

# Disaggregation of Cloud Networking via SmartNICs

Gerald Degrace

SONiC Pioneer & Co-DASH Creator

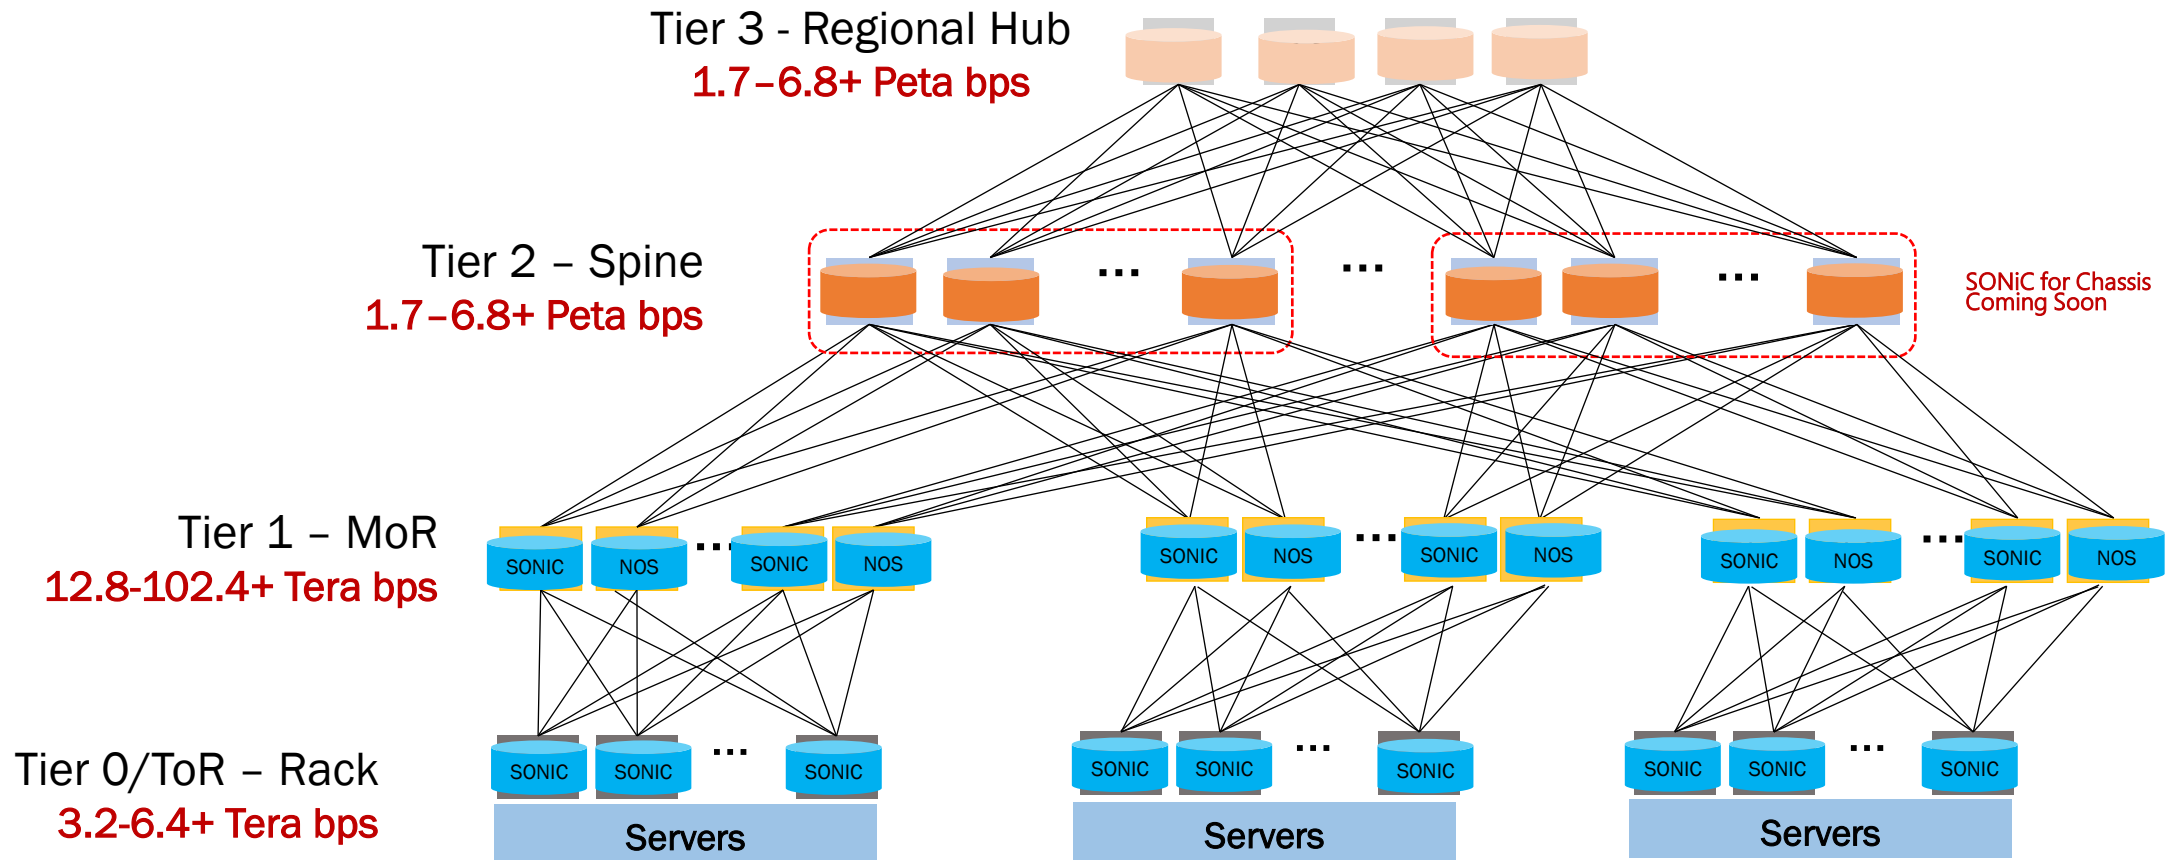
Microsoft Azure

# Introducing – SONiC (The foundation for DASH)

## **SONiC (Software Open Source for Networking in the Cloud)**

- Created by Microsoft and its partners then migrated to Linux Foundation
- Removes the requirement to run the switch suppliers NOS
- Containerized solution offering consistent operations, maintenance, monitoring, telemetry streaming, and security and more
- Running on 100s of 1000s of SONiC switches powering Azure's Cloud
- **Bringing 161+ Network Operating Systems -> 1 NOS**

# Azure's Cloud is composed on SONiC enabled Switching (Constructed from “any” Switching Hardware Platforms)



# SONiC Community Partners

## Merchant Silicon



## Switch Platform



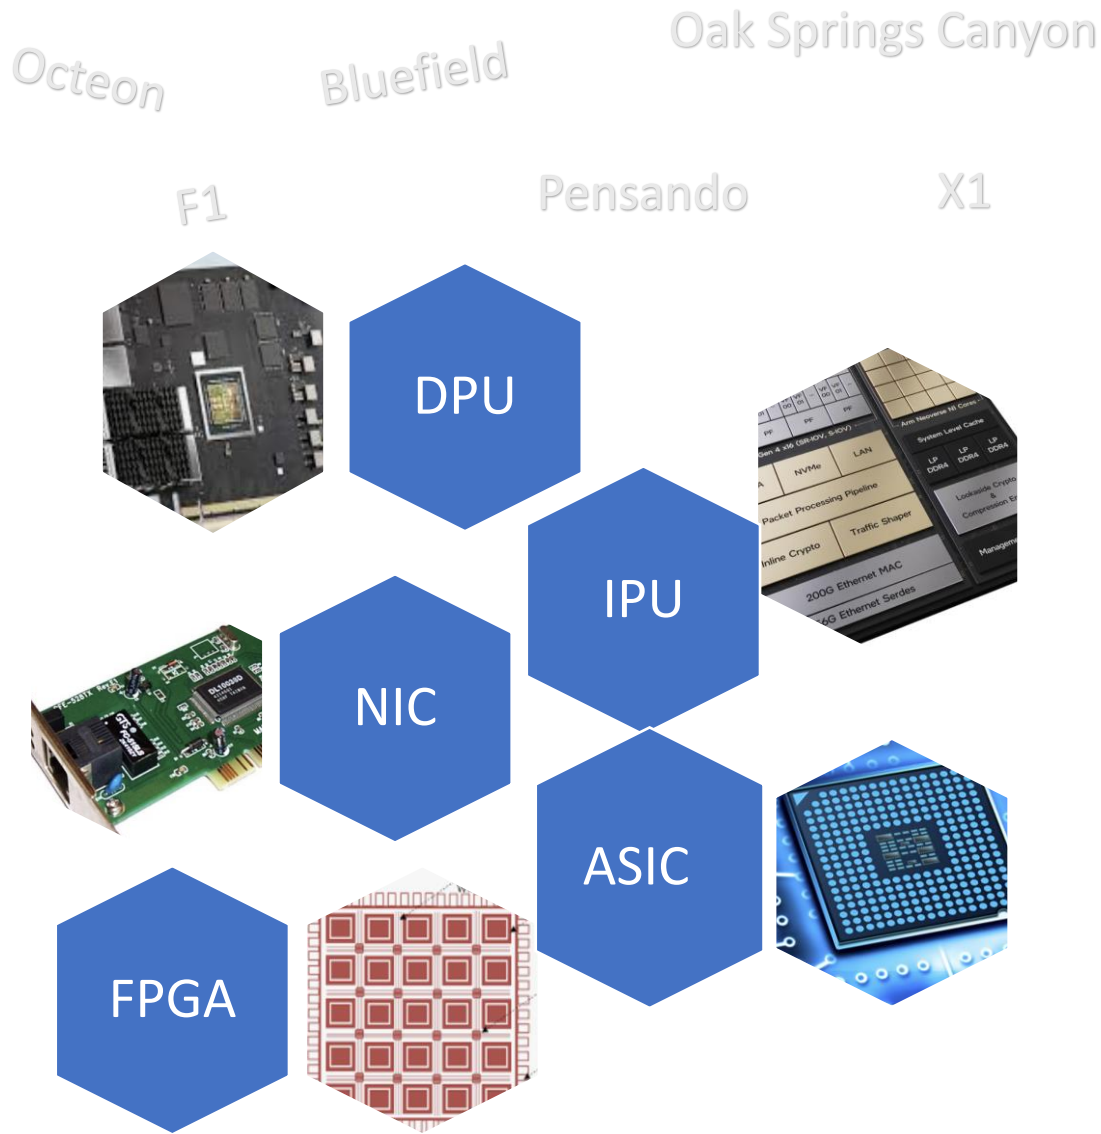
## Adoption




## System/Service



# So What's Next and Exciting ? – (Cloud based Connection Management)



Connectivity 200G - 800G  
PPS: 400M+ Packets per Second



3nm - 5nm low power  
P4 programmable or library based



256M+ connections at **Millions of Connections Per Second**



Thousands of rules per connection  
100s of Gbps+ Crypto/Security  
IPSec: 16M+ connections



