Get ready for the 800GE reality

Jonas Vermeulen
Technical Sales
EMEA Webscale Business

Nokia is taking care of your interconnectivity needs

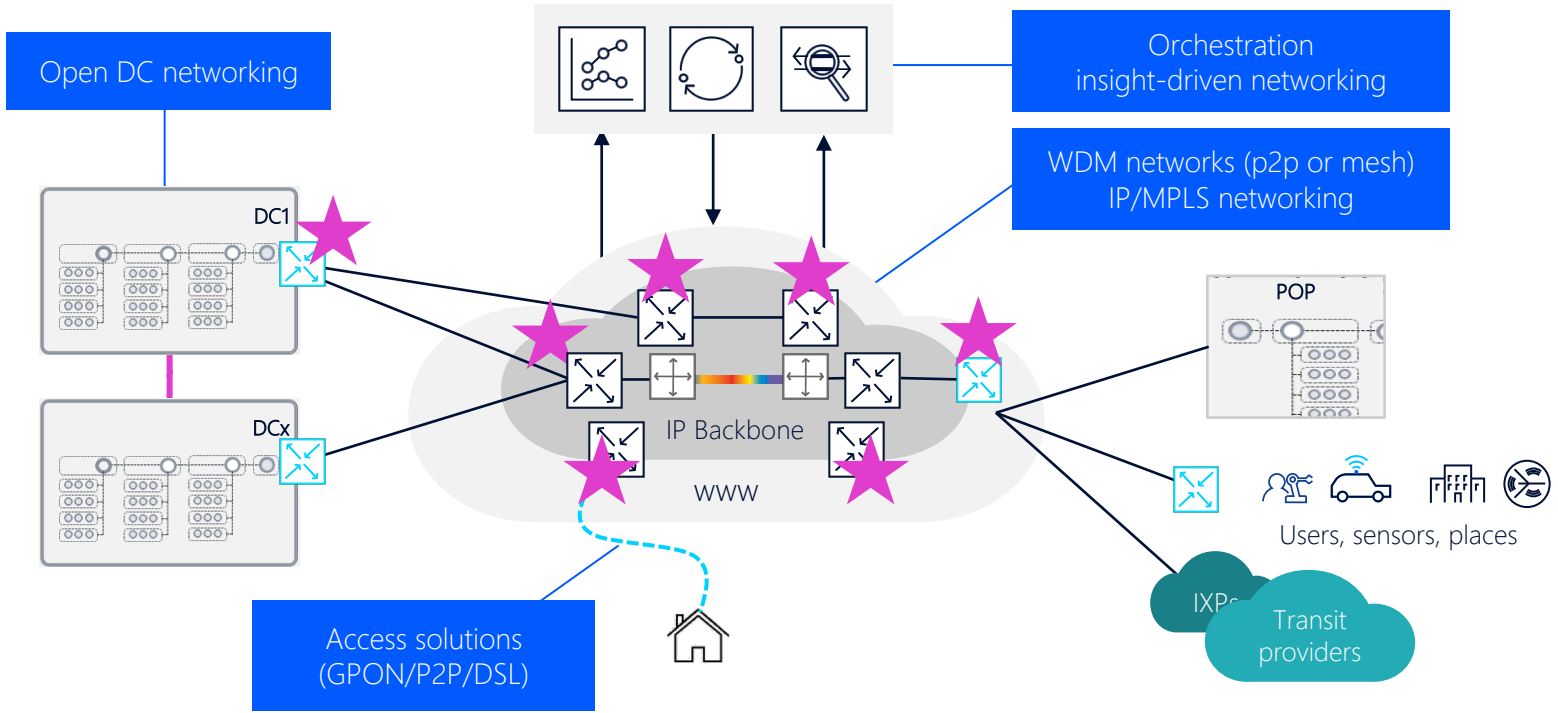

New interconnectivity needs driven by…

3x Acceleration in global bandwidth consumption 2022

Accelerated **data consumption** during and after COVID

1T Once rare 1 Tbs attacks now daily occurrence!

DDoS traffic is exponential – growing faster than video or any other form of Internet content

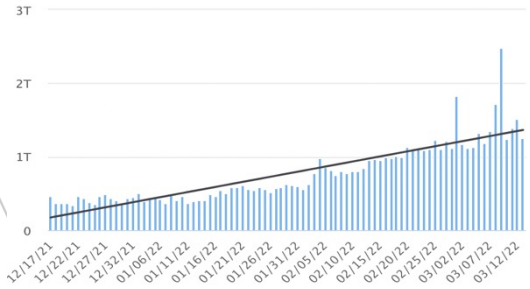

50% Emissions reduction by Nokia products & operations by 2030

Increased focus on power efficiency and sustainability

How are routers (re)designed to cope with this?

Routers are simple

…. or not so much?

