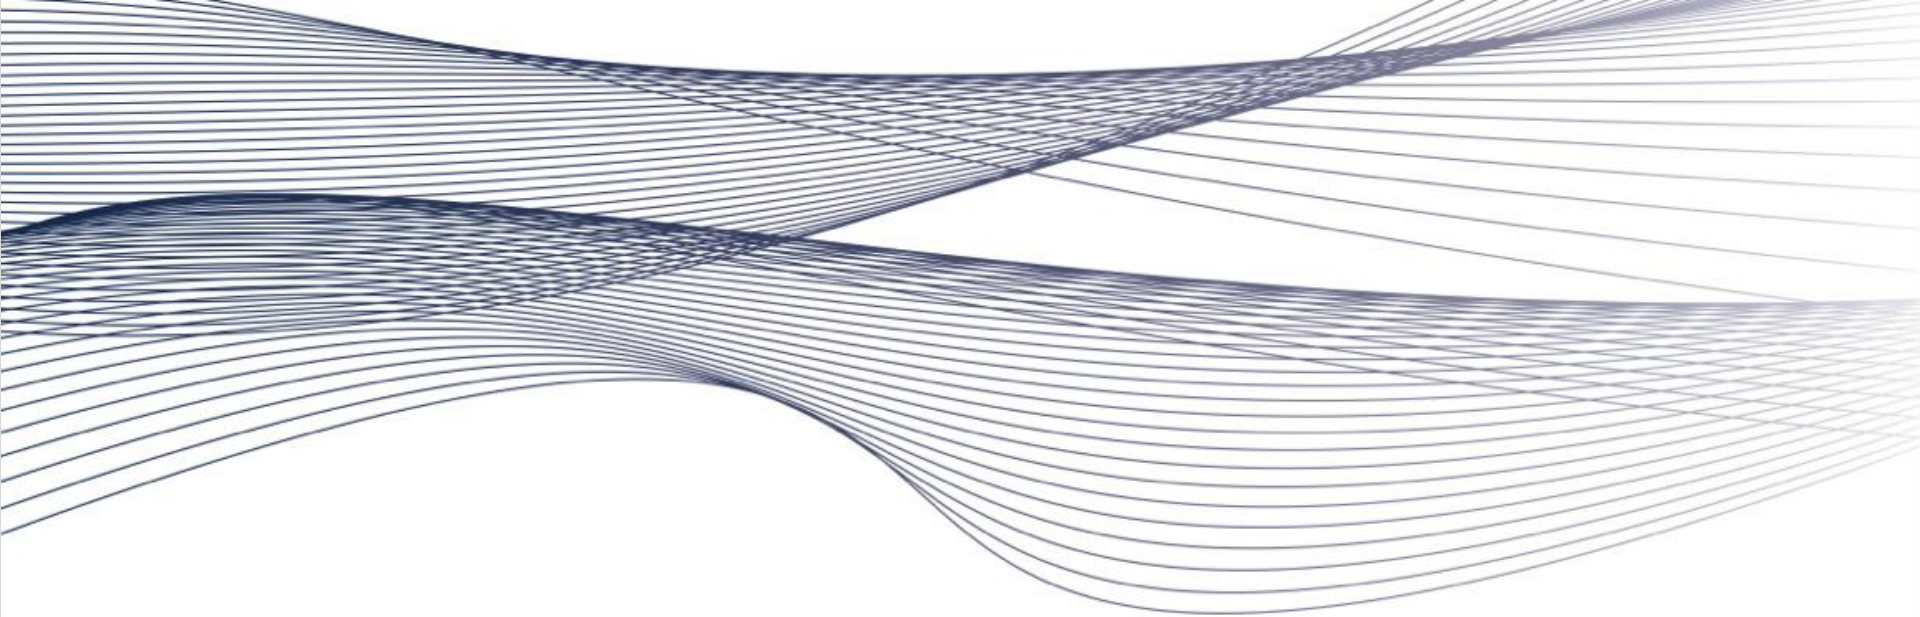

Disruptive Technologies

Marius Keown – Systems Engineer *@* Arista Networks

Moore’s Law

- Dr Gordon E. Moore – Co-Founder of Intel
- Predicted in 1965 the doubling of components per integrated circuit every year
- In 1975 revised the forecast to doubling every 2 years
- Used for decades as the guide for the industry for new products
- Nothing like this in the history of mankind

