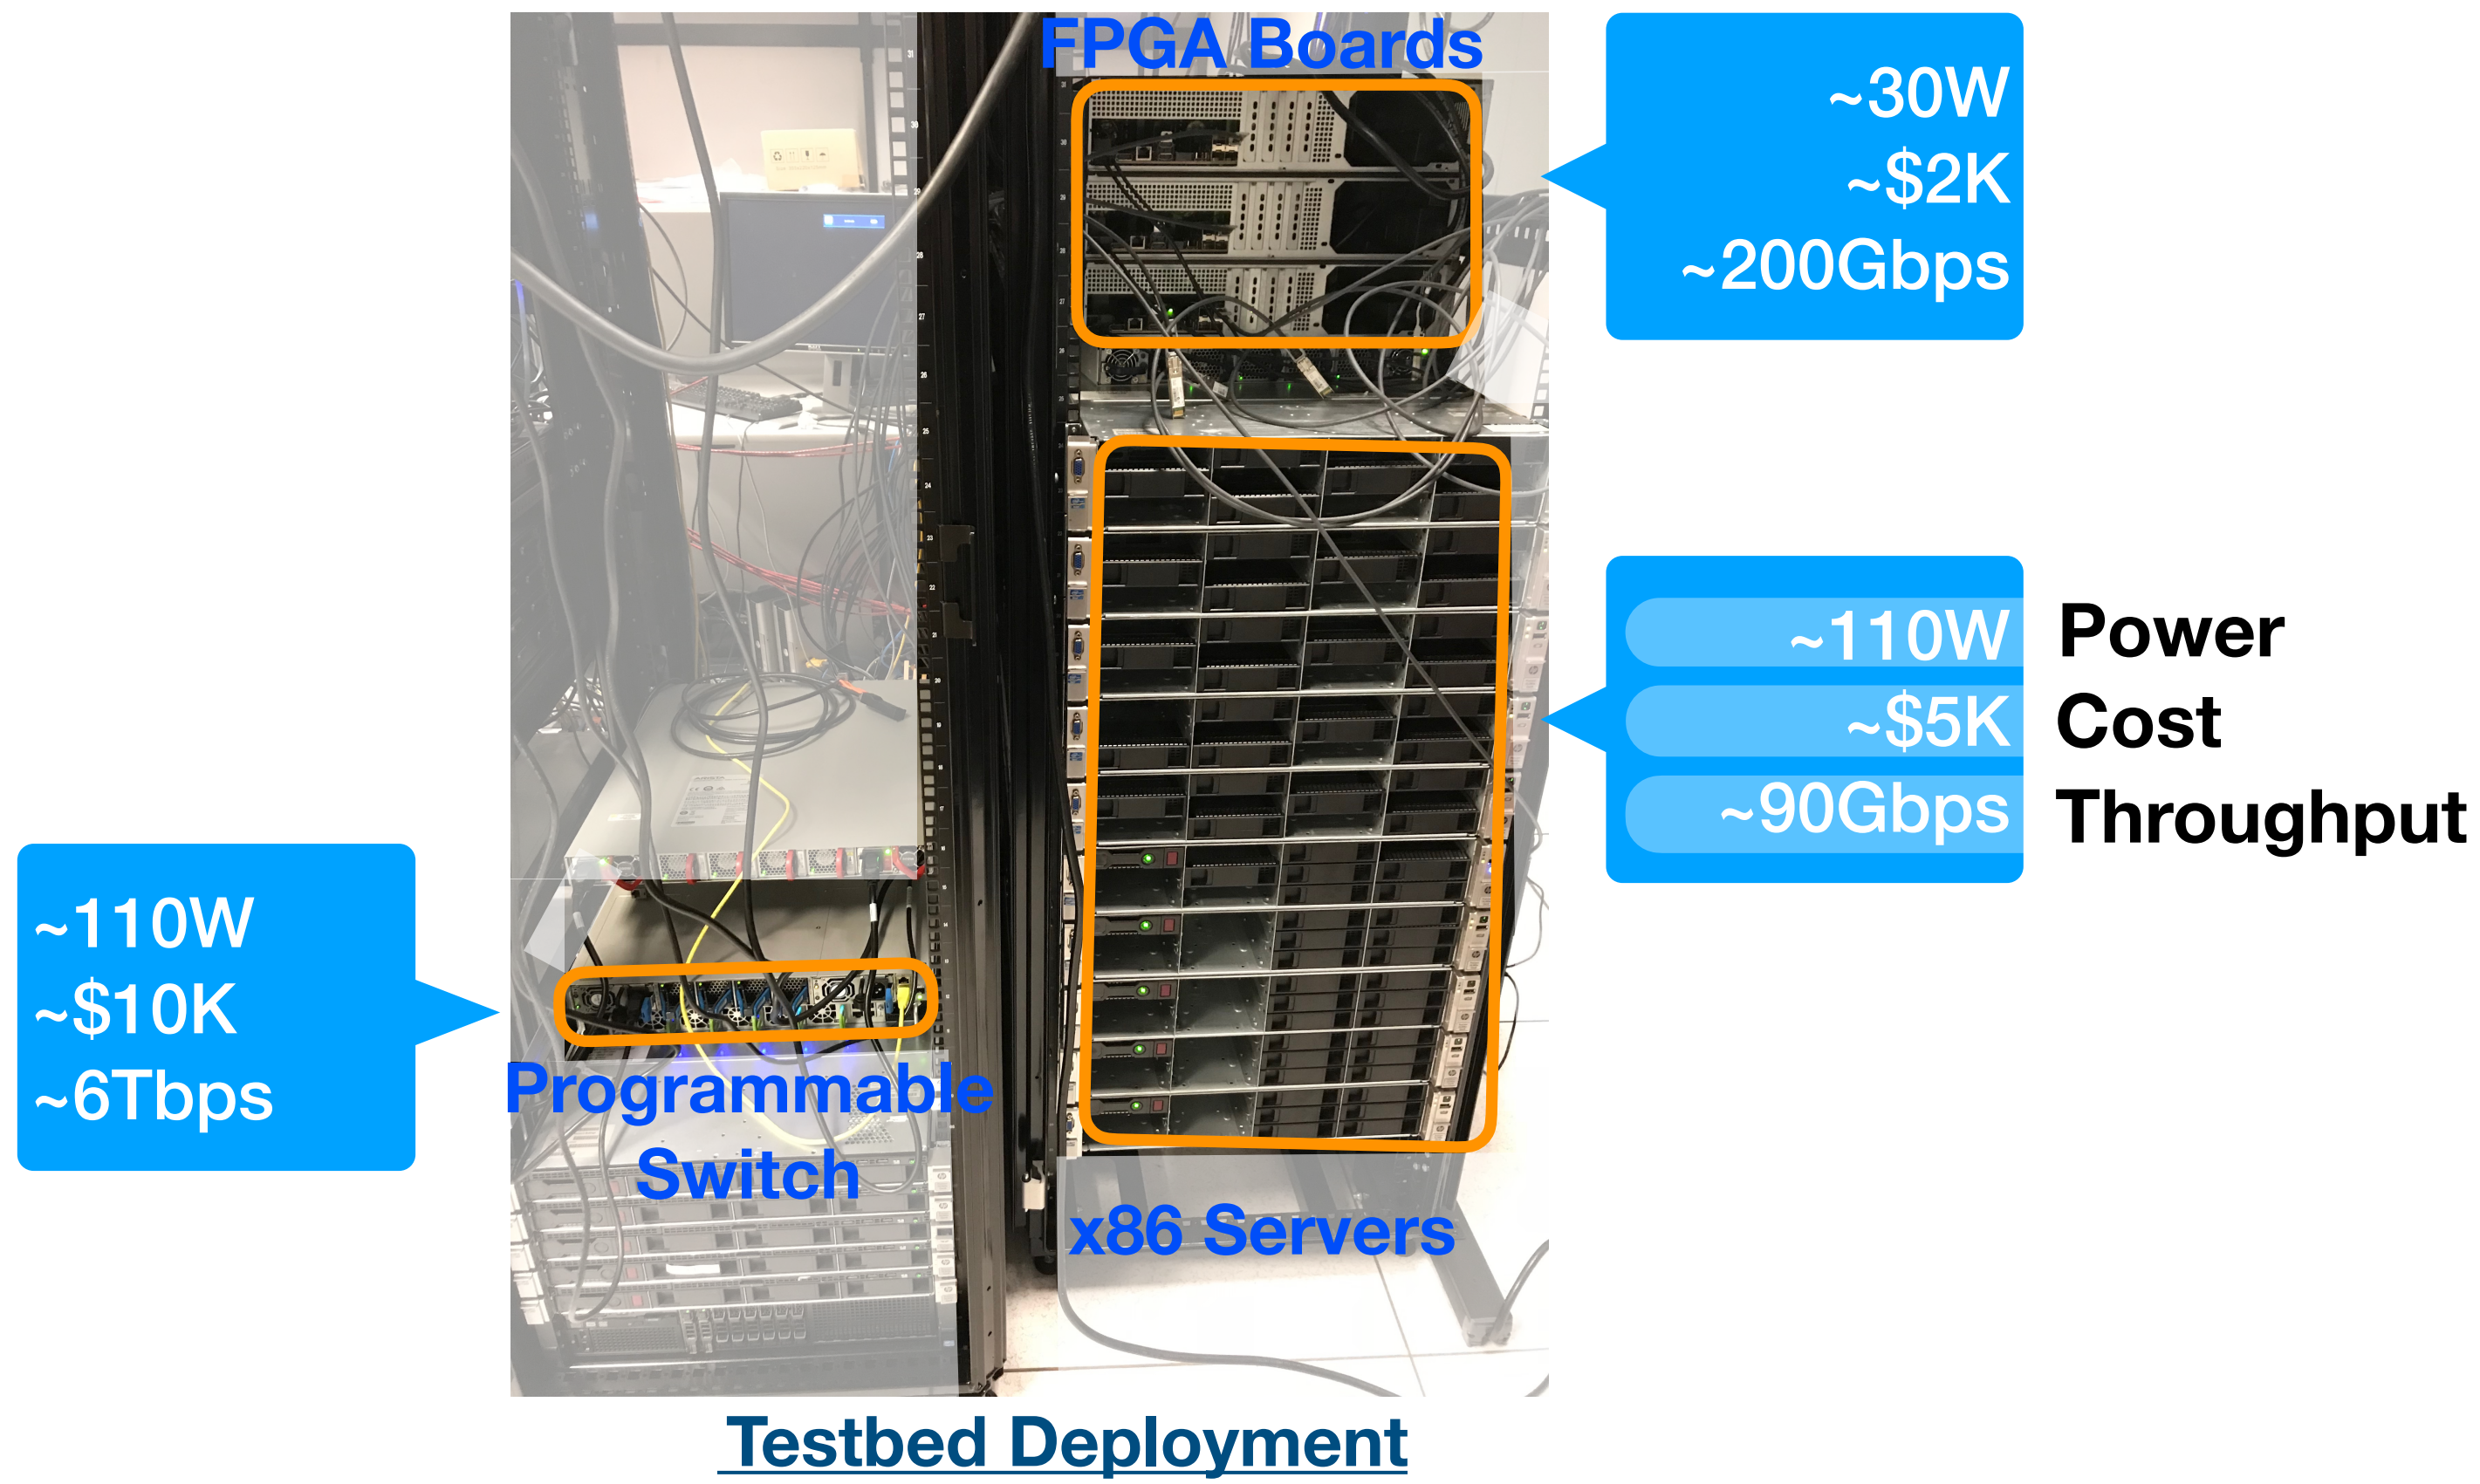


Testbed Deployment

# Heterogeneity



# Heterogeneity



Current paradigm: dataplane program → one dataplane



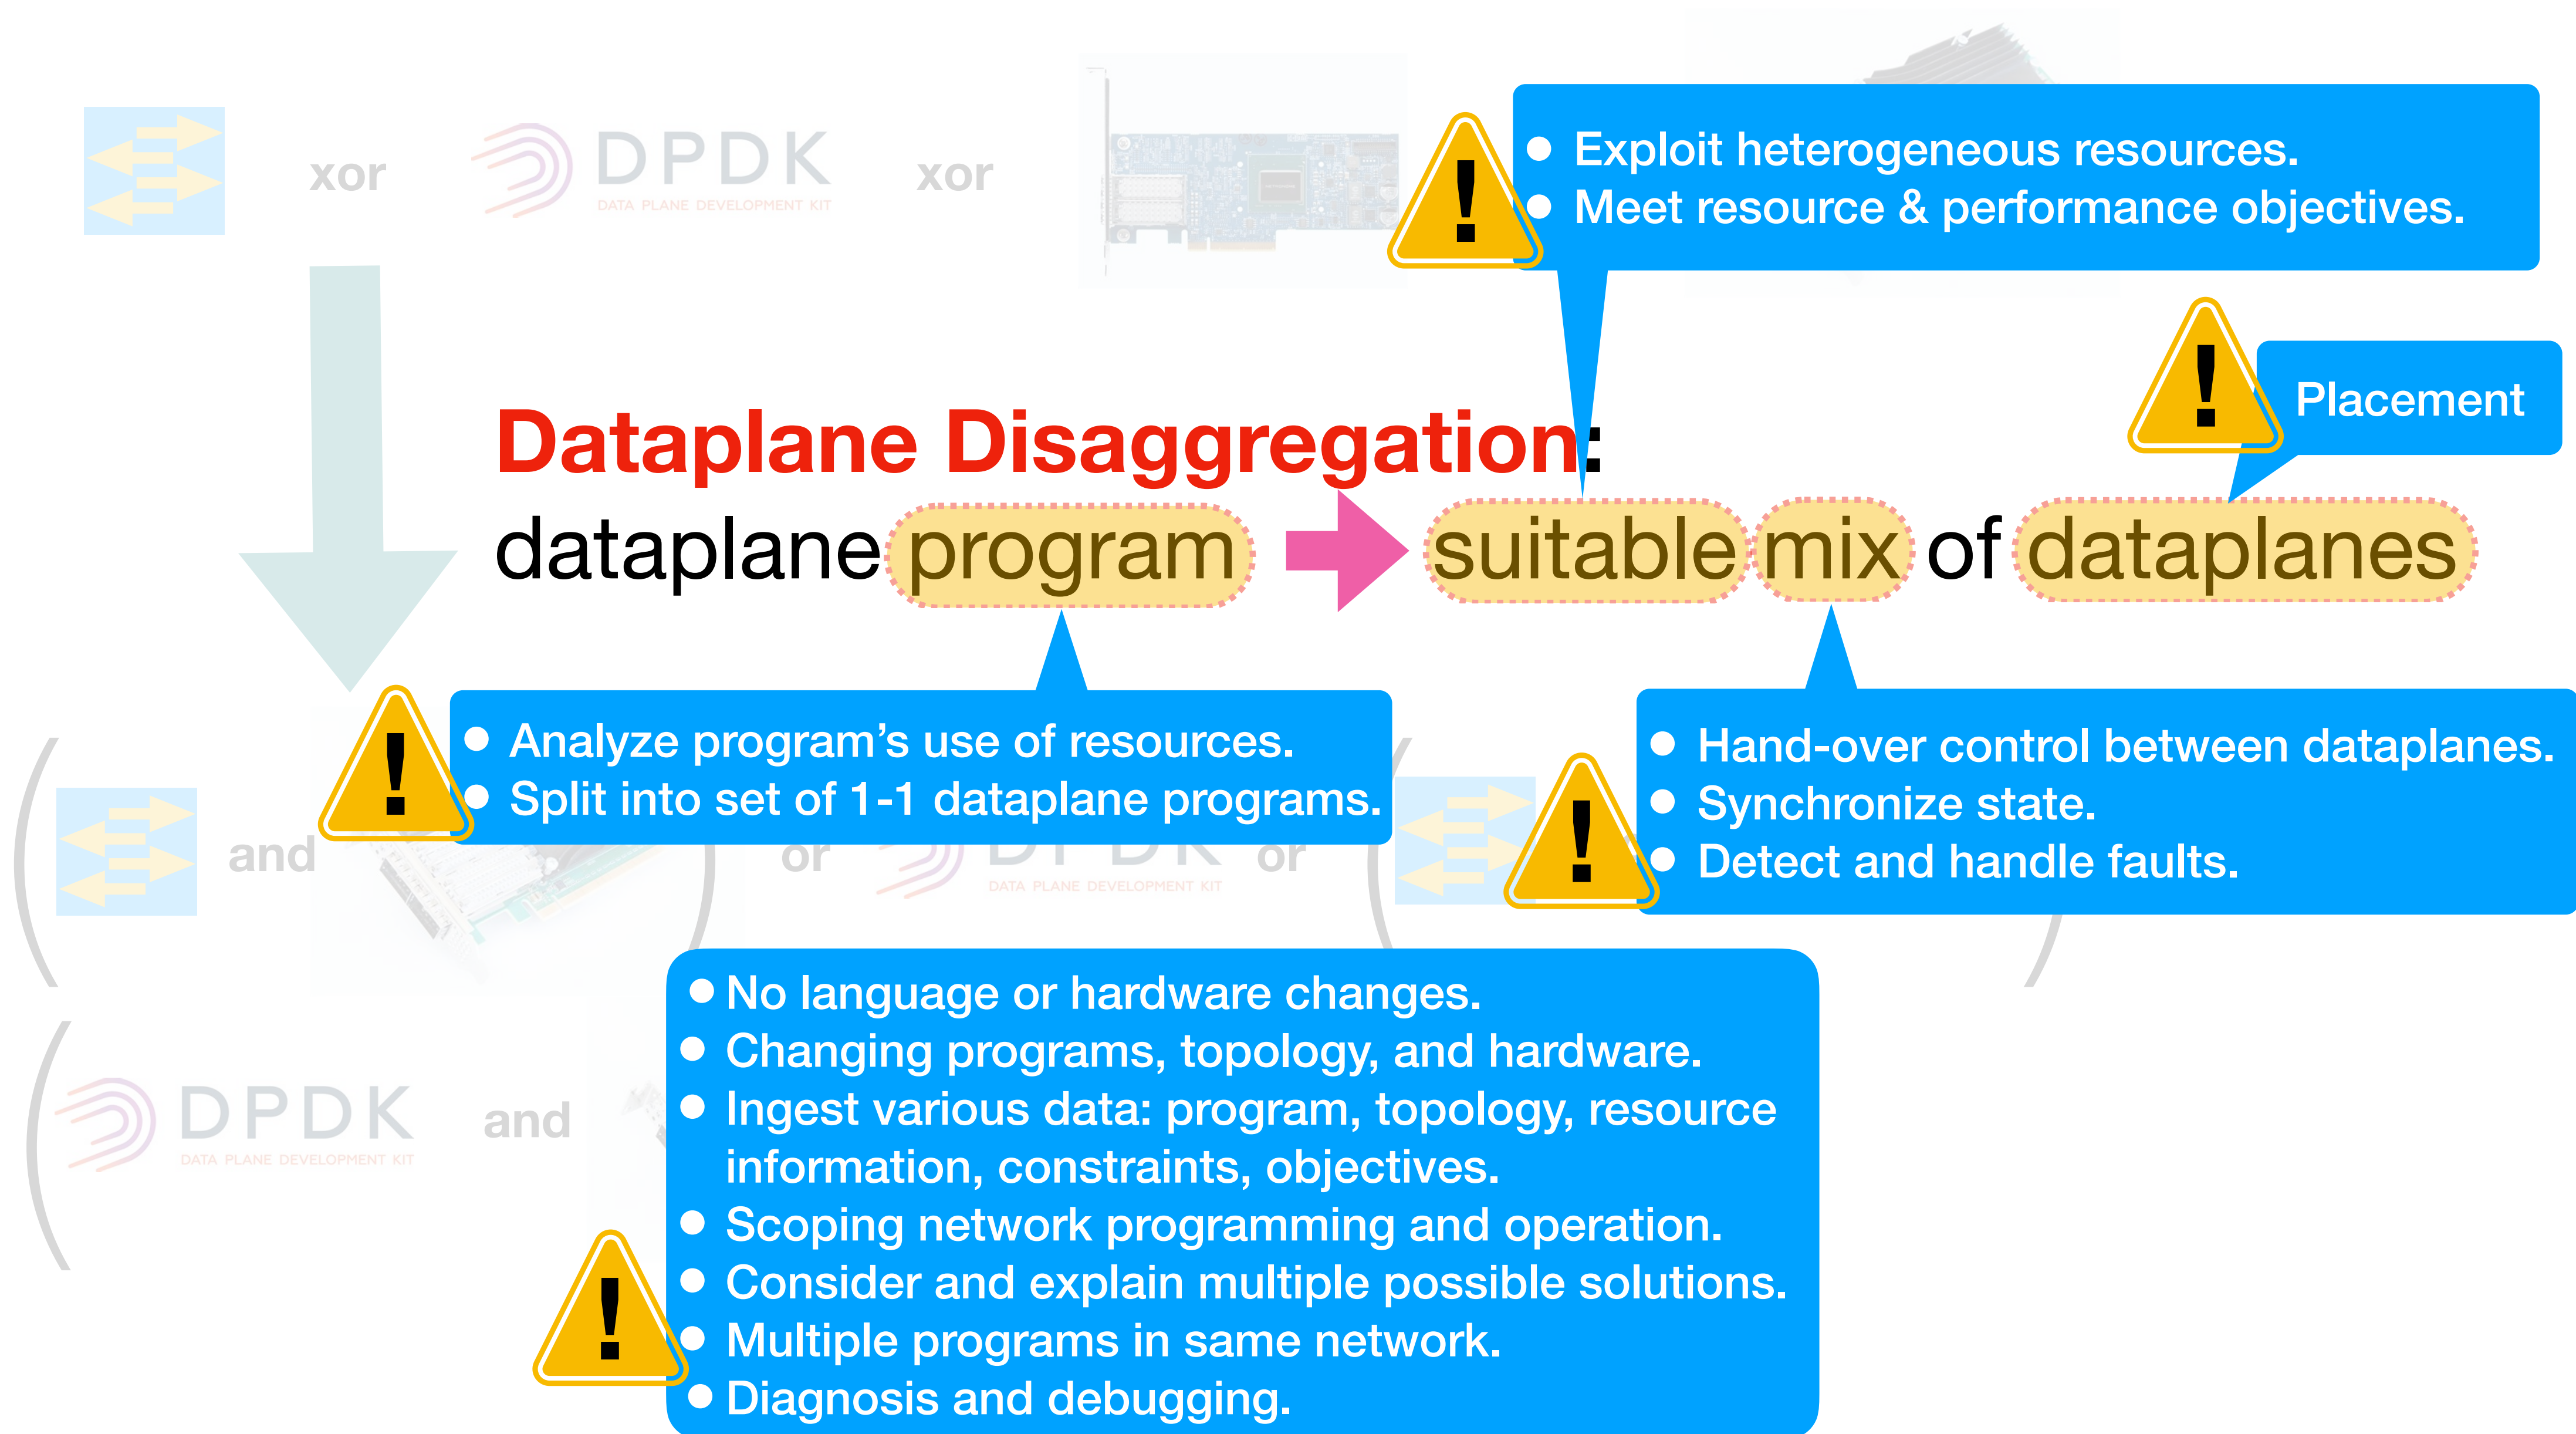
Current paradigm: dataplane program → one dataplane



Current paradigm: dataplane program → one dataplane



Current paradigm: dataplane program → one dataplane



Code, tests, scripts, data, documentation:  
<https://flightplan.cis.upenn.edu/>

program → one dataplane

# Flightplan

## Dataplane Disaggregation:

dataplane program → suitable mix of dataplanes


- Exploit heterogeneous resources.
- Meet resource & performance objectives.

Placement

- Analyze program's use of resources.
- Split into set of 1-1 dataplane programs.

- Hand-over control between dataplanes.
- Synchronize state.
- Detect and handle faults.



- No language or hardware changes.
- Changing programs, topology, and hardware.
- Ingest various data: program, topology, resource information, constraints, objectives.
- Scoping network programming and operation.
- Consider and explain multiple possible solutions.
- Multiple programs in same network.
- Diagnosis and debugging.



