Meta Data Centers Heterogenous Integration Driven by AI/ML and Network Applications

Dr. Ravi Agarwal

Advanced Packaging, Technical Sourcing

Agenda

AI/ML Applications

Network Power

Chiplet


Meta






WHATSAPP

65 BILLION

Messages sent per day

>2 BILLION

Minutes of voice and video calls per day

MESSENGER

81 BILLION

Messages sent between businesses and customers each month

VIDEO

3.5 BILLION

FB Live Broadcasts

>10 MILLION

FB Live Broadcasts on New Year's Eve

ML/AI

50%

Of the data warehouse is used to feed ML workload

Source: Public: Vijay Rao OCP Keynote 2019


PAPILLION, NE


LOS LUNAS, NM


PRINEVILLE, OR


FOREST CITY, NC


FORT WORTH, TX


NEWTON, GA


NEW ALBANY, OH


ODENSE, DENMARK


LULEÅ, SWEDEN


ALTOONA, IA


CLONEE, IRELAND


HENRICO, VA


EAGLE MOUNTAIN, UT


HUNTSVILLE, AL


SINGAPORE

AI/ML APPLICATIONS

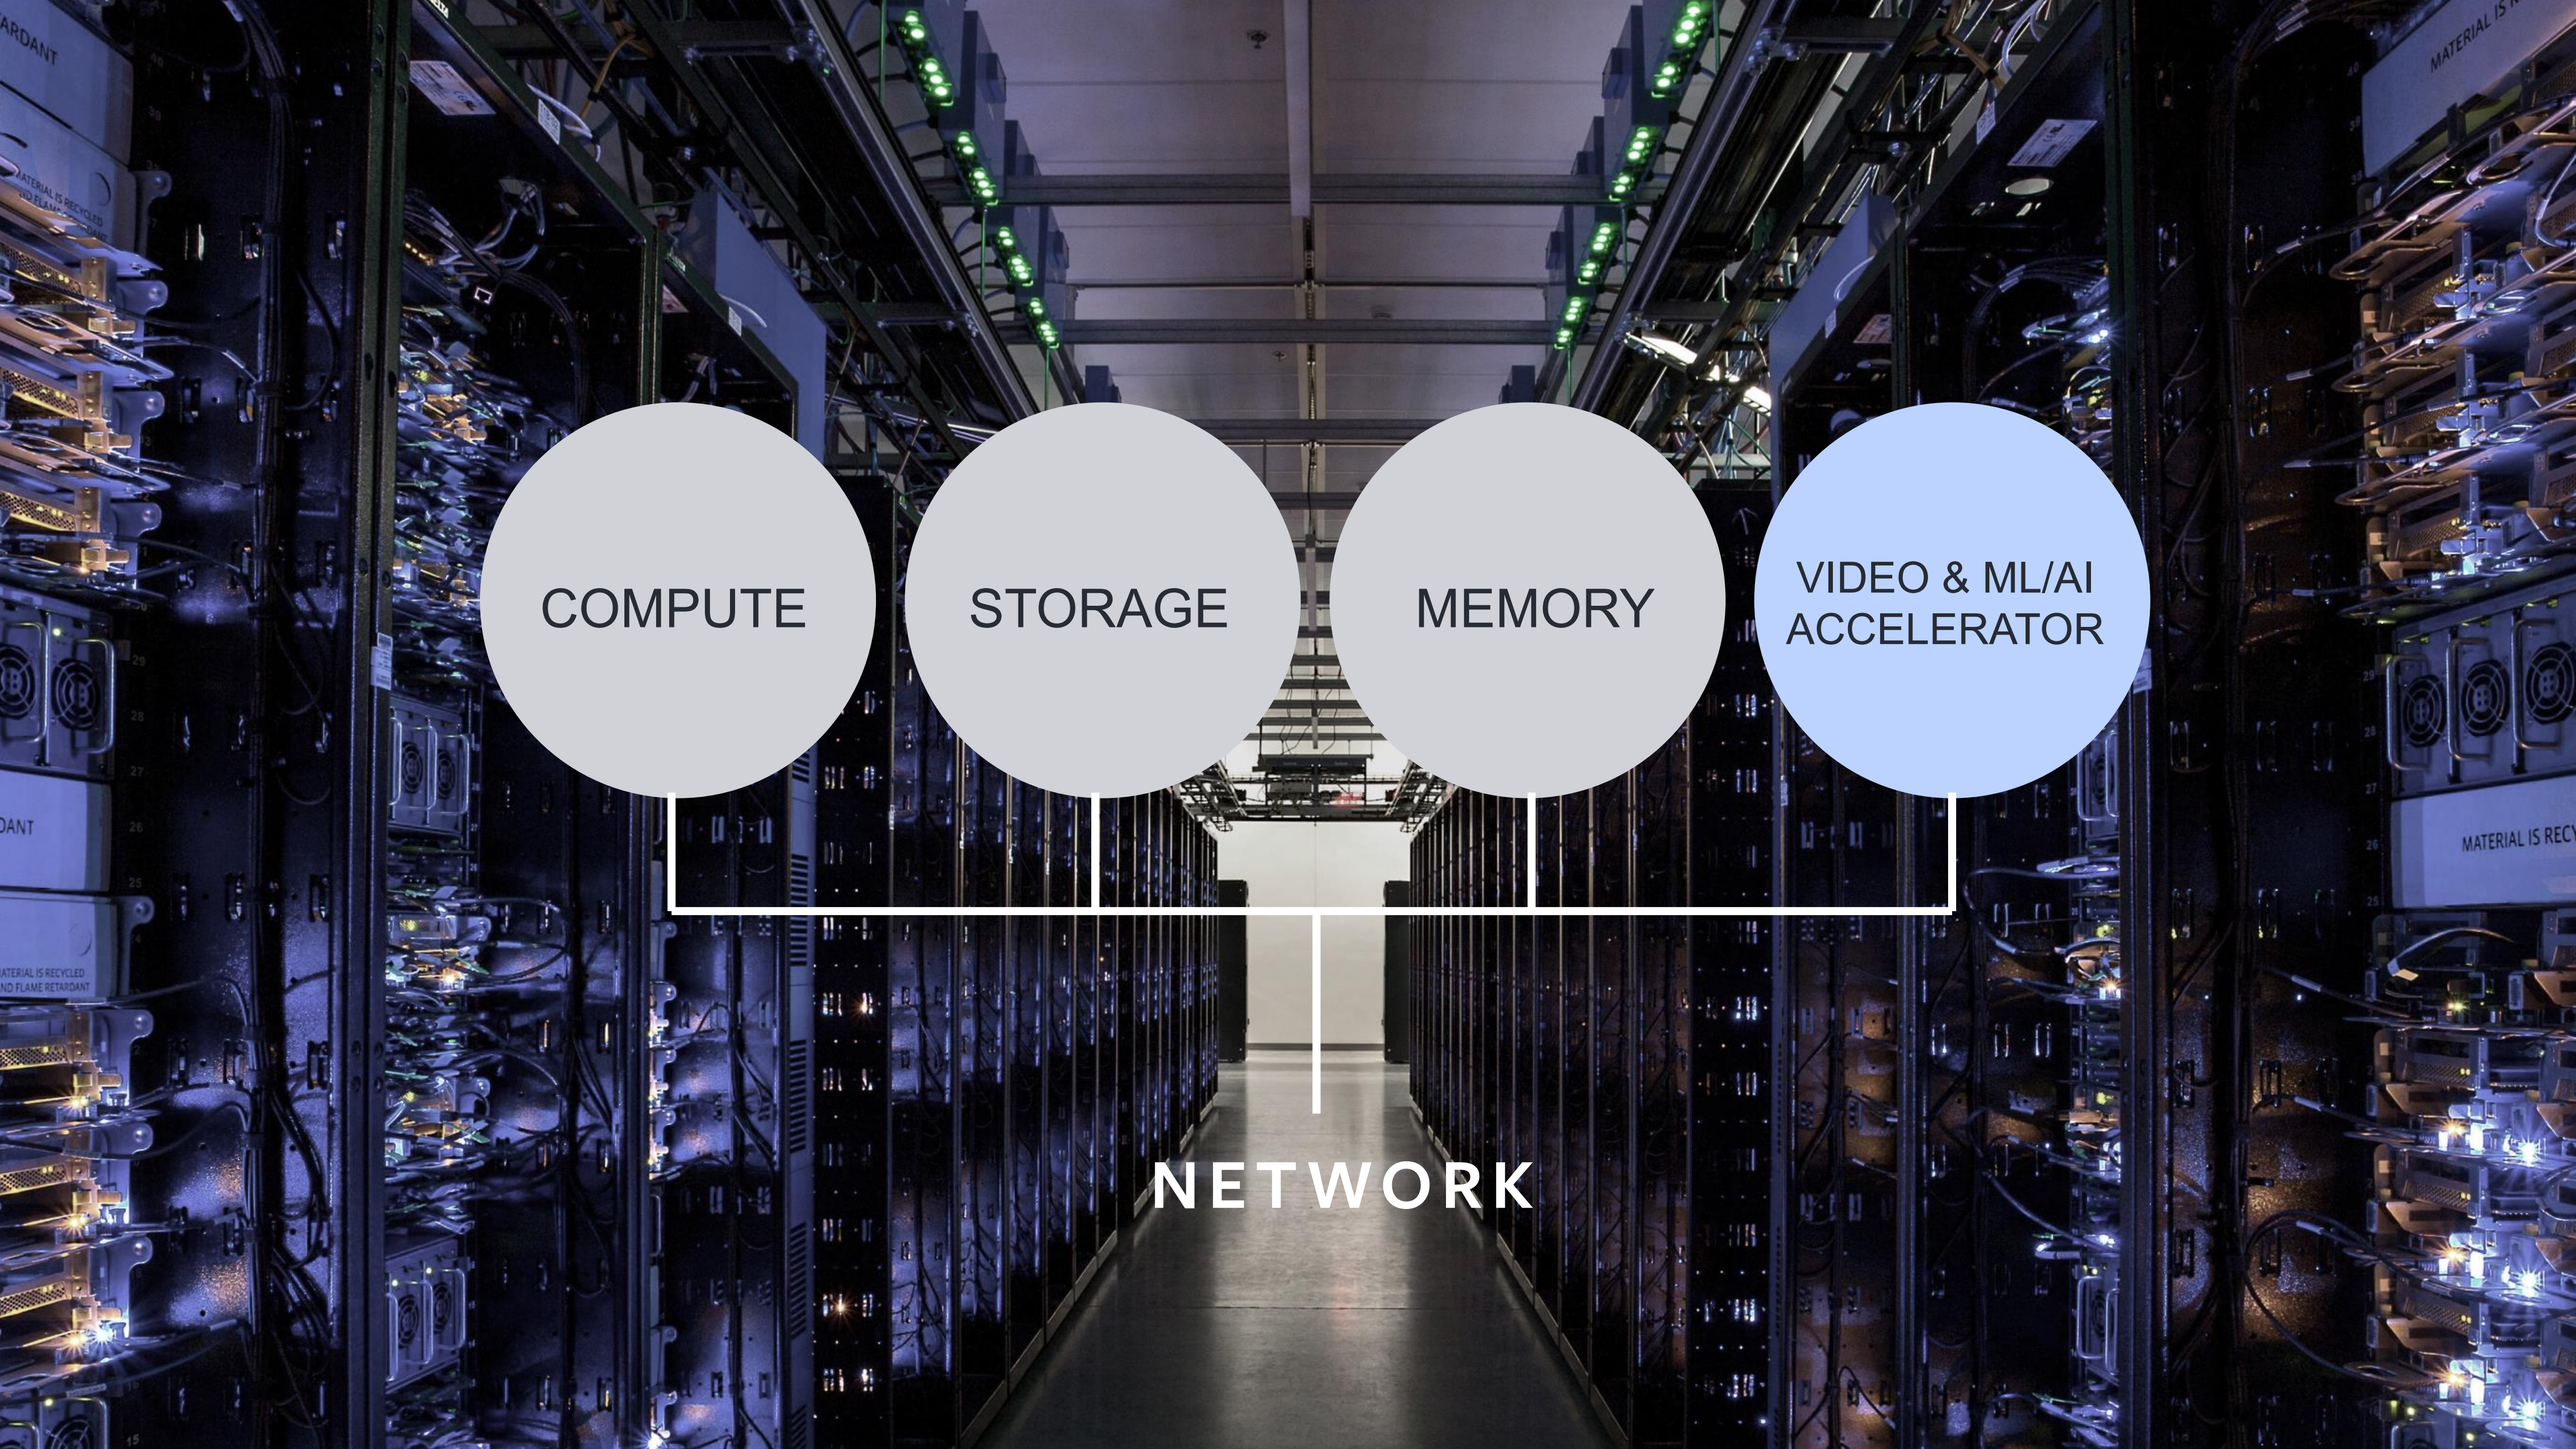
COMPUTE
STORAGE
MEMORY
VIDEO & ML/AI ACCELERATOR
NETWORK

