IP Routers: 400G and beyond

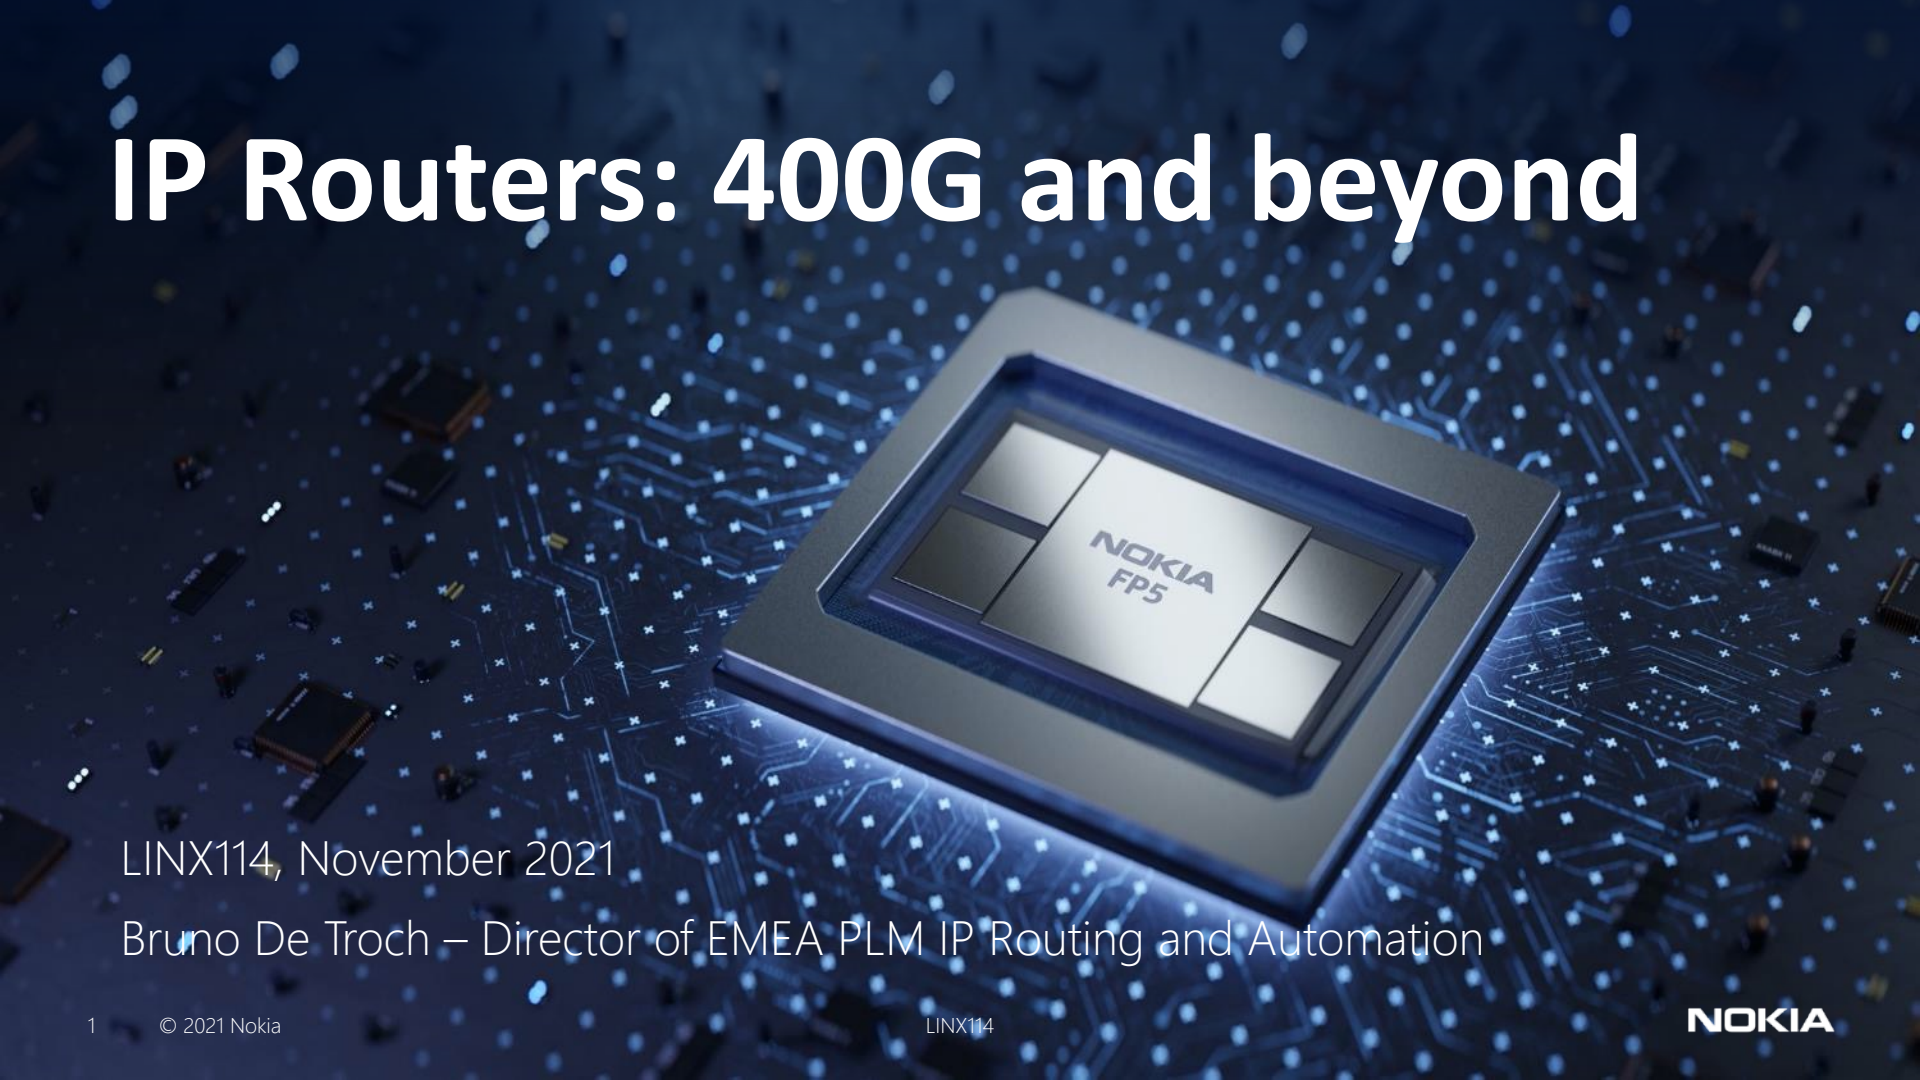

LINX114, November 2021

Bruno De Troch – Director of EMEA PLM IP Routing and Automation

NOKIA

Major industry trends

With COVID accelerating the impact

Digitalization of home, school and work has transformed our world and accelerated **data consumption**

3x Acceleration in global bandwidth consumption 2021

Today's fast growing network **threats** are just the tip of tomorrow's iceberg

100% YoY growth in DDoS traffic (2020)

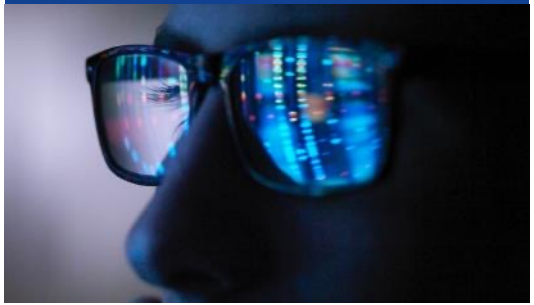

The planet is under stress and **sustainability** is now a pressing issue for global networks

