Tutorial : Network ASIC Architecture Review and Case Study

Fuguang Shi (ASIC Director)

NANOG 95 , 2025

JUNIPER NETWORKS | Driven by Experience™

- Network ASIC Basics Review
  - Packet Processing Architectures
  - Buffer Architectures
- Network ASIC Trends
  - Landscape
  - Case study

Quick review of Network ASIC architecture

- Following few slides serve as quick review of network ASIC architecture presented by my colleague Nancy Shaw at NANOG93
- Here's link to her presentation and video
  - https://storage.googleapis.com/site-media-prod/meetings/NANOG93/5312/20250202_Shaw_Network_Asics_v1.pdf
  - https://www.youtube.com/watch?v=HUAnjNmZ0OQ

Network ASIC In The Router (Chassis System)

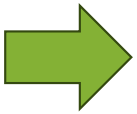

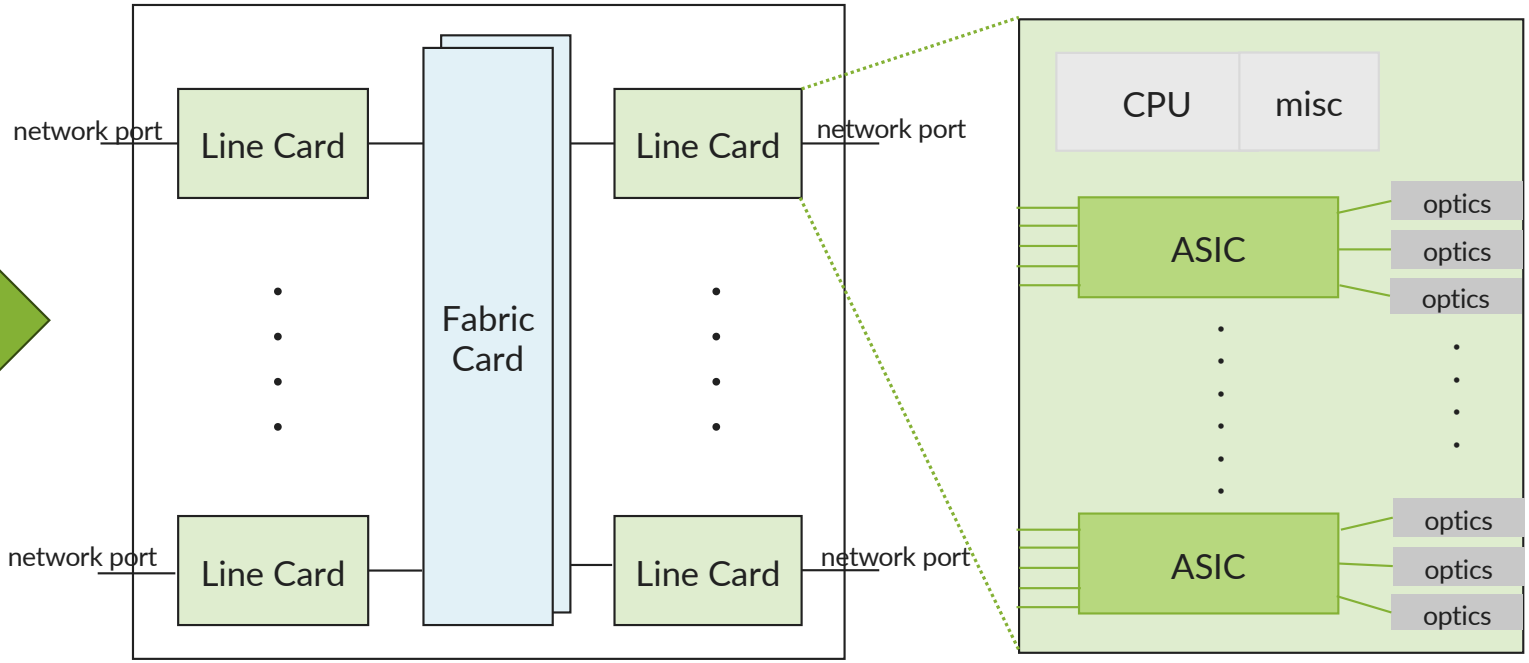

Network ASIC In The Router (Stand-alone Box)

