

Service Function Chaining on Programmable Data Plane

IEEE/IEIE ICCE-Asia 2020

Network and Security Research Lab.

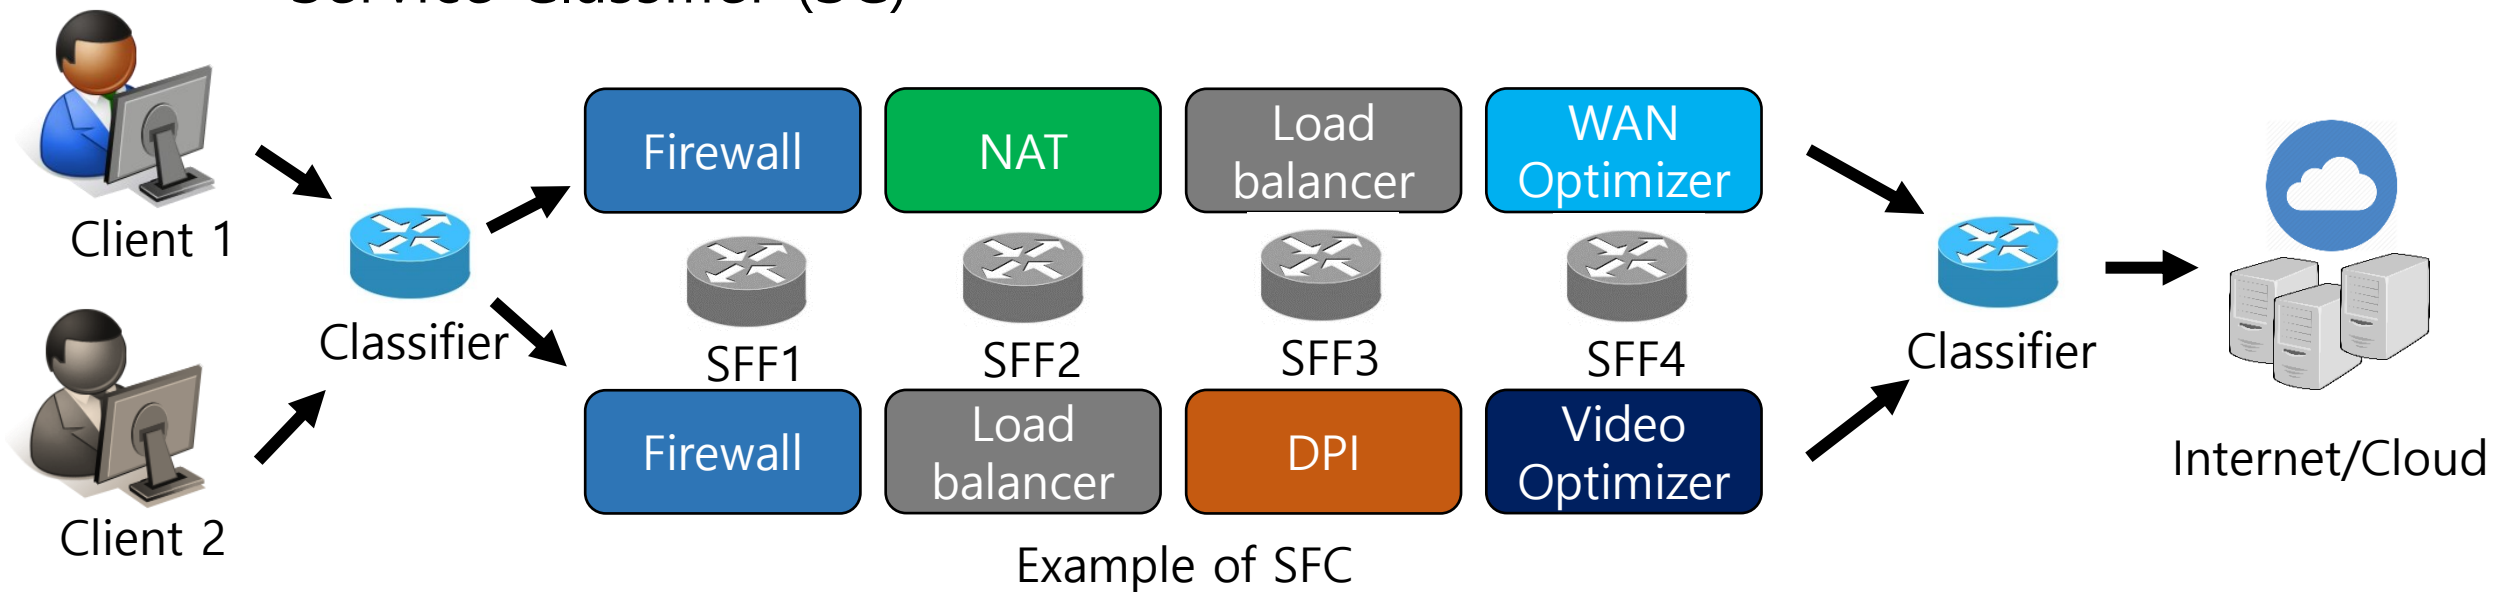
Korea University

Gyuyeong Kim and Wonjun Lee



Background

- Service Function Chaining (SFC)
 - Specifies the **set** and **order** of Service Functions (SFs) for a service
- Chaining Functions – SFC Components
 - Service Function Forwarder (SFF)
 - Service Classifier (SC)



P4-SFC Design

- Emerging Programmable switch (pSwitch) ASICs
 - We can customize the packet processing pipeline with high-level programming languages like P4
- Processing pipelines follow match-action rules

Match (if) \longrightarrow **Action (then)**

```
{  
  "table": "MyIngress.ipv4_lpm",  
  "match": {  
    "hdr.ipv4.dstAddr": ["10.0.1.1", 32]  
  },  
  "action_name": "MyIngress.ipv4_forward",  
  "action_params": {  
    "dstAddr": "00:00:00:00:01:01",  
    "port": 1  
  }  
}
```

