100 Gbps Open-Source Software Router? It's Here.

Jim Thompson, CTO, Netgate

@gonzopancho

THE LINUX FOUNDATION

Agenda

- Edge Router Use Cases – ‘Need for Speed’
- Cost, Flexibility, Control, Evolution
- The Engineering Challenge
- Solution Components
- Test Configuration & Results
- Productization
- Value Proposition
- Vendor Value Add

100 Gbps Edge Router Use Case #1

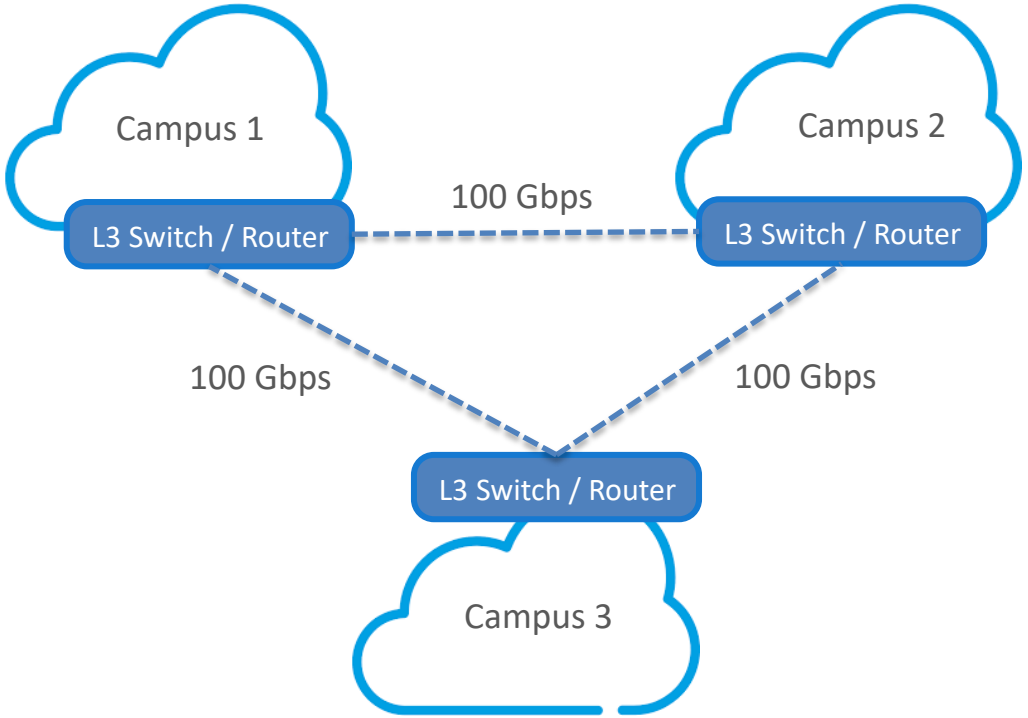

Need

- Campuses equipped with 100 Gbps white box h/w and fast NICs
- Secure routing software cannot keep up
- Limits campus interconnect to <10 Gbps of IMIX traffic