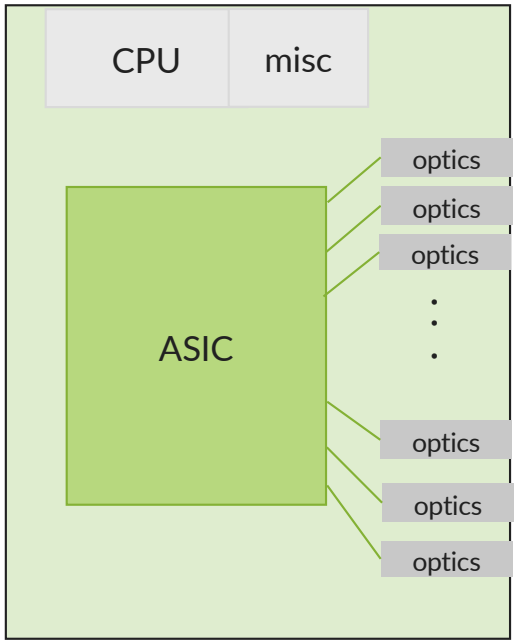

Packet Processing Architectures

Packet Processing Architectures

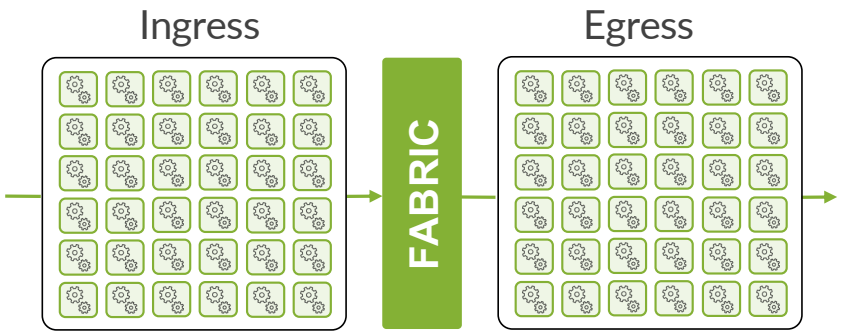

Multiple Packet Processing Engines (PPE) Architectures

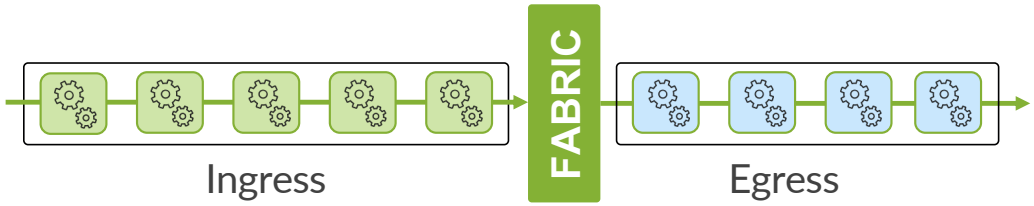

Pipeline Architectures

Packet Processing Architectures Pros and Cons

Multiple PPEs

- ✅ Infinite programmability
  - **All features** can be supported
- ✅ Highest level of flexibility
  - Each PPE can access a fungible memory
- ❗ Can't get the Power/Speed that pipelines offer

Pipeline

- ✅ Higher performance
  - Higher Speed / Lower power
- ✅ Lower latency
- ✅ Covers "most" of the major use-cases
- ✅ Can't reach the flexibility a multiple PPE architecture offers

Buffer Architectures

Buffering and Queues

- Buffer provides packet storage during periods of transient congestion
- Queues store packets according to priority or Quality of Service (QoS)

