Zebra 2.0 and Lagopus: newly-designed routing stack on high-performance packet forwarder

Kunihiro Ishiguro, Yoshihiro Nakajima,
Masaru Oki, Hirokazu Takahashi

NTT

Zebra 2.0

Open Source Routing Software

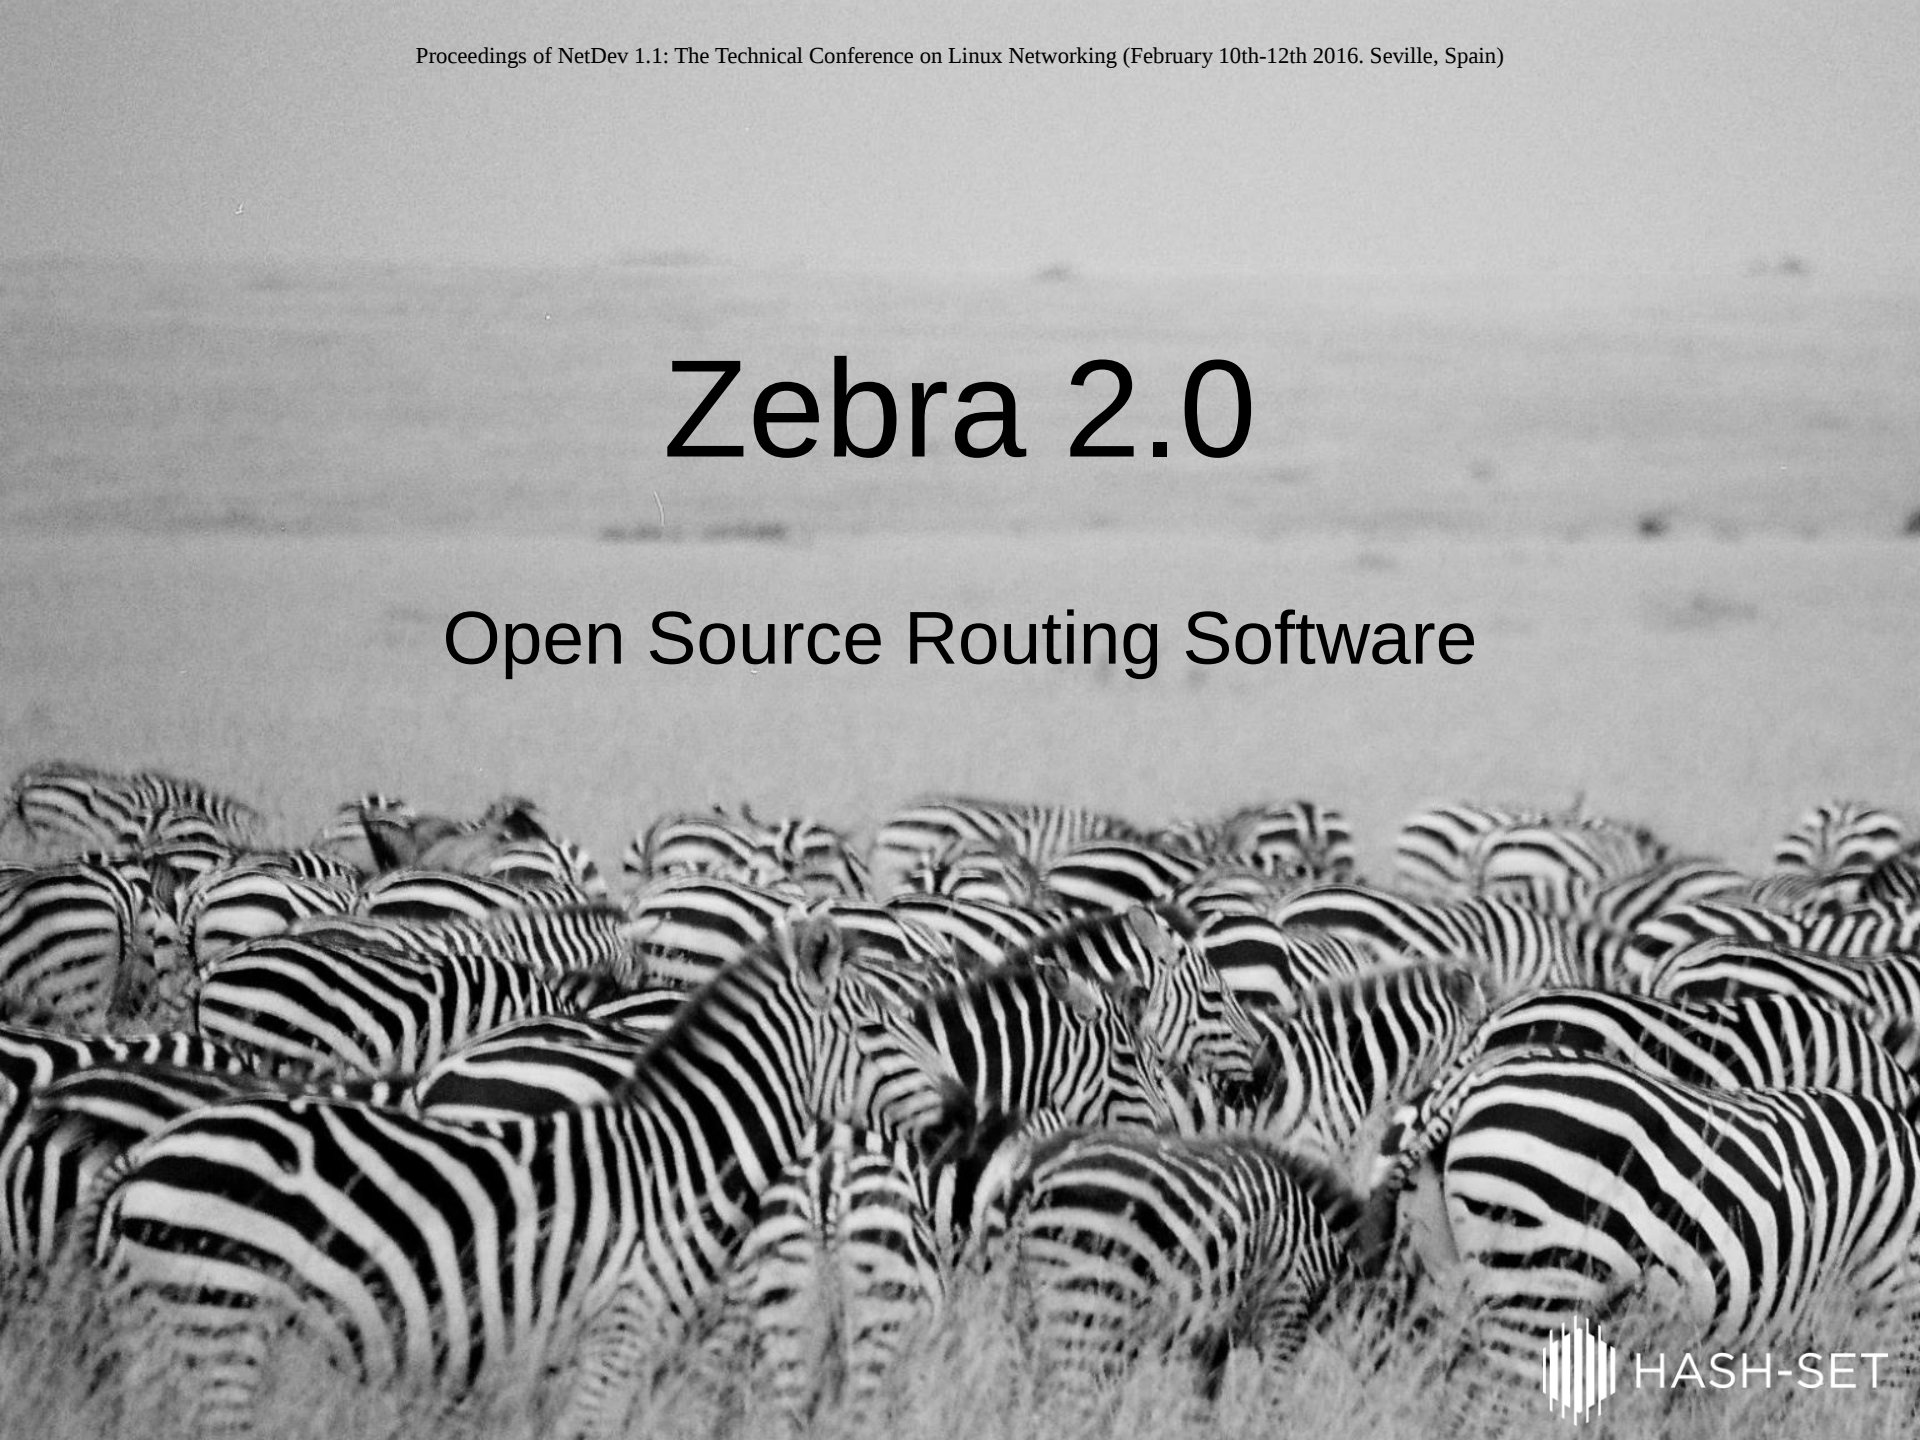

HASH-SET

Open Source Revisited

