DPDK

DATA PLANE DEVELOPMENT KIT

Lagopus Router

Tomoya Hibi, Hirokazu Takahashi

NTT Network Innovation Lab.

DPDK Summit - San Jose – 2017

Lagopus Project

- Develops SDN/NFV software node for carrier networks to realize
  - Differentiate services
  - Time-to-Market
  - Cost-efficiency
- Supported by
  - The Ministry of Internal Affairs and Communications of Japan (April, 2013 - March, 2016)
  - NTT (April, 2013 -)
- Release
  - Lagopus software OpenFlow 1.3 switch (July, 2014 -)
    - https://github.com/lagopus/lagopus
  - Lagopus software router (Aug., 2017 -)
    - https://github.com/lagopus/lagopus-router (alpha version)

Lagopus Switch

- Pure OpenFlow 1.3 Software Switch with Tunnel Extensions
  - Open soure under Apache v2 license (https://github.com/lagopus/lagopus)
  - High Performance with DPDK (over 20Mpps, over 1M flow entries)
  - Provides good programmability for stateless network functions (e.g. SDN-IX)
  - Difficult to handle stateful/complicate protocols such as
    - VRRP, ICMP(e.g. reply too large), IPv4 fragmentation, IPsec

Lagopus Router

- Modular architecture OSS router
  - https://github.com/lagopus/lagopus-router
  - Apache 2.0 License
  - Written in C + golang
    - High performance packet processing by C + DPDK
    - High extensibility of protocol handling by modular architecture and golang
  - Configuration is based on OpenConfig
  - Supports protocols
    - L2: VLAN
    - L3: IPv4/v6, ICMP, VRF, VRRP
    - Tunneling: IP/IP, GRE, VxLAN, IPsec(IKE v1/v2, AES256/SHA-2))
    - Other: Match-Actions tables module (like OF or P4)

※ under developing

Design

DPDK

- Loose coupled modular-based design (like a microservice architecture)
- Unified configuration datastore
  - Configuration defined by yang
  - Pub/Sub-like APIs, config datastore as a broker
- Using existing routing agents
  - Slightly modified for APIs
- High performance/extensibility dataplane
  - Fast path written in C + DPDK
  - Slow path written in go