100 Gbps Edge Router Use Case #2

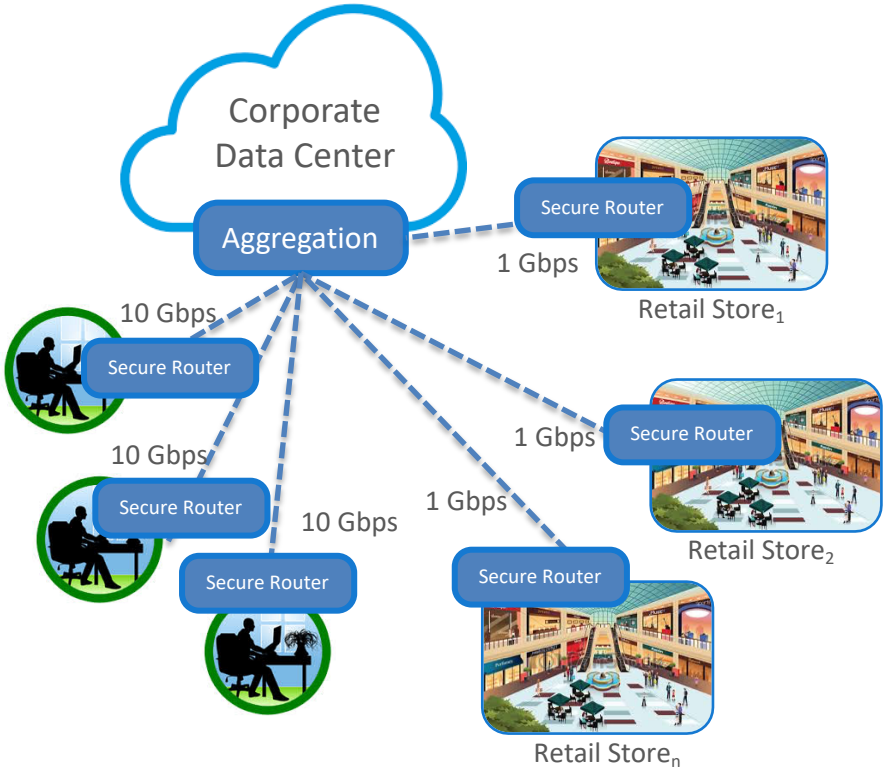

Need

- 1 – 10 Gbps secure routing needed for e-commerce, content access and developer productivity
- Big brand solution forces unnecessary features, and is expensive
- Customer wants API-based ACL mgmt

THE LINUX FOUNDATION

100 Gbps Edge Router Use Case #3

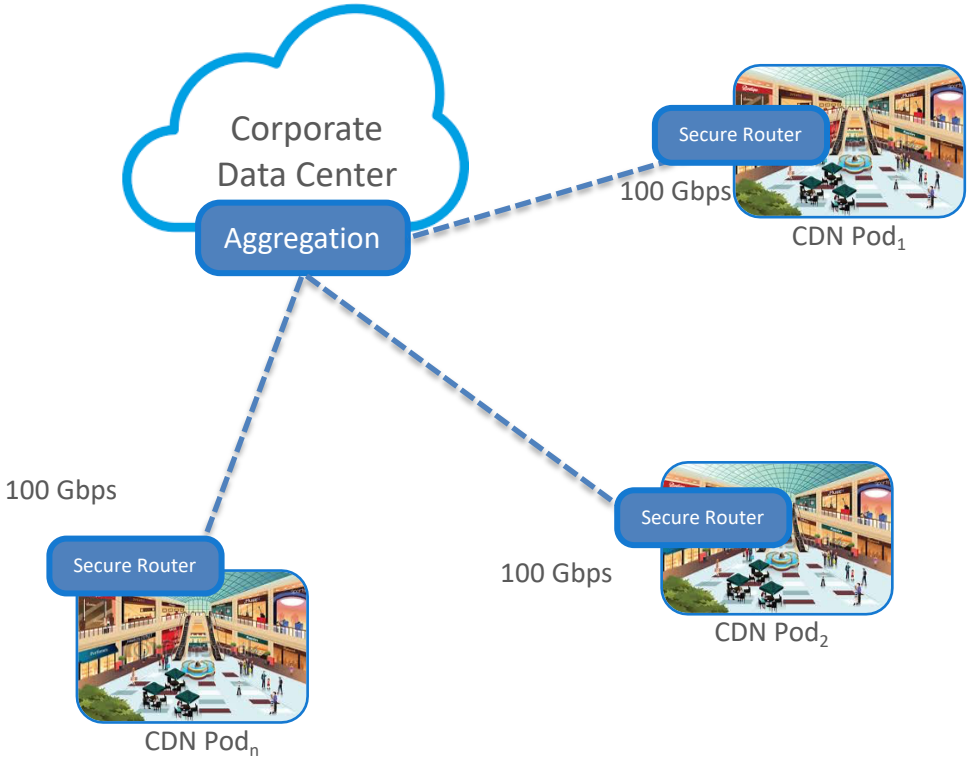

Need

- 100 Gbps large packet, IPSec content distribution
- Fast adaptation to network bandwidth costs
- Reduced per customer delivery costs to drive profit from competitive MRR

Evolution of Cost, Flexibility, & Control

$$$
Vendor Control

- Vendor Proprietary H/W
- Specialty Silicon
- Vendor Proprietary S/W
- Vendor Proprietary Mgmt

2005 - 2014

$$
Vendor Control

- **Whitebox H/W**
- **x86 Silicon**
- Vendor Proprietary S/W
- Vendor Proprietary Mgmt

2016-2017

$
User Control

- Whitebox H/W
- x86 Silicon
- **Open Source High Performance S/W**
- **Open Source Orchestration Mgmt**

2018+

THE LINUX FOUNDATION

Drivers | Enablers

Drivers

- Rapid movement of workloads and data to cloud
- Virtualization of everything
- Buyer demand for network architecture freedom

Enablers

- High powered, inexpensive commodity silicon
- Open source Vector Packet Processing (VPP)
- Open source Restful API + orchestration management

The Engineering Challenge

How do you fill a 100 Gbps pipe with small packets using s/w?