- Apache License
- Written From Scratch in C++
- Task Completion Model + Thread Model
- Single SPF Engine for OSPFv2/OSPFv3/IS-IS
- Forwarding Engine Abstraction for DPDK/OF-DPA
- Configuration with Commit/Rollback
- MPLS VPN Support

Architecture

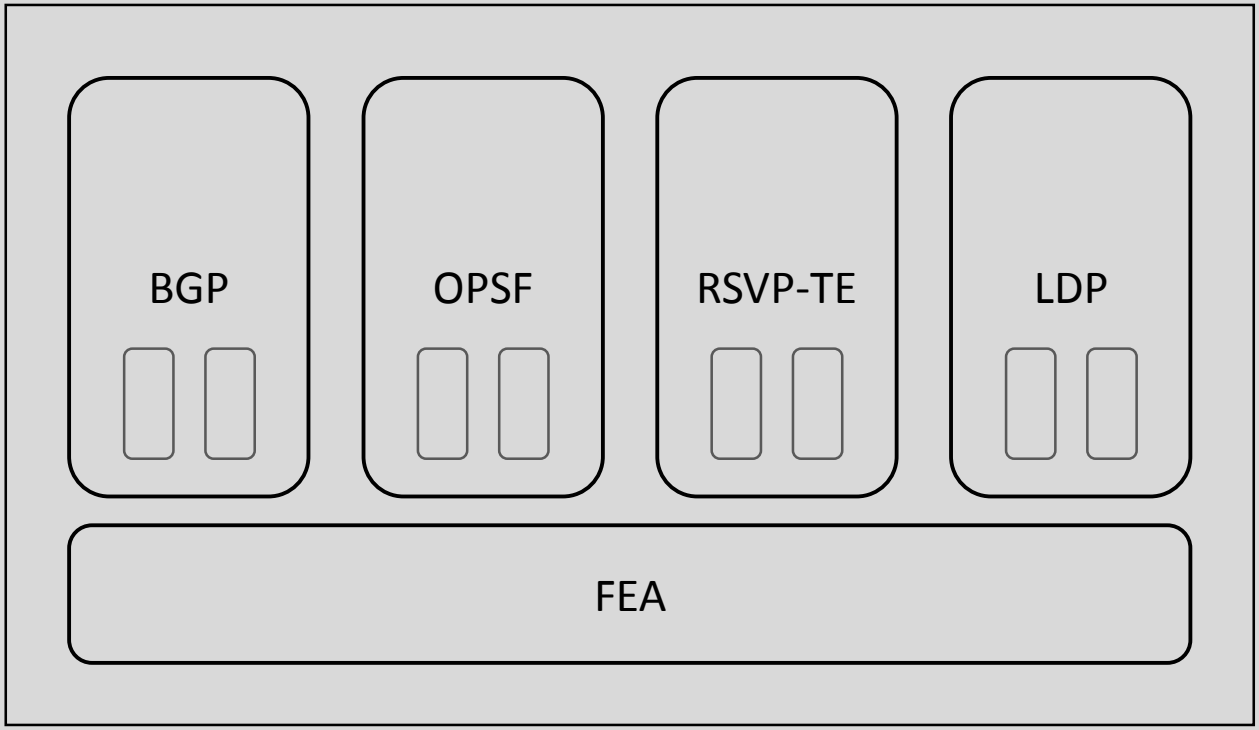

- Single Process/Multithread Architecture
- Protocol Module has 2 threads

OpenConfigd

- Commit & Rollback support configuration system
- Configuration is defined by YANG
- CLI, NetConf, REST API is automatically generated
- `confsh` - bash based CLI command
- OpenConfig is fully supported

