

### MOTIVATION

Can we build an inexpensive yet powerful router?

- Routers today are proprietary, closed-source, and expensive devices.
- Both routers and switches provide packet switching at line rate.
- **OpenFlow** [1] provides a programmable data-path.
- **Quagga** [2] implements an open-source software router.

### OUR APPROACH


FIBIUM [3] is a novel architecture for designing routers in an alternative way.

- Internet traffic is heavy-tailed among prefixes – efficient **prefix caching** might be feasible with a sufficiently smart caching scheme.
- An **OpenFlow**-capable device switches traffic contributed by cached top prefixes. A PC handles the remainder.
- **Quagga** peers with other routers and maintains the full routing information base (RIB).
- **FIBvisor**, our OpenFlow controller, implements a glue layer to bring together the software router and the switch.

### FIBIUM BUILDING BLOCKS

#### Design


- **Port mirroring**: Trunked VLAN to create virtual siblings of physical router ports, allowing Quagga to operate.
- **Fast path**: Pure hardware-based switching of traffic towards popular prefixes.
- **Slow path**: Unpopular prefixes take the detour through the PC to let Linux decide on routing.




#### FIBvisor

- **Prefix caching**: We propose **FIBpredict** as a caching strategy to keep a continuously updated set of popular prefixes not to overwhelm the slow path.
- **Traffic monitoring**: Keep traffic statistics for both, fast path and slow path.
- **FIB monitoring**: Monitor modifications on the Linux FIB by Quagga.

End of time interval:



Incoming BGP update:




#### FIBpredict

- **Select prefixes** for the fast path FIB (OpenFlow flow table) based on short, mid, and long-term traffic statistics to keep the adaptation time low while at the same time remain stable among the long-term top prefixes.
- **Reduce FIB churn** by dampening the fluctuation among less popular prefixes.
- **Update the fast path FIB** periodically, but react on BGP updates immediately.
- **Detect malicious activity** and consider countermeasures, e.g., discard traffic originated by a certain source using ACL's.

### FEASIBILITY AND PERFORMANCE EVALUATION

#### Baseline experiments


- Experiments based on 48hrs of anonymized packet-level data as seen by a customer aggregation router of a large European ISP.
- The optimal case with future knowledge supports our claim, that efficient prefix caching in theory is possible.




Baseline slow path requirements.

#### Prototype evaluation

- Simulating FIBpredict results in a behavior which seems close to the optimal case.
- LRU and LFU, both reacting on individual cache misses, still show an idealized behavior, which is unfeasible in practice.



Comparison of slow path requirements.



Churn for different caching schemes.

### FUTURE WORK

- Evaluate the performance of FIBpredict with more data from different network locations.
- Measure the performance of different OpenFlow-capable devices, e.g., quantify the inherent penalty induced by flow table modifications.
- Deploy FIBvisor in an experimental facility with real users.

### REFERENCES

- [1] MCKEOWN, N., ANDERSON, T., BALAKRISHNAN, H., PARULKAR, G., PETERSON, L., REXFORD, J., SHENKER, S., AND TURNER, J. OpenFlow: Enabling Innovation in Campus Networks. In ACM CCR (2008).
- [2] Quagga Routing Suite. <http://www.quagga.net>
- [3] FELDMANN, A., UHLIG, S., SHERWOOD, R., SARRAR, N. Route Caching System. Patent application 10007/TUB.