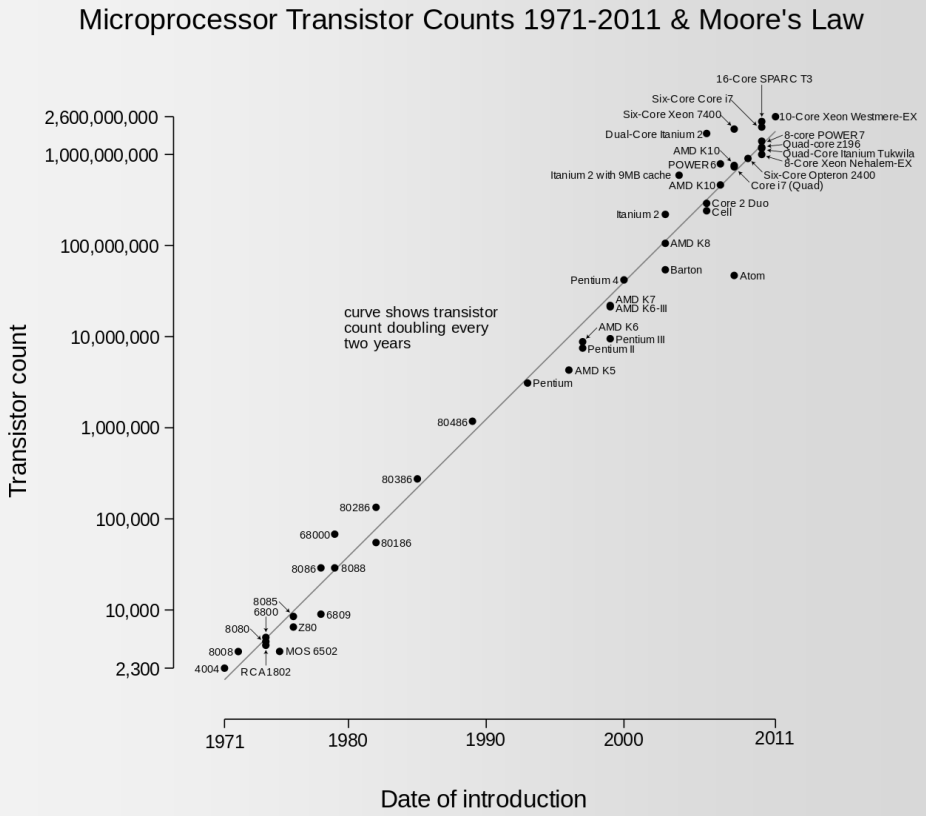

Impact on Technology Industry

- Economic Impact – Performance and Cost
- Modern Computing

Moore's Law and Networking

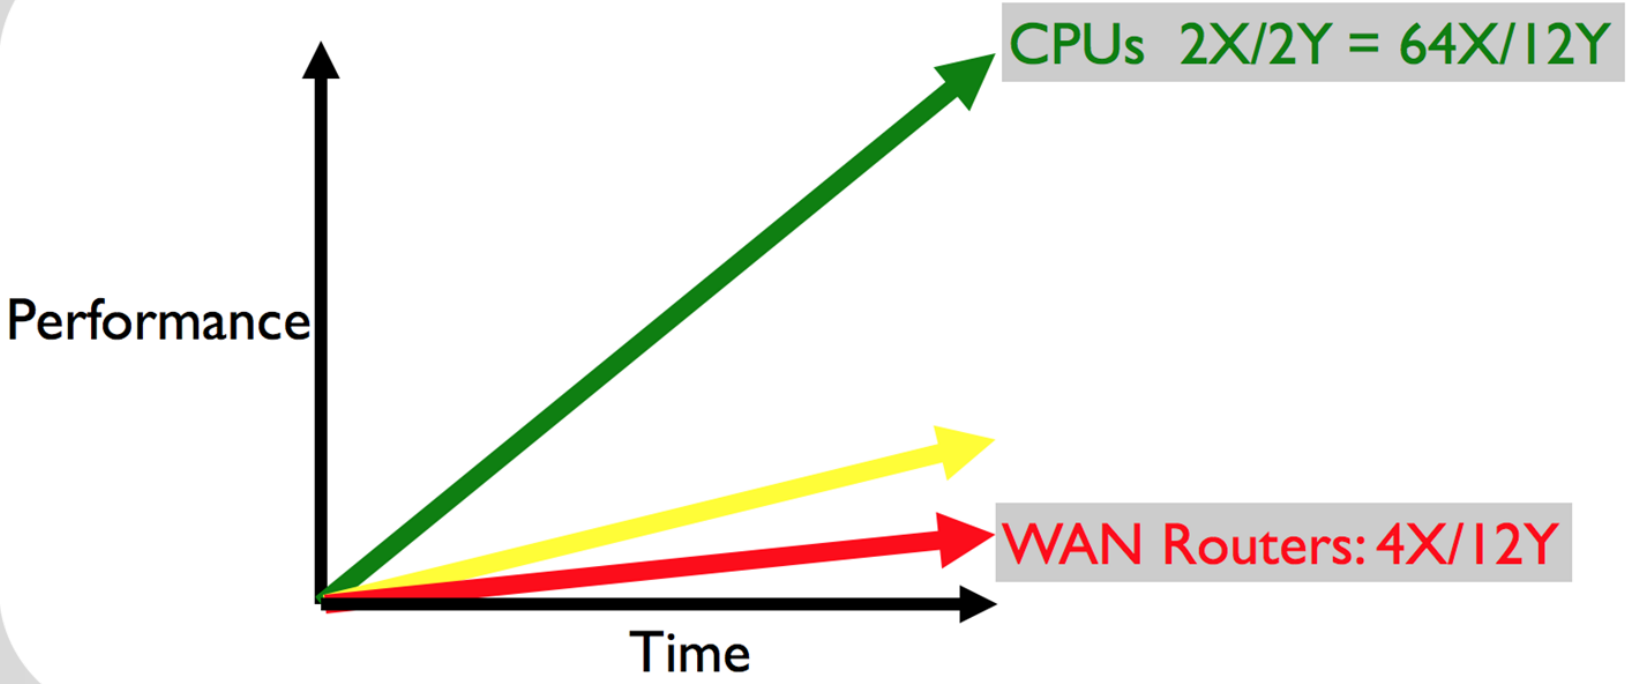

Why has Networking not kept up with Moore's Law?

Three main problems

- **Moore's Law applies to Transistors, not Speed**
  - Transistor count is doubling every 2 years
  - Transistor speed is only increasing slowly

- **Number of I/O pins per package basically fixed**
  - Limited by the area and package technology
  - Only improvement is increased I/O speed

- **Bandwidth ultimately limited by I/O capacity**
  - Throughput per chip = # IO Pins x Speed/IO
  - No matter how many transistors are on-chip

SERDES Speed (high density CMOS)