A router’s simple task

- Receive a packet
- Find the next-hop
- Send out the packet

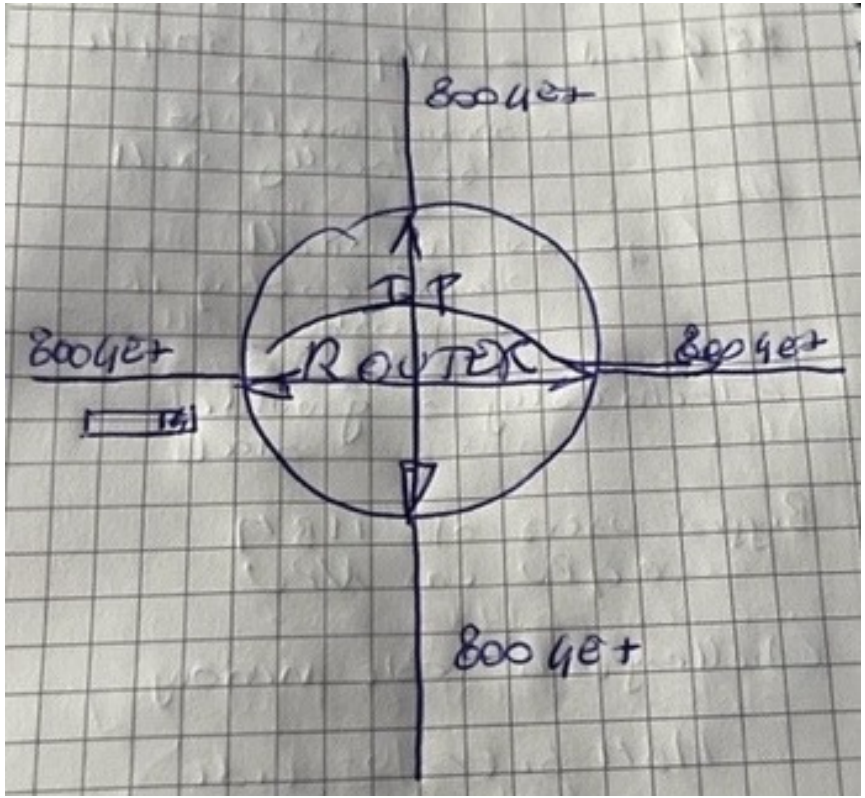

Routers are simple

…. or not so much?

But…

- Multiple billion times per second (and growing)
- Preferably without any hickups (aka packet drops)

And

- Do it for more than just goodput IPv4 (also IPv6, MPLS, VLAN, VPN, L2, PW, GRE, SDH, VXLAN, Video, MBH, Multicast, QoS, … [ever growing list])
- Deal with “not-so-good” traffic (aka DDoS)
- With practical constraints like manageability, cost, power, …

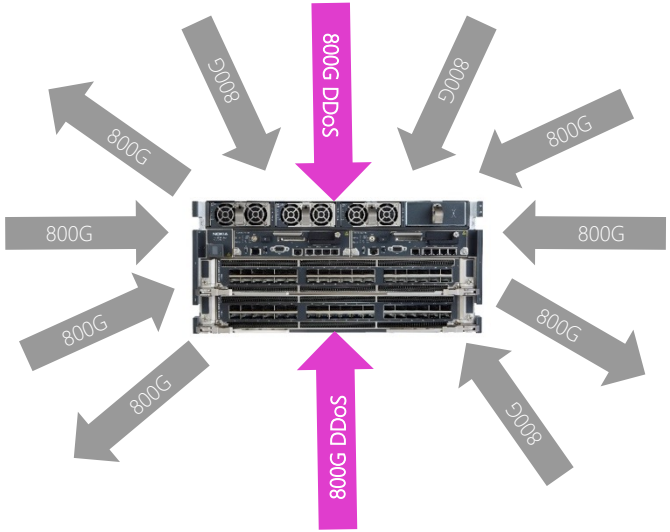

“From simple, uniform, best-effort IP forwarding to … Life-over-IP”

How do we build for scale and power efficiency?

Key technology evolutions on the router enabling 800GE

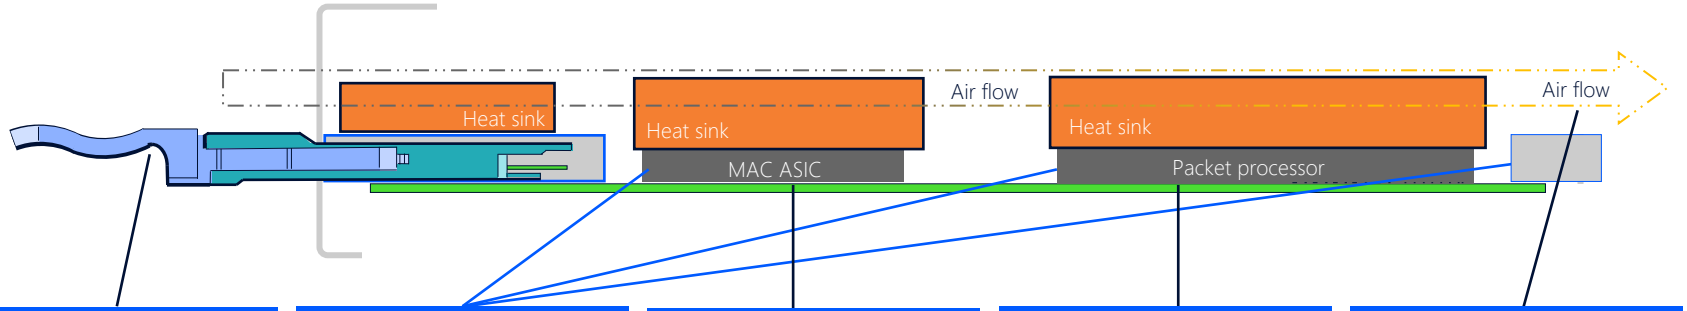

Pluggable optics and cage type
- Distance
- Density
- Compatibility

SERDES
- Coding
- Speed

MAC ASIC
- Single flow speed
- Flexibility
- Compatibility

Packet processor (NPU)
- Forwarding speed
- Scale
- Buffering

Air flow and power
- Effectiveness
- Design longevity
- Power budget

How do we build for scale and power efficiency?

