



Sustainable 400G Data Center Networking

Cisco Knowledge Network

Max Alvarado Brenes, Worldwide Technical Solutions Architect

Usha Andra, Leader, Customer Solutions Marketing

Errol Roberts, Distinguished Architect

Yoav Schreiber, Optics Customer Solutions Marketing

June 29, 2023

Agenda

- Why sustainability?
- Technology enablers
- Nexus platform evolution
- Rack space design
- Summary

Why Sustainability?

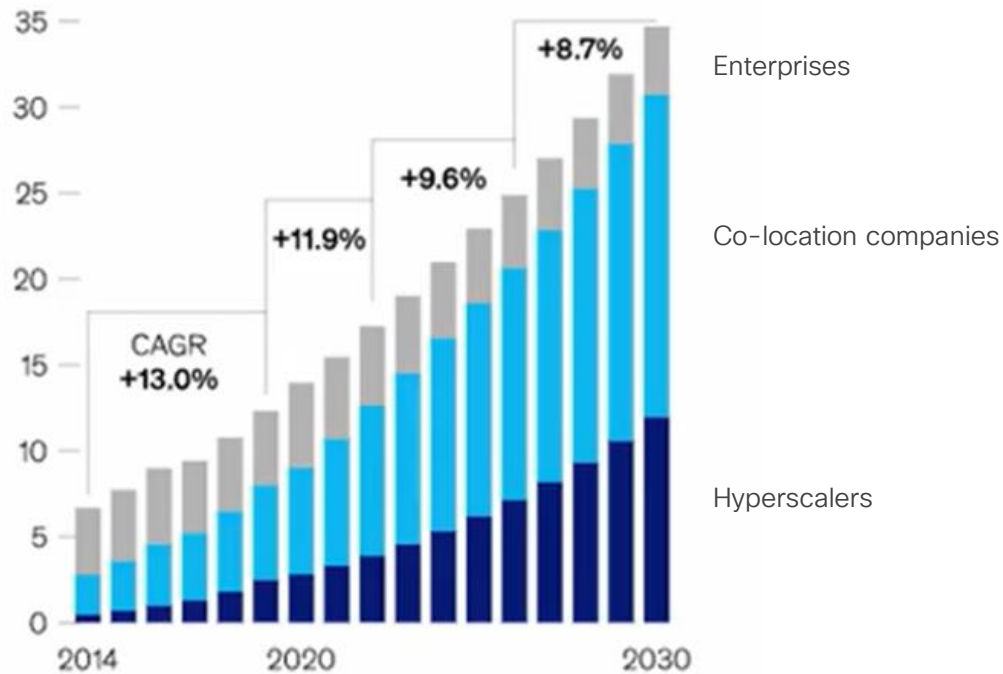


Demand for data center power and speed


US data center power demand forecast to grow 10% a year through 2030

Global spend on data center construction to reach \$49B by 2030

Data Center Power Consumption by Providers/Enterprises (gigawatts)



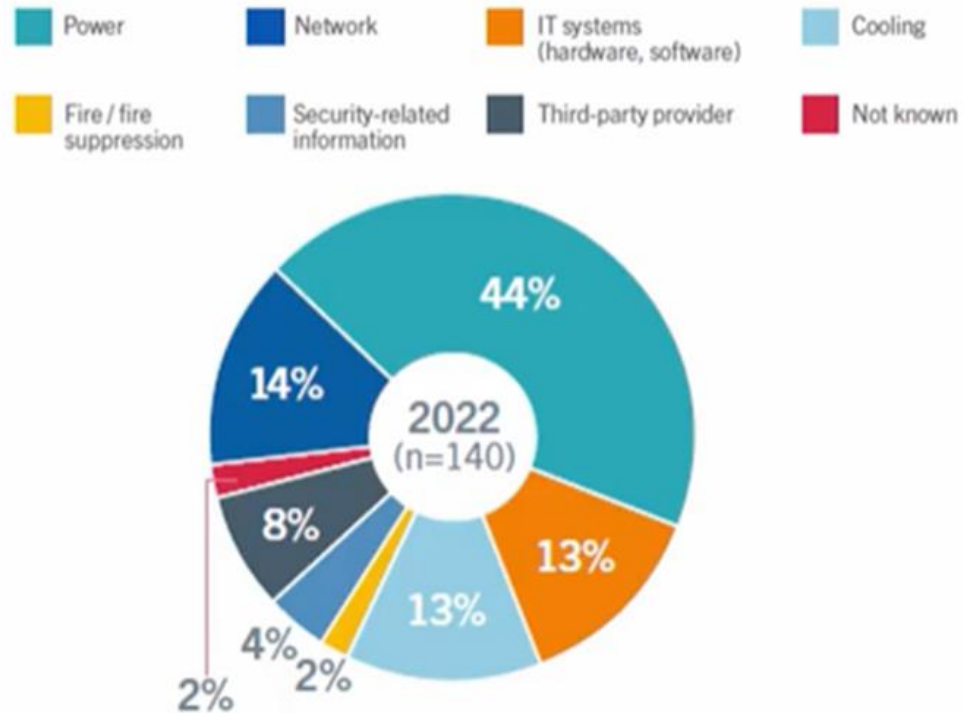
Data Center Construction Spending (Billions)



Source: Investing in the Rising Data Center Economy. Mckinsey & Company


Data center uptime and sustainability priorities

Power accounts for 44% of data center outages




Source: [2023 Outage Analysis Report](#). Uptime Institute

Ranked benefits of data center sustainability programs



Source: [Unlock the business benefits of sustainable IT infrastructure](#), Gartner May 2023