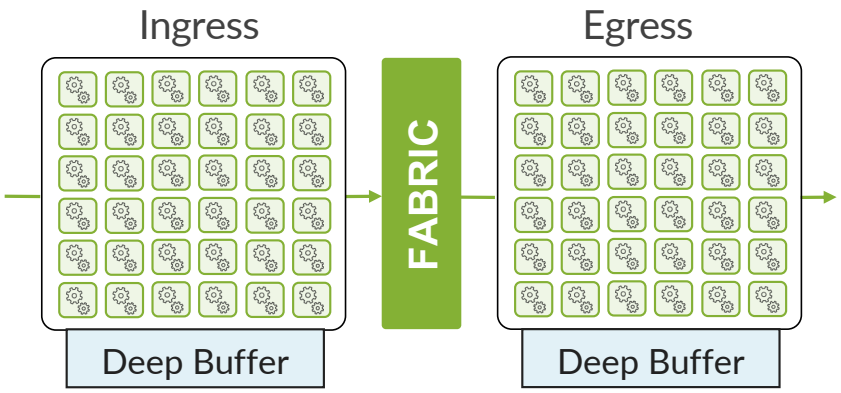

Two-Stage

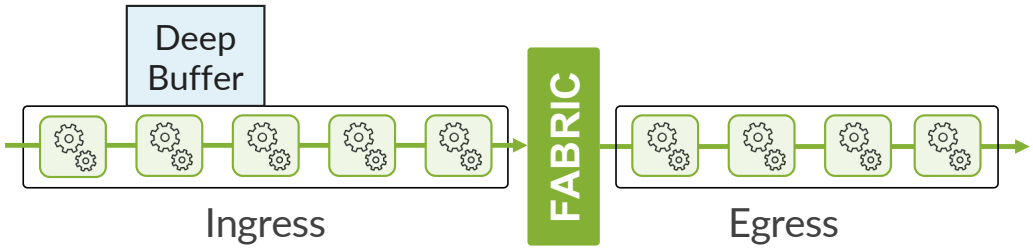

Ingress-Only Single-Stage

Two-Stage Buffering Architecture

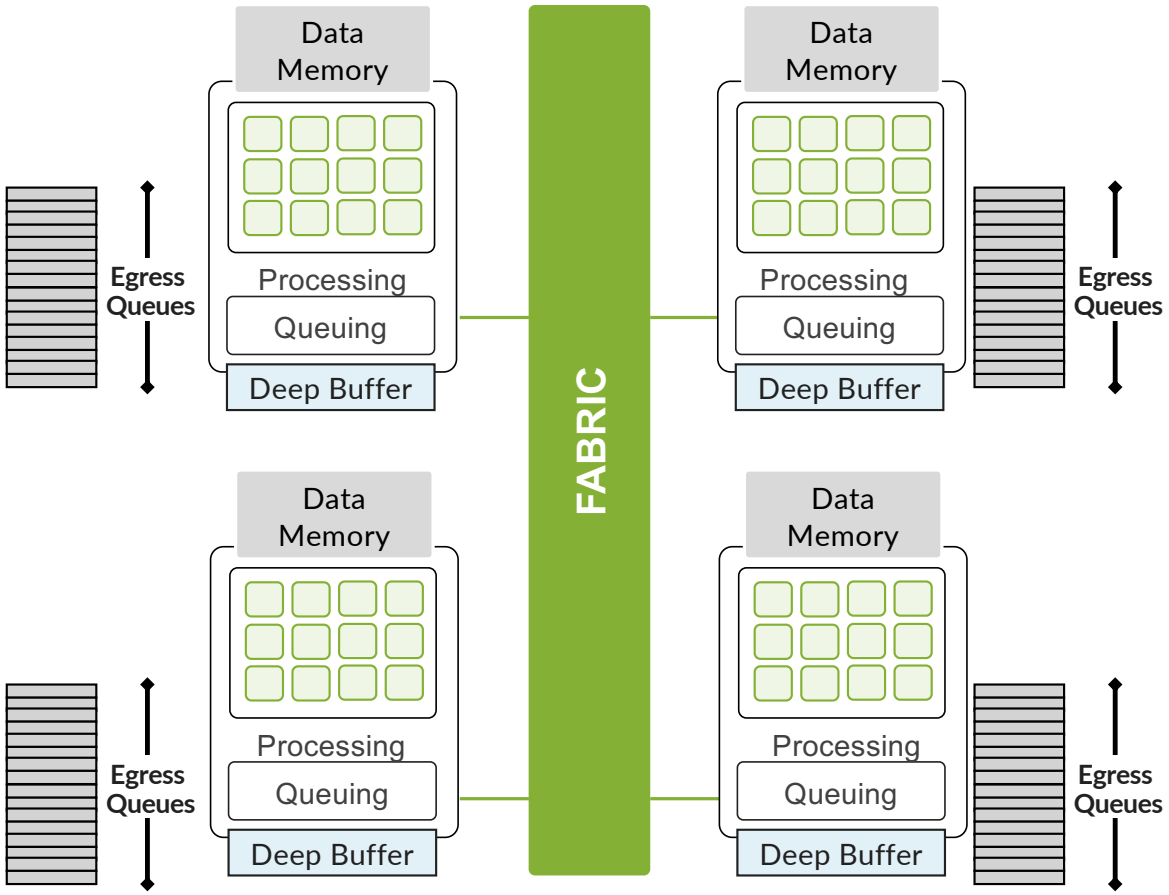

- Packet gets buffered first on ingress to be scheduled across fabric
  - Low number of queues
- Packet gets buffered a second time on egress to be scheduled out the local port
  - Large number of queues
  - All queues are only for “local” egress ports
- **More line cards, more ASICs, more queues!**
  - Total queues = #cards x #ASICs x #queues/ASIC

