



# The New Era of Open Networking

Tetsuya Murakami  
[tetsuya@arrcus.com](mailto:tetsuya@arrcus.com)

# 自己紹介



- 古河電気工業（'98/4 – '05/7）
  - WDM, SDH, SONETのソフトウェア開発
  - BGP, MPLSのソフトウェア開発
- Cisco Systems（'05/7 – '09/7）
  - Japan Development Centre
  - P2MP Label Multicastの開発（IOS-XR）
- IP Infusion（'09/7 - '18/10）
  - 6RD開発
  - MAP開発・標準化
  - ZebOS/VirNOS/OcNOS開発
- Delta Network Inc（'18/11 – '19/5）
  - Whitebox OpenSource Software開発支援
- Arrcus Inc（'19/6 – 現在）
  - ArcOS MPLS開発
  - VPP for SRv6 mobileの開発支援

# Company Overview



## Mission

- Software-powered network transformation for the interconnected world

## Company

- Founded 2016, Headquarters: San Jose, CA.
- Team: ~50 people

## Target Customers

- Fortune 100, Telecom/Cloud & CDN/Service Providers

## Board Investors

- **Guru Chahal:** Lightspeed Venture Partners
- **Steve Herrod:** General Catalyst
- **Chris Rust:** Clear Ventures

CLEAR

GENERAL CATALYST

Lightspeed

pillsbury

## Advisors

- **Kelly Ahuja:** ex-SVP of Service Provider; Cisco
- **Fred Baker:** ex-Chair IETF/Co-Chair IPv6, Cisco Fellow
- **Amarjit Gill:** SiByte, PA Semi, Maginetics, Viptela etc
- **Nancy Lee:** CHRO, Lime; ex-VP of People, Google
- **Farzad Nazem:** ex-VP of Eng./Oracle CTO/Yahoo
- **Pankaj Patel:** ex-EVP of Cisco Systems
- **Rajiv Patel:** ex-VP of Engineering, Juniper
- **Shawn Zandi:** Network Director, LinkedIn/MSFT



# Leadership Team



**Devesh Garg**  
*Founder & CEO*

EZchip  
Bessemer Venture Partners  
Broadcom, Synopsys, LSI Logic



**Keyur Patel**  
*Founder & CTO*

Cisco  
Distinguished Engineer (Routing & VPNs)  
Interdomain Segment Routing Lead



**Murali Gandluru**  
*VP, Product Management*

Cisco & Oracle  
Insieme/Datacenter  
NX-OS, IOS-XR, ACI, Linux



**Arthi Ayyangar**  
*VP, Customer Engineering*

Arista & Juniper  
Product Mgmt/SW Dev.  
EOS, JUNOS, NX-OS



**Brad Bratten**  
*VP, Business Development*

Broadcom  
Drove MSDC/SP Business  
Annual Revenues ~\$500M



**Abhay Roy**  
*VP, Engineering*

Cisco  
Cloud, Virtualization, & Manageability  
BGP, OSPF, IS-IS, BFD, Segment Routing  
Co-Chair IETF OSPF WG



**Derek Yeung**  
*Founder & Chief Architect*

Cisco & Procket – MTS (~# 7)  
Development Lead for OSPF; IS-IS  
Routing/NX-OS; EVPN & YANG

# Disaggregation to Open Integration...

Arrcus, Inc. Confidential - Do Not Distribute or Reproduce



# Information Revolution: Massive Scale, Agility & Intelligence



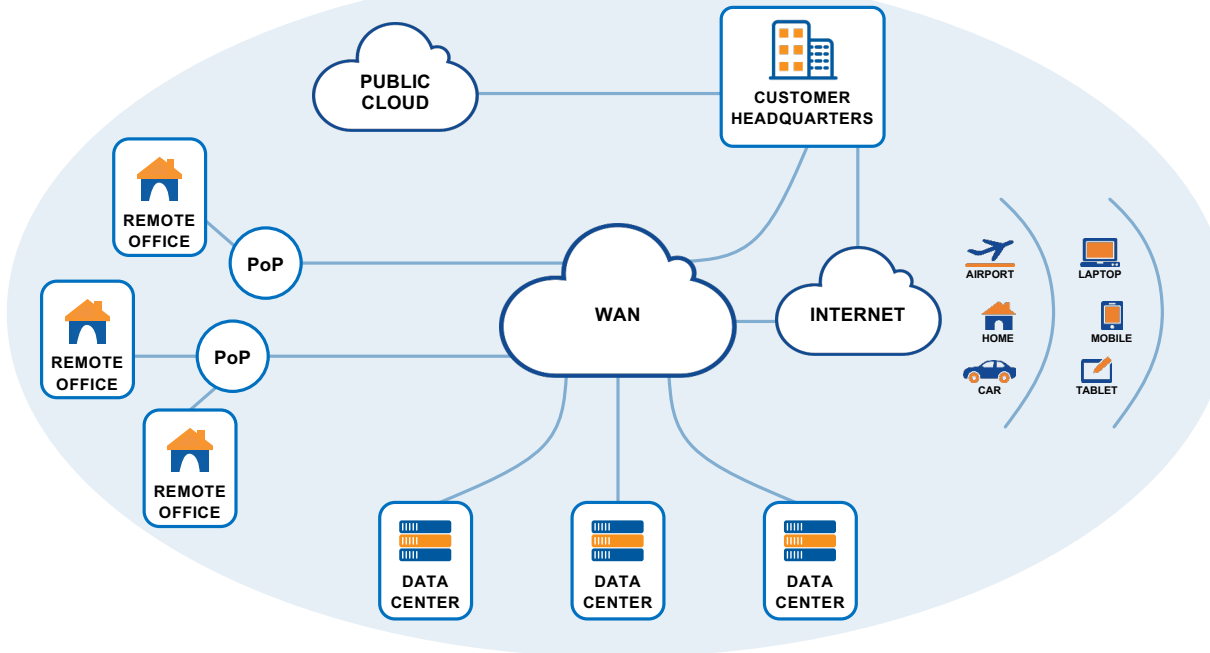
## Immersive, Data Intensive Experiences

Video is 82% of traffic by 2021, AR, VR, Intelligent 'things', Blockchain  
More intelligent, interconnected network edge



### 5G Networks

Higher data speeds,  
Lower latency,  
Accelerates network relocation to distributed locations



### Cloud-to-Edge

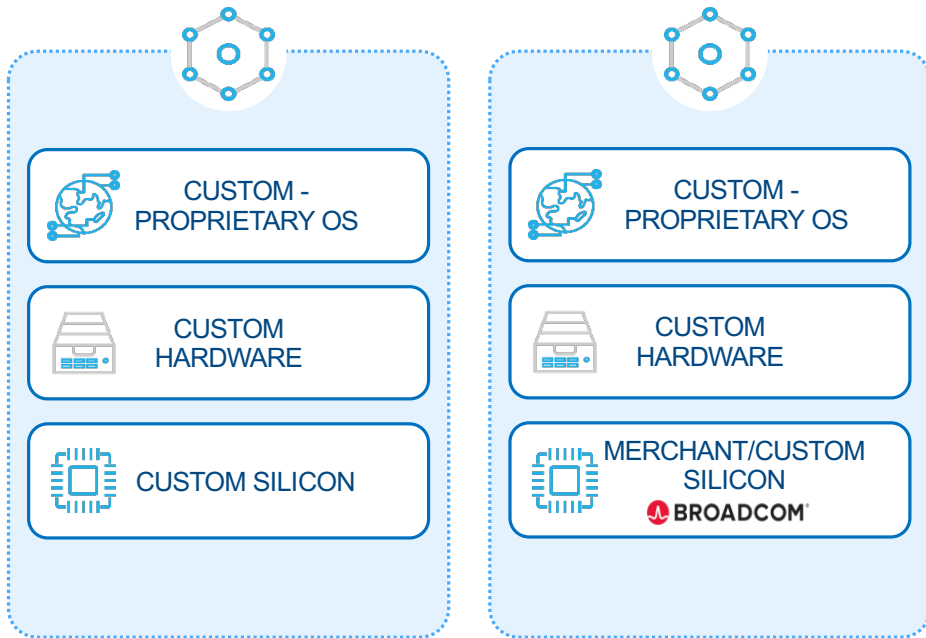
Data is Closer to User  
Distributed DCs & Interconnect,  
Accelerating 40G-  
>100G...400G, SAAS based models require stricter SLAs



# Customers are Constrained by OEMs



## VERTICAL INTEGRATION



**Legacy Adoption:**  
Enterprises, Commercial

**1st Generation Adoption:**  
Hyperscale/Telecom/ SaaS/CDN  
Providers, Global Enterprise

Rigid Deployment Model(s)