OpenConfigd Architecture

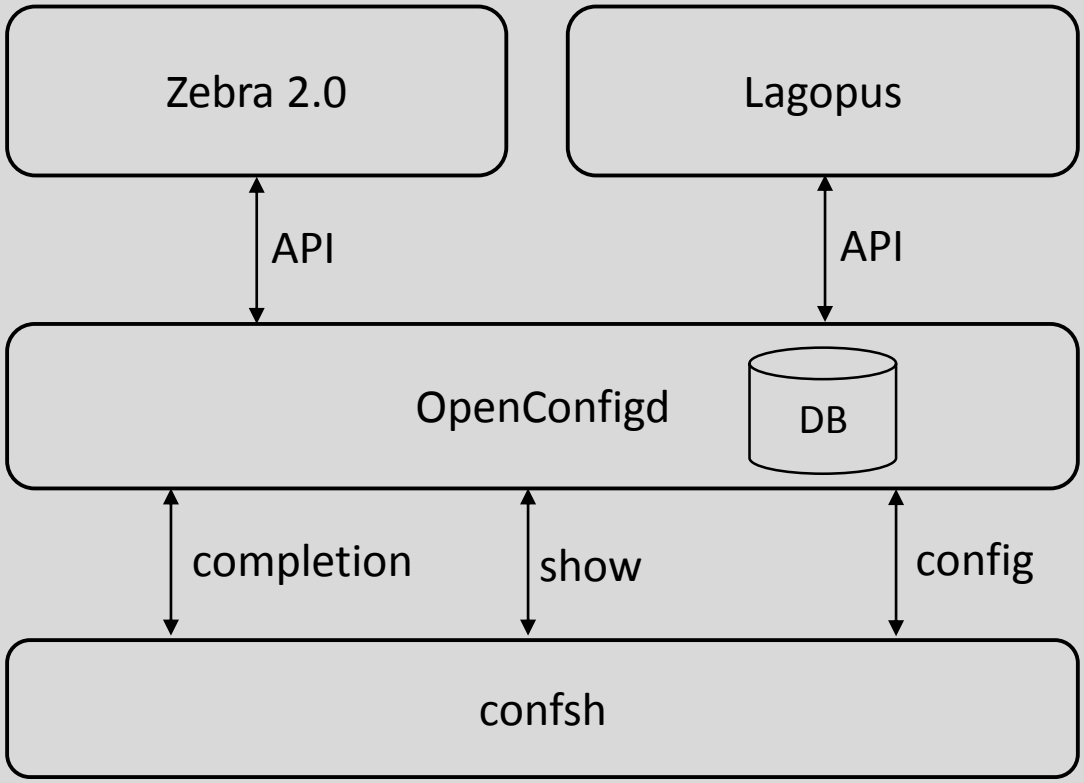

- gRPC is used for transport
- completion/show/config APIs between shell

Language Bindings

- The module is implemented as C++ Class
- Other Language Bindings with SWIG
  - Python
  - Java
  - Ruby etc…

Forwarding Engine Abstraction

- Various Forwarding Engine Exists Today
  - OS Forwarder Layer
  - DPDK
  - OF-DPA
- FEA provides Common Layer for Forwarding Engine
- FEA provides
  - Interface/Port Management
  - Bridge Management
  - Routing Table
  - ARP Table

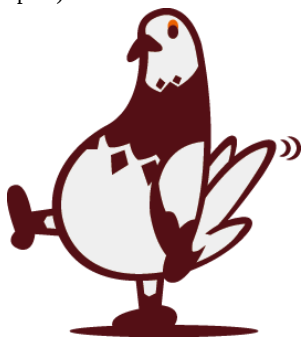

lagopus

Innovative R&D by NTT

Lagopus: Multicore CPUs-aware fast & programmable packet forwarder

https://lagopus.github.io

Lagopus is
a genus of bird in the grouse subfamily

NTT

Lagopus overview (1/2)