Key technology evolutions on the router enabling 800GE

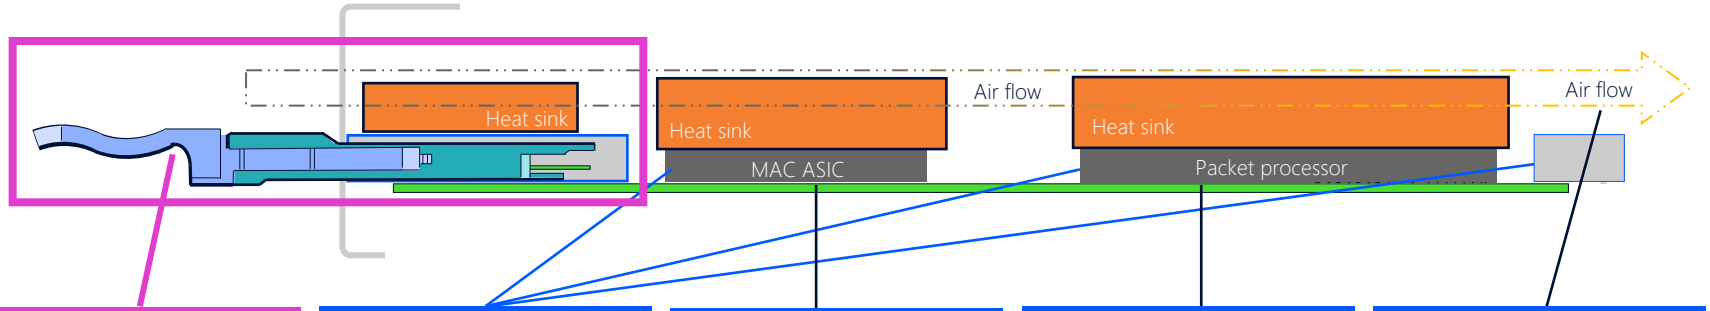

Pluggable optics and cage type
- Distance
- Density
- Compatibility

SERDES
- Coding
- Speed

MAC ASIC
- Single flow speed
- Flexibility
- Compatibility

Packet processor (NPU)
- Forwarding speed
- Scale
- Buffering

Air flow and power
- Effectiveness
- Design longevity
- Power budget

Pluggables for 400G and beyond…

Optical interface technology enablers

Formfactor

- Packaging
- Mechanics and cooling
- Routerinterface density

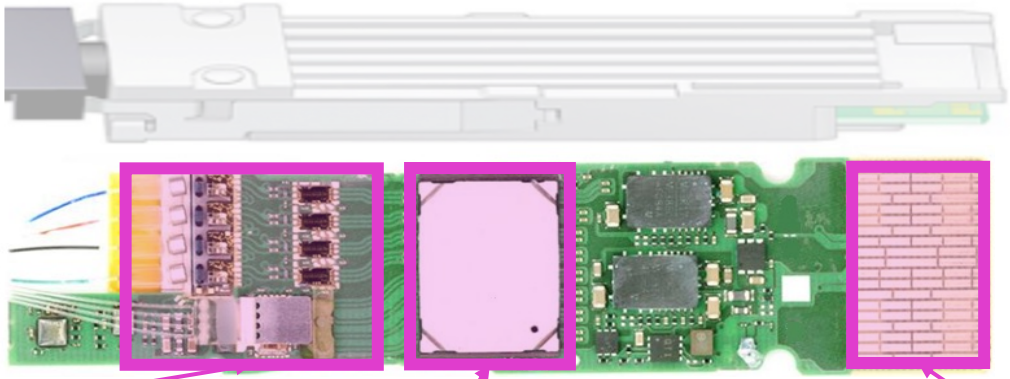

Arbitrary example photo of an optical pluggable PCB for illustration purpose

Photonics and drivers

- Optical modulation and number of wavelengths (lambda's) are the key factor affecting cost and performance.

DSP/Multi-link gearbox

- Modulation/Demodulation digital signal processing
- One of KEY factors in defining power/thermal envelopes of the module

Attachment Unit Interface (AUI)

- Data transmitted over Electrical SerDes links from module to chipset

Diversity and uniformity

Pluggable optics and cage types

Cage types becoming universal

- **Narrow (100G-)**: SFP, SFP28, SFP56, SDF-DD, SFP112, SFP112-DD
- **Wide (100G+)**: QSFP28, QSFP28-DD, QSFP56, QSFP112, QSFP-DD 400, QSFP-DD 800

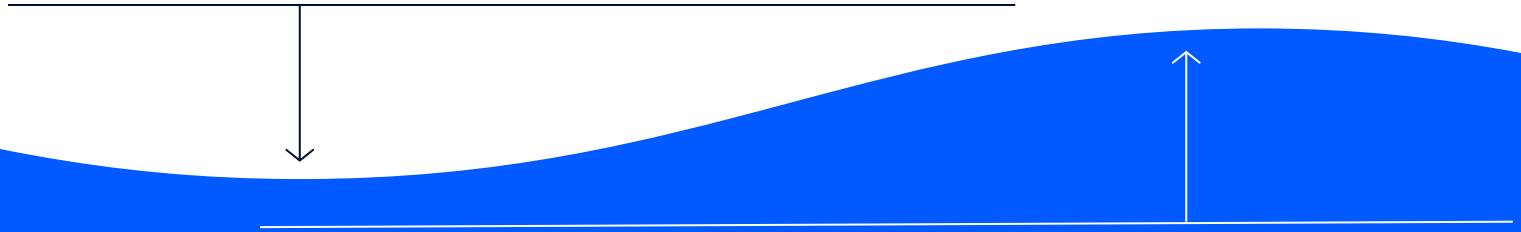

Interesting evolutions…

- **100G Single Lambda**: 800G, higher 100G/400G density, cost reduction
- **400G ZR/ZR+**: Coherent to enable "Pragmatic IPoDWDM" designs

SFP ‘Narrow’ cages

1 channel

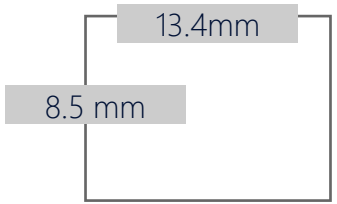

	SFP+	SFP 28	SFP-DD	SFP 112
Ethernet speed	10G	25G	100G	100G
Interface to ASIC (AUI)	1 x 10G XAUI	1 x 25G CAUI-1	2 x 50G 100GAUI-2	1 x **100G** 100GAUI-1
Modulation	10G NRZ	25G NRZ	**100G PAM4**	**100G PAM4**
Typical optical connection	LC (1 l)	LC (1 l)	LC (1 l PAM4 or BiDi) or MPO LR1, FR1, DR1, SR1.2, SR2	LC (1 l PAM4) LR1, FR1, DR1
Power consumption	1.5W	1.5-3W	4W	<3.5W

