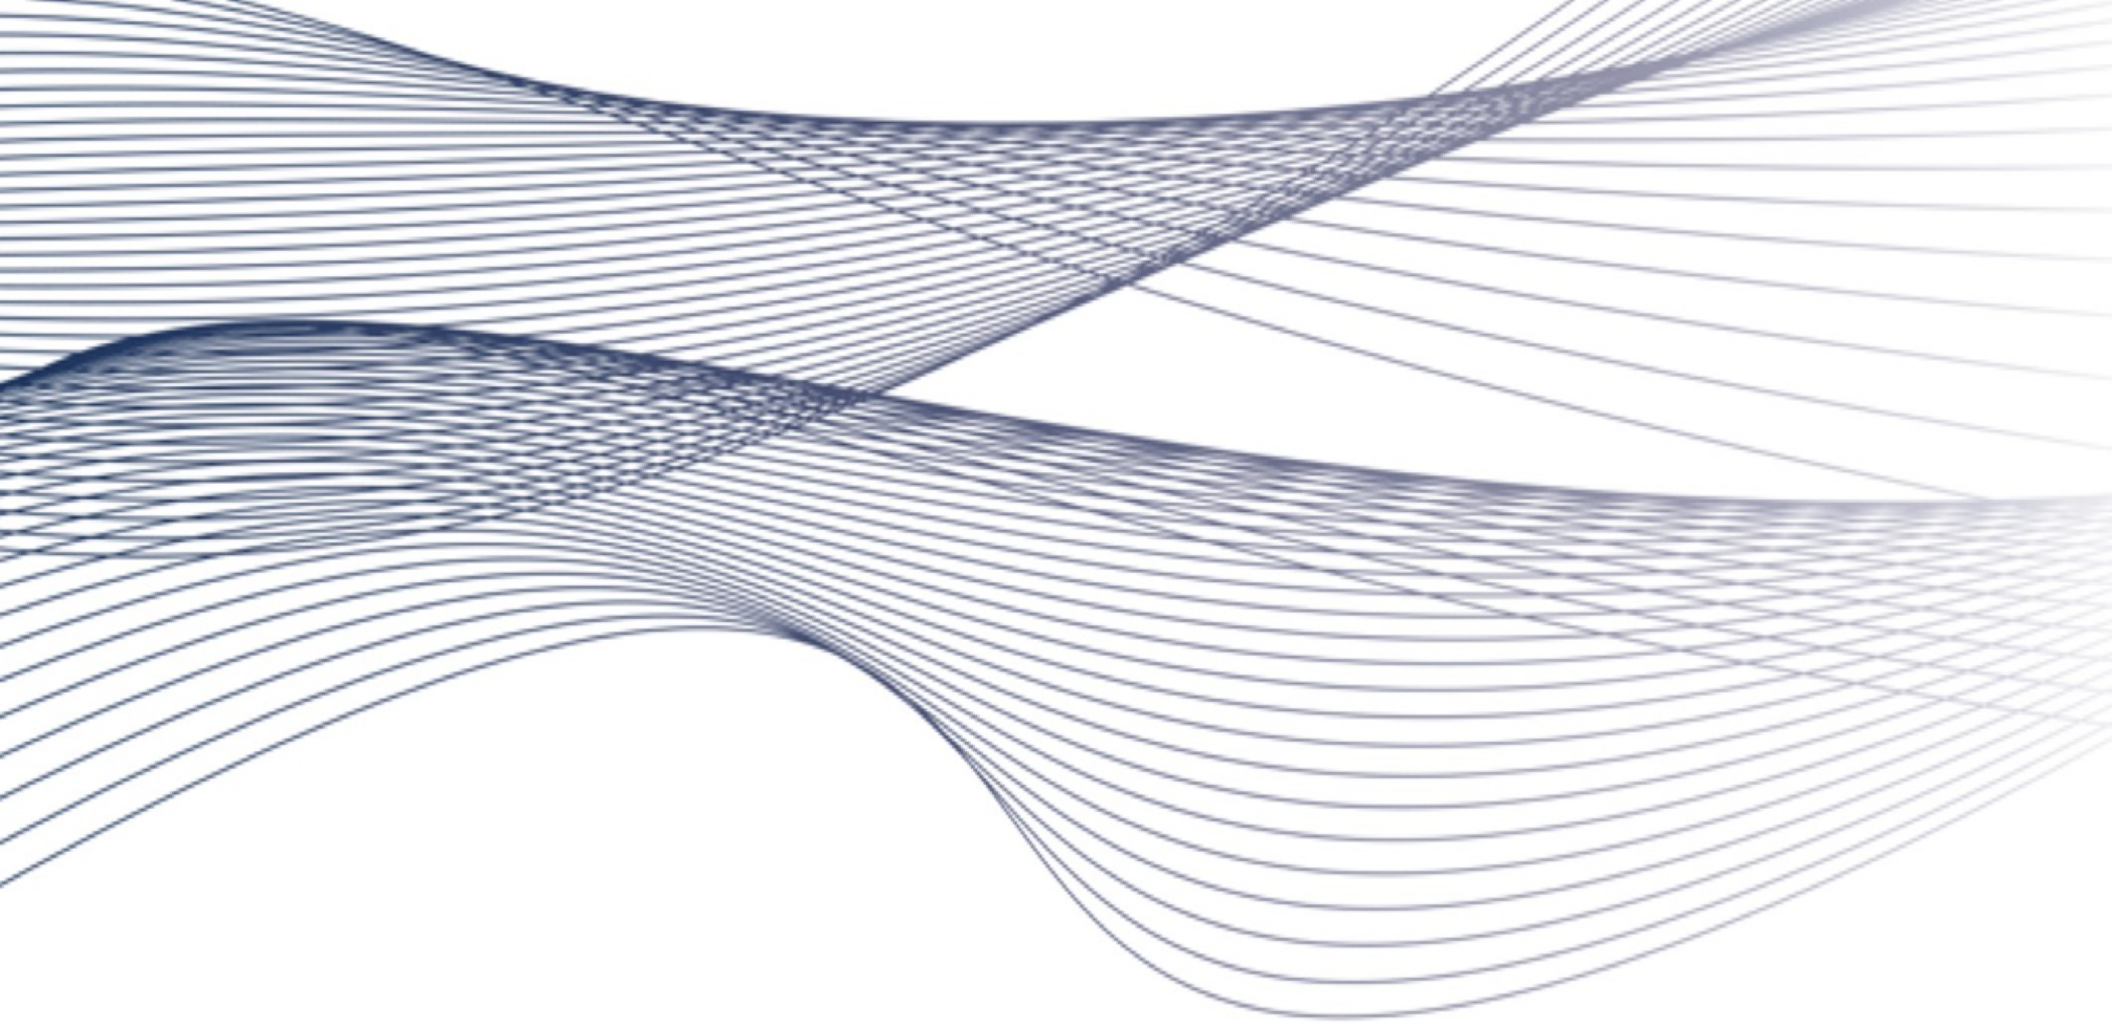

Merchant Silicon – Are switches… routers ?

Alex Nichol, Arista Technical lead

Arista Overview

- Company founded 2004, first product launch 2008
- Public trading company, IPO June 2014- $5B market cap
- 5 Million+ 10Gbe ports shipped, 3000 + customers
- Leading the GbE to 10GbE data center transition
- 2nd in Data Center switching/routing based on Market Share

