Intel® Intelligent Fabric Vision and P4 Technology

P4 Workshop Taiwan, Dec 21, 2021

Todd Koelling, Director of Marketing and Technology Planning, Barefoot Switch Division (BXD)

AGENDA

- Vison: Intel® Intelligent Fabric for the Data Center of the Future
- Key Technologies
  - P4 Programming Language
  - In-band Network Telemetry
  - Congestion Control with Remote Priority Flow Control (Remote PFC)
- Products and Applications
- Call to Action

Data Center of the Future: Intel® Intelligent Fabric

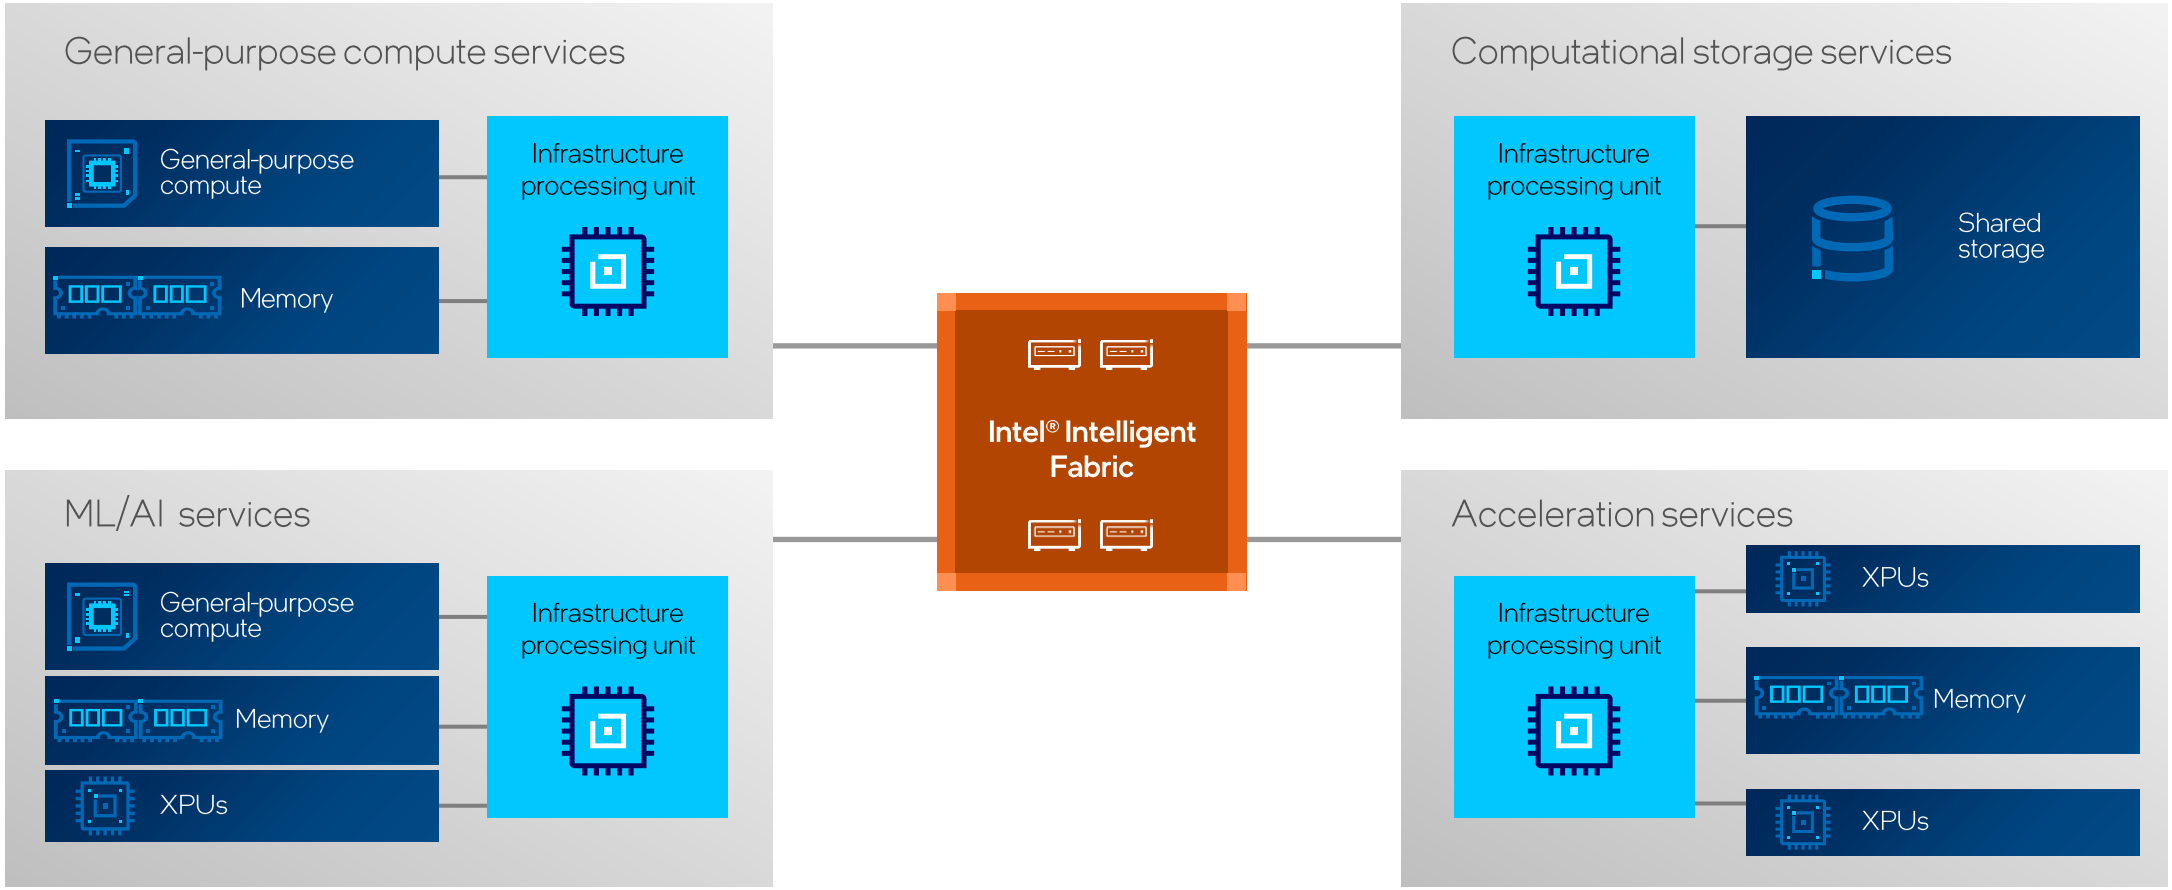

Networking Challenges

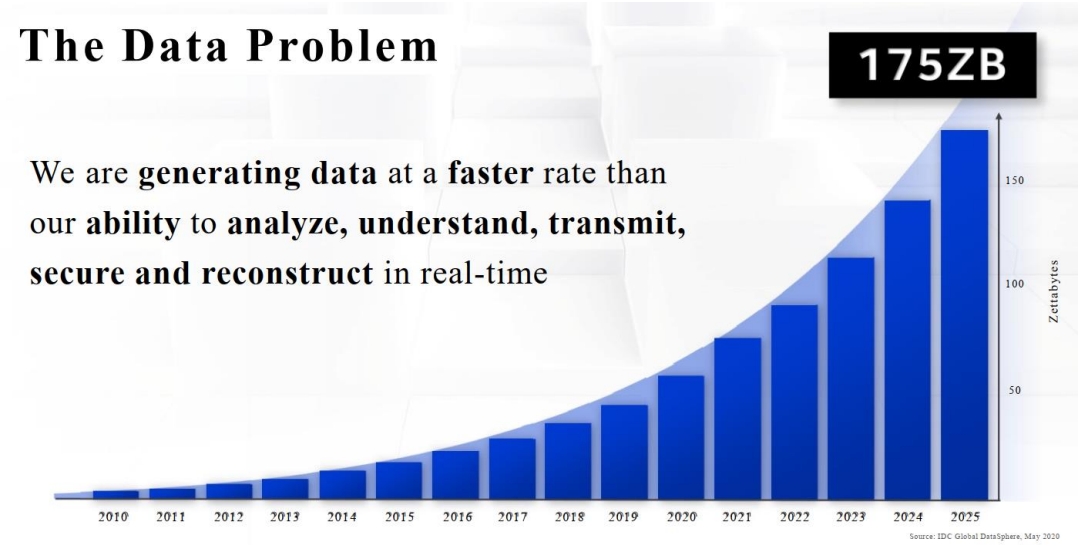

- Move to Cloud-native architecture with container-based processing, orchestration and automation
- AI and changing workloads spawn need for growing network optimizations
- Distributed, scale-out world is changing architecture
- Operational challenges of root causing network slowdowns
- End-to-end Security
- Increasing network CAPEX and OPEX investment for service providers