L3 forwarding in P4



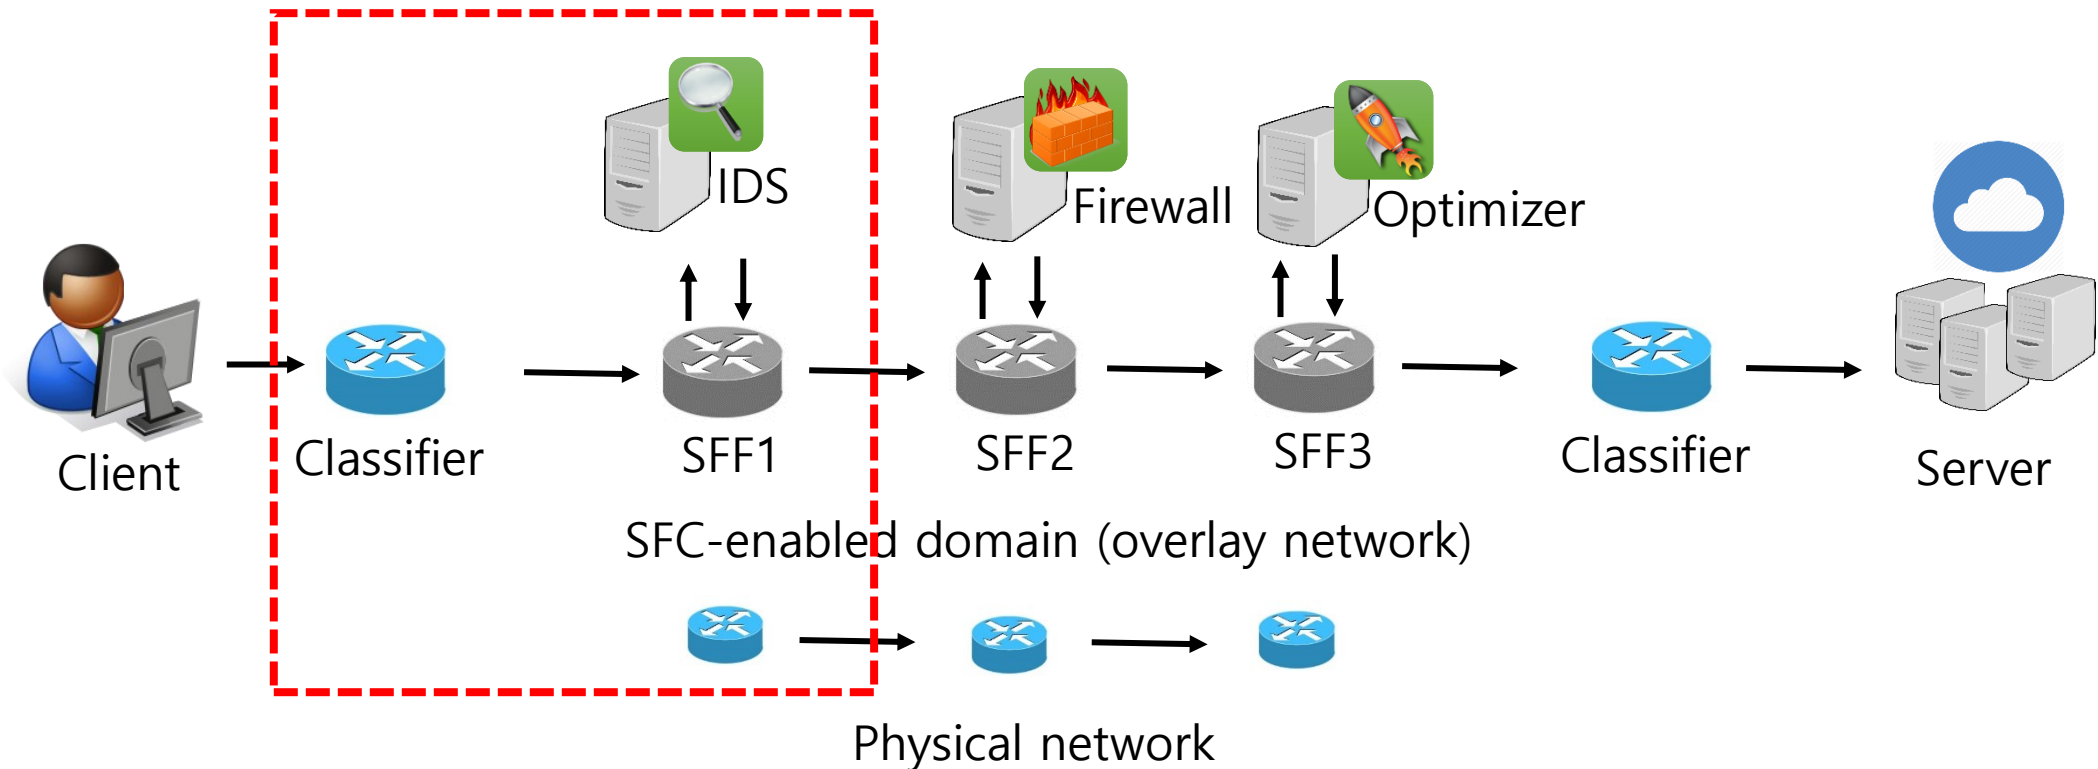
Barefoot Tofino



Cavium XPliant

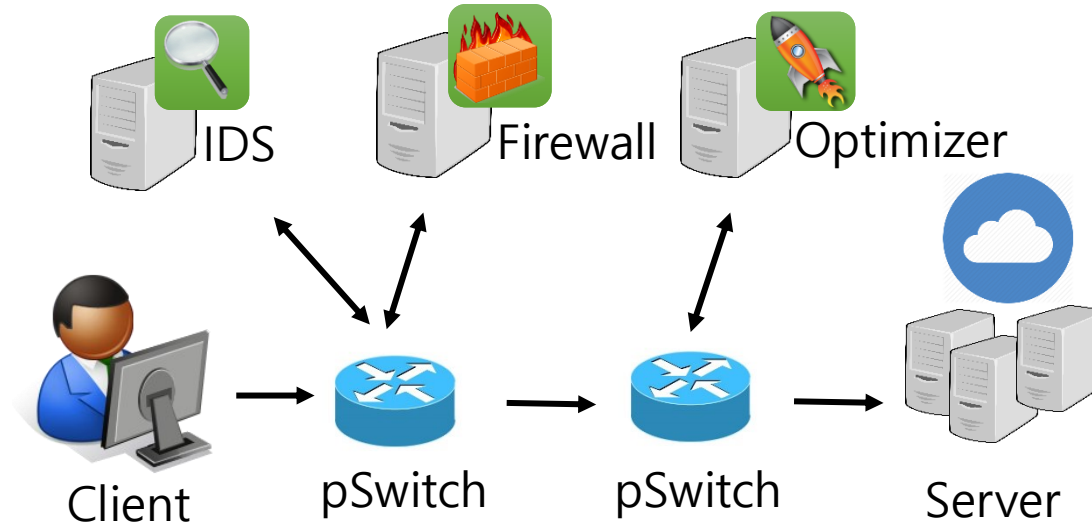
P4-SFC Design

- We propose P4-SFC, a pSwitch-based SFC design
- Possible scenario: Classifier + SFF + Switch