IOPs: >40M/s

# Introducing – Microsoft Inspired DASH

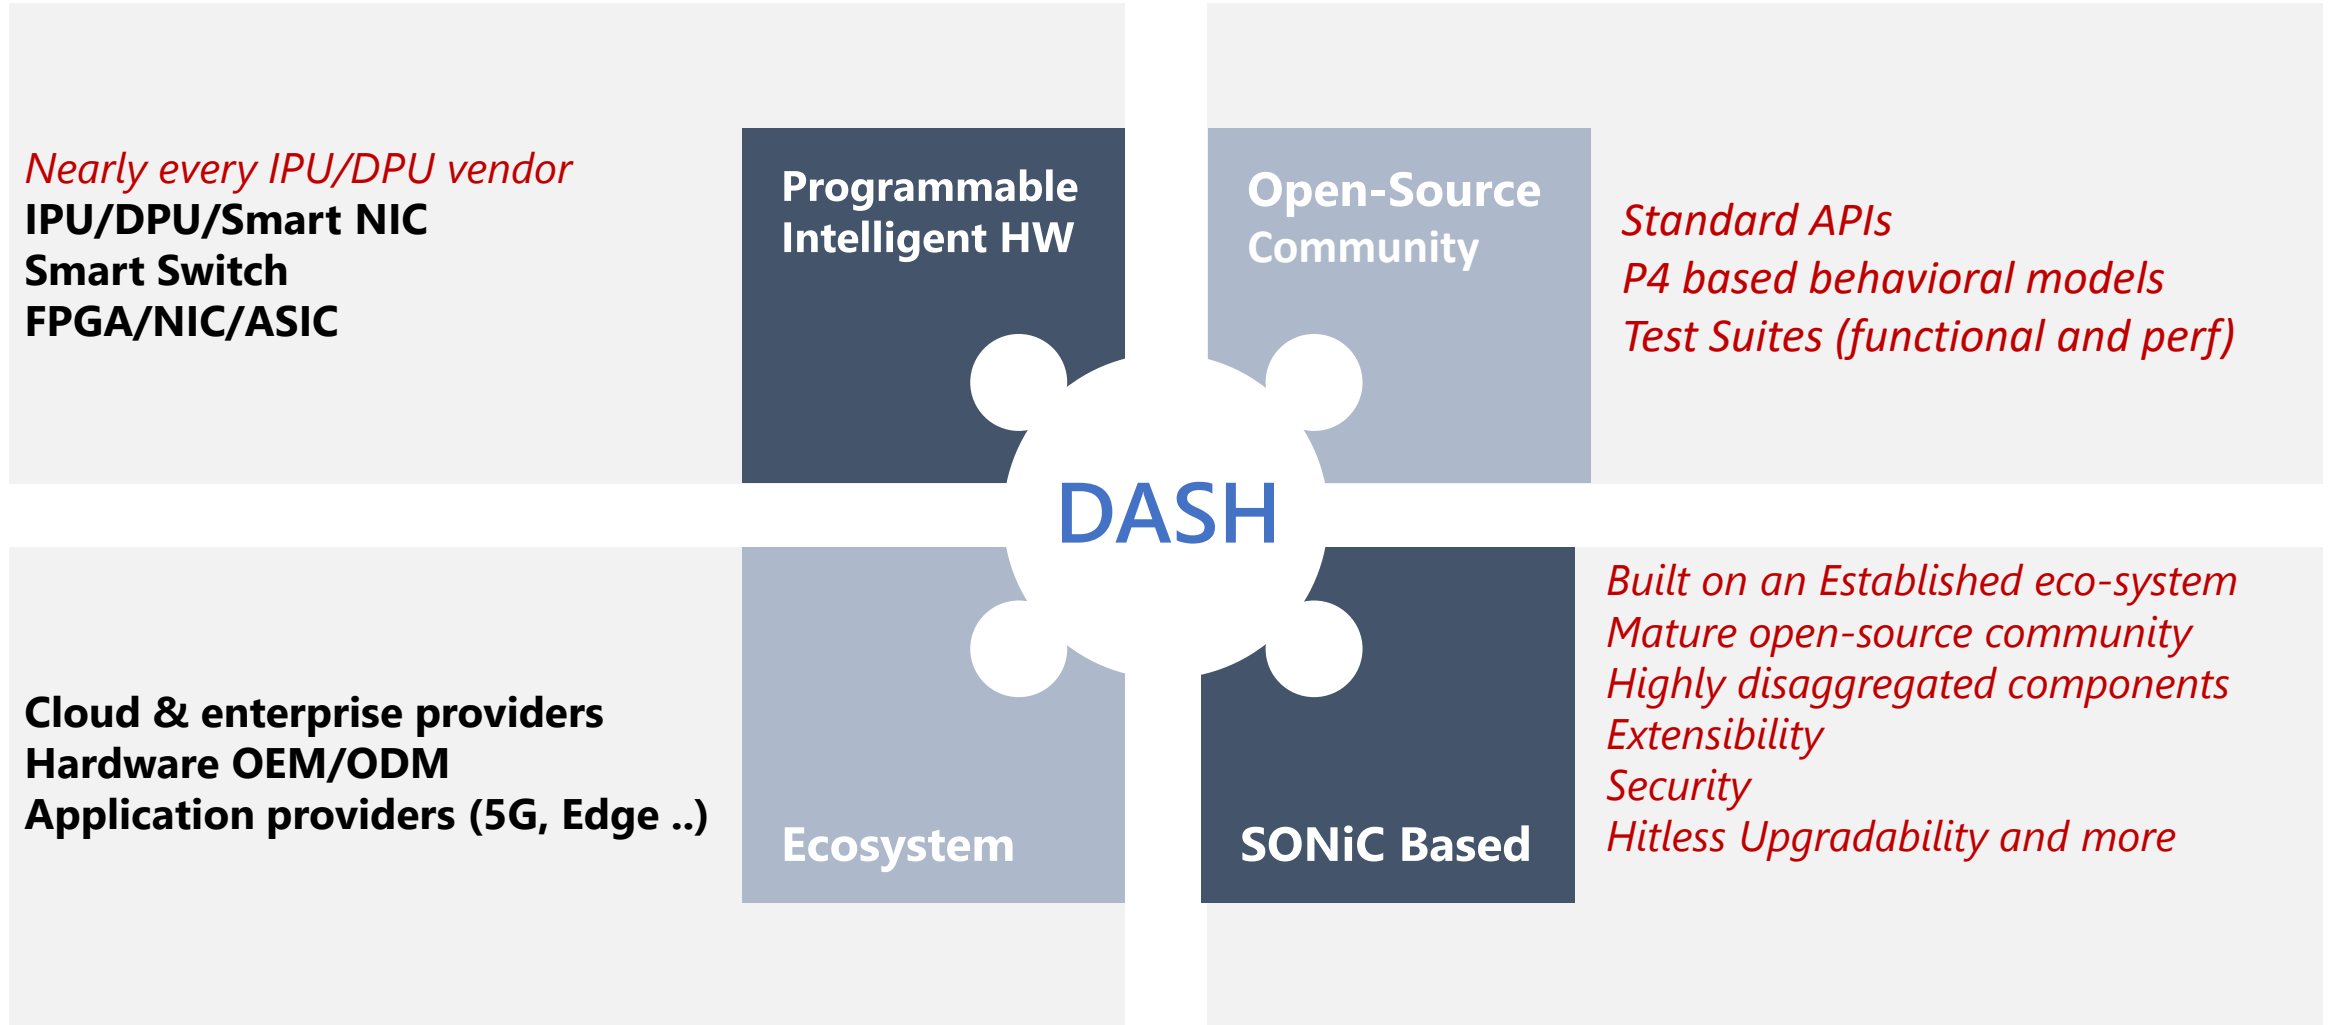
**DASH** (**D**isaggregated **APIs** for **SONiC** **H**osts)