Full Custom
Merchant Silicon

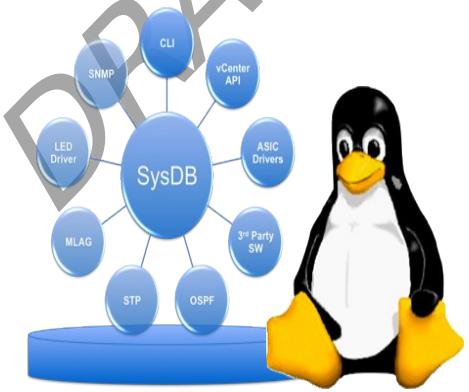

Single OS
Fully programmable
LINX OS

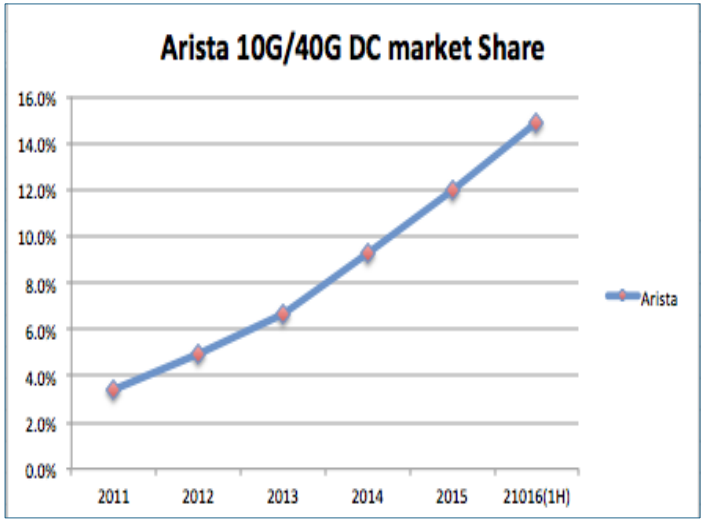

YoY market Share growth
Crehan Research Inc – Q2 2016

What do we mean by Merchant silicon ?

- ASIC Design approach
  - Top down approach
  - Focusing on the silicon functionality rather performance
  - Designed around pre-defined building block and libraries
  - Designed for functionality, not optimized for performance and cost

- Full Custom Design
  - Bottom up approach
  - Designed at the component level,
  - Allowing optimal clock rate/performance, lowest power, highest density
  - Delivers the optimal performance but can be costly process
  - Speed time cycle and reduce cost
    - Development done by third-party silicon manufactures not networking vendors
    - Broadcom, Intel, Dune, Fulcrum etc
  - Hence the term Merchant silicon

Merchant Silicon – Design trade-offs

- Designing merchant silicon – multi-dimension challenge trade-off within a fixed footprint
  - Pins 10G (for 10G/40G support) or 25G (for 10G/25G/50G/100G) – defining the throughput requirements
  - Layer 2/3 forwarding logic, tunnel support (GRE. MPLS, VXLAN etc)
  - Number of Pins on the silicon, defining the number ports, off- chip resources and fabric links if it's a multi-chip fabric design
  - Table sizes, MAC table (Exact match table), ARP tables (LEM), LPM for Layer 3 tables