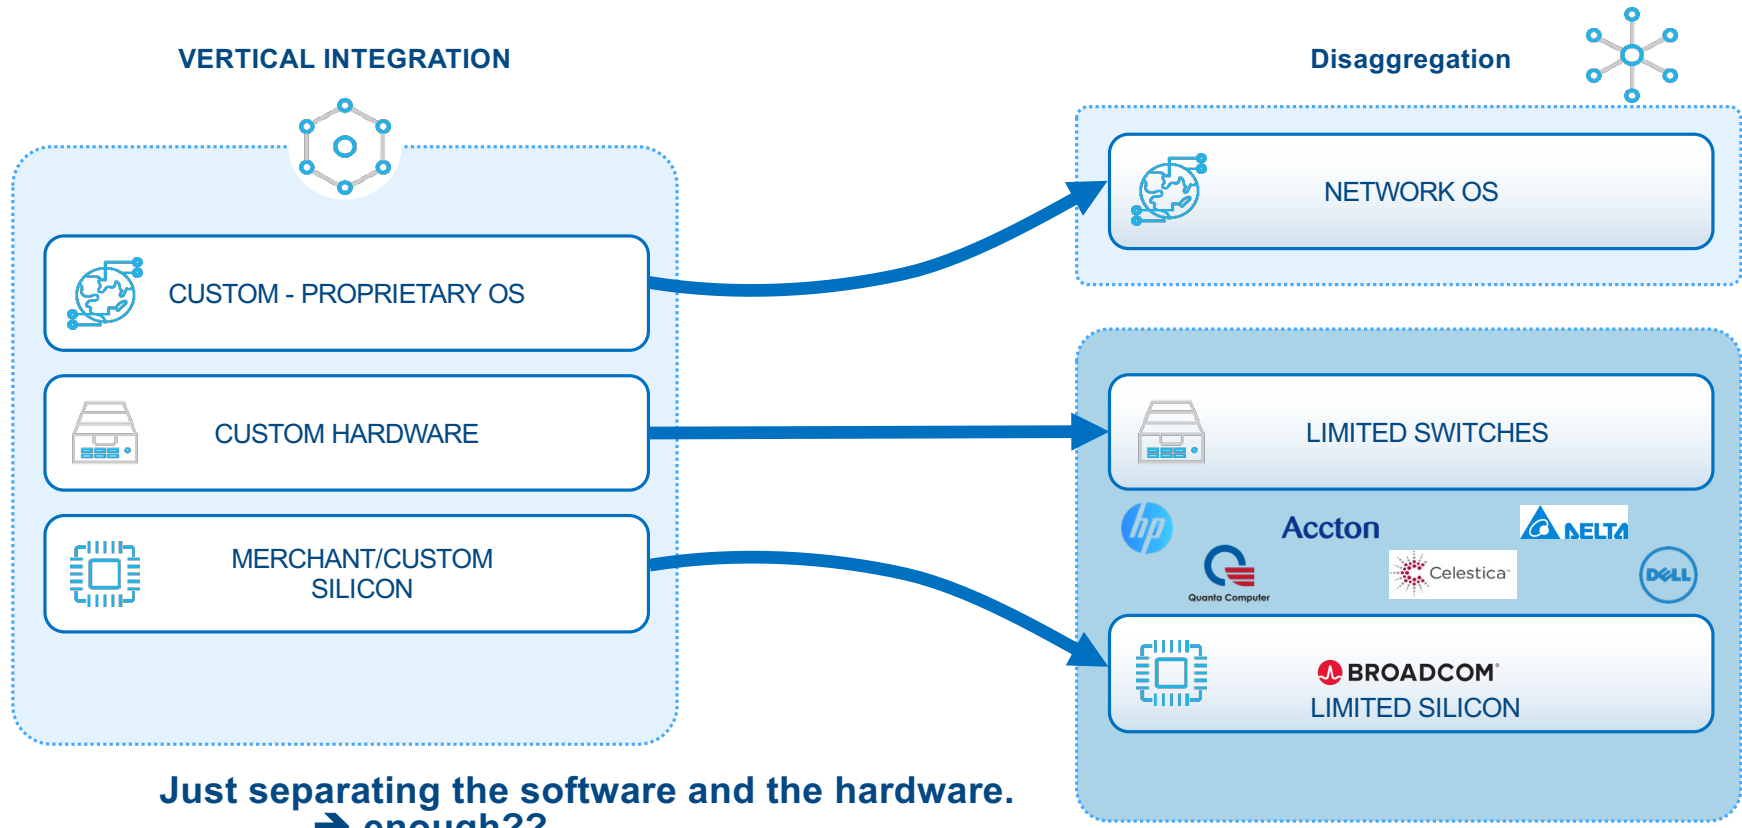
Limited Choices Dampen Competition & Innovation

Ever Increasing Network/Business Interruption:  
Poor Software Quality ("Bloat")

Inconsistent Data Models

Legacy Procurement Model & Vendor Lock-In

# Legacy approach: Disaggregation





# Disaggregation Vendors Tried and Failed



- **Whitebox could separate the software and the hardware for Networking appliance....**
  - **Hardware**
    - Using commoditized Processor and Broadcom Networking silicon
    - Limited switches provided by ODM hardware vendors
    - **Good:**
      - Possible to choose a suitable hardware based on the required scalability, performance, etc.
      - Increase the choice of the switch hardware
  - **Software**
    - Software vendors can deliver the software for the bare metal solution only
    - The networking software provided by software vendors includes everything like protocol stack, kernel, driver, etc
    - Only one software vendor's solution can run on a given device
    - Using non commoditized/vendor specific data model
    - **Still Not Good:**
      - **Inconsistent CLI, access models etc based on the software vendor's solution**
      - **Software vendor's lock-in (no space to run other applications to fill in the missing functions)**
      - **Limited capability provided by a single software vendor**
      - **Support model could be complicated....**

# So, What needs to Happen to Support Network Transformation?

Arrcus, Inc. Confidential - Do Not Distribute or Reproduce

 **ARRCUS**  
NETWORK DIFFERENT

# Routing-Centric Network Transformation



## ARCHITECTURAL TRANSFORMATION

## OPERATIONAL TRANSFORMATION



SCALE-UP/OUT  
CONTROL PLANE

PROGRAMMABLE  
DATA PLANE  
(1G - 400G)

FLEXIBLE HW CHOICE  
(FIXED/MODULAR)

OPEN  
STANDARDS

NETWORK  
TELEMETRY &  
ANALYTICS

PROGRAMMATIC  
APIs



DATA CENTER



CLOUD NETWORKING



ROUTE REFLECTOR (RR)



PEERING

# High-Performance Routing is Tablestakes for Networking!



## FLATTENING OF INTERNET

- Video is 82+% of traffic by 2021
- Subscriber bandwidth growing at 30%+ CAGR
- More direct connections and less transit back-ups



- Smaller footprints, distributed sites
- Greater automation due to greater number of remote sites
- Resiliency required to avoid costly back-ups
- Flexible port density



## PEERING MOVES DEEPER INTO THE INTERNET

## NEW PARADIGMS

- Video dominates but.....
- SaaS based business models & stricter SLAs
- 5G accelerates relocations of network functions, data storage to multi-edge compute locations
- Shift to regional and metro peering



- IPv4 and IPv6 networks continue to grow
- 5G, IOT, Remote PHY increase device endpoints by 5-10X
- Internal routing continues to grow as more services are distributed and cloud-native
- Router scale requires a no-compromise router longevity approach



## EVER GROWING INTERNET ROUTING TABLE

# Foundational Elements of Network Transformation



## SIMPLE




- Simplified Protocols (e.g. Common BGP control plane)
- Standards based Automation (e.g. OpenConfig/YANG)
- Easy Manageability and Debuggability