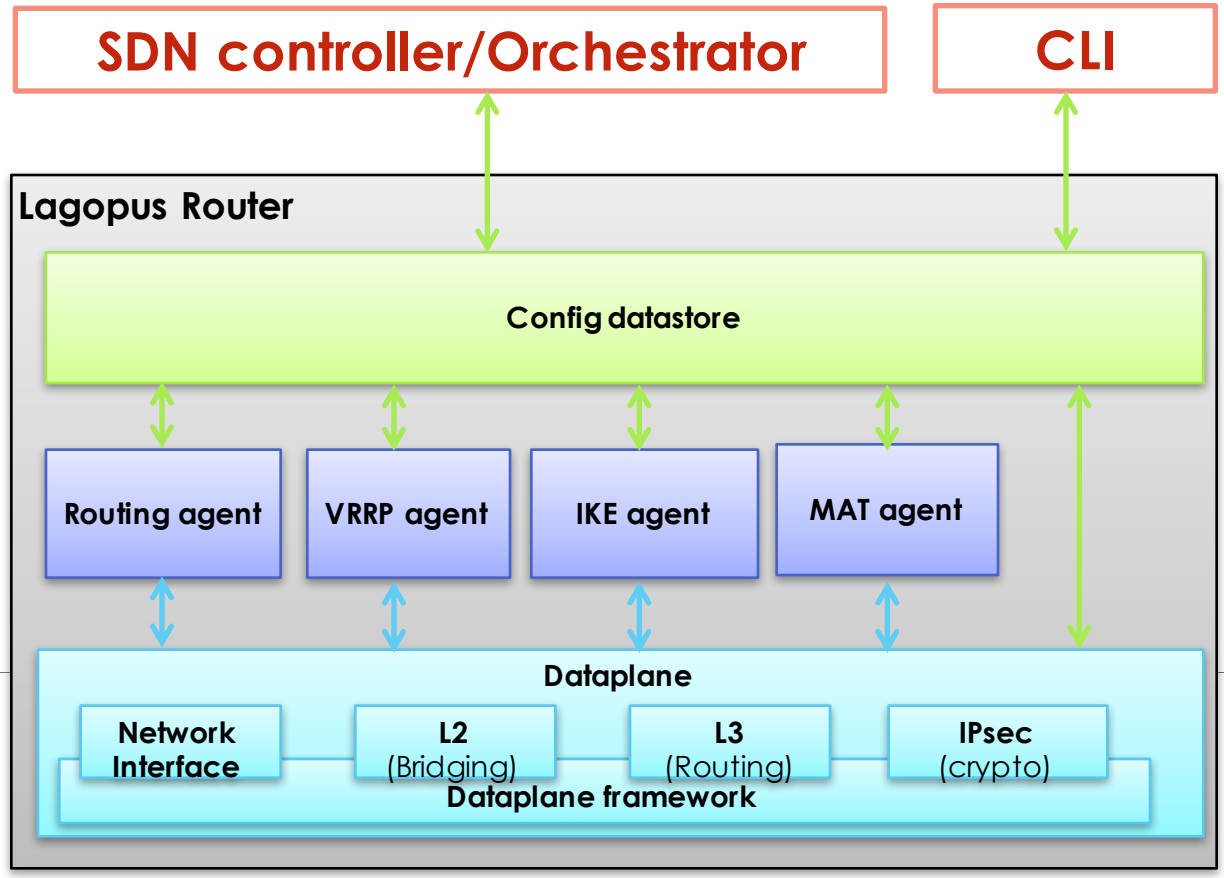

Architecture of a dataplane module

- Written in C + golang
- Packet Processing
  - Written in C with DPDK for the performance
- Module managing and Handle control packets
  - Written in golang
  - Subscribe configuration from the config datastore
  - Protocol handling, likely ARP, ICMP, BGP with router agens.
- More details
  - Takanari Hayama, Using DPDK with Go, 2016 Dublin DPDK userspace
  - https://dpdksummit.com/Archive/pdf/2017Userspace/DPDK-Userspace2017-Day2-11-using_dpdk_with_go.pdf

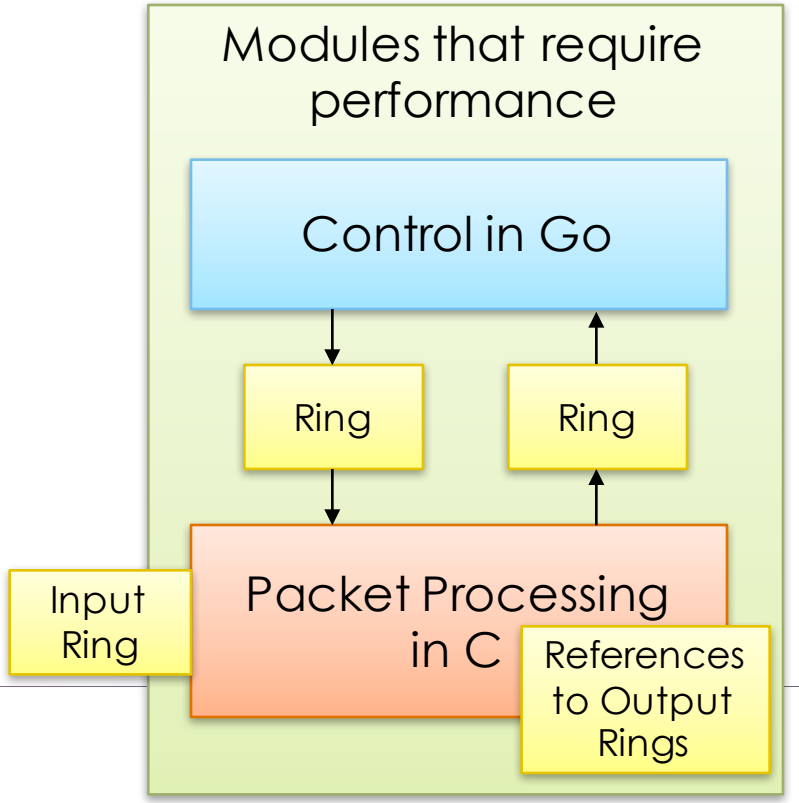

Config Datastore

- Openconfigd
  - https://github.com/coreswitch/openconfigd
  - Commit & Rollback support configuration system in Zebra2.0
  - Configuration is defined by YANG model
  - Automatically generates APIs
    - gRPC
    - CLI with completion/show/config