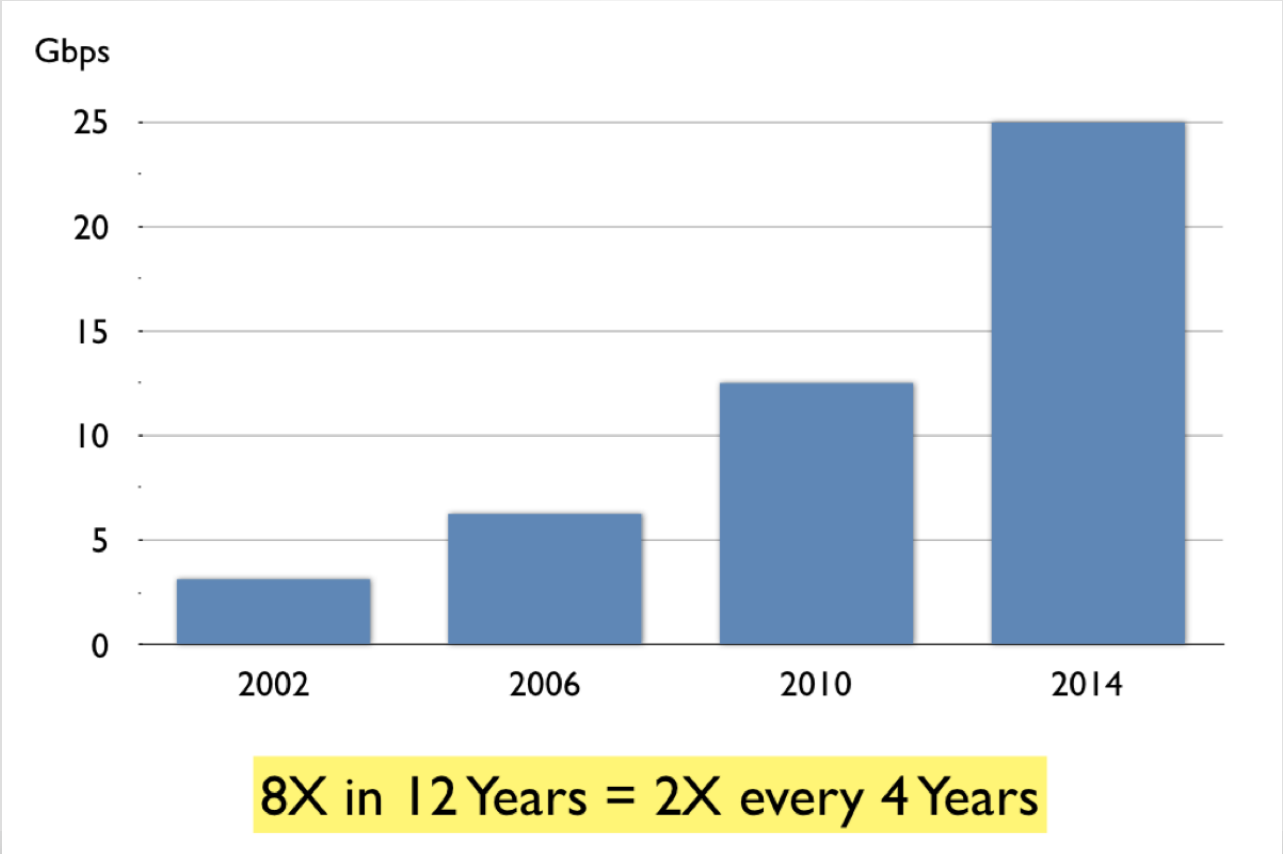

Number of SERDES per Package

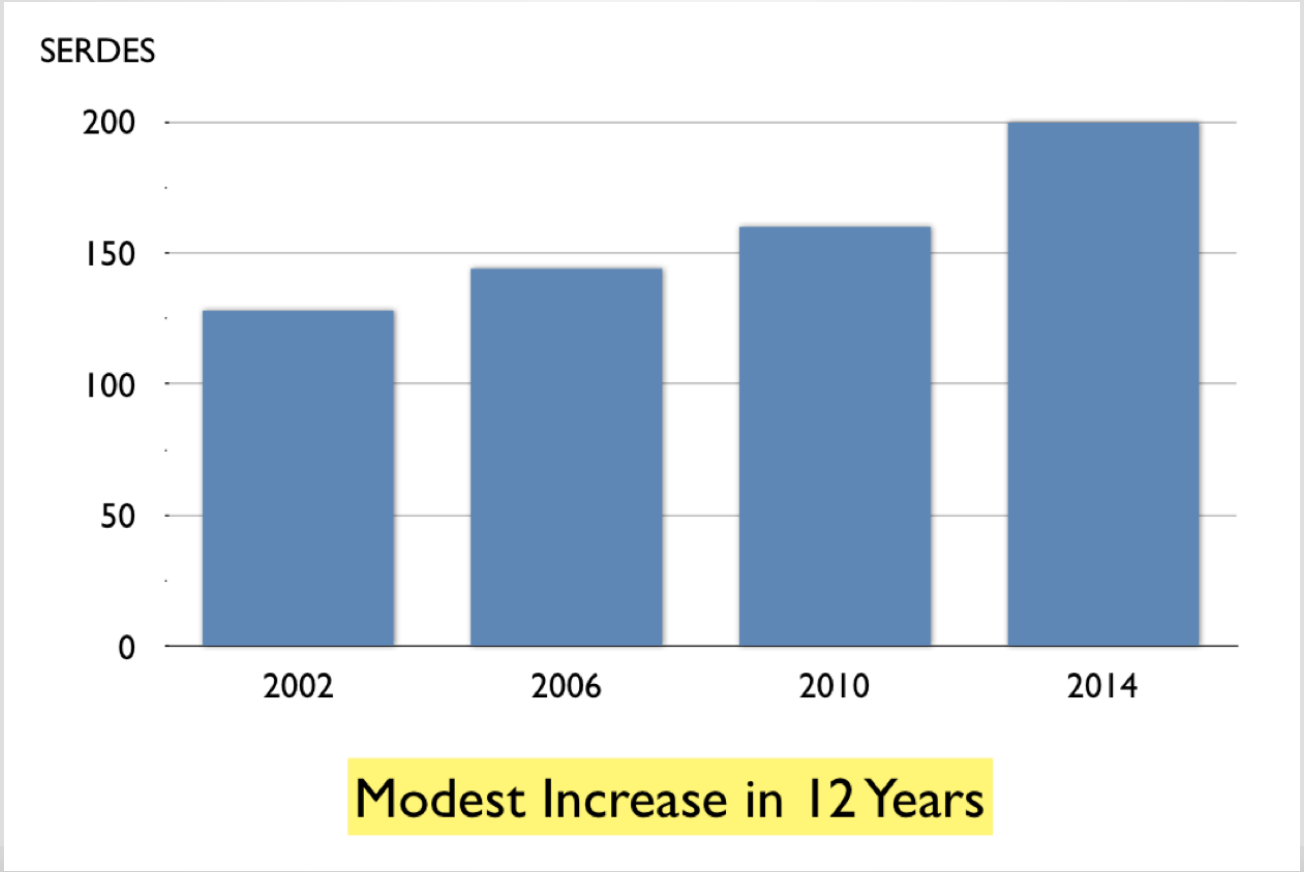

Maximum Throughput per Chip

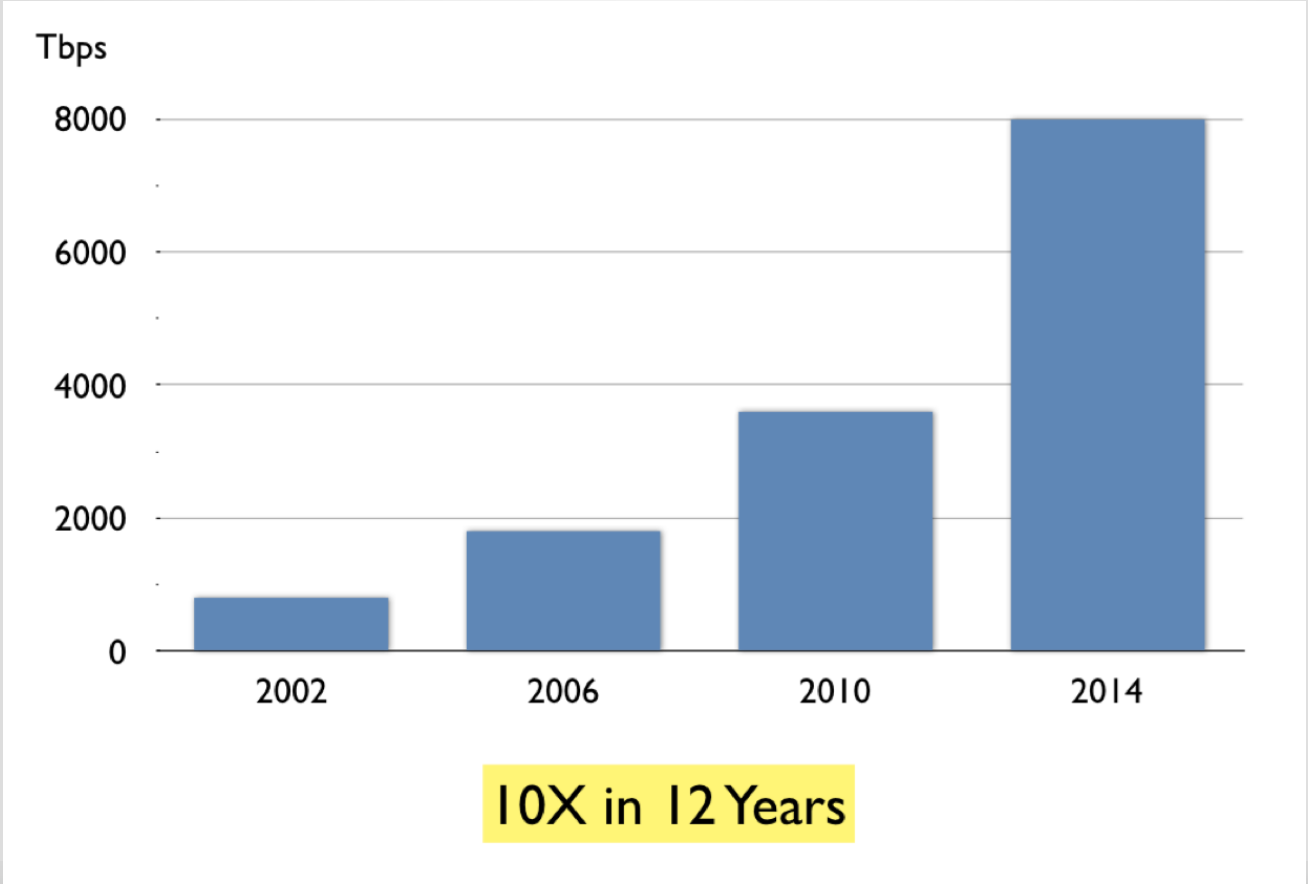

Moore's Law and Networking

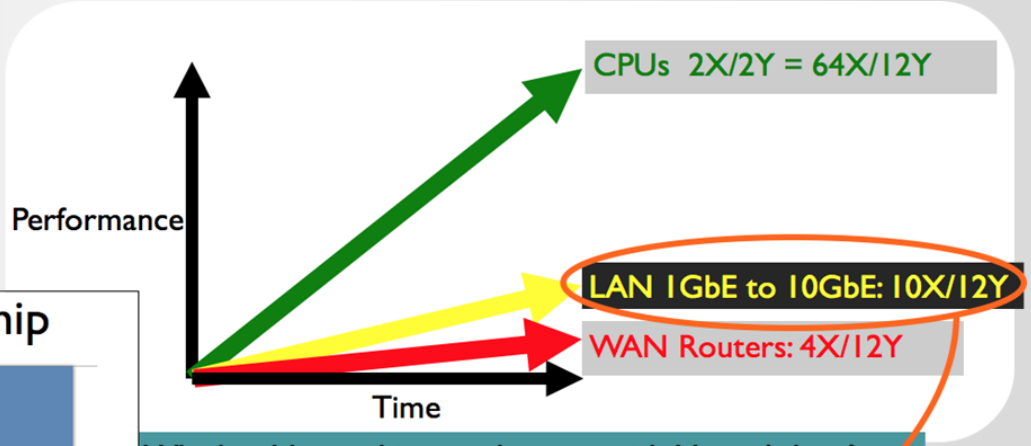

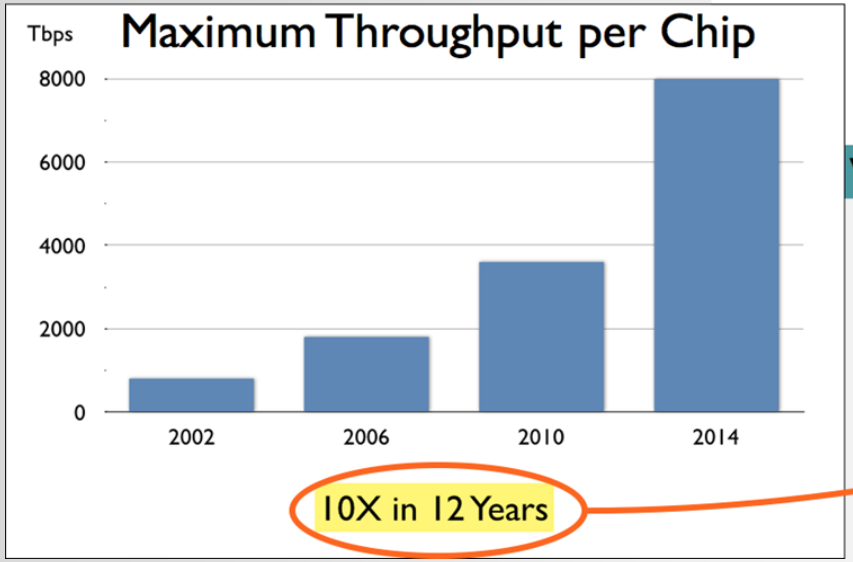

Network Switching Industry

Network ASIC performance has not increased like CPU performance.

In a 12 year span:

- Network ASIC increased: 10x
- CPU perf has increased: 64x
- Investment vs. ROI
- Low speeds, low port density, high power consumption
- Long and slow development cycle
- Inflexible to market changes

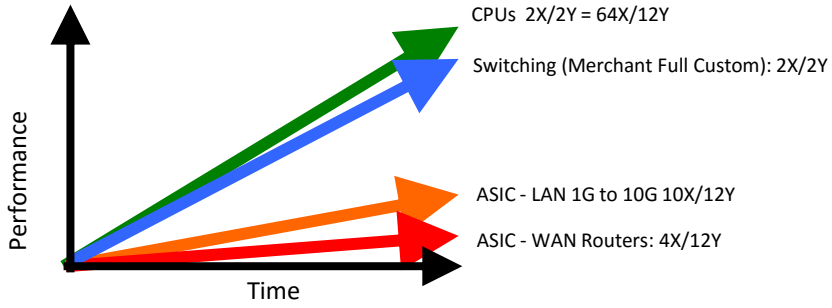

Why has Networking not kept up with Moore's Law?