Network needs to get smarter while increasing bandwidth

VISION: Intel® Intelligent Fabric

End-to-End Co-Optimizations

Photonics Integration

OPTICAL MODULES

High-bandwidth connectivity at 100G, 400G and beyond

ETHERNET SWITCH

P4-programmable scale-out fabric with uncompromising performance

ETHERNET IPUs and NETWORK ADAPTERS

Programmable infrastructure acceleration for demanding data movement

CPUs & xPUs

Fabric-enabled endpoints aligned to accelerators & software pipelines

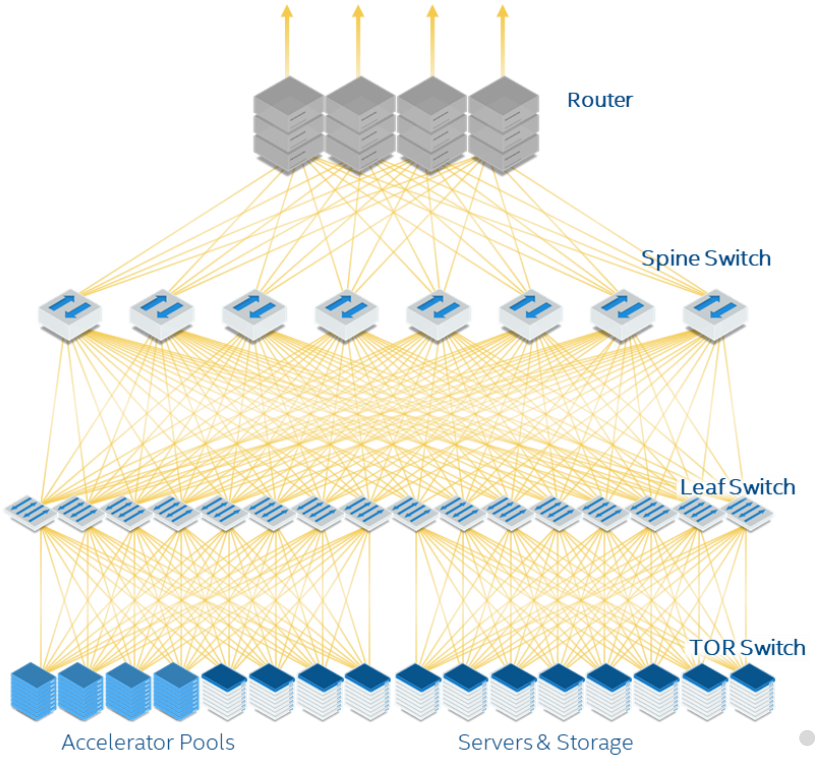

Intel Portfolio

- Ease of Use
- Massive Bandwidth
- AI-driven Self-monitoring / Self-analyzing / Self-healing
- Enhanced Security
- End-to-end Optimization w/Fabric Telemetry
- Improved Density / Power / Cost

Industry Standards

Intel® Intelligent Fabric Key Benefit Vectors

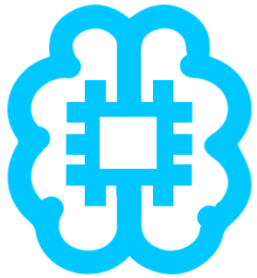

INTELLIGENCE

- Fully Customizable P4-Programmable Pipeline
- Intelligent Packet Processing for Accelerating AI/ML Workloads
- Expandable Table and Buffer Sizes with Intel® FPGAs
- Enhanced Security with Intel® Software Guard Extensions (Intel® SGX) and Intel® Trust Domain Extensions (Intel® TDX)

PERFORMANCE

- 6.4/12.8/25.6 Tbps Total Throughput
- 112G/56G SerDes for high speed and easy migration
- High speed Intel® Silicon Photonics
- Power-optimized Hyperscaler Use Cases for Intel® Tofino™ Intelligent Fabric Processors

VISIBILITY & CONTROL

- Enhanced Congestion Control
- Identify delays or hotspots with real-time In-band Network Telemetry (INT)
- Analyze packet flows with Deep Insight Network Analytics Software
- Remedy using AI or Deep Insight reports
- Traffic monitoring and steering for enhanced security and reliable transport
- Increase INT data available with Intel® IPUs and Ethernet Network Adapters

At the Heart of Intel® Intelligent Fabric
Intel® Tofino™ 3 Intelligent Fabric Processor