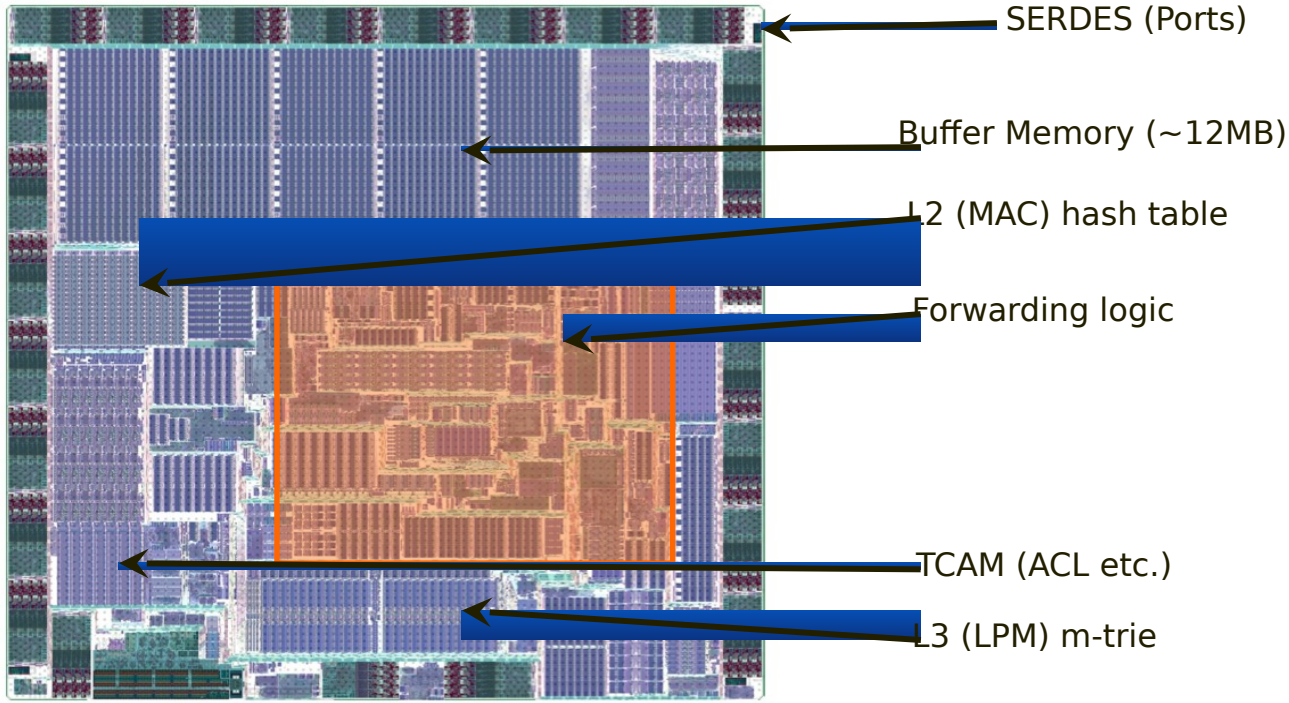

Merchant Silicon Landscape, its Evolution

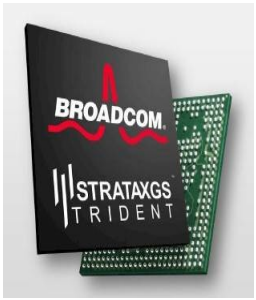

T+ → T2 → T2+, TH

10G SerDes 25G SerDes

64x10G 32x40G 32x100G

- 1RU
- SoC architecture
- High density 10/25G

→ T3?

→ TH2?

64x100G

DUNE NETWORKS

Petra → Arad → Jericho

8x10G 6x40G 6x100G

10G SerDes 25G SerDes

- Multi-chip
- Lossless Cell-based
- & VoQ architecture
- Deep buffers

→ Jericho2?

14x100G

intel

FULCRUM microsystems

FM4000 → FM6000

8x10G 64x10G

- Ultra-low latency (HFT)
- NAT

→ ?

Merchant Full Custom approach – Moore's law

Network ASIC performance has not increased like CPU performance.

In a 12 year span:

- Network ASIC increased: 10x
- CPU perf has increased: 64x

Performance is based on the **design process**

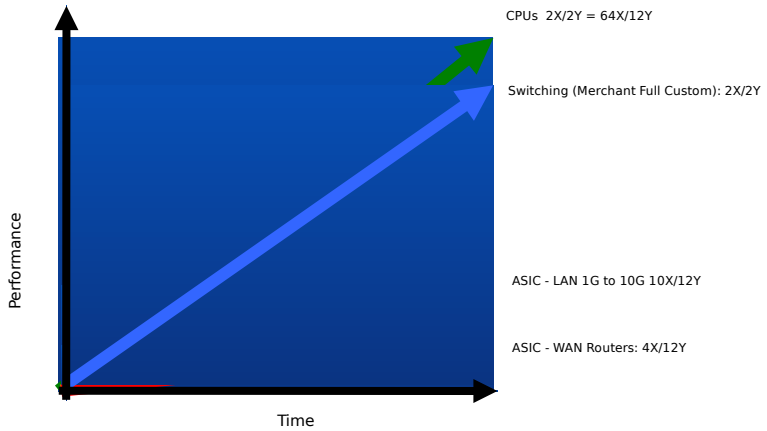

Why has Networking not kept up with Moore's Law?

ASIC = Application Specific Integration Circuit

- Top down design, independent of the layout
- Network Vendor focusing on the functionality not the implementation
- ASIC supplier does the physical implementation
- Difficult to achieve high clock rates and scale

Full Custom design flow

- Bottom up approach, chip vendor focus on potential implementation
- Chip design starts with the clock rate objective
- Data paths optimise to achieve the clock rates
- Only way to achieve high clock rates

Only the Full custom Chip will allows us to scale for the future

Bringing Merchant Silicon to the Routing Market

- Look at the routing market
  - The domain of the latest vendor ASIC
  - Changing with the latest – Merchant silicon