Ingress-Only Buffering: VOQ system

- VOQ = Virtual Output Queue
- No head of line blocking

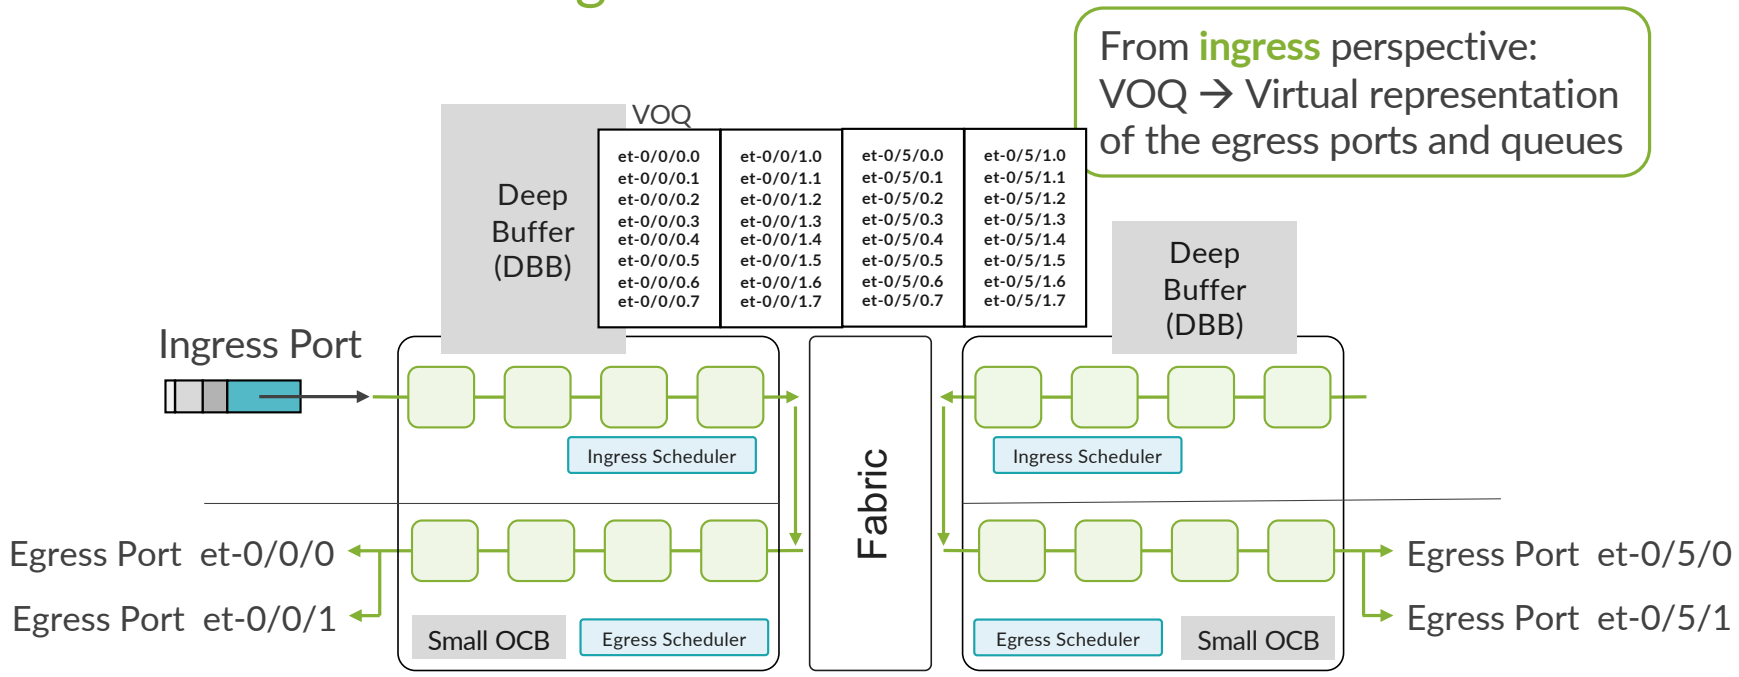

Buffering Architectures Pros and Cons

Two-Stage Buffering

- ✅ Many more queues per egress port
- ✅ Supports complex Subscriber and QoS management
- ❗ Head of Line Blocking
- ❗ Packet write/read to external memory twice in the system

Ingress-Only Single-Stage Buffering

- ✅ No Head of line blocking