SFP112, SFP-DD and QSFP28: 100G Single Lambda

Complementary deployment options

Single Lambda portfolio on the fiber

- 100G SR1.2 100M MMF
- 100G DR 500 m SMF
- 100G FR 2 km SMF
- 100G LR 10 km SMF

Compatible with 4x100G QSFP

- 4x100G DR/FR/LR
- 8x100G DR/FR/LR

100G future proof investment

- PAM4 forward compatible

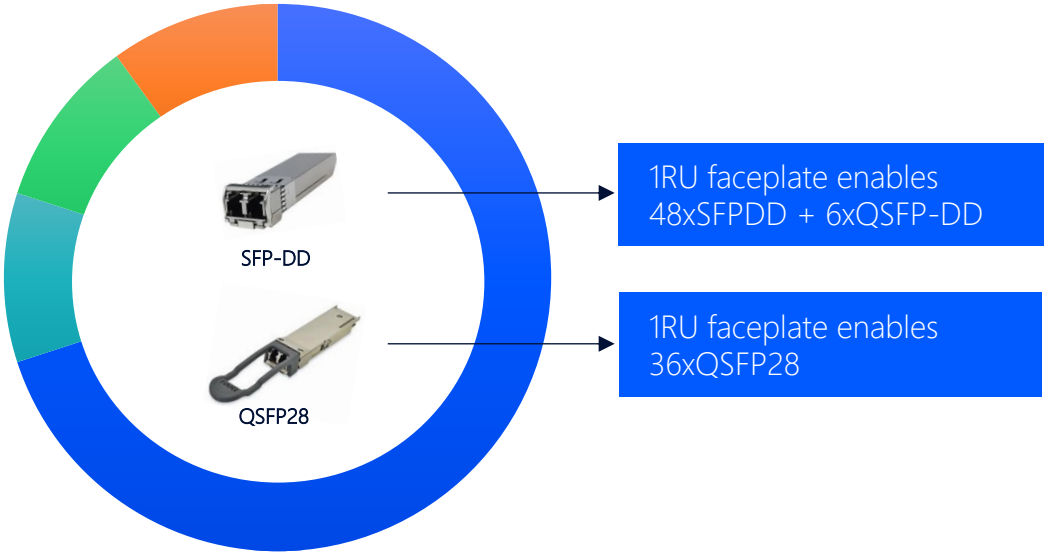

QSFP ‘Wider’ cages

4 channel

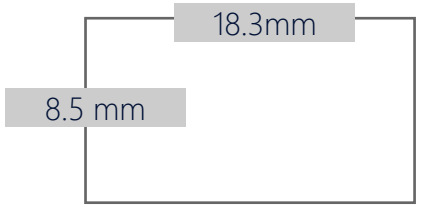

	QSFP+	QSFP28	QSFP-DD	QSFP112	QSFP-DD800
Ethernet speed	40G	100G	400G	400G	800G
Interface to ASIC	4 x 10G XLAUI	4 x 25G CAUI-4	8 x 50G 400GAUI-8	4 x **100G** 400GAUI-4	8 x **100G** 800GAUI-8
Modulation	10G NRZ	25G NRZ	**100G PAM4**	**100G PAM4**	**100G PAM4**
Typical optical connection	MPO LC (4 l MUX)	MPO LC (4 l MUX, or 1 l)	MPO LC (4 or 8 l MUX)	MPO LC (4 l MUX)	MPO LC (8 l MUX)
Power consumption	3.5W	3.5-5.5W	14W to 23.3W (coherent)	< 8W	16-18W

QSFP-DD800G

Introducing 800G

QSFP-DD MSA

Formal announcement of QSFP-DD800 MSA May 24, 2021

- Designed to be backwards compatible with existing QSFP-DD
- Point-to-point or breakout for high density 100G applications

Pluggable modules

Available

- QSFP-DD 800G DR8 (500m)
- QSFP-DD 800G DR8+ (2km)
- QSFP-DD 2x400G FR4 (2km)
- QSFP-DD 8x100G LR1 (10km)

100G SerDes

Architected to leverage the IEEE 802.3 ck work to fit 800G

- 100G PAM4 electrical, using 50Gbaud signaling

Transceiver

QSFP-DD form-factor

- 0/70°C case temperature
- 25% -43% power savings over 400G
- Price neutral to 400G

Clear economic and power advantages to 800G

NOKIA

What's the status of coherent transceivers?

800D QSFP-DD ZR/ZR+

- 800G @120GBaud
- Modulation and power
- 400G QPSK (LH) (~24w)
- 400G PCS (Regional) (~24w)
- 800G 16QAM (LR) (~22w)
- 800G 16QAM (OIF 800GZR) (~24w)
- 800G 16QAM (ZR+) (~27w)

Available Mid-2024
QSFP112-DD
QSFP-DD 800

400G QSFP-DD ZR+ 0 dBm

- Higher Tx power
- ~100km unamplified
- Up to 750 km (pt-pt) or 6-7 ROADMs*
- OpenZR+ MSA compliant
- OpenROADM 5.0 compliant
- FlexGrid (6.25, 12.5, 25, 50, 75, 100 GHz)

Available now
QSFP56-DD

400G QSFP-DD ZR/ZR+ -10 dBm

- OIF-400G-ZR compliant
- Up to 500km (pt-pt) or 3-4 ROADMs*
- OpenZR+ MSA compliant
- FlexGrid 6.25,12.5,25,50,75 & 100GHz.

Available now
QSFP56-DD

100G QSFP ZR

- Up to 300km amplified and up to 120 km unamplified reach
- Full C-Band Tunable 50/100 GHz grid
- ~6w

Available 2H-2023
QSFP28

*Typical performance only; actual capacity-reach performance dependent on specific link characteristics

How do we build for scale?