## SCALABLE



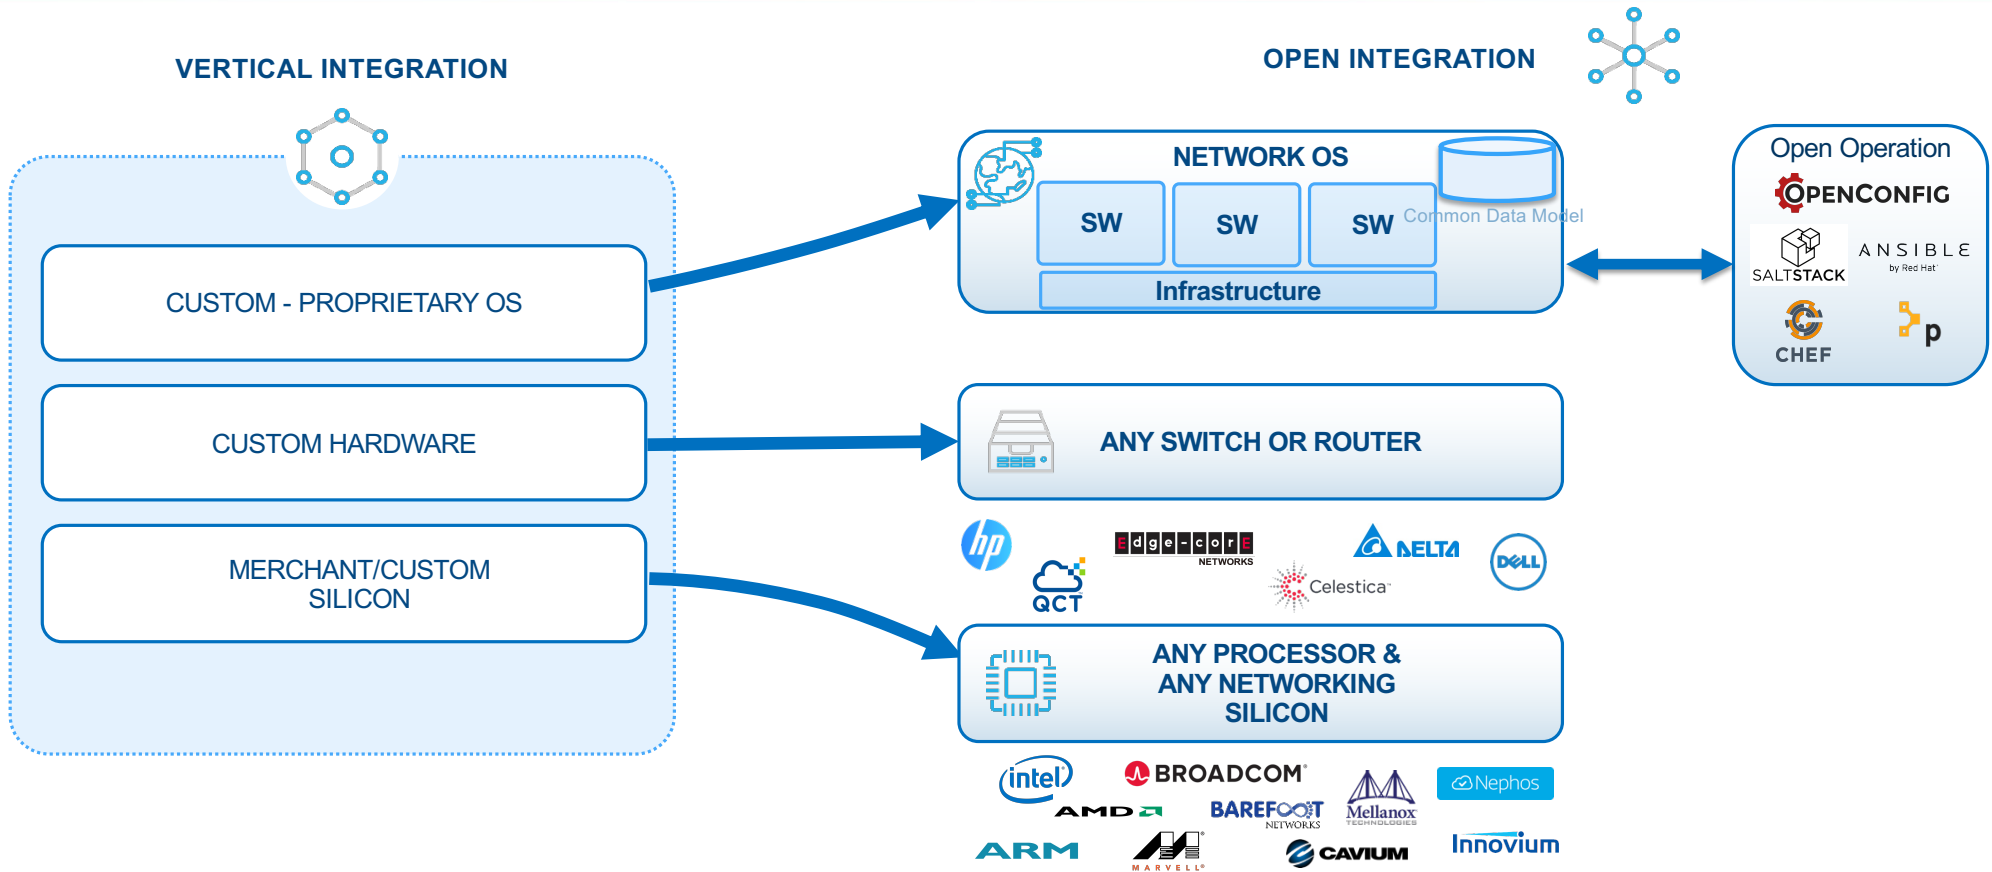
- Architectural Scale (e.g. Internet-scale routing w/ fast convergence)
- Network Scale (e.g. 100G/400G high-perf, low latency platforms)
- Operational Scale (e.g. Resilient, per-process restartability)

## SECURE

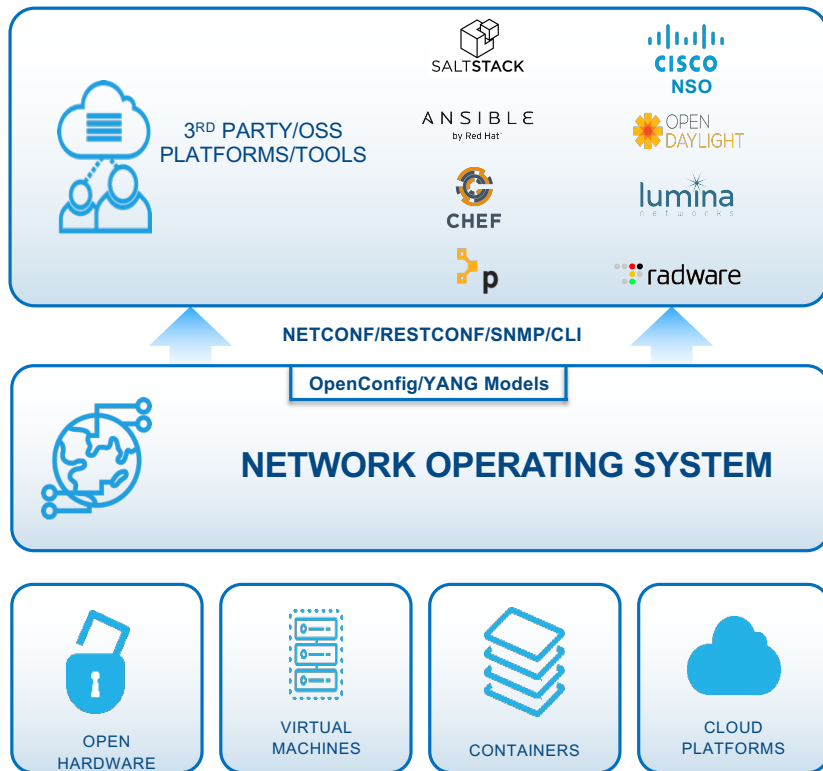


- Real-time Network Telemetry & Analytics
- Routing Security (ROV)
- DDoS Mitigation

# Open Integration is the Way Forward



# Open Integration Overview



## ▪ Networking Operating System

- Micro-service architecture
- Docker/container
- Open Software framework
- Common Operating System (Open Network Linux)
- Multiple applications/software solutions running on a single platform.
- Open Adaptive Cloud/Industry standard interface, i.e, OpenConfig, REST, Netconf, etc.

## ▪ Composable, Microservices-based Architecture

- Flexible form factors – physical, virtual, container, and cloud
- Dataplane Adaption Layer

## ▪ Automation: Standards-based APIs

- OpenConfig/YANG models for standardized access
- 3<sup>rd</sup> party/Operations Support System (OSS) integration
- Real-time streaming telemetry

# Open Integration cont'd



- **Open Network Software framework to eliminate the software vendors' lock-in....**
  - Using Open Network Linux as the base operating system.
  - Using Docker/container and/or VM to achieve the flexible software framework.
    - Need to provide not only bare metal solution but also Docker/Container or VM solution.
  - But this might be causing the complicated operation/support model if there is no common framework...
    - Need the common data model (i.e., OpenConfig)
    - Need the common tools/methods to manage the lifecycle for each application. (i.e., Kubernetes)
- Software vendors should support various installation model like bare metal, Docker/Container, VM...
- Software vendors should support common data model like OpenConfig.
- There are several tools to manage the lifecycle of docker/container, vm but need to define/develop the consolidate framework.

