Improving Data Center Resiliency and Availability through Path Load Balancing Strategies

Champagne Loïc — 2023
lchampagne@uliege.be

Context: DCN

Modern data center requirements

- **High-bandwidth** connectivity to support the transfer of large amounts of data between servers and storage systems.
- **Low-latency** networking to ensure fast response times for real-time applications like video conferencing and online gaming.
- **High-capacity servers** to handle millions of requests per second from users accessing websites, applications, or other online services.
- **Robust software security measures** to protect against cyber attacks and unauthorized access to data, including regular software updates and patches, multi-factor authentication, firewalls, intrusion detection and prevention systems, and data encryption

Context: DCN (2)

Modern data center requirements

- **Redundant** power supplies, cooling systems, and networking equipment to ensure high availability in case of hardware failures or other disruptions.
- **Scalable infrastructure** that can be easily expanded or upgraded as demand for computing resources grows over time.
- **Efficient use of resources** to minimize power consumption and reduce the carbon footprint of the data center.

Problematic

- Data centers are a target of choice for a variety of malicious actors, including hackers, cybercriminals, and state-sponsored attackers.
- Malicious attacks on data centers can result in a range of negative consequences, including the shutdown of important business functions or services, and may cause substantial financial harm.
- Due to the crucial role they play in supporting the operations of multiple companies, data centers represent a high-value target.