ASIC = Application Specific Integrated 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 optimize to achieve the clock rates
- Only way to achieve high clock rates

Only the Full custom Chip will allow us to scale for the

Merchant Silicon 64-ports 10G Switch Chip

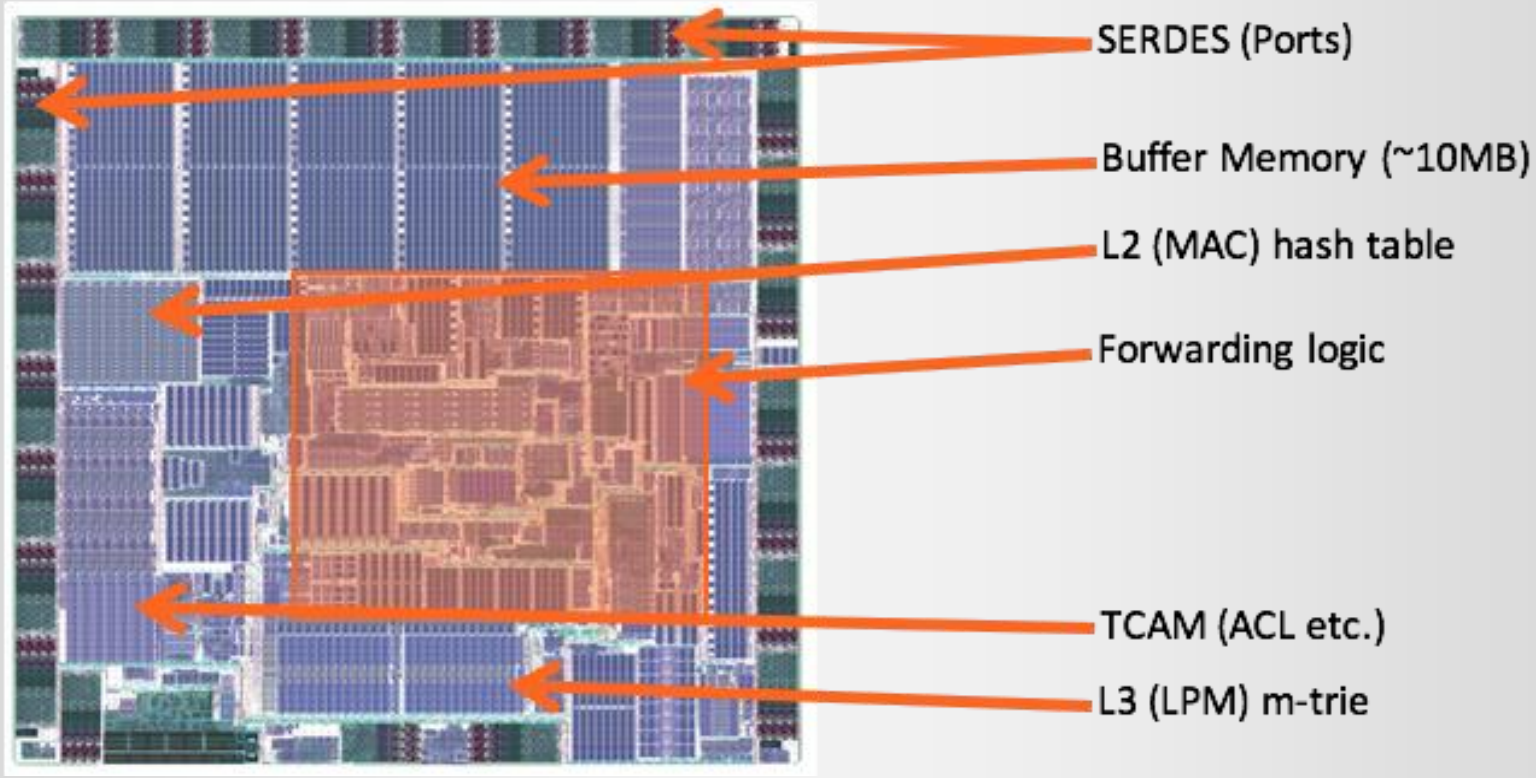

Port Density on Merchant Silicon

- **Broadcom Trident 2**
  - 128 x 10G
  - 32 x 40G
- **Broadcom Tomahawk**
  - 32 x 100G
  - 128 x 25G
- **Broadcom Jericho**
  - 6 x 100G, but with Big Buffers and large routing tables

Advantages of Merchant Silicon

- More ports per chip, increased throughput
- More room for additional logic/processing/functionality
- Less Chips:
  - Increased reliability, reduced complexity
  - Reduced latency (fewer chip crossing)
  - Consume less power (less chips less power draw)

Custom Design vs. ASIC Design

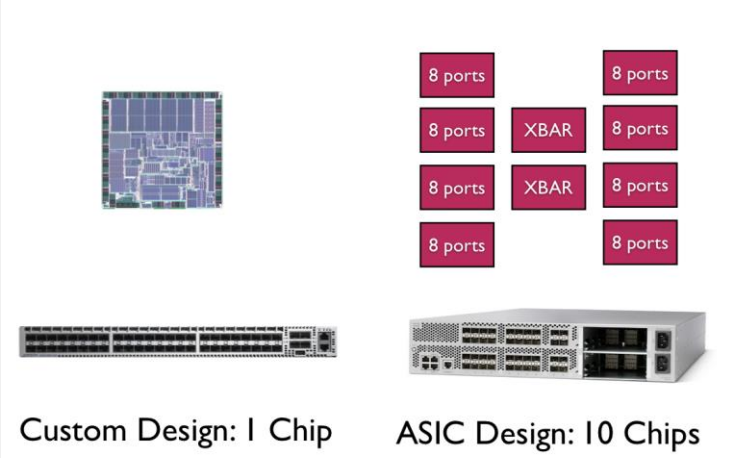

- Merchants' full custom chips are now on Moore's Law growth rate
- ASIC designs are NOT on Moore's Law growth