New! Announced at Intel Innovation in Oct '21

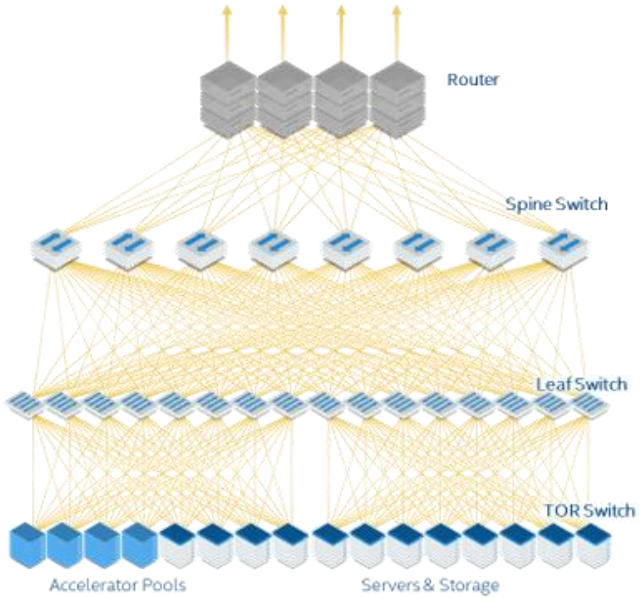

Optimized for:
Cloud and Edge Data Centers
HPC
Comms moving to Cloud technologies

Challenge: Explosion in 5G and IoT data, distributed workloads; requires AI and cloud-to-edge visibility

Solution:
Intel® Tofino™ 3 Intelligent Fabric Processor (IFP)