xor




xor



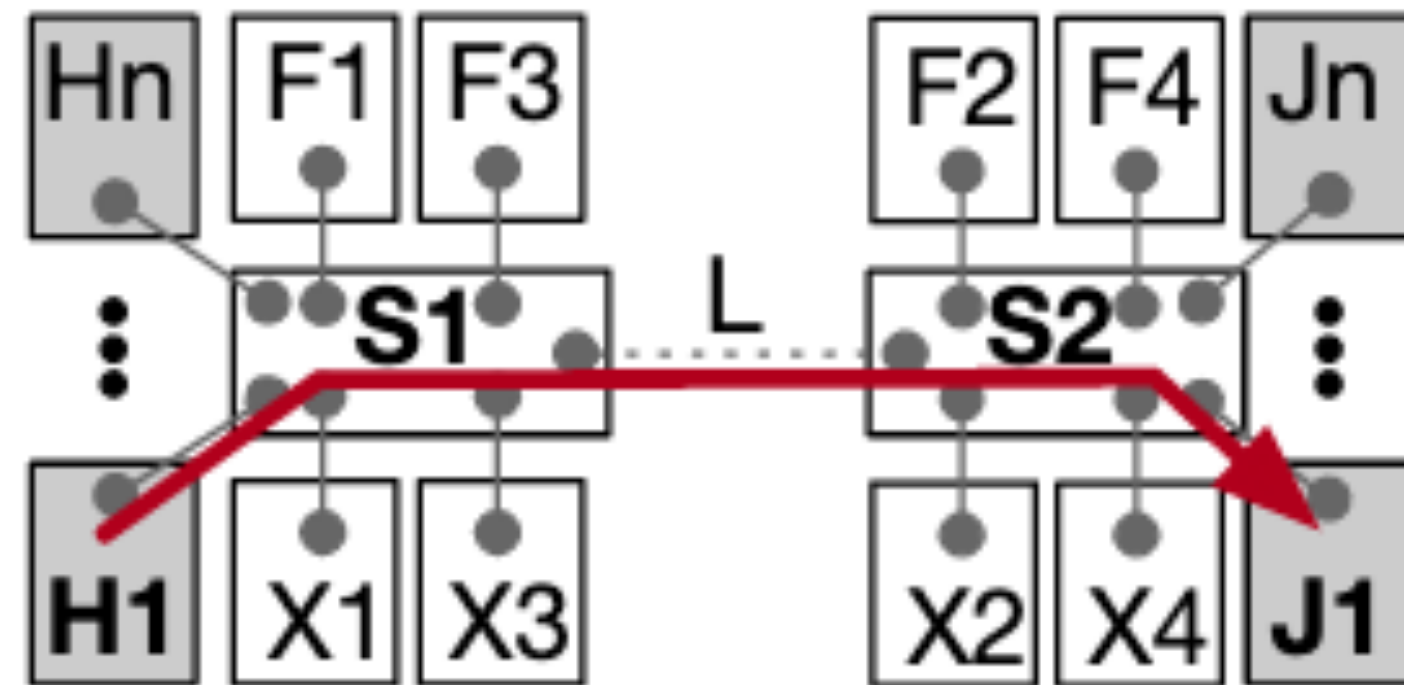
and




and



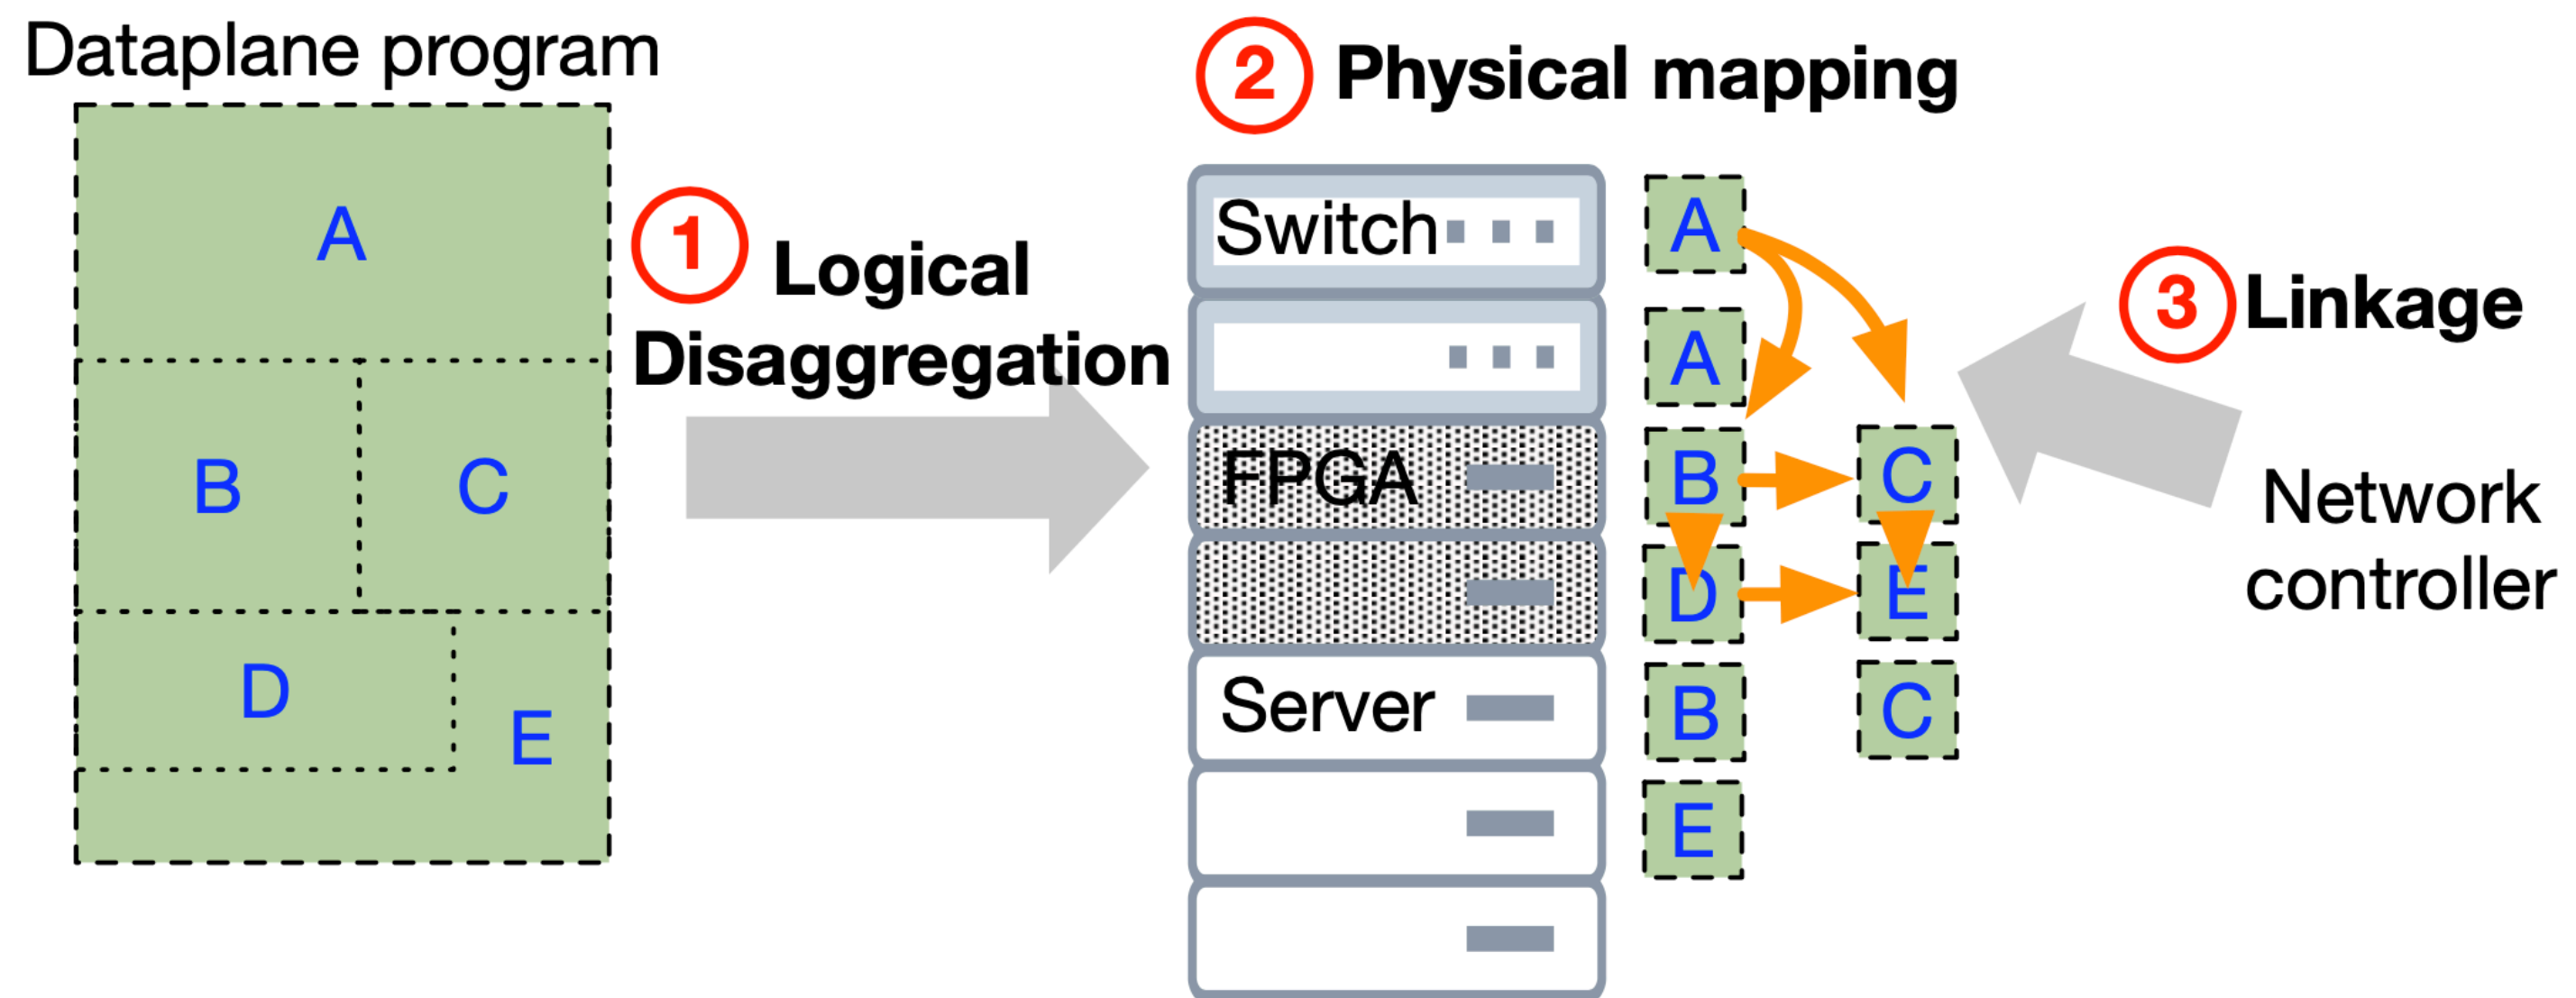




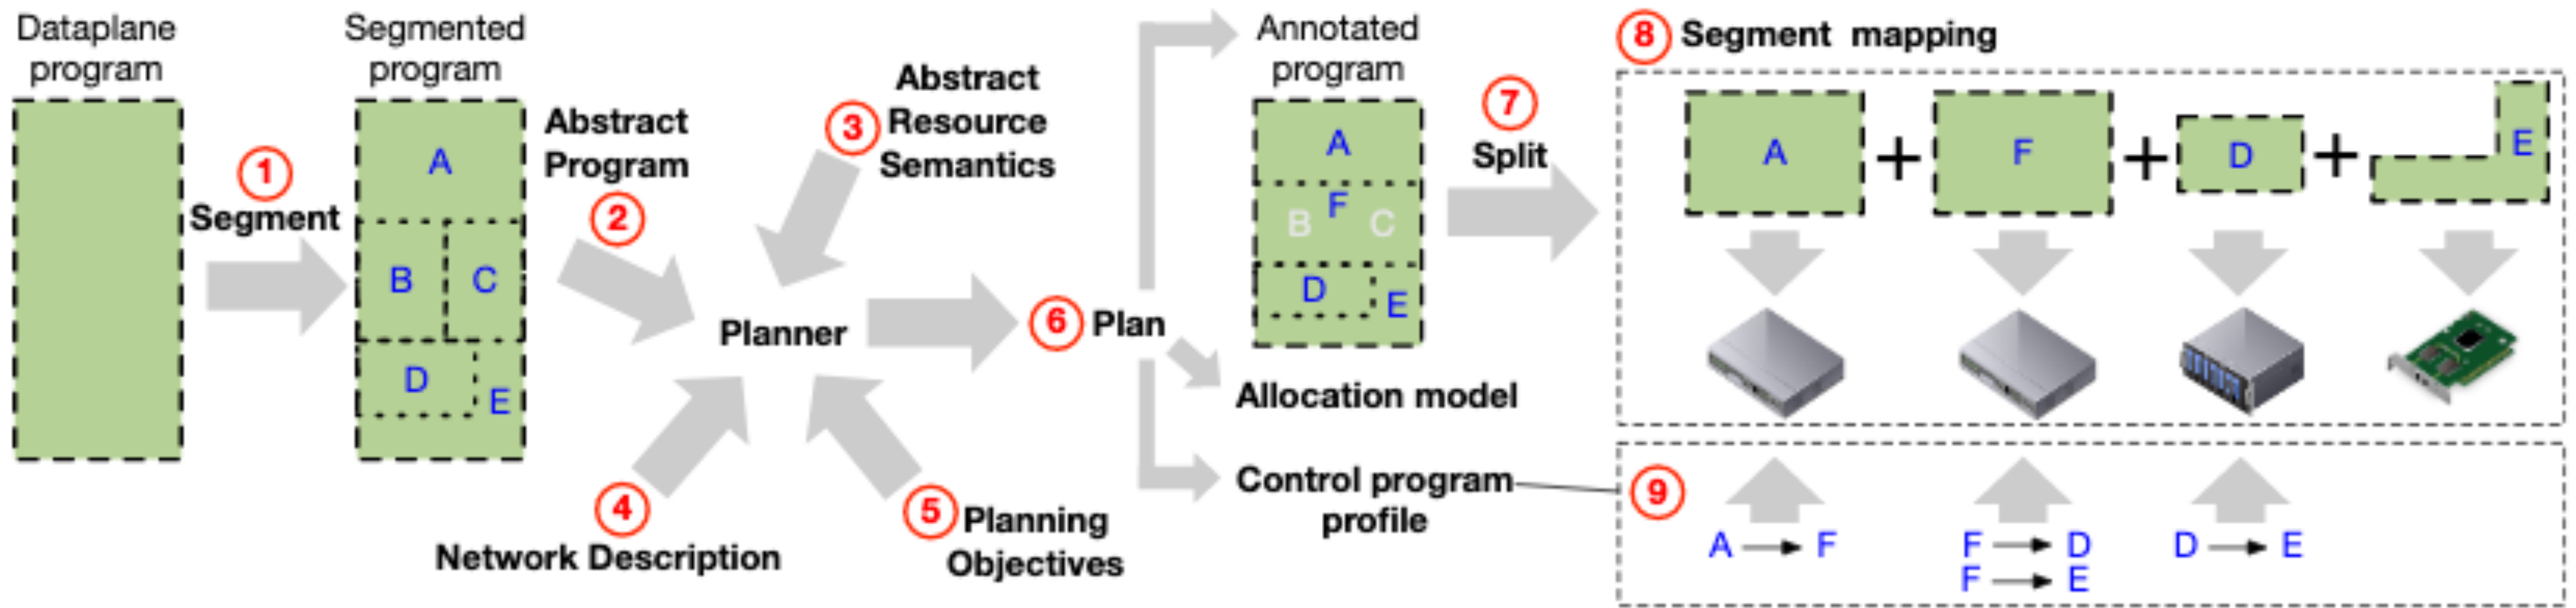
# Flightplan



# Flightplan

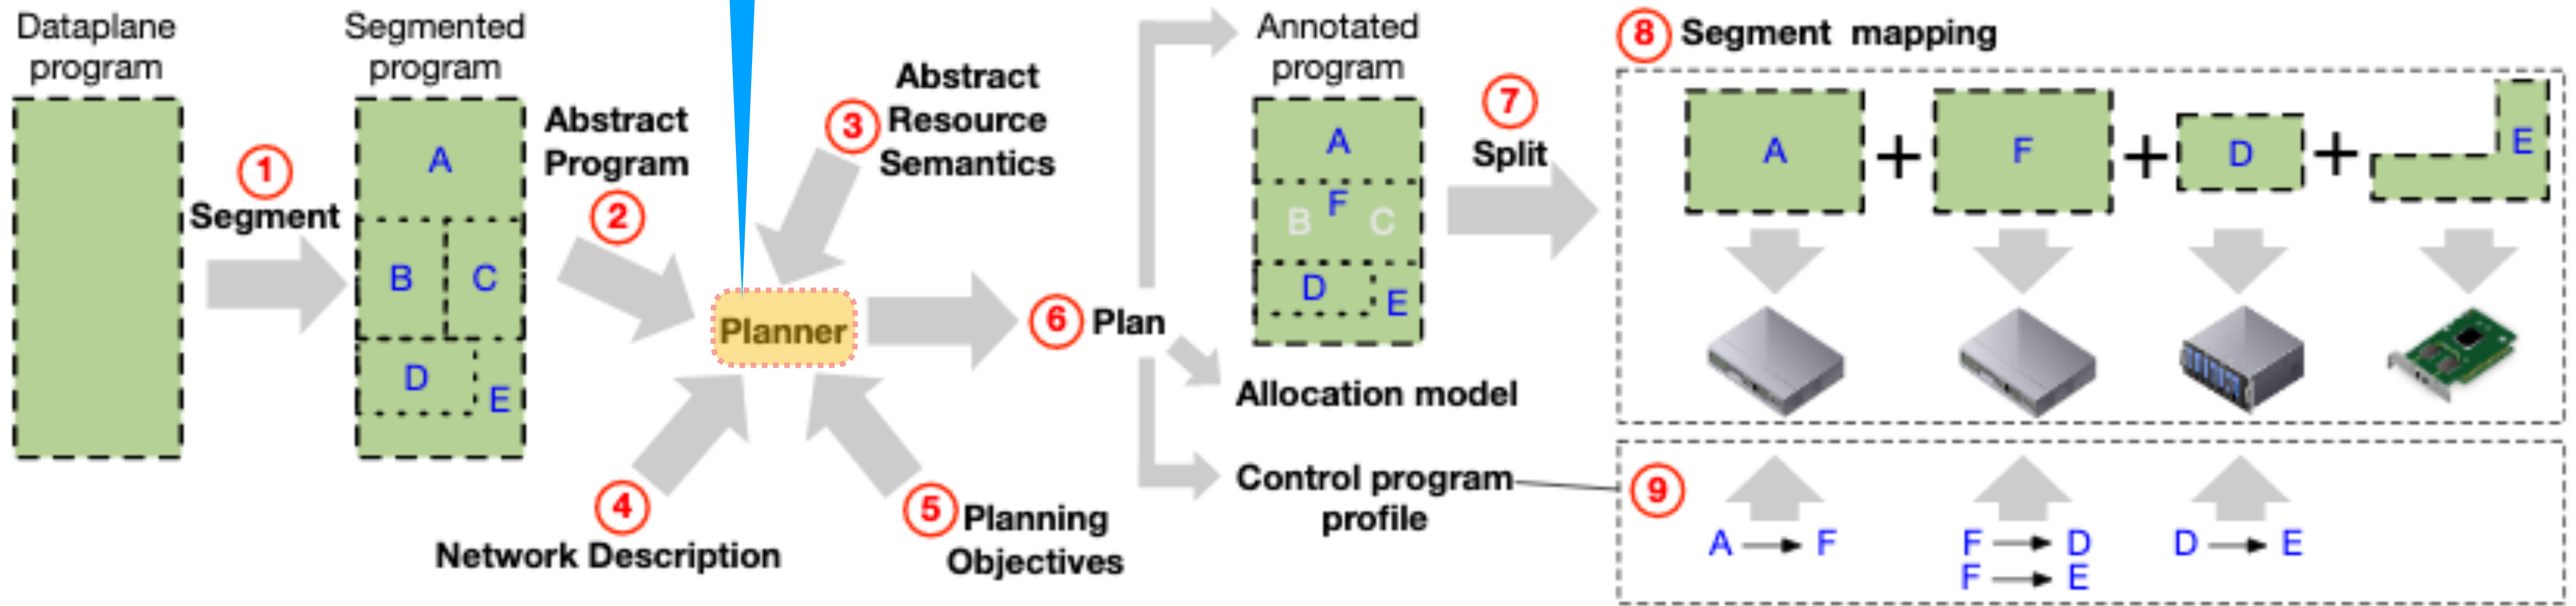


# Flightplan




# Flightplan

Idea from symbolic AI:  
rule-based search.