Workloads @ Meta

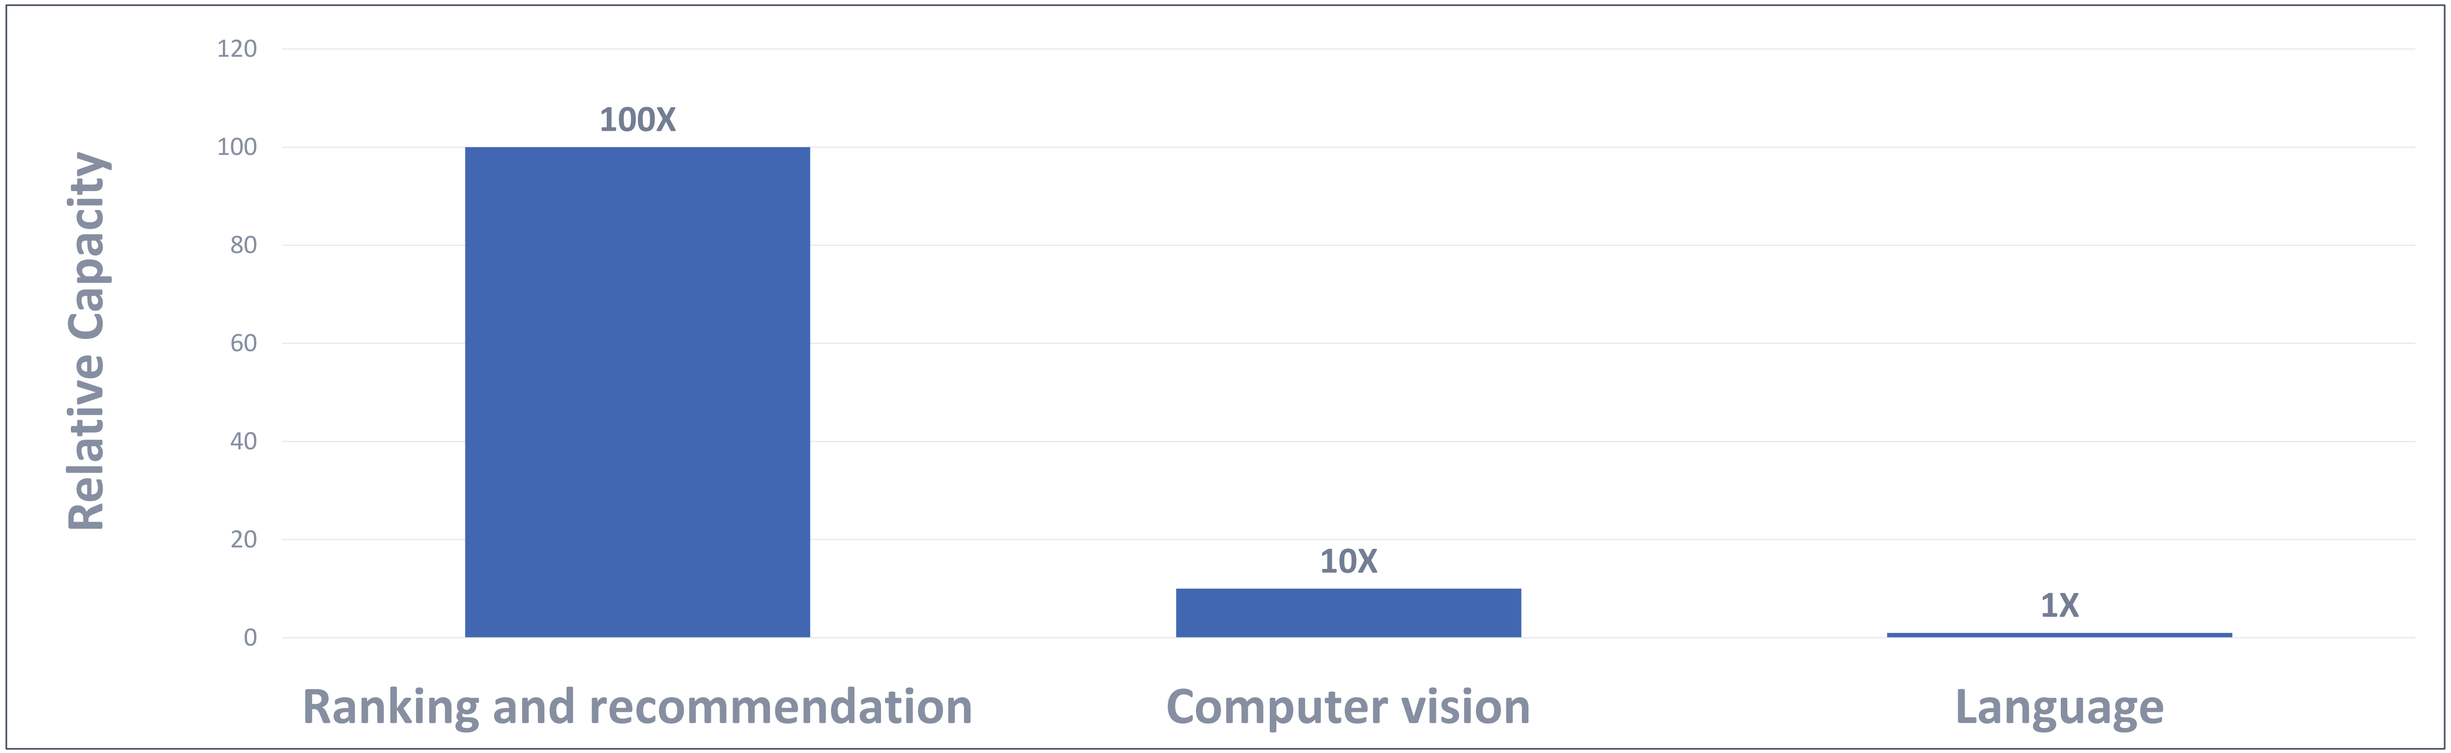

- Recommendation models are among most important models
- Meta shared DLRM model https://github.com/Metaresearch/dlrm