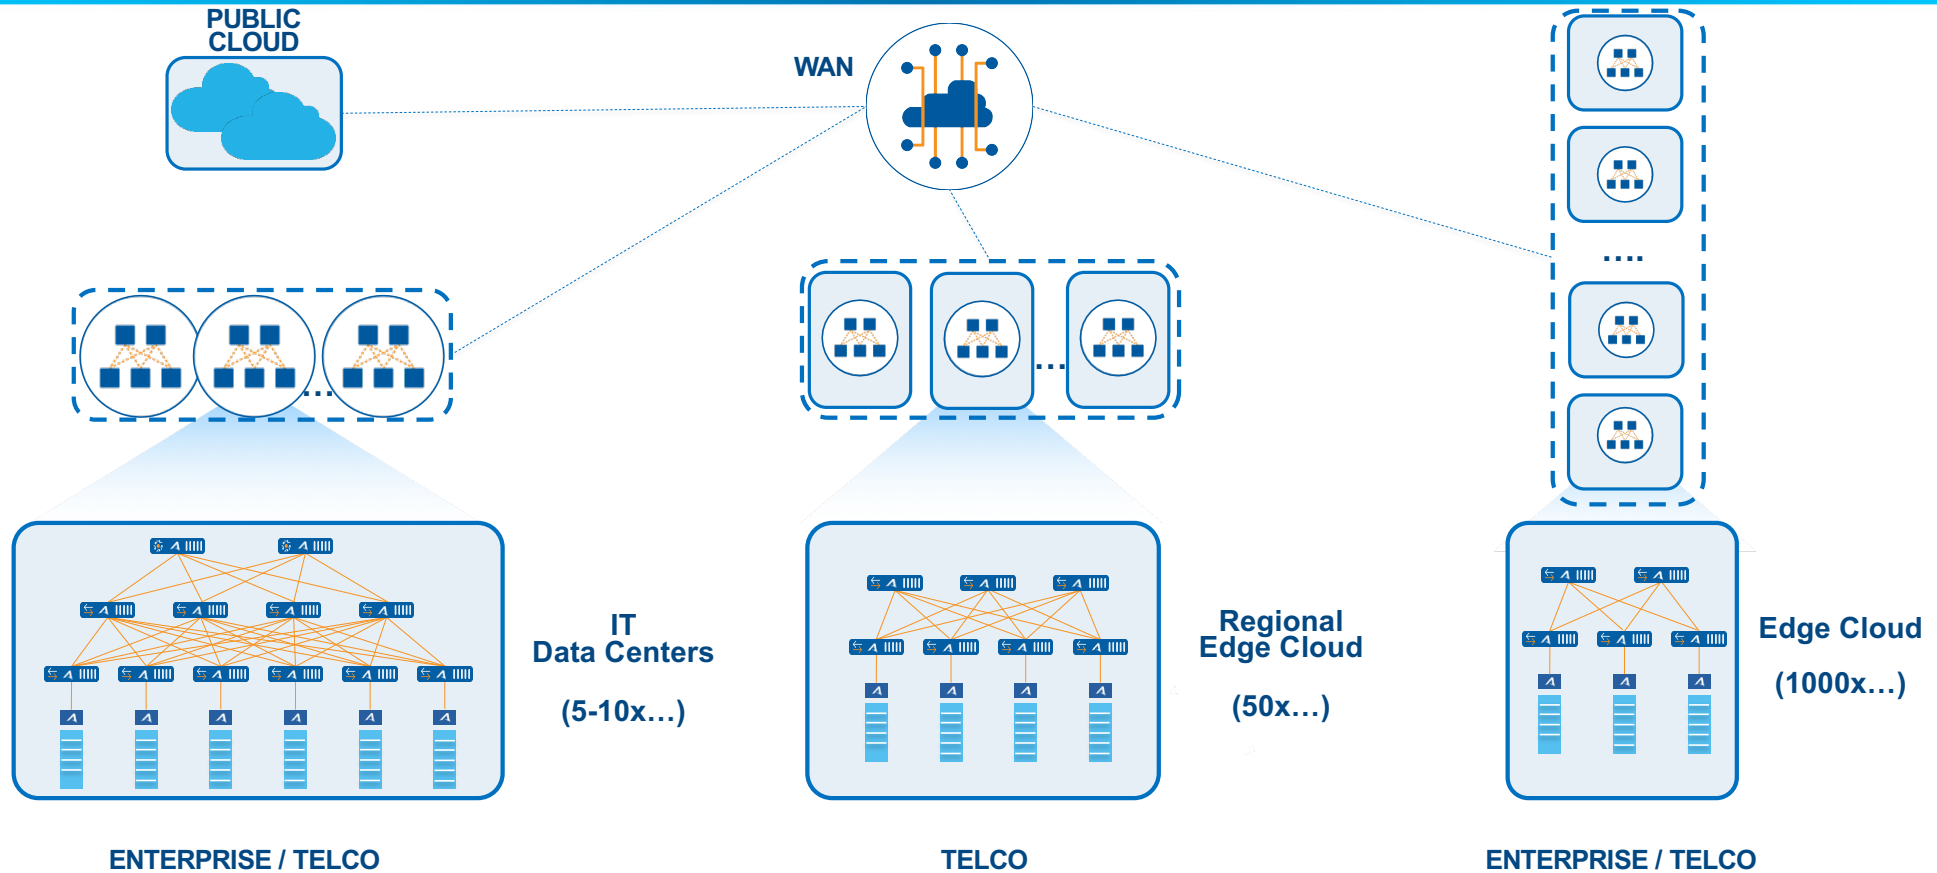


# Use Cases The Arrcus Solution

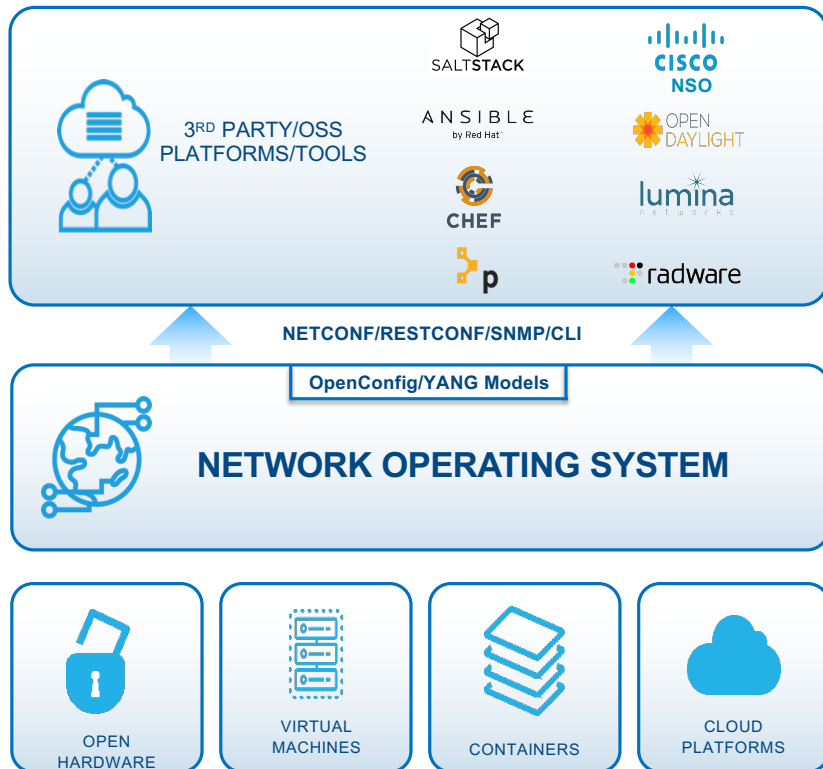
Arrcus, Inc. Confidential - Do Not Distribute or Reproduce



# IP CLOS Architecture: Spreading ToR to Super Spine, Switching to Routing, DC to Edge



# Built for Performance & Open Integration for ODM Platforms



## ▪ Best-in-Class Protocols

- BGP
- IPv4/IPv6/Label Unicast/MPLS/SR/EVPN support
- Process restartability for routing, RIB, and FIB

## ▪ Massively Scalable Architecture

- 64-bit user space processes on a 64-bit OS
- Multi-threaded with minimal locking
- Control and Data Plane independently scale, Datastore

## ▪ Composable, Microservices-based Architecture

- Flexible form factors – physical, virtual, container, and cloud
- Native patching capability for easy upgrades
- Data plane Adaption Layer