- Intelligence
- Performance – up to 25.6 Tbps
- Visibility & Control

Switch to Intelligence with Intel Tofino 3 IFP

End-to-End Software Architecture

Programming Model Across Servers, IPUs, FPGAs & Switches

New at OCP Summit:
SAI Packet Test
Framework Contributions

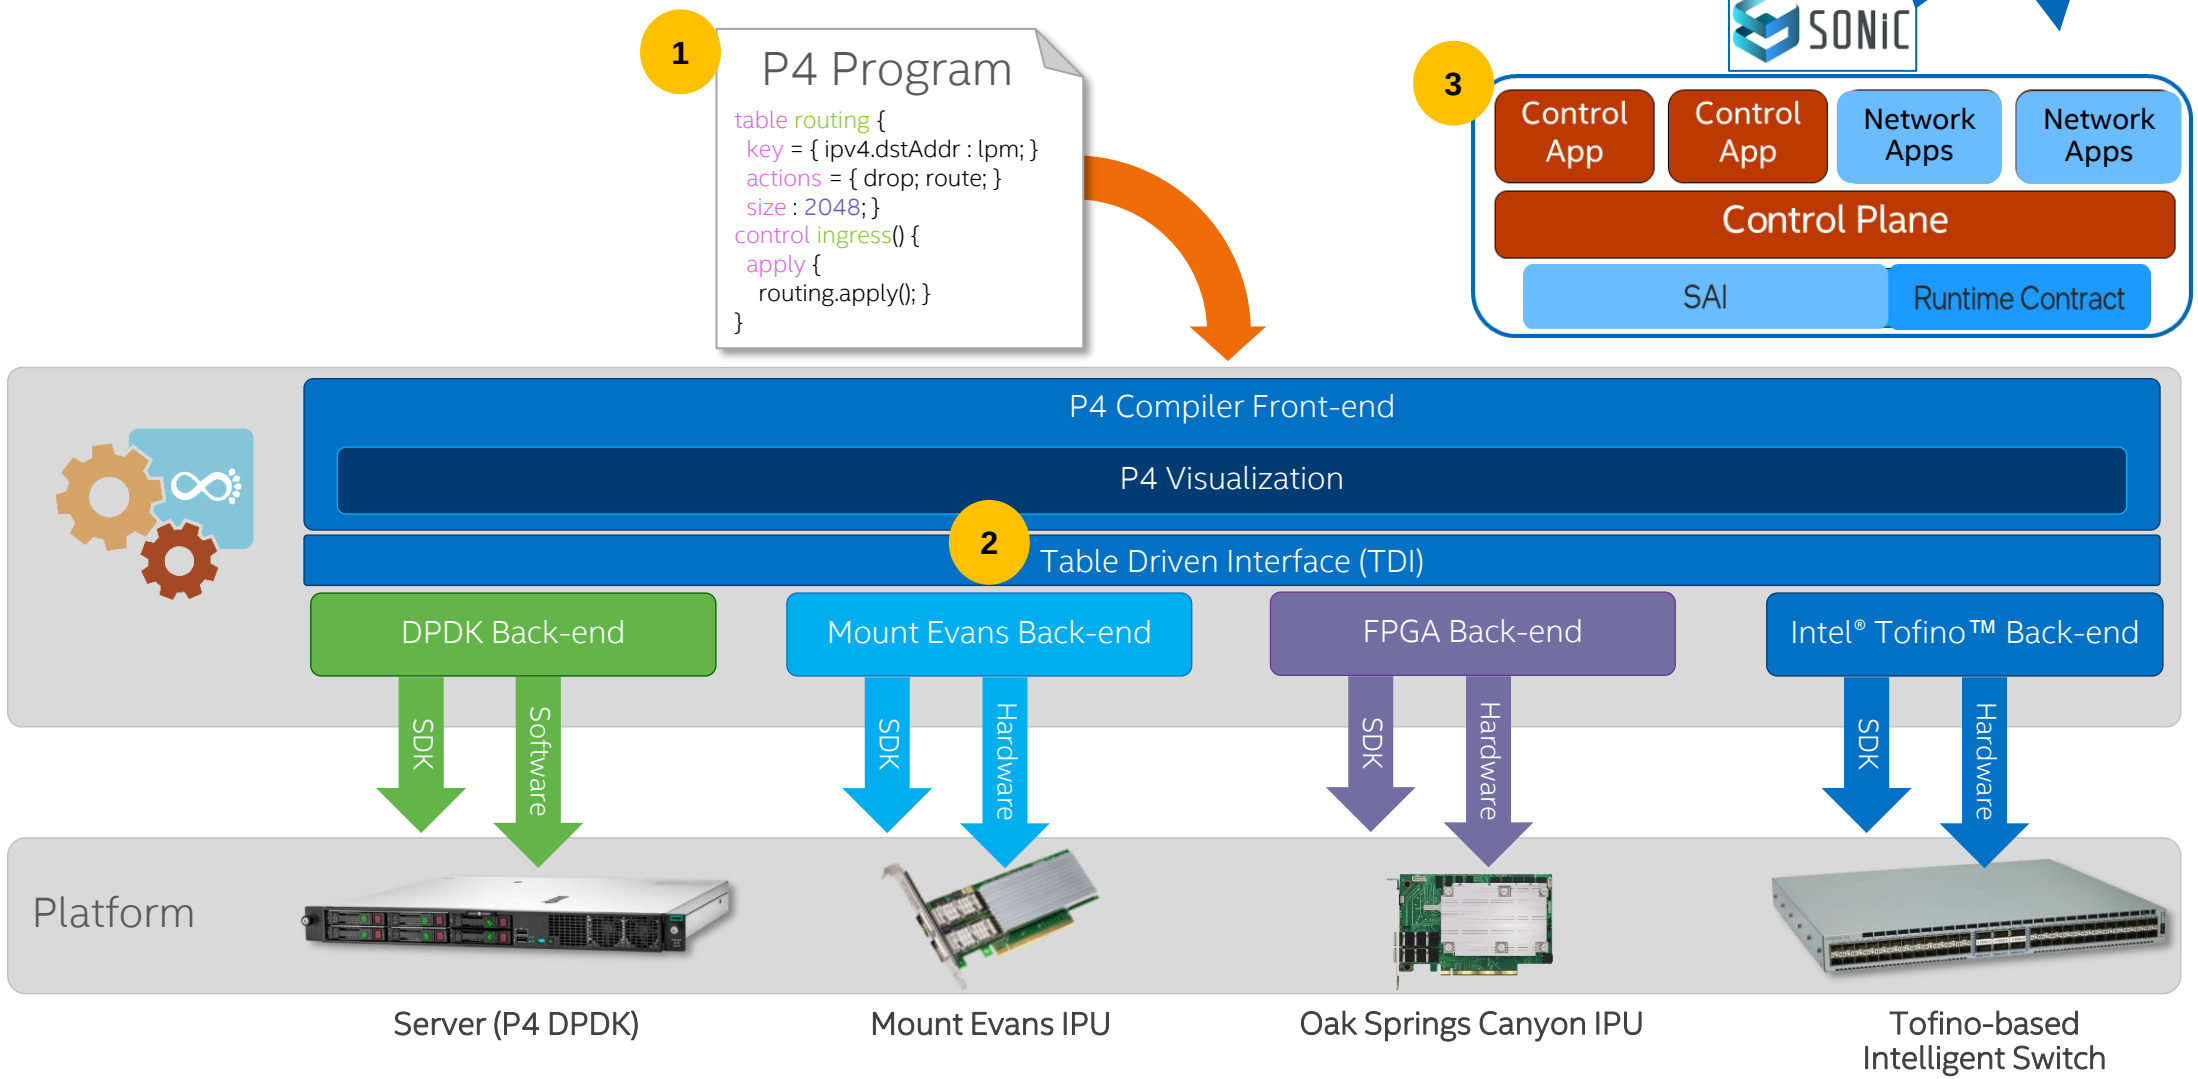

Key Technologies

P4: Built with community support

Intel is an active member in the P4 community, providing knowledge and expertise to end users from all walks of life.