50% Emissions reduction by Nokia products & operations by 2030

The 400G Wave

Satisfying the increased capacity requirements

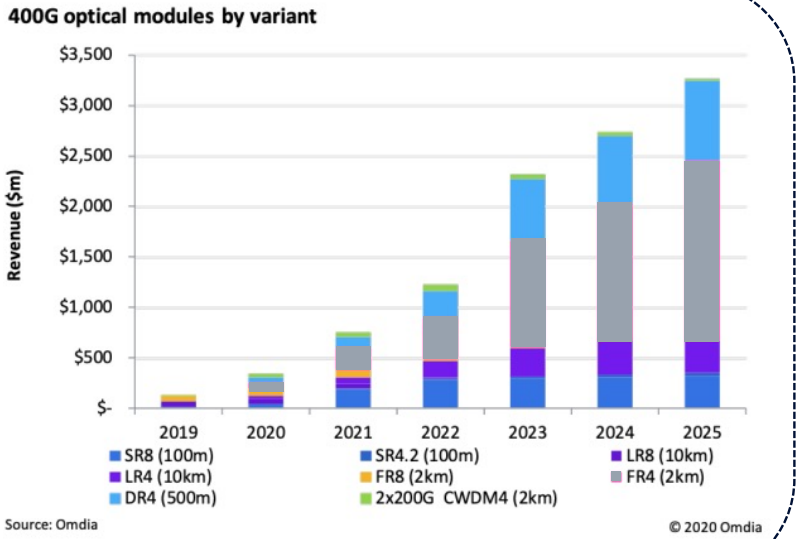

- Capacity requirements drive 400G optical module revenue
- Many module variants depending on the use case, cost and technology evolution
- 800G optical modules are around the corner (with benefits to early adopters)

400G+ wave driven via router innovation

Enabling the 400G+ Wave
Key technology evolutions

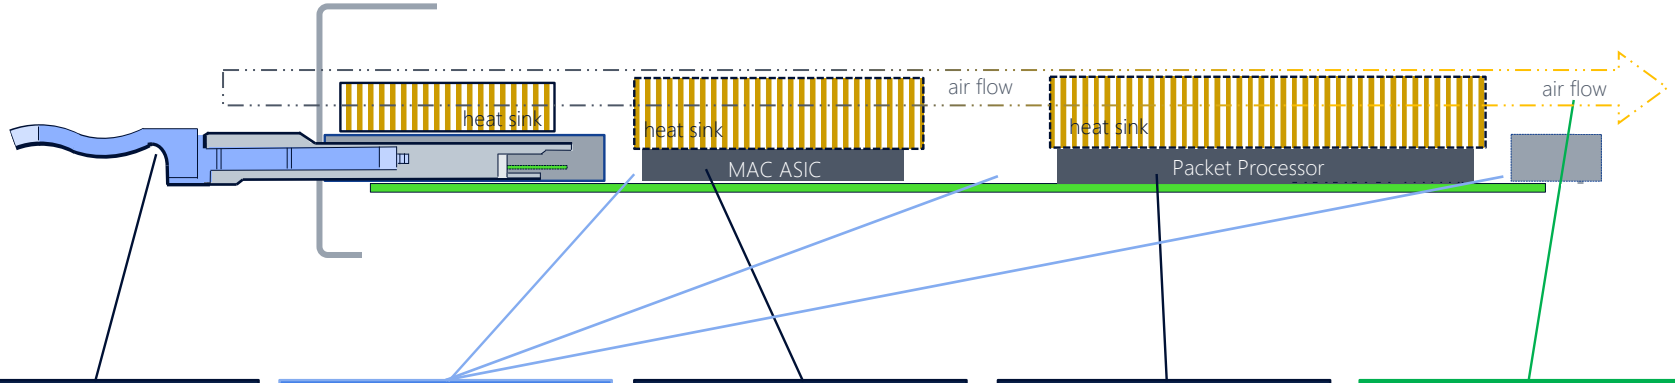

Pluggable Optics & Cage Type
- Distance
- Density
- Compatibility

SERDES
- Coding
- Speed

MAC ASIC
- Single Flow Speed
- Flexibility
- Compatibility

Packet Processor (NPU)
- Forwarding Speed
- Scale
- Buffering

Air Flow
- Effectiveness
- Design longevity

NOKIA

Pluggable optics and cage types

Diversity and uniformity

Cage types becoming universal

- **100G+**: QSFP28, QSFP28-DD, QSFP56, QSFP-DD 400, QSFP-DD 800
- **100G-**: SFP, SFP28, SFP56, SDF-DD

Pluggable optics in many shapes and flavors

- **100G Lambda**: PAM4, 800G, higher 100G/400G density, cost reduction
- **400G ZR/ZR+**: Coherent to enable "Pragmatic IPoDWDM" designs