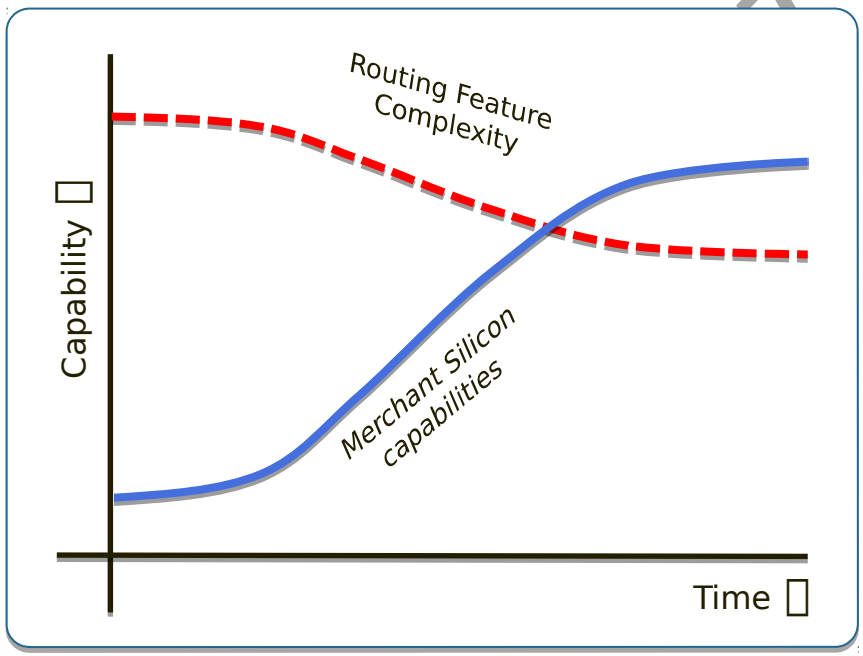

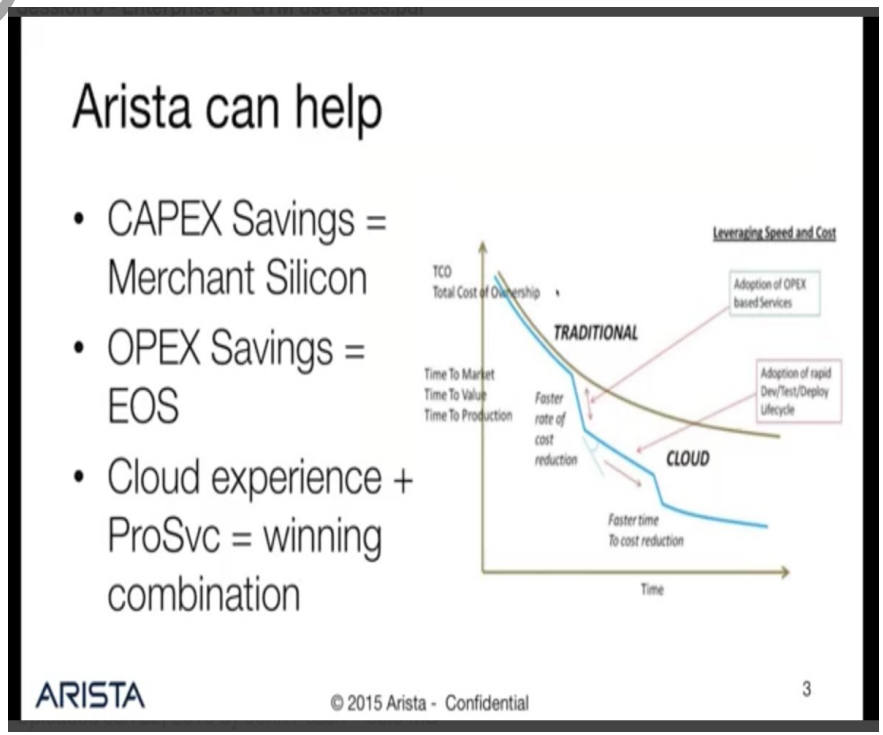

Evolution of 100GE Densities

96 x 100GE, 3W per 10G port

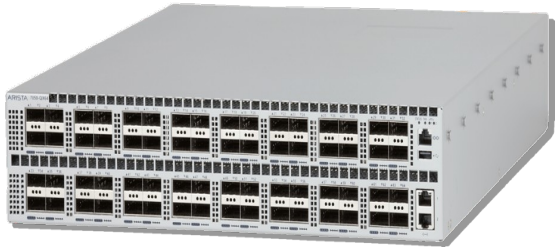

64 x QSFP 40GE in 2RU

48x10GE - 2x100GE

2015 → Today

36 x QSFP 100GE per LC
432 x 100G per chassis

48 x QSFP 100GE & 4 x 40G

12 x QSFP 100GE & 24 x 40G

Evolution of 100GE Densities

96 x 100GE, 3W per 10G port

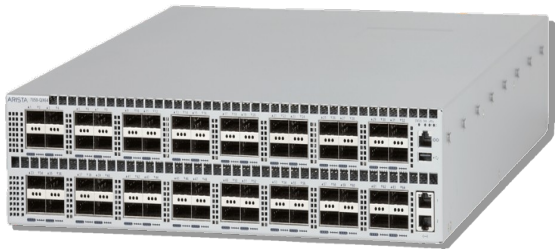

64 x QSFP 40GE in 2RU

48x10GE - 2x100GE

2015

Today

✓ Full Internet table

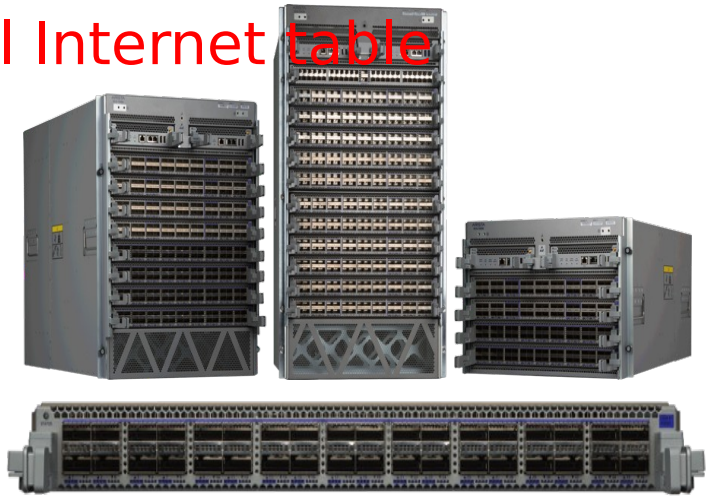

36 x QSFP 100GE per LC