AI Inference Demands Across our Datacenters

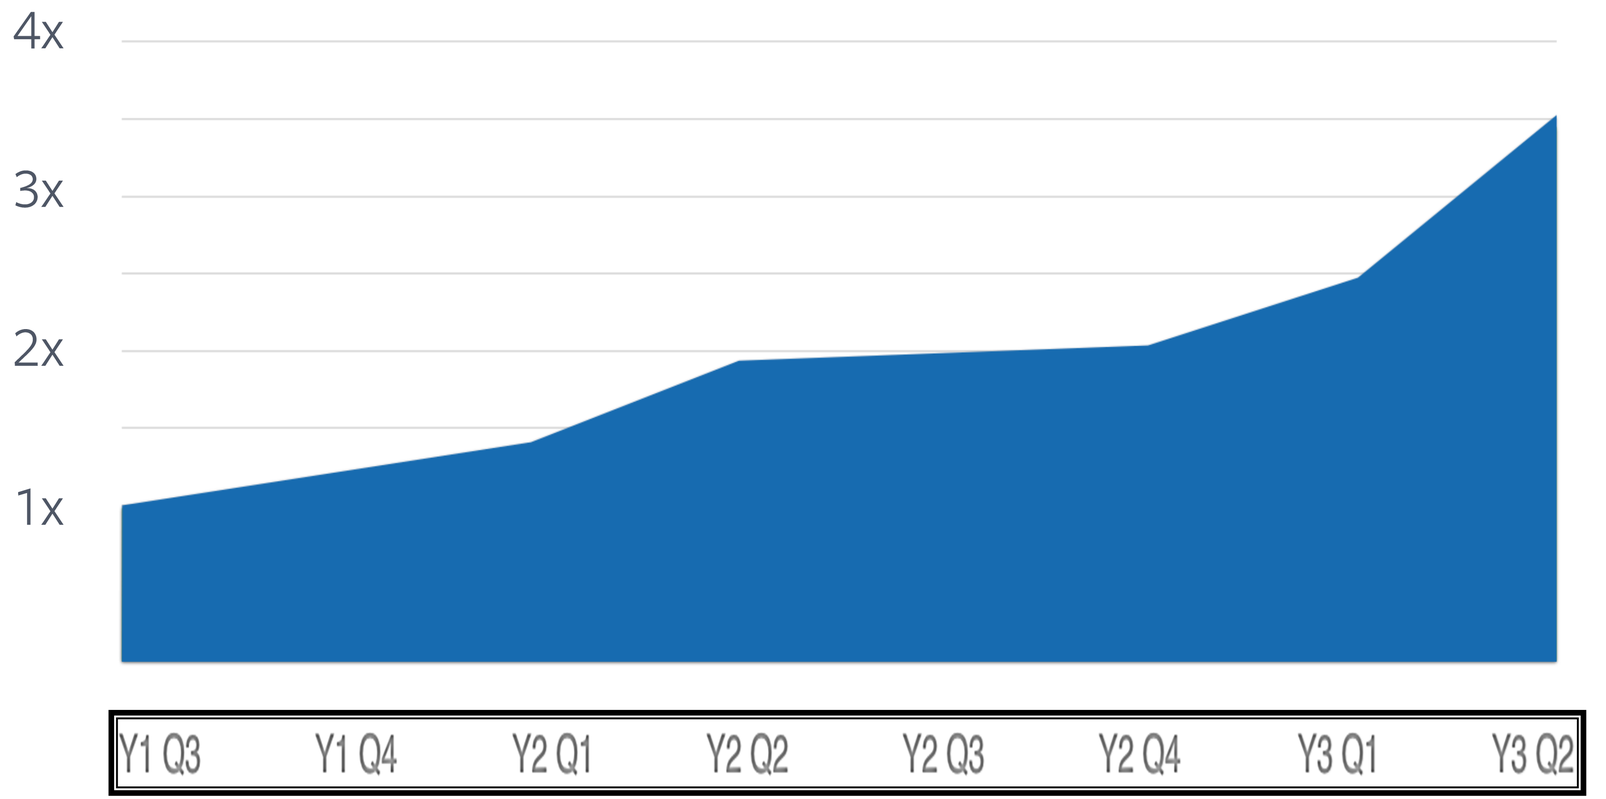

ML Growth and Scale at Meta

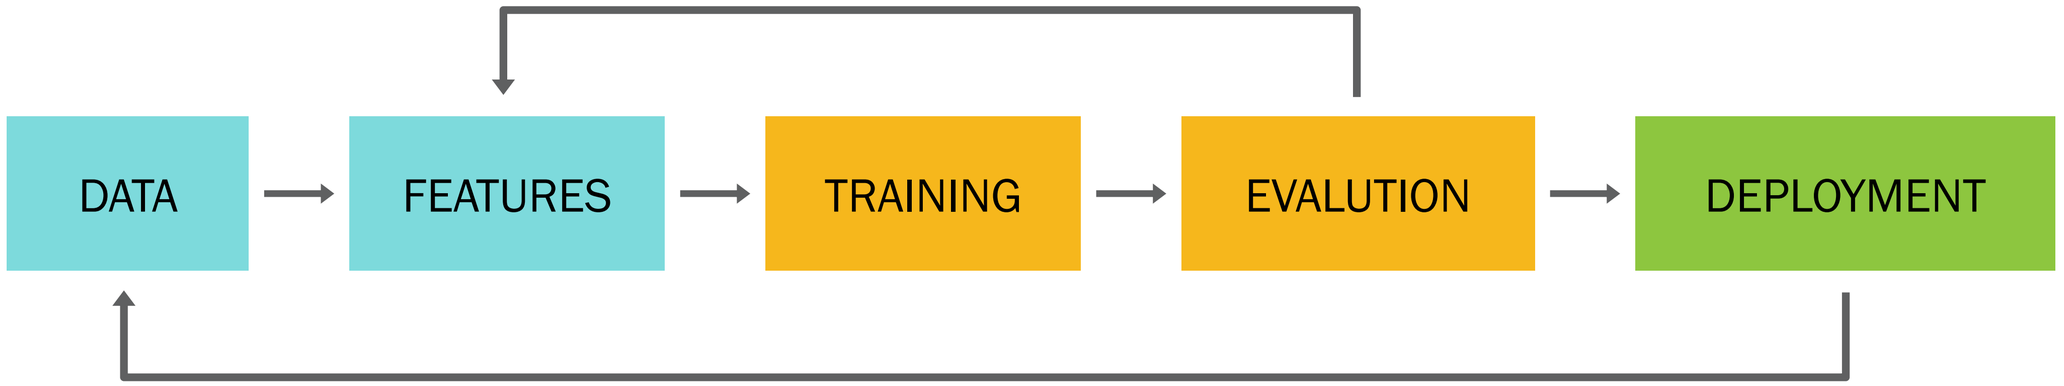

ML pipeline data growth

- Usage in 2018: 30%
- Usage 2019: 50%
- Growth in one year: 3X

1-year training growth

- Ranking engineers: 2X
- Workflows trained: 3X
- Compute consumed: 3X

Inference scale per day

- # of predictions: 400T
- # of translations: 6.5B
- Fake accounts removed: 99%

Accelerator System Logical View

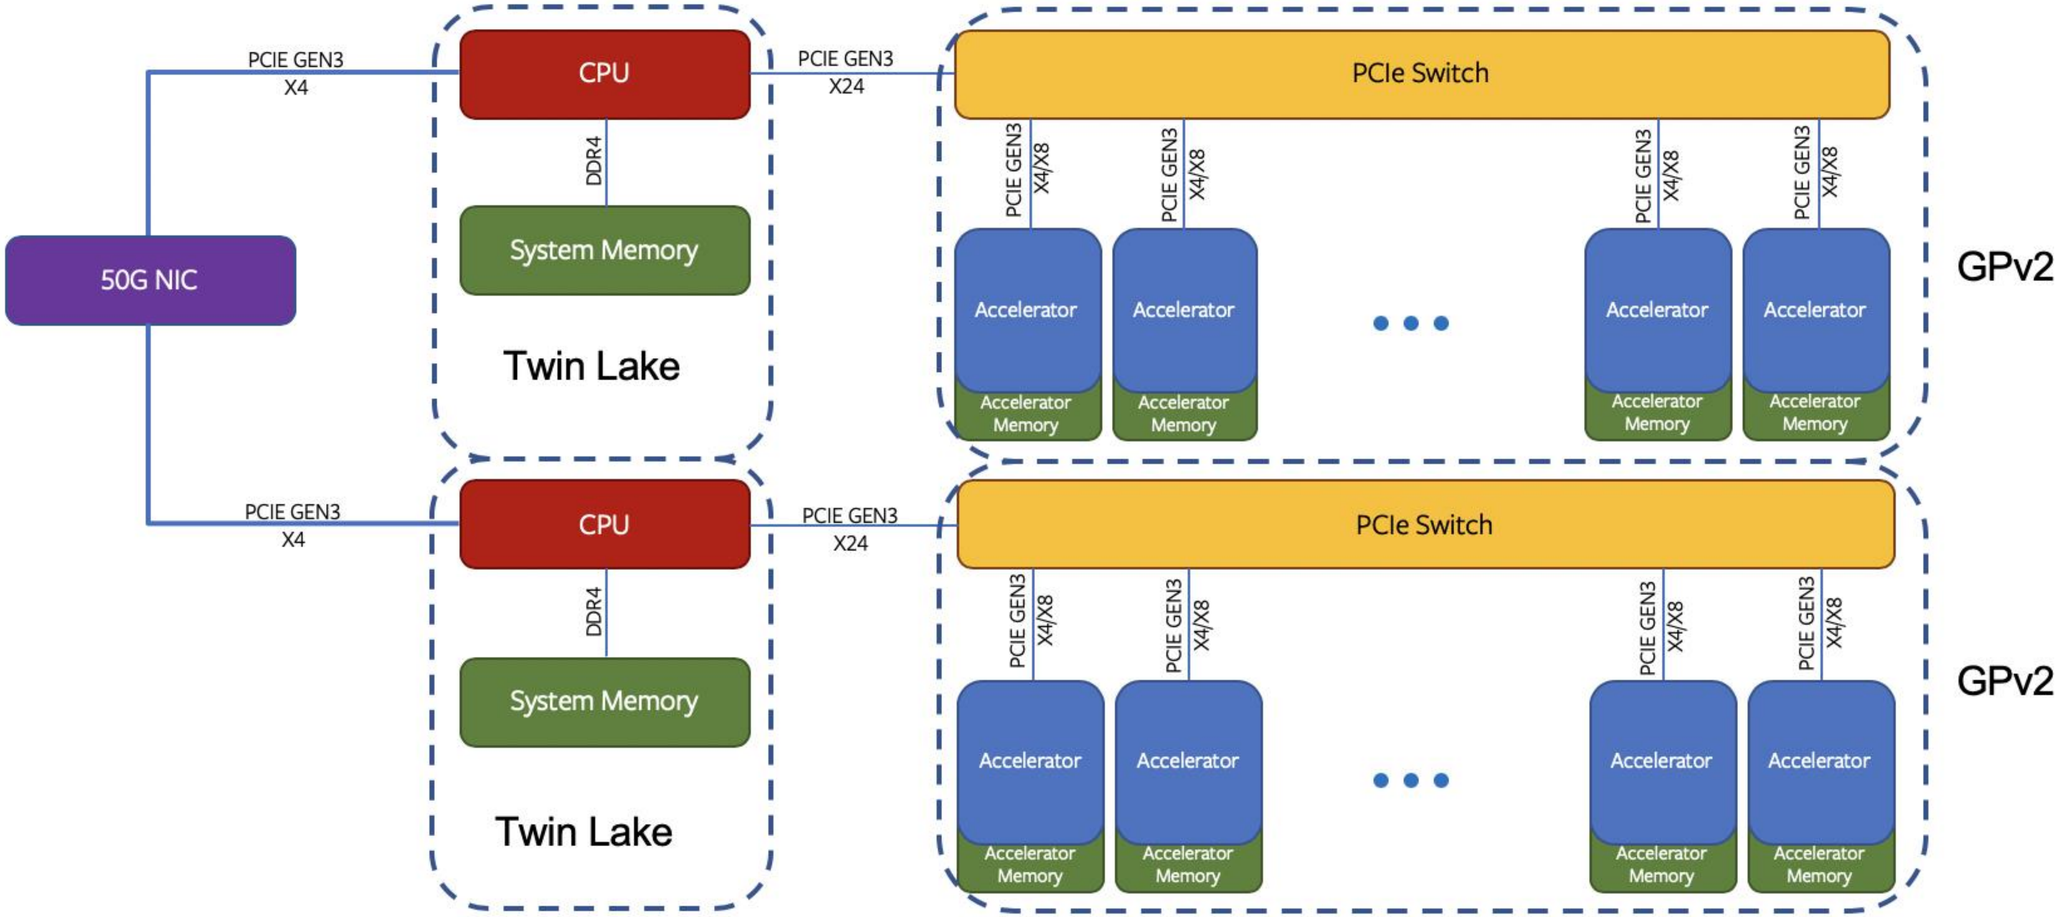