800G

Around the corner

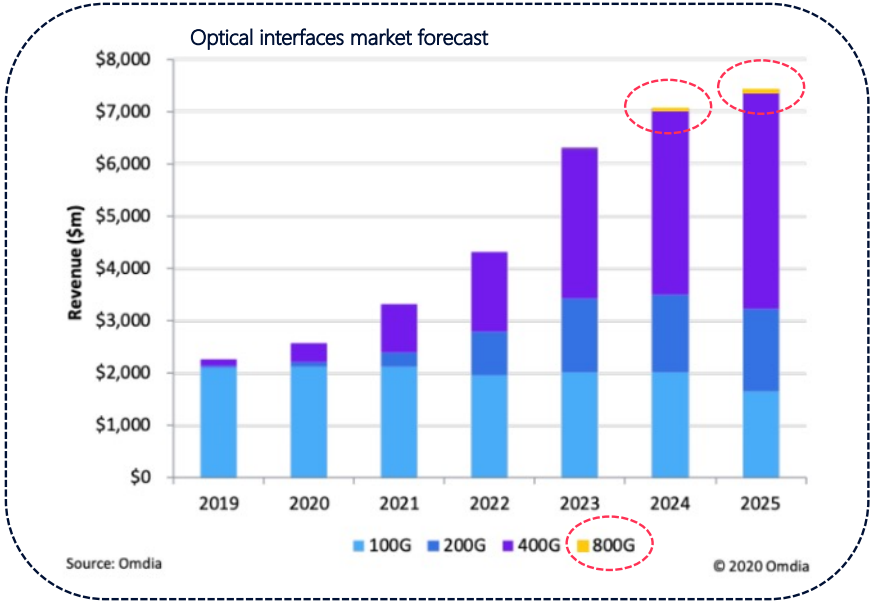

- First QSFP-DD 800 enabled router ports in 2022
- 1x800GE clear-channel standard in process
- First 800G optics mid-'22
  - 2x400 & 8x100 (16-18W)
  - 25% -43% power savings over 400G
  - Price neutral to 400G

800GPluggable MULTI-SOURCE AGREEMENT

IEEE 802

Clear economic and power advantages to 800G

Pluggables for 400G and Beyond …

Optical interface technology enablers

Formfactor

- Packaging
- Mechanics and Cooling
- Router Interface density

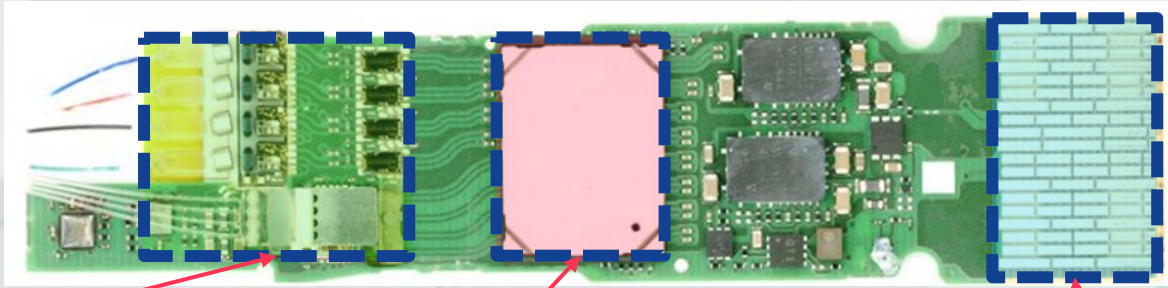

Arbitrary example photo of an optical pluggable PCB for illustration purpose

Photonics & Drivers

- Optical modulation and number of wavelengths (lambda's) are the key factor affecting cost and performance.
- 100G Lambda is BARE-MINIMUM for 800G

100G Lambda MULTI-SOURCE AGREEMENT

DSP/MLG

- 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
- 400G today relies on 56G SerDes
- 100G SerDes is KEY for 800G

IEEE P802.3ck

SerDes

- Serializer/Deserializer
  - Connection between ASICs and towards cage
  - Increasing speeds of an individual lane: 10G, 28G, 56G
- Latest specification: 100G SerDes (802.3ck)
  - Use of PAM4 modulation
  - Well-aligned with optics evolution (100G Lambda)
- Benefits
  - Higher I/O possible
  - Better power characteristics and cost
- Complex, but necessary evolution

MAC ASIC and Packet Processor (NPU)

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

Enabling 400G+ interfaces requires an evolution across the main forwarding components* of the router

MAC

- 100GE, 400GE and evolution to 800GE (and higher)
- 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
- QoS support