Merchant Silicon for SP

Hyper Connected World

We are in an age of exponential growth

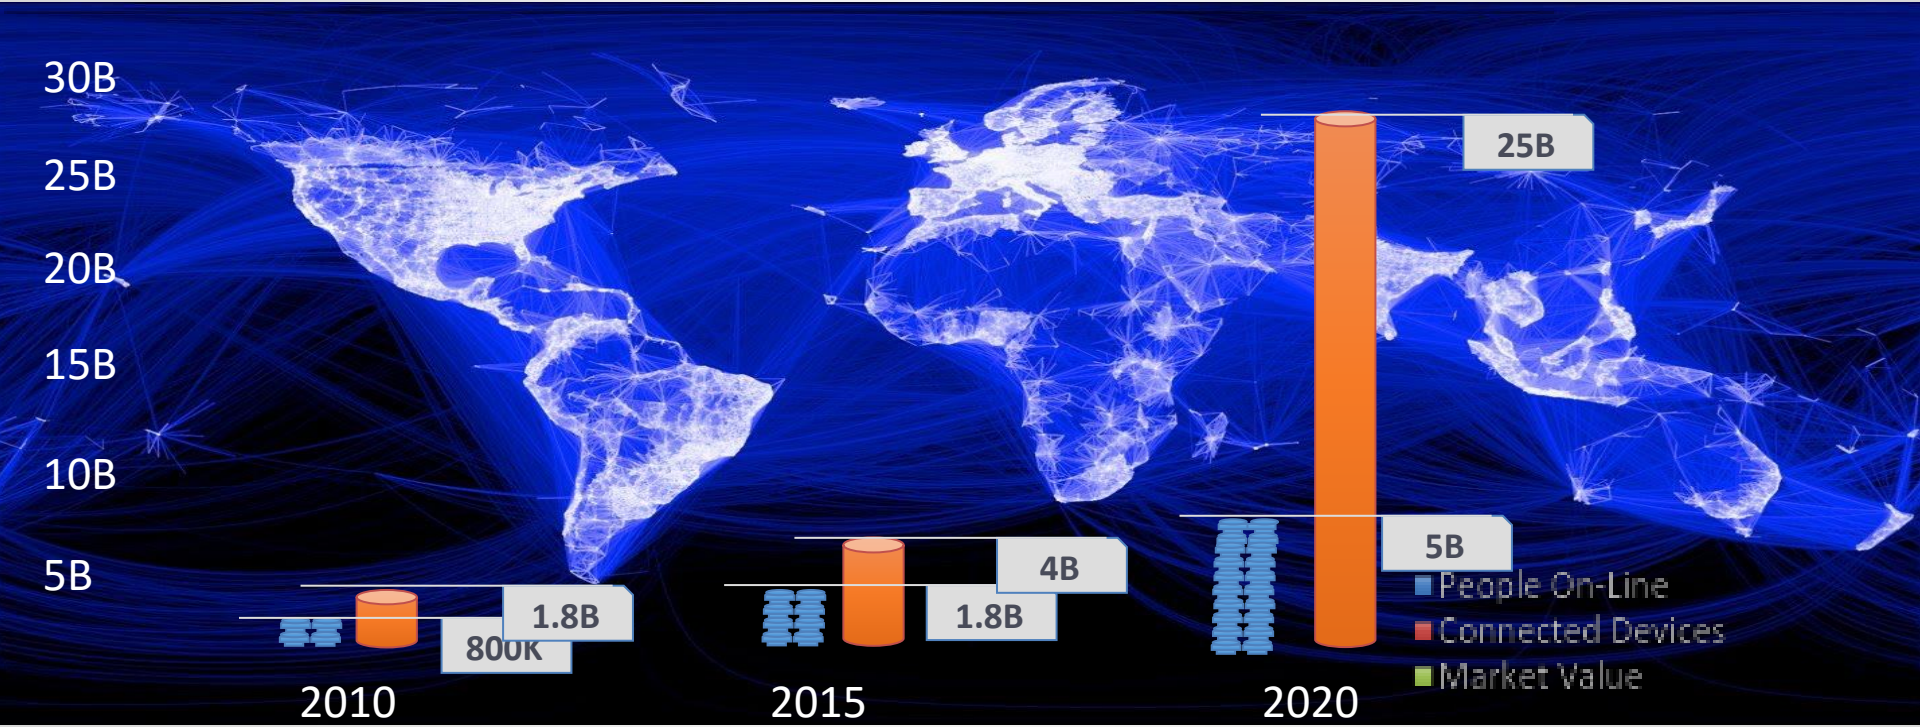

Sources: Cisco ,Gartner, CMA Research, EdgeConneX

SP Market Dynamics

Over-the-Top Video

Mobility

NFV

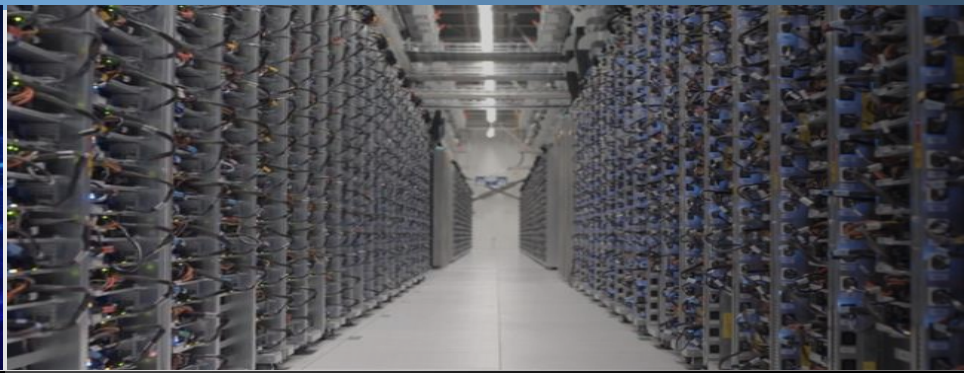

Hyper Scale DC

Emerging SP Challenge and Opportunity

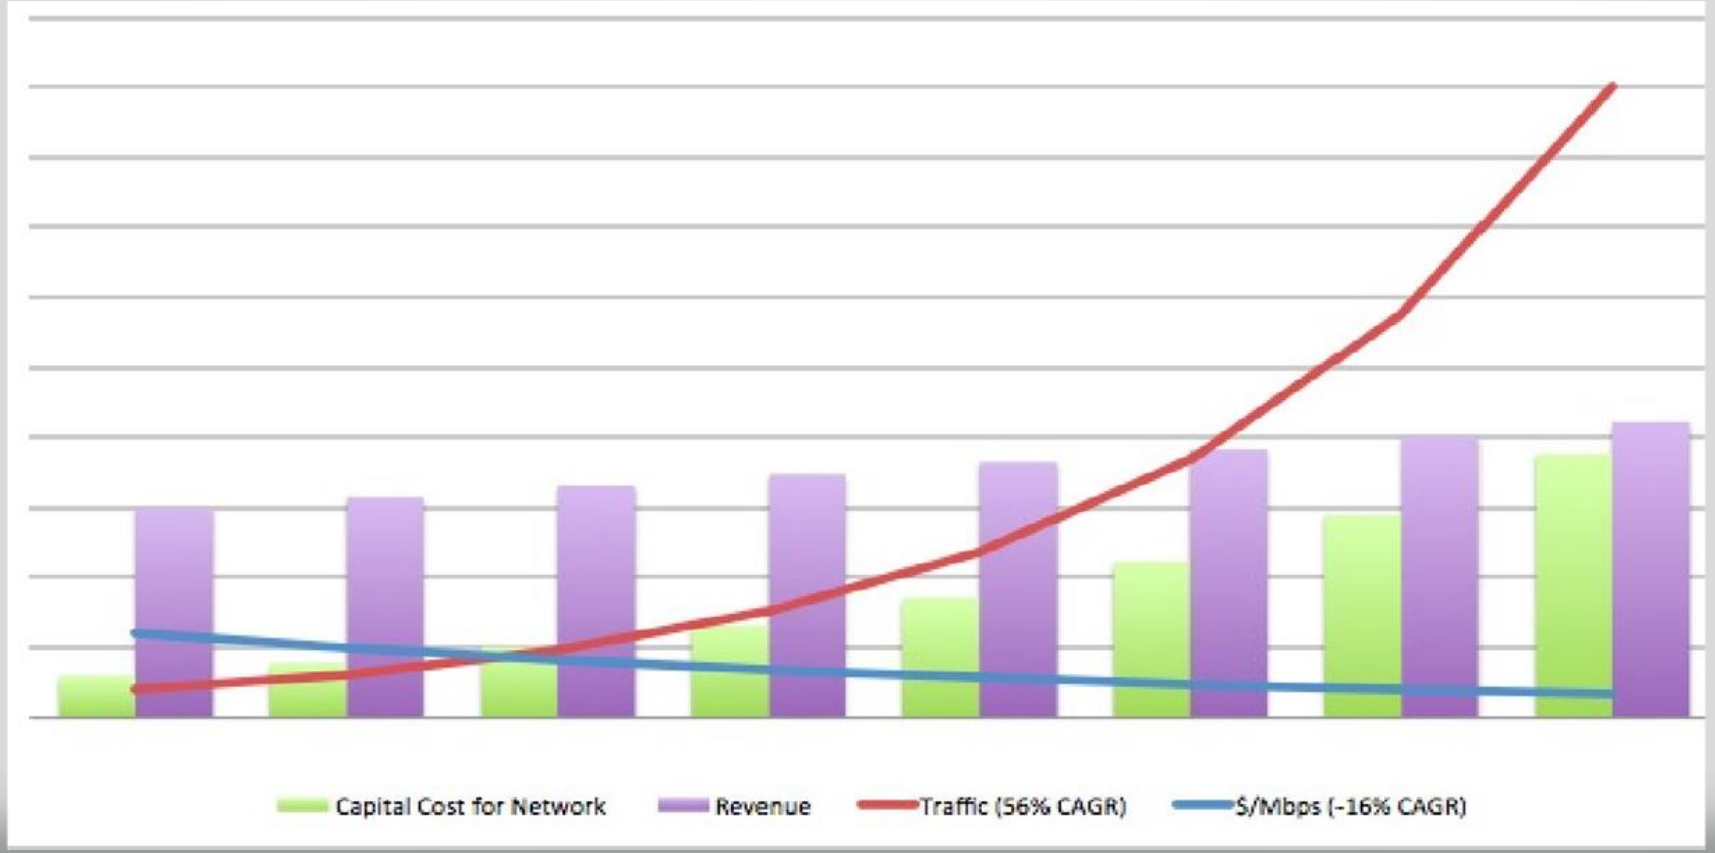

SP Networks Transformation

OBJECTIVE

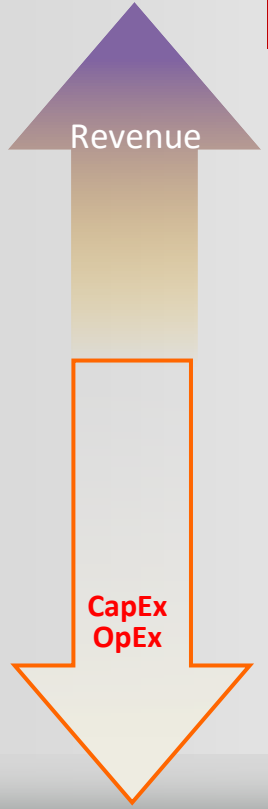

Requirement

- Urgent Network C/O-EX Optimization
- Operational Efficiency
- Improved Revenue Stream