# Flightplan



# Flightplan

# Example program (Crosspod)

Segment annotation

Resource  
dependence

```
1  bit<1> compressed_link = 0;
2  bit<1> run_fec_egress = 0;
3  ...
4  flyto(Compress);
5  // If heading out on a multiplexed link, then header compress.
6  egress_compression.apply(meta.egress_spec, compressed_link);
7  if (compressed_link == 1) {
8      header_compress(forward);
9      if (forward == 0) {
10         drop();
11         return;
12     }
13 }
14 flyto(FEC_Encode);
15 check_run_FEC_egress.apply();
16 // If heading out on a lossy link, then FEC encode.
17 if (run_fec_egress == 1) {
18     ...
19     classification.apply(hdr, proto_and_port); // Sets hdr.fec.isValid()
20     if (hdr.fec.isValid()) {
21         encoder_params.apply(hdr.fec.traffic_class, k, h);
22         update_fec_state(hdr.fec.traffic_class, k, h,
23                         hdr.fec.block_index, hdr.fec.packet_index);
24         hdr.fec.orig_ethertype = hdr.eth.type;
25         FEC_ENCODE(hdr.fec, k, h);
26     }
27     ...
28 }
```

# Abstract program


```

flyto(Compress);
// If heading out on a multiplexed link, then header compress.
egress_compression.apply(meta.egress_spec, compressed_link);
if (compressed_link == 1) {
    header_compress(forward);
    if (forward == 0) {
        drop();
        return;
    }
}
flyto(FEC_Encode);
    
```

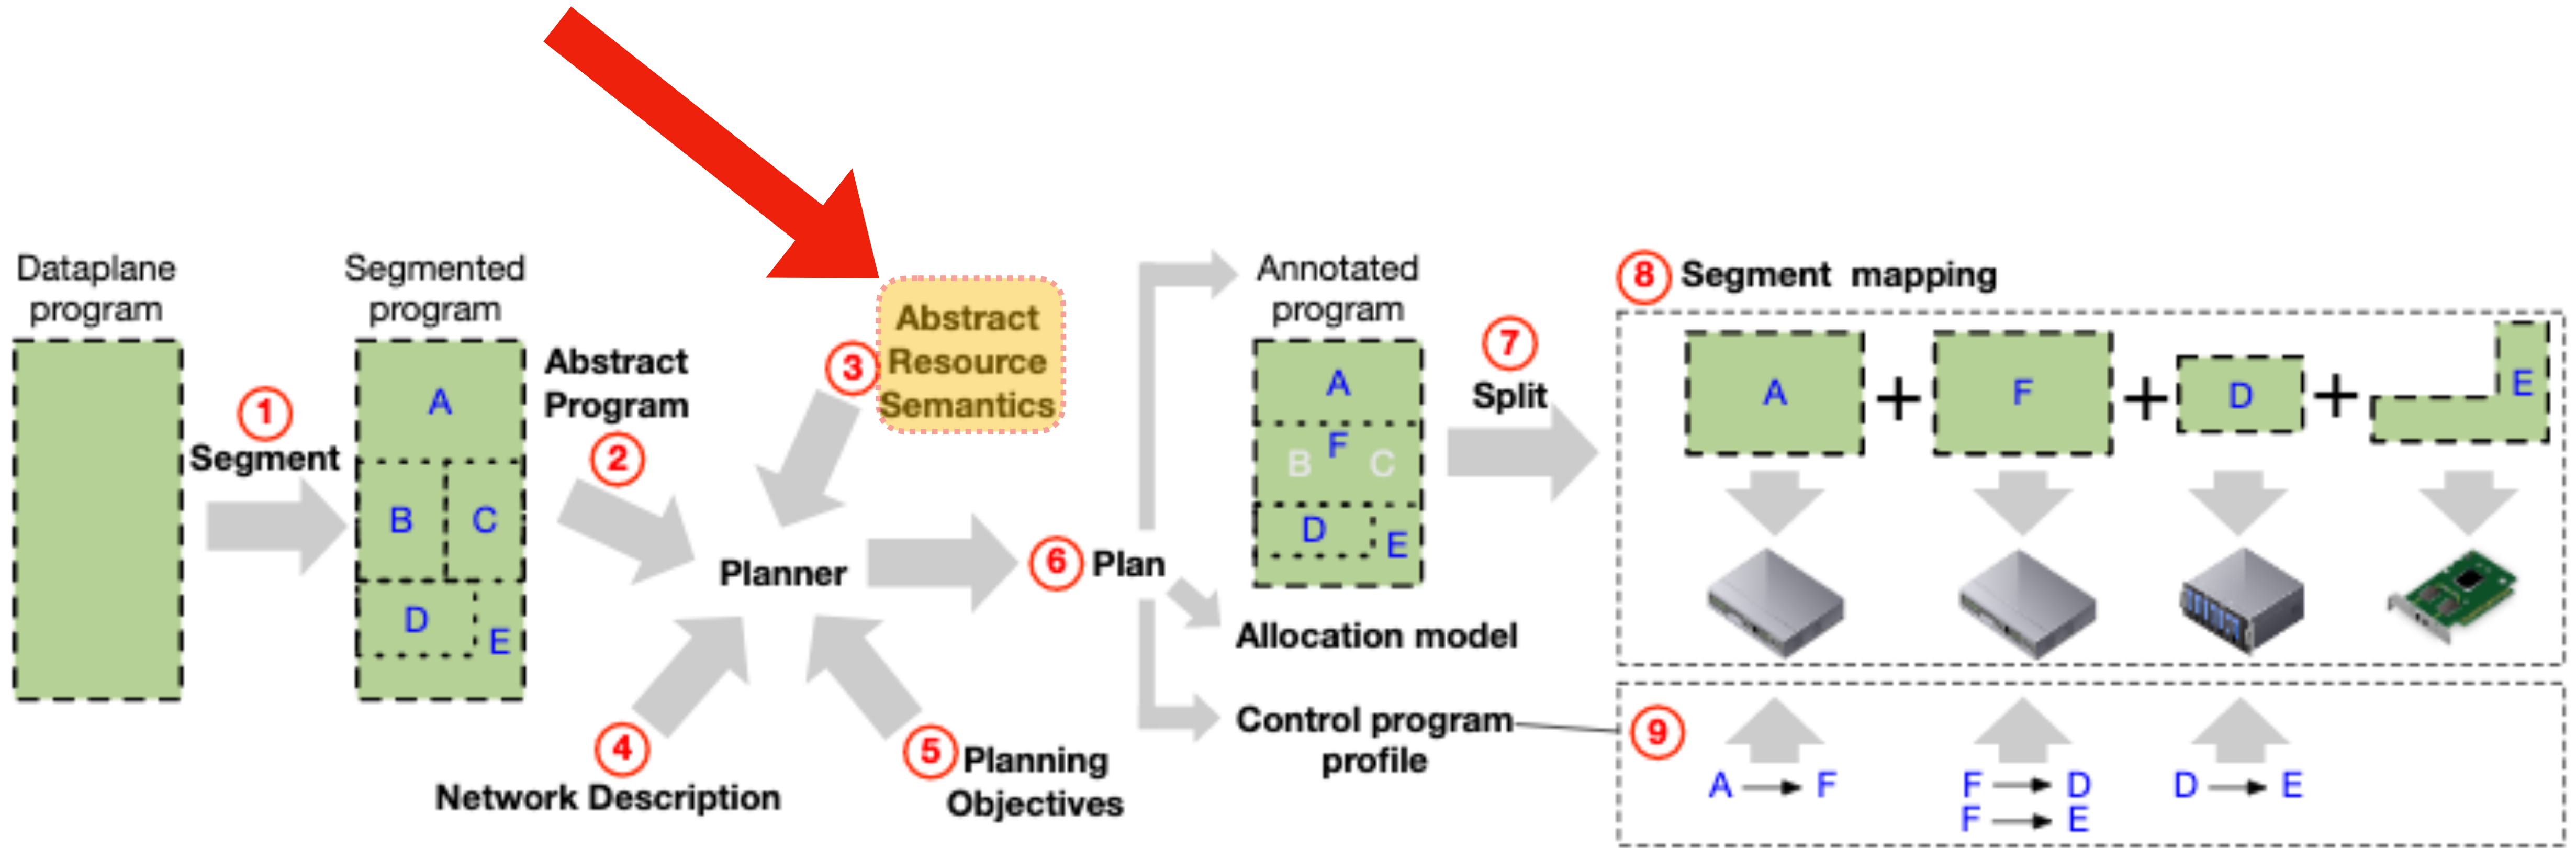
(Rule generation is fully automatic)

egress_compres.	header_compress	drop
R2	Compress	






# Flightplan




# Flightplan

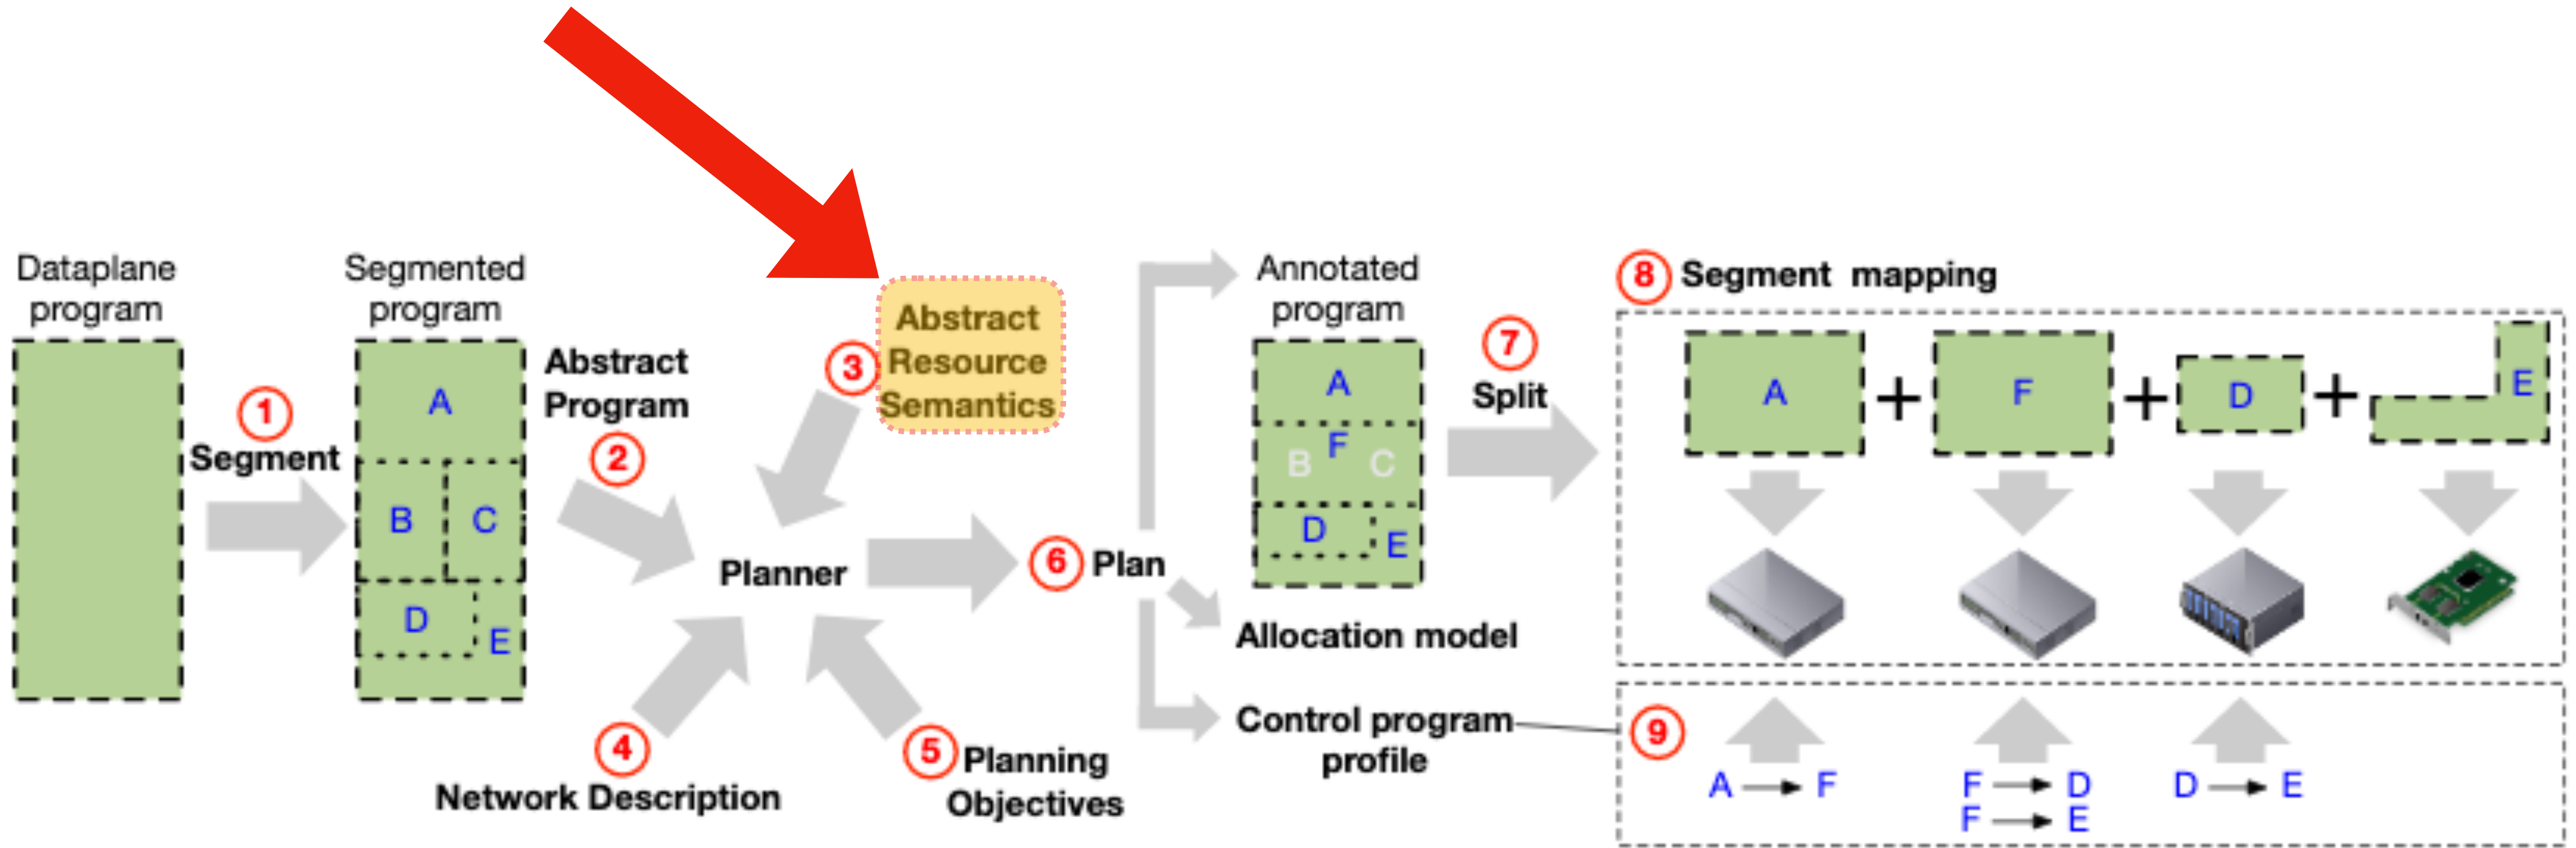
# Abstract Resource Semantics



$$\frac{\text{CPU} \quad \text{Rate} < 2 \times 10^8 \quad \text{PacketSize} > 1000}{\text{header\_compress}} \left[ \begin{array}{l} \text{Lat.} \mapsto \text{Lat.} + 7.4 \times 10^{-3} \\ \text{Rate} \mapsto \text{Rate} \times \frac{189.9}{194.75} \\ \text{once Power} \mapsto \text{Power} + 150 \text{ W} \\ \text{once Cost} \mapsto \text{Cost} + 5 \end{array} \right]$$