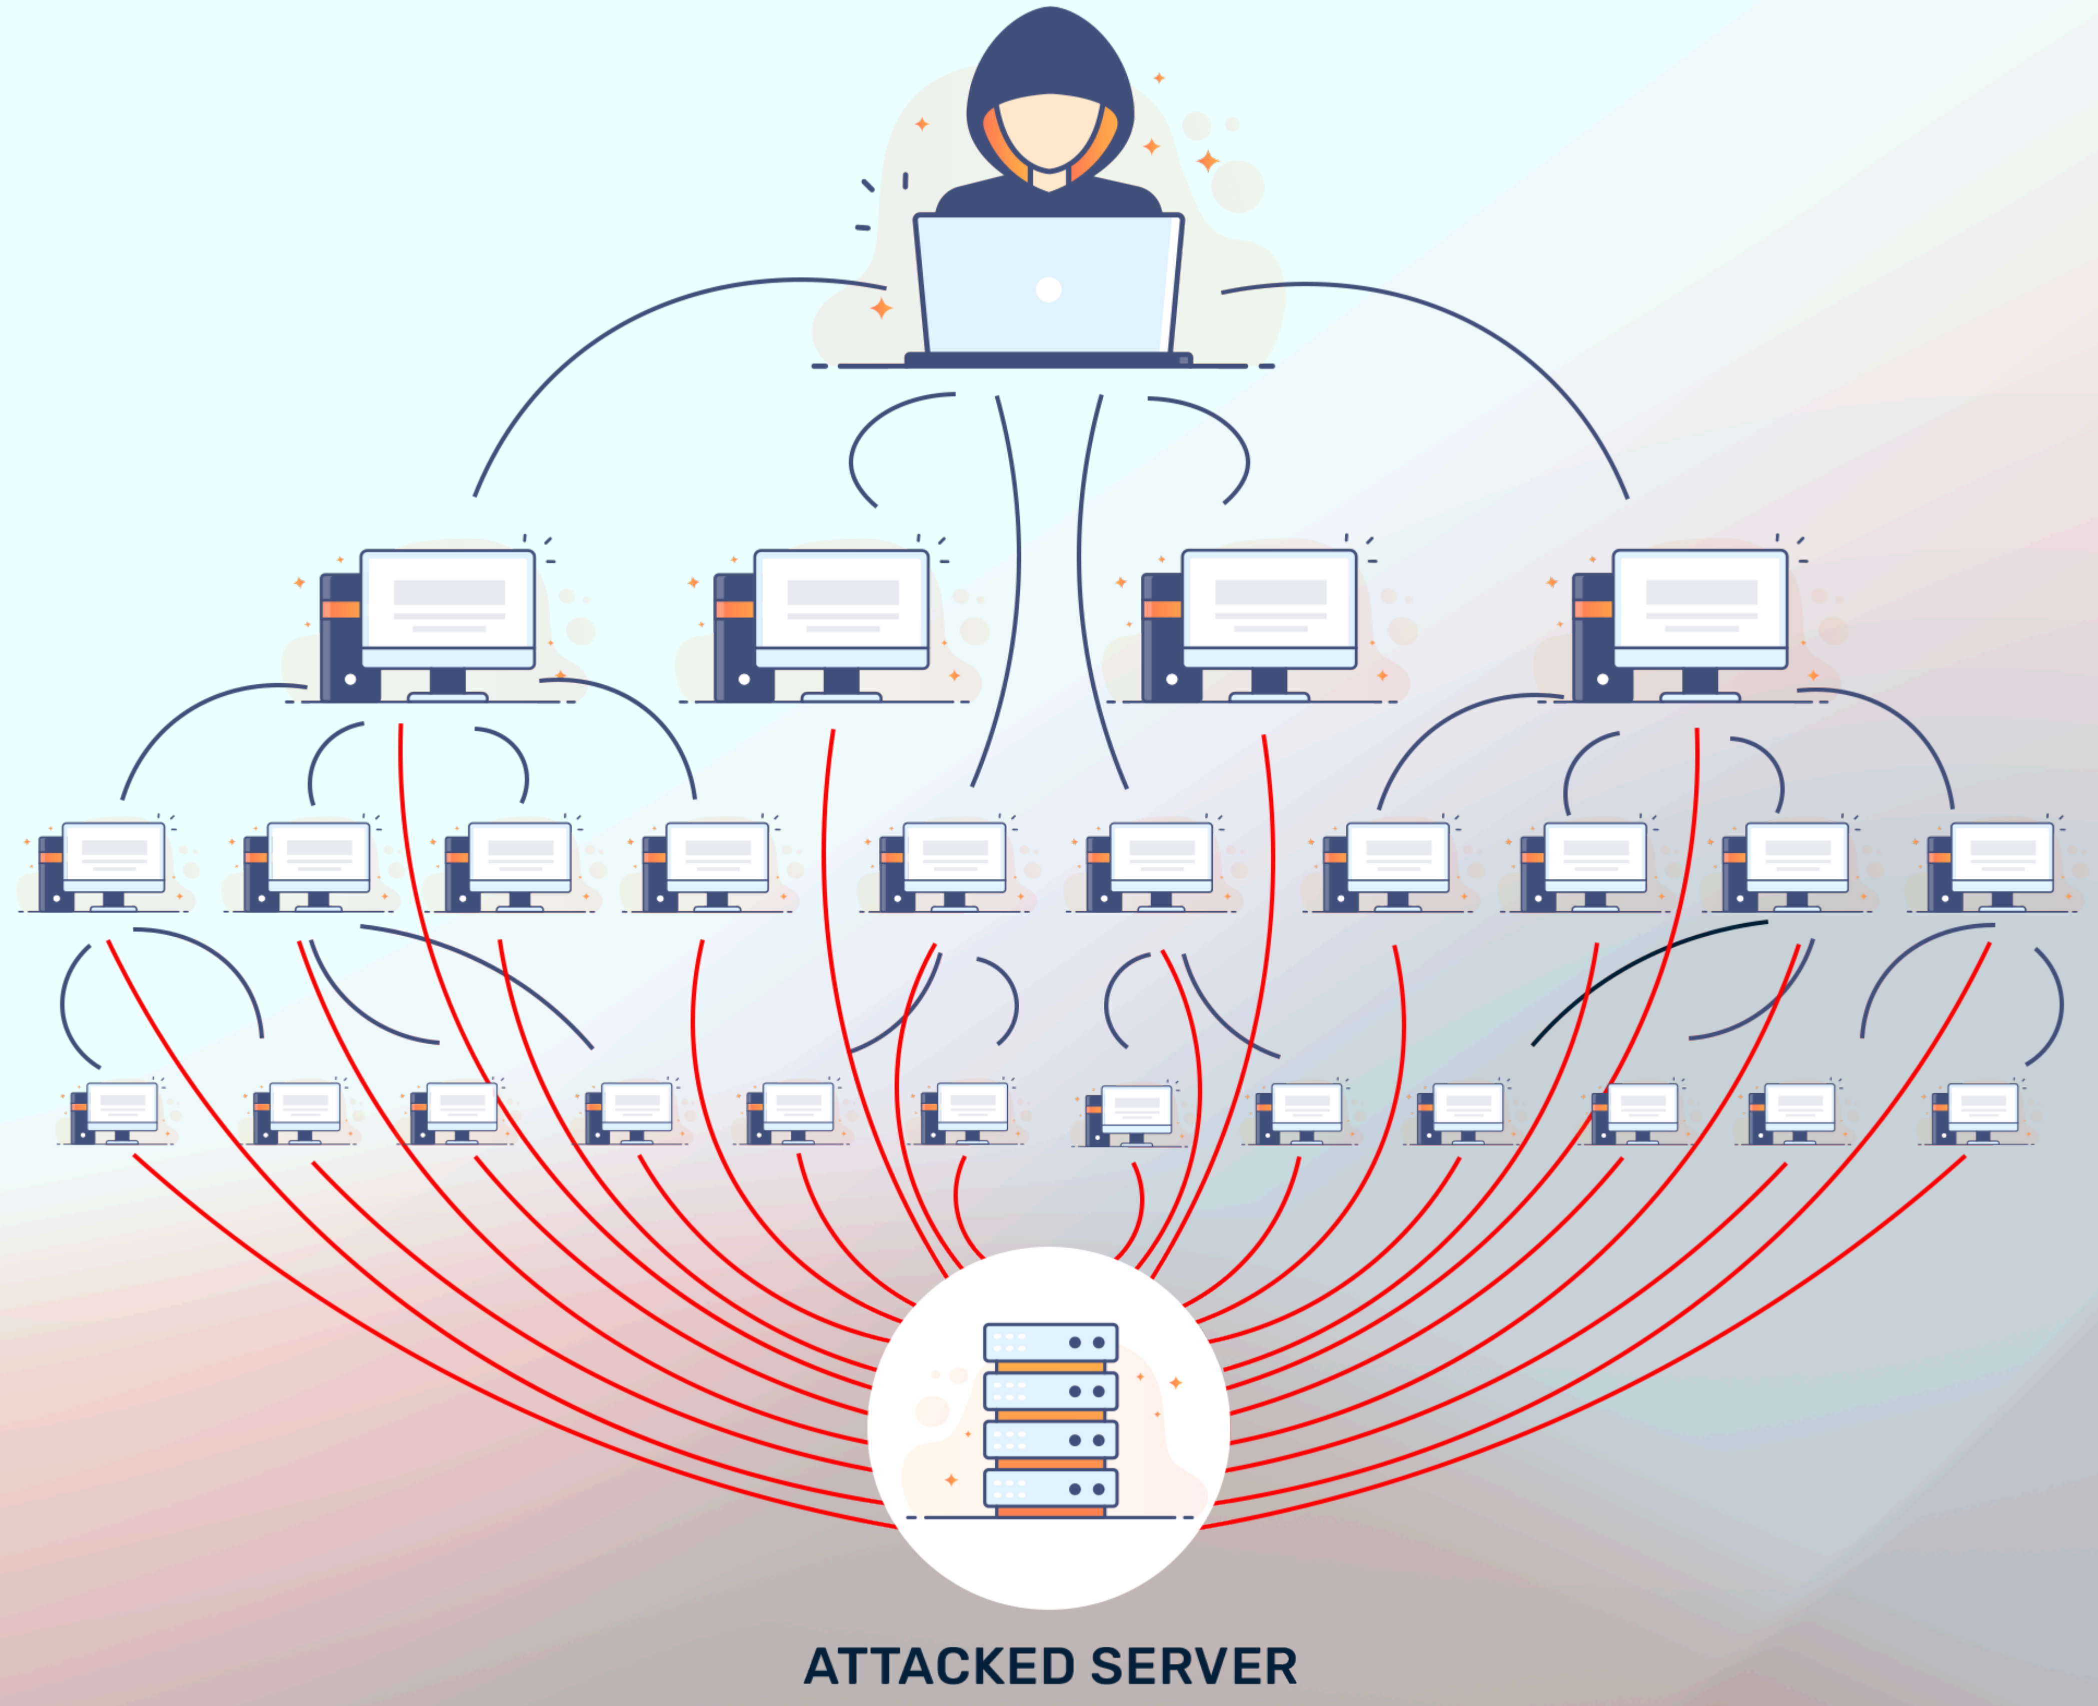

Coping Mechanisms and Countermeasures

Modern Data Center architectures

A first step towards resilient data centers

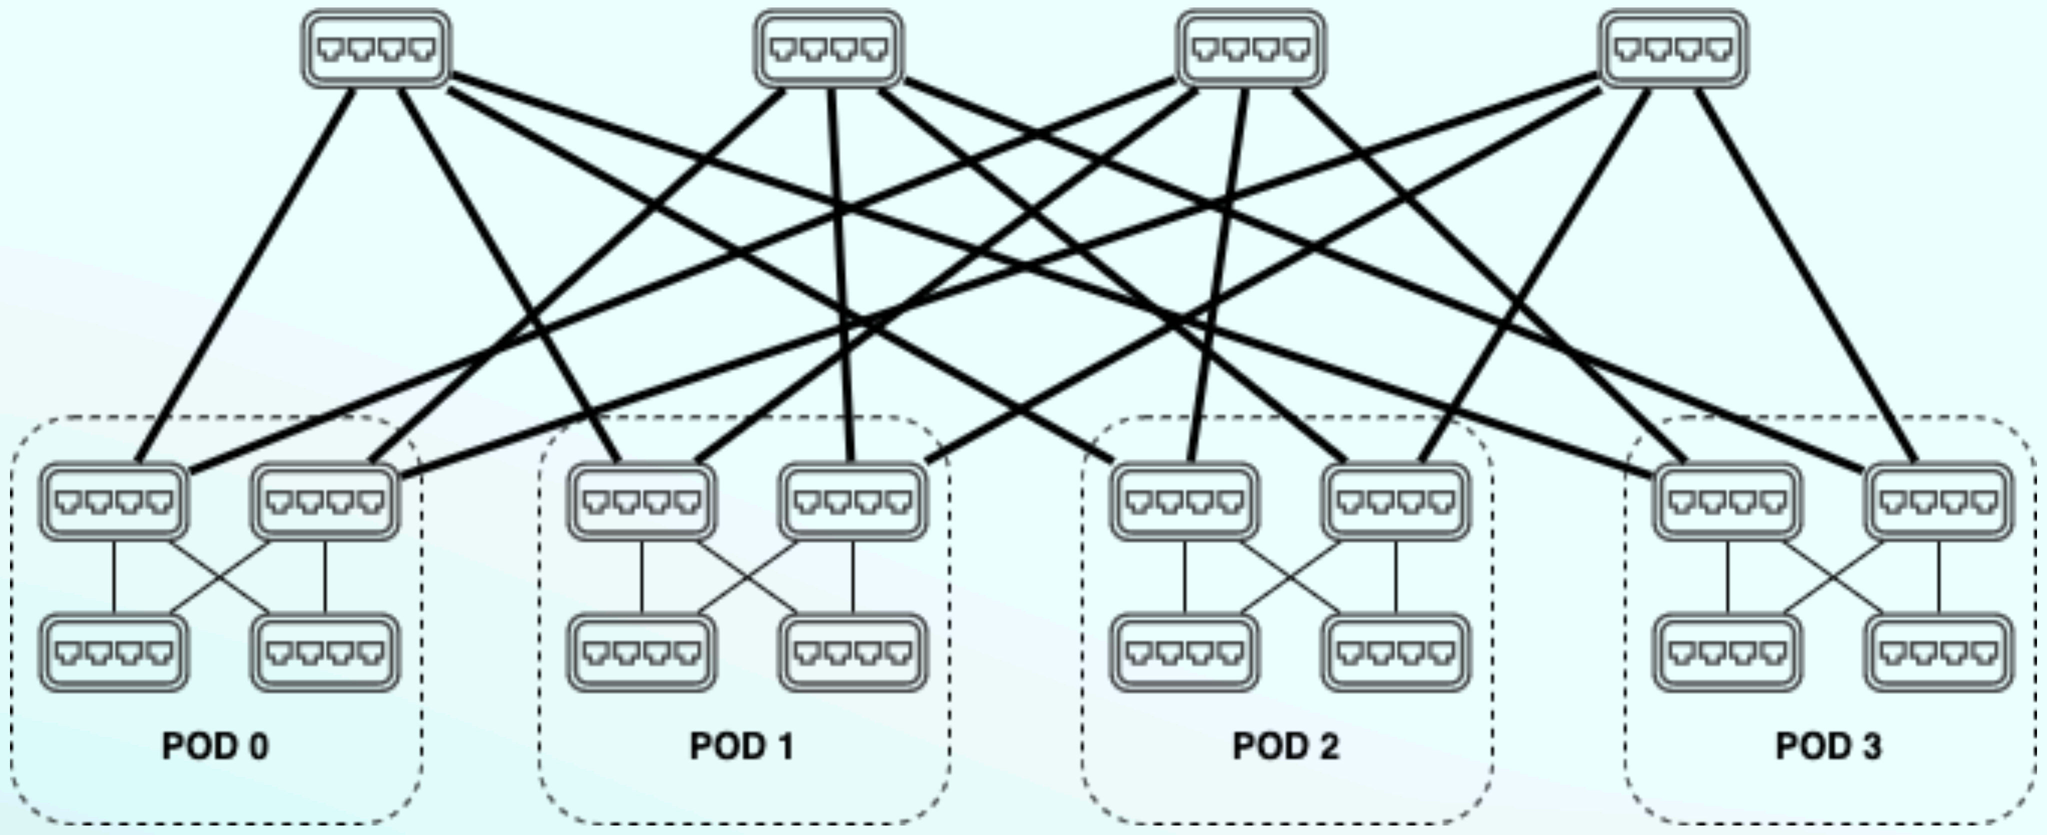

K-pod Fat Tree

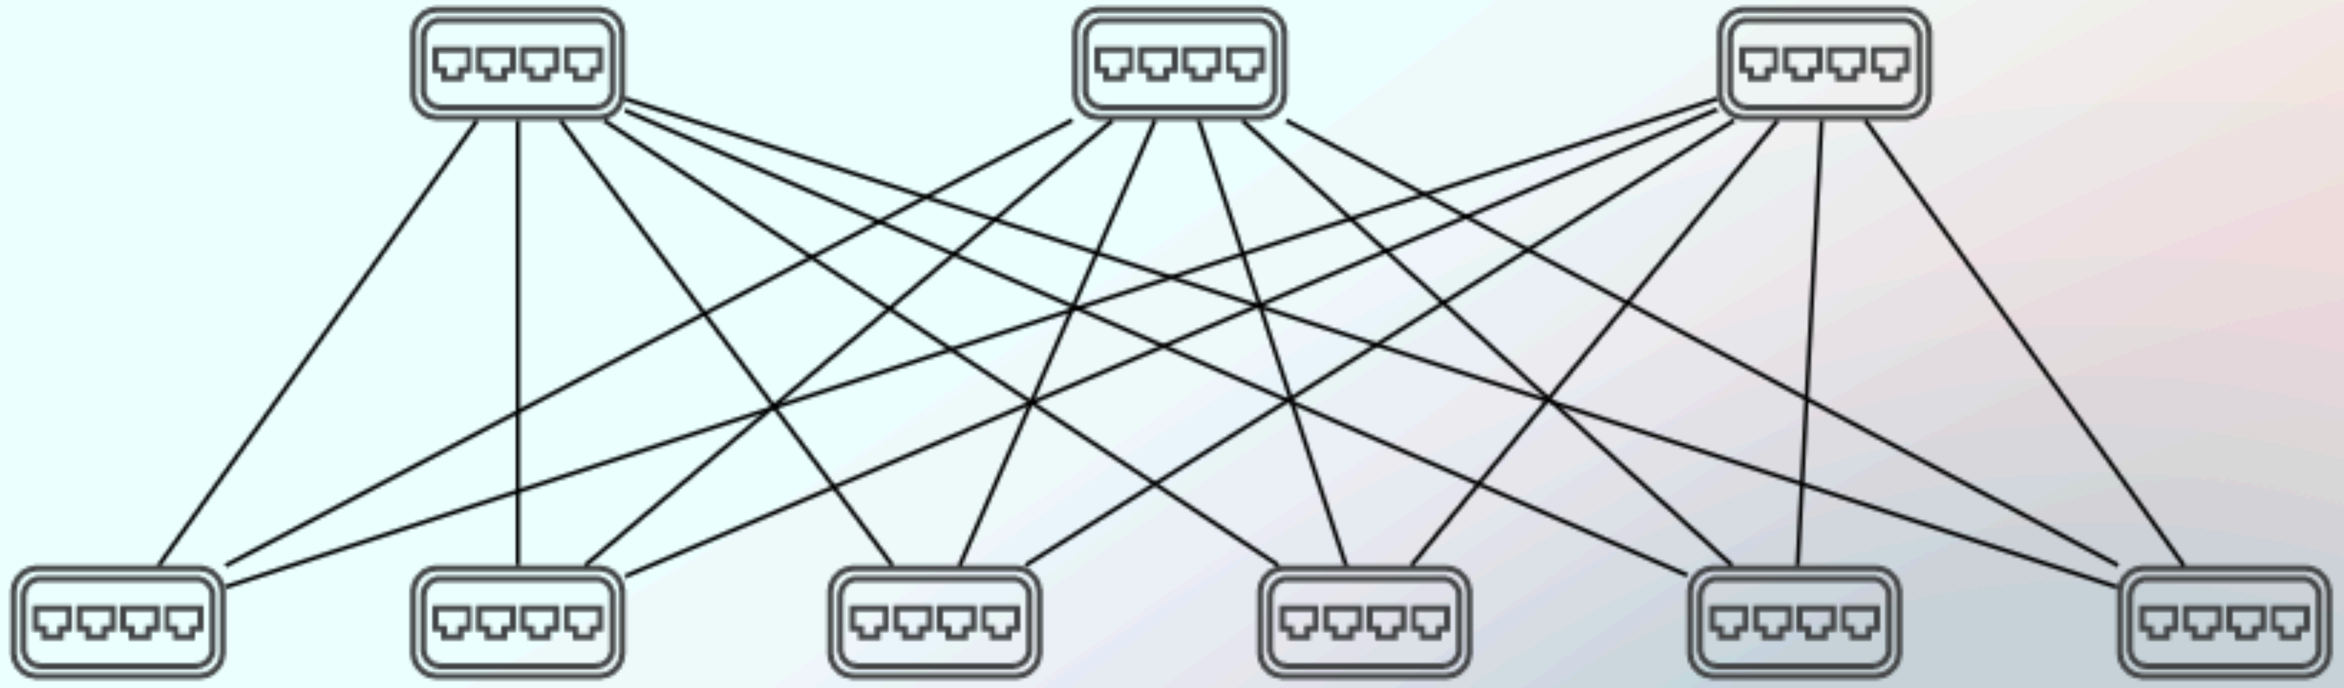

Spine Leaf

Both architectures exhibit a high level of bisection bandwidth to satisfy the necessary throughput demands, and they also incorporate alternative paths to cope with potential failures.

Data Center Observability

Obtaining comprehensive insights into the system's behavior, performance, and potential issues.