APPROACH

- Sustainable OPEX compression
  - Reliable and Cost Optimized Hardware
  - NFV
- Operational re-engineering
  - Operational Automation Programmable Network
- Network Re-architecture
  - Re-Architect for the Least Common Denominator

Bringing Merchant Silicon to the WAN Edge

WAN Edge (CE) routers represent highest CapEx investment in infrastructure today

Challenges:

- Many niche features
- Full internet routing table in hardware

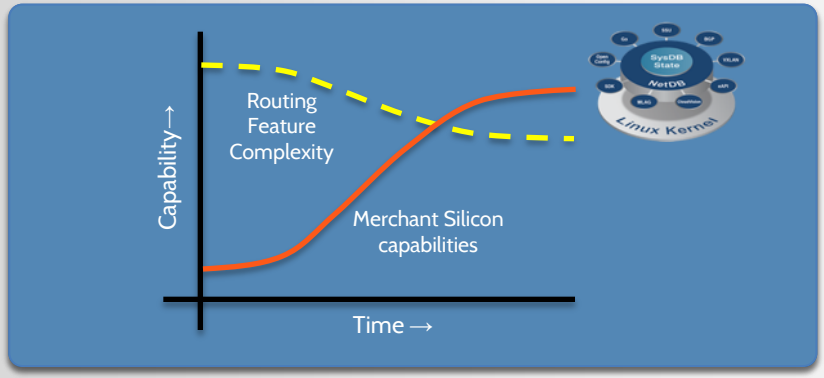

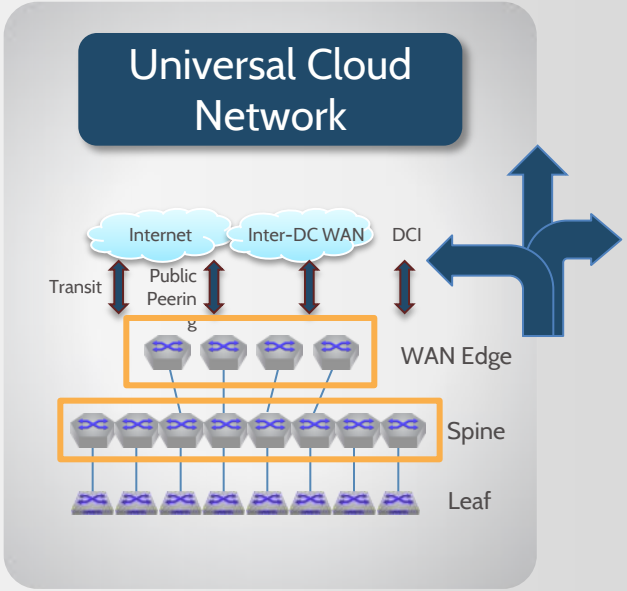

Less Than Half the Power Consumption

Merchant Silicon | ASIC Vendors

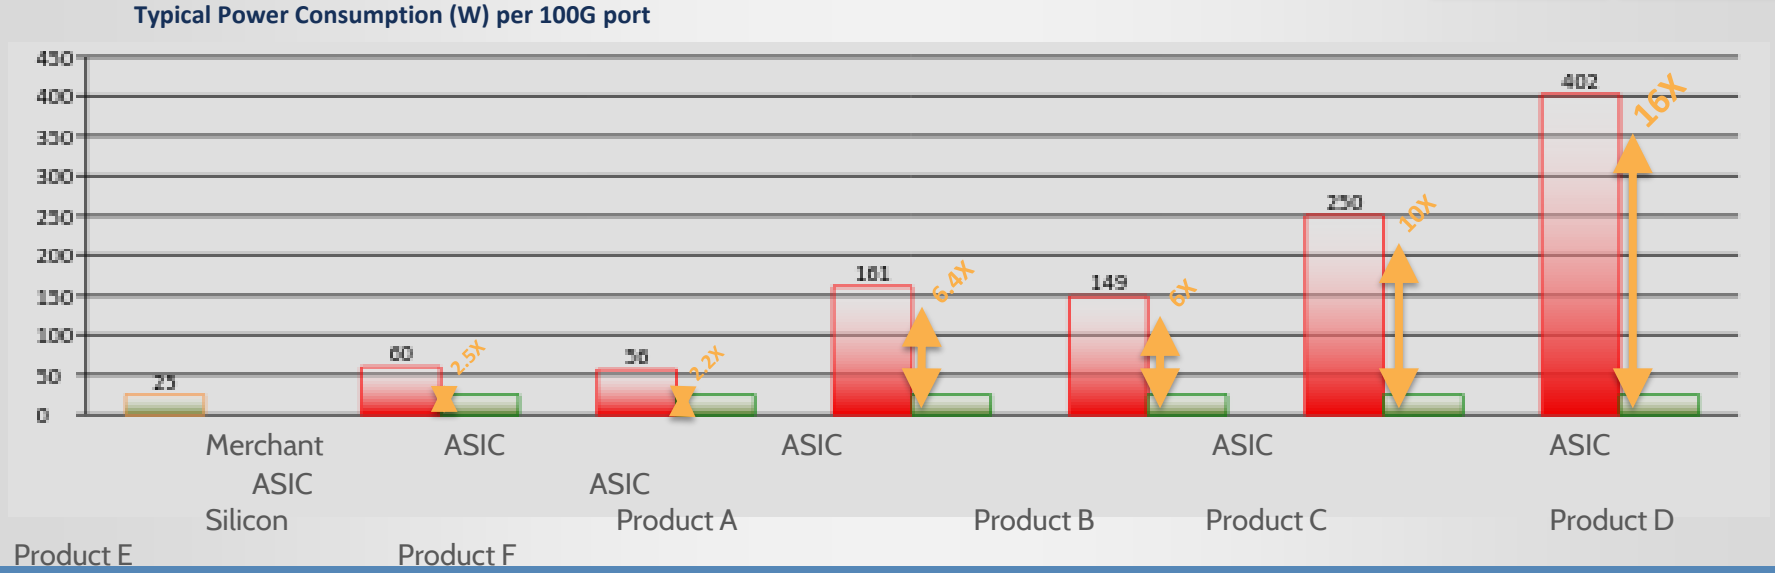

- **Half the power compared to closest competitor**
- **>6X more power efficient than legacy**

Merchant Silicon vs. Legacy Router Price

Merchant Silicon | ASIC Vendors

- **Legacy Routing Platforms heavy on features, power, and price**
- **Expect new Routing switch platforms to disrupt installed base**