- ✅ Packet write/read to external memory only once in the system
- ❗ Low Scale Queueing

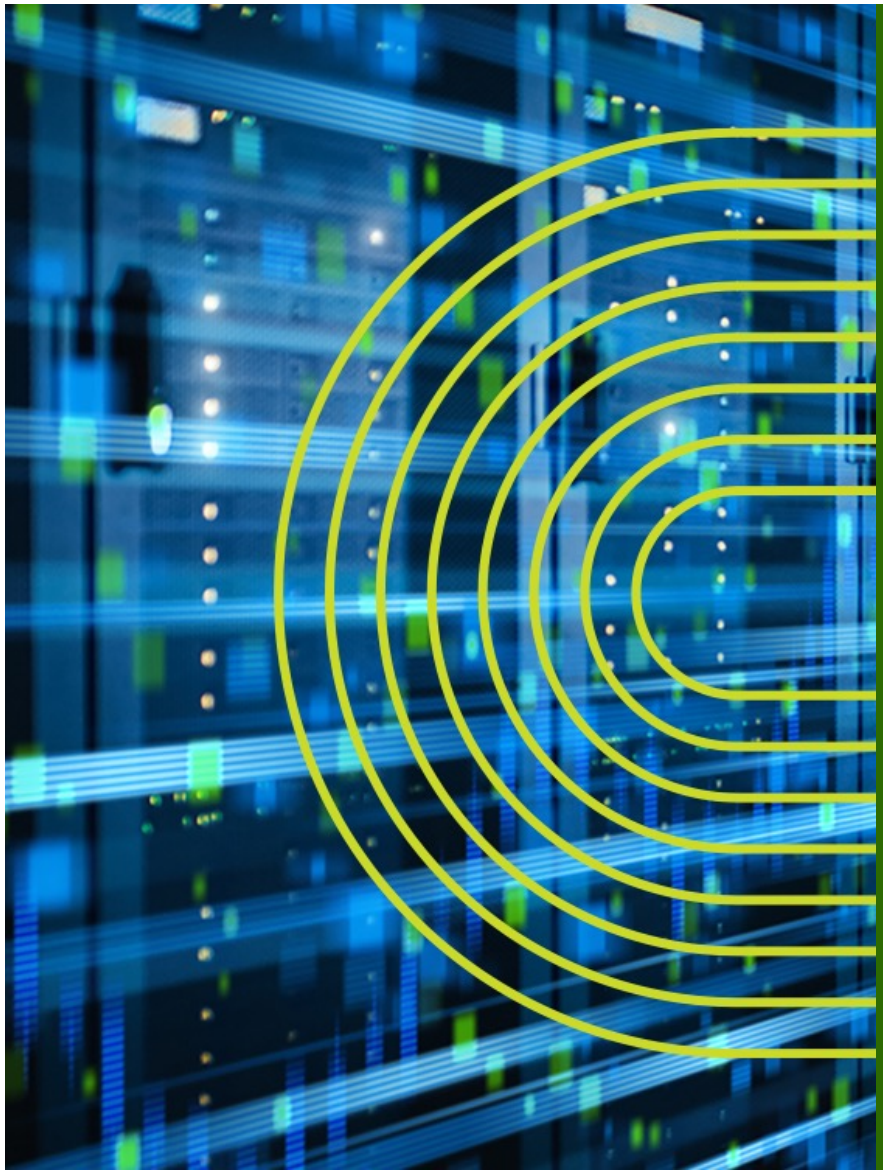

NETWORKING INDUSTRY TRENDS AND MARKET SEGMENTS WITH DIFFERENT ARCHITECTURES

Juniper Networks

Recent Networking Industry Trends

- Bandwidth growth continues with no end in sight!
- AI is main force today to drive some key technology for networking
- Increased network specialization with different equipment requirements
- Leads to network silicon domain segmentation