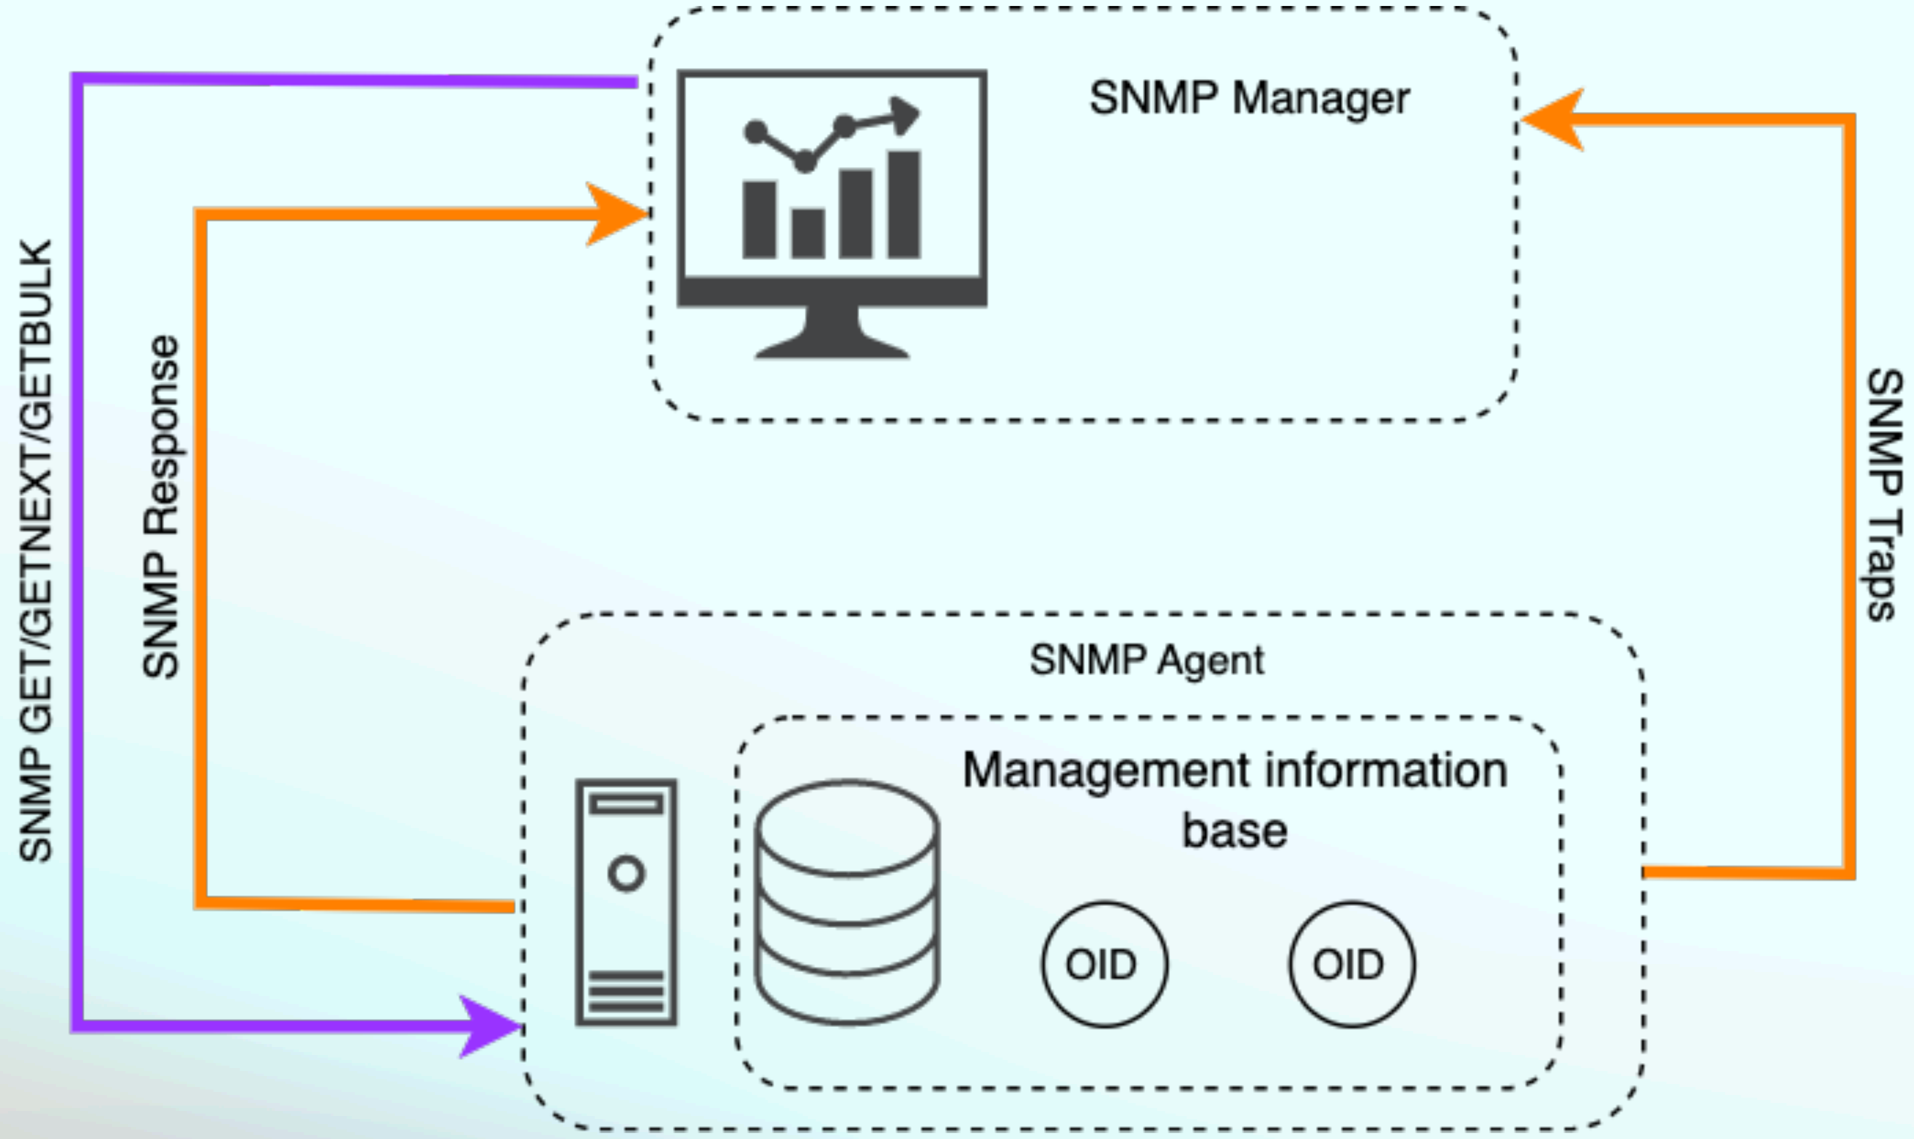

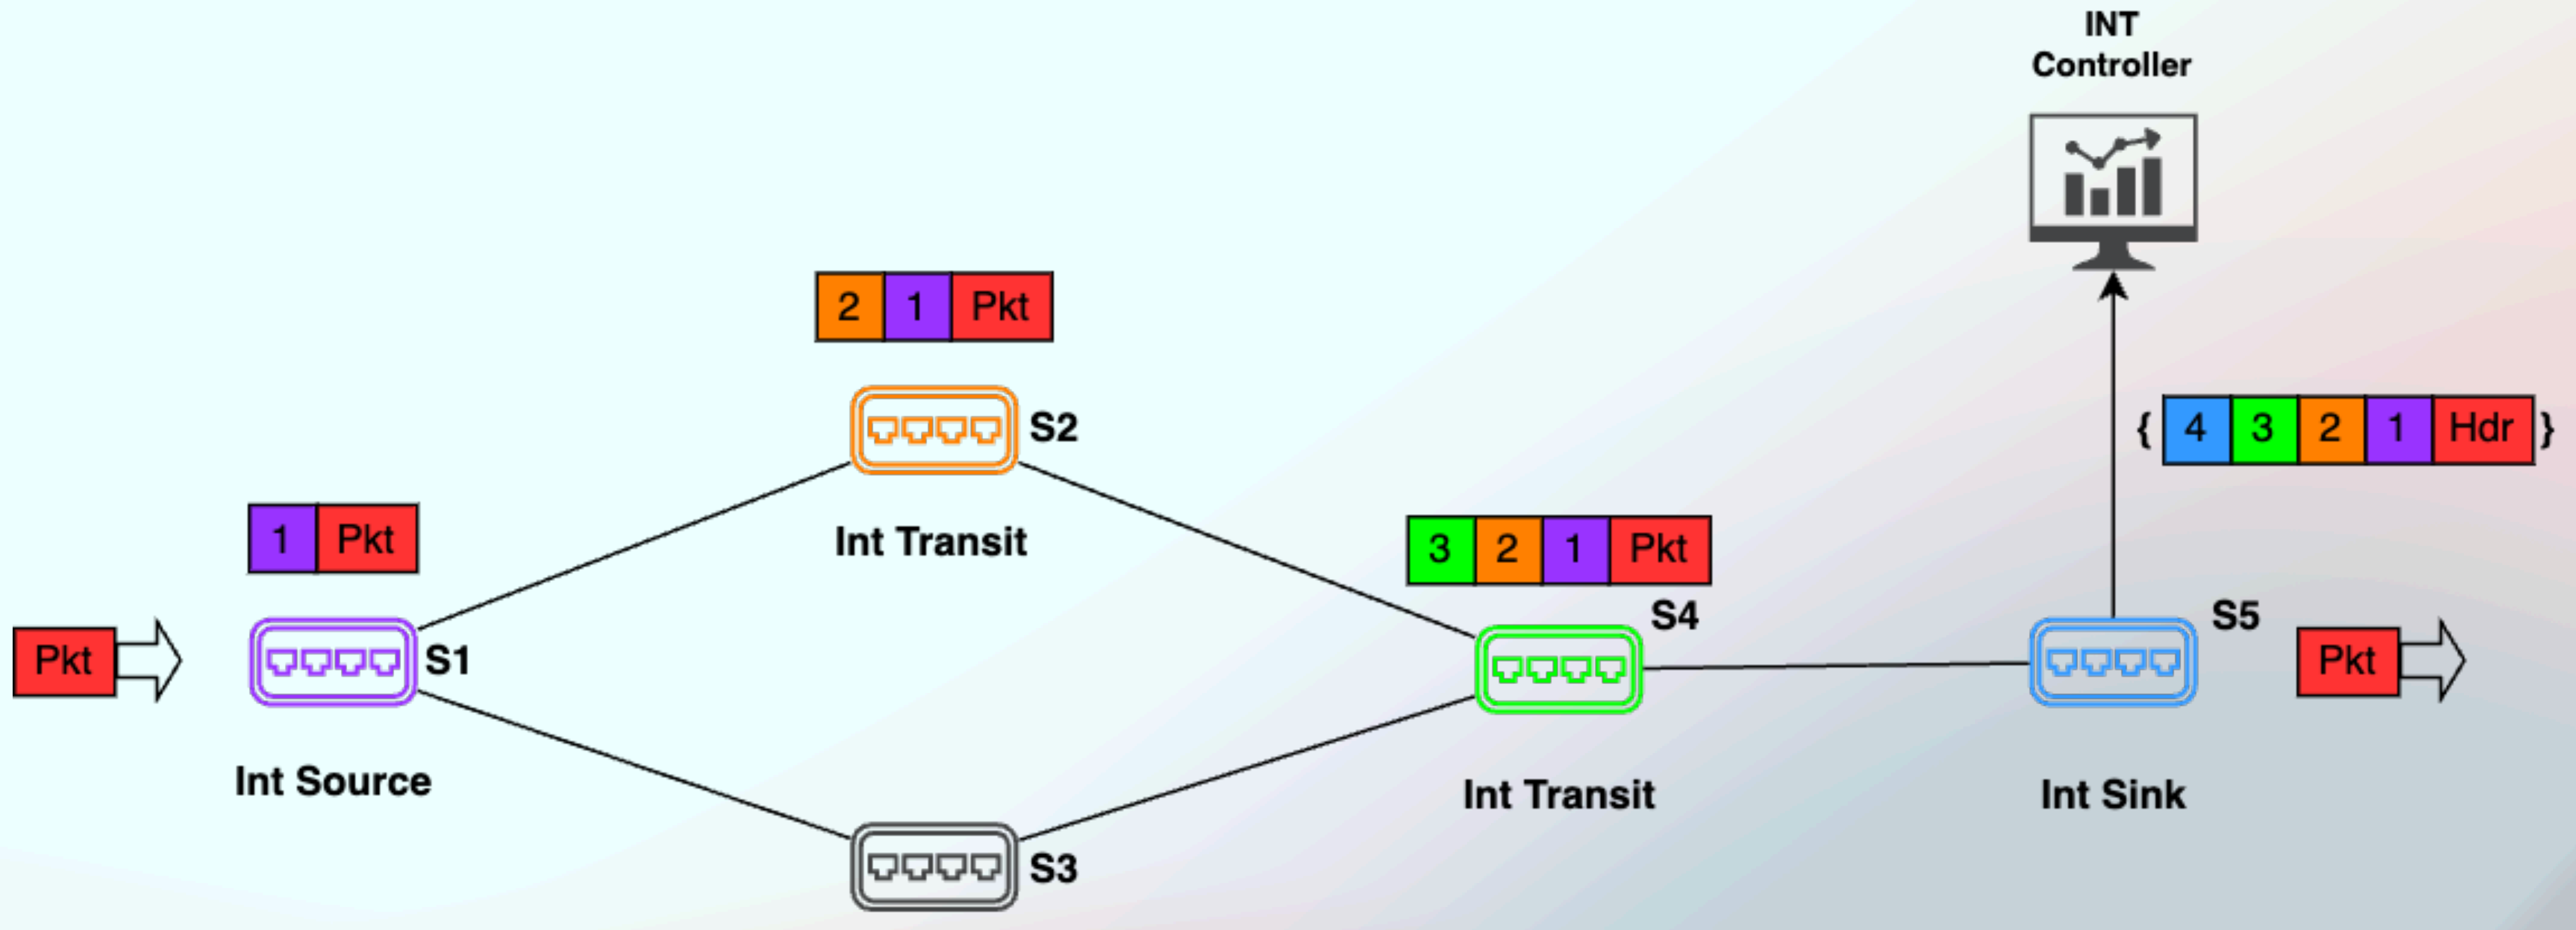

Simple Network Management Protocol based monitoring

Telemetry status rate: 1 - 5 min(s)

Inband Network Telemetry[1]

Telemetry status rate: 10 - 100 ms

1. https://p4.org/p4-spec/docs/INT_v2_1.pdf

Data Center Observability (2)

Inband Network Telemetry

- Gather Telemetry and OAM (Operations, Administration, and Maintenance) information along the path within the data packet, as part of an existing/additional header
  - No extra probe-traffic (as with ping, trace, ...)

- Header location:
  1. INT over IPv4/GRE
  2. INT over TCP/UDP
  3. INT over VXLAN
  4. Many others

- INT Data: Node id, Ingress interface identifier, Ingress timestamp, Egress interface identifier, Egress timestamp, Hop latency, Egress interface TX Link utilization, Queue occupancy