Merchant Silicon for SP Use Cases

1: DC CORE NEEDS A SPINE

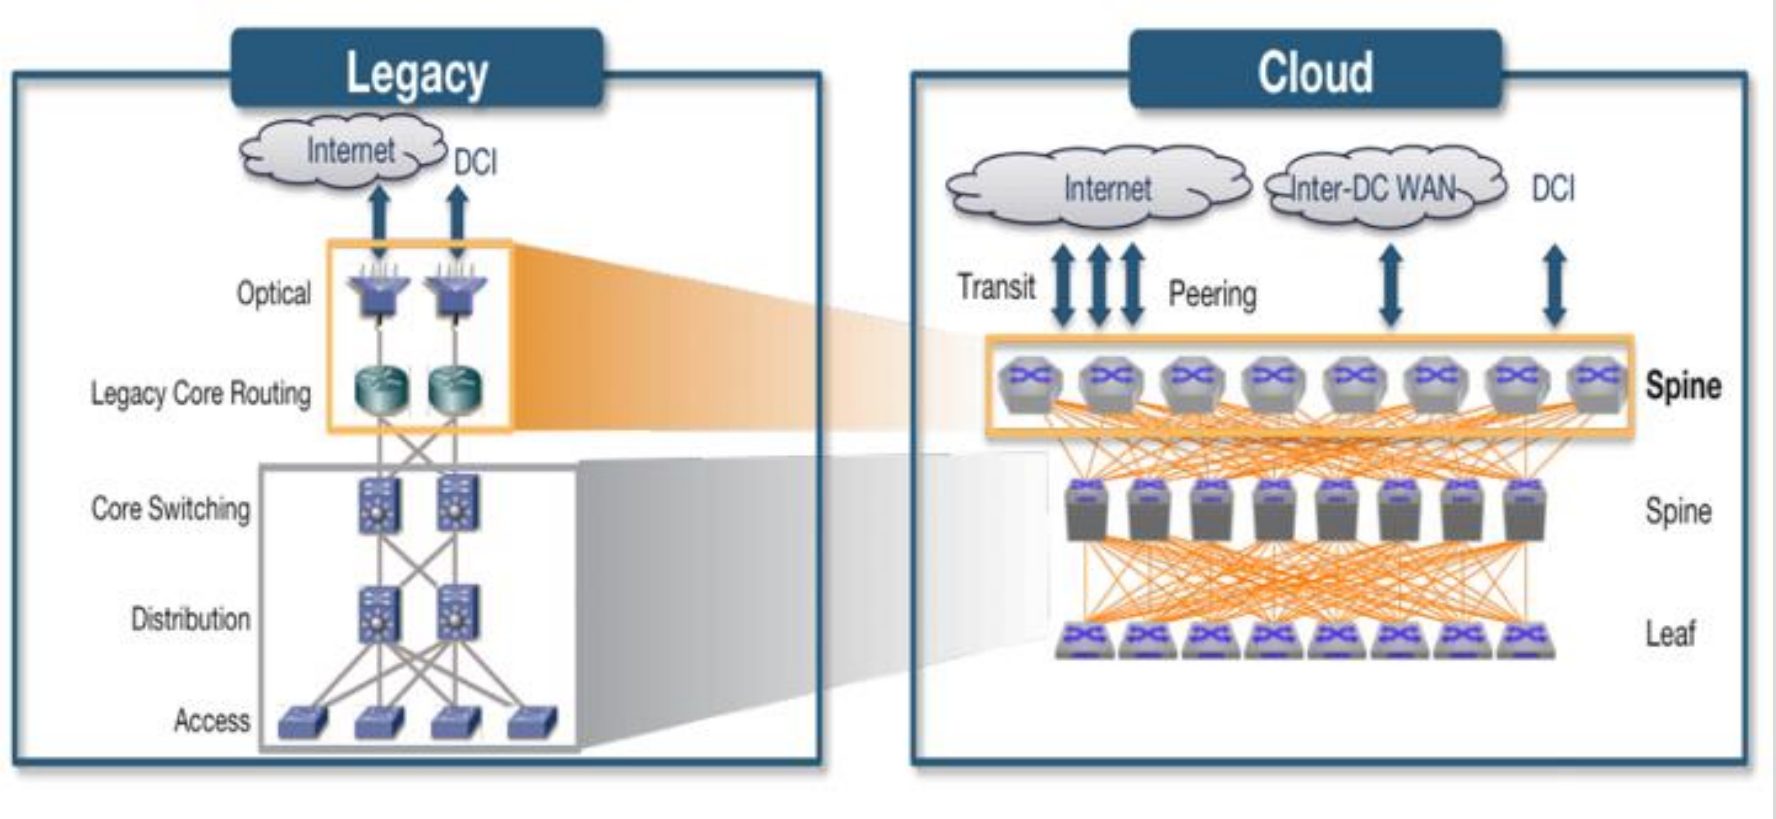

2: Internet Peering – Evolution to Content Peering

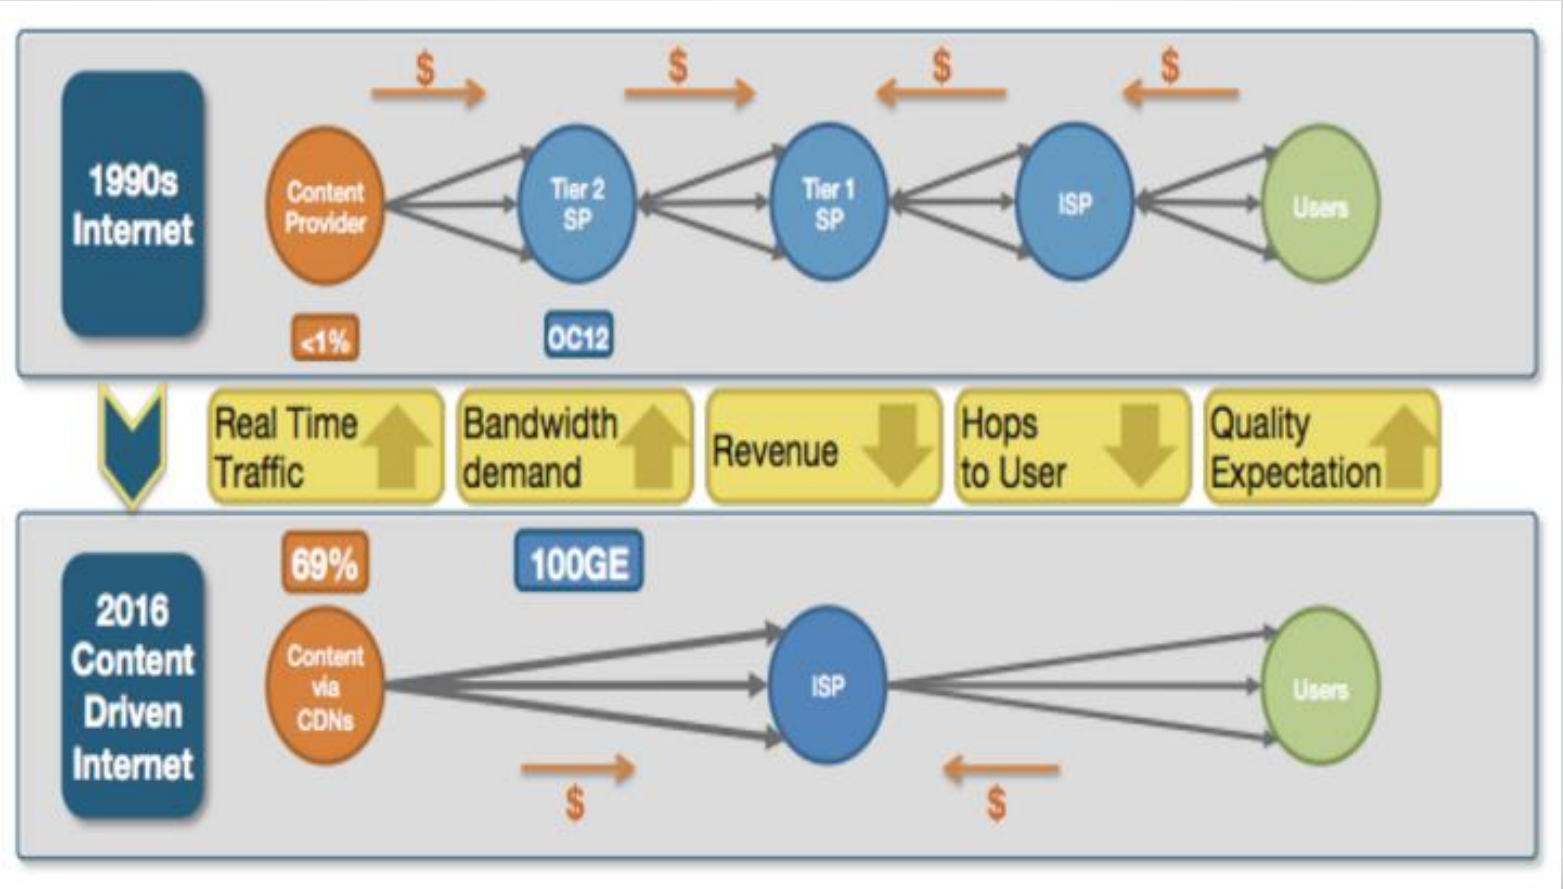

3: Cloud and WAN Segment Routing

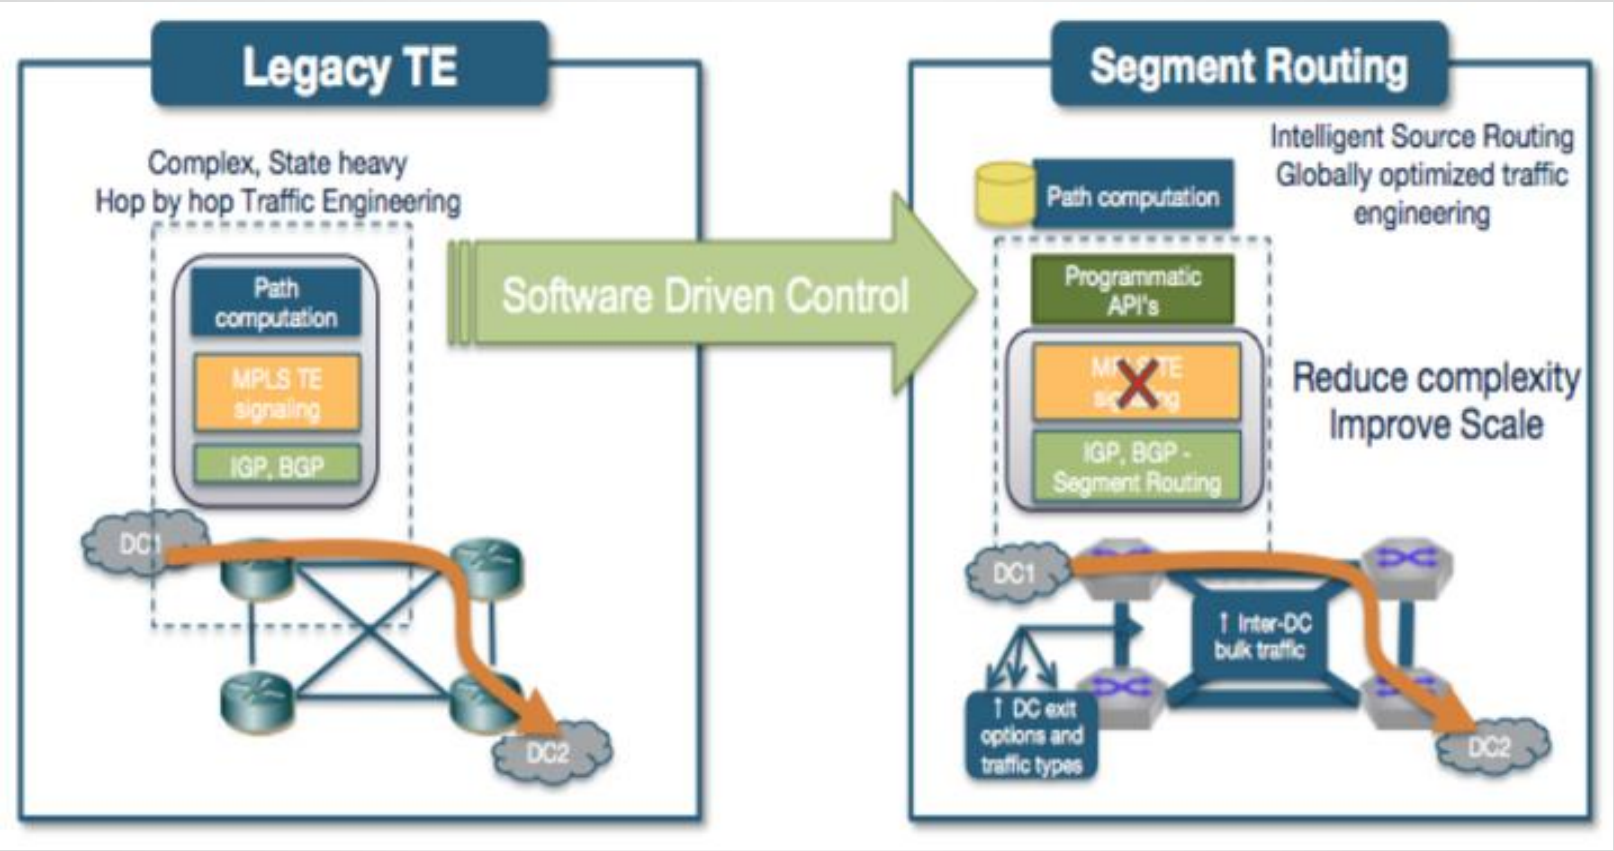

4: Telco Transformation - Service Provider NFV

SP Service Edge Evolution

Time to Market for new Services Increase	Software Strategy - Orchestration, Service instantiation