10 Gbps Line Rate

- Shortest possible Ethernet packet
  - 46 byte payload
  - 18 byte Ethernet header & CRC
  - 20 byte preamble, start-of-frame delimeter + inter-frame gap (IFG)
  - 84 byte total (672 bits)
- True line-rate is 10,000,000,000 bits per second, 672 bits at a time
- **14,880,952 packets per second**
- 67.2 ns to process one packet
- A 2 GHz clock CPU core → 1 core clock cycle of 0.5 ns
- 67.2/.5 = 134 CPU clock cycles/packet (CPP)

40 Gbps Line Rate

- Per packet processing budget is 67.2/4 = 16.7 ns
- Equals 134/4 = 33.5 CPP
  - Receive the packet on a given interface
  - Process the packet
  - Transmit the packet out a (presumably) different interface

100 Gbps Line Rate

- Per packet processing budget is 67.2/10 = 6.7 ns
- Equals 134/10 = 13 CPP
- **148,809,523 packets per second**
- **Secure routing s/w cannot get there with kernel-based, single packet per time processing**
- **Enter Vector Packet Processing (VPP)**

Key Solution Components

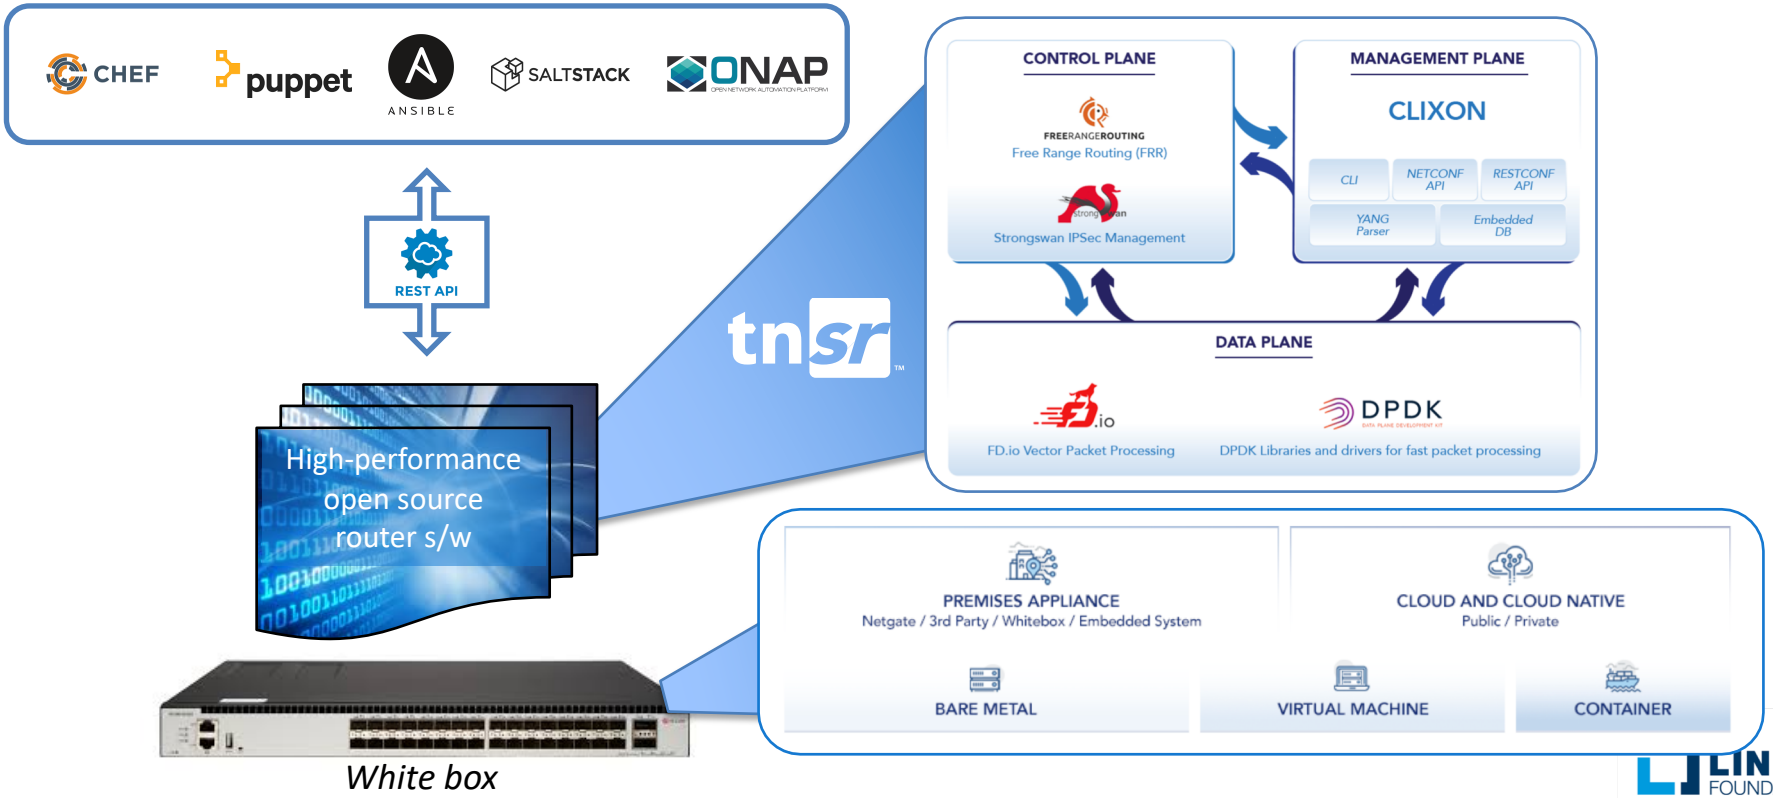

White box

Key Software Components

TNSR: Open Source-based Secure Networking Software Platform

Routing Stack

Intelligence as routing protocols to forward packets as well as a control plane

Forwarding Plane

Raw muscle to get packets through the system

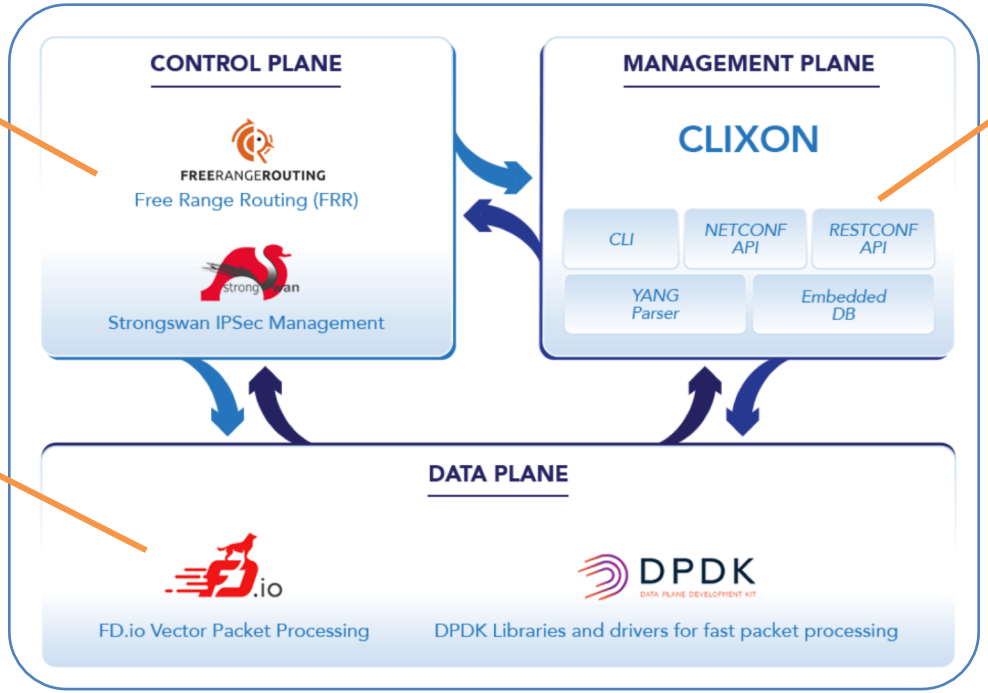

Discrete Mgmt

Orchestration Management Interface enables external policy configuration and visibility