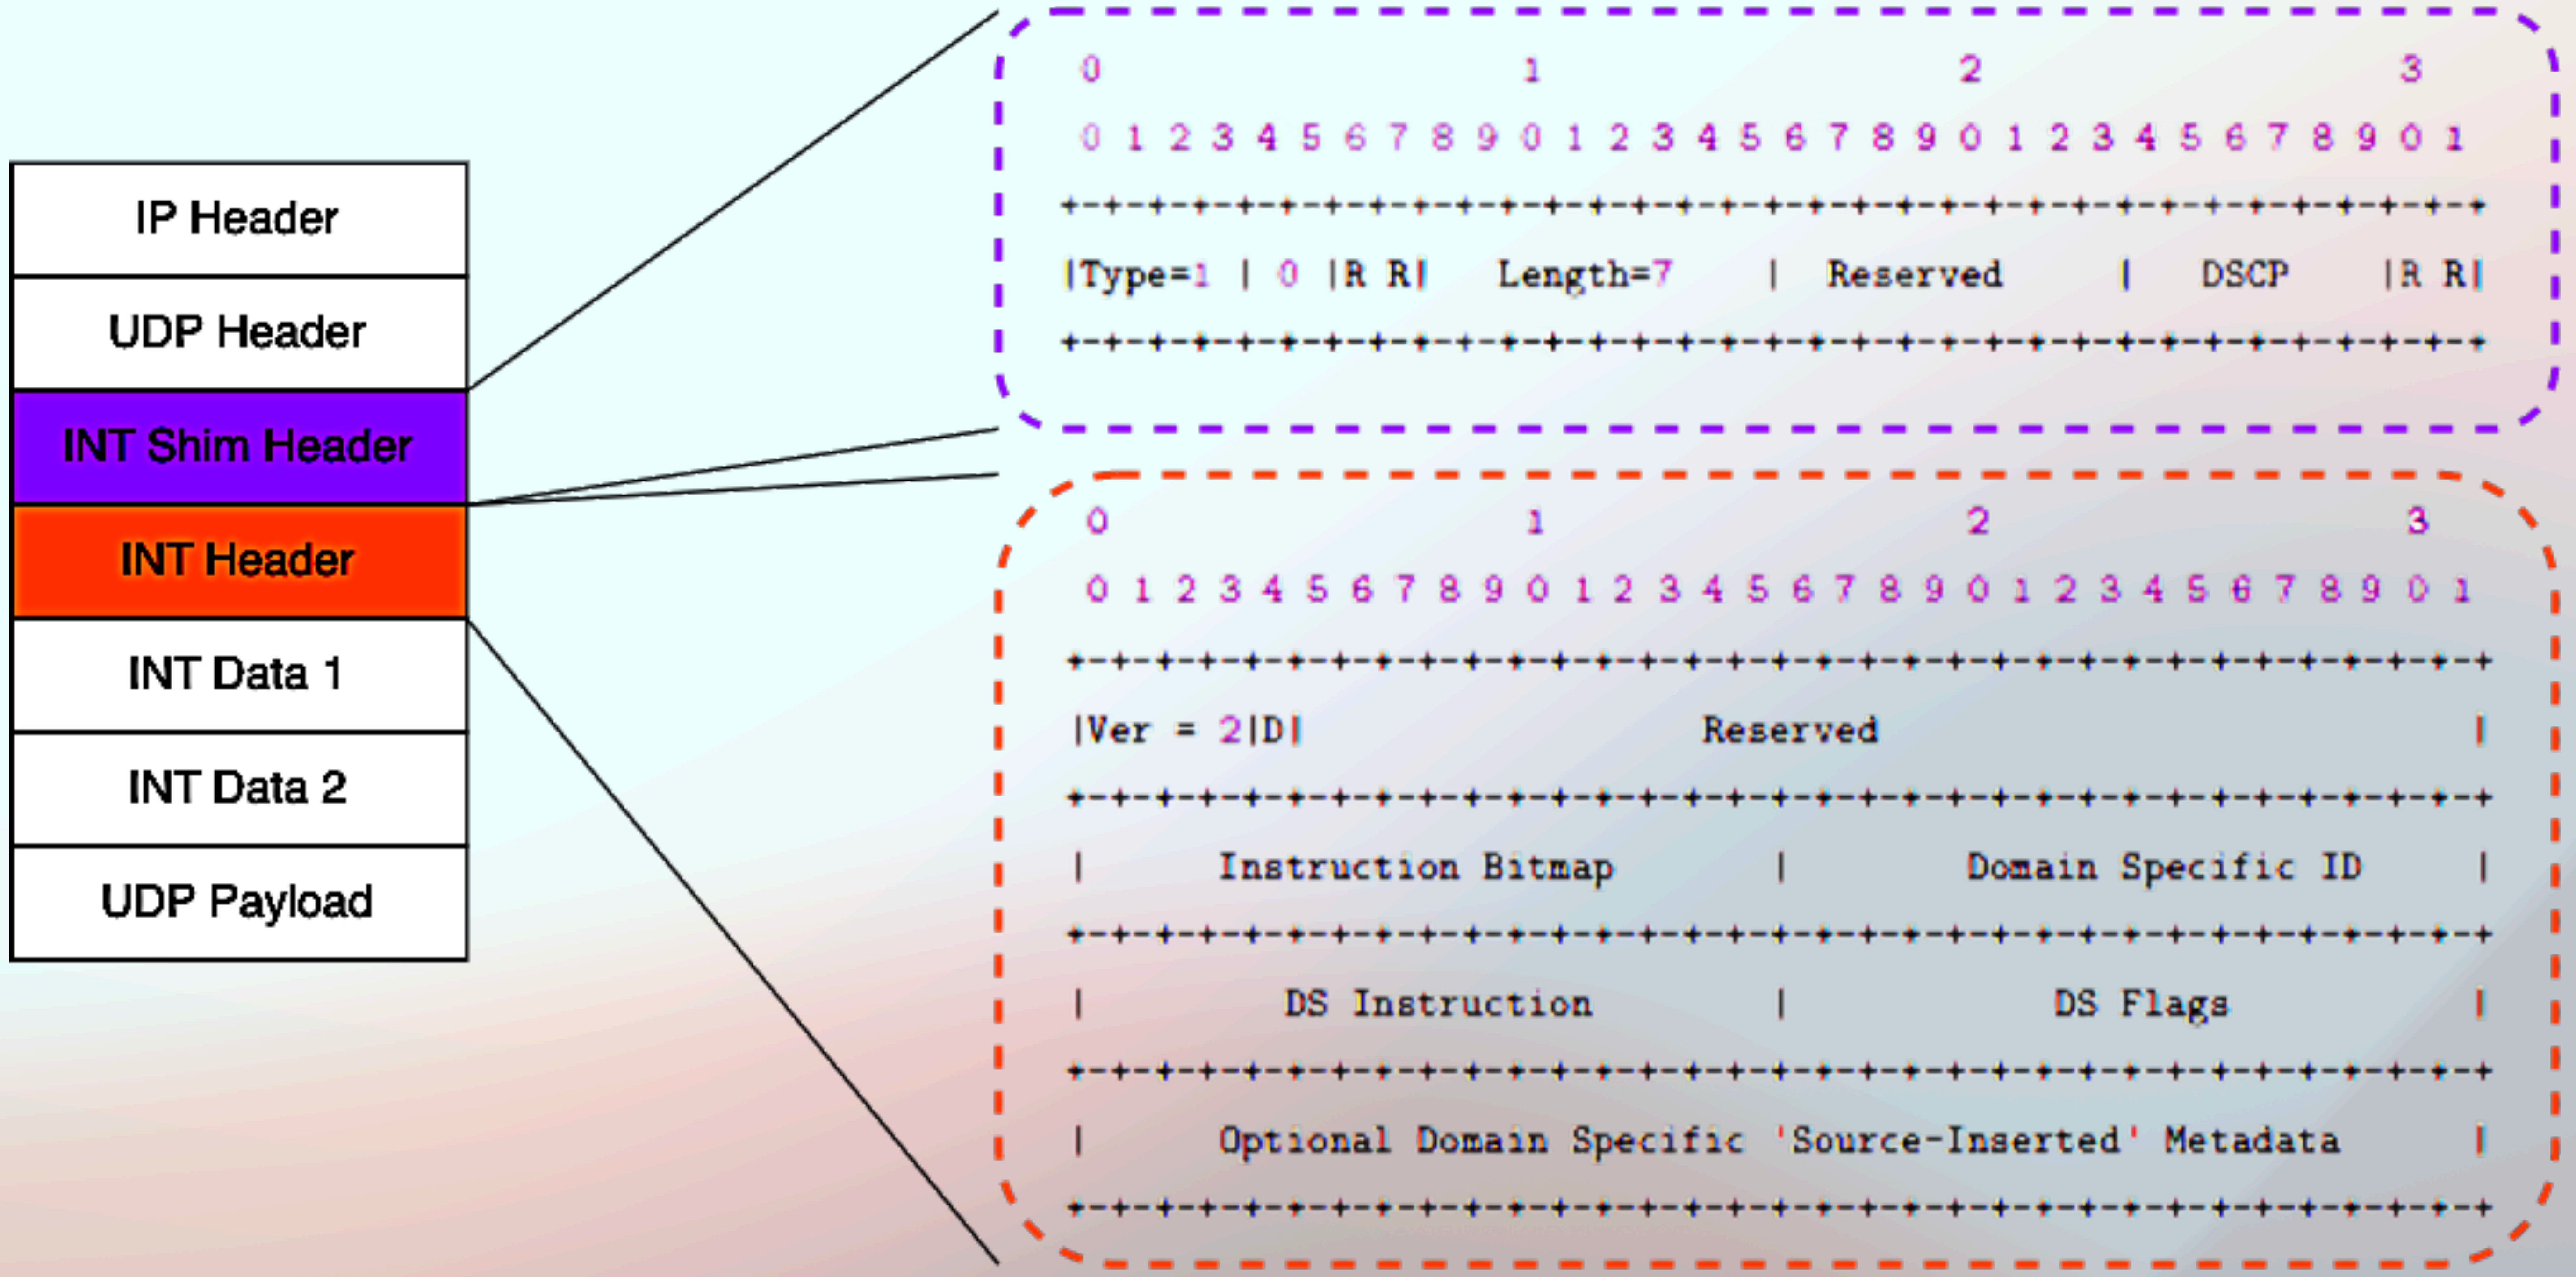

Data Center Observability (3)

Smoothie

Weight matrix					
	S1	S2	S3	S4	S5
S1	0	0	1/64	0	0
S2	4/64/1/3	0	0	0	0
S3	4/64	0	0	0	0
S4	0	0	1/64	0	0
S5	0	0	0	0	0