* Different implementations/combinations possible

System architecture
Design Considerations

Mechanical design of huge significance

Midplane vs. Orthogonal Direct Cross Connect

Line card pitch & orientation

Cooling design

Power design

Impacts

- Density
- Power consumption
- Optics support

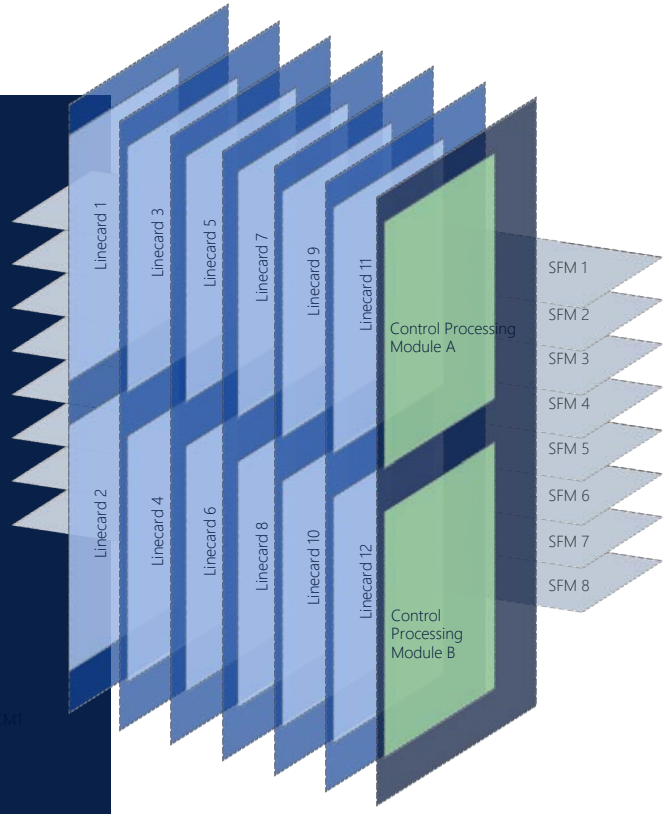

Optics power evolution

The challenge of cooling today's and tomorrow's optics

LINX114

Optics cooling design

Single sided PCB

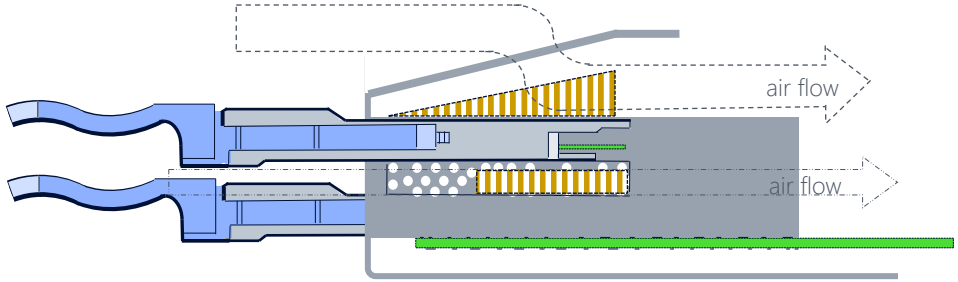

Dual sided PCB

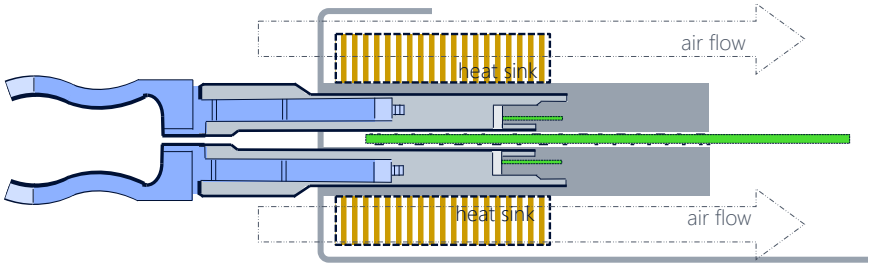

Stacked SFP Cages

- 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 40C
- Limits applicability to future optics
- Fans might have to run faster

Belly-to-Belly SFP Cages

- Future proof design
- Large dedicated heat sink per cage
- Even cooling to all cages
- Cooling to 26W+ in all cages at 40C

Enabling 400G and beyond on IP routers

Design choices along the datapath

Platform

Mechanical design
Power
Cooling

Dataplane & chipset interconnect

Forwarding
MAC
SERDES

Pluggable Optics

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

NOKIA