Cisco public sustainability commitments



Technology Enablers of Improved Sustainability in the Data Center



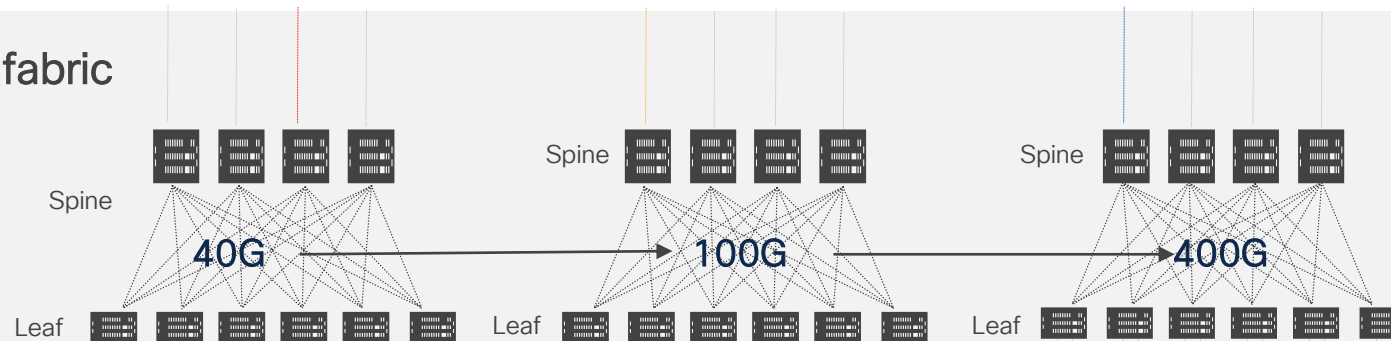
Speed evolution in the data center

Inter-Data Center



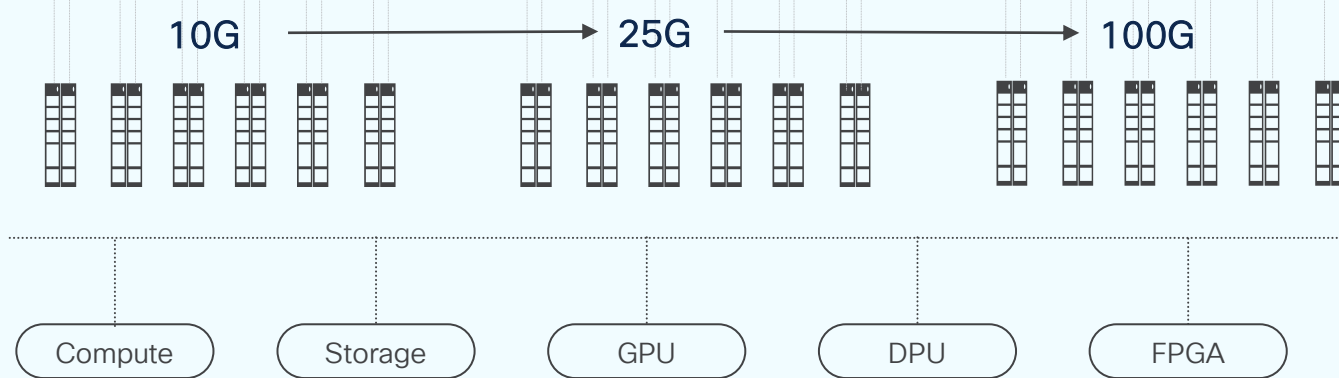
- Transitioning to pluggable DCI (DWDM coherent)
- Open Line System

Switch fabric



- Switch silicon bandwidth growing due to higher Radix and faster Serdes speeds
- Switch ASIC throughput growing: 6.4 Tbps to 12.8 Tbps to 25.6 Tbps to 51.2 Tbps (future)
- Optics increasing from 40Gbps to 100G Gbps to 400Gbps to 800Gbps


Access



- Server network connectivity evolves with server processor upgrade cycles as data center traffic grows
- Server port speed is transitioning from 1/10 Gbps to 25 Gbps to 100 Gbps
- Storage, GPU, DPU, FPGA driving connectivity bandwidth, PCIe speed increase

Increasing ASIC capacity drives higher efficiencies


Improving power/watt



QSFP-DD innovations for 400G efficiencies

Same faceplate with 2nd row of pins. Backwards compatibility with QSFP+, QSFP28, and QSFP56

Support for coherent modules with over 25w of power dissipation in riding heatsink on platform




Higher speed interfaces adopted PAM4 modulation. Ubiquitous use of FEC.