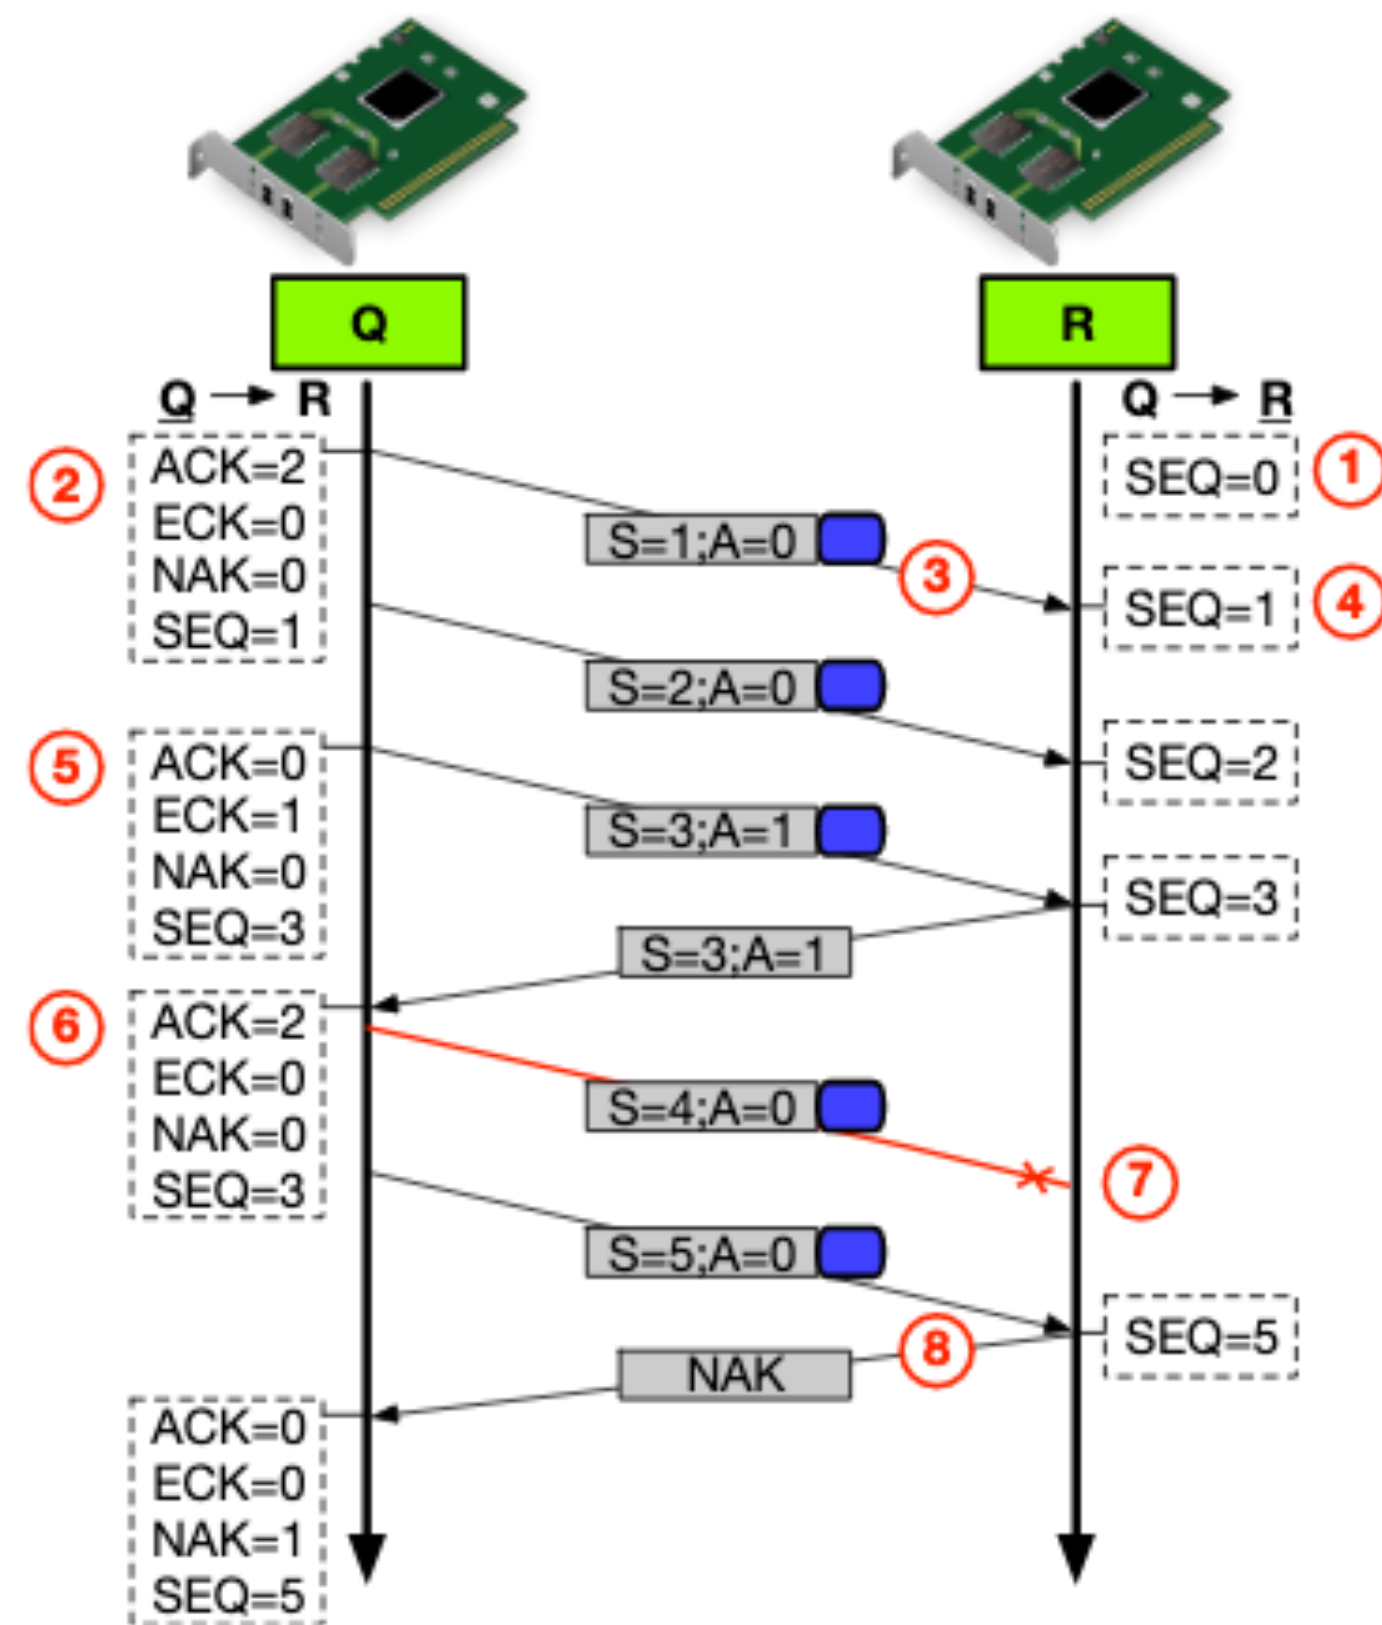


$$\frac{\text{FPGA} \quad \text{Rate} < 9.5 \times 10^9 \quad \text{PacketSize} > 100}{\text{header\_compress}} \left[ \begin{array}{l} \text{Lat.} \mapsto \text{Lat.} + 6.44 \times 10^{-6} \\ \text{Rate} \mapsto \text{Rate} \times \frac{9.15}{9.3} \\ \langle \text{LUTs} \rangle \mapsto \text{LUTs} + 24.4\% \\ \langle \text{BRAMs} \rangle \mapsto \text{BRAMs} + 54.4\% \\ \langle \text{FF} \rangle \mapsto \text{FF} + 15.8\% \\ \text{once Power} \mapsto \text{Power} + 30 \text{ W} \\ \text{once Cost} \mapsto \text{Cost} + 2 \end{array} \right]$$



# Flightplan

# Runtime: Fault Detection + Handling

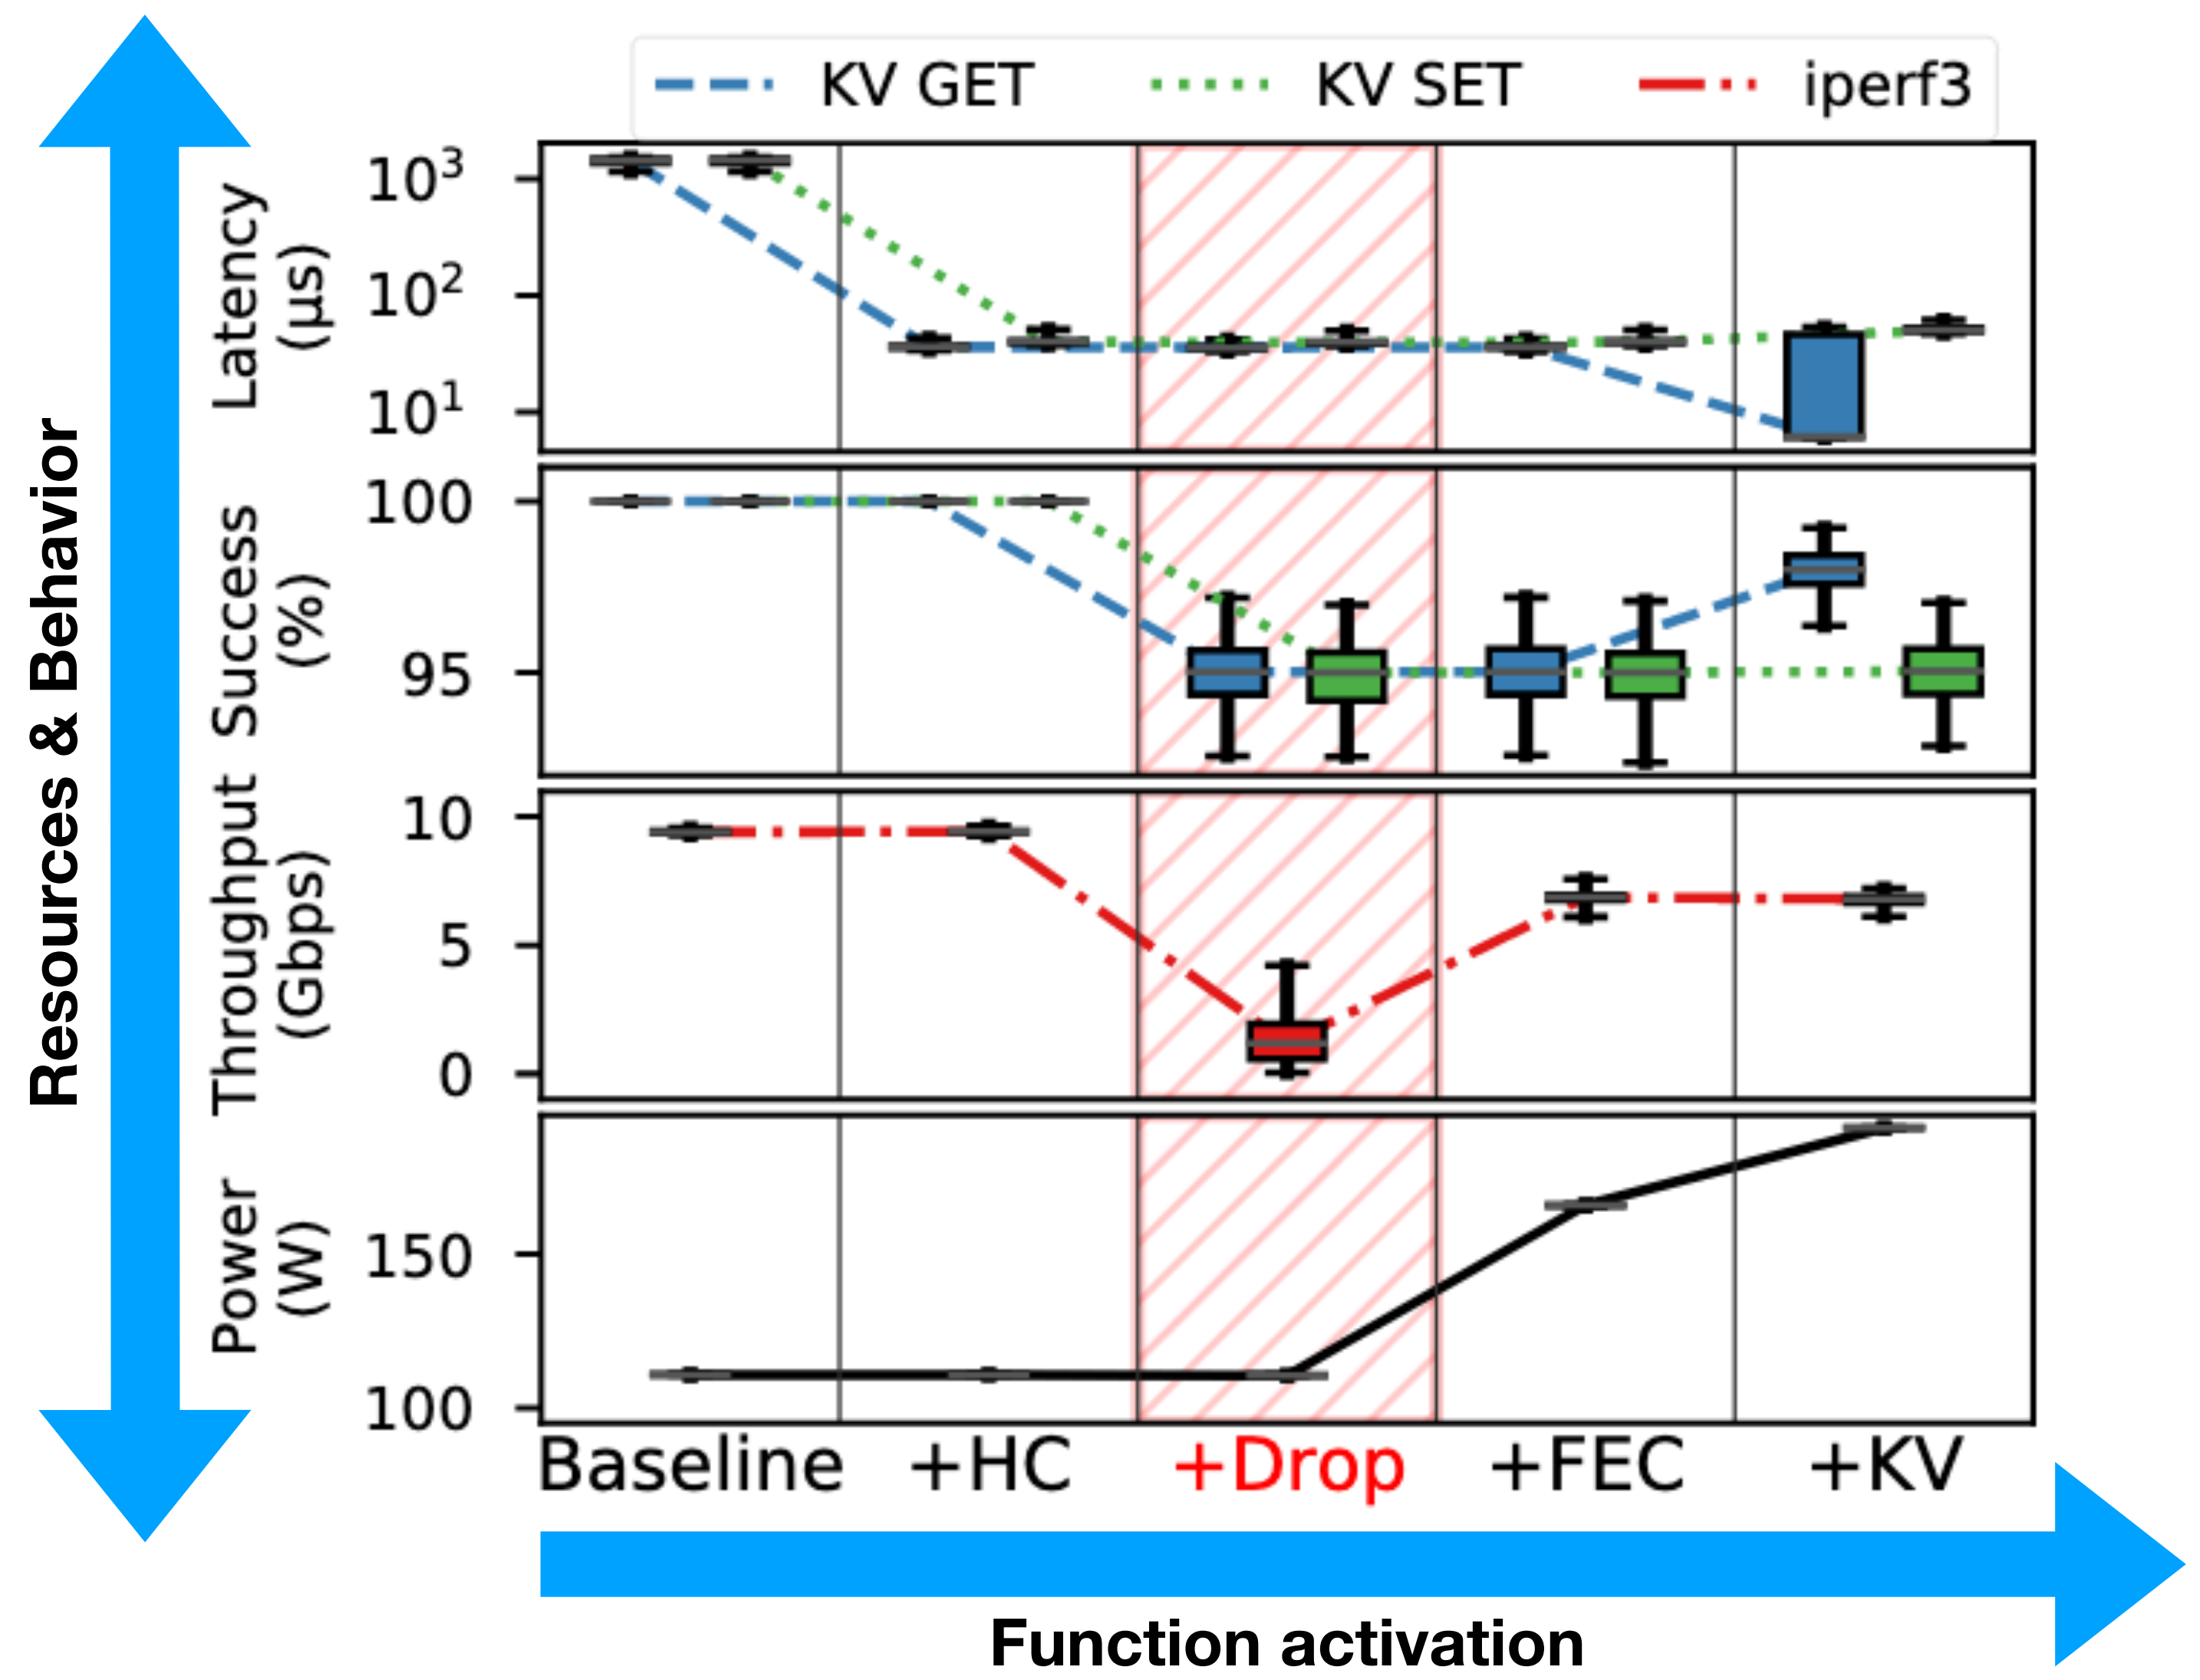


- Two mechanisms**
- In-dataplane: +ve and -ve Acks.
  - Strokes from control program.