Innovative R&D by NTT

- **High performance userspace packet forwarder**
  - Designed for multicore-CPUs architecture
    - Over-40Gbps and over-20MPPS performance
    - Micro-sec order low latency packet processing
  - Designed for large scale RIB/FIB environment
    - Over-1M routes, Over-1M FIB entry are supported
  - Leverage high-performance network I/O library
    - DPDK, packet processing offload/accleratator
- **High-programmable packet forwarder**
  - Many protocol frame matches and actions
    - Ethernet, VLAN, QinQ, PBB, MPLS, IPv4, IPv6, UDP, TCP
  - Tunnel protocol frame match and actions
    - MPLS, IPv4, VxLAN, GRE
    - NSH, Geneve, GTP (in future)
  - QoS actions
    - Meter, Policer, Shaper

NTT

lagopus

Lagopus overview (2/2)

- **Various dataplane control APIs**
  - JSON
  - Netlink-compatible RIB/FIB API
  - OpenFlow 1.3 API
  - SwitchDev API (in future)
- **Switch configuration APIs**
  - JSON, Switch DSL
  - Netconf, OVSDB (Limited)
- **Management protocol and tool**
  - CLI, SNMP, JSON
- **Cooperate with Linux kernel network stack**
  - C-plane packet escalation using tap interface

Lagopus architecture

Innovative R&D by NTT

SDN switch Agent
- Best OpenFlow 1.3 compliant
- Netlink-compatible
- L2/L3 APIs

management API
- OVSDB, REST, OF-CONFIG
- CLI

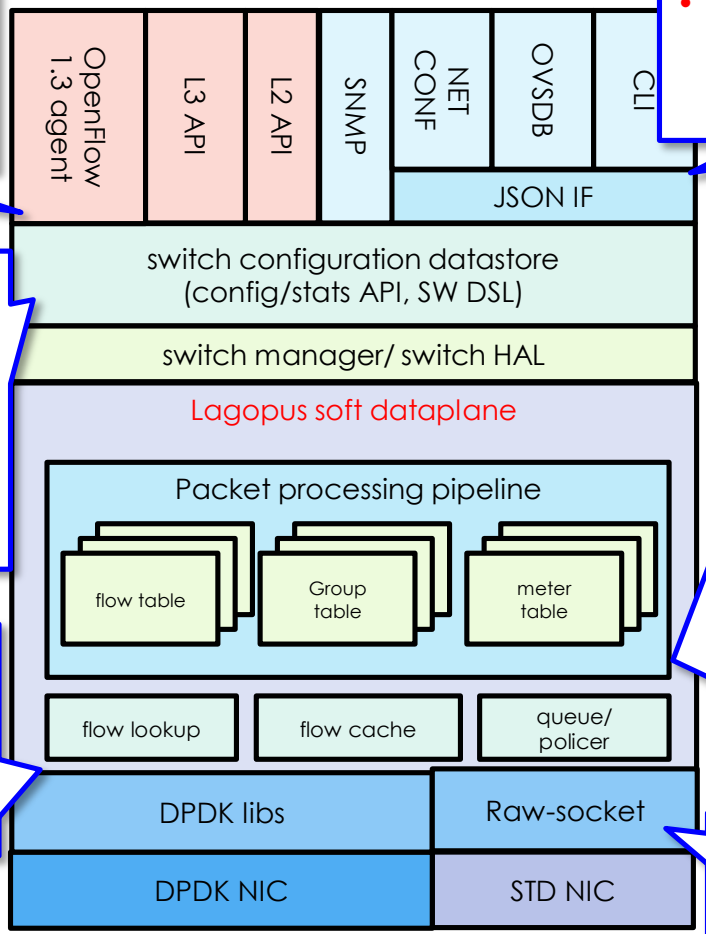

Switch configuration datastore
- Pub/sub mechanism
- Switch config/managment DSL
- Stats database, SNMP
- JSON support

Multicore CPU-aware dataplane
- Over-20-Gbps and 20-MPPS performance
- Low latency packet processing
- high performance multi-layer flow lookup

Cooperation with Linux OS network stack
- Tap support

Various platform and NIC support

NTT

lagopus

Intel Data Plane Development Kit

- **x86 architecture-optimized data-plane library and NIC driver**
  - Memory structure-aware queue, buffer management
  - packet flow classification
  - polling mode-based NIC driver
- **Low-overhead & high-speed runtime optimized with data-plane processing**
- **Abstraction layer for hetero server environments**
- **BSD-license :)**
  - Good for commercial use