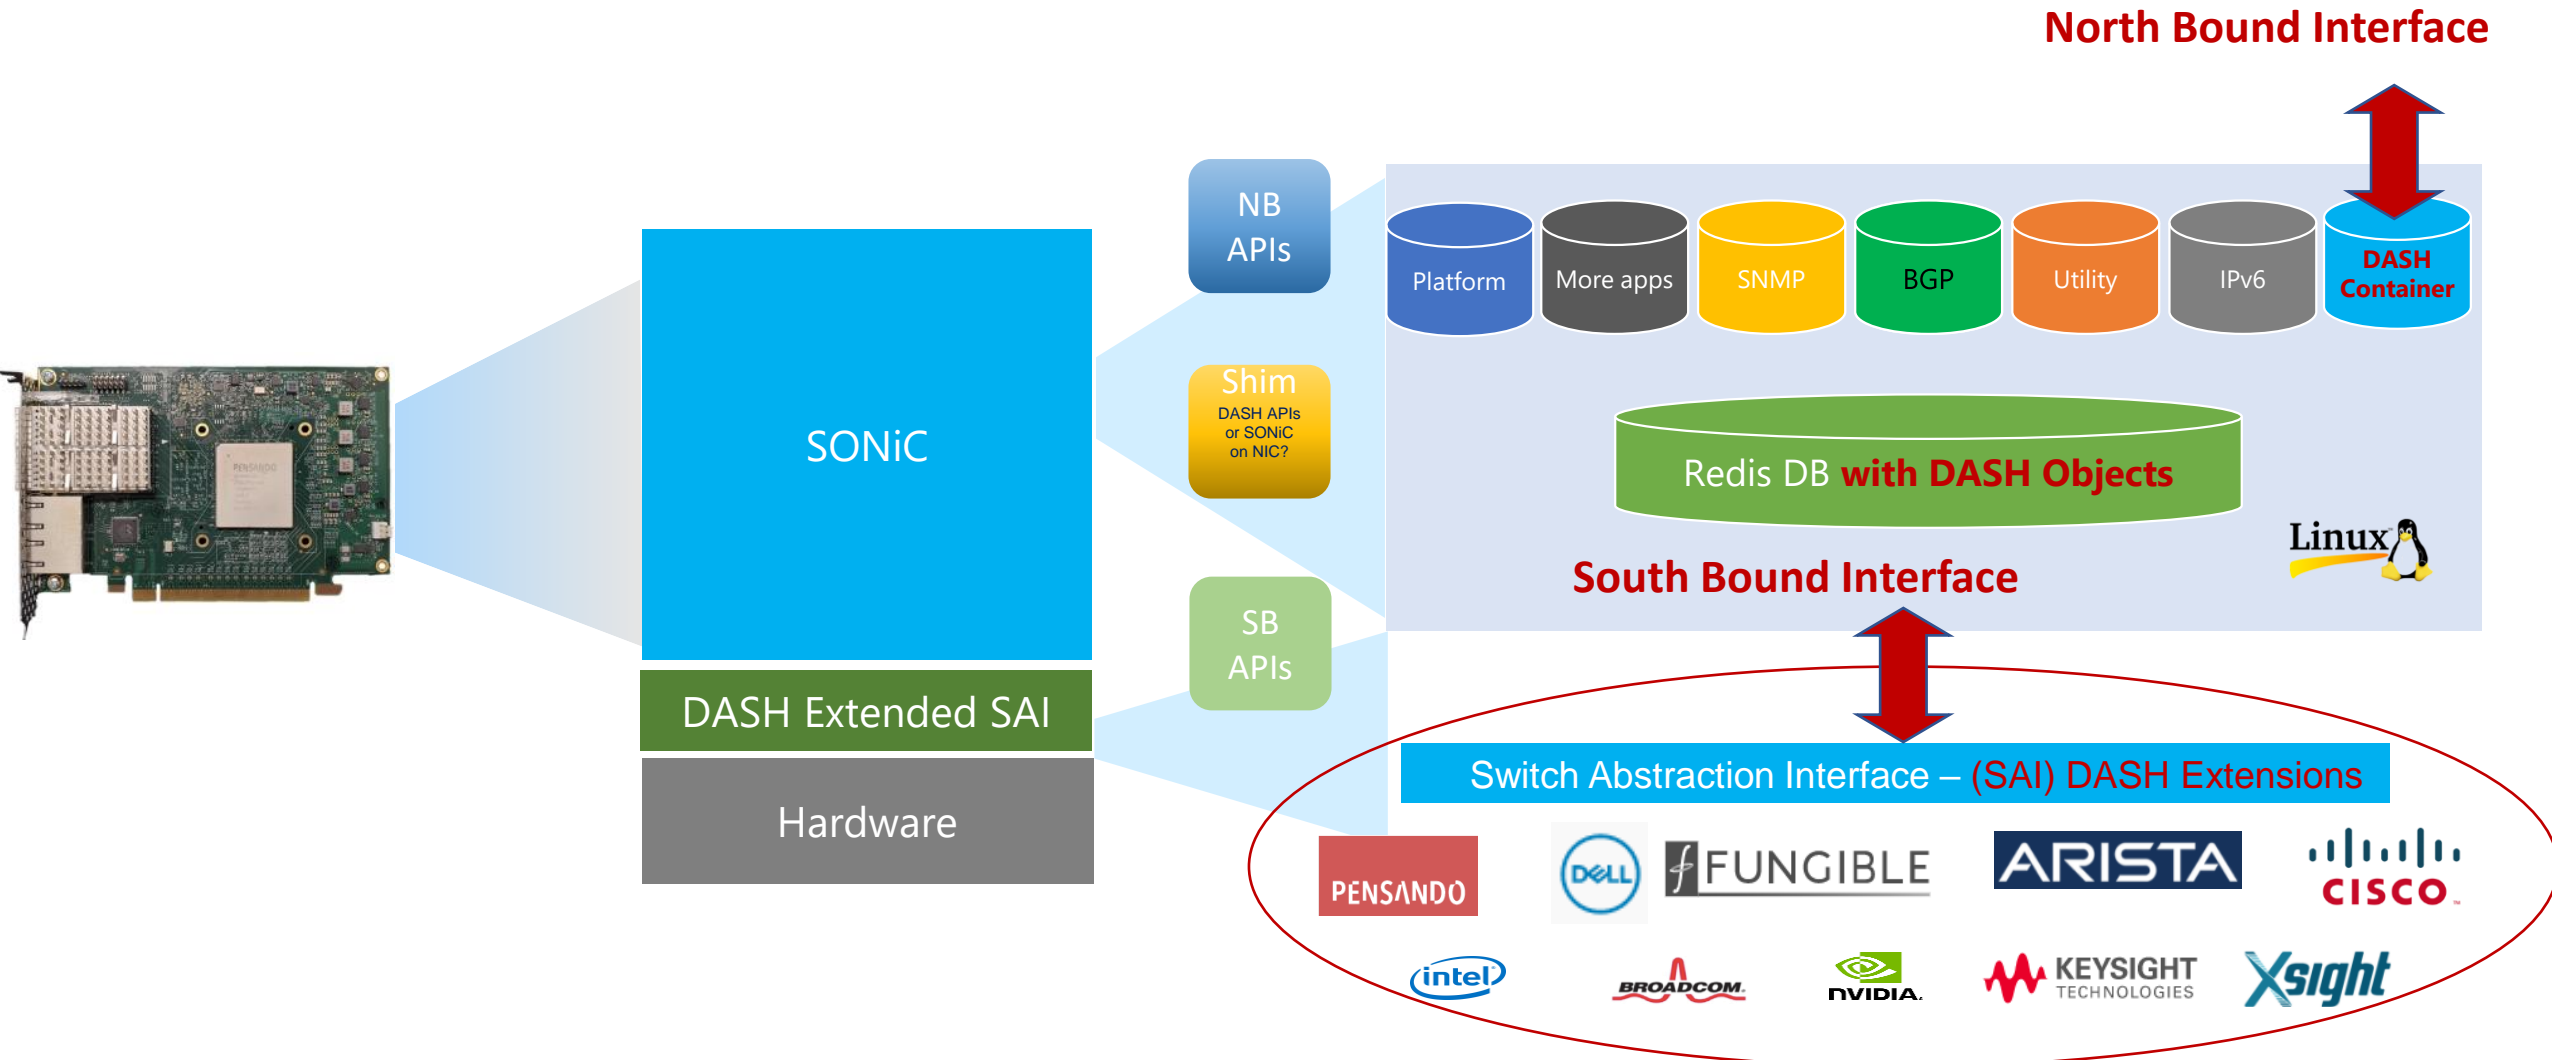
- An open source project now part of Linux Foundation
- To achieve 10-**100x perf** in **cloud connection management** scale than offered in today's Cloud
- By maximizing innovation across Switch and DPU/IPU in a complementary manner where 1+1 is even greater than 3

# SONiC DASH Components

(100x Cloud Networking and Connection Management Performance over Software Only)



# SONiC DASH – Some New Players <-> Many SONiC Pioneers





# Let's start with an Cloud Networking Example

(10s of thousands of connections/sec)

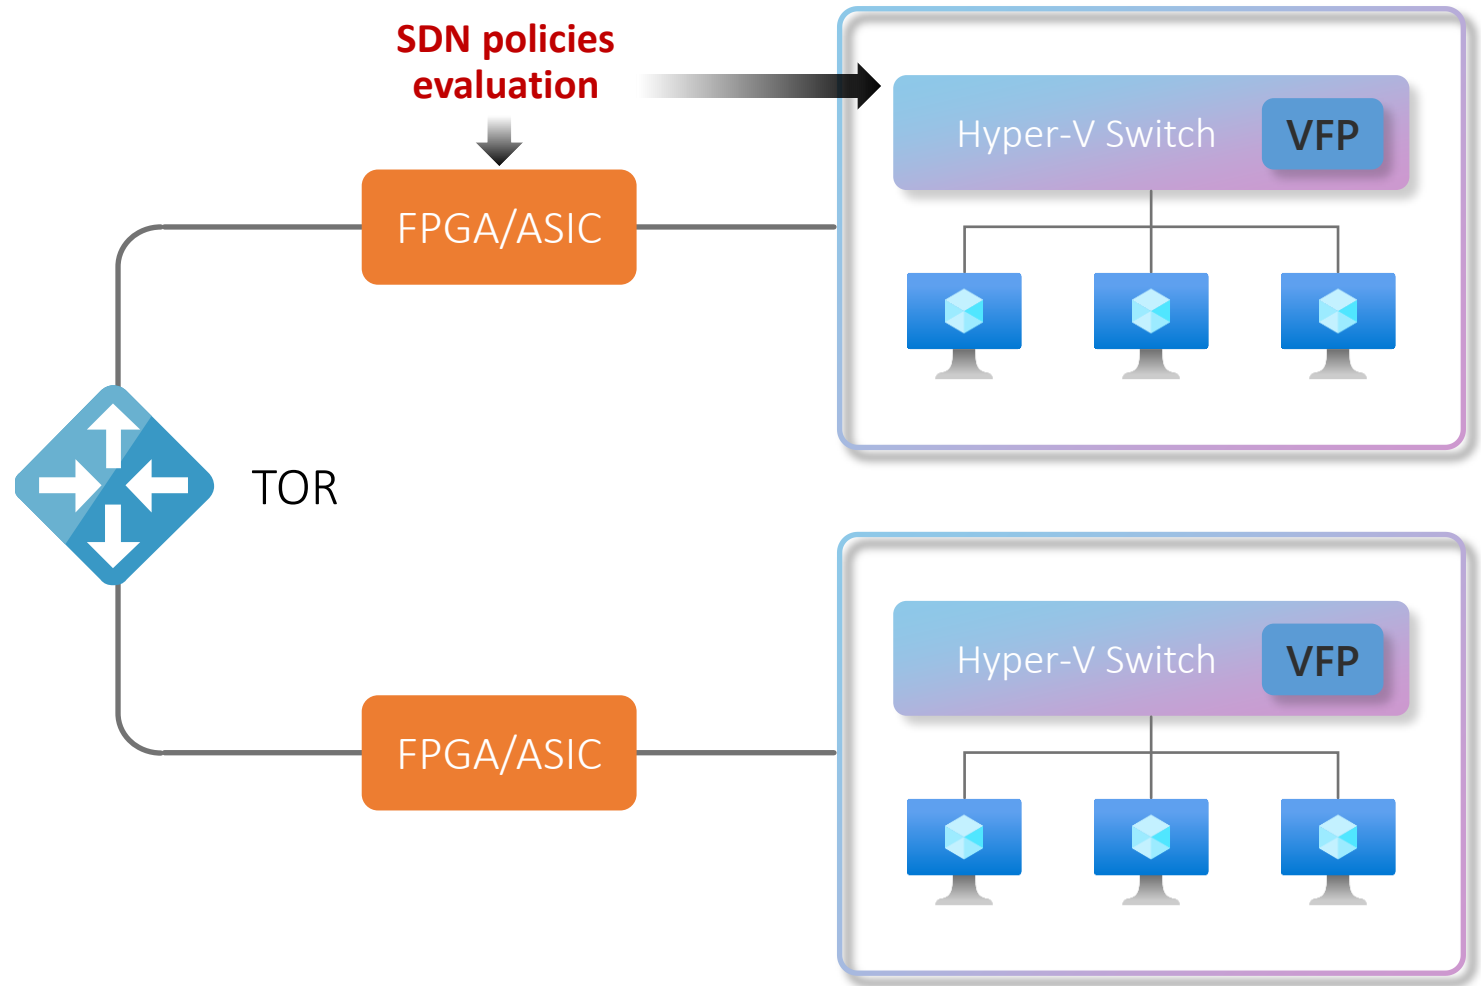
Network policy processing today is performed on host node in virtual switch

