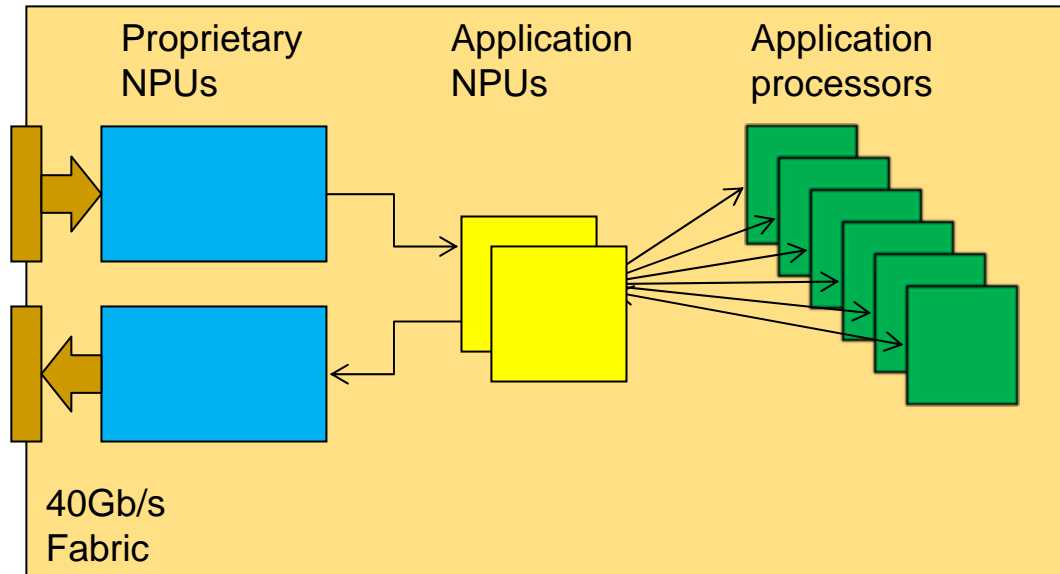


# Open routers?

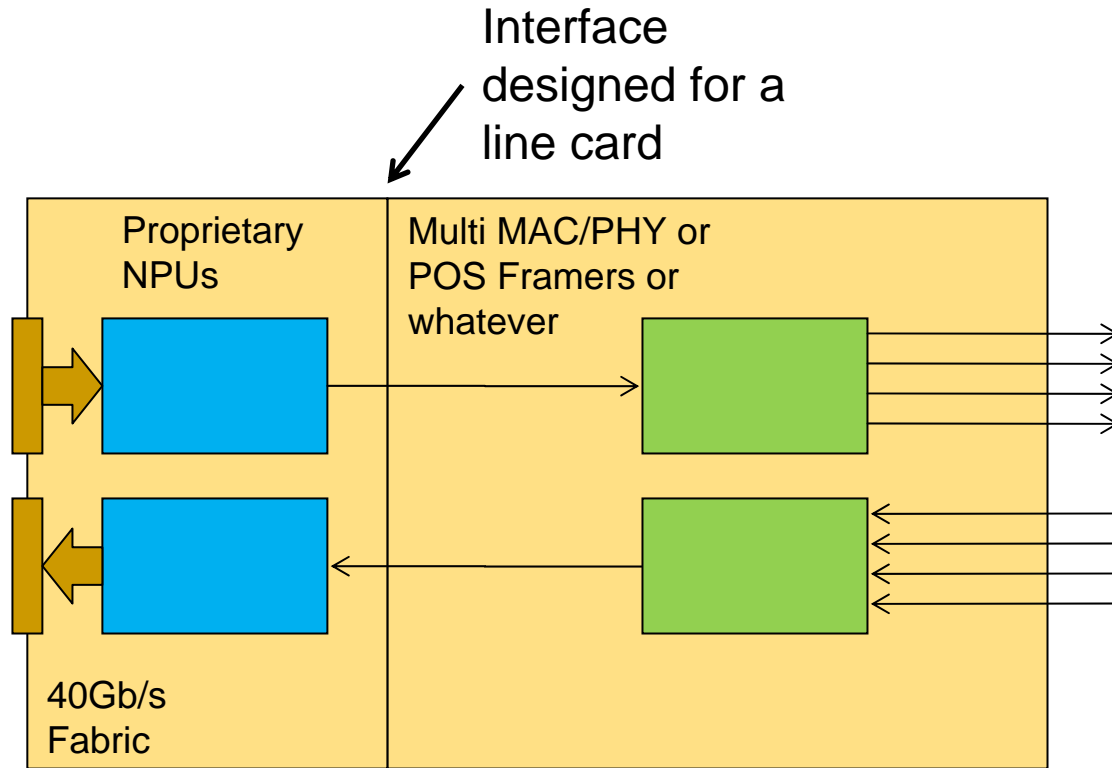
- A real nasty example from last Thursday
  - Why are there two 40Gbps NPUs and several application NPUs?
  - Why can I not run my code in those application processors?