THE LINUX FOUNDATION

Test Configuration

- Xeon Gold 6130
- 32GB RAM
- Mellanox 100 Gbps NIC
- 1U server

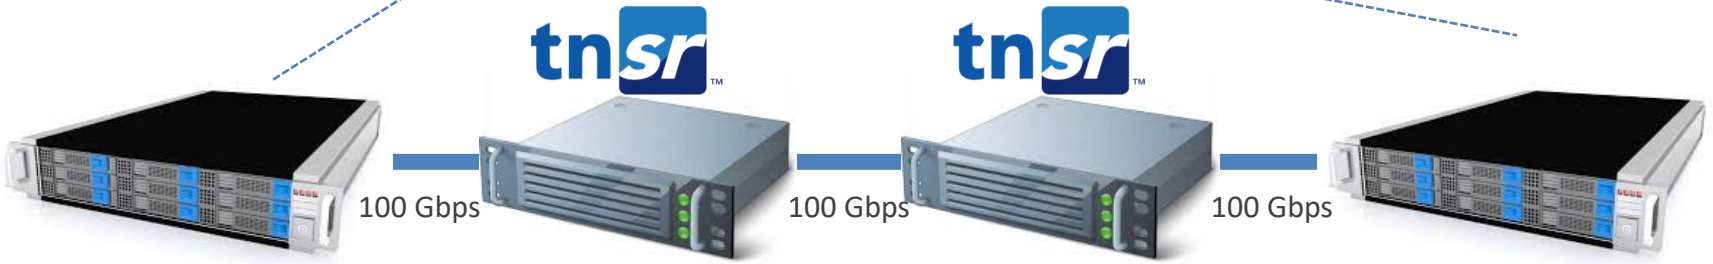

- I7-6950X CPU, overclocked to 3.5GHz (10C)
- ASUS X99 board
- Water-cooled to avoid thermal throttling
- 32GB RAM
- Mellanox 100 Gbps NIC
- Intel DH8950 (Coletto Creek) QAT

Packet Processing Test Results

Mpps				
# Streams / Cores	1/1	2/2	8/8	8/256
Stream Type				
64 byte	14.1	28.3	67.6	56.3
64 byte; AES-CBC-128	4.5		10.7	
64 byte; AES-GCM	4.3		9.0	
128 byte	12.6	22.6	58.0	48.0
128 byte; AES-CBC-128	4.5		10.5	
128 byte; AES-GCM	4.2		8.9	
256 byte	6.3	11.2	30.7	34.1
256 byte; AES-CBC-128	3.7		10.0	
256 byte; AES-GCM	3.5		8.6	
512 byte	6.0	12.0	23.5	21.7
512 byte; AES-CBC-128	3.5		7.7	
512 byte; AES-GCM	3.4		8.6	
1380 byte; AES-CBC-128	3.5		3.5	
1380 byte; AES-GCM	2.8		3.8	

Gbps				
# Streams / Cores	1/1	2/2	8/8	8/256
Stream Type				
64 byte	9.48	19.02	45.43	37.83
64 byte; AES-CBC-128	3.02		7.19	
64 byte; AES-GCM	2.89		6.05	
128 byte	14.92	26.76	68.67	56.83
128 byte; AES-CBC-128	5.33		12.43	
128 byte; AES-GCM	4.97		10.54	
256 byte	13.91	24.73	67.79	75.29
256 byte; AES-CBC-128	8.17		22.08	
256 byte; AES-GCM	7.73		18.99	
512 byte	25.54	51.07	100.02	92.36
512 byte; AES-CBC-128	14.90		32.77	
512 byte; AES-GCM	14.47		36.60	
1380 byte; AES-CBC-128	39.20		39.20	
1380 byte; AES-GCM	31.36		42.56	

- **100 Gbps at 512 byte frames**
- With these limits…
  - 10 cores
  - 40 Gbps QAT card
  - Single socket i7
- Lab cost (S/W not incl.)…
  - $2,000 box (mostly CPU)
  - $700 NIC
  - $650 CPIC (QAT) card
- …and the H/W can be boosted
  - 2/4 socket boxes w/ NUMA
  - More cores
  - Faster NICs

Productization

White Box Edge Router

Cloud VM Router

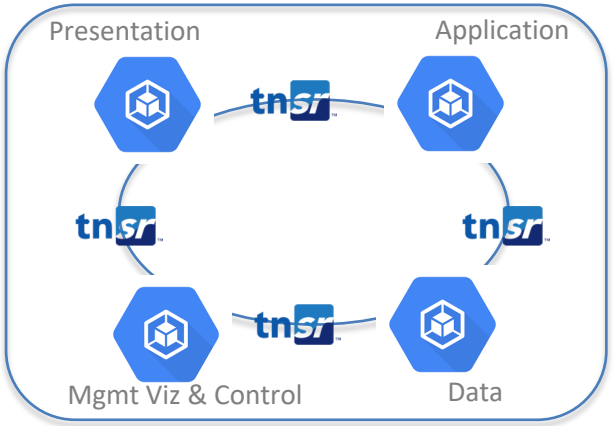

Cloud Container Interconnect

Speed, Scale, Flexibility – when and where you need it.