432 x 100G per chassis

✓ Full Internet table

48 x QSFP 100GE & 4 x 40G

✓ Full Internet table

12 x QSFP 100GE & 24 x 40G

Evolution of 100GE Densities

96 x 100GE, 3W per 10G port

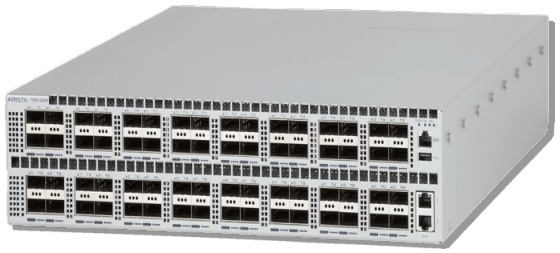

64 x QSFP 40GE in 2RU

48x10GE - 2x100GE

✓ < $3000 per 100G port

2015 → Today

✓ Full Internet table

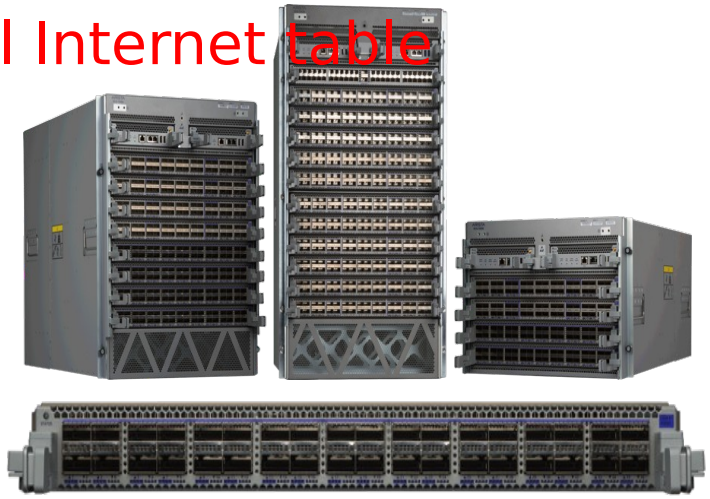

36 x QSFP 100GE per LC
432 x 100G per chassis

✓ Full Internet table

48 x QSFP 100GE & 4 x 40G

✓ Full Internet table

12 x QSFP 100GE & 24 x 40G

Merchant Silicon Use Cases for Service Providers

NFV / Cloud / DC