Connect to multiple lower speed interfaces

System design optimization for higher power dissipation


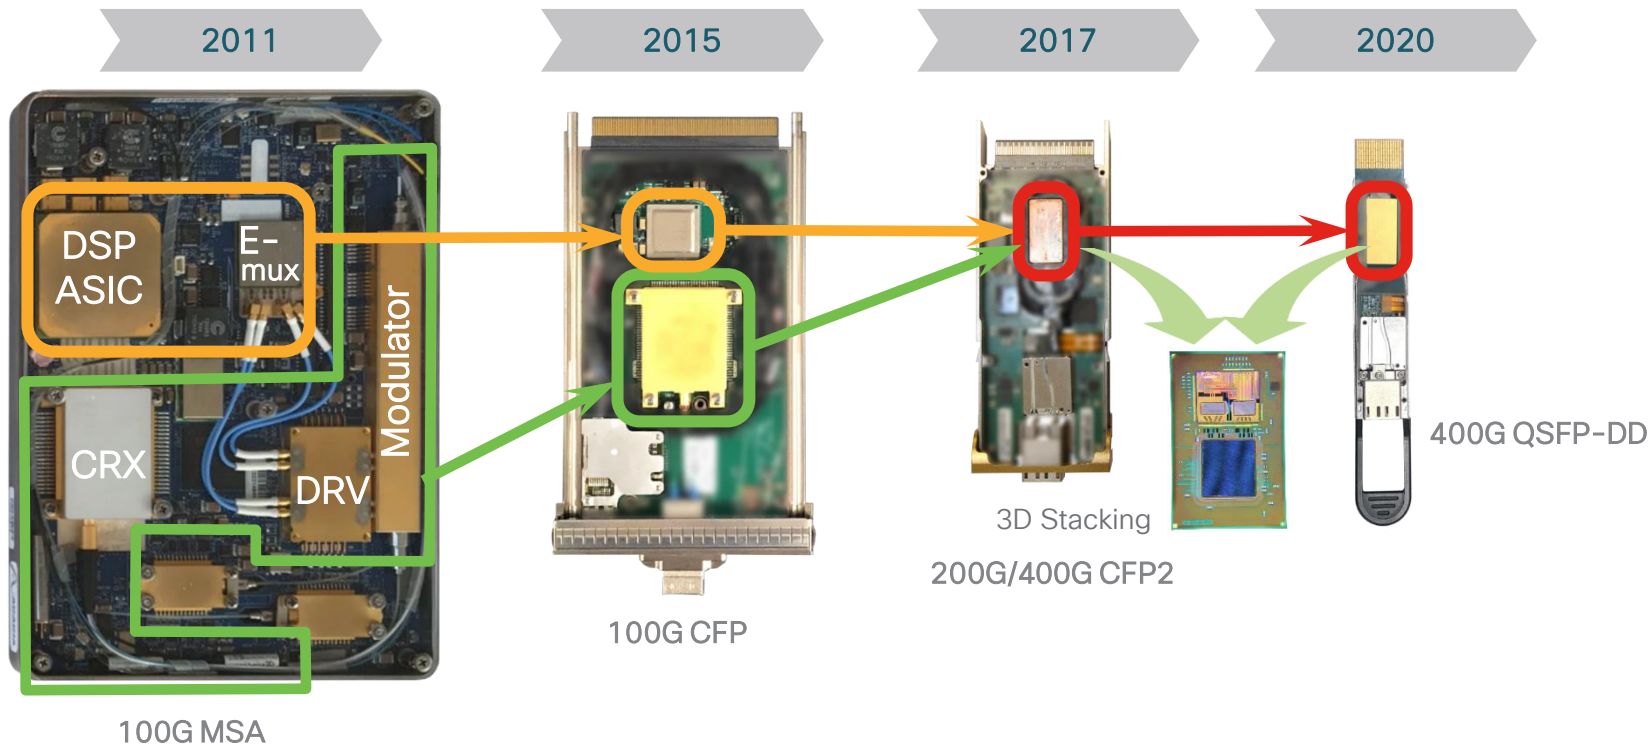
QSFP-DD riding heatsink



- Thermal design can be adapted to the system application
 - Optimize design around overall power consumption
 - Not constrained by pluggable with fixed heatsink
- Add heatsinks of different sizes
 - Extend heatsinks up into available space
 - High-power row, low-power row


Reducing power, size, and cost with increased integration

400G coherent optics and silicon photonics



Increase sustainability with high density 400GbE systems

25.6T user capacity using multiple switches with 12.8T ASICs (32x 400 GbE)



50 Gb/s ASIC IO (SerDes)
32 ports of 400GbE
(128 ports of 100 GbE)

~3000 Watts
26,280 kWh/year

25.6T user capacity using single switch with 25.6T ASIC (32x 800G)




100 Gb/s ASIC IO (SerDes)
32 ports of 800G
(64 ports of 400 GbE
256ports of 100 GbE)

~400 Watts
3,504 kWh/year


Up to **87%**
Energy Savings

83% less space/fans

Increase sustainability with 400G data center fabric



High Scale Leaf/Spine based Designs



Common network architecture between 100G, 400G and 800G

Same physical port densities, same media reaches

Continued investment in fiber plant – SM & MM fiber

Flexibility to adopt 400G breakout for high radix 100 GbE design

Connectivity to 100 GbE equipment

Design flexibility (Switch Platform)

High bandwidth, high port density platform flexibility w/ fixed, modular

Link bandwidth distribution

Port flexibility – non-coherent / coherent use cases, and mixed data rates

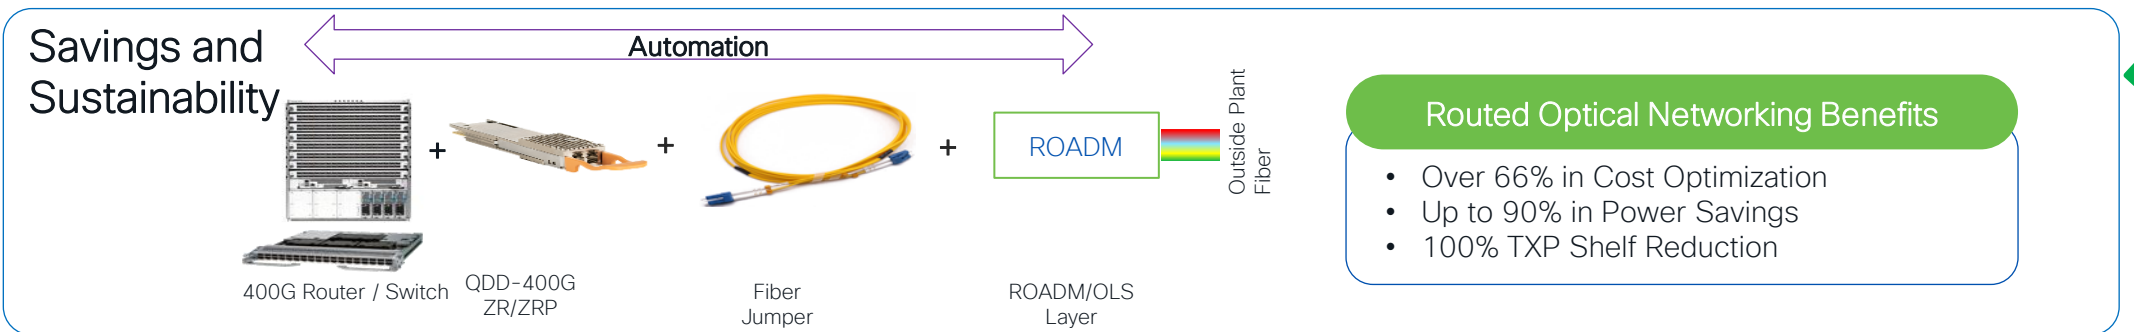


Cabling flexibility

Backward and forward compatibility (QSFP, QD-DD)