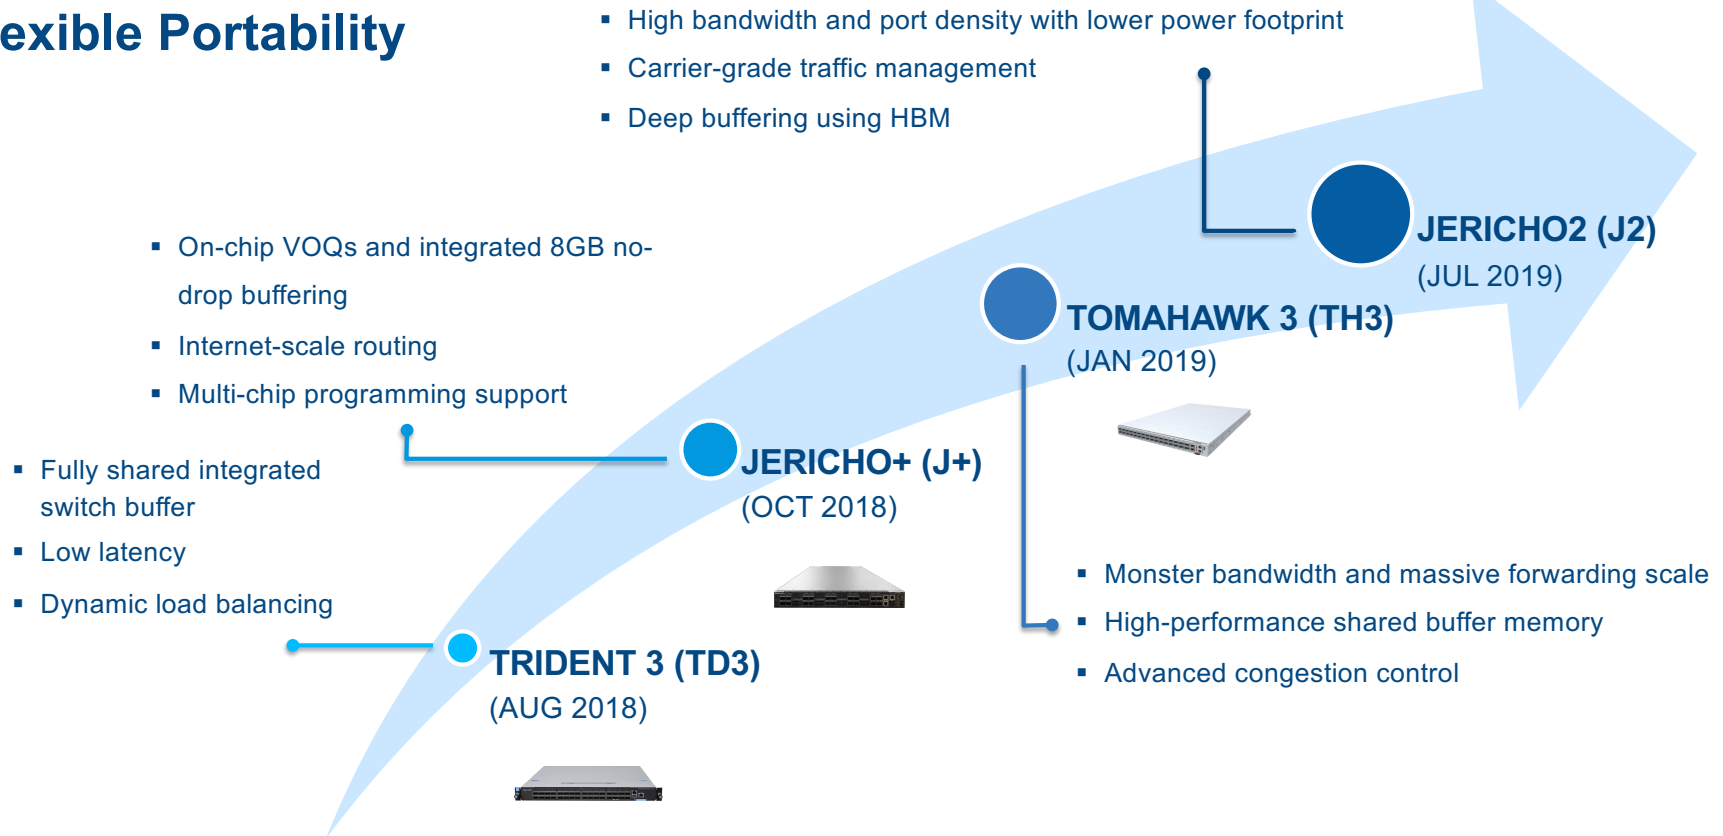
## ▪ Automation: Standards-based APIs

- OpenConfig/YANG models for standardized access
- Easy 3<sup>rd</sup> party/Operations Support System (OSS) integration
- Real-time streaming telemetry

**Require Carrier-Grade, Internet Scale Performance, Usability and Flexibility**

# Newer Chipset for Large Scale Network: TD3, J+, TH3, & J2

## Flexible Portability



# Powering 100G/400G Routing w/Jericho2



24 ports 100GbE +  
6 ports 400GbE

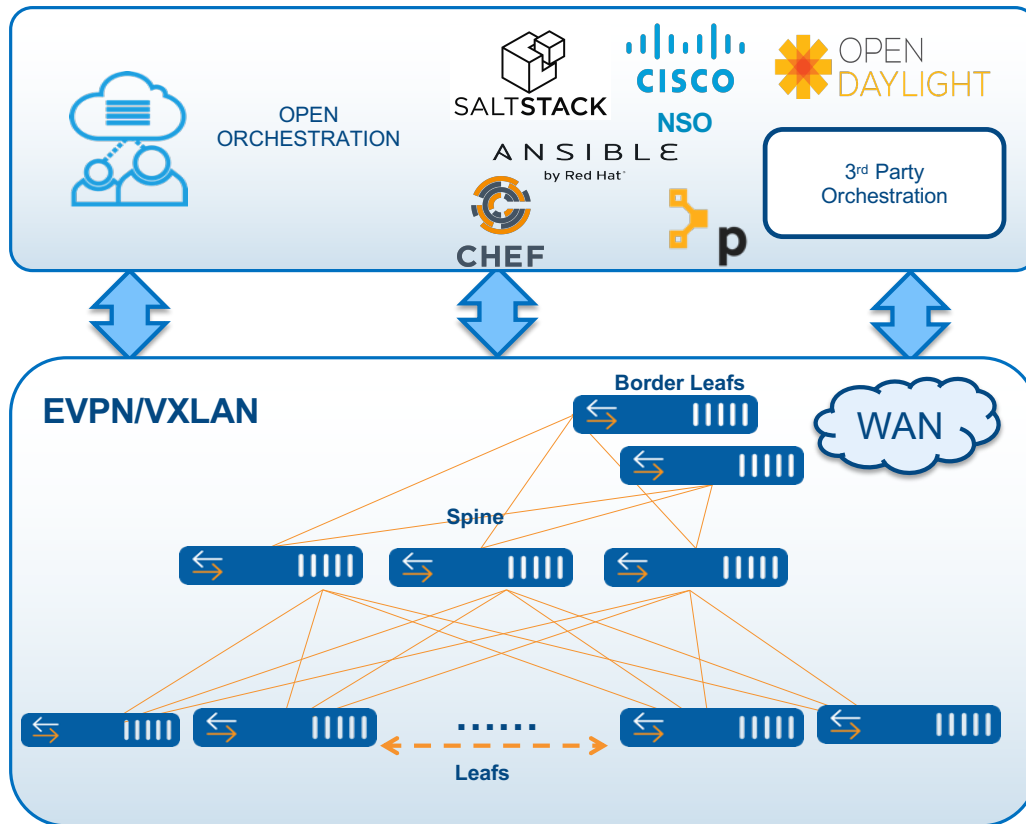
40 ports 100GbE

80 ports 100GbE

96 ports 100GbE

- 10Tbps (5X higher bandwidth) with over 2M IPv4 routes (1M IPv6)
- 70% lower power per gigabit (vs. Jericho+)
- Efficient traffic management with scalable packet buffer memory
- Optimized Hyperscale Cloud, Edge, and 5G Networks
- Superior cost/performance
- Multi-vendor hardware options

# The Programmable Multi-Tenant Data Center Fabric



## Multi-Tenancy at Scale

- Integrated L2, L3 EVPN with automated peer discovery
- Process-Restartability for robust, resilient fabrics
- Per-Tenant Fabric Visibility with Streaming Telemetry

