

## **P4-Protect: 1+1 Path Protection for P4**

Steffen Lindner, Daniel Merling, Marco Häberle, Michael Menth


*<http://kn.inf.uni-tuebingen.de>*



- ▶ Motivation
- ▶ Mechanism
  - Overview
  - Packet selection
- ▶ Evaluation
  - Processing time
  - TCP Throughput
  - Influence on Jitter
- ▶ Conclusion



- ▶ Critical traffic requires 100% reliability
  - Classical FRR mechanisms/reconvergence have too long response times
  - Reconvergence happens at time scale of a seconds
  - FRR detection may take a few 10s of milliseconds
  
- ▶ 1+1 Protection protects 100% at SLF/SNF without packet loss
  
- ▶ Classic field of application: Optical Networks, MPLS, Ethernet
  
- ▶ Our solution: 1+1 Path Protection for IP networks with P4
  
- ▶ Requirement: Start and end point are P4 switches (e.g. Tofino Edgecore Wedge)

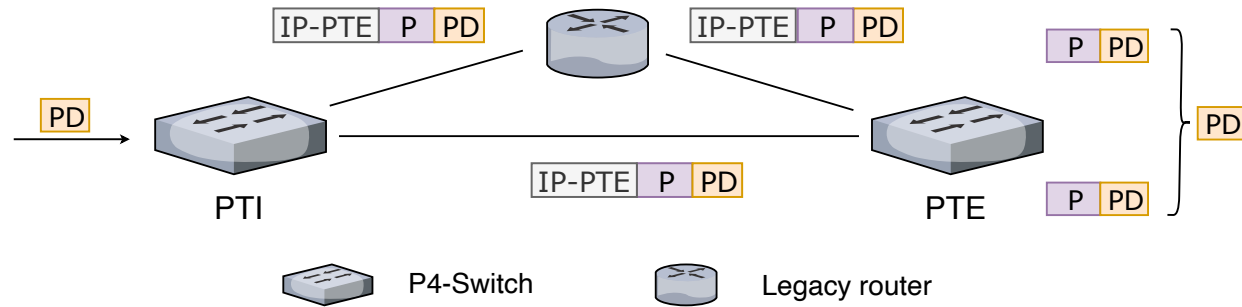


## ► Setup

- Two P4 switches are connected via a network with several, preferably disjoint paths
- The network can be composed of P4 switches and legacy devices
- It is also possible that the P4 switches are connected to each other via a foreign network

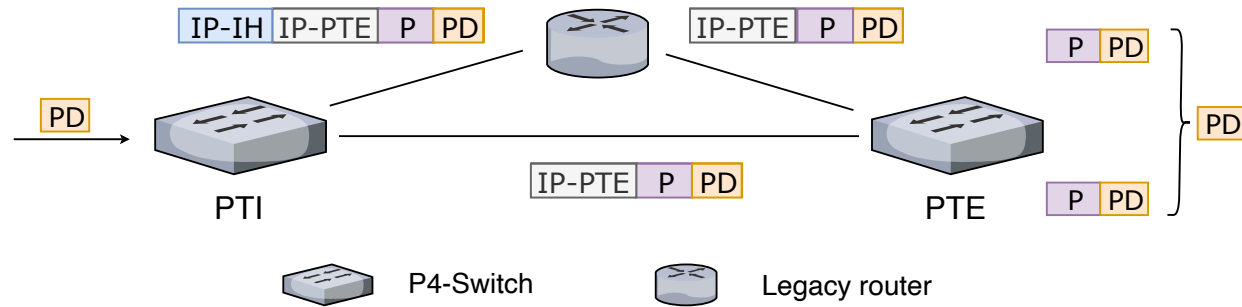
## ► Objective

- Redundant transfer of data between the two P4 switches




## ► Mechanism

- Upon packet arrival, ingress node (PTI) decides if traffic should be redundantly transmitted
- Redundant traffic gets replicated onto two different paths
- Traffic is tunneled to egress node (PTE)
- The PTE decides which packet version gets accepted



## ► Path selection

- In the own network, traffic engineering can be used to ensure disjoint paths between PTI and PTE
- To secure traffic over a foreign network, e.g. the internet, we suggest the use of triangular routing
- One path is build using an intermediate IP hop to enforce as disjoint paths as possible



