Packet Broker: A Network Appliance Based on P4

SWITCH

Alexander Gall
alexander.gall@switch.ch

TNC 2021

P4 and Data Plane Programming Bof

18.6.201

Motivation: Netflow Export

- Unsampled Netflow: can’t use Netflow on routers
- Optical taps on external interfaces to copy packets
- “Packet Broker” to aggregate packets onto 2x100Gbps links
- Exporter creates and exports flows

Per-PoP Exporter Architecture

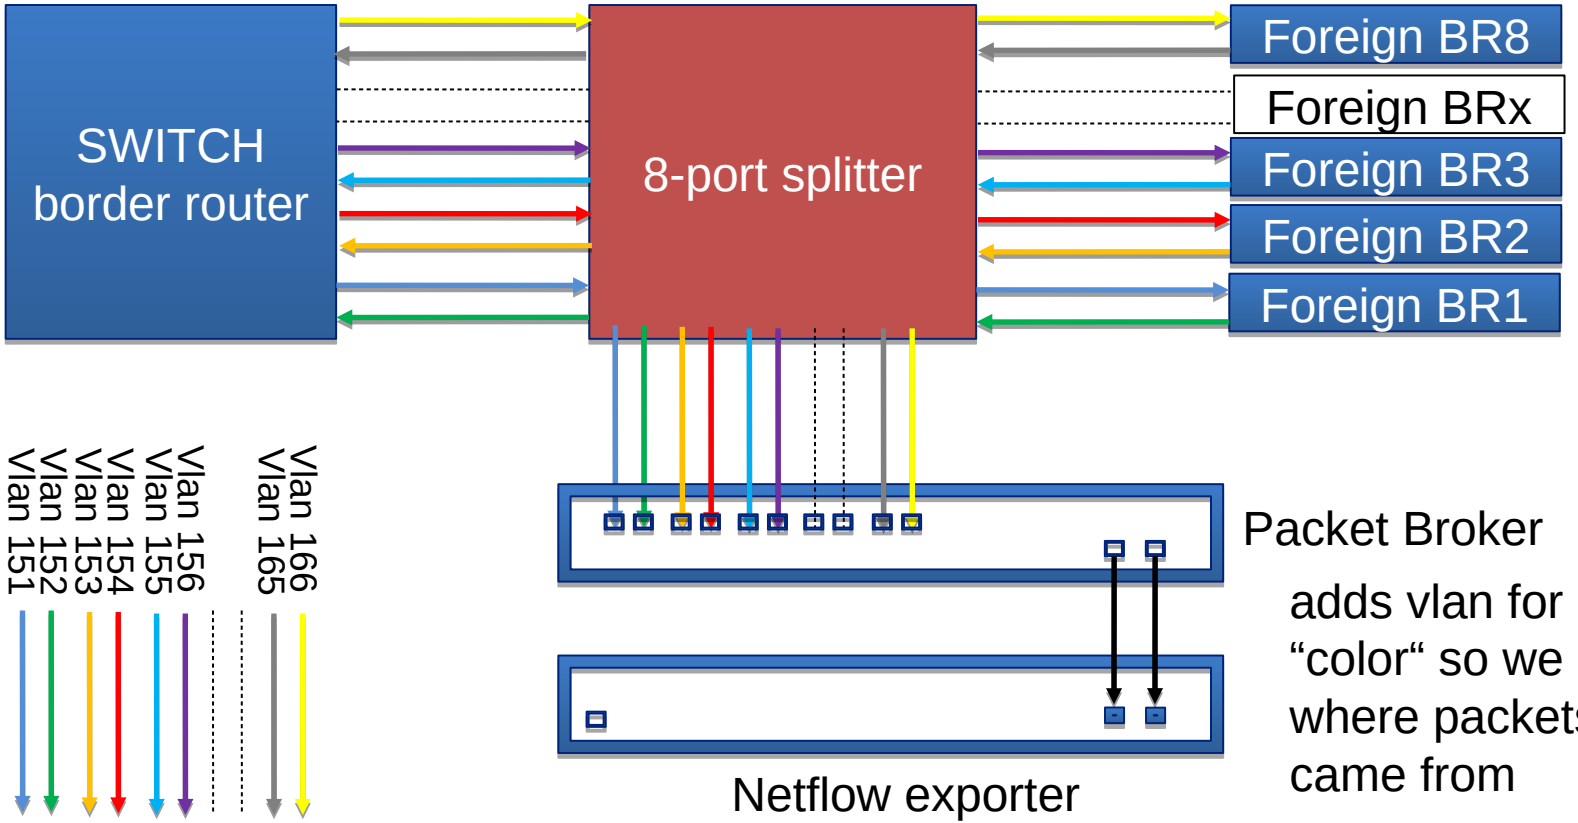

Functionality

- Flow-aware aggregation towards exporter (via hash over flow-fields)
- Add or rewrite VLAN tags on ingress to identify original links
- Rewrite source/destination MAC addresses
- Filter packets based on source address
- “Flow mirror”: mirror packets based on flow signatures
- Redirect dropped packets to arbitrary port (“deflect on drop”)

Hardware

- P4-programmable, based on Tofino ASIC
- Device from Edgecore, 32xQSFP (WEDGE 100BF-32X), ~6k EUR

Implementation

- Tofino-specific P4 program: https://github.com/alexandergall/packet-broker
- Requires Tofino SDE from Intel to build (under NDA)
- Easy to extend
- Simple control-plane
  - Python script runs as daemon, talks to `bf_switchd` via gRPC to program the tables from a config file
  - SNMP agent provides ifMIB for Tofino ports
  - Command-line tool (`brokerctl`) to interact with daemon

Complete Appliance

- Packaging using the Nix package manager
  - Easy release management with rollback capability
- ONIE-based installer
  - Installation of fully functional system in a single step
- Not yet public due to legal constraints from Intel :(