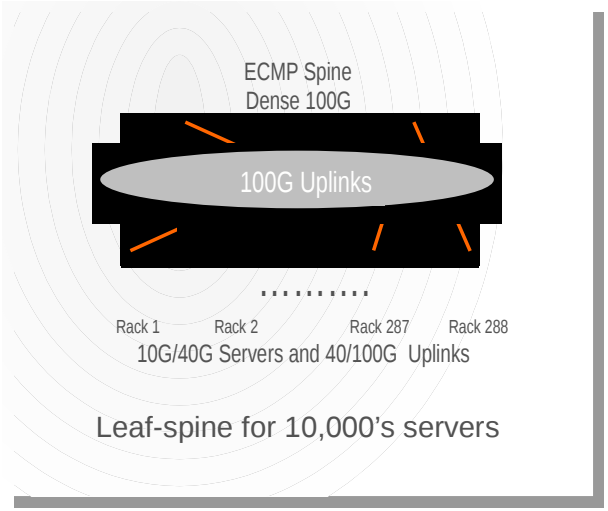

- Wire speed L2 & L3
- VXLAN and eVPN
- Flexible 40/100G uplinks
- Ease of migration
- Ultra deep buffers

Internet / content peering

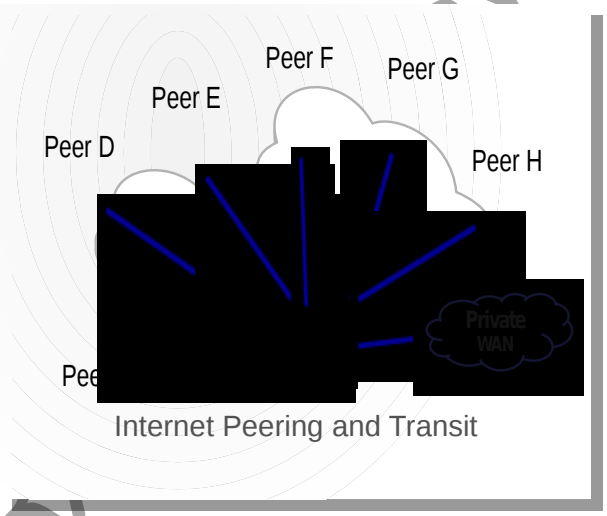

- Non-blocking 10G & 100G
- Large forwarding tables
- Alternative to legacy
- Dense 100G Capacity
- Ultra deep buffers
- Power efficient

WAN / Aggregation

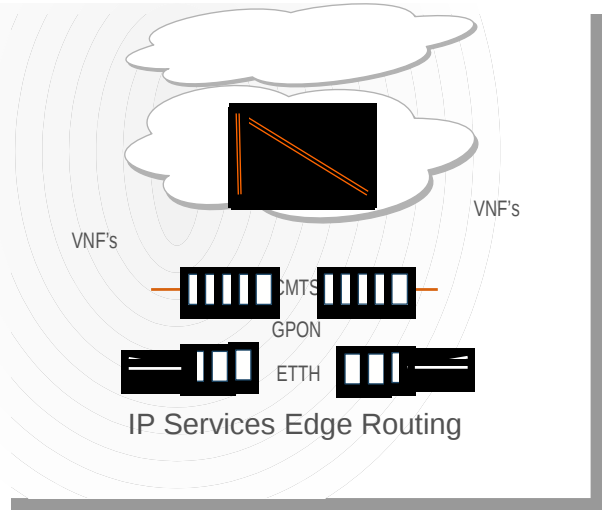

- L2, L3 Services, MPLS-Lite
- Segment Routing
- 10/40/100G Aggregation
- Power efficient
- Very cost effective bandwidth
- Ultra deep buffer architecture
- DC and NEBS