---

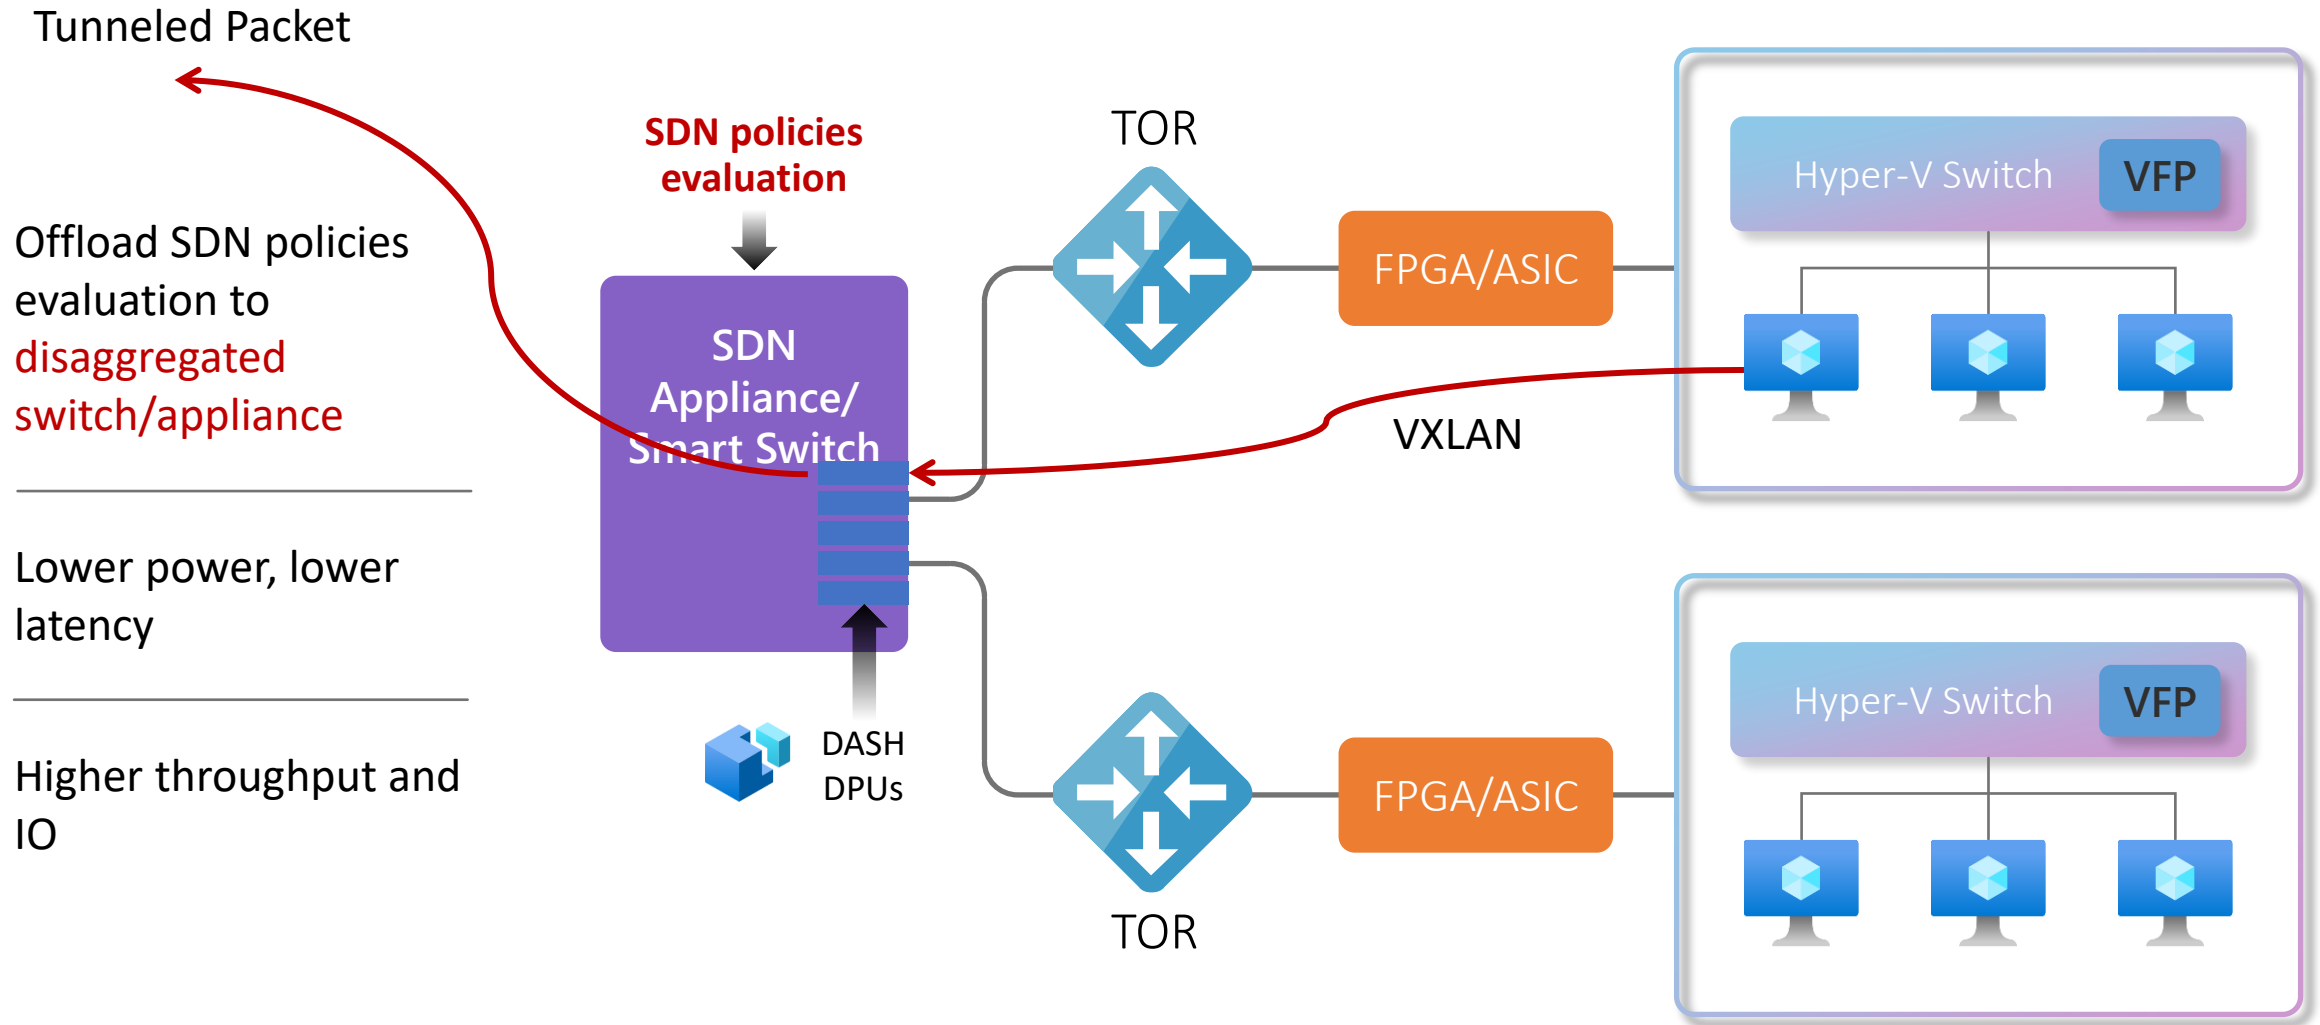
Demanding customers use high number of SDN rules

---

SDN policy evaluation can be resource intensive when configurations are complex




# Introducing Appliance/Switch - Powered by DASH (10s of Millions of Connections/Sec!)




# DASH Challenge!

## Solving the toughest cloud networking problems!




### Connections per Second

95% require < 10s of k/sec  
<5% require 100s of k/sec  
<1% require 1M+/sec



### Active Connections

10% of connections active  
30% of connections active  
100% of connections active