# Why not?

- Why are there two 40Gbps NPUs and several application NPUs?
  - Product Group A own the fabric NPUs
  - Product Group B have to deliver an application blade
  - A does not trust B nor will they tie their release cycles together
  - If they don't trust each other they won't trust any customers
- Why can I not run my code in those application processors?
  - Application platform (OS/App NPU/driver) is only tested with equipment provider's applications
  - Customers will break the system without extensive platform testing ala Microsoft, RH, Intel, AMD, VMware etc

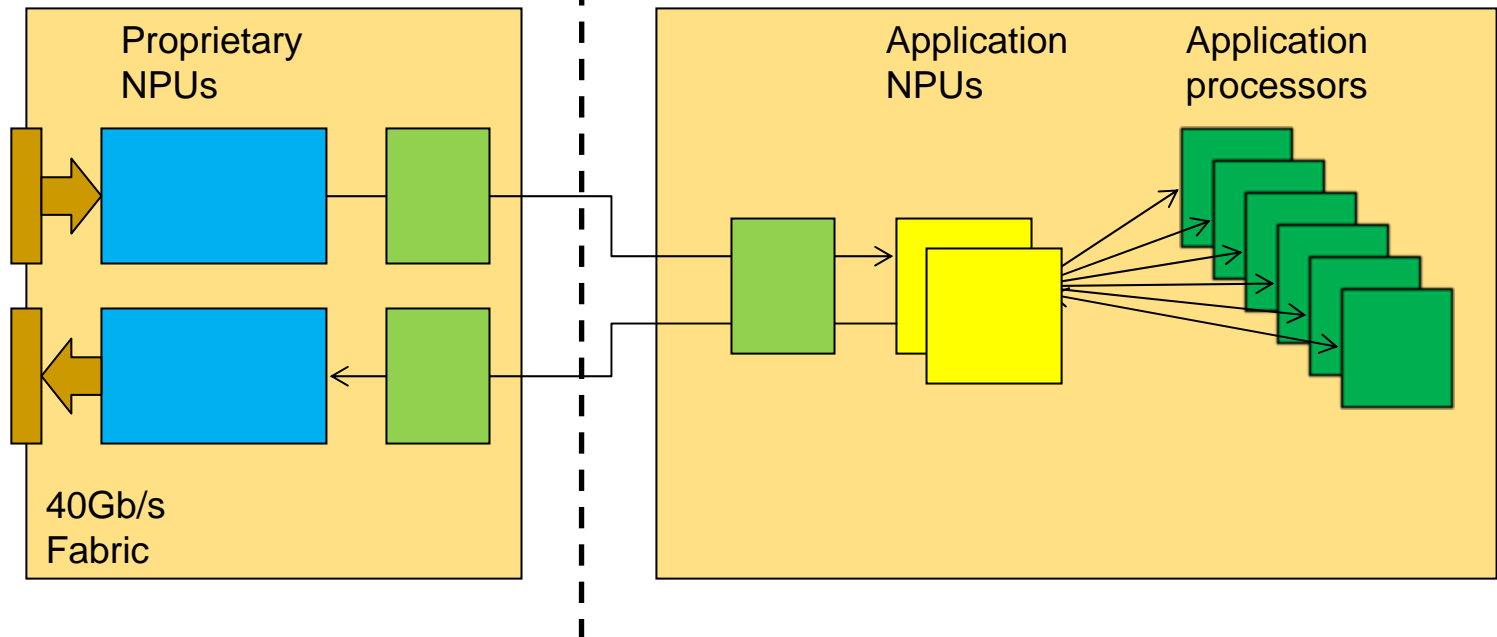
# But this is a “cut and shut”