- Accelerator considerations include they can plug into existing DC/edge server platforms

Hardware System Considerations

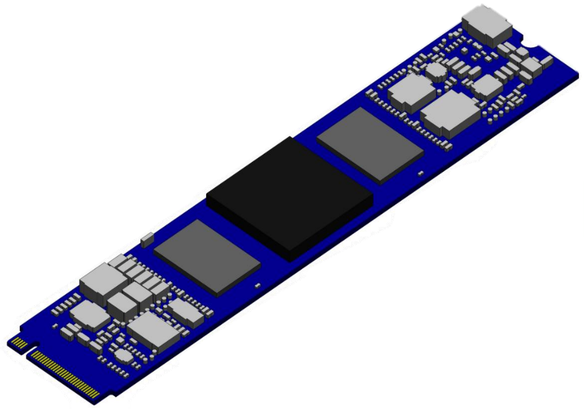

Single M.2

Module

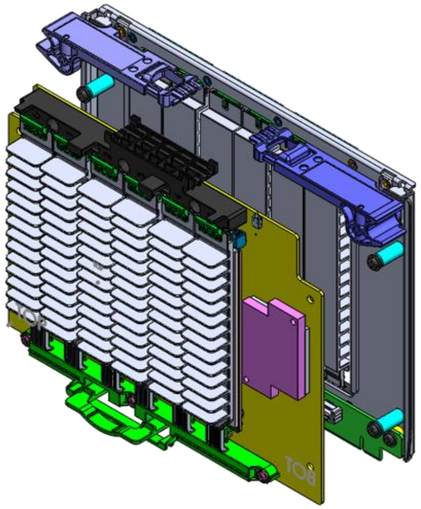

12x Modules in GPv2 Card

Carrier

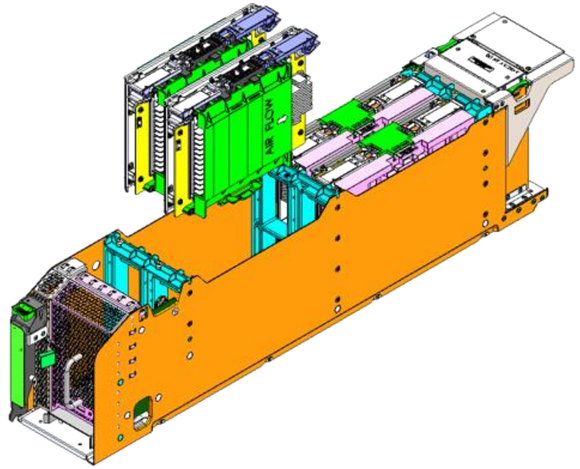

2x CPU cards and 2x GPv2 cards using 1x shared NIC

Sled

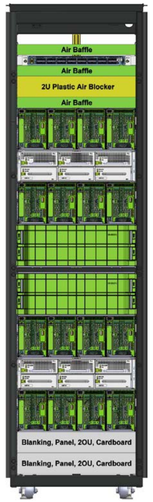

16x Sleds

Rack

- Next level considerations include how the module is integrated into the sled/rack