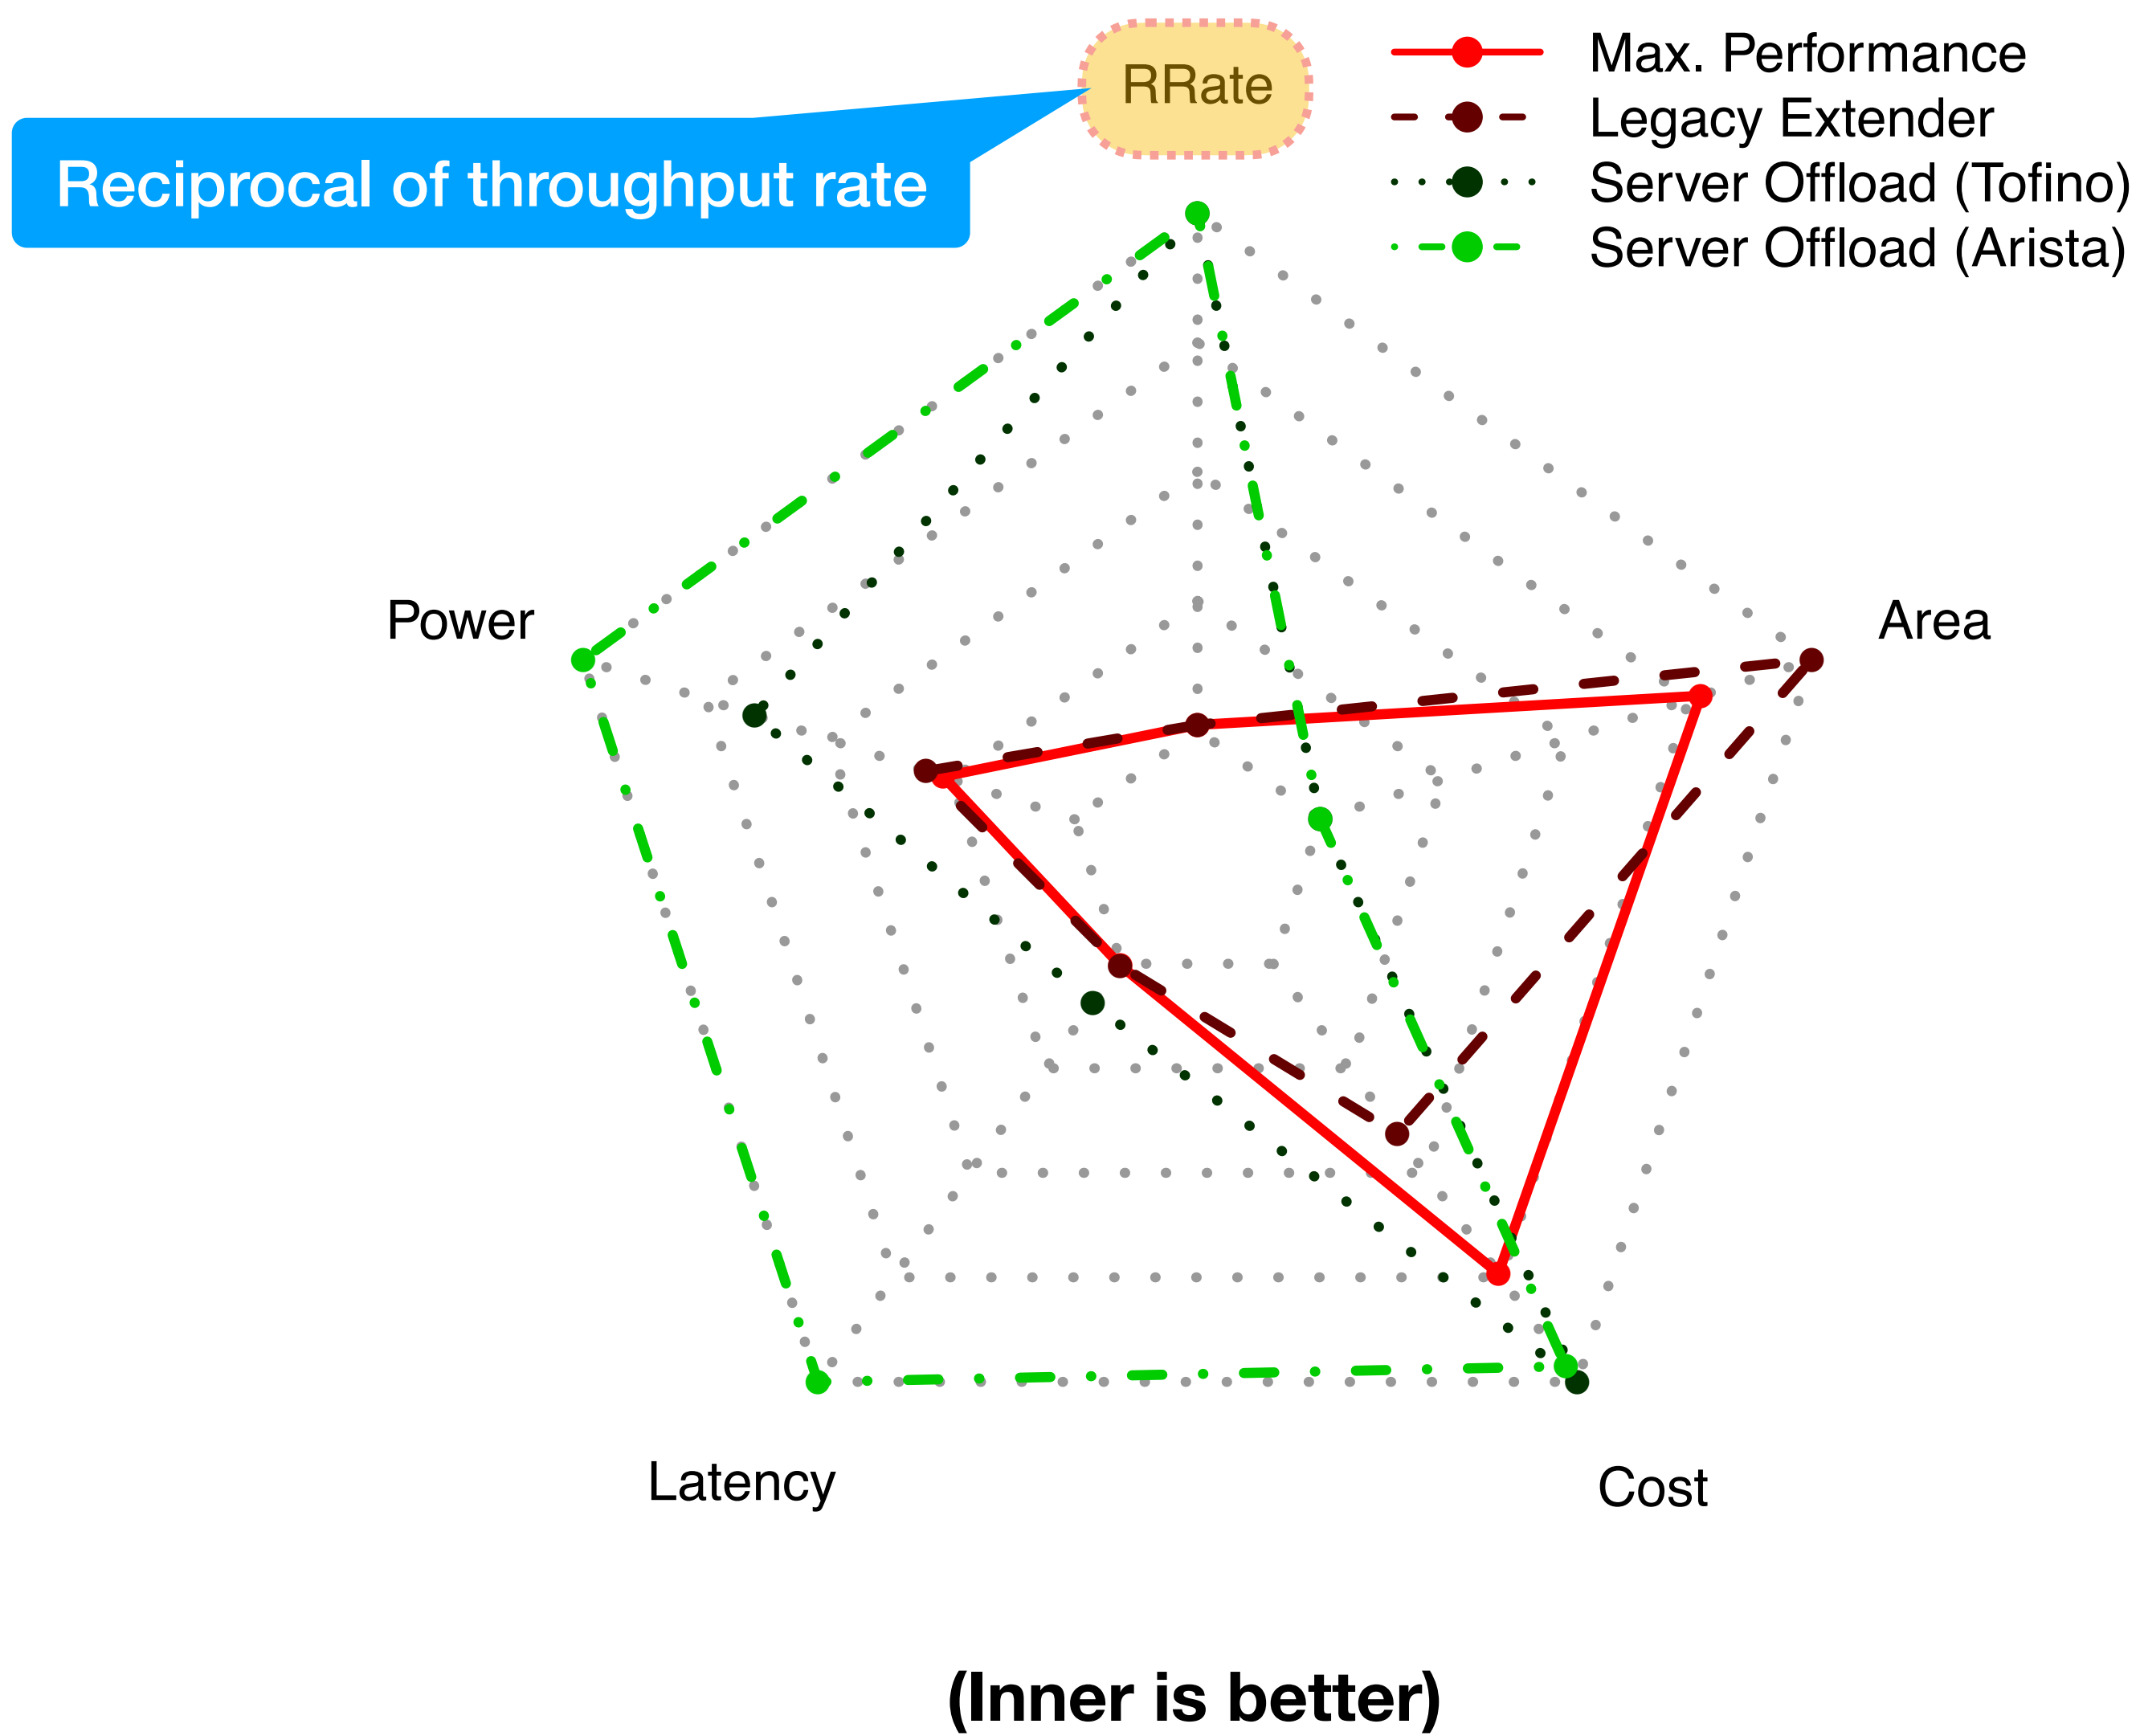
# Evaluation

- **Simulation:**
  - Scale of the network (featuring various programs)
  - Overhead
  - Disaggregation (different programs split in different ways)
  - Fail-over
- **Test-bed:**
  - Throughput, latency, power, resource utilization
  - Plan comparisons for hardware alternatives
  - Single-feature evaluation