40G, 100G, 400G, 800G

Increase sustainability with 400G data center interconnect

Simplification, savings and sustainability



Improve Sustainability with Nexus Platform Evolution



It starts with telemetry

Software telemetry

Insights



Custom applications






```
0 1 0 1 0 0 0
1 0 1 1 0 1 0
0 1 1 1 0 0 0 0
1 0 0 0 0 1 0
1 0 0 0 1 0 1
0 0 1 1 0 1 1
0 1 0 0 0 0 1
0 1 1 0 1 0 1
0 0 0 0 0 0 0
1 0 1 1 0 1 0
0 1 0 1 1 0 0
0 1 1 1 1 0 1
1 0 0 0 1 0 1
0 0 0 0 1 0 1
1 1 0 1 0 1 1
0 0 0 1 0 0 0
0 1 0 1 0 0 0
1 0 1 1 1 0 1
0 0 0 1 0 0 0
```



```
0 1 0 1 0 0 0
1 0 1 1 0 1 0
0 1 1 1 0 0 0 0
1 0 0 0 0 1 0
1 0 0 0 1 0 1
0 0 1 1 0 1 1
0 1 0 0 0 0 1
0 1 1 0 1 0 1
0 0 0 0 0 0 0
1 0 1 1 0 1 0
0 1 0 1 1 0 0
0 1 1 1 1 0 1
1 0 0 0 1 0 1
0 0 0 0 1 0 1
1 1 0 1 0 1 1
0 0 0 1 0 0 0
0 1 0 1 0 0 0
1 0 1 1 1 0 1
0 0 0 1 0 0 0
```



```
0 1 0 1 0 0 0
1 0 1 1 0 1 0
0 1 1 1 0 0 0 0
1 0 0 0 0 1 0
1 0 0 0 1 0 1
0 0 1 1 0 1 1
0 1 0 0 0 0 1
0 1 1 0 1 0 1
0 0 0 0 0 0 0
1 0 1 1 0 1 0
0 1 0 1 1 0 0
0 1 1 1 1 0 1
1 0 0 0 1 0 1
0 0 0 0 1 0 1
1 1 0 1 0 1 1
0 0 0 1 0 0 0
0 1 0 1 0 0 0
1 0 1 1 1 0 1
0 0 0 1 0 0 0
```



Hardware telemetry

It continues with visibility



Visibility into network energy consumption and cost



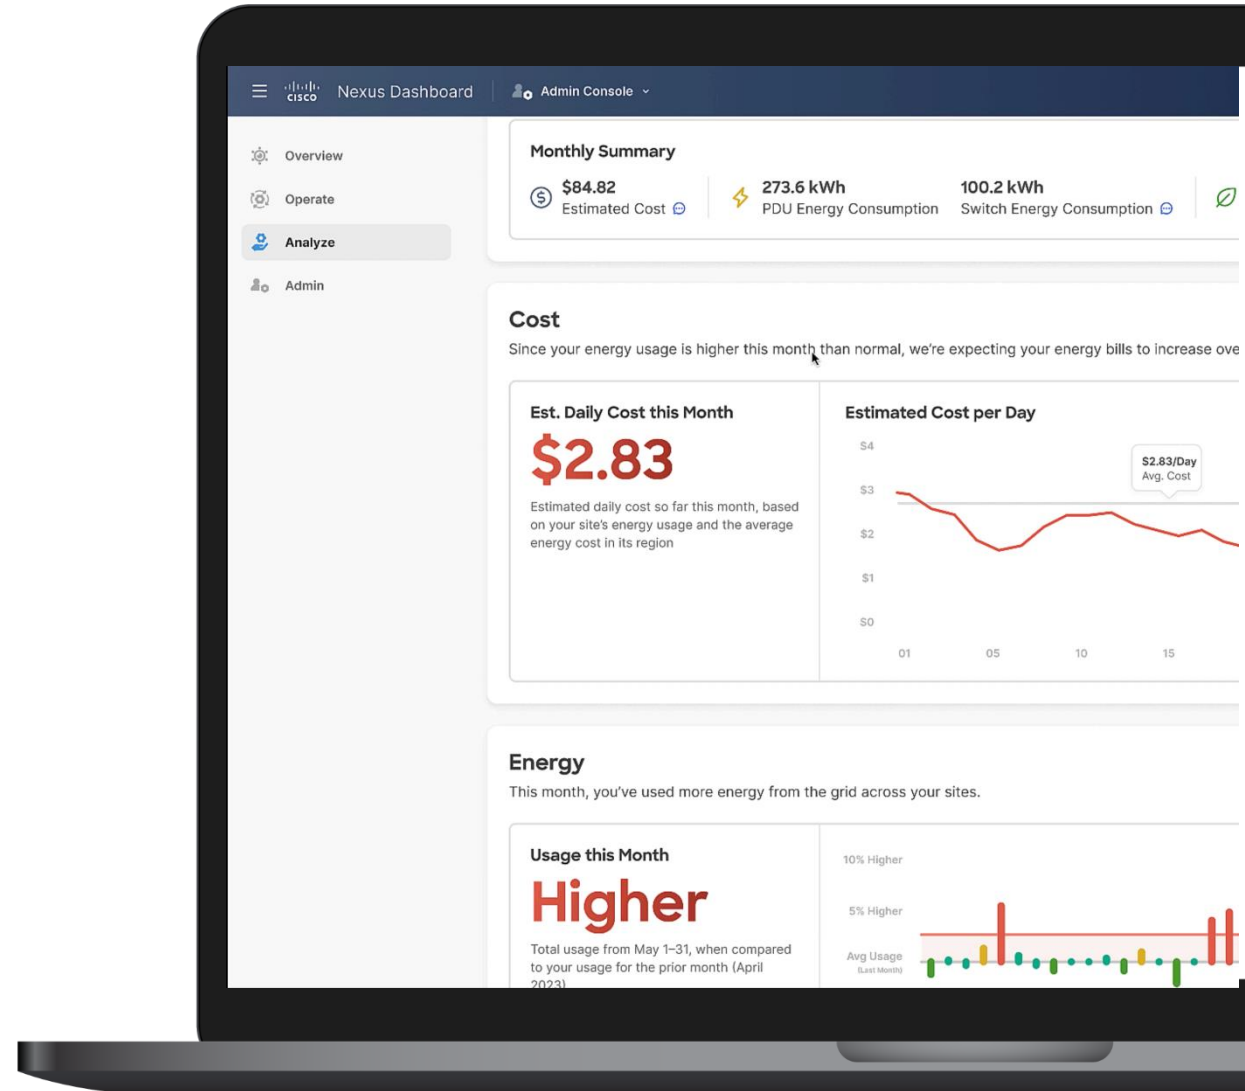
Key sustainability metrics, Top impacting devices