Networking Silicon Domains

One Size Does Not Fit All

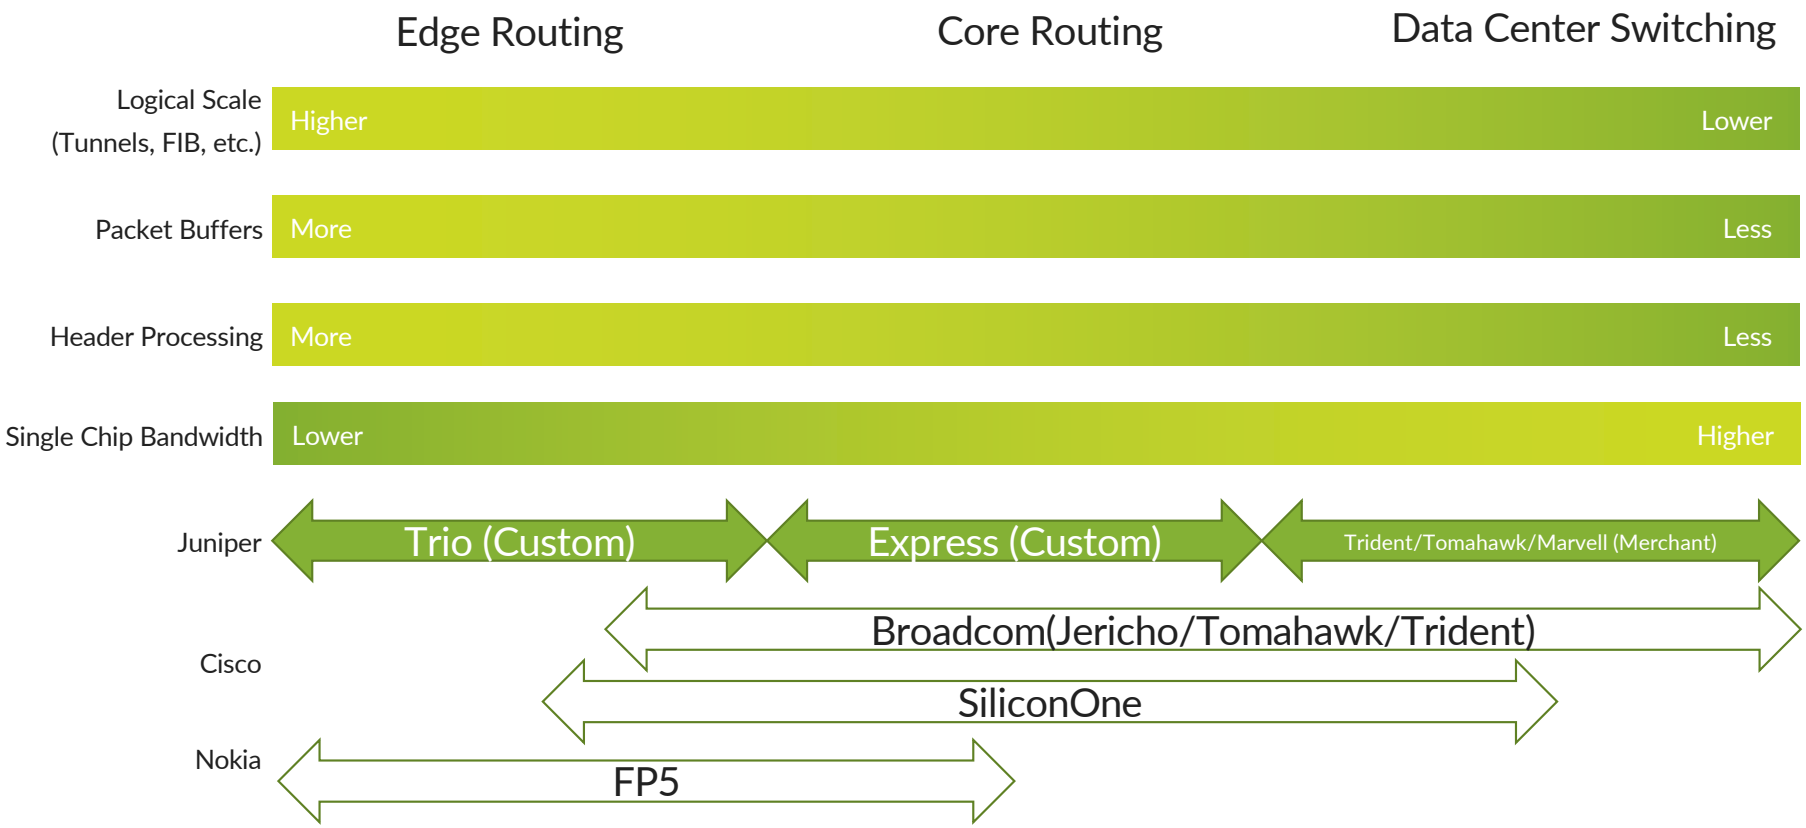

Two different Architectures

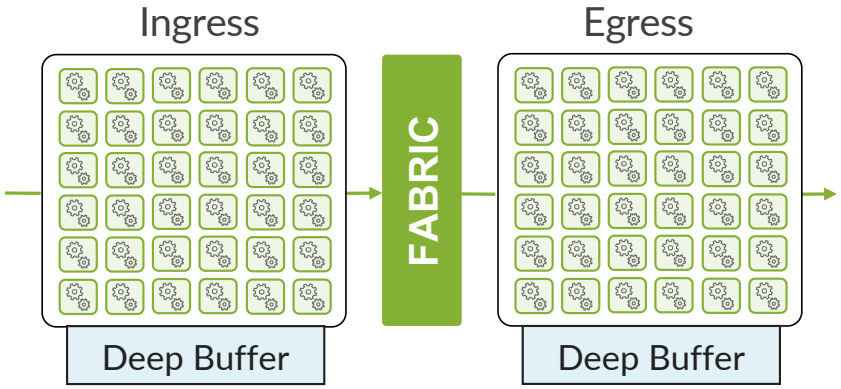

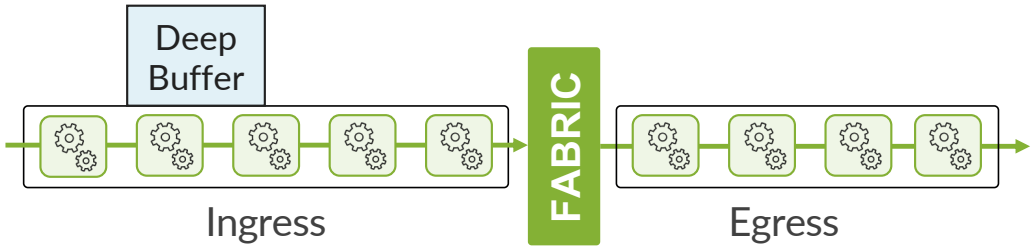

Multiple Packet Processing Engines
– Juniper Trio Architecture
- Nokia FP5

Pipeline Architectures
– Juniper Express Architecture
- Cisco SiliconOne
- Broadcom Jericho family