## ► Protection header

- PTE needs to forward packets without duplicates
- Protection header contains connection identifier and sequence number
- Sequence number is used to accept the first version of each packet



## ► Packet Selection

- Duplicates must be recognized at the PTE and must not be forwarded
- Due to packet loss and latency differences packets with sequence numbers within a given range should be accepted
- We leverage a receive window to compensate packet loss and latency differences
- Acceptance criterion with window width  $W$ :


$$SN_{last}^{PTE} < SN \leq SN_{last}^{PTE} + W$$

- Due to cyclic sequence number range wrap-around must be considered



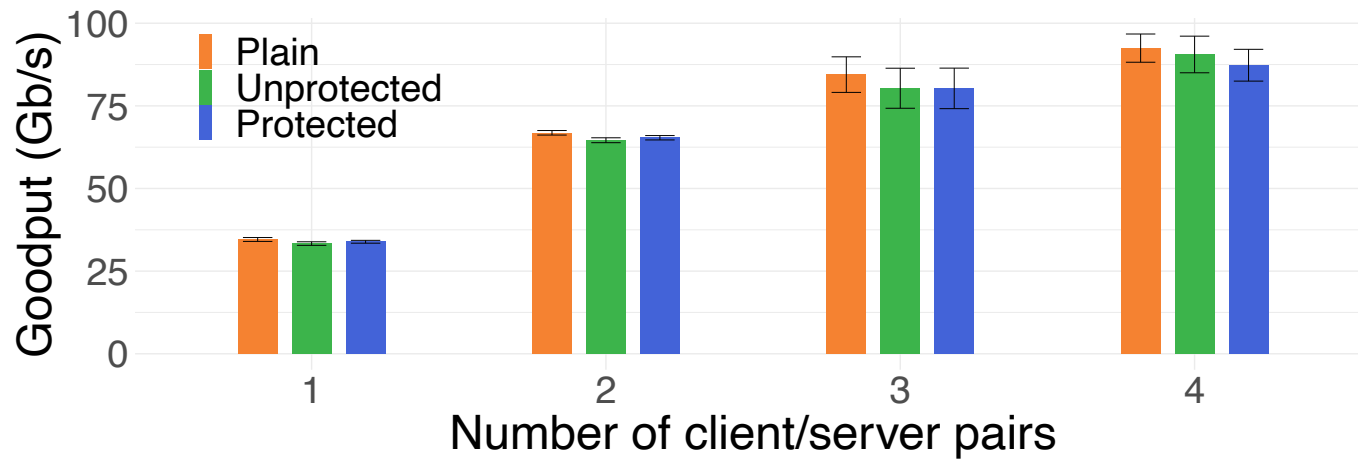
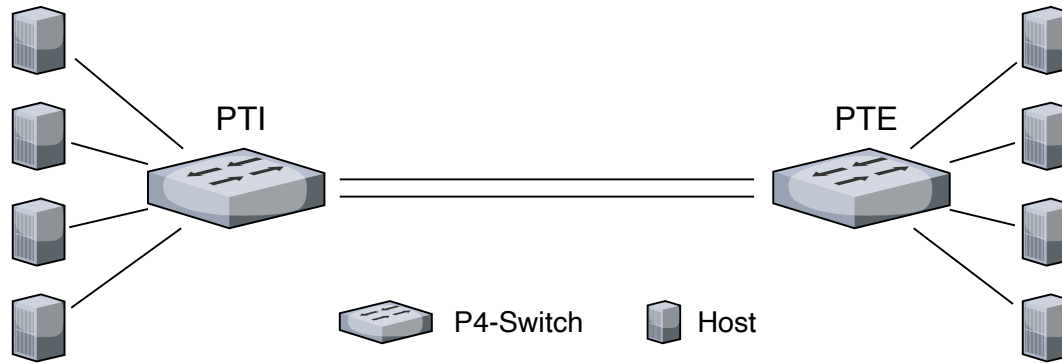


# Evaluation – Processing time







# Evaluation – TCP Throughput





# Evaluation – Influence on Jitter





- ▶ P4-Protect enables 100G 1+1 protection without performance loss compared to legacy implementations
- ▶ Only two 100G P4 switches are required (PTI and PTE)
- ▶ With the help of triangular routing, P4-Protect can also be used across foreign networks



## **Steffen Lindner**

University of Tuebingen

Faculty of Science

Department of Computer Science

Chair of Communication Networks