### Total Connections/Flows

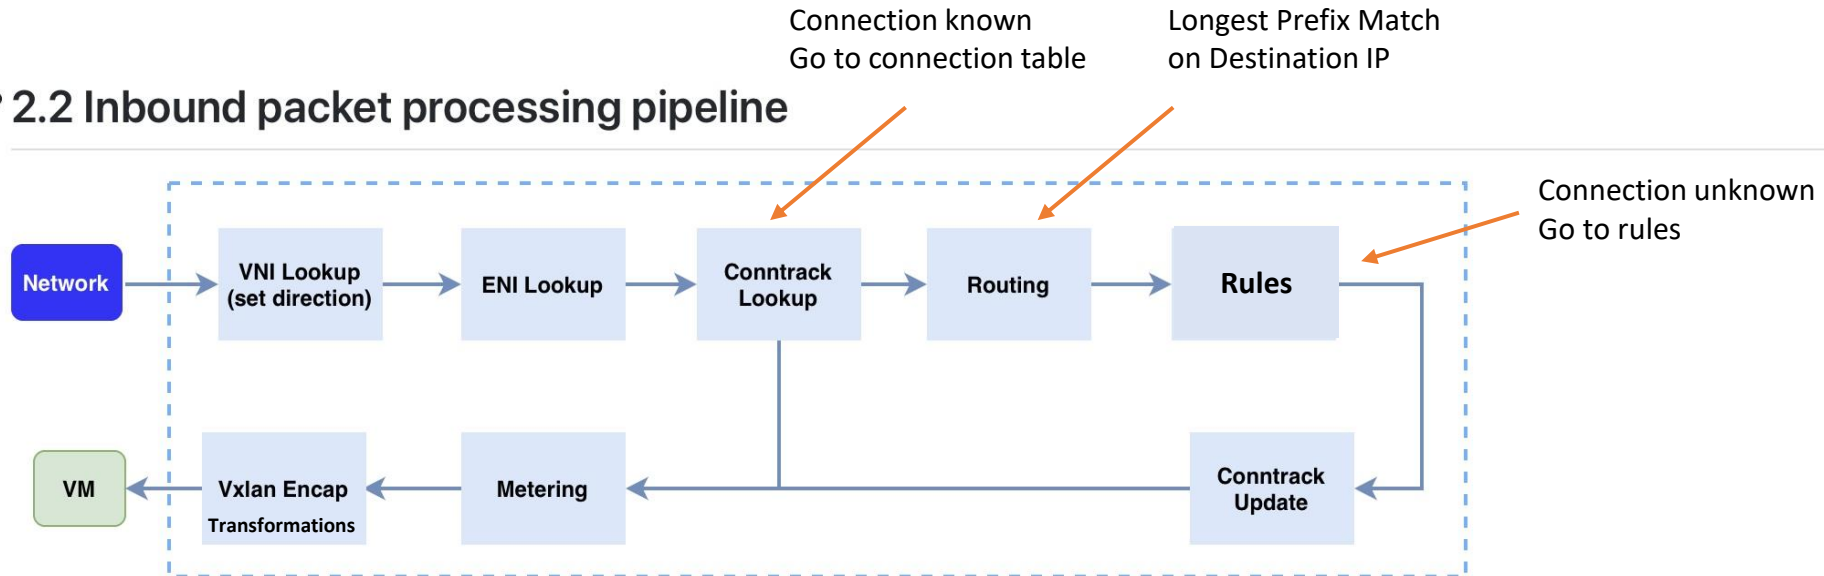
100k connections  
1M connections  
10M connections

# DASH Behavioral Models

[Azure/DASH: Disaggregated APIs for SONiC Hosts \(github.com\)](https://github.com/Azure/DASH)

[sonic-dash@googlegroups.com](mailto:sonic-dash@googlegroups.com)

## 2.2 Inbound packet processing pipeline



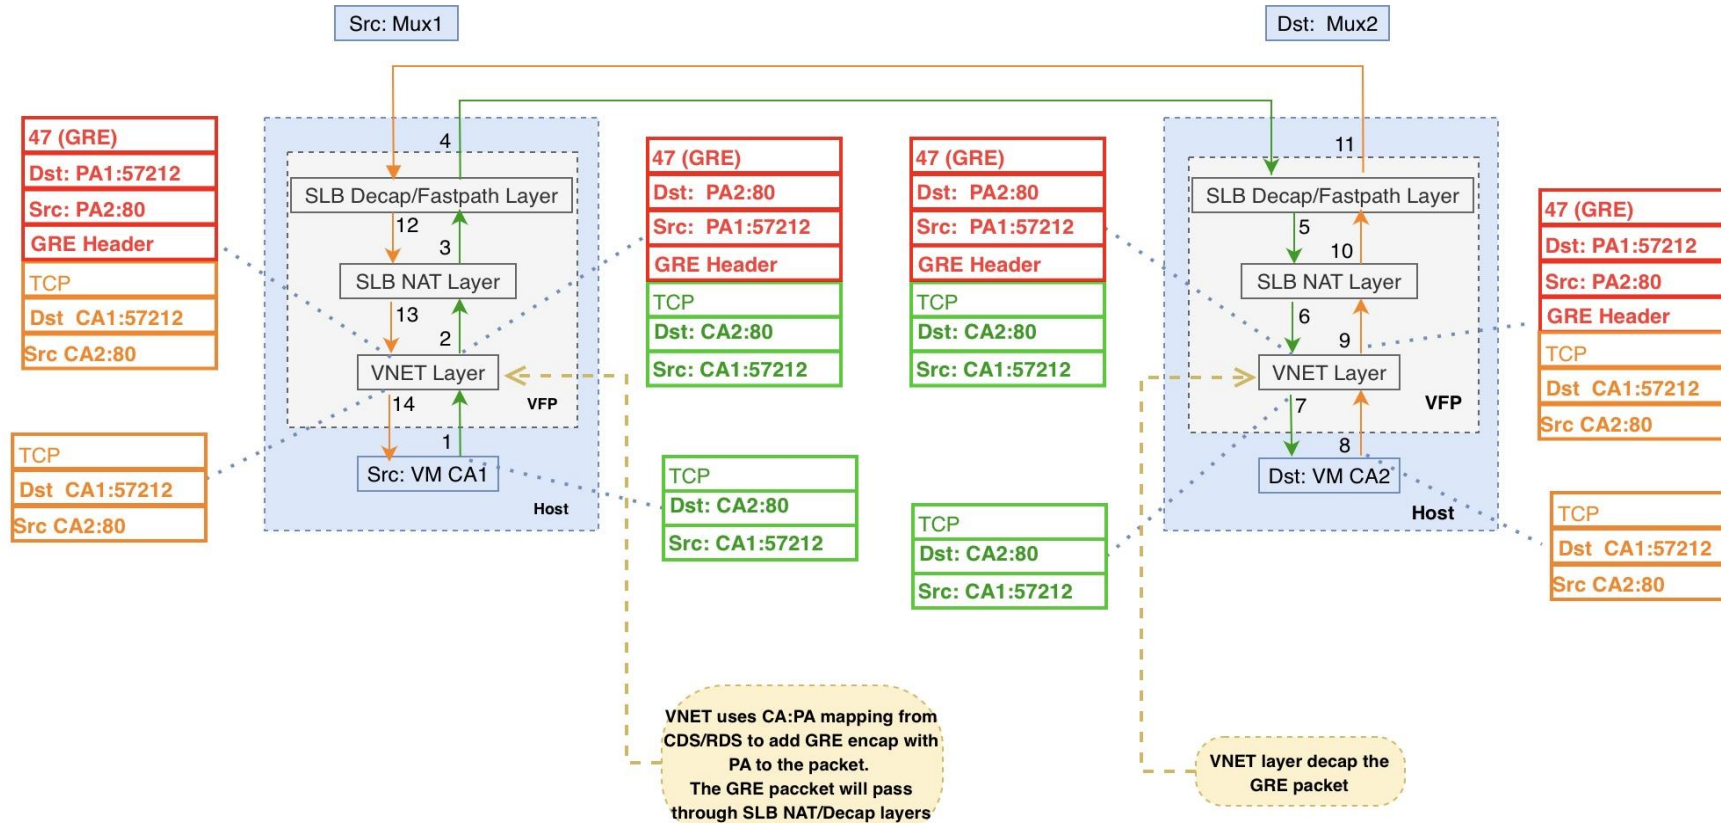
- Behavior is described in documents but also described in P4 along with matching test cases.
- DPU suppliers have full access to the behavioral model and test suites to validate their implementations
- DASH working group reviews new and existing services twice weekly