Widespread adoption:

Growing community of 4000+ developers, 100+ member organizations

- Telcos
- Cloud operators
- Researchers
- OEMs
- System integrators

Visit p4.org for more information

Follow each packet's journey from beginning to end

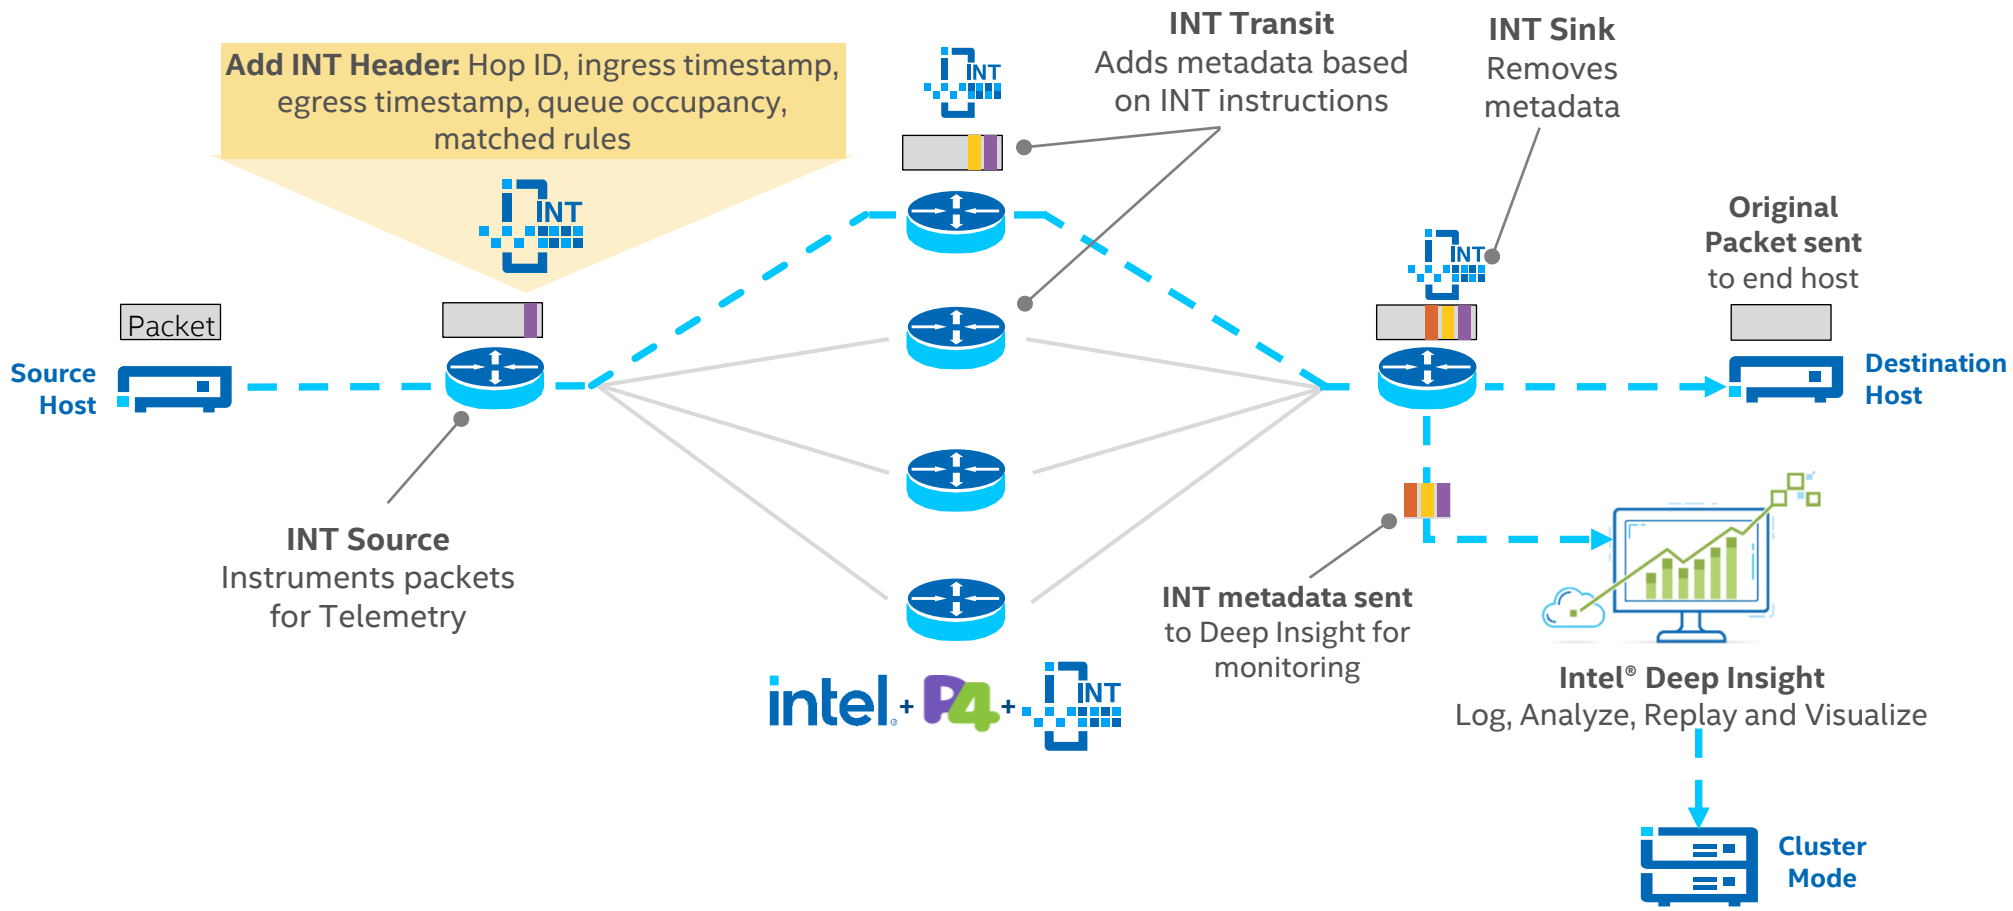

Data-plane Telemetry
In-band network telemetry (INT.P4)
Intelligent deduplication and triggers
Hardware primitives
Line rate monitoring

Answer for **every** packet... 1 Why is it here? 2 How did it get here? 3 How long was it delayed? 4 Why was it delayed?

Remote Priority Flow Control (Remote PFC) Edge-to-Edge View

What is Remote PFC?

- Edge-to-Edge signaling of congestion
- Flow control that instantly 'flattens the curve'
- Signaling + 'source' flow ctrl all in sub-RTT

Remote PFC d

- ~~aim 100% le~~
- ~~e2e congest~~
- ~~Pause Agg/C~~
- ~~Need greenfield de~~ → ToR-only upgrade

Example of Benefits of Intelligence
Next Session:
Improved Congestion Control with P4-based Remote Priority Flow Control (Remote PFC)
By Jeremias Blendin