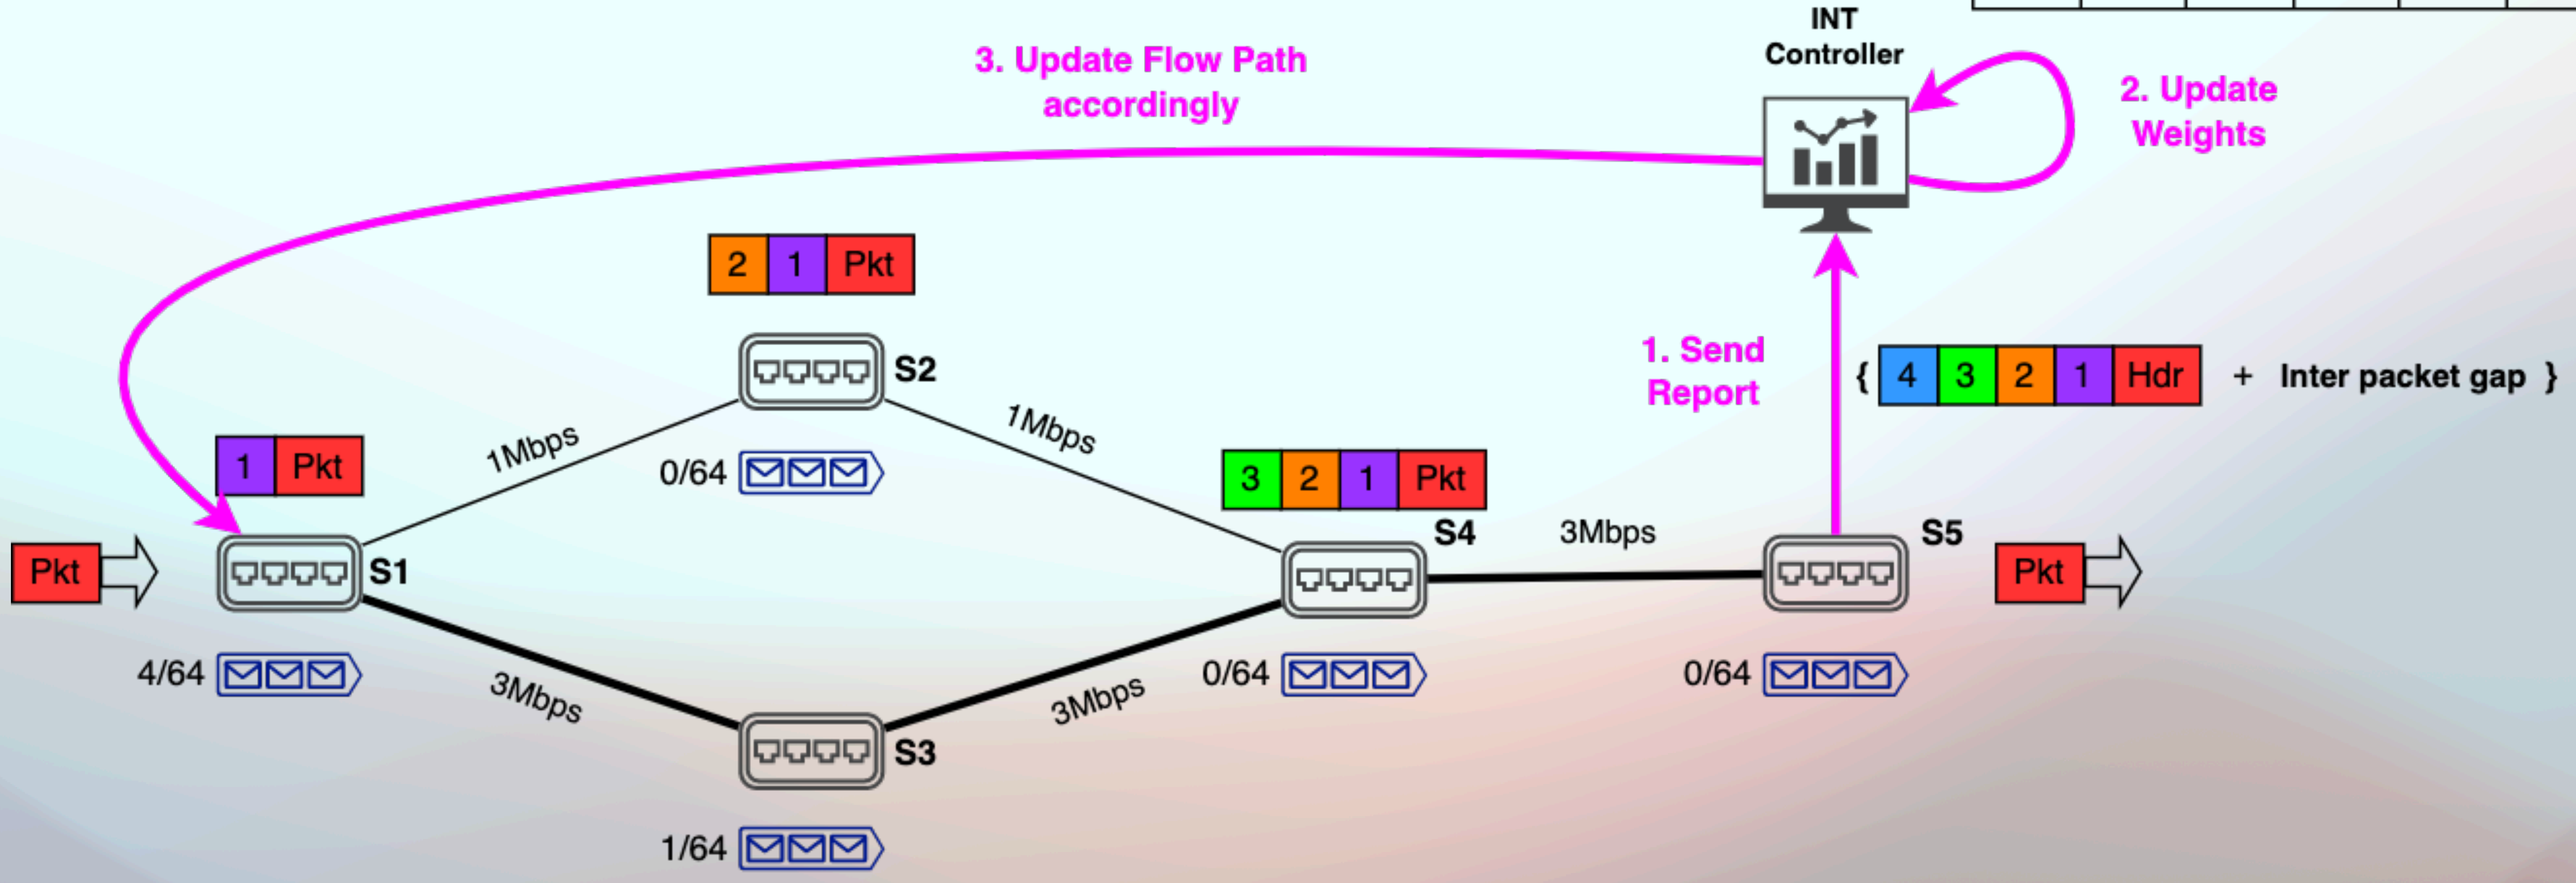

Flow to path dynamic allocation with Smoothie

Data Center Observability (4)

Scenario Timeline:

1. Initially, Flow 1 and Flow 2 are assigned to the same path.
2. Path start to become congested.
3. Smoothie notices the sustained congestion.
4. Flow 2 is routed to a new path.