# DASH Transformations for every Cloud Networking Service

VM to VM  
(in VNET)  
communication

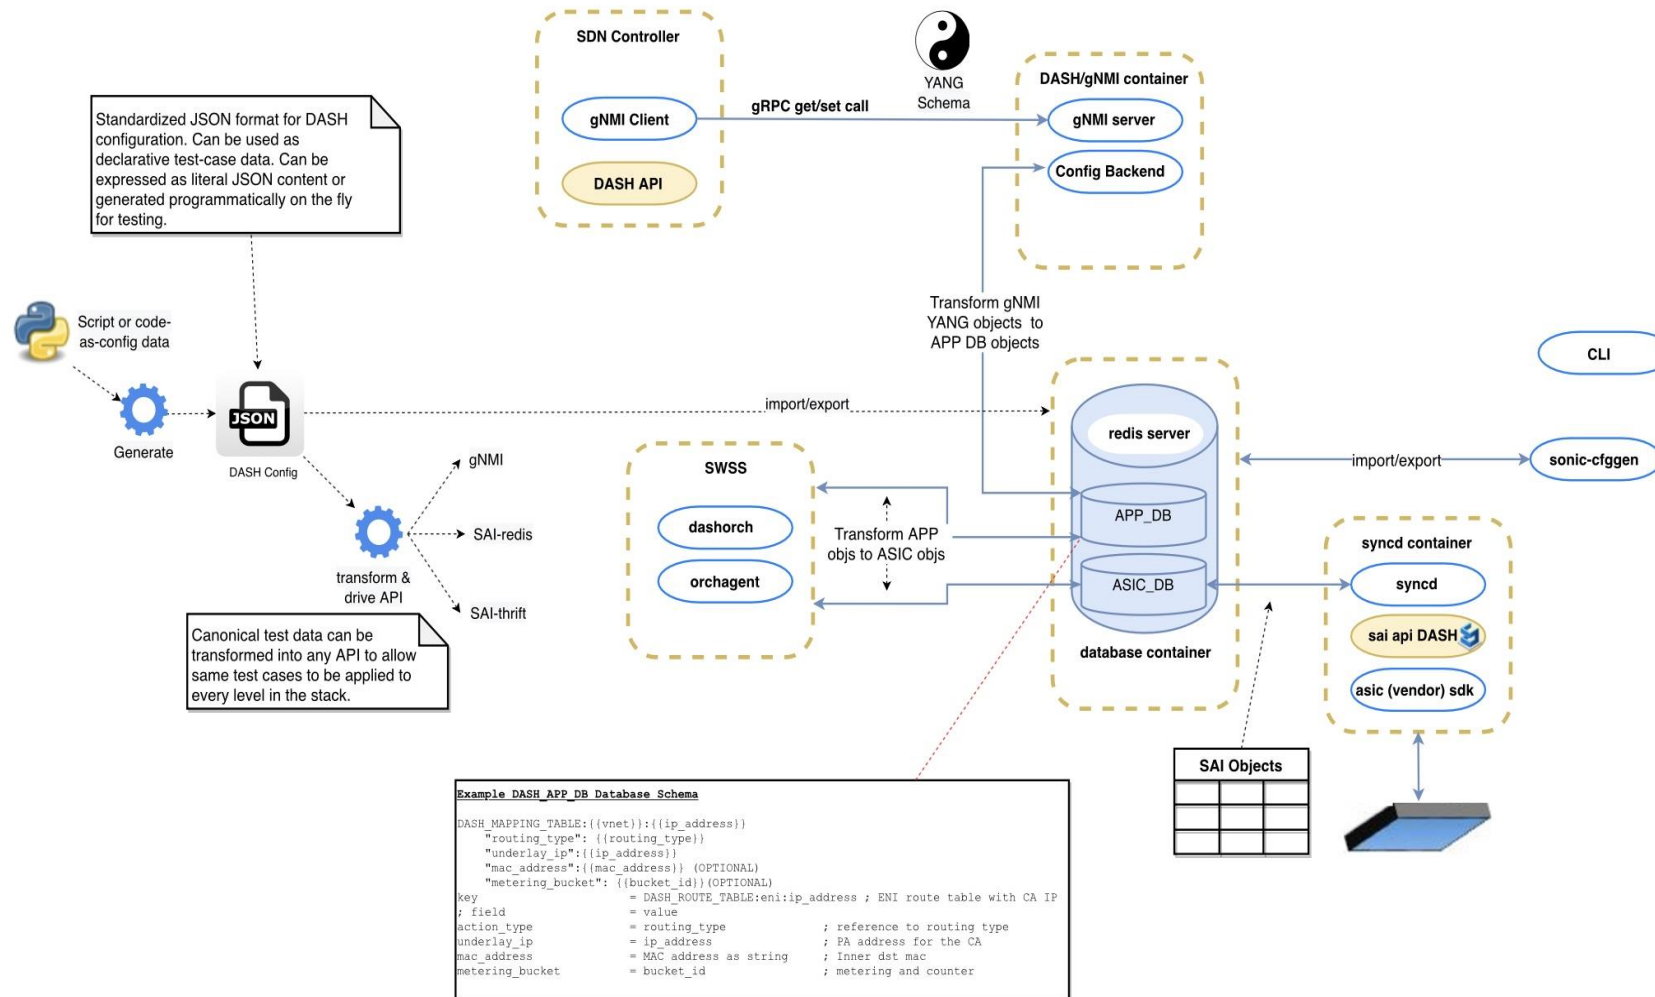
Encapsulation

47 (GRE)
Dst: PA1:57212
Src: PA2:80
GRE Header

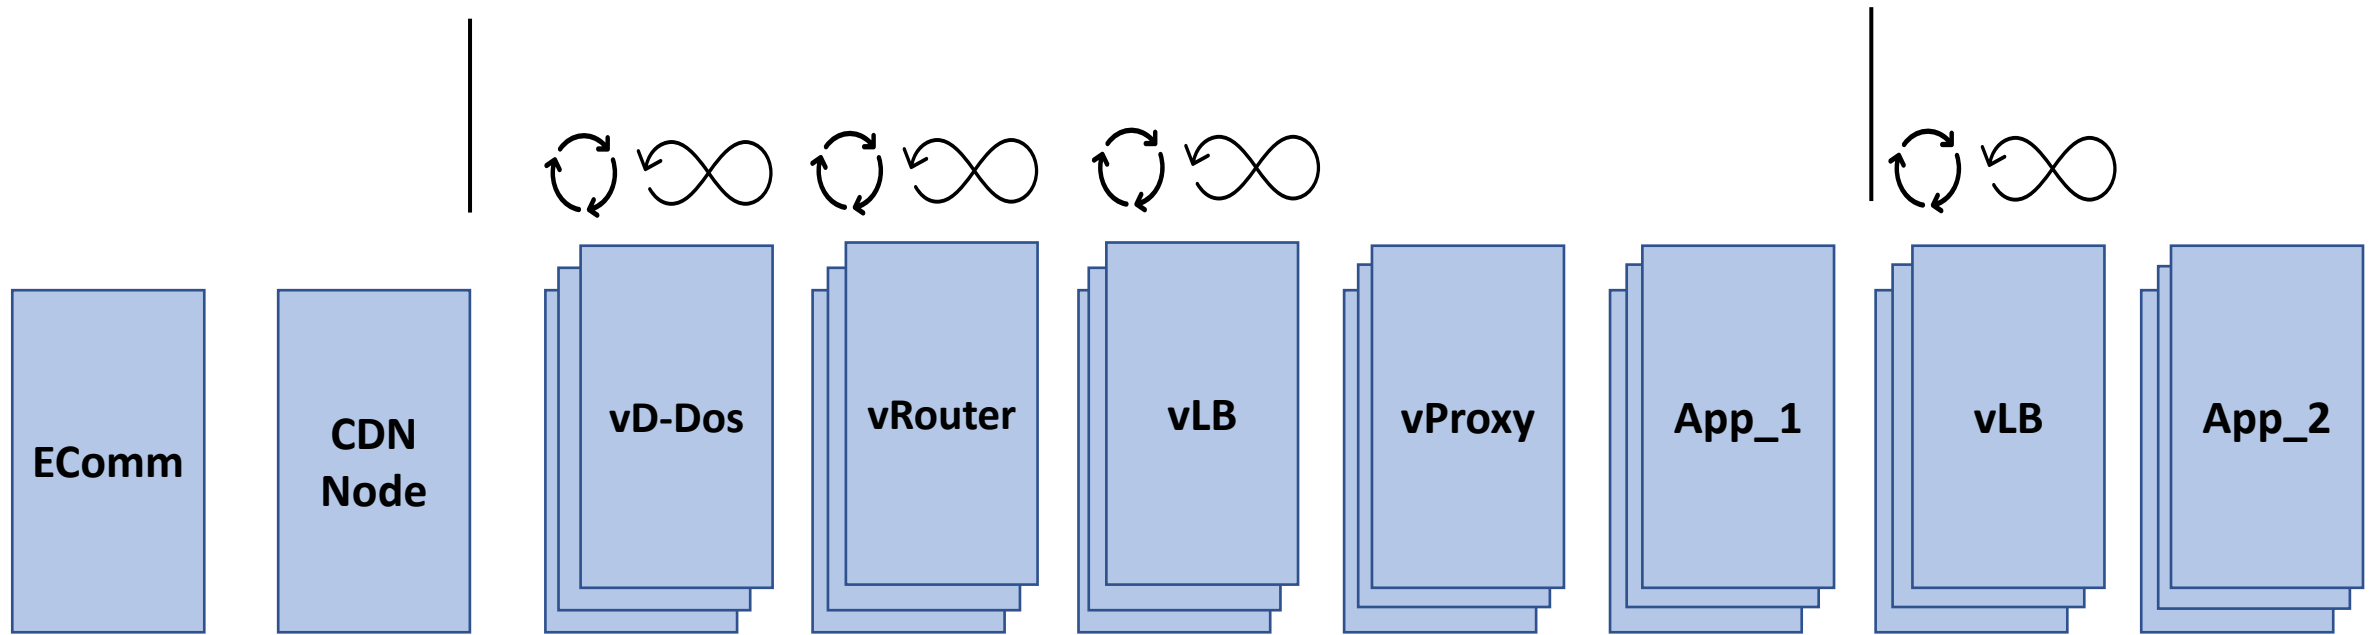


# (DASH) Open APIs

Figure - Schema Relationships



# Cloud Example of Virtual Appliances Turbo Charged via DASH (Processing Millions of Simultaneous Users)

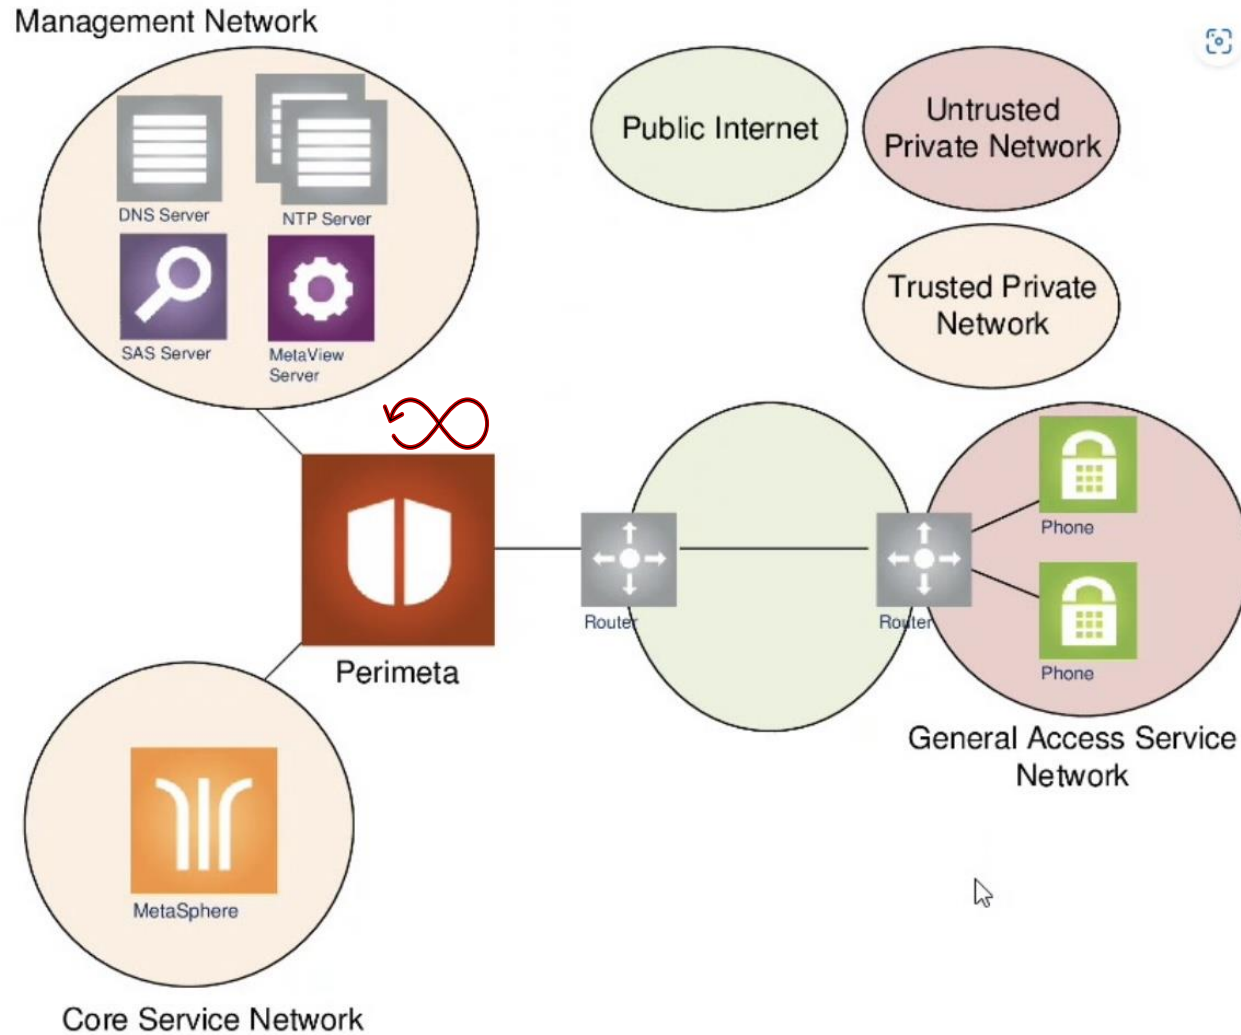


Millions of users are directed through layers of Content Caching, Load Balancers, Firewalls, DDoS, Authenticators, Billing Systems, and Applications