Understand **carbon footprint (GHG)** of managed devices



Green path determination (future)



Cisco Nexus 9800 series

Performance, flexibility and efficiency

8-slot shipping

4-slot mid CY23



12.8 Tbps per slot | 36 400G ports per slot

New PSUs

Higher efficiency, improved single-feed power

Heat sink innovation

54V → 3.3V

Removed intermediate voltage conversion

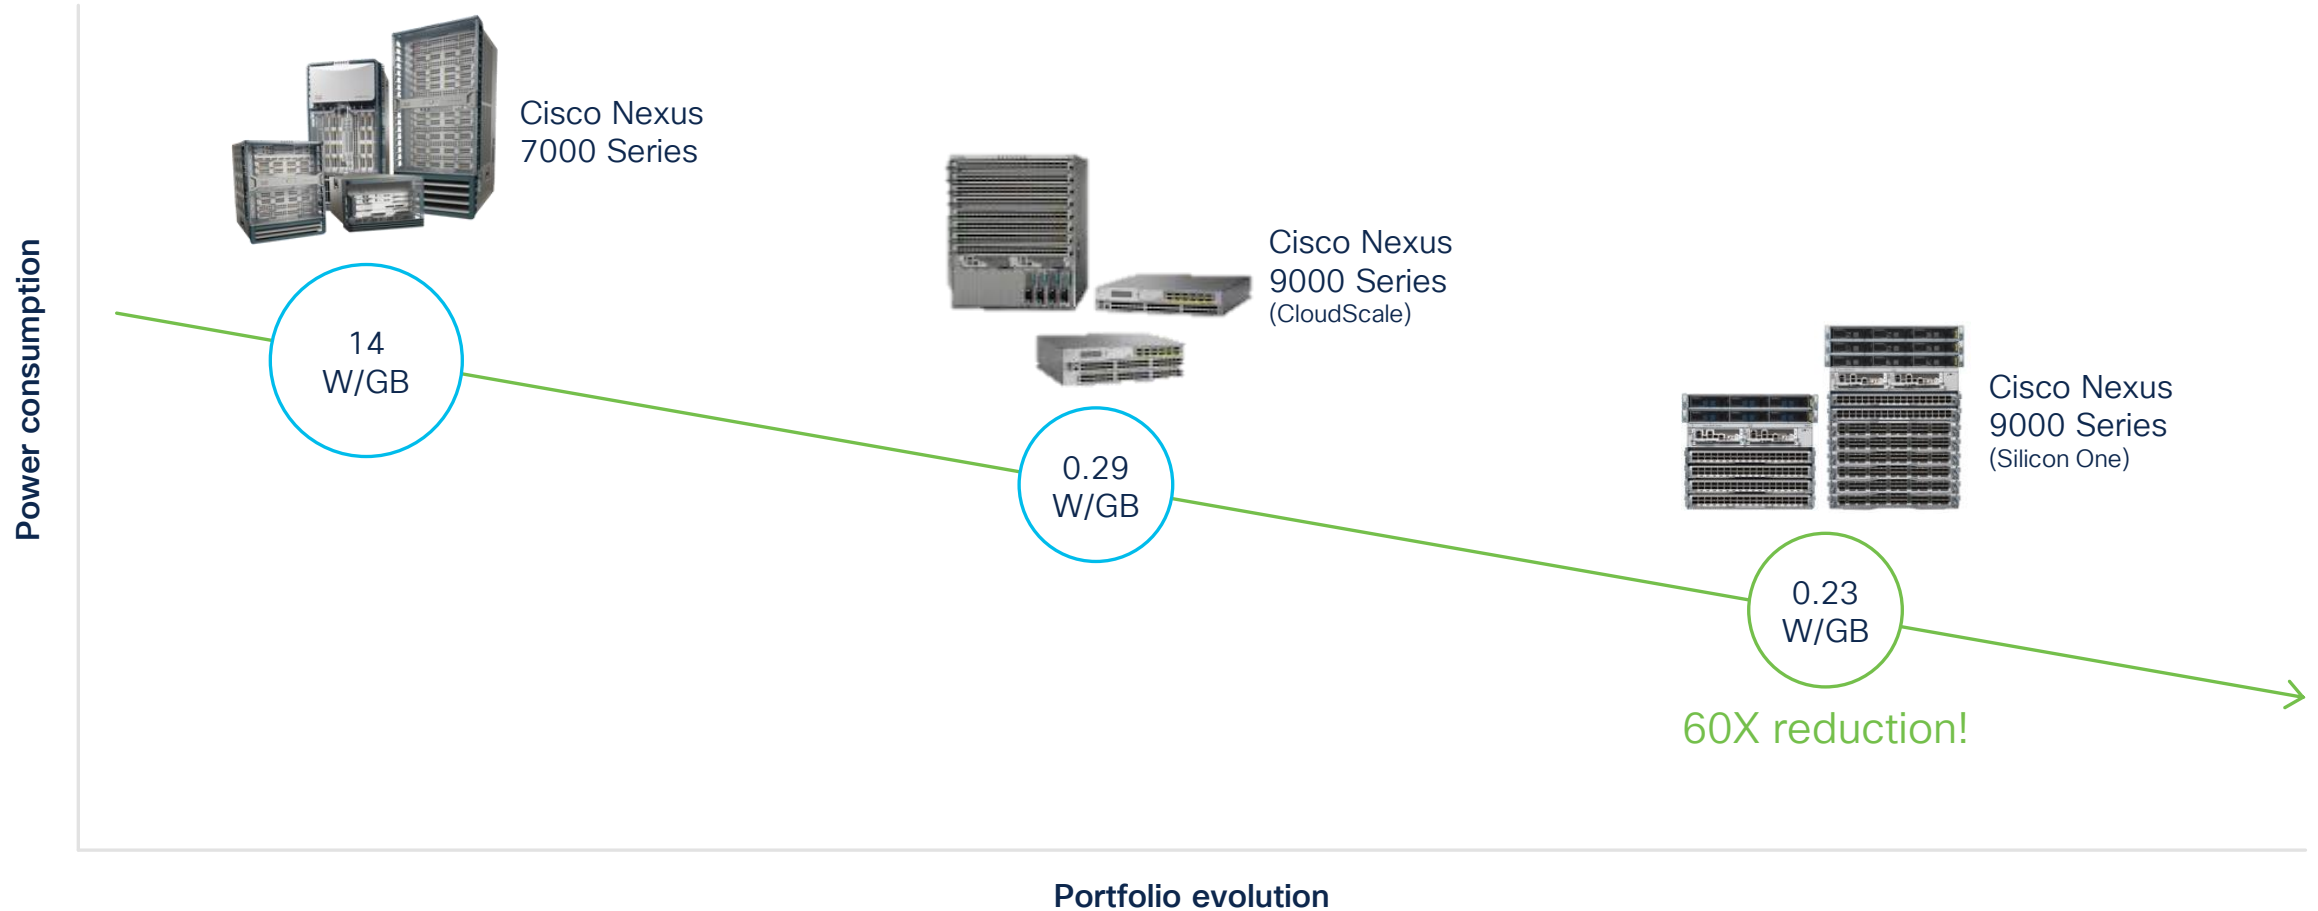
Our hardware difference

Cisco Nexus 9400

- **Single 25.6T Cisco ASIC**
- **Scalable | Compact 4 RU design, 24in depth**
Max 64 400G ports or 128 100/200G ports
- **Low power**
7.5W/100G
- **Field replaceable switch and supervisor**