NETWORK POWER CHALLENGES

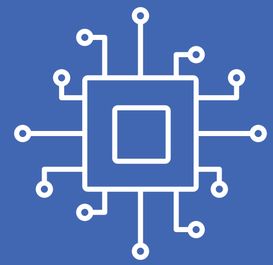

Growing Power Allocation for Networking

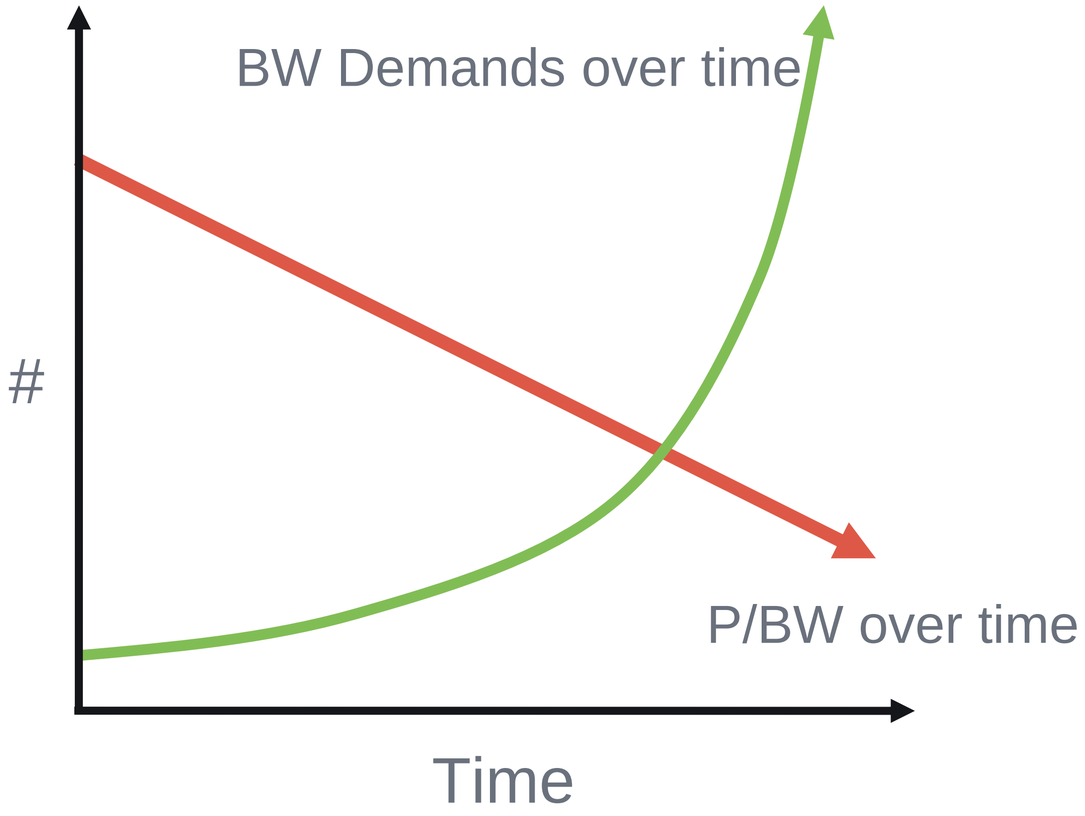

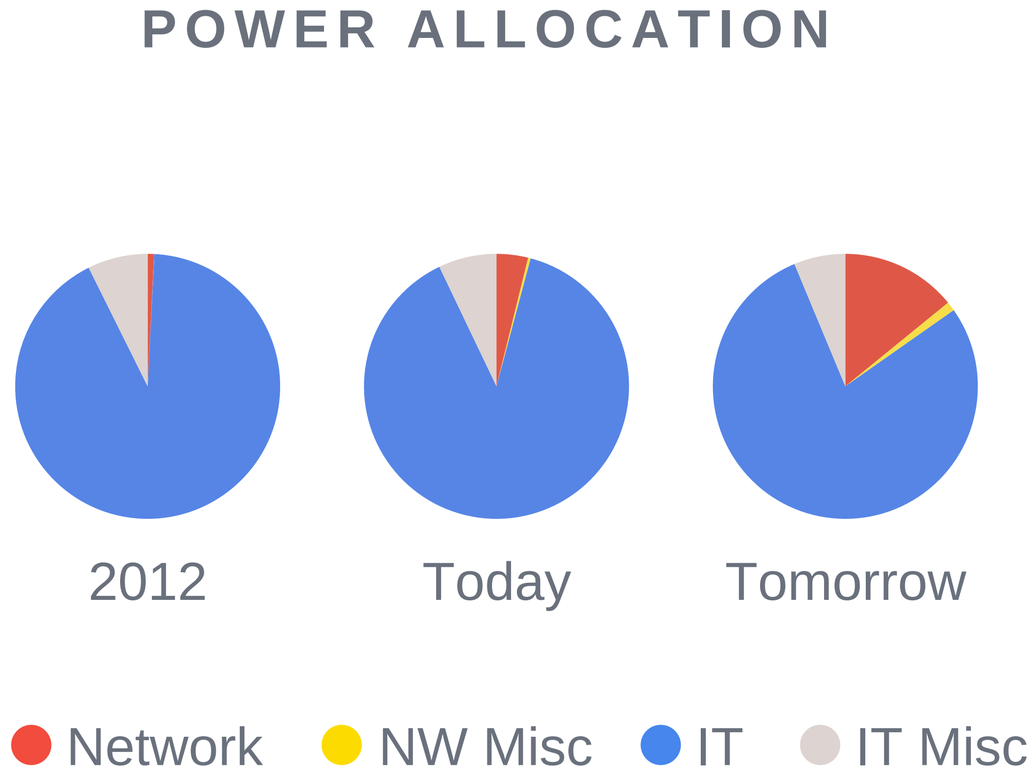

Network Switch Power Per Function

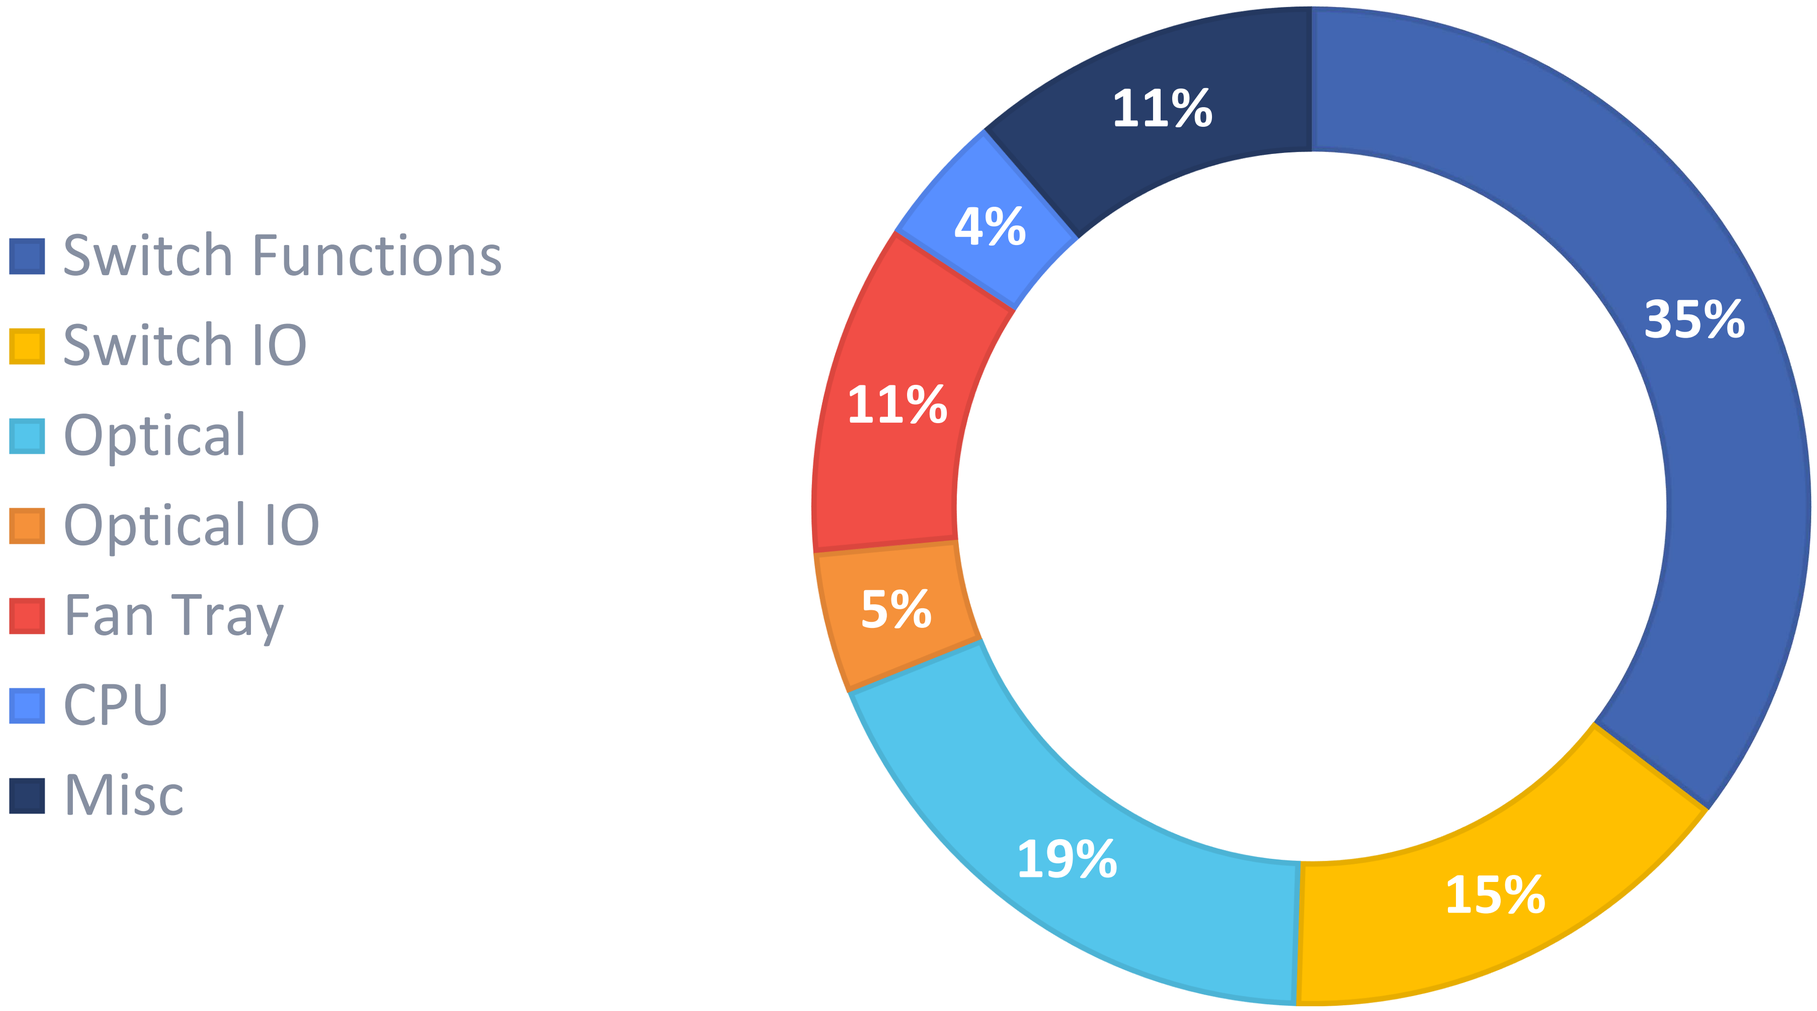

Evolution of Optical Interconnect

Pluggable Transceiver

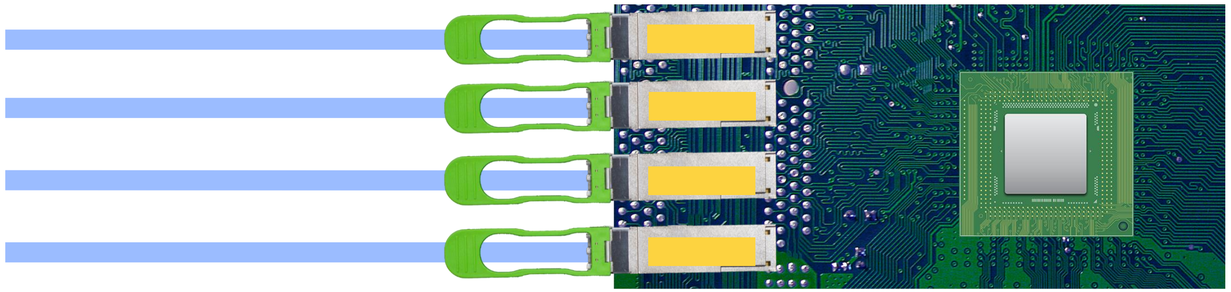

Mid-board Optics

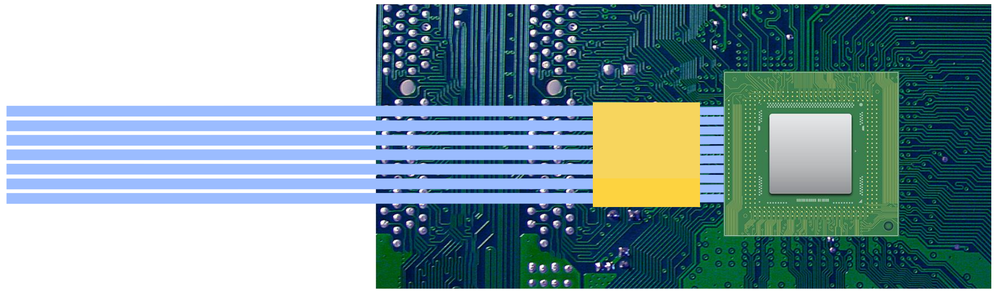

Co-packaged Optics

CPO: http://www.copackagedoptics.com/

CPO Packaging Architecture

Section below

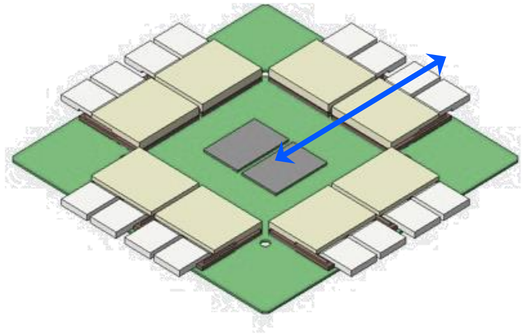