**DASH DOES NOT require changes to the Virtual Appliance VM  
(Already tested with F5, Palo Alto, Cisco, Alkira and more)**

# Cloud Example of 5G Edge Turbocharged via DASH

(Processing 100s of thousands of Continuous Transfers from 5G Devices)



Truly diverse and active connections result in PU (Processor Unit) cache misses forcing external DDR access and finally leading to reduced PPS or throughput

Software based connection/rule policy enforcement can lead to call quality especially in setup times

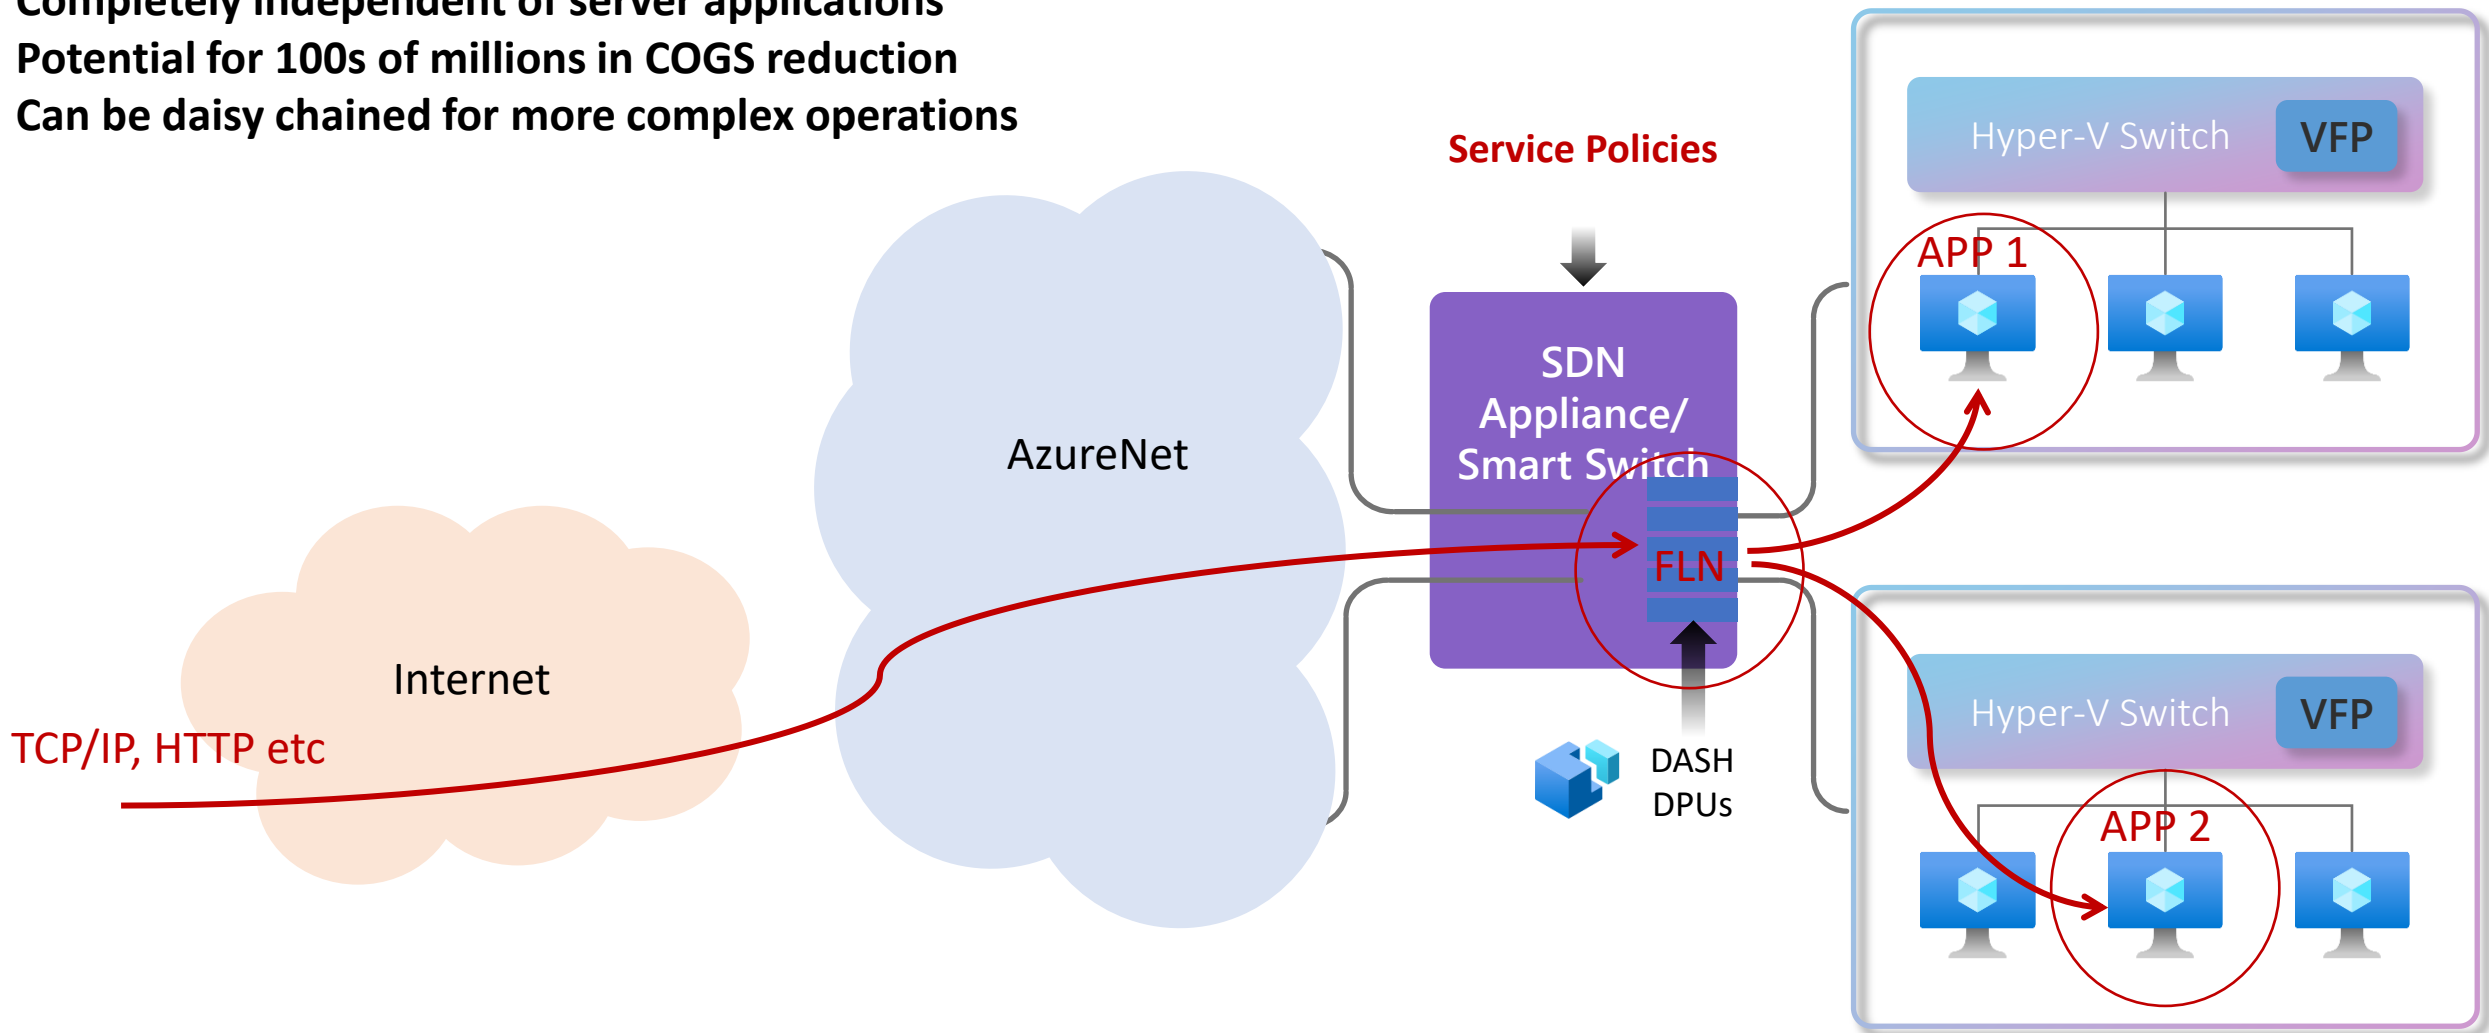
DPU are typically good at providing large internal cache and good techniques for pre-fetching data