P4-SFC Design

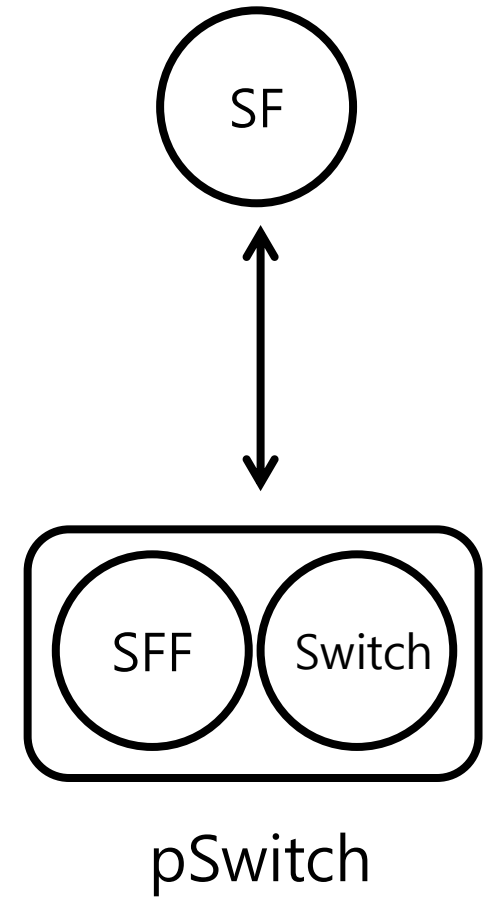
- Possible scenario: Classifier + SFF + Switch



P4-SFC Design

- Benefits of P4-SFC

1. We inherit **high throughput of switch hardware**
 - Billions of SFC packets can be processed per second with a sub-microsecond per-packet processing delay
2. Since the chaining functions run on the switch, **the packet does not experience RTTs to process packets** in the chaining function server
3. P4-SFC does not place chaining functions and SFs together. Therefore, **BM SFs are also compatible**