Key technology evolutions on the router enabling 400G – 800G - beyond

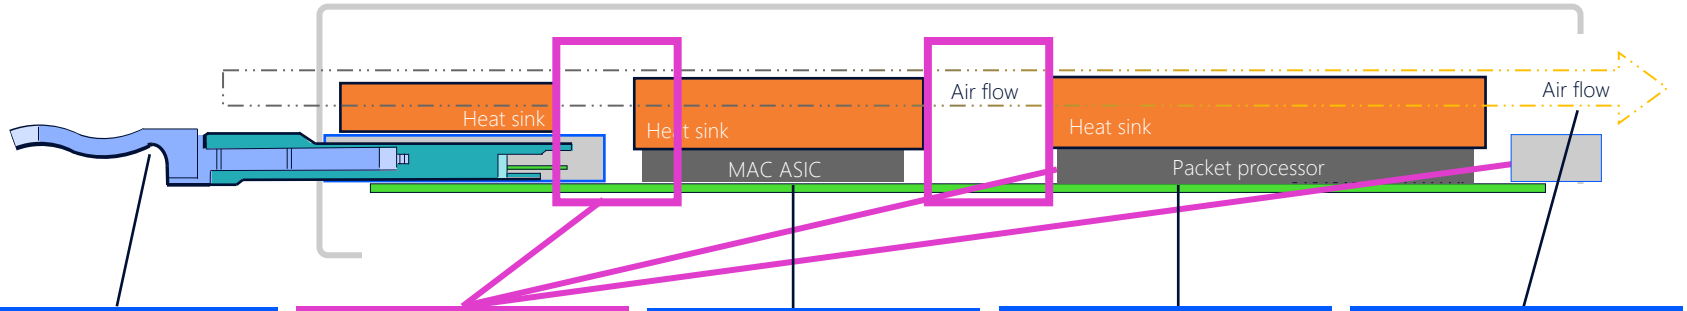

Pluggable optics and cage type
- Distance
- Density
- Compatibility

SERDES
- Coding
- Speed

MAC ASIC
- Single flow speed
- Flexibility
- Compatibility

Packet processor (NPU)
- Forwarding speed
- Scale
- Buffering

Air flow and power
- Effectiveness
- Design longevity
- Power budget

SerDes

- Serializer/deserializer
  - Connection between ASICs and towards cage
  - Increasing speeds of an individual lane: 10G, 28G, 56G
- Latest specification: 100G SerDes (802.3ck) for chip-to-chip or chip-to-module communication
  - Use of PAM4 modulation
  - Well-aligned with optics evolution (100G Lambda) – resulting in less active components and complexity in the transceivers
- Benefits
  - Higher I/O possible
  - Better power characteristics and cost
- Complex, but necessary evolution

How do we build for scale?

Key technology evolutions on the router enabling 400G – 800G - beyond

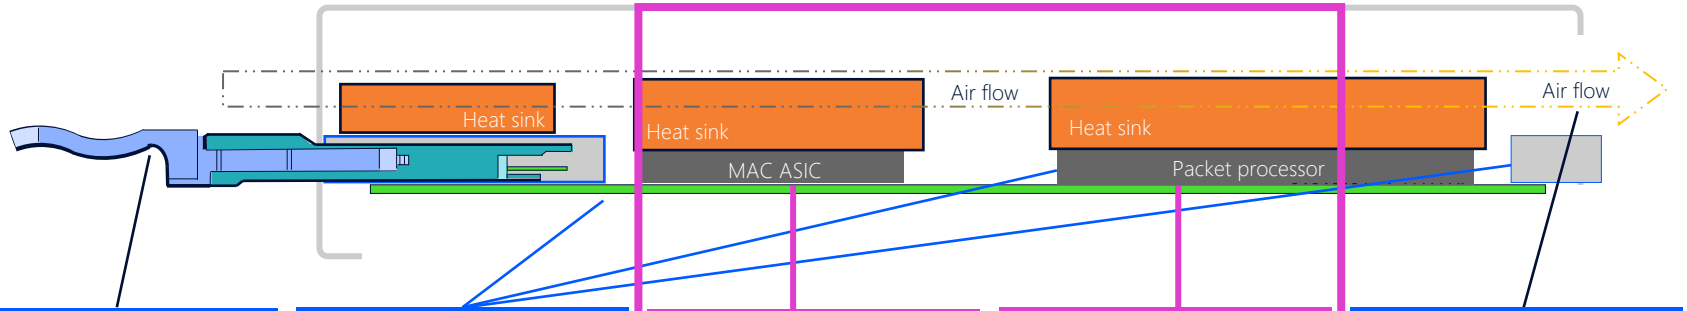

Pluggable optics and cage type
- Distance
- Density
- Compatibility

SERDES
- Coding
- Speed

MAC ASIC
- Single flow speed
- Flexibility
- Compatibility

Packet processor (NPU)
- Forwarding speed
- Scale
- Buffering

Air flow and power
- Effectiveness
- Design longevity
- Power budget

MAC ASIC and Packet Processor (NPU)

Evolving the router's data-plane to higher speeds, scale and capabilities

MAC

- 800GE and higher
- Enabling 'universal ports'
- Optional support for
  - MACSec
  - Flex-E
  - Intelligent aggregation

Store

- Buffer characteristics
  - Location (ingress, egress, both)
  - Size
  - Bandwidth (full vs. partial)

Forward

- Lookup/forwarding speed
- Scale
  - FIB scale
  - ACL scale
  - uRPF impact
- Scale

Enabling 800G+ interfaces requires an evolution across the main forwarding components* of the router