Network Taxonomy and Product Mapping

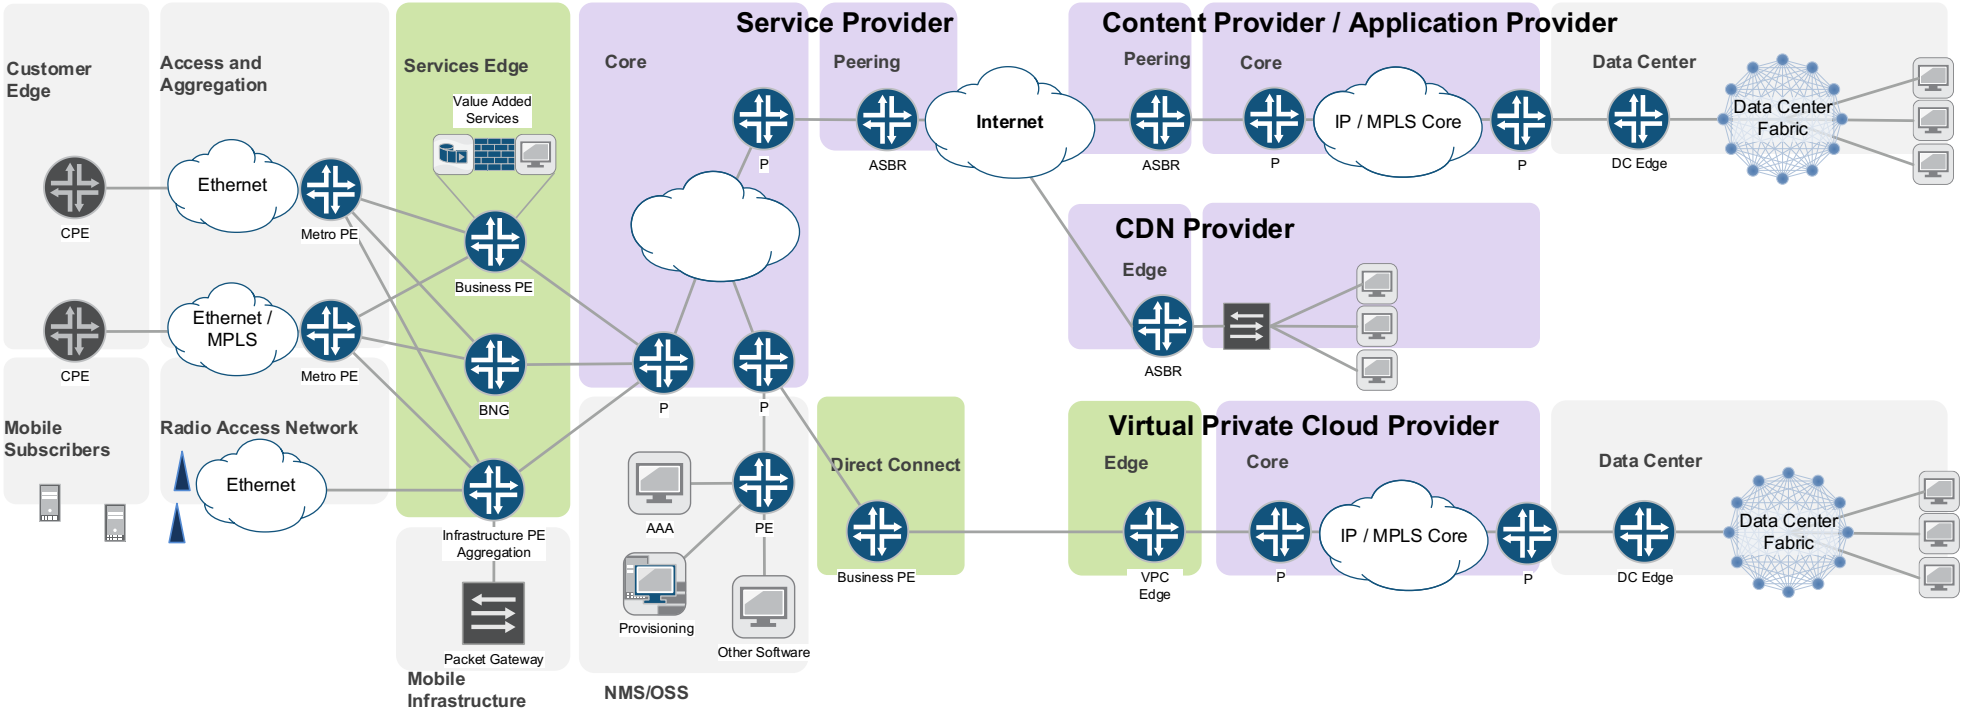

Pipeline architecture

PPE architecture

Merchant

AI Driving New Network Demand

Scale Out

Using Many Small-Form-Factor Systems

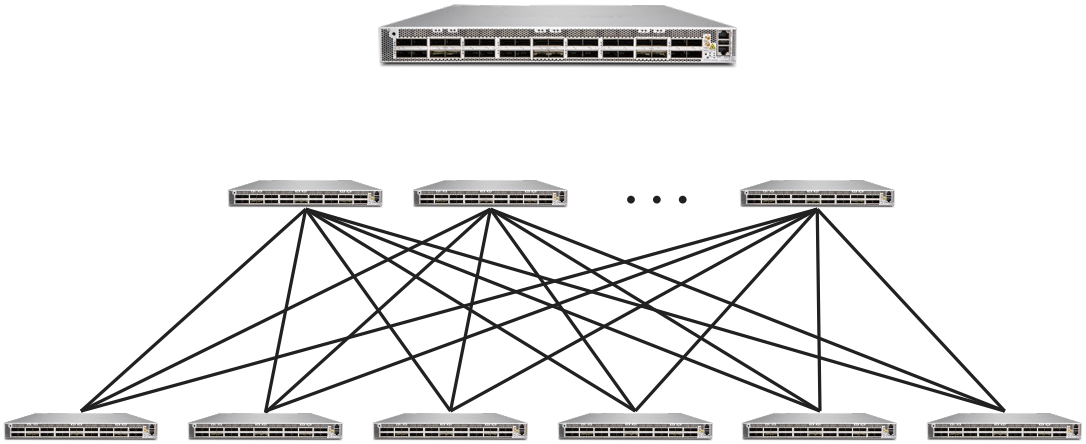

Scale Up & Scale out

Using Fewer Large Modular Systems

Midplane-less chassis

Efficient cooling

Architecture choices

- Multiple packet processing engines vs Pipeline architectures
- Two stage buffering vs Ingress-only VOQ scheme
- Deep buffer with external DRAM (HBM) vs Shallow buffer without external DRAM
- New emerging silicon technology
  - Advancing Silicon technology 3nm → 2nm → 1.4nm
  - New 2.5D/3D package technology – Chip on Wafer on Substrate (CoWoS)
  - Co-packaged optics (CPO)

Case Study

Express ASIC Evolution

	Express 1	Express 2	Express 3	Express 4	Express 5	Express 6
Slot Capacities, Gbps						
	120G/chip 1T / slot	500G/chip 3T / slot 3T / RU - fixed	1T/chip 8T and 16T fixed	3.6T/chip 14.4T / slot 9.6T fixed	28.8T/chip 28.8T / slot 28.8T fixed Designed for 800GE	57.6T/chip Design for 1600GE
Features						
						Distributed Chassis
					SRv6, HQoS, BIER	
				8M Statistics Counters		
			MACSEC up to 100GE	MACSEC up to 400GE	MACSEC up to 800GE	MACSEC up to 1600GE
		Telemetry			INT-MD	
		IPFIX / JFlow			Native IPFIX Export	
		2M+ FIB			10M+ FIB	
		Flexible Packet Processing				
	Lean Forwarding					
	Optimized for Density, Power and Scale					
	Broadway	Paradise	ZX	BT	BX	KX
	2012	2016	2019	2020	2022	2025