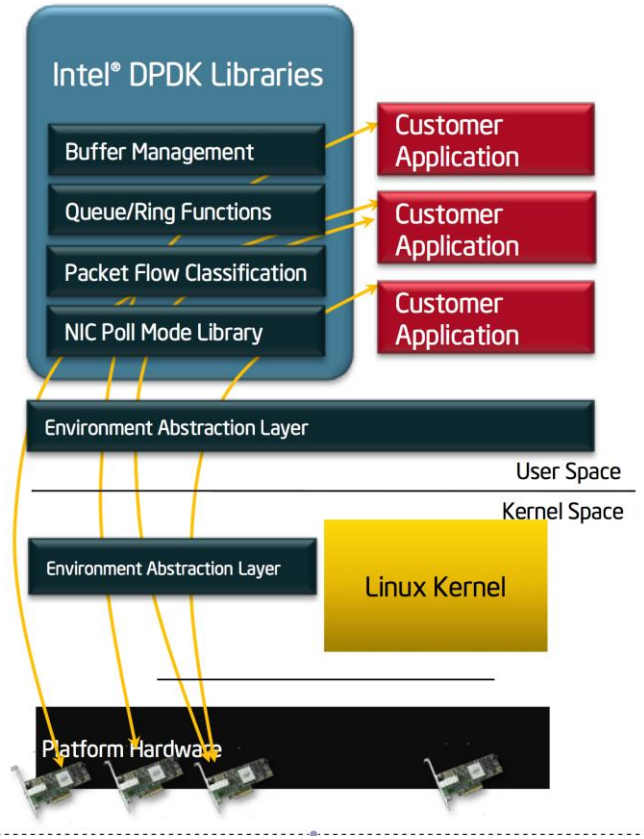

NTT

lagopus

Multi core CPU aware packet processing

Innovative R&D by NTT

- **Exploit many core CPUs**
  - Reduce data copy & move (reference access)
  - Explicit thread assign to CPU core
  - Packet classifier /loadbalancer for parallel processing
- **Decouple I/O processing and flow processing**
  - Improve D-cache efficiency
- **Packet batching**
  - Reduce lock and Improve I-cache and D-cache efficiency