*Different implementations/combinations possible

The ease of having universal ports

MAC

QSFP+ QSFP-DD QSFP-DD800

QSFP28 QSFP112

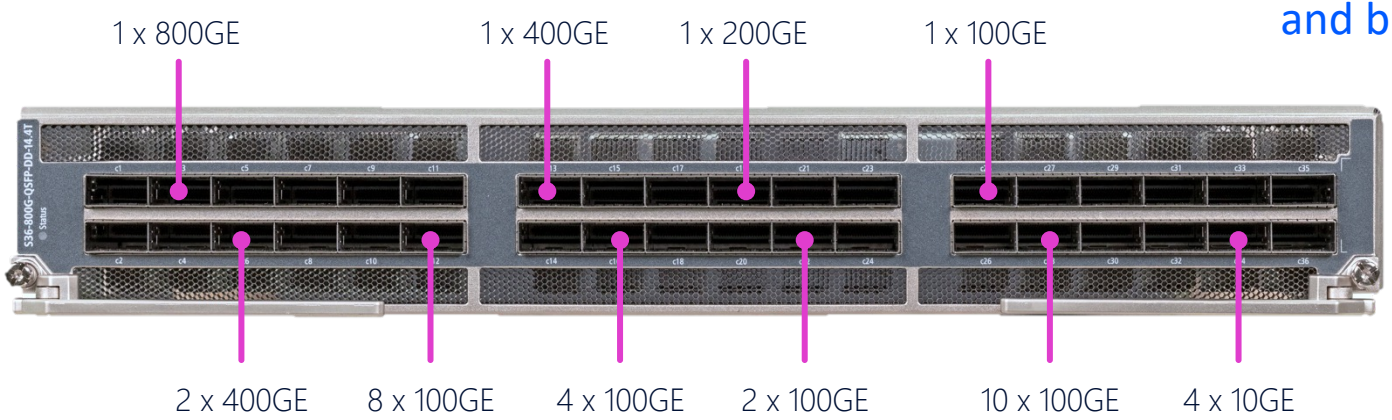

Your choice to use each port at whatever speed and breakout you like

How do we build for scale?

Key technology evolutions on the router enabling 400G – 800G - beyond

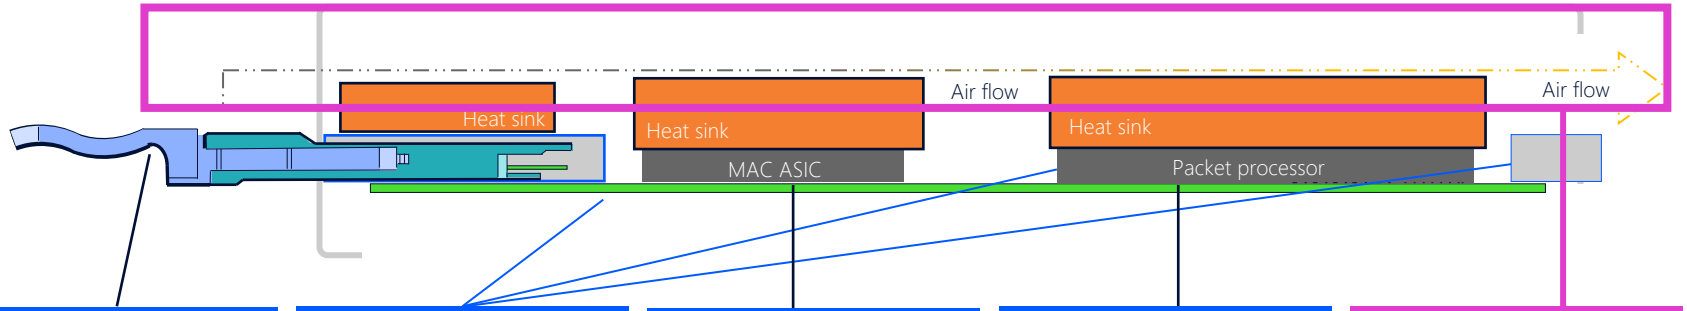

Pluggable optics and cage type
- Distance
- Density
- Compatibility

SERDES
- Coding
- Speed

MAC ASIC
- Single flow speed
- Flexibility
- Compatibility

Packet processor (NPU)
- Forwarding speed
- Scale
- Buffering

Air flow and power
- Effectiveness
- Design longevity
- Power budget

800G optics evolution

Cooling today's and tomorrow's optics

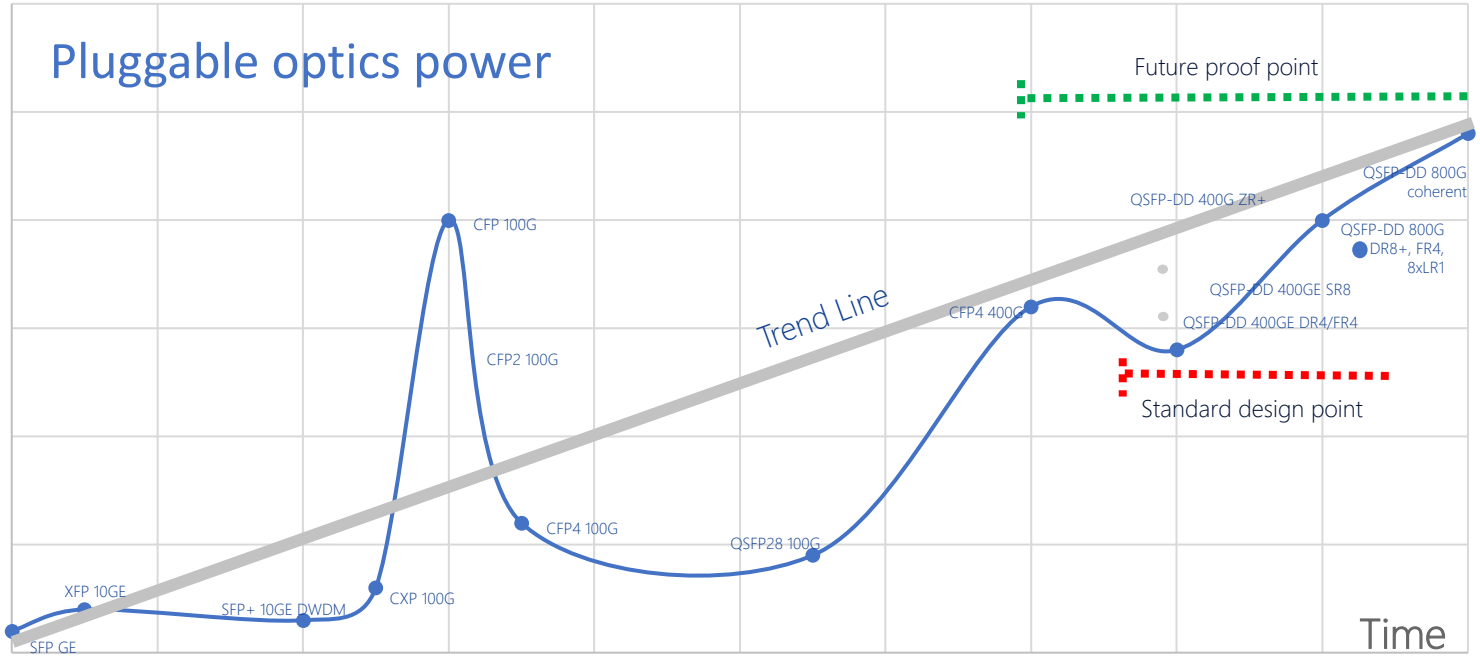

Stacking more linecards into a chassis => Chassis system architecture

Design considerations

- Mechanical design of huge significance
- Midplane vs. orthogonal direct cross connect
- Line card pitch and orientation
- Cooling design
- Power design
- Impacts
  - Density
  - Power consumption
  - Optics support