- Multiple dies integration
- 2.5D / 3D advanced packaging

- Thermal management
- Package warpage

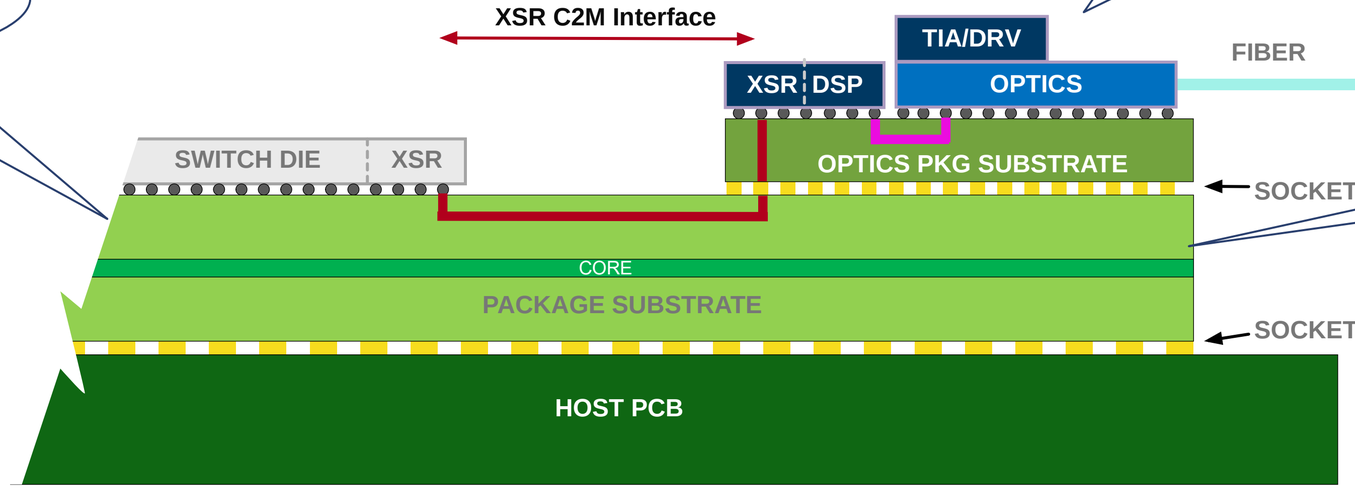

- Largest package
- Substrate size

Industry Initiatives

- CPO Demonstrations / Concepts

Cisco / Luxtera Mechanical Concept

OCP Summit 2018

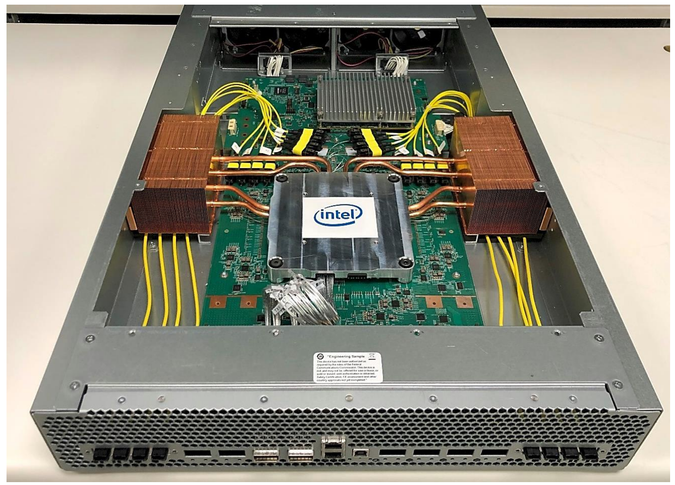

Intel CPO Live Demo

OFC 2020

Rockley Photonics Mechanical Concept

OFC 2020

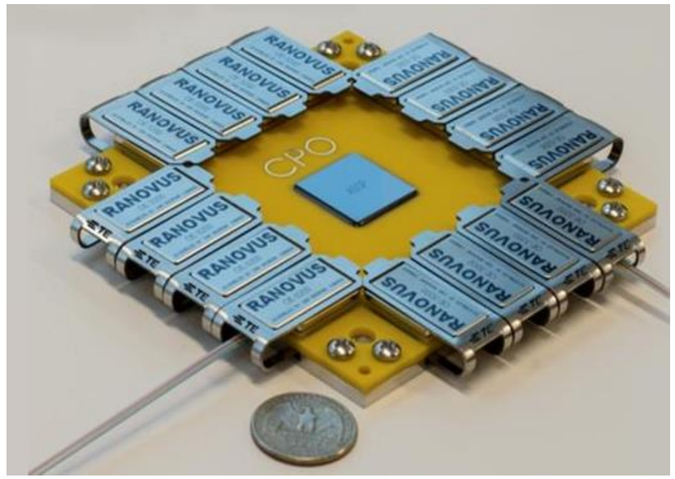

Ranovus Mechanical Concept

OFC 2020

CHIPLETS

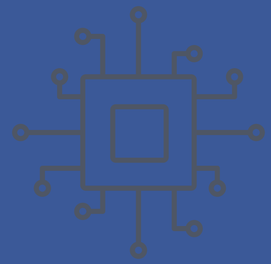