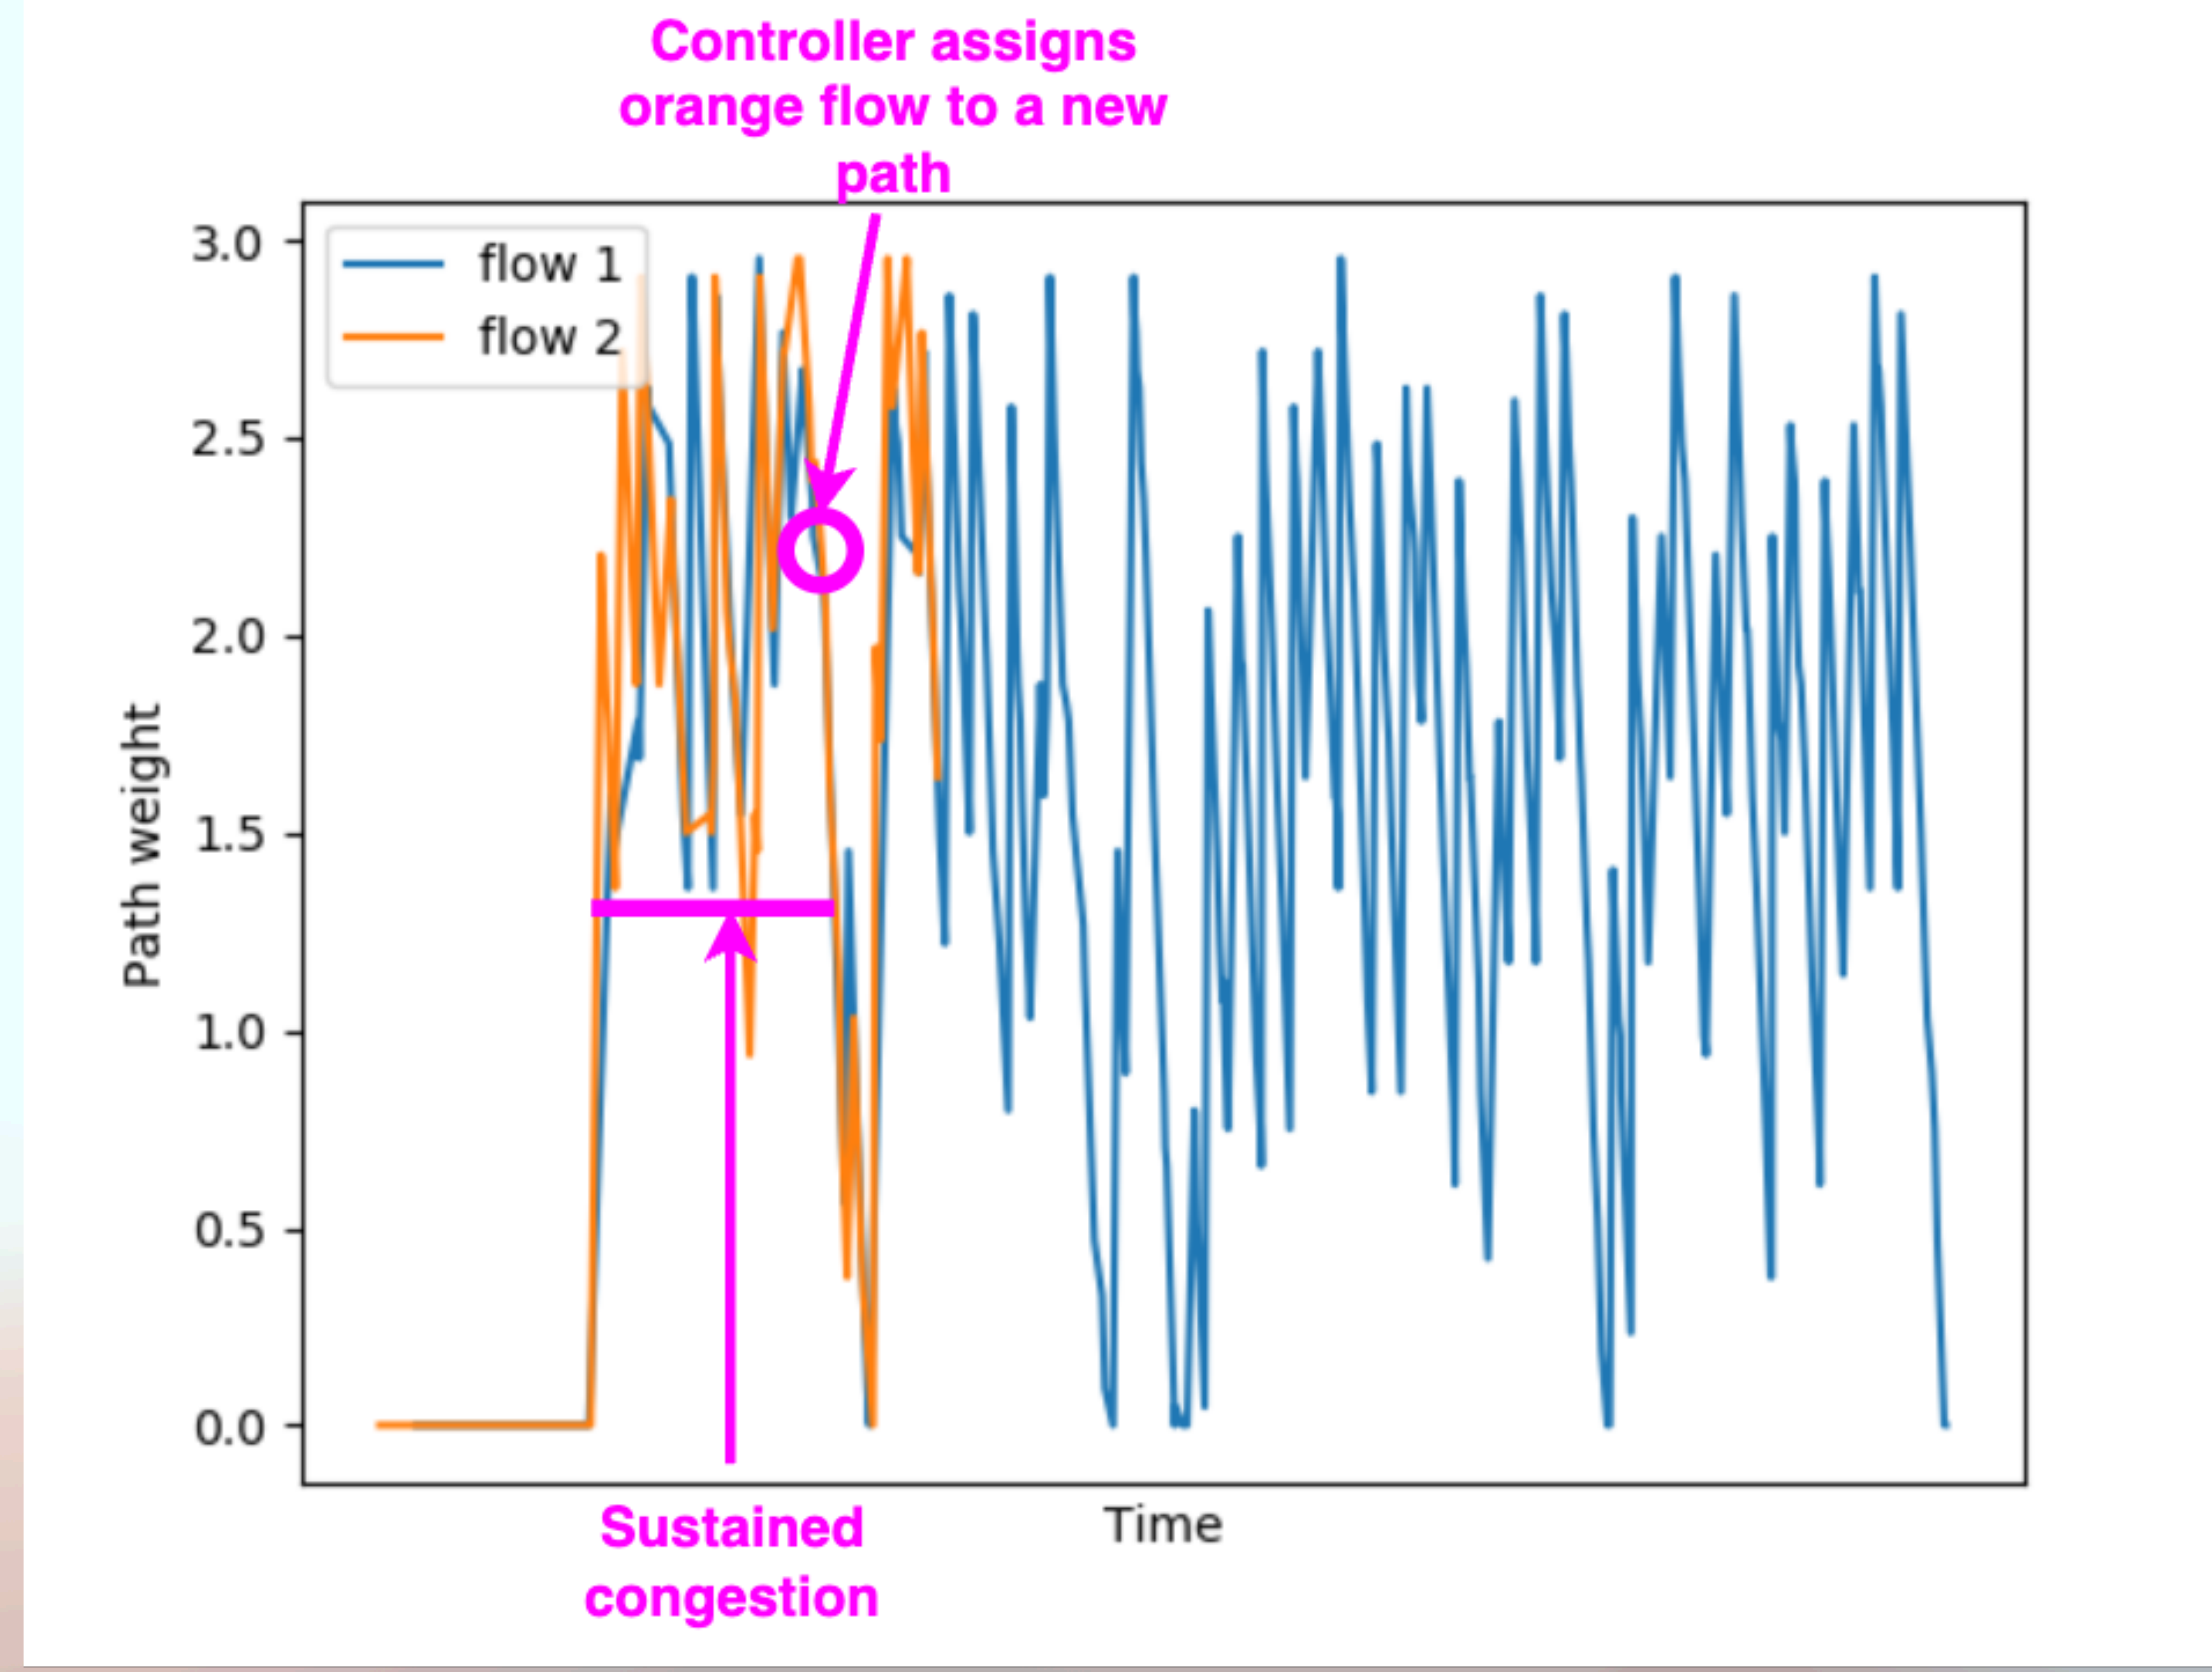

Improved resiliency with Smoothie

- **Algorithmic Collision Attacks:** In a static path load balancer, an attacker can use a Algorithmic Collision Attack to manipulate the path assignment function and force all traffic to flow through a specific path, causing a denial of service. (Ex: forge packets so that they have the same hash and end up on the same path with ECMP) ==> Solved by dynamically modifying the path to flow assignment.

- **Link Failures:** a person or group can disrupt the communication between two networked devices by disabling or damaging the physical connection between them. ==> Solved by **quickly** rerouting traffic.

- **(D)DOS:** (distributed) denial of service ==> Mitigated by spreading the traffic on all available paths.

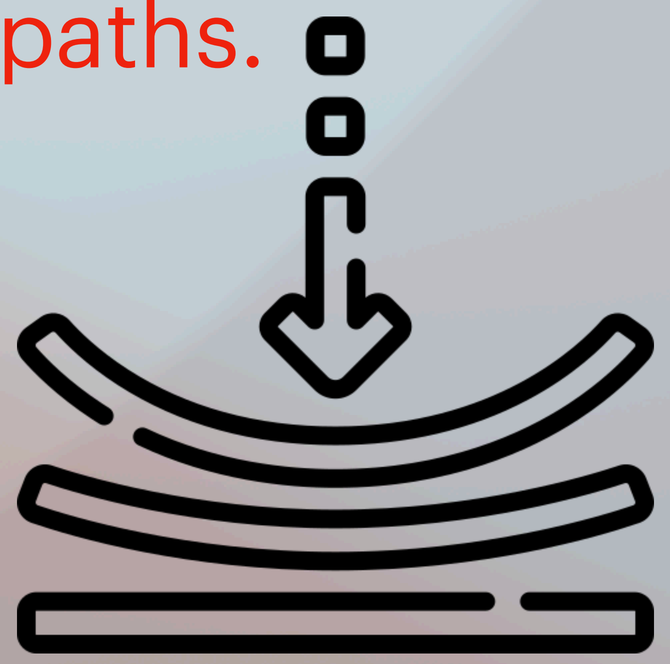

Collateral advantages of incorporating Smoothie

- Should improve[1] Flow completion time under load.
- Should improve[1] Service Level Agreement (SLA) percentage.
- The proof of concept is developed in P4 which is vendor agnostic.
- Fine grained load balancing.
- Reconfigurable on the field.

1. Still need to be proven by experiments

Take home message

- Dynamic path load balancers can improve the resiliency of a data center against cyber attacks.
- They distribute traffic, reducing the risk of (D)DoS attacks overwhelming any one server.
- They automatically reroute traffic away from servers under attack, ensuring critical applications remain available.
- They provide visibility and control over network traffic, allowing to quickly identify and respond to potential attacks.
- Dynamic path load balancers are an important part of a comprehensive cybersecurity strategy in the face of increasingly frequent and advanced cyber attacks.