Not all implementations are equal and require careful testing

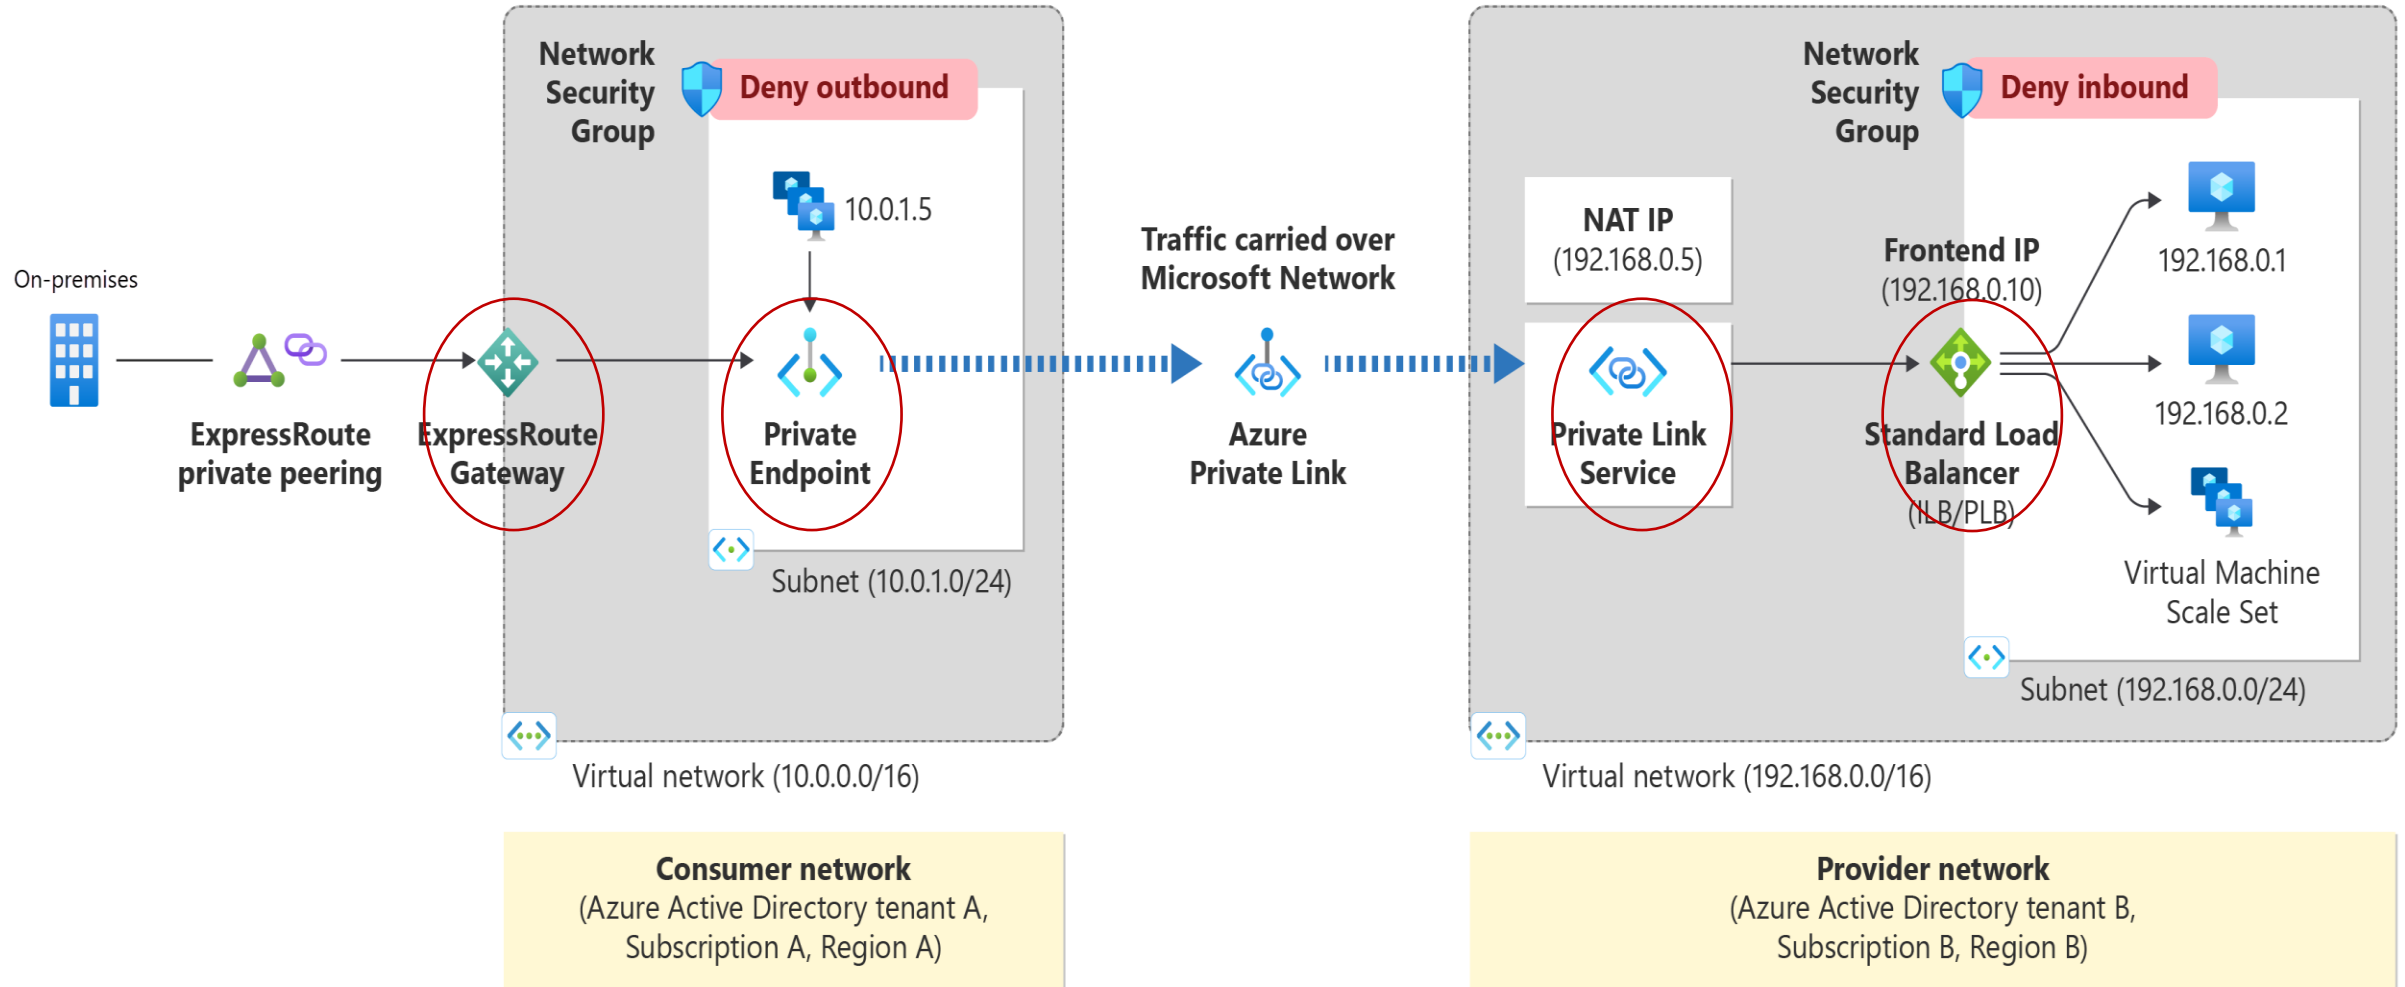


# Introducing the Floating NIC (FLN)

Full service at a fraction of the cost with much higher performance  
Completely independent of server applications  
Potential for 100s of millions in COGS reduction  
Can be daisy chained for more complex operations



# Floating NIC (FLN) Examples



# Where to find out much more detail

DASH Github: [Azure/DASH: Disaggregated APIs for SONiC Hosts \(github.com\)](https://github.com/Azure/DASH: Disaggregated APIs for SONiC Hosts)

SONiC Github : <http://sonicfoundation.dev/>

Mailing list: [sonic-dash@googlegroups.com](mailto:sonic-dash@googlegroups.com)  
[sonicproject@googlegroups.com](mailto:sonicproject@googlegroups.com)

## Inviting contributions in all areas

- Azure's documentation for all cloud services deployed and enabled by DASH
- DASH/SONiC/SAI/P4/Behavior Models
- Hardware platform
- New features, applications and tools
- Download, test, deploy!
- Training, Hackathon, FAQ


# Customer Benefits with NVA Accelerated Connections

## 3rd Party SLB with Accelerated Connections

### Decreased CPU Utilization

22

Core savings



30

8

CPU Cores

Consumed on x86

### Customer benefits

Frictionless Insertion

>10x performance versus previous thresholds

Total active connections increasing up to 2x

Increased connections per second (CPS)

Consistent active connections

Increased CPU stability for high traffic network optimized VM

Reduced jitter