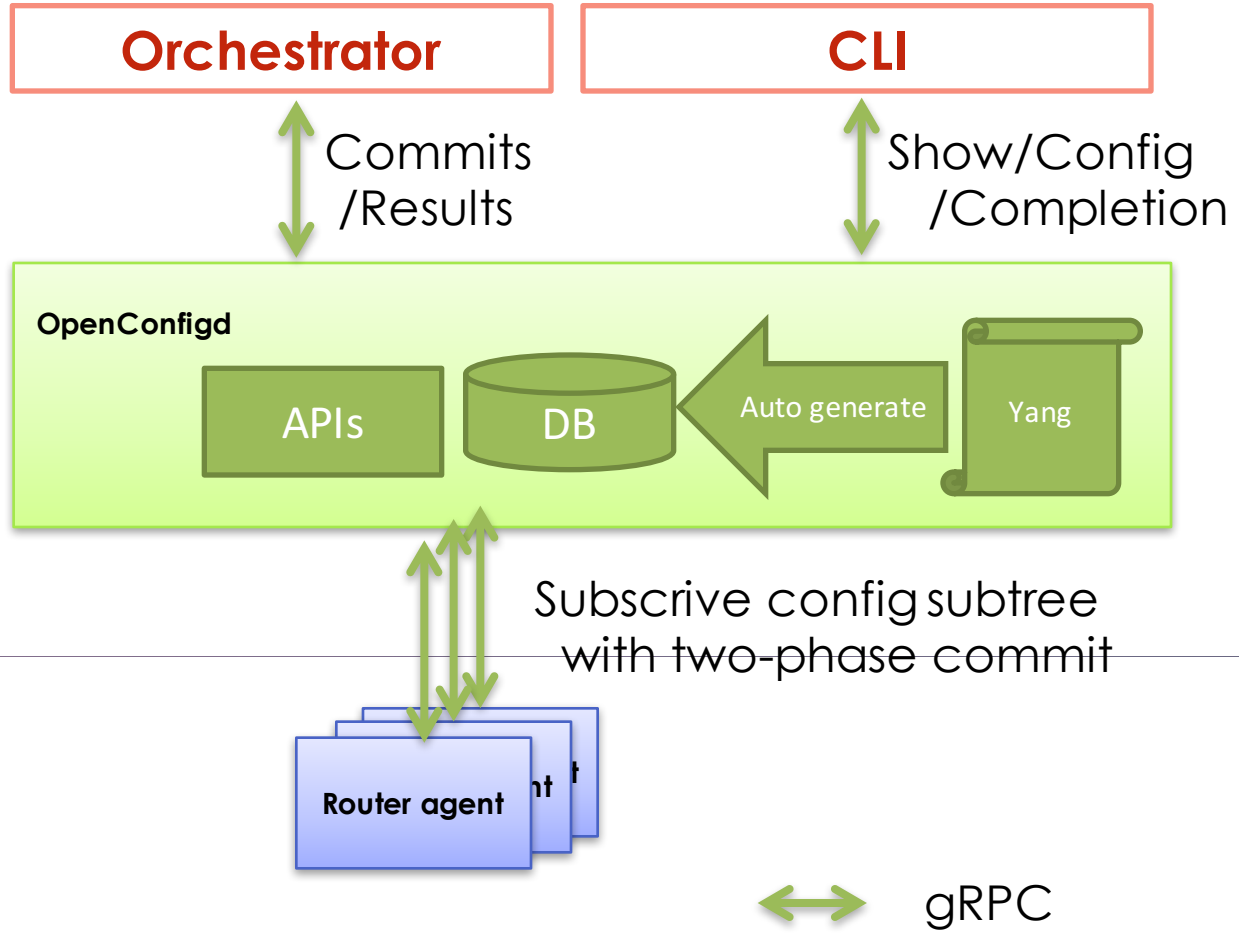

Performance

- For the most simple rules (two routing rules)
  - Over 10Gbps (for long packets)
  - About 1 Mpps
- Performance tuning and detail performance measurement and tuning future works.

Roadmap

- Aug. 2017
  - Alpha release
  - Simple L3 module with VRF and VRRP
- 2018.1Q
  - Beta release
  - L2 module
  - L3 Tunnel module (IP/IP)
  - Match and Actions Tables
  - Stats module
  - Performance tuning
- 2018.2Q
  - IPsec module(IKE v1/v2, AES256/SHA-2)
  - L2 Tunnel module (VxLAN)
- Future works
  - NAT
  - IPv6
  - MPLS
  - BGP Flowspec
  - ACL
  - xFlow

Summary

- Lagopus Router
  - Modular architecture OSS router
    - https://github.com/lagopus/lagopus-router
  - Written in C + golang
    - High performance packet processing by C + DPDK
    - High extensibility of protocol handling by module architecture and golang
  - Supports many protocols including OF like tables

Reference

- Lagopus
  - https://lagopus.org
- Github
  - OF1.3 Switch: https://github.com/lagopus/lagopus
  - Router: https://github.com/lagopus/lagopus-router
- Lagopus books
  - OF1.3 switch: http://www.lagopus.org/lagopus-book/en/html/
  - Router: Under working
- Mailing list
  - https://lists.sourceforge.net/lists/listinfo/lagopus-devel
- Slack
  - https://lagopus-project-slack.herokuapp.com/