Over The Top Traffic Increase	Universal Cloud Network Leaf-Spine Architecture
Reliance on expensive HW Service Edge Routers	Virtualize and Scale Out

Next Gen Telco NFV Cloud

SDN Controller

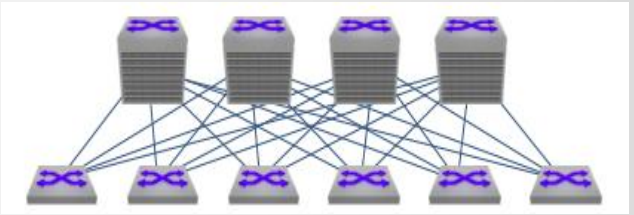

Deep Buffer Leaf-Spine Architecture

Summary

- **Following Moore's Law**
- **Higher Port Density**
- **Lower Price per Port**
- **Lower Power Consumption**

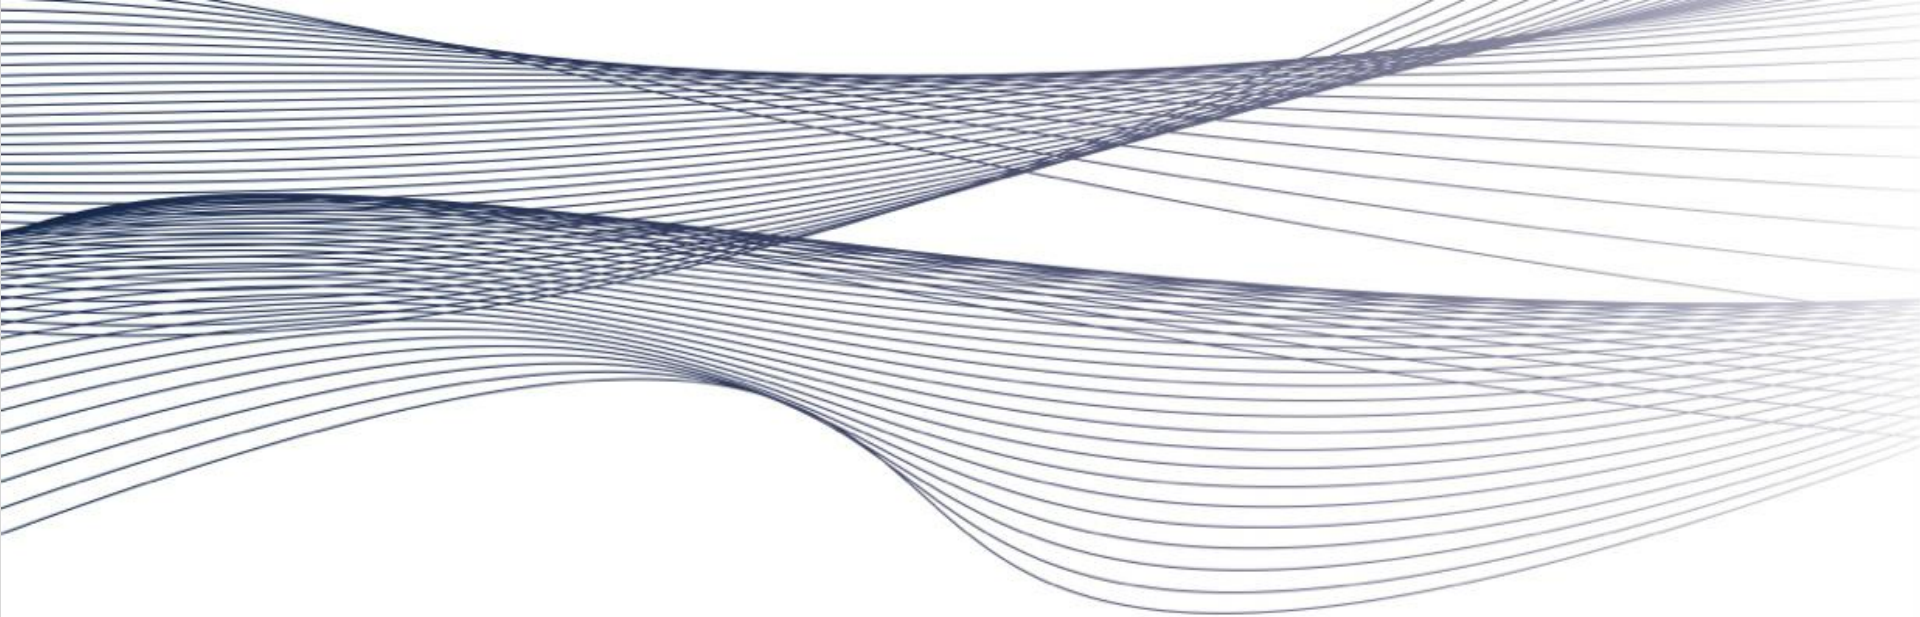

Thank You