Technology Innovation - Die Disaggregation

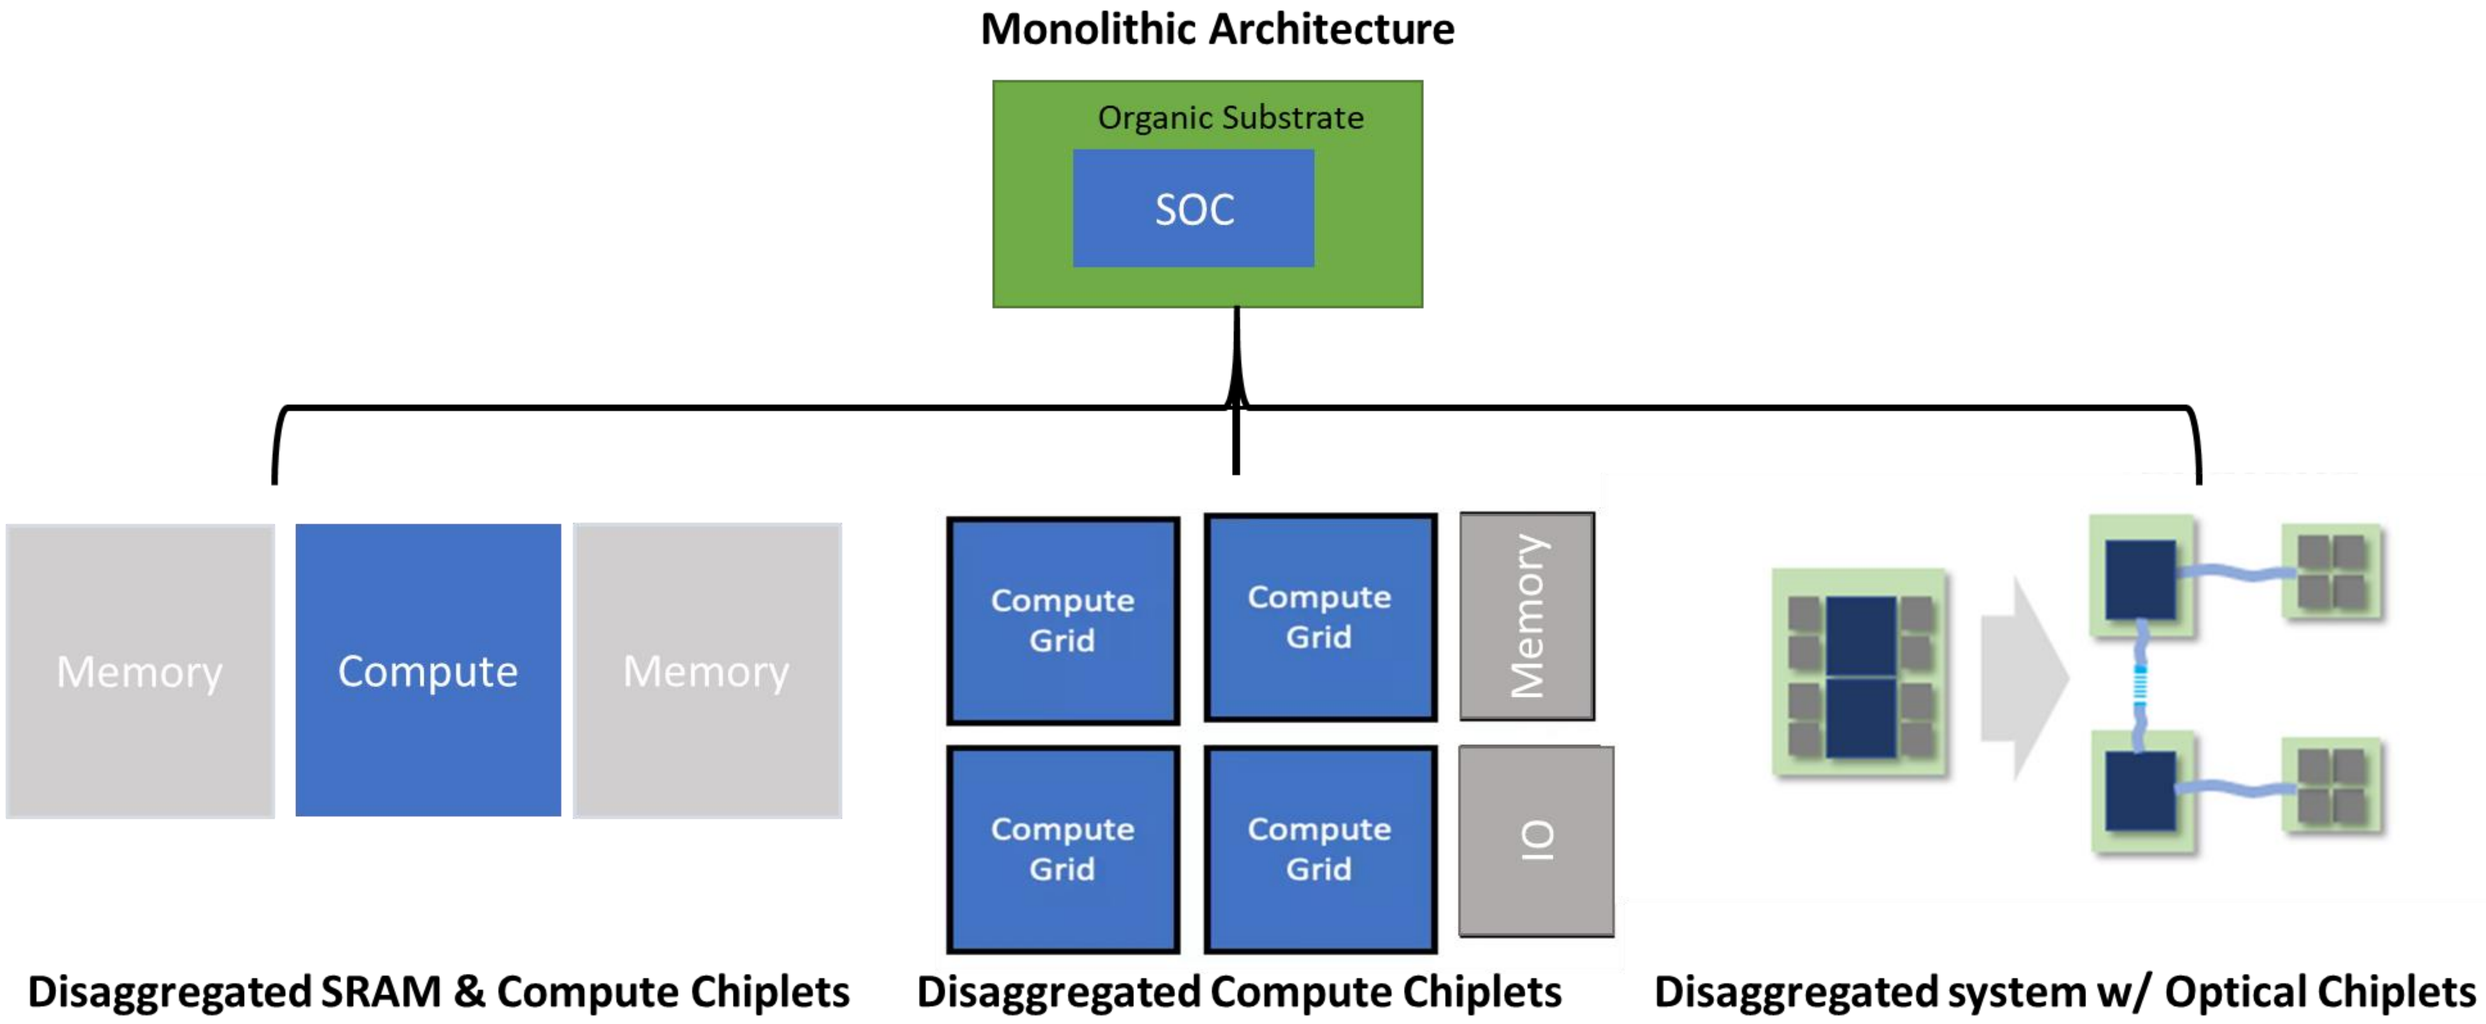

Open Domain Specific Architecture (ODSA)

- ODSA seeks to develop Chiplet an open eco-system marketplace
- Focus on three (3) use cases:
  - IO Disaggregation
  - Core Disaggregation
  - System Integration

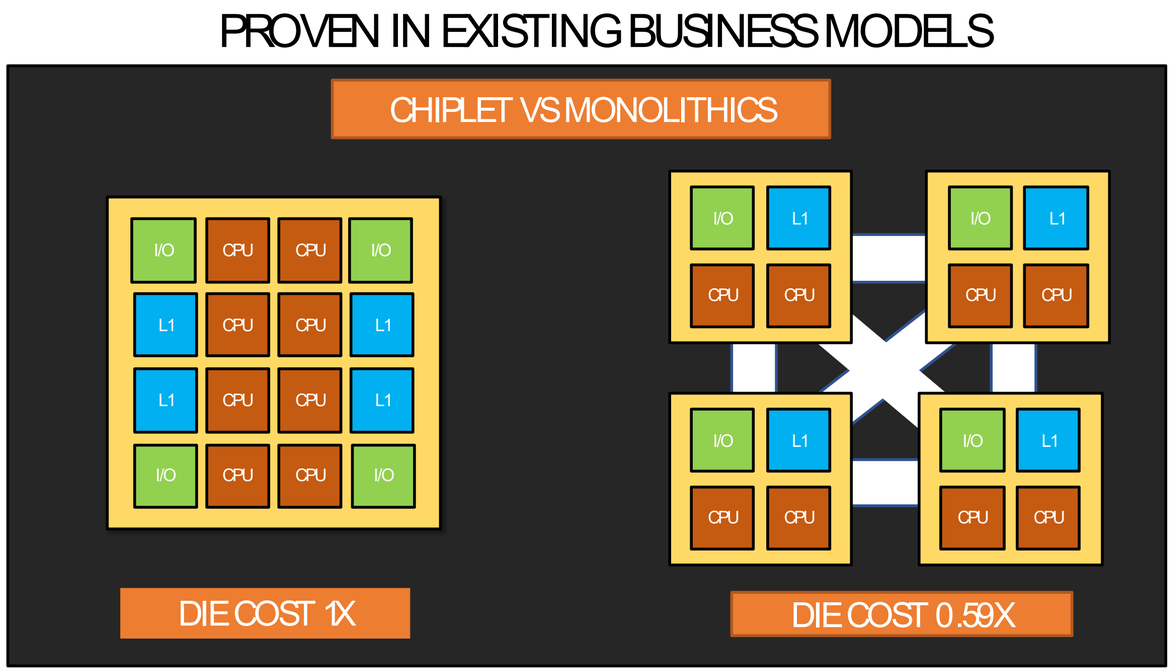

Chiplets Business Challenges

Cost Reduction

- Re-use existing silicon
- Reduce R&D cost
- Fast product SKU development

Standardization

- D2D interface supporting variety of packaging options
- Pin templates for multi-vendor integration