Trio ASIC Evolution

70Gbps -> 3.2Tbps ~46X

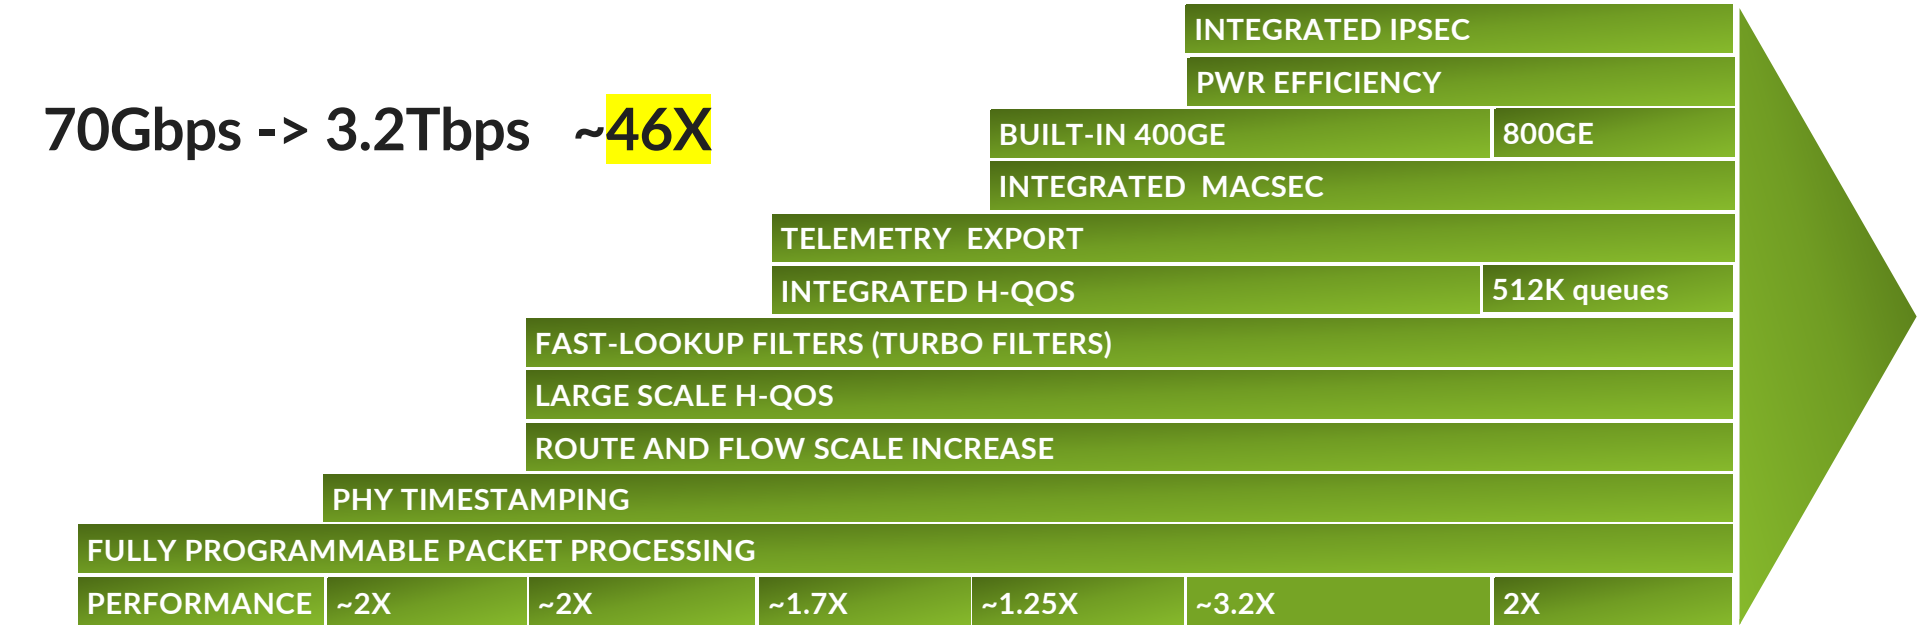

70Gbps	130Gbps	240Gbps	400Gbps	500Gbps	1.6Tbps	3.2Tbps
TRIO 1	TRIO 2	TRIO 3	TRIO 4	TRIO 5	TRIO 6	TRIO 7
2009	2012	2014	2016	2019	2021	2025
Multiple chips	Multiple chips	Multiple chips	Multiple chips, HMC	Single chip, 2.5D package, HBM	Single chip, 2.5D package, HBM	Single chip, 2.5D package, HBM

Broadcom Tomahawk Evolution

Feature	Tomahawk 1	Tomahawk 2	Tomahawk 3	Tomahawk 4	Tomahawk 5	Tomahawk 6
Max Bandwidth	3.2 Tb/s	6.4 Tb/s	12.8 Tb/s	25.6 Tb/s	51.2 Tb/s	102.4 Tb/s
Key Innovation	High-density 25/100G Ethernet for cloud scale networks.	Doubled performance while maintaining software compatibility.	56G-PAM4 SerDes, 400GbE, and advanced telemetry.	100G-PAM4 SerDes, higher performance, and broader interconnect options.	Monolithic 5nm die, AI/ML acceleration, and Cognitive Routing.	3nm chiplet design, Co-Packaged Optics, and 200G/100G PAM4 SerDes.
Notable Feature	Early cloud-scale adoption.	Maintained backward compatibility.	Enhanced telemetry, elephant flow detection, and shared buffers.	Expanded family with different SerDes variants.	Unmatched power efficiency for 51.2T, optimized for AI/ML.	Multi-die architecture, lowest power per bit, and support for CPO.
Process Node	28nm	16nm	16nm	7nm	5nm	3nm
SerDes Speed	25G NRZ	25G NRZ	50G PAM4	50G PAM4 / 100G PAM4	100G PAM4	100G PAM4 / 200G PAM4
year	2014	2016	2017	2019	2022	2025

Thank you