Groundbreaking energy efficiency innovation



Network design difference with consolidation

Networking for AI/ML clusters, HPC, IP Storage



3456 x 25G access layer ports

72 x 93180YC-FX3

56,907W

16.5 watts/25G access port


124 RU



25G → 100G

61% less
total power

71% less
total space



3584 x 100G access layer ports

32 x 9332D-GX2B

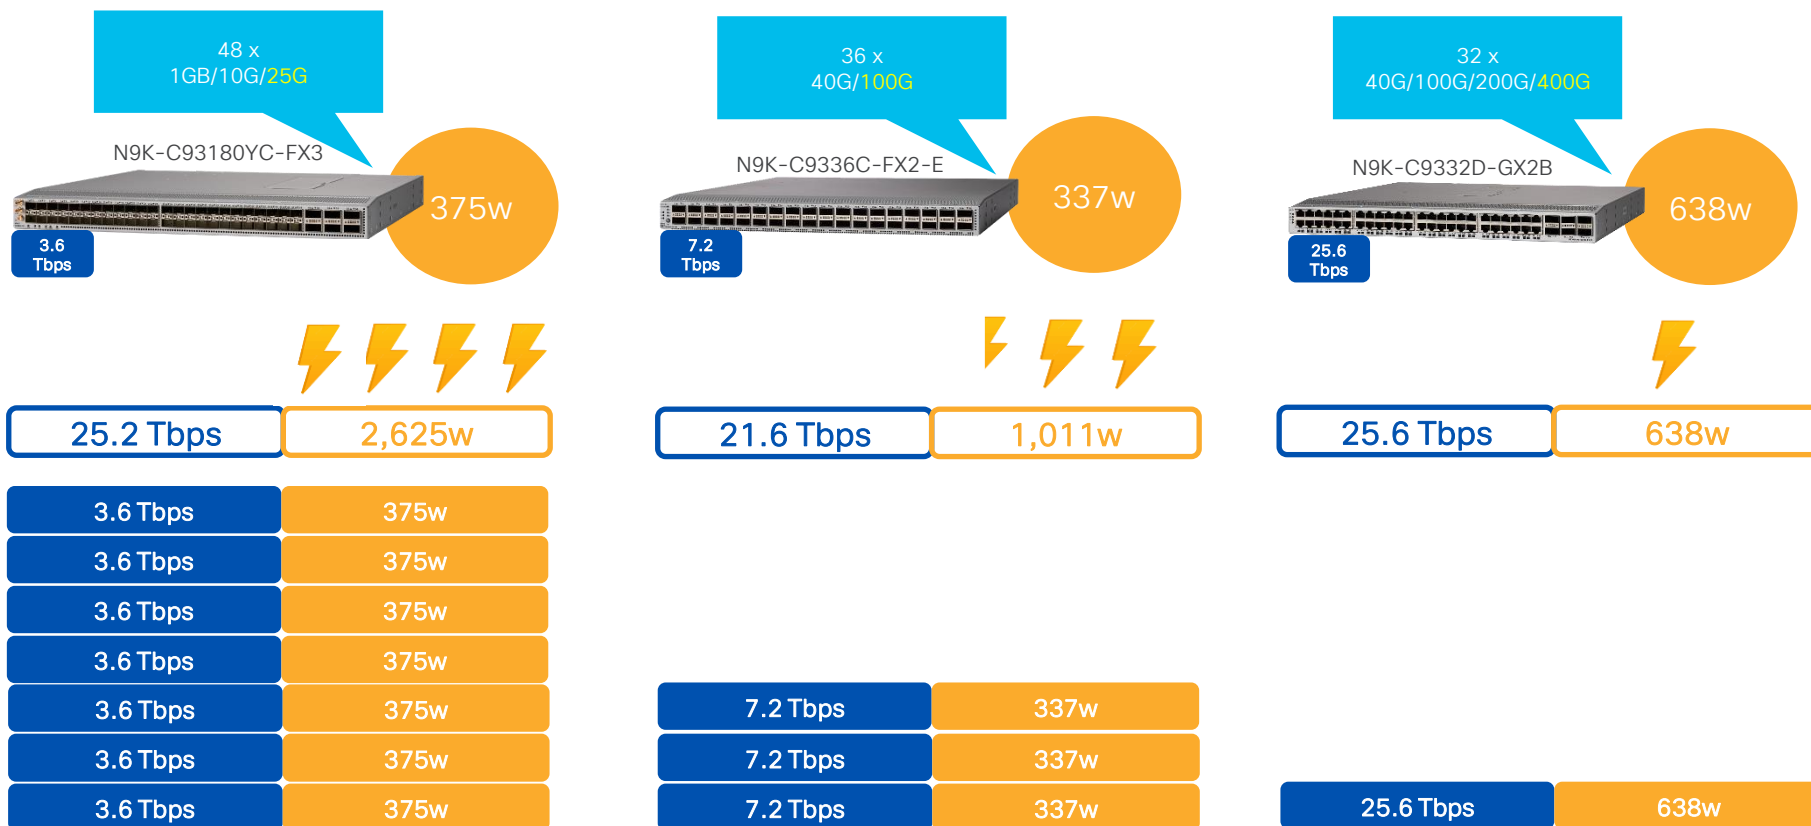
22,968W

6.41 watts/100G access port

36 RU

Nexus platform evolution

Save power while increasing bandwidth



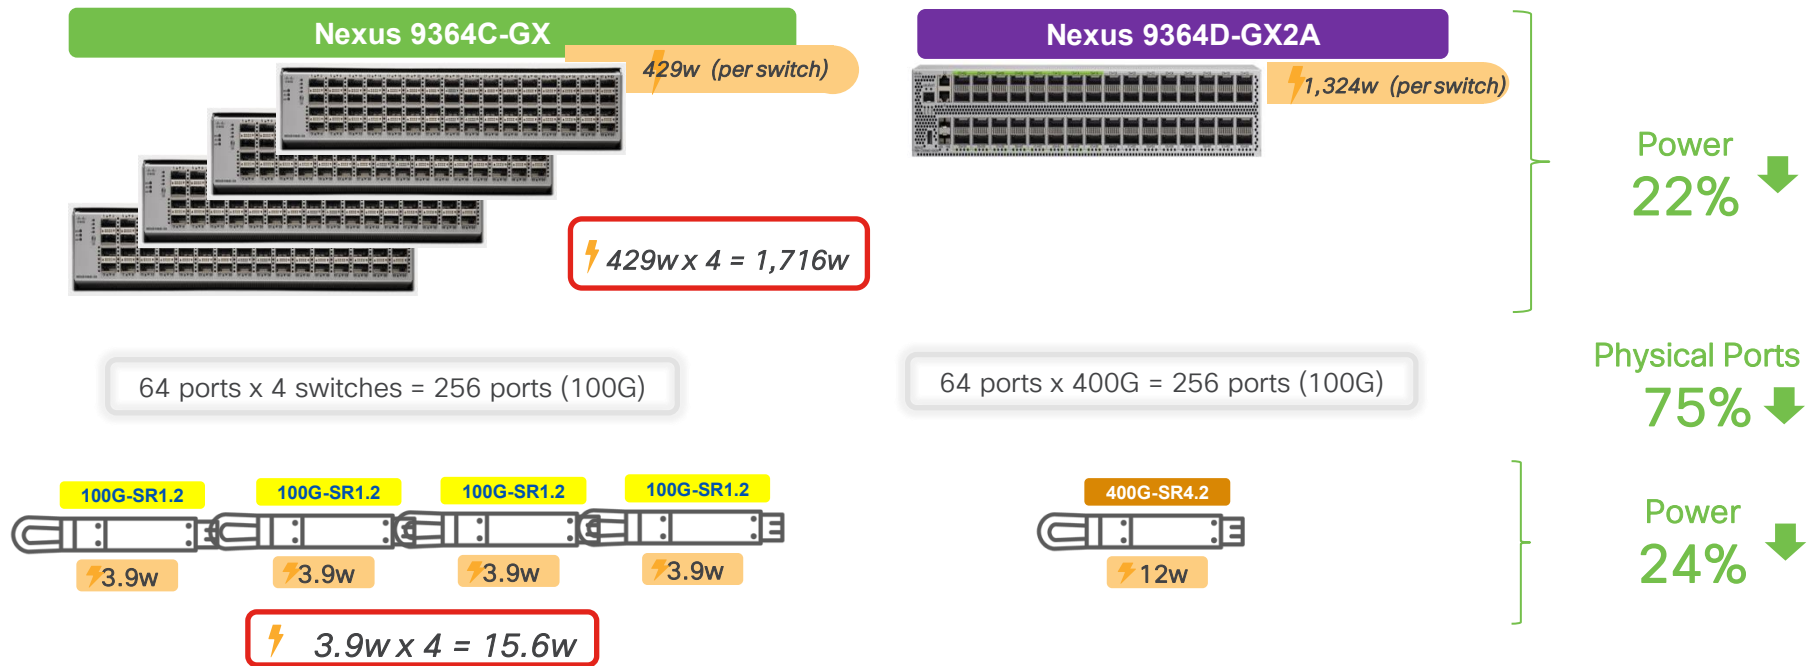
Bandwidth

Typical Power



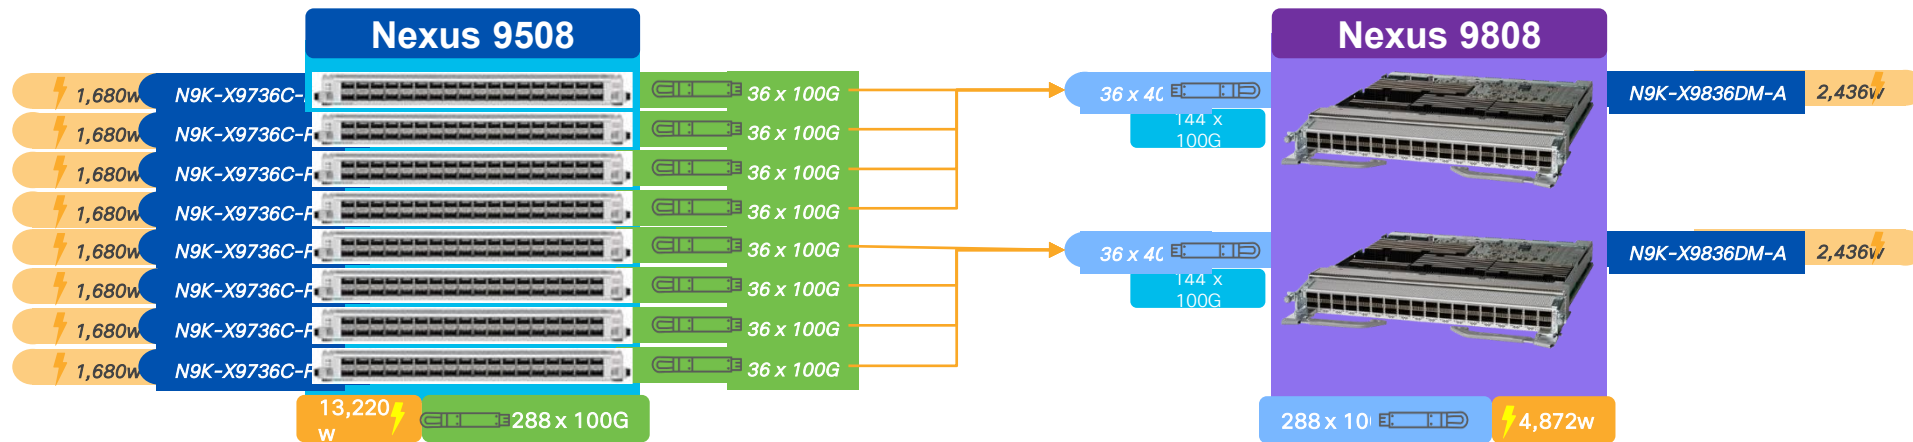
Nexus platform evolution

Increase port density



Nexus platform evolution

Consolidate with modular line cards



36 ports (100G) x 8 linecards = 288 ports (100G)

13,220w

36 ports (400G) x 2 linecards = 288 ports (100G)

4,872w

Optics and Ports
75% ↓


Linecards Power
63% ↓
(Based on Max. Power)

Nexus platform evolution

Move to higher speed optics




$$3.9w \times 4 = 15.6w$$



Power
24% ↓




36
Optics



$$3.9w \times 36 = 140.4w$$



9
Optics




$$12w \times 9 = 108w$$

Power
24% ↓

Rack Space and Cabling Design for Data Center Sustainability

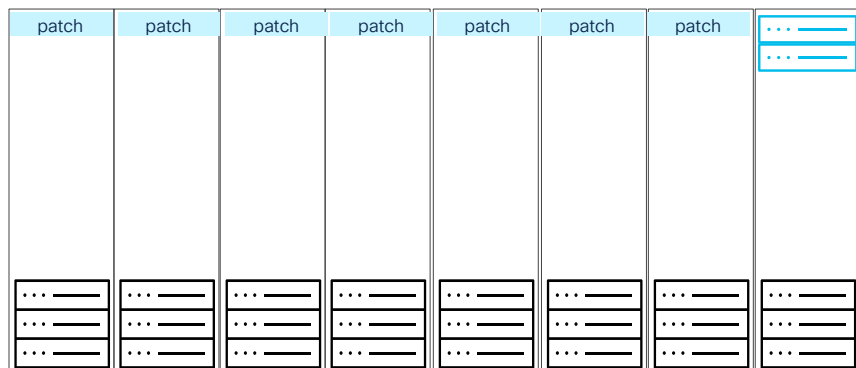
Improve power efficiency with 400G

Consolidate switches with high density breakout across data center racks



Current Architecture