P4-SFC Design

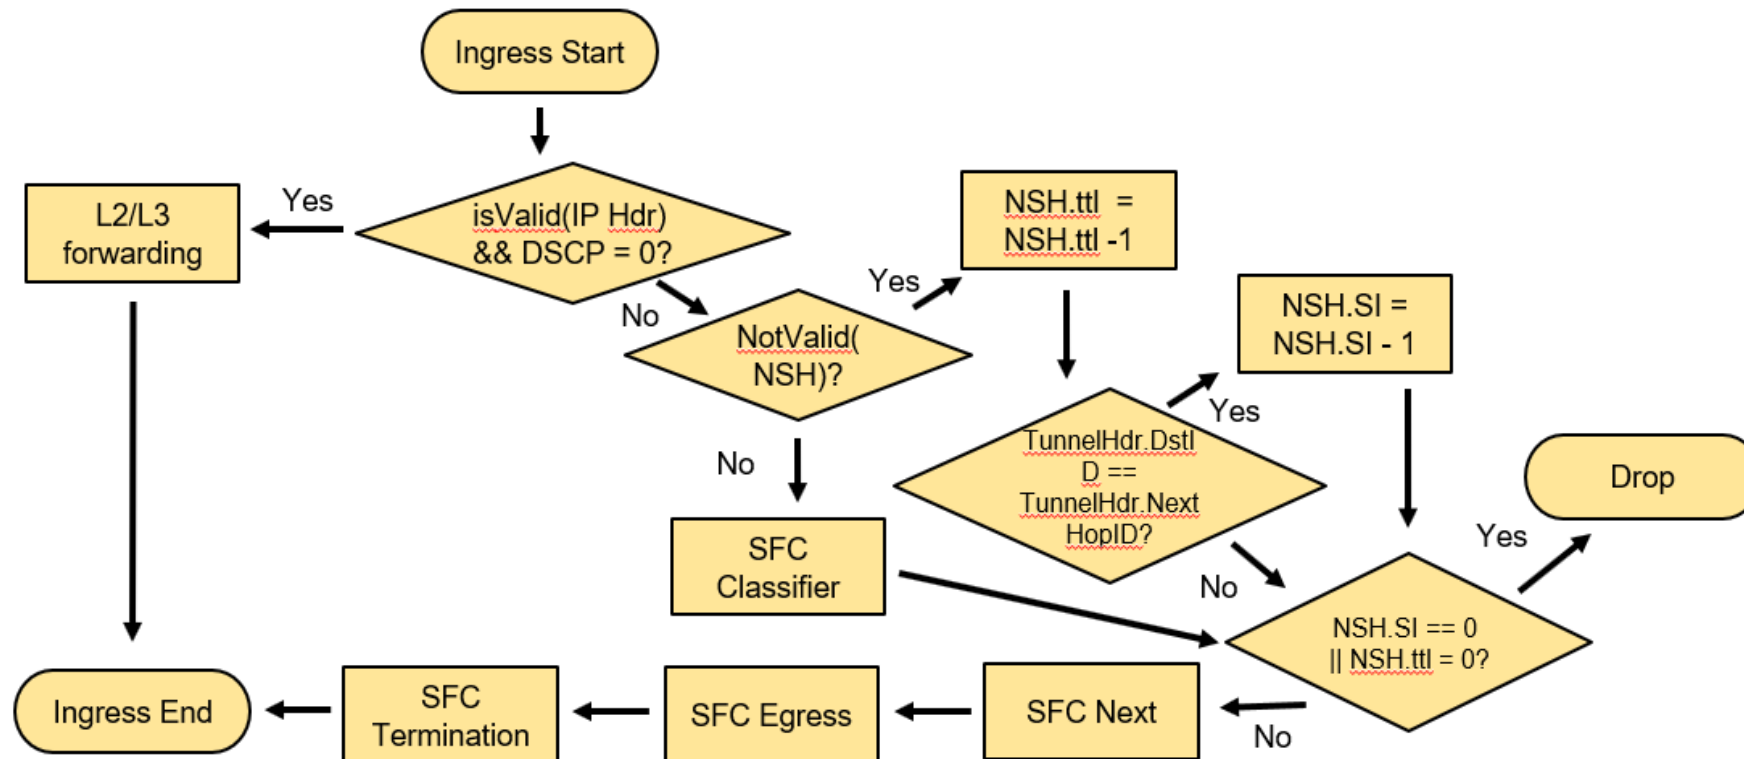
- Comparison P4-SFC with the existing SFC designs
 - Commodity server: Implementing chaining functions in commodity servers
 - Virtual switch: co-existing chaining functions and SFs in commodity servers together

| | Commodity server | Virtual switch | Programmable switch (P4-SFC) |
|------------------|------------------|----------------|---|
| Low latency | X | O | O |
| High throughput | X | X | O |
| BM compatibility | O | X | O |

*Bare-Metal (BM) SFs: SFs that use dedicated ASIC hardware

P4-SFC Design

- P4-SFC tables in Ingress pipeline



P4-SFC Design

- Performance evaluation
 - We implemented P4-SFC using P4 language
 - We conduct a Mininet-based experiment using P4 BMv2 software switches
 - Ping with 1.5KB packet
 - The results show that P4-SFC can reduce the average packet latency by half

| | SFC | P4-SFC |
|----------------|---------------|---------------|
| Max | 17.27ms | 17.06ms |
| Min | 6.64ms | 2.88ms |
| Average | 8.45ms | 3.88ms |

Conclusion

- We present P4-SFC, a pSwitch-based SFC design that provides
 - Low latency
 - High throughput
 - Reduced hardware and operating costs
- P4-SFC can be implemented at ingress pipeline using P4 language
 - NSH + Transport header
 - 4 Tables, 4 Actions

Thank you!

For more information

- Network and Security Research Lab., Korea University
- <http://netlab.korea.ac.kr>
- gykim08@korea.ac.kr