# Fig 10: Resources & Behavior vs Functions (Testbed experiment)

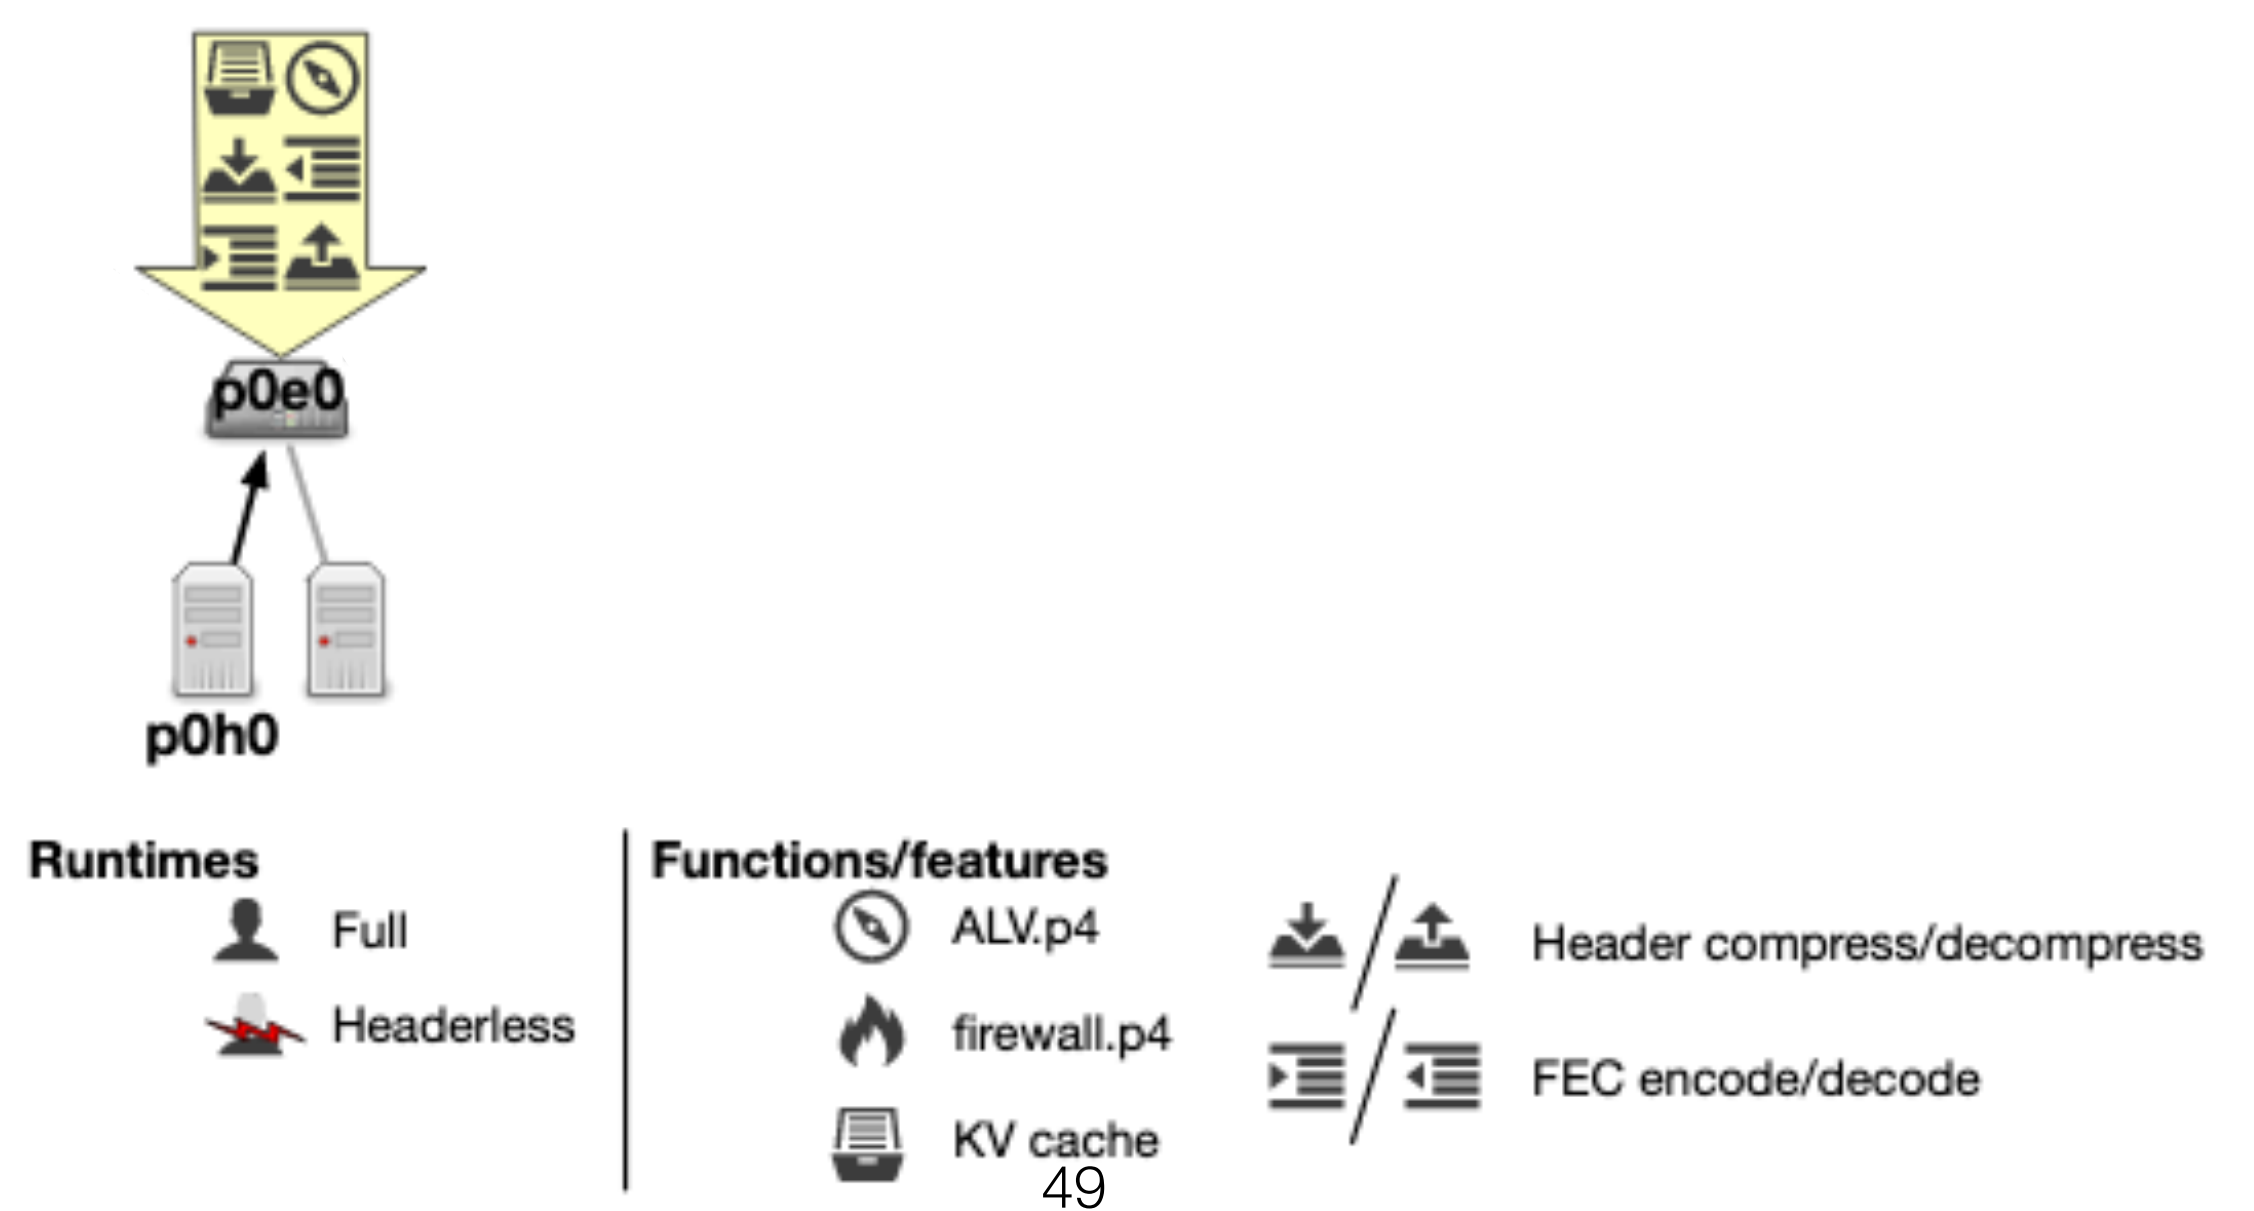


# Fig 9: Plan comparison (Testbed + Simulation)

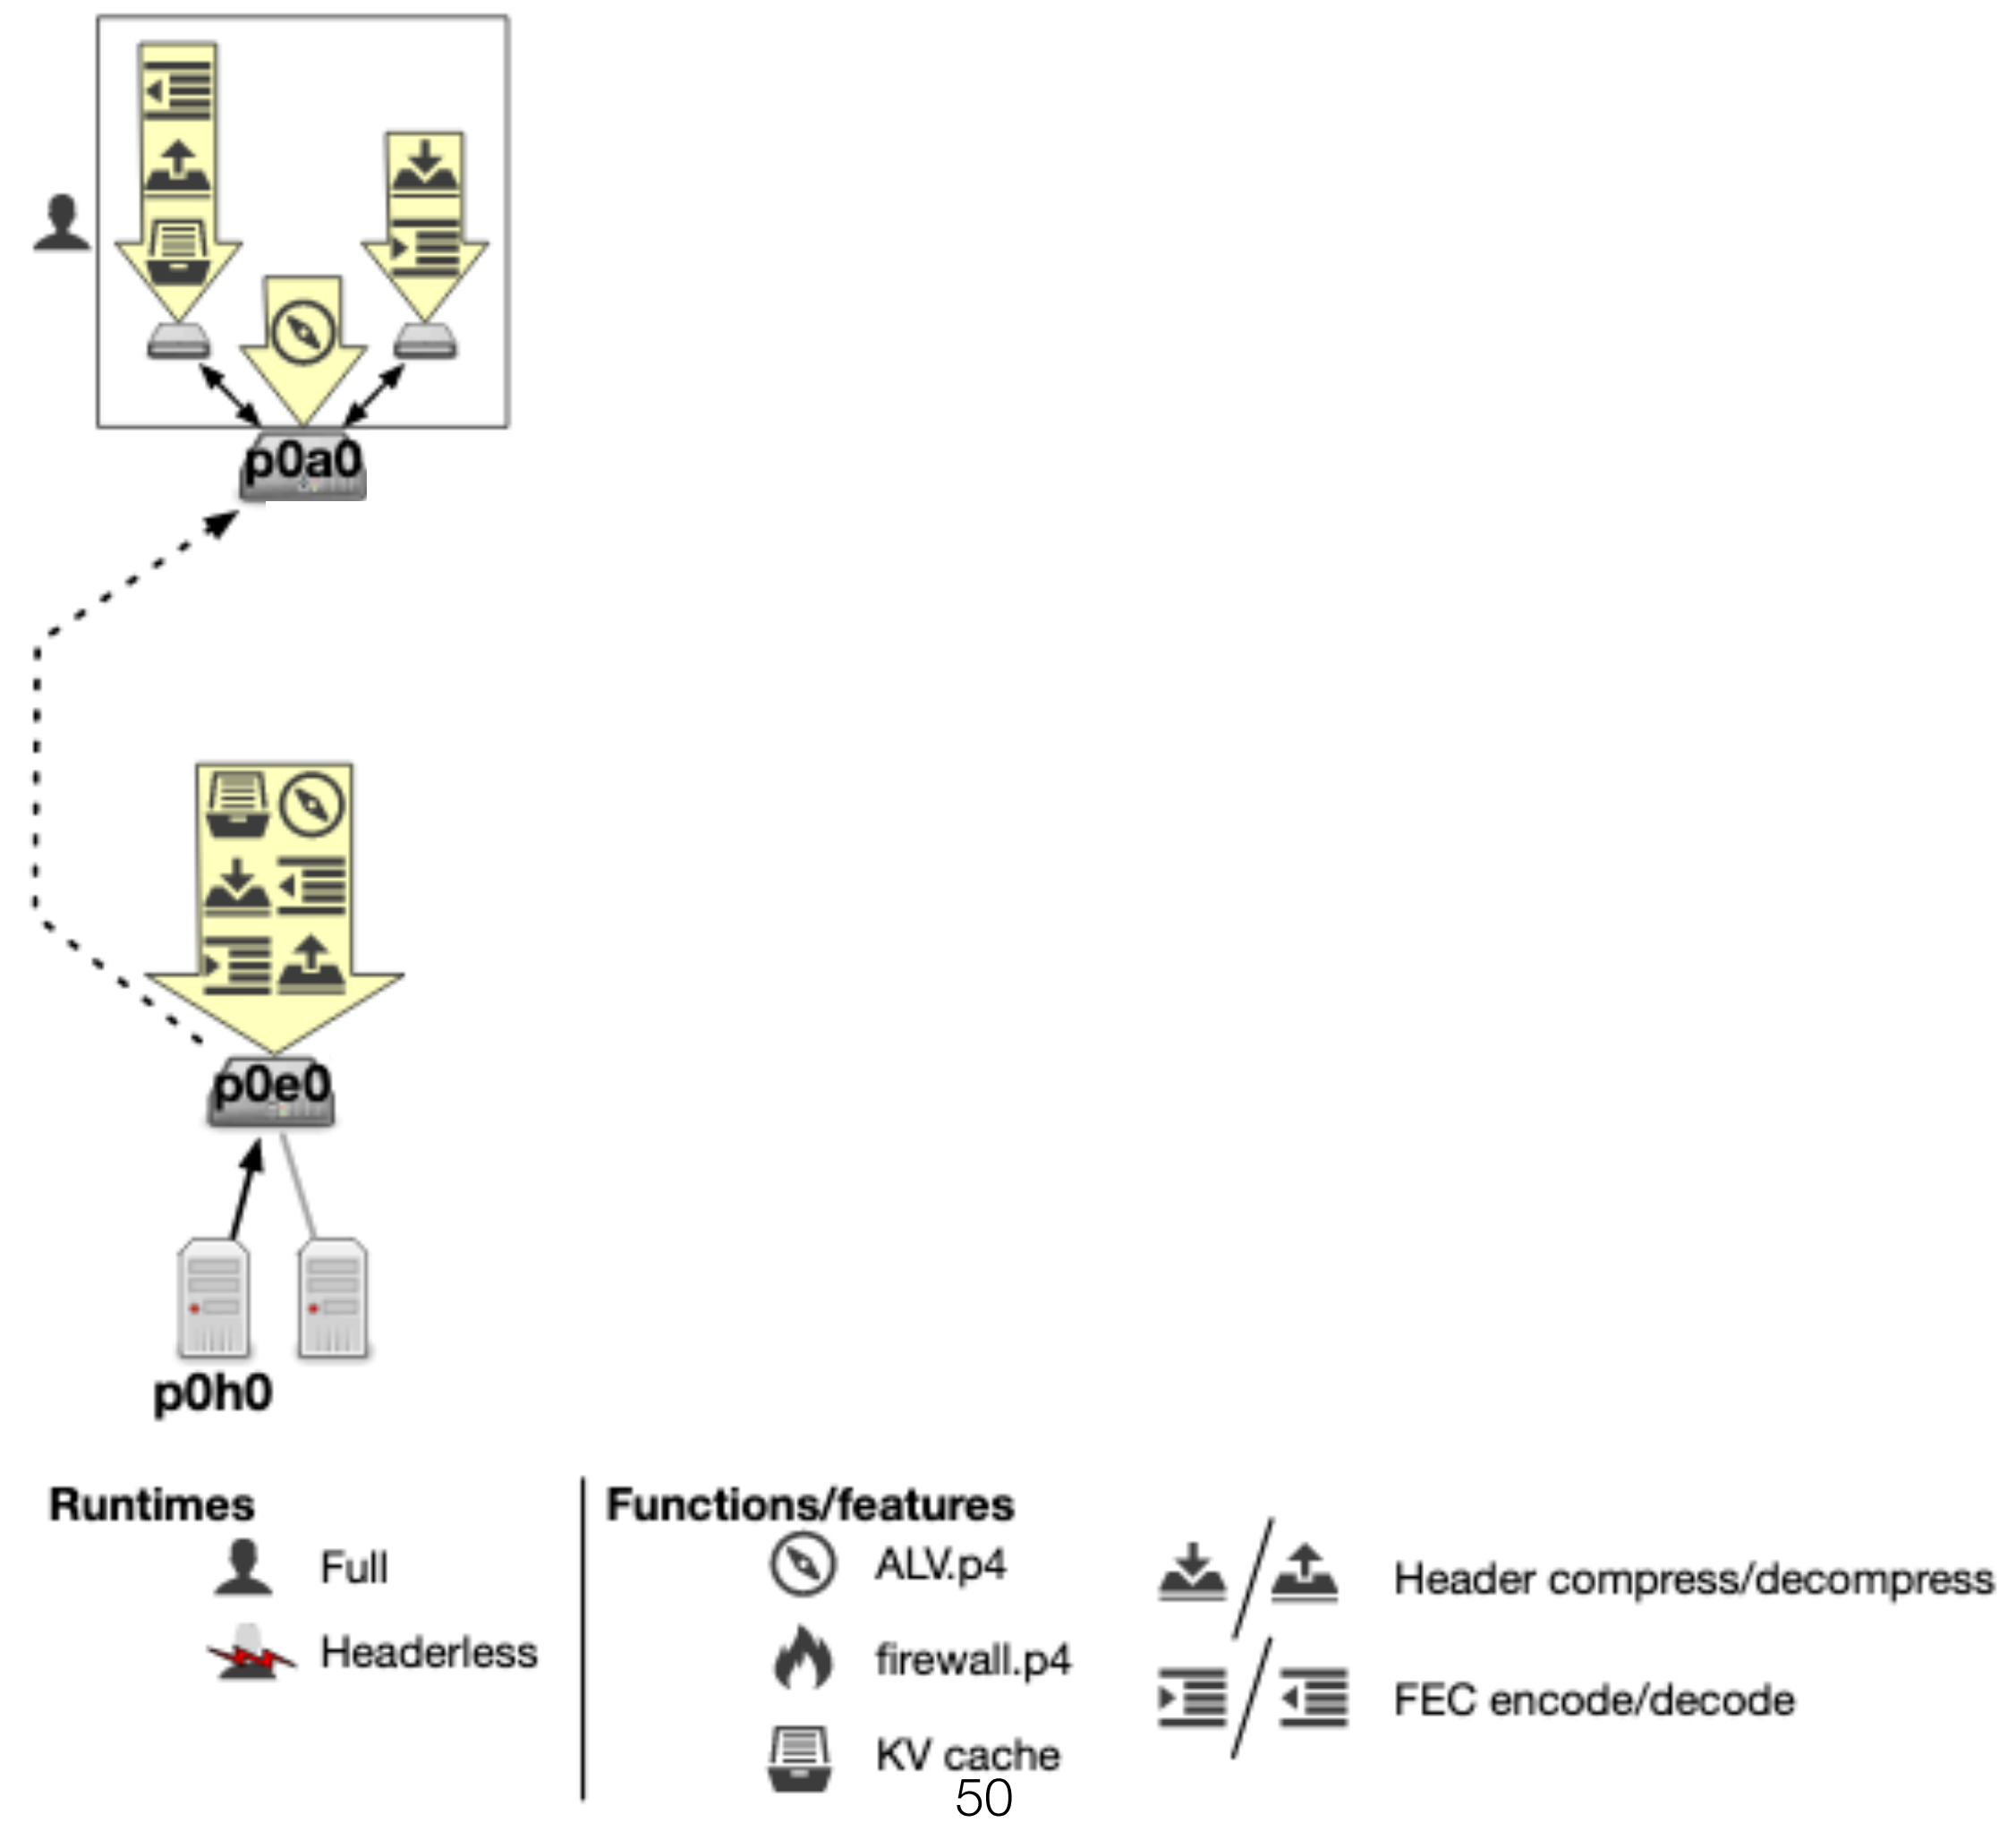




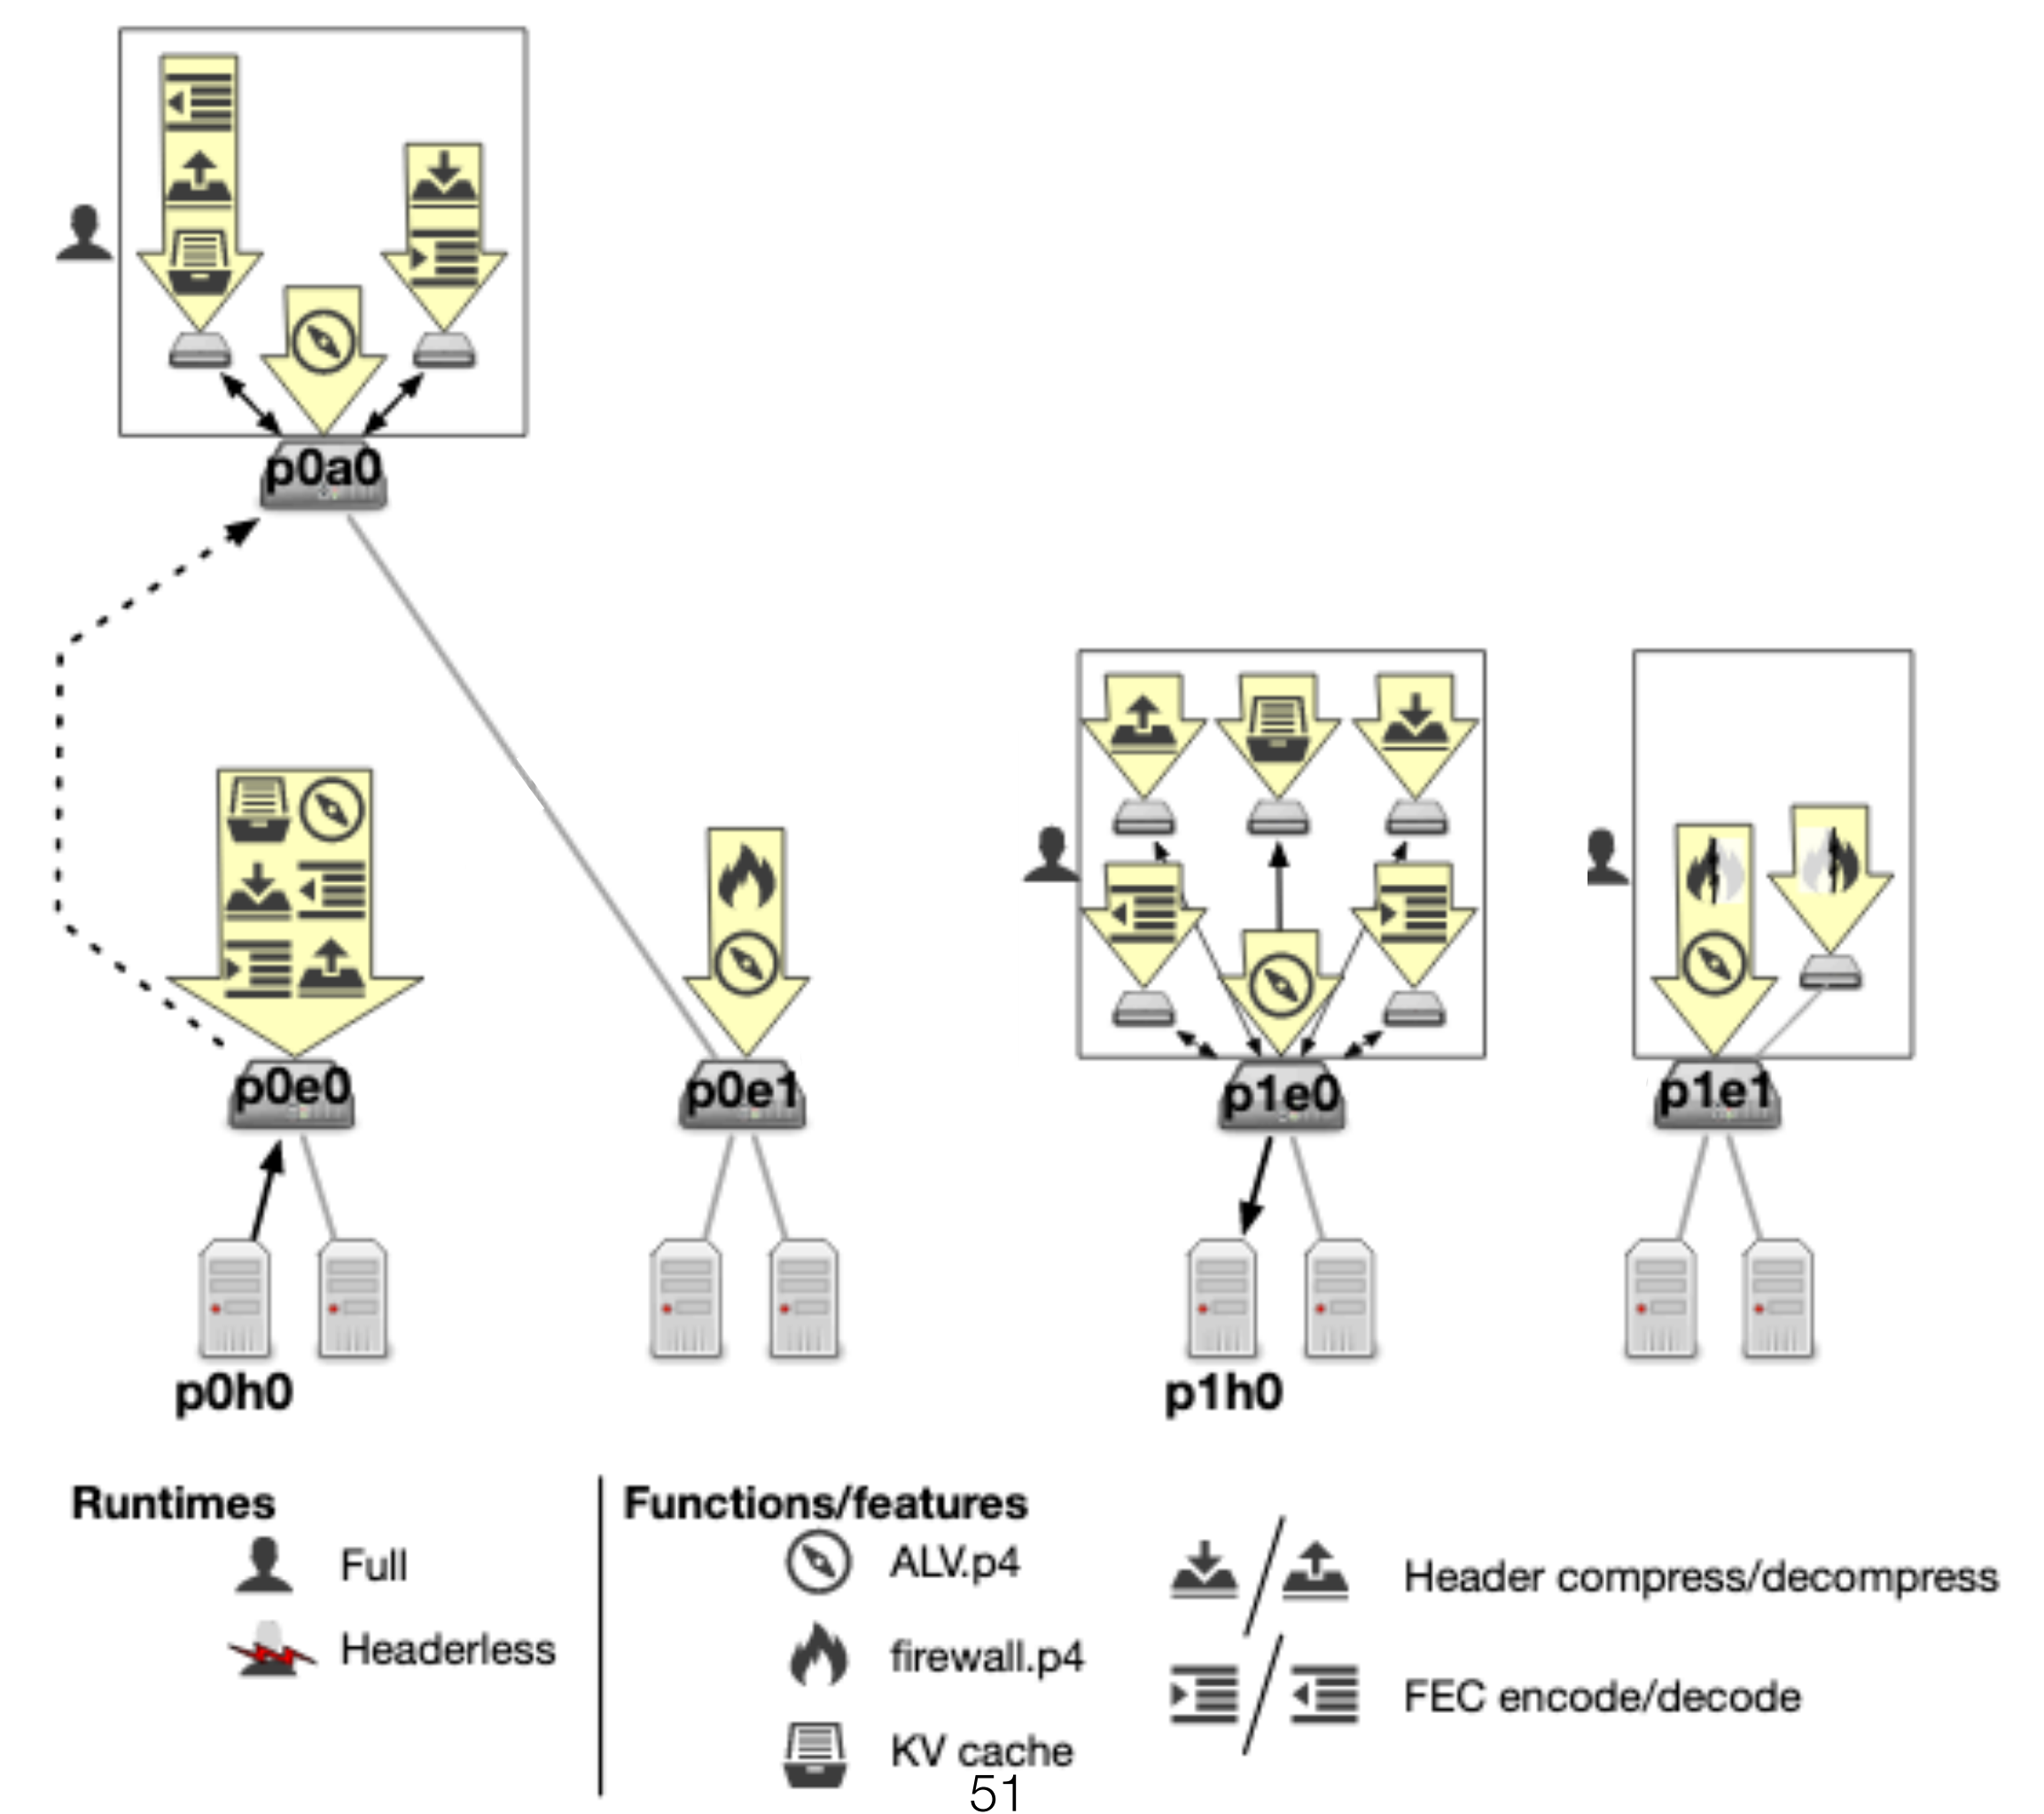
# Fig 7: Multiple Programs vs Runtimes vs Splits in same network (Simulation)




# Fig 7: Multiple Programs vs Runtimes vs Splits in same network (Simulation)



# Fig 7: Multiple Programs vs Runtimes vs Splits in same network (Simulation)





**Runtimes**

- Full
- Headerless

**Functions/features**

- ALV.p4
- firewall.p4
- KV cache
- Header compress/decompress
- FEC encode/decode

# FLIGHTPLAN DEMO

Choose an Experiment...

Start

About

**MSc students:** Heena Nagda (GATech), Rakesh Nagda (Penn)

**Other features:** graphs, multimedia cues (e.g., icons, packet structure), ...

<https://flightplan.cis.upenn.edu/>

Flightplan demo

# Flightplan: Dataplane Disaggregation and Placement for P4 programs

Nik Sultana   John Sonchack   Hans Giesen   Isaac Pedisich   Zhaoyang Han  
Nishanth Shyamkumar   Shivani Burad   André DeHon   Boon Thau Loo

*University of Pennsylvania*

<https://flightplan.cis.upenn.edu/>

NSDI'21