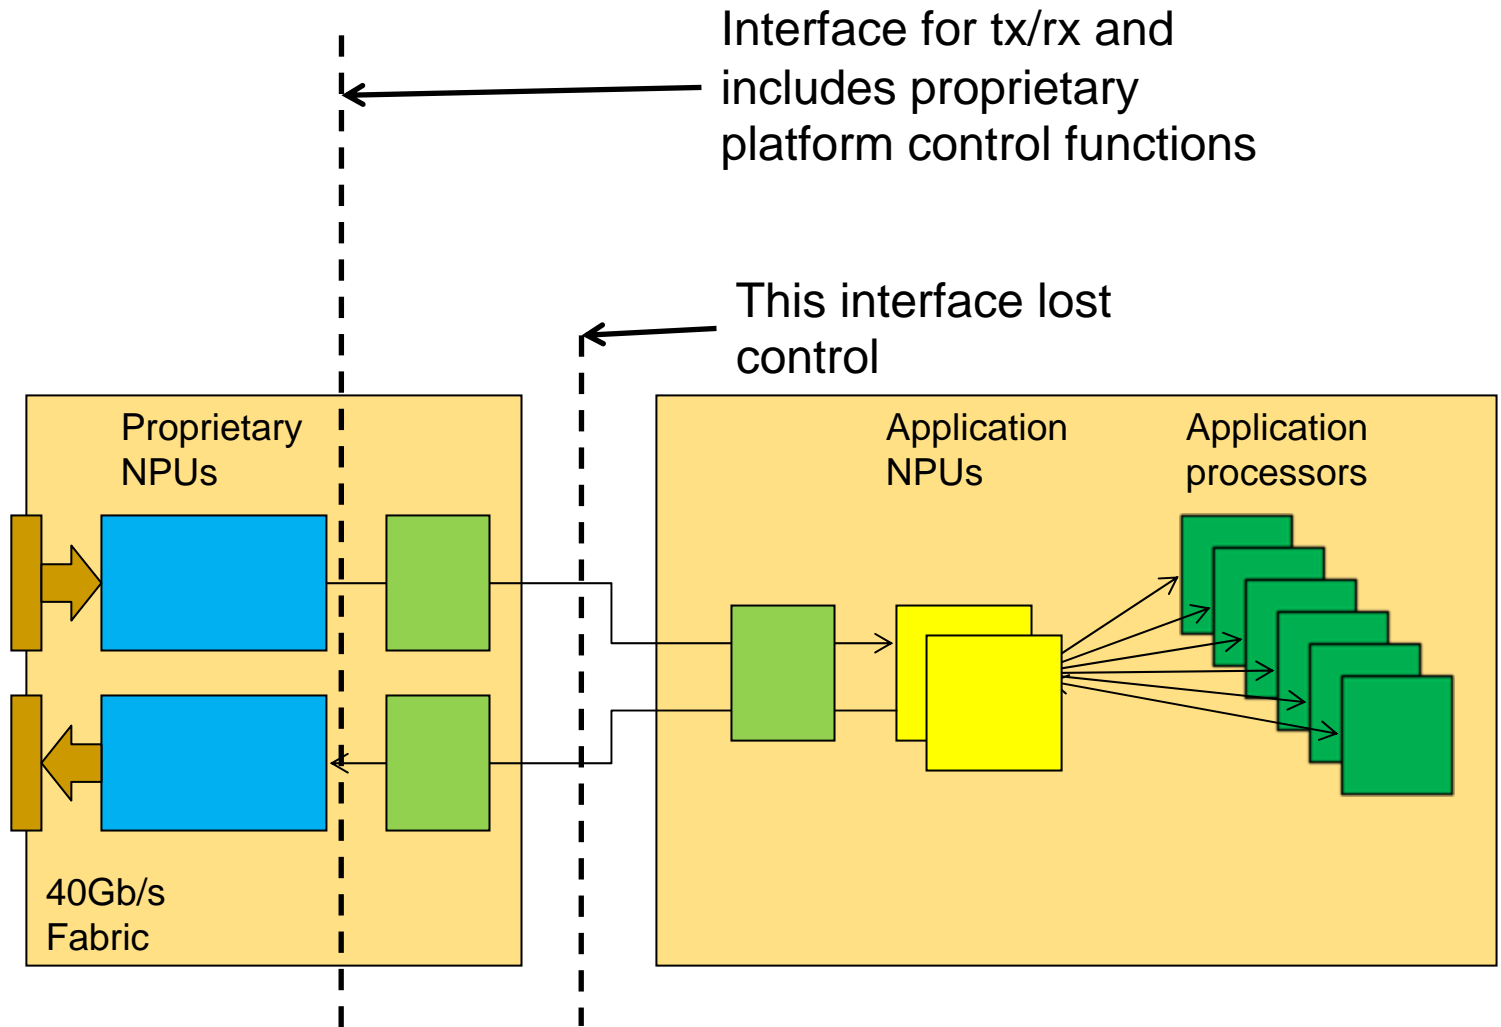
# So why not application on a stick?



Why is this not a std “open platform” which is independently implemented and tested?



# So where's the problem?



# What's the problem

- There's a proprietary control interface not exposed outside the box
- Simply put it is the interface between:
  - Routing and forwarding
  - Spanning tree and bridging
  - Flow decisions and flow routing
  - "Control plane and data plane"
- We've seen this before:
  - Failed start ups include: Ipsilon, Cplane
  - IXA; Not sufficiently generic? Not a protocol?

# A great man speaks

- Paraphrasing:

“We failed to define the interface between routing and forwarding in the Internet and we’ve been living with the consequences ever since.”

Dave Clarke

- Go fix it with a research programme, lets not sacrifice any more startups