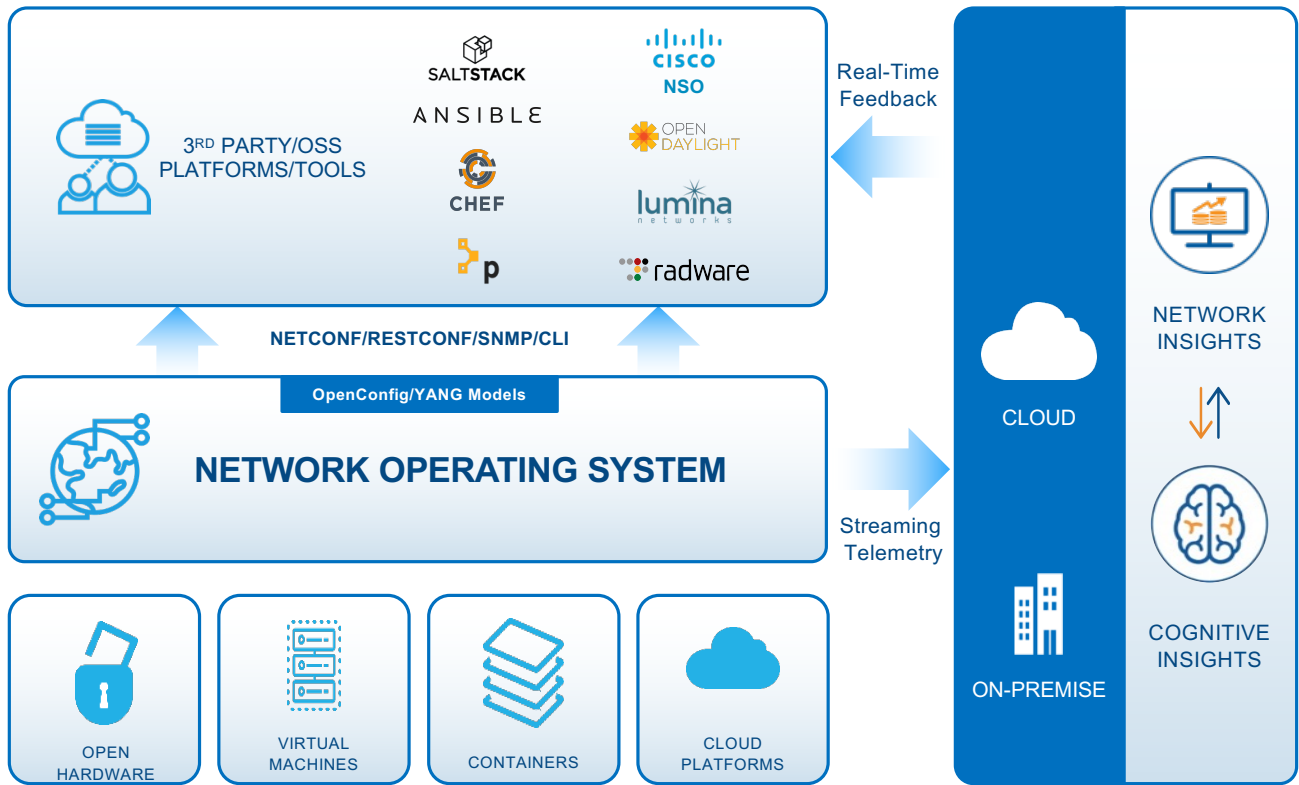
## Operational Simplicity

- Rapid, Automated Host Discovery and Mobility
- Innovative simplified control plane option w/BGP LSVR
- Easy-to-Deploy with intuitive Config Model

## Open Integration across HW and Vendors

- Support for TD3 across multiple ODM vendors
- Standards-based interop with vendors

# Streaming Telemetry for Programmable Networks



## Network Security

- Control plane, RIB, FIB, interface stats
- BGP topology, peers, & events
- ACLs (most used ACLs, least used ACLs)

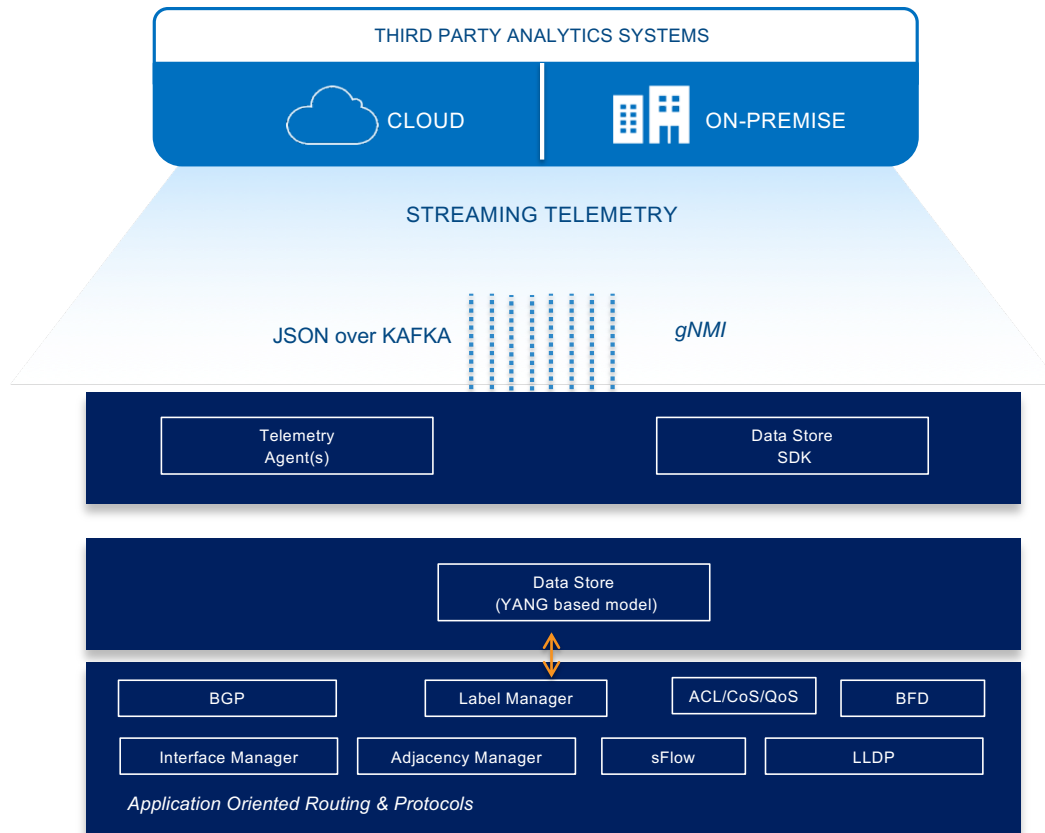
## Network Health

- Per-device platform resource state/usage
- Service insights: Platform software version tracking
- Process blacklist, whitelist

## Workload Mobility

- Tracking workload of VMs/containers
- Desired state queries (production vs development)


# Requirement for Network Telemetry & Analytics



- **Highly Available Data Store**
  - Resilient store for routing, infrastructure data/events
  - Building block for HA and streaming telemetry
  - Support for 3<sup>rd</sup> party integration via SDK
- **Streaming Telemetry**
  - Uses data store SDK to collect and securely stream the data out
  - JSON encoding over secure streaming platform (i.e., Kafka with gNMI)
  - Decoders to translate JSON to OpenConfig models