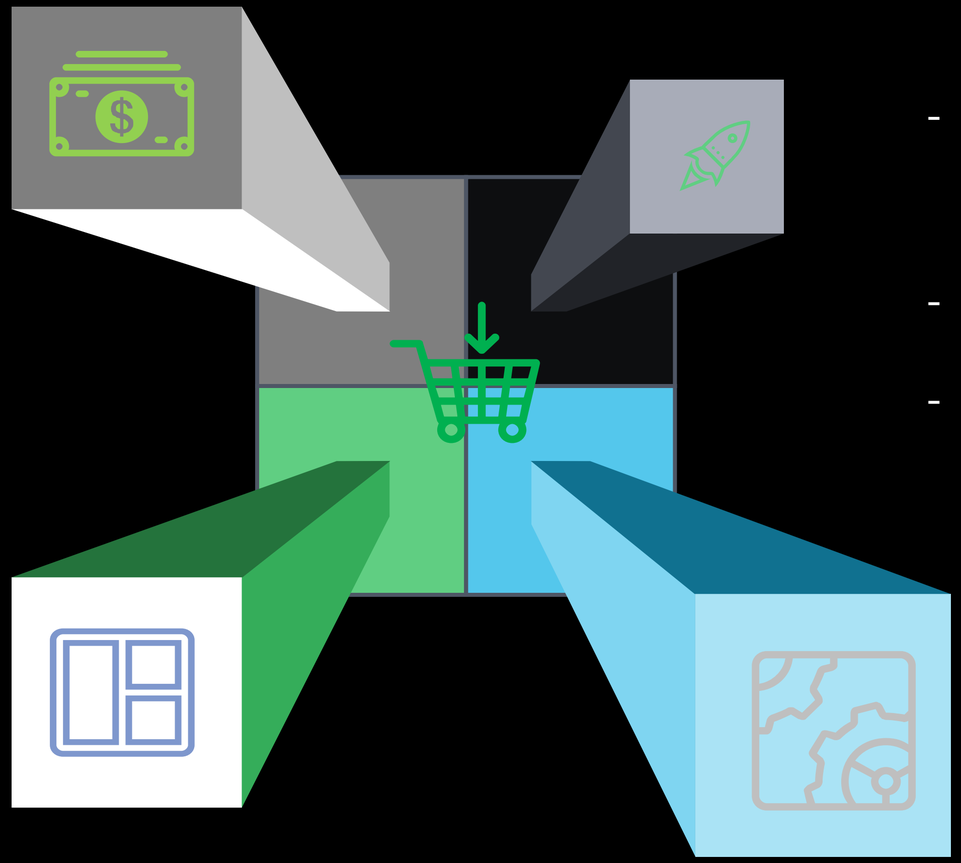

Fast Innovation

- Ability to innovate at different process nodes – no soup to nuts in 3nm
- Impact operation costs via reusability
- Adding 3rd party IP into existing IP library

Operations

- Agreed upon KGD tests to allow foundries to ship to OSATs for assembly and integration
- Common security protocols across supply chain

Supply chain across multiple vendors that allows more players in the market


IF YOU WANT TO GO FAST, GO ALONE
IF YOU WANT TO GO FURTHER, GO TOGETHER
AFRICAN PROVERB

Meta

Thank You!