Value Proposition

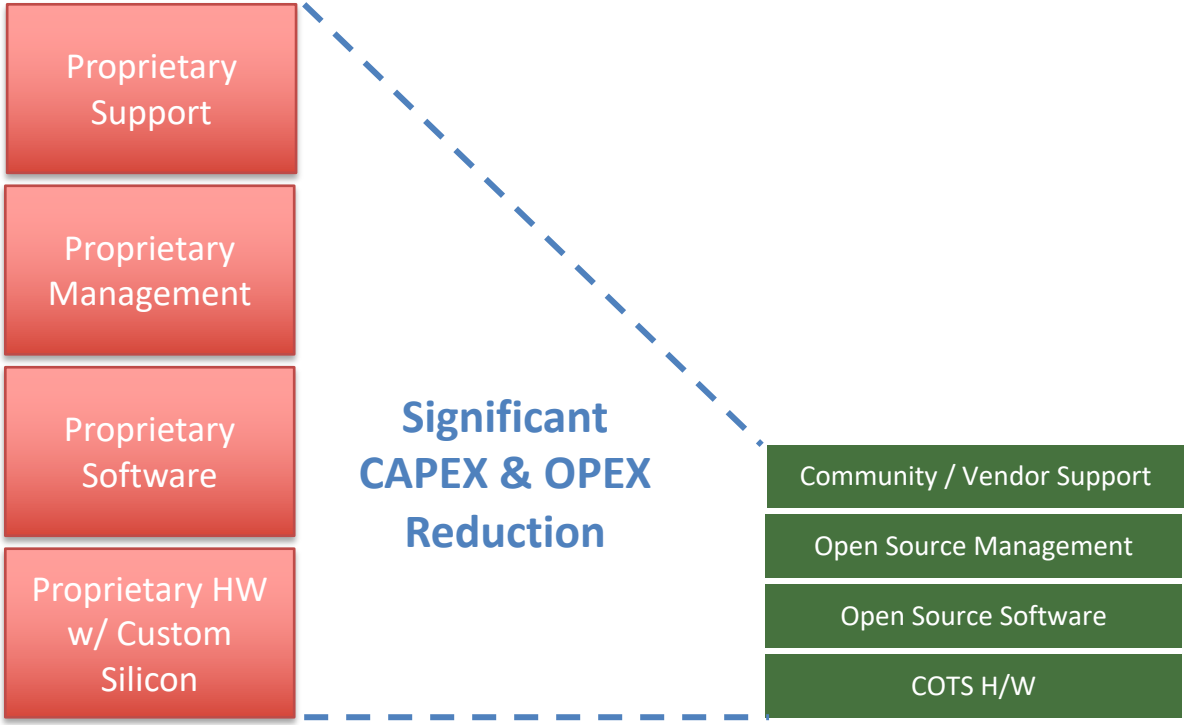

- 40 Gbps IPsec
- 100 Gbps Routing
- $7,500 / instance
- Including support

THE LINUX FOUNDATION

Who Needs an Open Source Vendor?

FLEETWOOD MAC

"You can go your own way…"

- 'Free like a puppy'
- All software evolves, has issues
- Requires understanding, software integration development, productization, testing, verification, packaging, distribution, support
- Is it worth your time?

THE BEATLES

"Won't you please, please help me…"

- Open source software productization vendor
- Moved from project to product
- Development, productization, testing, verification, packaging, distribution, support provided
- Fraction of the cost of proprietary vendor model

THE LINUX FOUNDATION

Summary

- Superfast, highly scalable, open-source based router
- Fast, easy and cheap to reconfigure or add new services
- At a fraction of today's prices
- Utility-based secure networking affordable for all
- Open source vendors make it ready for prime time

THE LINUX FOUNDATION

OPEN SOURCE SUMMIT

EUROPE

THE LINUX FOUNDATION