# AI-Driven Network Operations Center (AI-NOC) Platform



## Multi-Cloud. Edge. Multi-Vendor

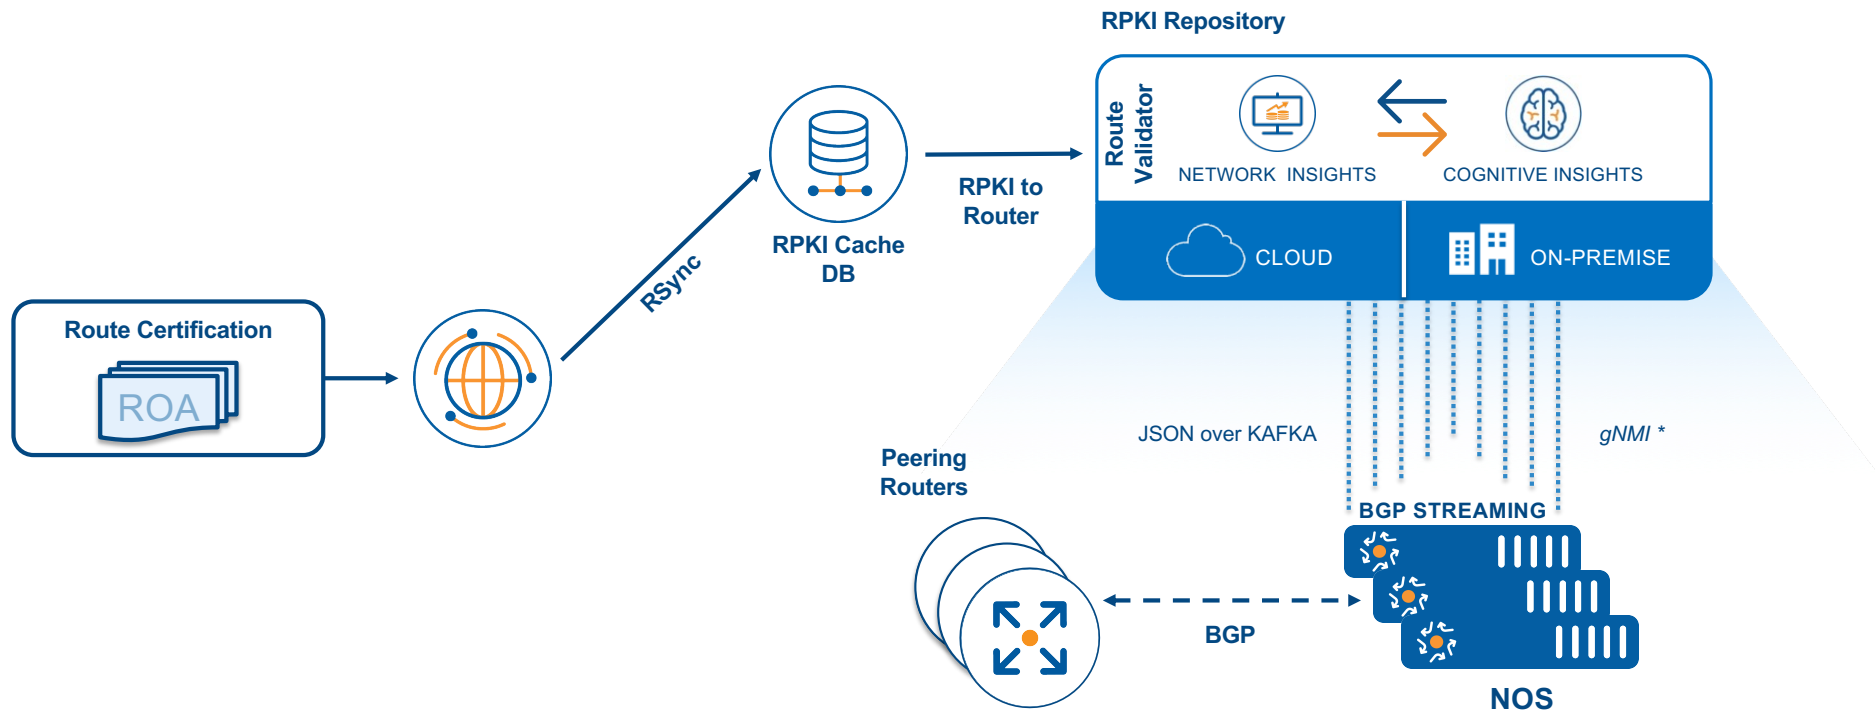
- **Visibility**
  - Open and standards-based
  - Network Health across DC, Cloud and Edge
  - Asset Management (Peripherals, Environmentals..)
- **Control**
  - AI/ML Driven Network Computation Engine
  - Intelligent Traffic Management (Policy Service Insertion)
- **Security**
  - Threat management, ROV Alerts, DDOS Mitigation
  - Partner Integration