Questions ?

Champagne Loïc — 2023
lchampagne@uliege.be

Data Center Observability

Smoothie

$$Link\_weight = (Q\_occupency / Q\_size) / (Link\_BW / Max\_BW)$$

$Link\_weight \in [0,1]$ (the Smaller the better)

$$Path\ weight = \sum_{n=1}^{Nbr\_links} weight(n)\ *\ f(n)$$

$$where\ f(i) = \begin{cases} 1 & \text{if link i on the path} \\ 0 & \text{otherwise} \end{cases}\quad and\ weight(i)\ is\ the\ weight\ of\ the\ i^{st}\ link$$

Data Center Observability (4)

Smoothie

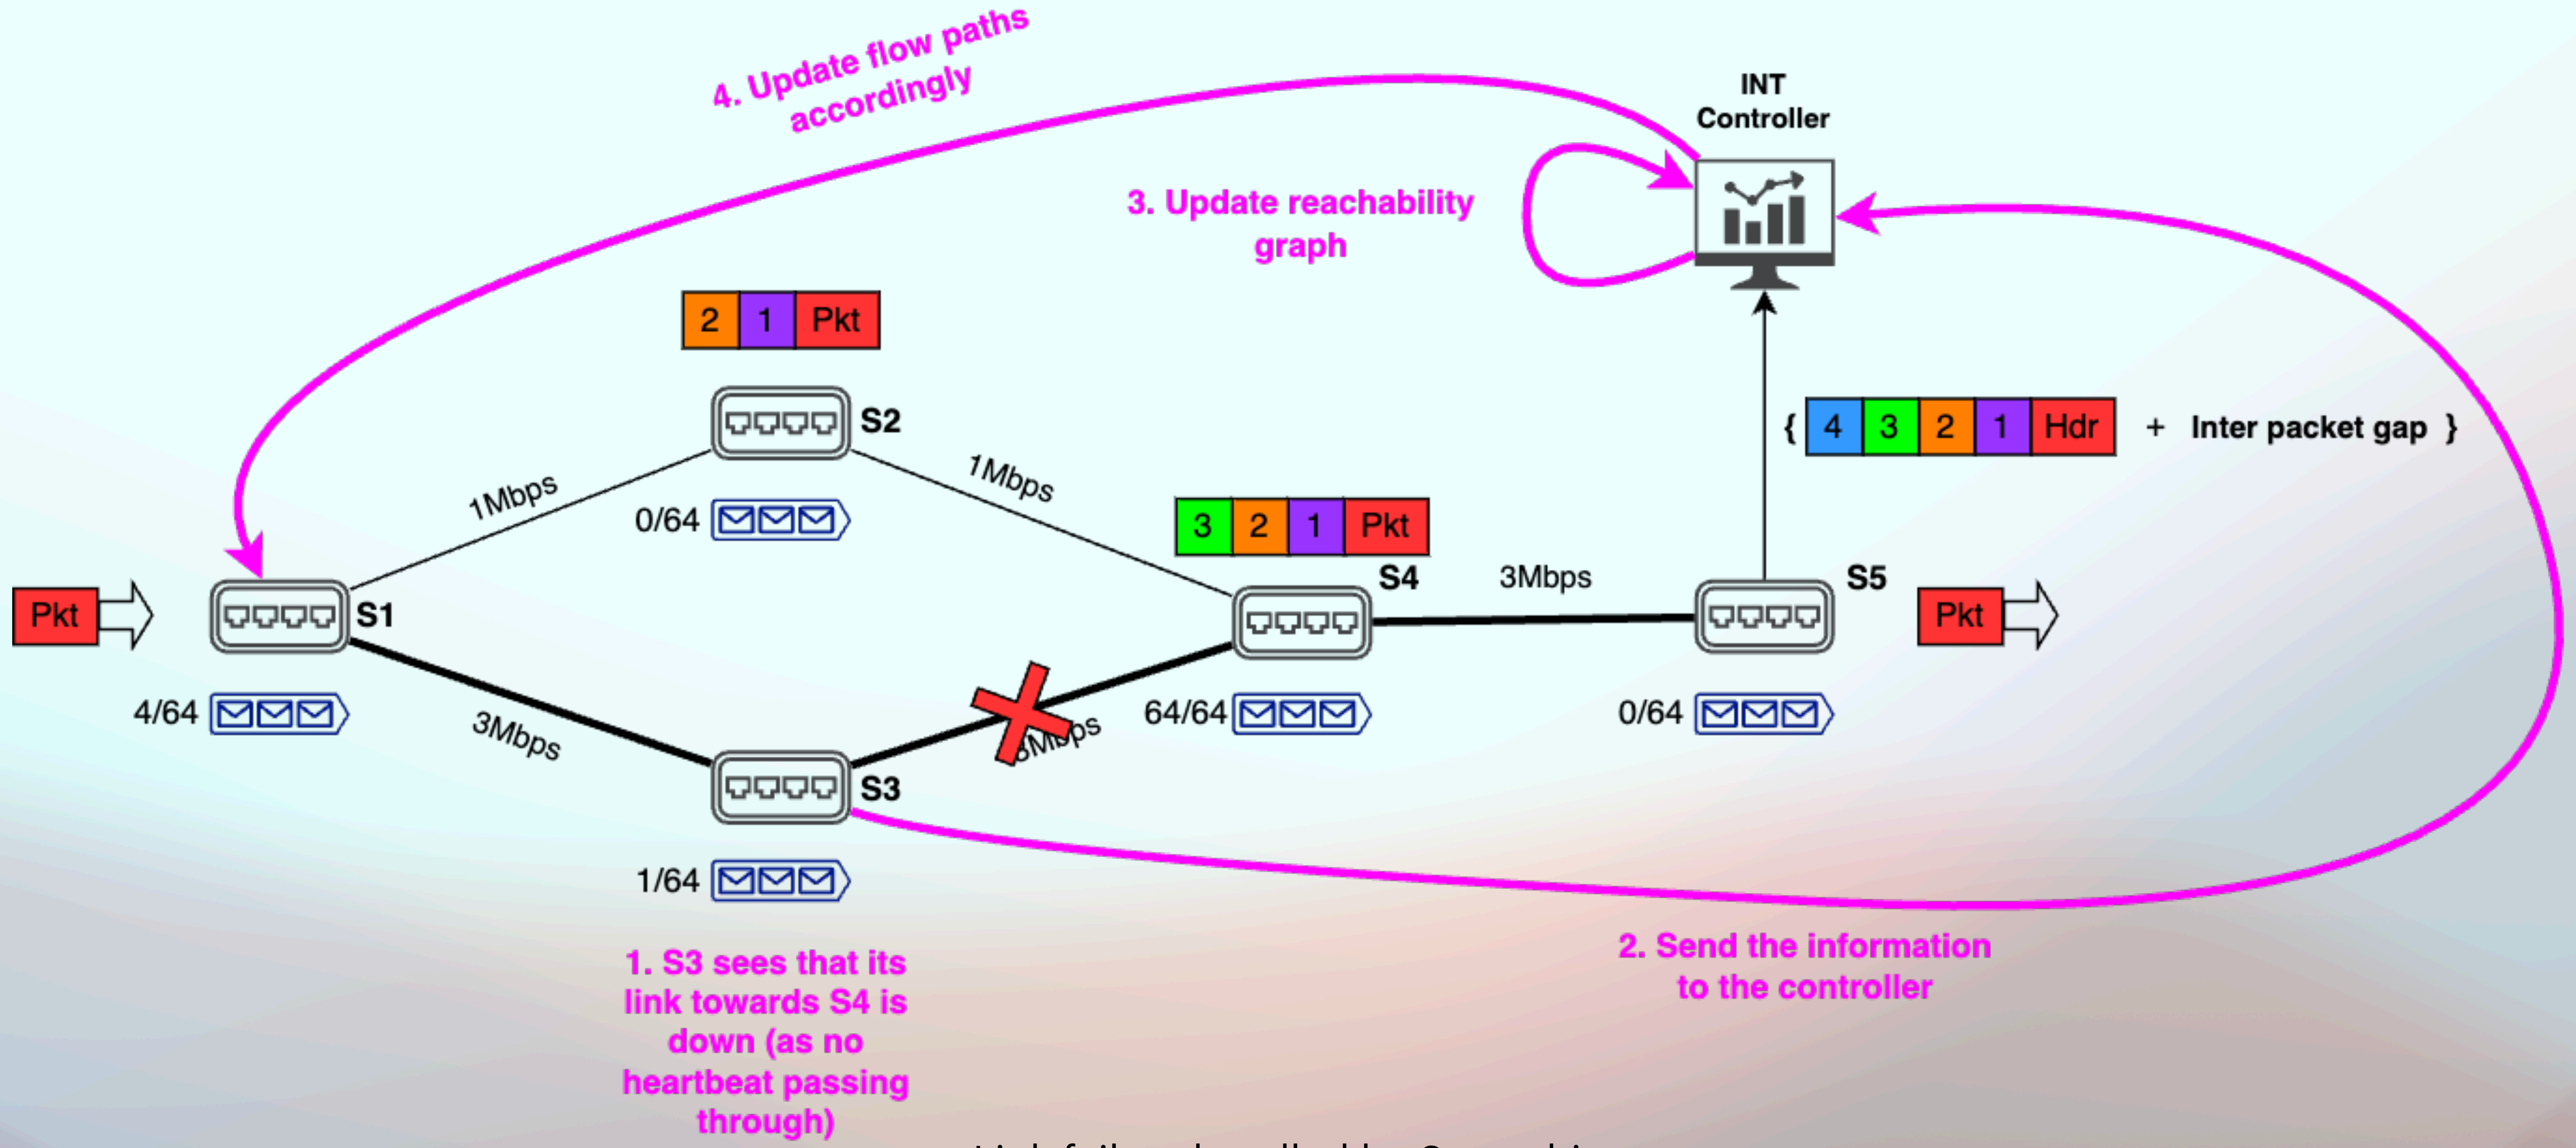

Link failure handled by Smoothie