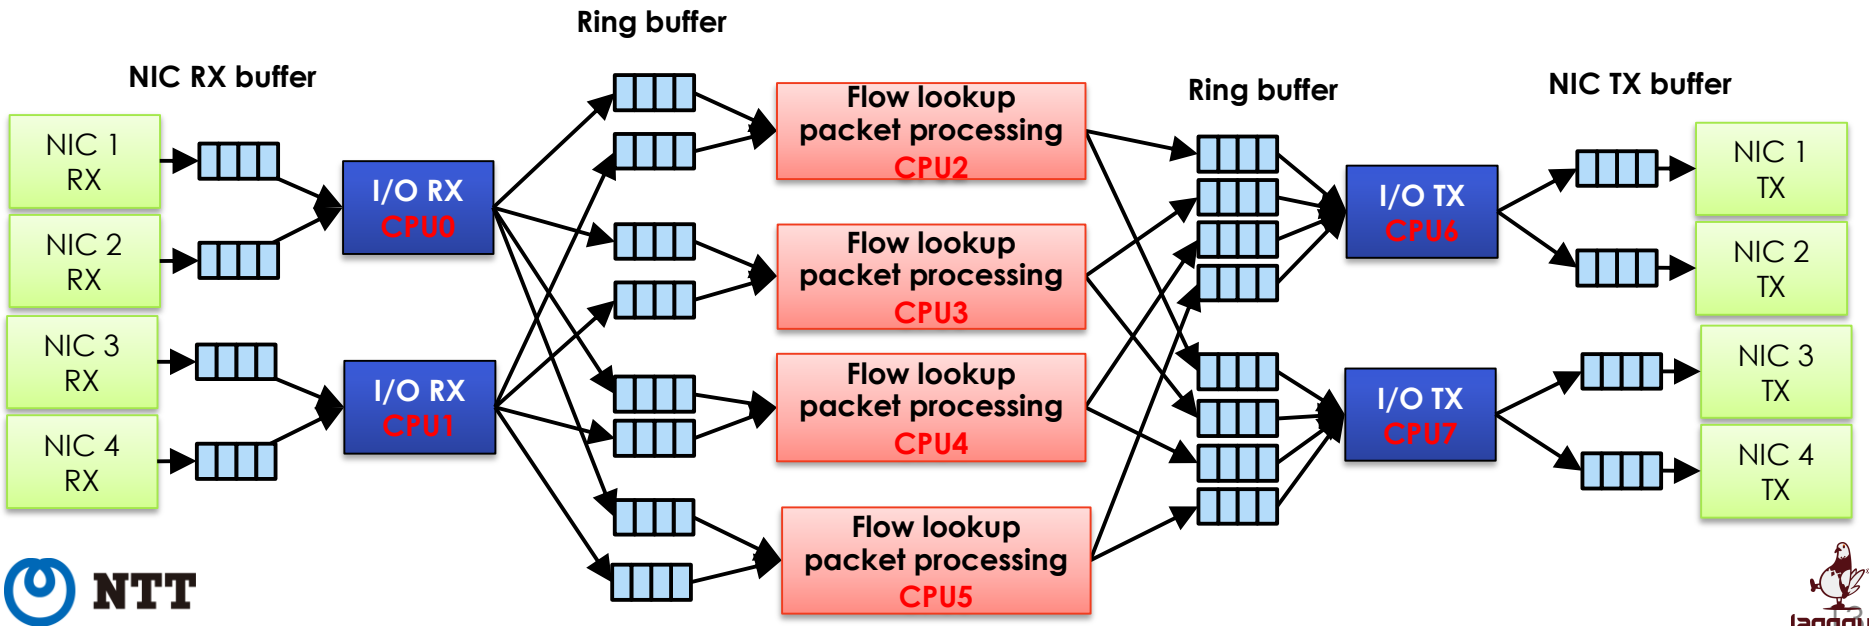

NTT

lagopus

Forwarding Engine Integration

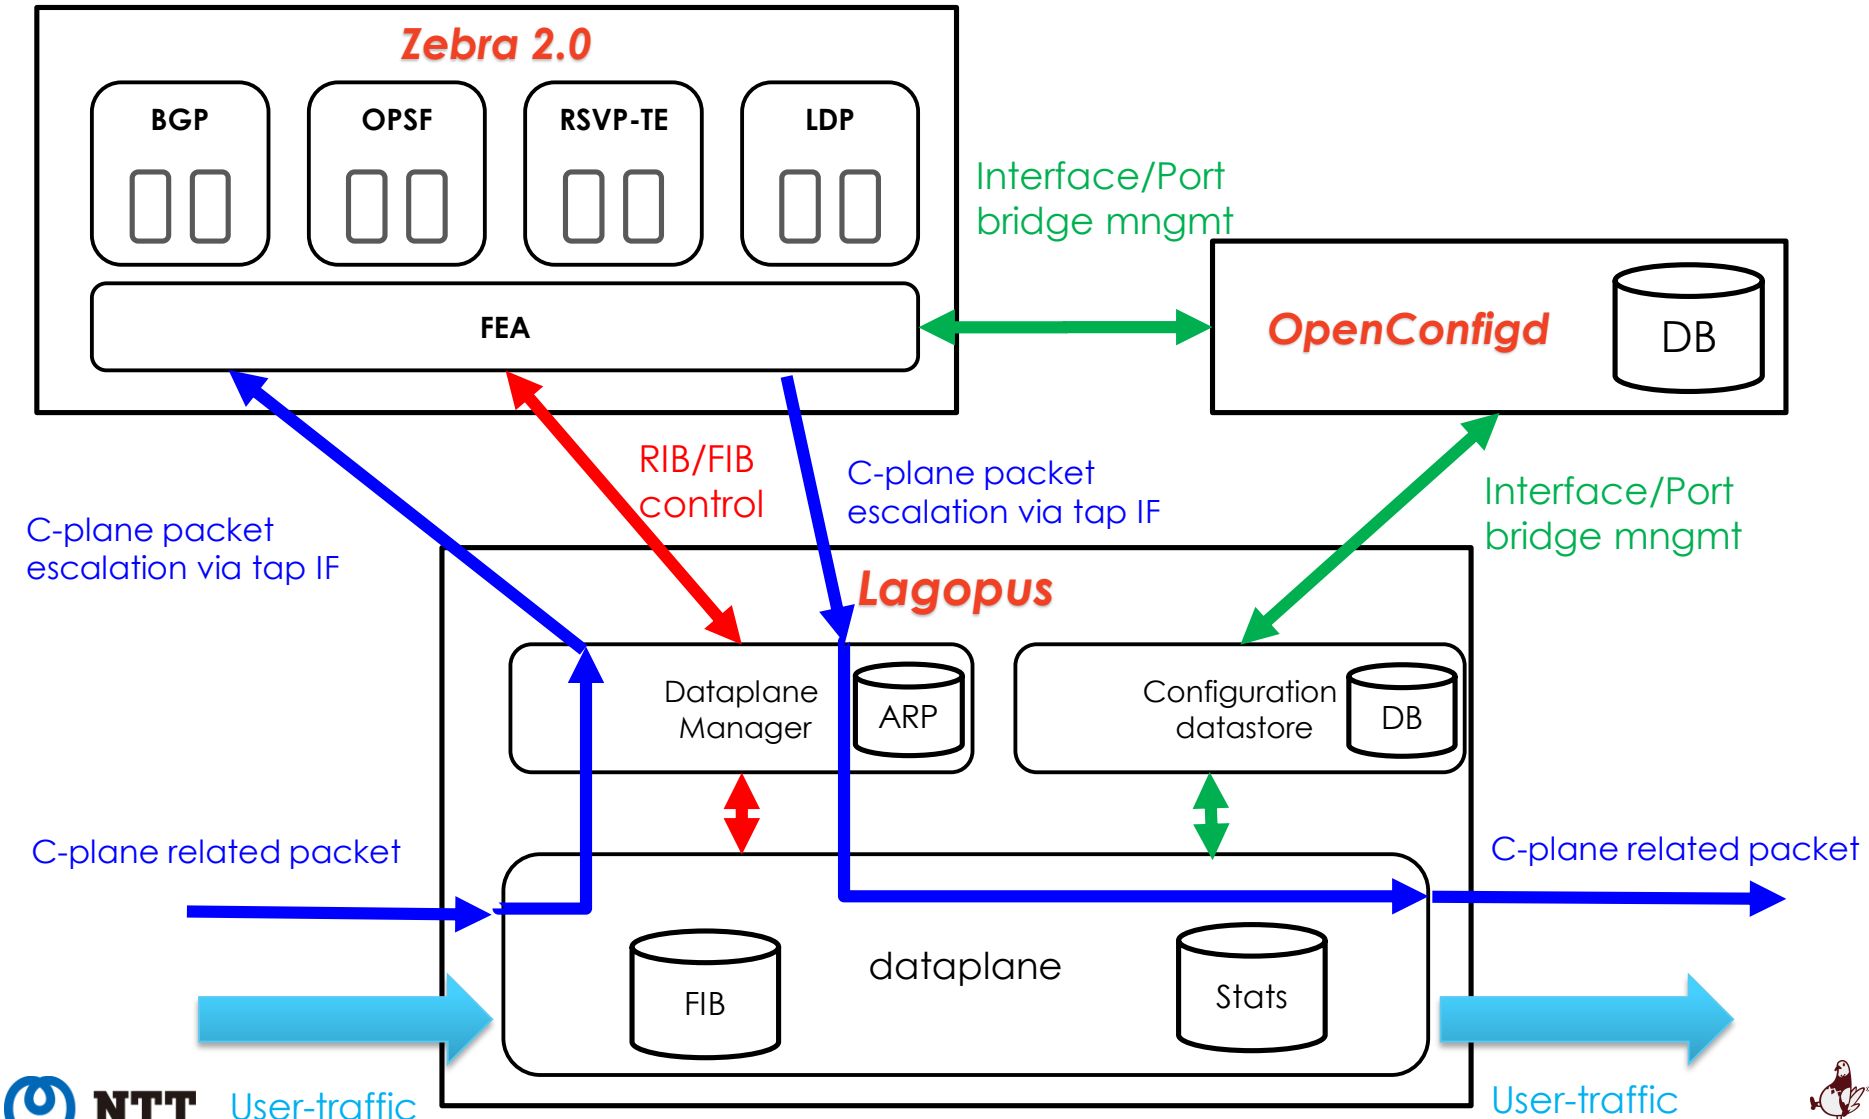

Conclusion

- **Lagopus**
  - https://lagopus.github.io/
- **Zebra 2.0**
  - https://github.com/hash-set/Zebra-2.0 (coming soon)
- **OpenConfigd**
  - https://github.com/hash-set/OpenConfigd (coming soon)