OPEN SOURCE NETWORKING DAYS Accelerate Cloud Native with FD.io Naoyuki Mori, Ping Yu, Kinsella Ray, Hongjun Ni Intel



- FD.io*: Cloud native acceleration framework
- Acceleration of Envoy with FD.io* TCP and QAT
- Acceleration of Load Balancing with FD.io* LB
- Summary



FD.io* VPP Overview

SDN Integration



*Other names and brands may be claimed as the property of others.

FD.io VPP – The "Magic" of Vectors

Compute Optimized SW Network Platform



Packet processing is decomposed into a directed graph of nodes ...





... packets move through graph nodes in vector ...





Vector on Vector for amplified performance (?)

graph node implements a "micro-NF", a "micro-NetworkFunction" processing packets

Makes use of modern Intel® Xeon® Processor micro-architectures. Instruction cache & data cache always hot -> Minimized memory latency and usage.

*Other names and brands may be claimed as the property of others.

Network Function Virtualization to Cloud



FD.io Cloud Native

Strategic Engagements



Envoy - TCP/TLS Integration



Commercial Product

- Cisco Container Platform
- More ...

Cloud Native Technologies



FD.io Cloud Architecture: Contiv

- Can deliver complete container networking solution entirely from userspace.
- Replace all eth/kernel interfaces with memif/userspace interfaces.
- Apps can add VCL library for Higher Performance (bypass Kernel host stack and use VPP TCP stack)
- Legacy apps can still use the kernel host stack in the same architecture



FD.io Cloud Architecture: Ligato



FD.io Cloud Architecture: TCP Acceleration



VPP TCP Host Stack

VPP Hoststack is a high performance L4 implementation in SW.

- Best-in class TCP/UDP Performance in Software; applying 'DPDK' SW optimization techniques.
- Features:
- TCP in client & server (active/passive) modes.
- Common TCP options; MSS, timestamp, wnd scaling, s/ack.
- Common TCP features; ddos protection, delayed ack, congestion.
- Support for Cut-Through, Namespaces, Session Tables and TLS.
- Also support for other Layer 4 protocols such as UDP, SCTP.

New novel approaches to accelerating TCP!

*Other names and brands may be claimed as the property of others.

Istio/Envoy for cloud native

- Industry is moving toward a highly distributed microservices architecture.
- Network should be transparent to applications
- Services are decoupled from each other and communicate with a common service language
- Data plane is composed of a set of intelligent proxies (Envoy) deployed as sidecars. These proxies mediate and control all network communication between microservices, a generalpurpose policy and telemetry hub.
- The control plane manages and configures the proxies to route traffic.







Re-architect Envoy network stack with Intel ingredient



Transparent TLS

- TLS engine between App session and transport layer.
- TLS is transparent to application just like one of session protocol such as TCP, UDP or SCTP.
- TLS application registers key and certification via API and requests TLS as session transport. Besides that, every goes like TCP.





Building a Scalable Crypto Engine

- Vector Packet Processing makes TLS transparent to the application
- QAT Accelerates the Crypto Algorithm
 - Quick Assist Technology
 - QAT Intel hardware acceleration engine to offload crypto
- Asynchronous enabling to maximize HW performance



Network Functions with FD.io VPP *

Discrete Appliances, VNFs & CNFs

Universal	Data Plane	Control Plane
CPF	L2 Switch	Netconf/Yang
	VLAN / Q-in-Q	BGP
<i>∎ "))</i>	L3 Router	TR-069
	NAT	IKEv2
	ACL (mac, ip, port)	
	IPSEC	

Cloud	Load
Bala	ncer



Data Plane	Control Plane	
Bonding	OpenStack LBaaS	
VLAN / Q-in-Q		
NAT	Kes Kuba Drawy	
ACL (backlist)	Kas Kube-Proxy	
TM (Policing, Metering)		
L4 Load Balancer		

Broadband	Data Plane	Control Plane
Network	L2 Switch	Netconf/Yang
Gateway	L3 Router	BGP
Gutemay	Classification	
	hQoS	
	ACL	
	TM (Policing, Metering)	
Intrusion	Data Plane	Control Plane
Intrusion Prevention	Data Plane	Control Plane Netconf/Yang
Intrusion Prevention System	Data Plane L2 Switch L3 Router	Control Plane Netconf/Yang BGP
Intrusion Prevention System	Data Plane L2 Switch L3 Router Classification	Control Plane Netconf/Yang BGP
Intrusion Prevention System	Data Plane L2 Switch L3 Router Classification NAT	Control Plane Netconf/Yang BGP

External Load Balancer



VPP Service Proxy

- Router, Load Balancer and Service Proxy are supported on VPP.
- On Router, will enable ECMP feature.
- VPP Load Balancer distributes traffic and encapsulates packets via GRE tunnel.
- On K8s node, it removes GRE tunnel and goes through Service Proxy to distribute traffic to chosen pod.



Multithread Support: Ingress traffic



- RSS enables traffic associated with one connection to a given thread.
- Load balancing and connection track redirects traffic to a chosen pod.

Check out Yahoo! JAPAN* VPP LB use case session!





- FD.io* is robust and commercially deployed networking stack for cloud native acceleration
- Intel IA platform is devoted to contribute and accelerate FD.io*