- One ToR Switch (1RU) per Cabinet
- Provides connectivity to 16-32 servers
- One port per server



New Architecture

- One or two ToR Switches (2RU) per (8) Cabinets
- Provides connectivity to 128-256 servers
- One port can service 4 or 8 servers with 4 to1 (SR4.2) or 8 to1 (SR8) breakouts

Simplify cable management with patch panel cassettes



6 Cassettes
(fully removable front & back)

SOC™ Sleeve
Cable Organizer
Mechanical Protection & Strain Relief

Cable Hanger Brackets
Slack Management

Front view
of 1RU
chassis with
6 cassettes



MPO-12


Duplex LC

Top view
of cassette

1 cassette supports (3)
MPO-12 in rear that break
out into 12 duplex LC ports
in front

Data center rack consolidation with patch panel

51.2 Tb capacity, 140 x 60amp DC feeds




CRS 8+1
9 full Racks 512 x 100Gb ports

172.8 Tb capacity, 24 x 100amp DC feeds



8812 Front w/ Cable Mgmt
1728 x 100Gb ports



8812 Rear w/ Cable Mgmt

Summary



Cisco Green Pay


Cisco Green Pay Circular IT payment solution

5%

95%

Predictable fixed payments

Includes option to blend software and services into payments



Up to 20 quarters



Incentive

Up to 5% on hardware



Upgrade

To latest sustainable technologies



Return

Free shipping costs, sustainability certificate




Extend

Keep current equipment month to month or for one year

Cisco product takeback and reuse program

Free removal and transport of equipment at customer end-of-use



Simple

Cisco offers **various tools** to help you create a return request and schedule the free pickup of your used Cisco equipment



Secure

Returned equipment is stored in a secure location and hard drives are cleared according to the **U.S. NIST 800-88 guidelines**



Sustainable

Cisco can help you reach your sustainability goals – we **reuse and recycle 99.9%** of what is returned to our facilities

Cisco Nexus Dashboard PDU integrations

PANDUIT[®]



Cisco Nexus
Dashboard



Sustainability insights

Real-time and historical insights into the energy consumption, energy costs, GHG emissions

Monitor the ambient temperature to help improve cooling efficiency of the data center

Support for any IT equipment (servers, switches, storage etc.)

Emission forecast and optimization

Estimate future energy consumption, energy cost, GHG emissions

Provide recommendations to help reduce energy consumption

Get started today

3

Learn more about Cisco sustainability initiatives

cisco.com/c/en/us/about/csr/environmental-sustainability.html

2

Evaluate optical transceivers compatibility and interoperability

www.tmgmatrix.com

1

Explore the Cisco portfolio for Nexus and Optics

www.cisco.com/go/nexus

www.cisco.com/go/optics



The bridge to possible