Classic backplane/ midplane approach

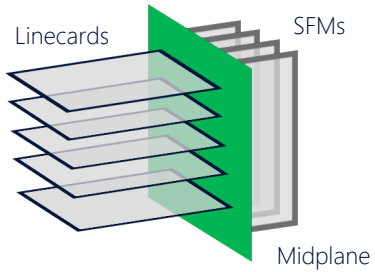

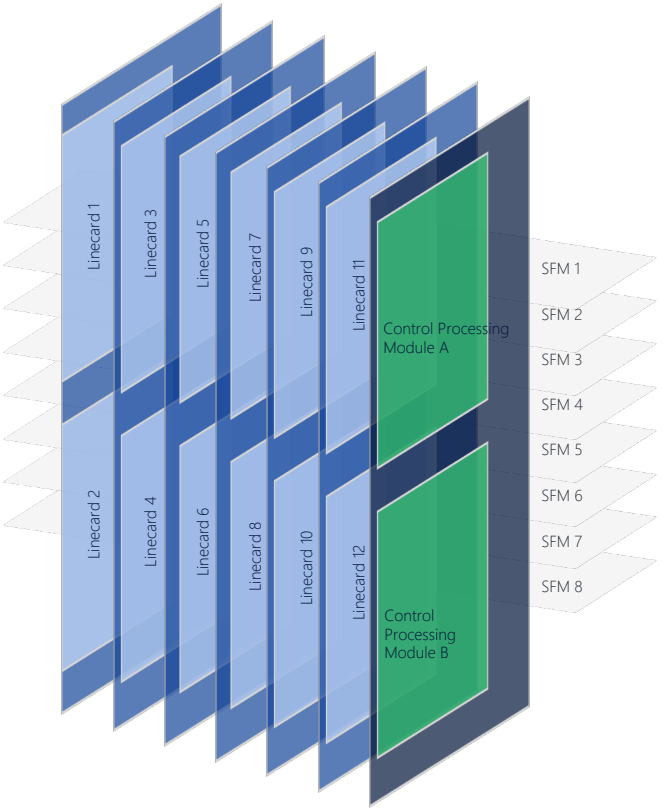

Optics cooling design

Single-sided PCB

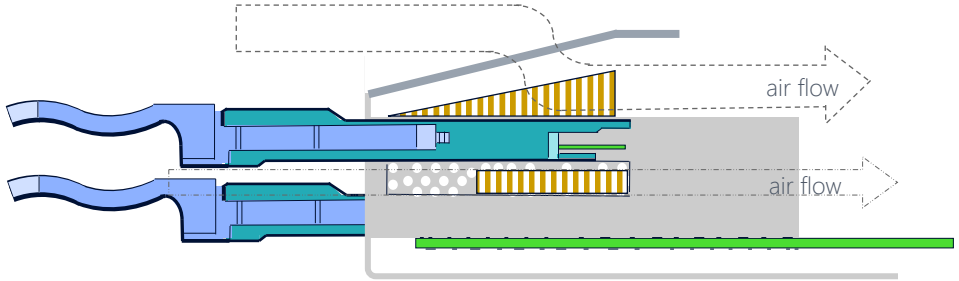

Stacked SFP cages

- Classic DC design
- Large heat sink only on top cage
- Bottom cage always hotter - imbalanced optical performance
- DD design point ~13W optics in all cages at 40°C
- Limits applicability to future optics
- Fans might have to run faster

Dual-sided PCB

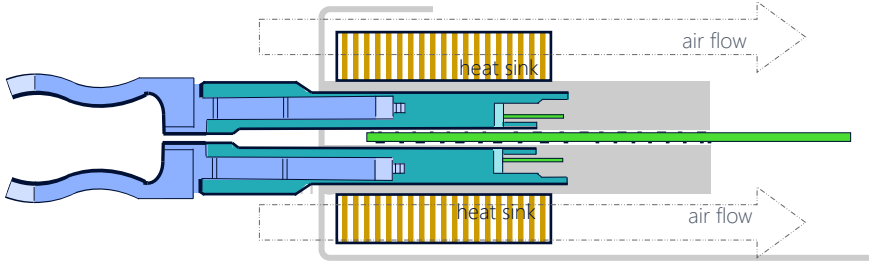

Belly-to-belly SFP cages

- Future proof design
- Large dedicated heat sink per cage
- Even cooling to all cages
- Cooling to 28W+ in all cages at 40°C