**ENABLING SELF-HEALING NETWORKS**

# Network Telemetry and Analytics Use Case: Secure Control Plane Network Solution



## RPKI-Based Route Origin Validation (ROV)

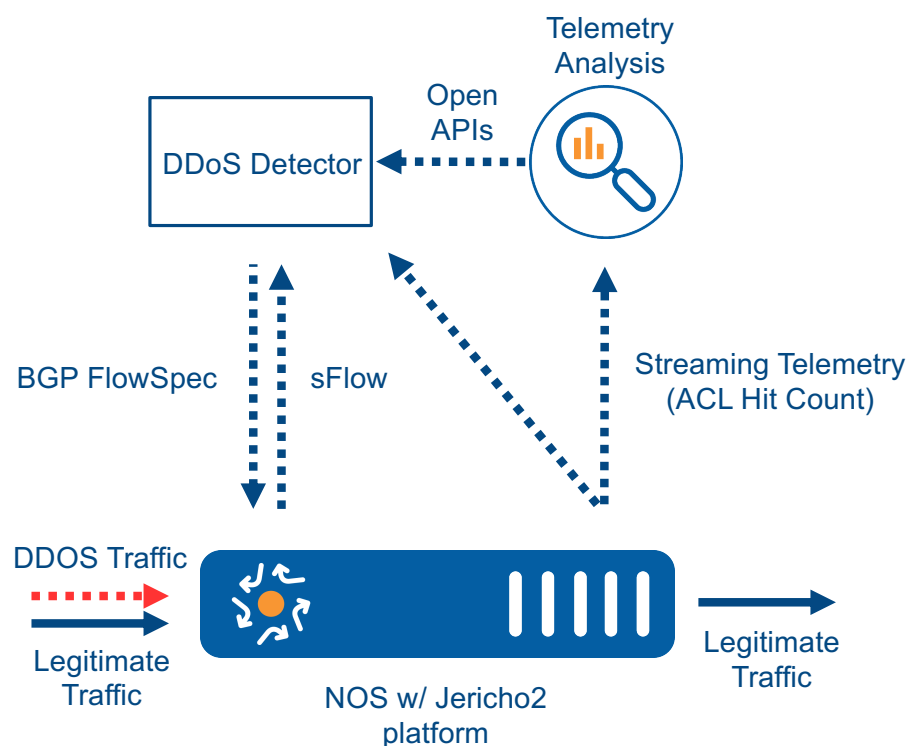


\*Roadmap

# Network Telemetry and Analytics Use Case: Secure Data Plane Network Solution

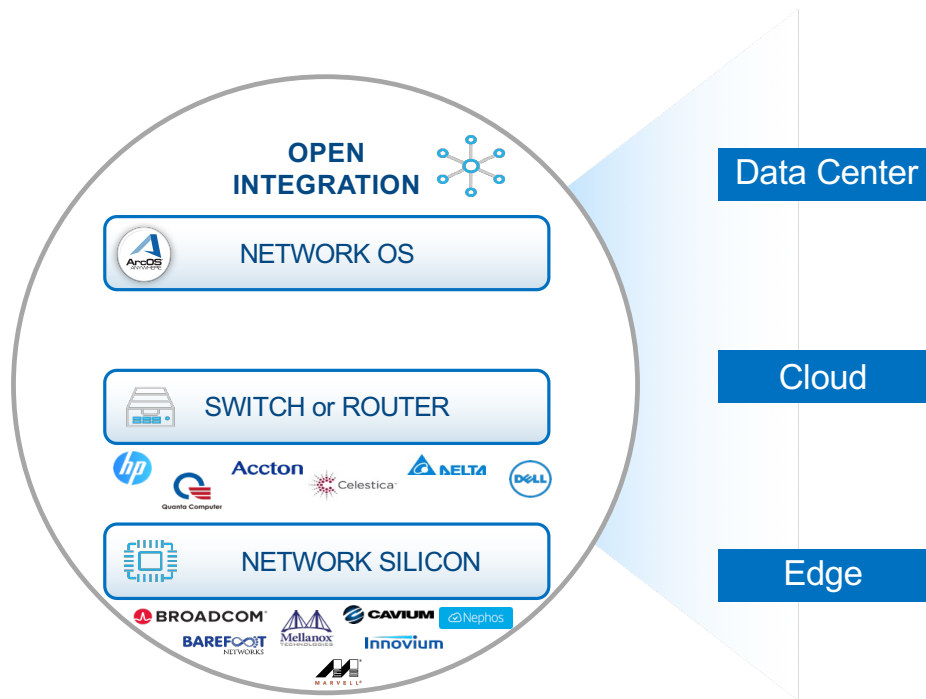


## BGP FlowSpec-Based DDoS Mitigation



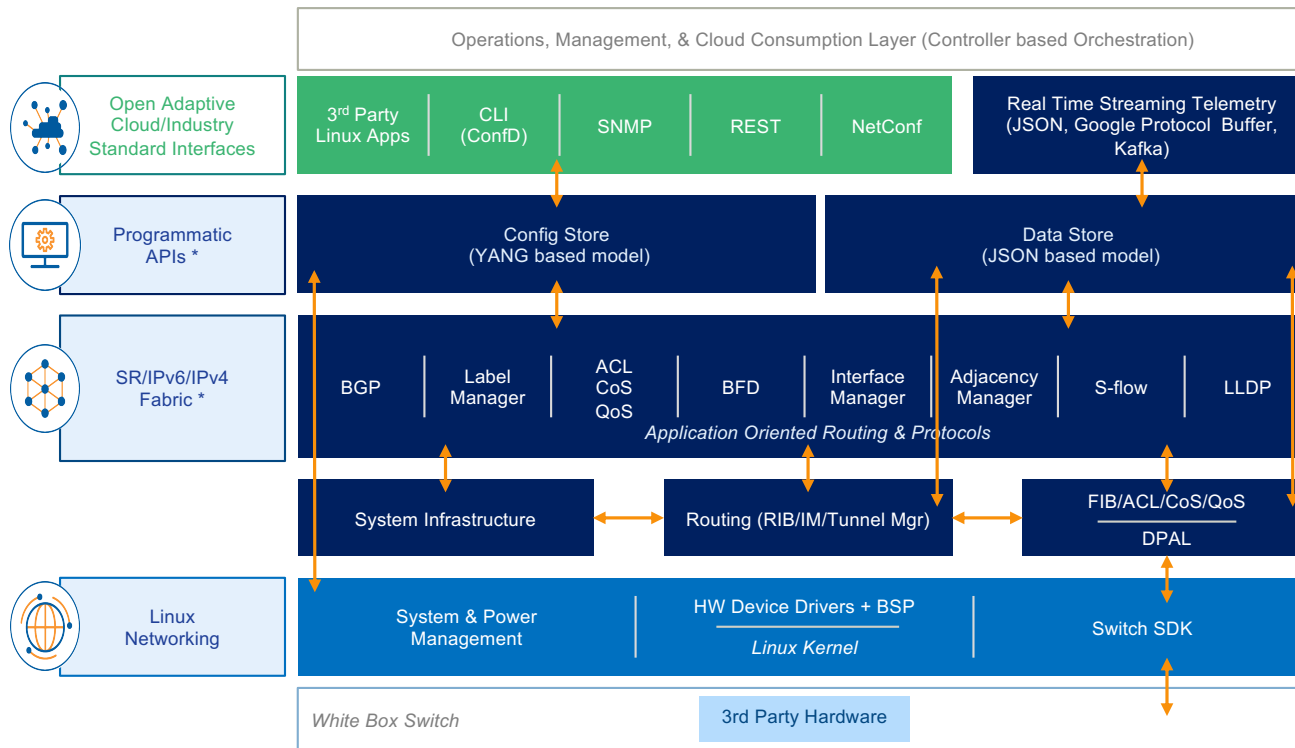
- **Real-Time Visibility**
  - Line-rate sFlow streaming
  - Resource view before, during, and after DDoS attack
- **Dynamic Control**
  - Granular ACL rules
  - Real-time feedback to DDoS detector
- **Security Automation**
  - BGP Flowspec-based signaling w/ DDoS controller
  - Open standards-based APIs

# Why Open Integration for Networking ?



- **Simplicity**
  - L3 Clos topology proven architecture of MSDCs
- **Operational Control & Flexibility**
  - Not flying blind... troubleshooting resource access
  - Automation built-in (provisioning, config changes & maintenance)
  - Minimize code bloat
  - Increased uptime and HA
- **Accelerate Innovation & Time to Deployment**
  - Streaming telemetry: Control and Data Plane events
  - Proactive measures in managing infrastructure
- **Supply Chain Flexibility**
  - Vendor independence
  - No forced upgrades & hardware obsolescence
- **TCO Savings (CapEx & OpEx)**
  - Realized cost leverage

# The ArcOS<sup>®</sup> Architecture



## ▪ Micro Services Architecture

- Components implemented in user space
- Minimal required support of Linux Kernel
- Container ready

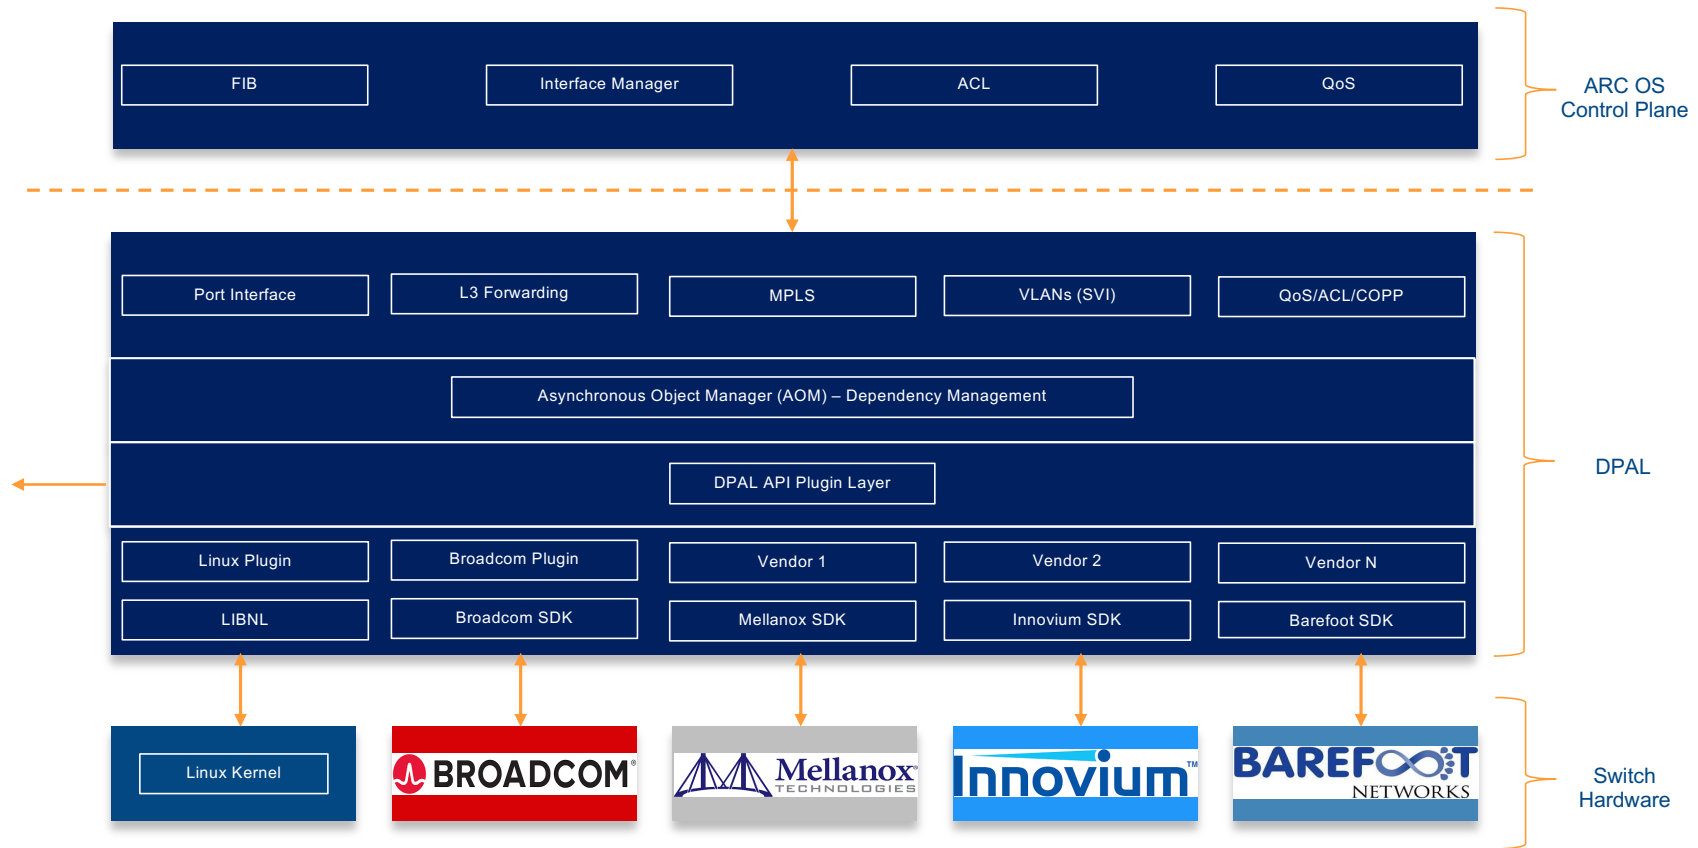
## ▪ Modularity

- Replace/Add/Use any component with minimal impact
- Every component implemented as a multi-threaded separate process

## ▪ Hardware Agnostic

- Switch and/or Server ready
- Intel and/or ARM processor support
- Dataplane Adaptation Layer (DPAL)

# ArcOS<sup>®</sup> Data Plane Adaptation Layer (DPAL<sup>™</sup>)



# Arcus Ready to Scale



- **ArcOS: Network Operating System (NOS)**
  - Carrier Grade at Internet Scale
  - Open, Composable and Micro-Services
  - Physical, Virtual & Cloud Environments Supported
- **ArcIQ: AI-Driven Network Operations Center (AI-NOC)**
  - Highly Correlated Data Analytics
  - Network Security, Traffic Engineering & Health
  - SaaS/Cloud Services based Offering



Thank You!