Enabling 800G and beyond on IP routers

Design choices along the datapath

Platform

- Mechanical design
- Power
- Cooling

Dataplane and chipset interconnect

- Forwarding
- MAC
- SERDES

Pluggable optics

- SFPDD-100, QSFP28, QSFP56-DD, QSFPDD-800

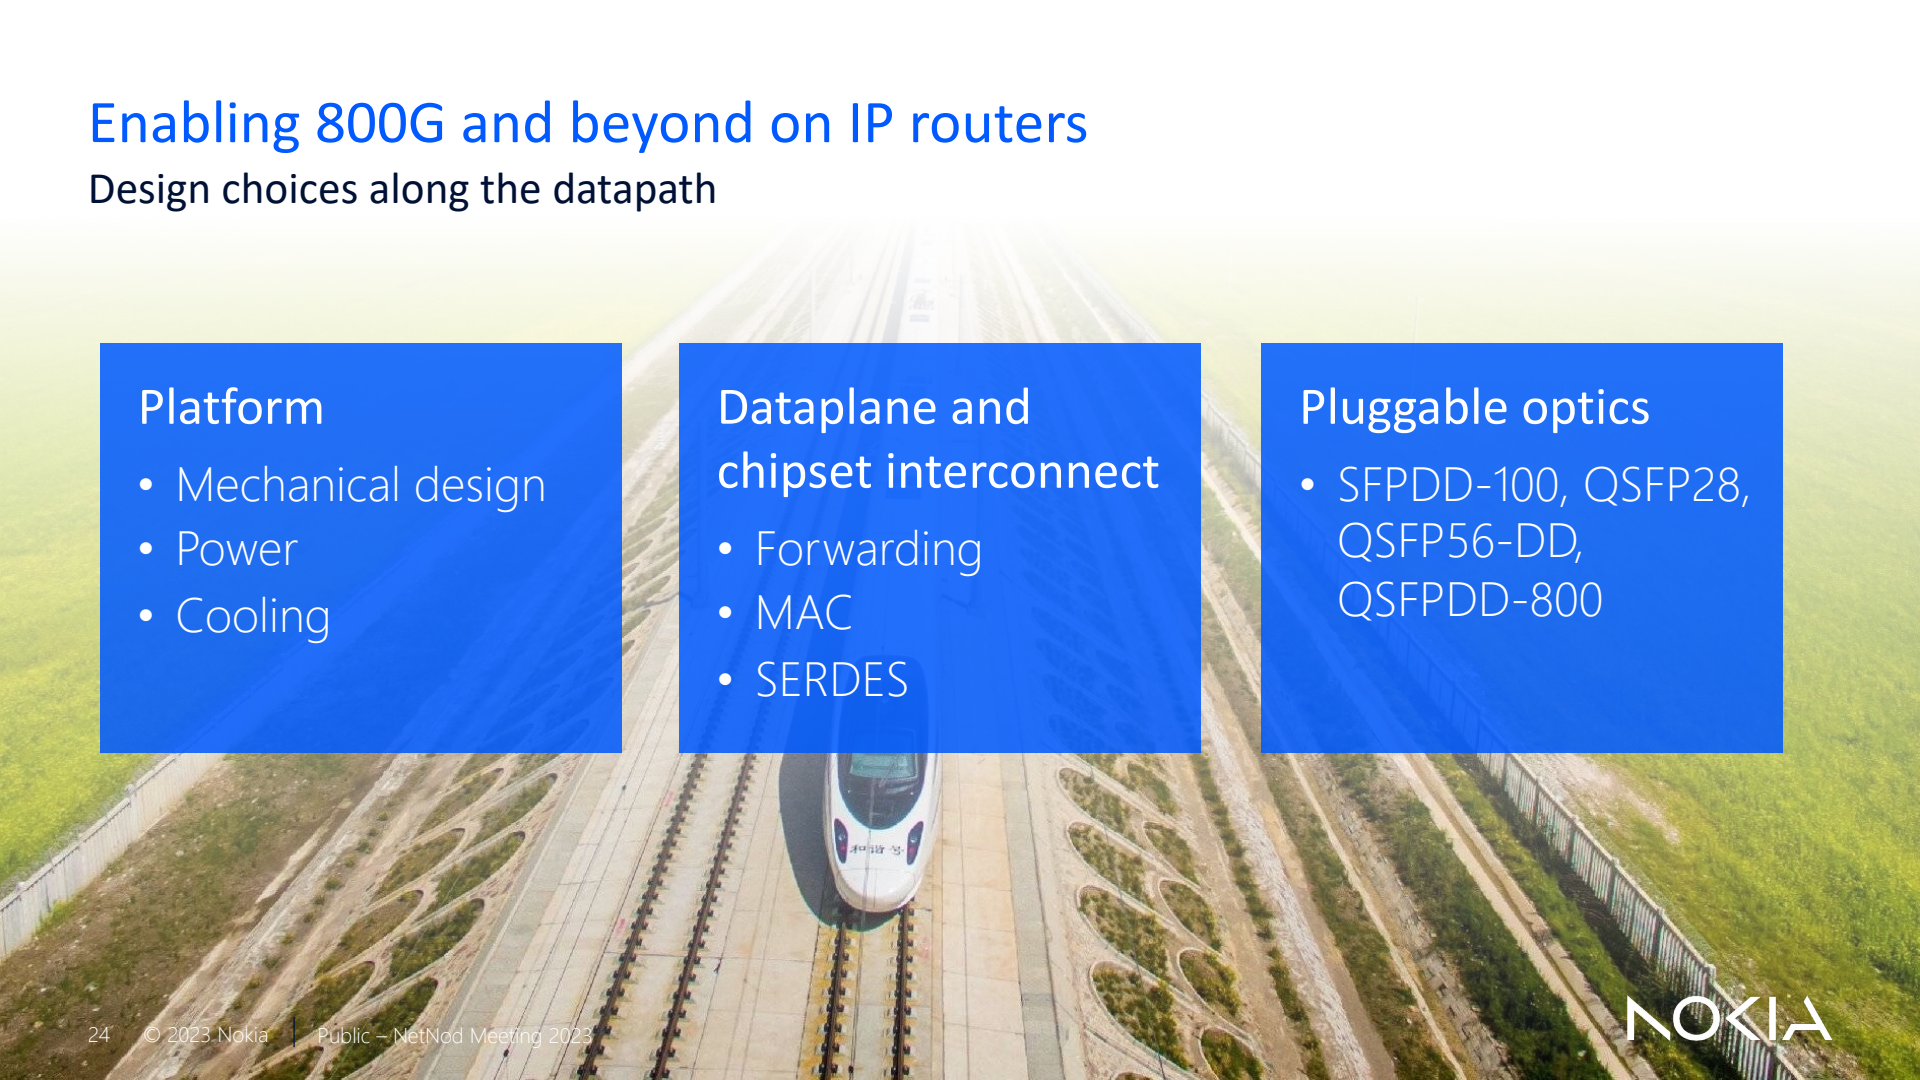


NOKIA