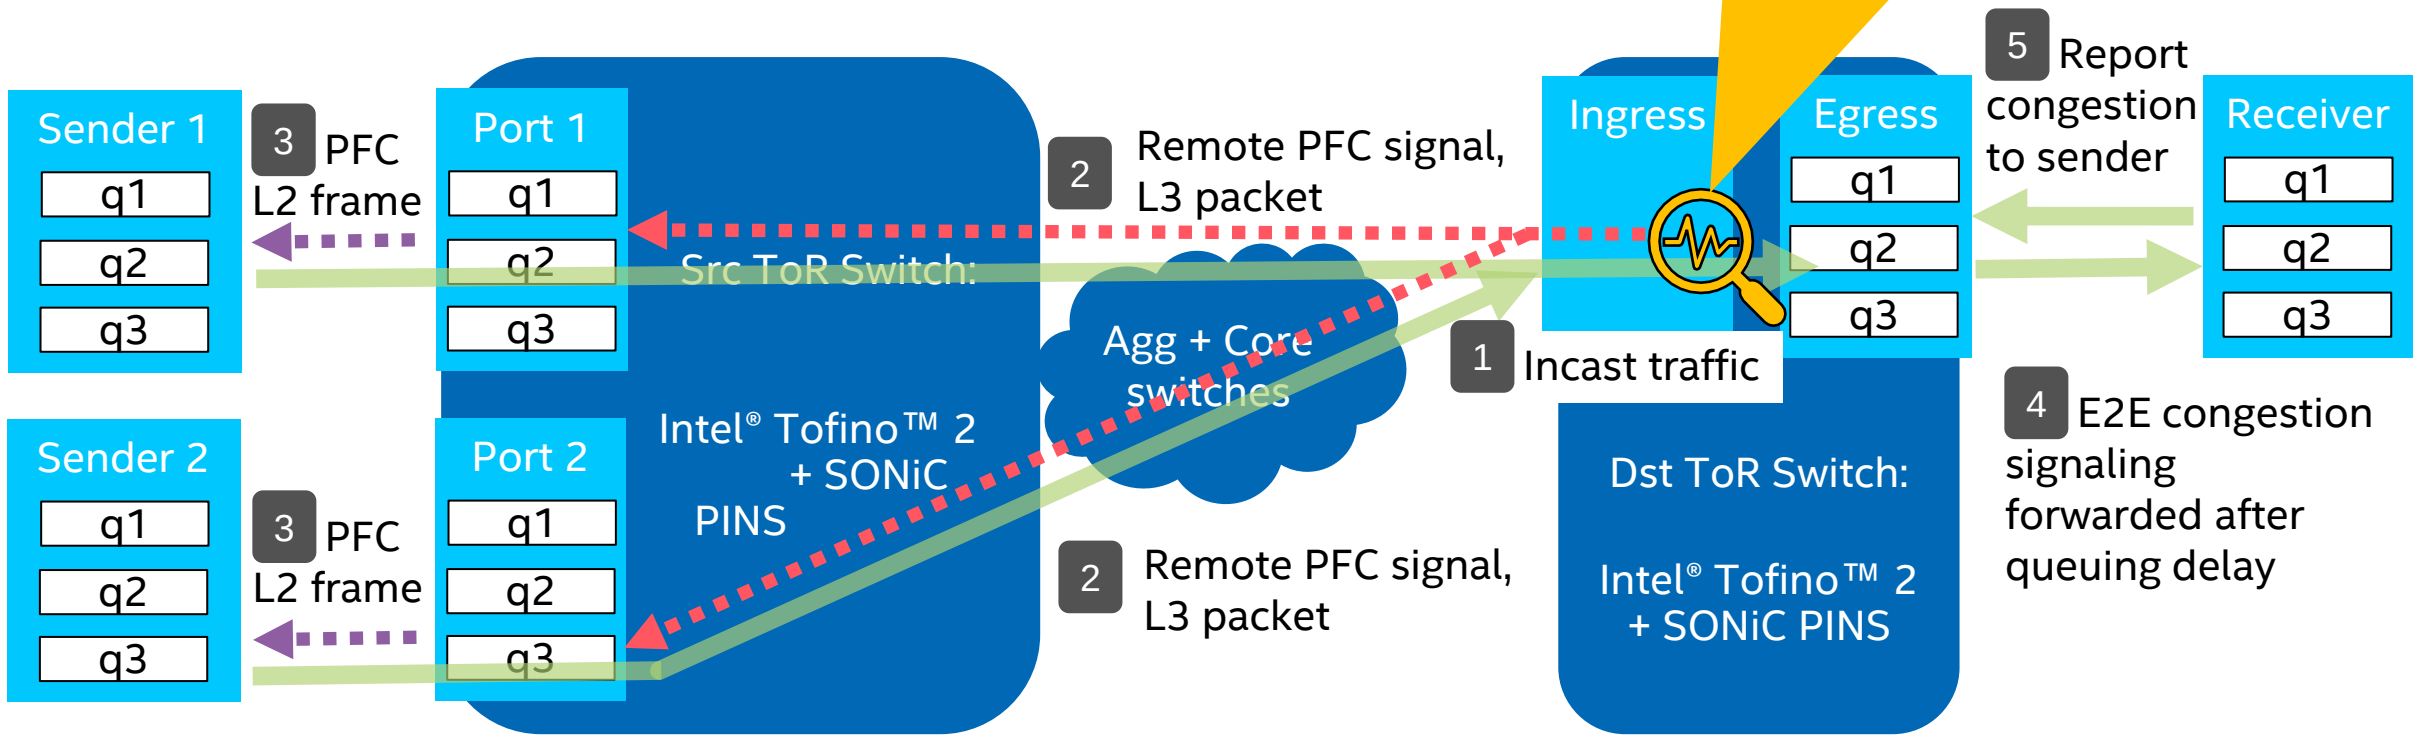

Products and Applications

Intel® Tofino™ Intelligent Fabric Processors

Tofino (16nm)

- 1.8 to 6.4 Tbps
- 25G SerDes

Tofino 2 (7nm)

- 3.2 to 12.8 Tbps
- Modular Chip Design
- 56G SerDes

Tofino 3 (7nm)

- 12.8 to 25.6 Tbps
- Modular Chip Design
- 112G/56G SerDes

Status: Production now — Production now — Future

Intelligence

- P4 Programmable
- AI/ML Acceleration
- Highly-Secure

Performance

- Up to 25.6 Tbps throughput
- 112G/56G SerDes
- Power-Optimized Use-Cases

Visibility and Control

- Edge-to-cloud real-time telemetry
- Enhanced congestion control
- Self-healing network capabilities

Explore and innovate with new P4 Applications

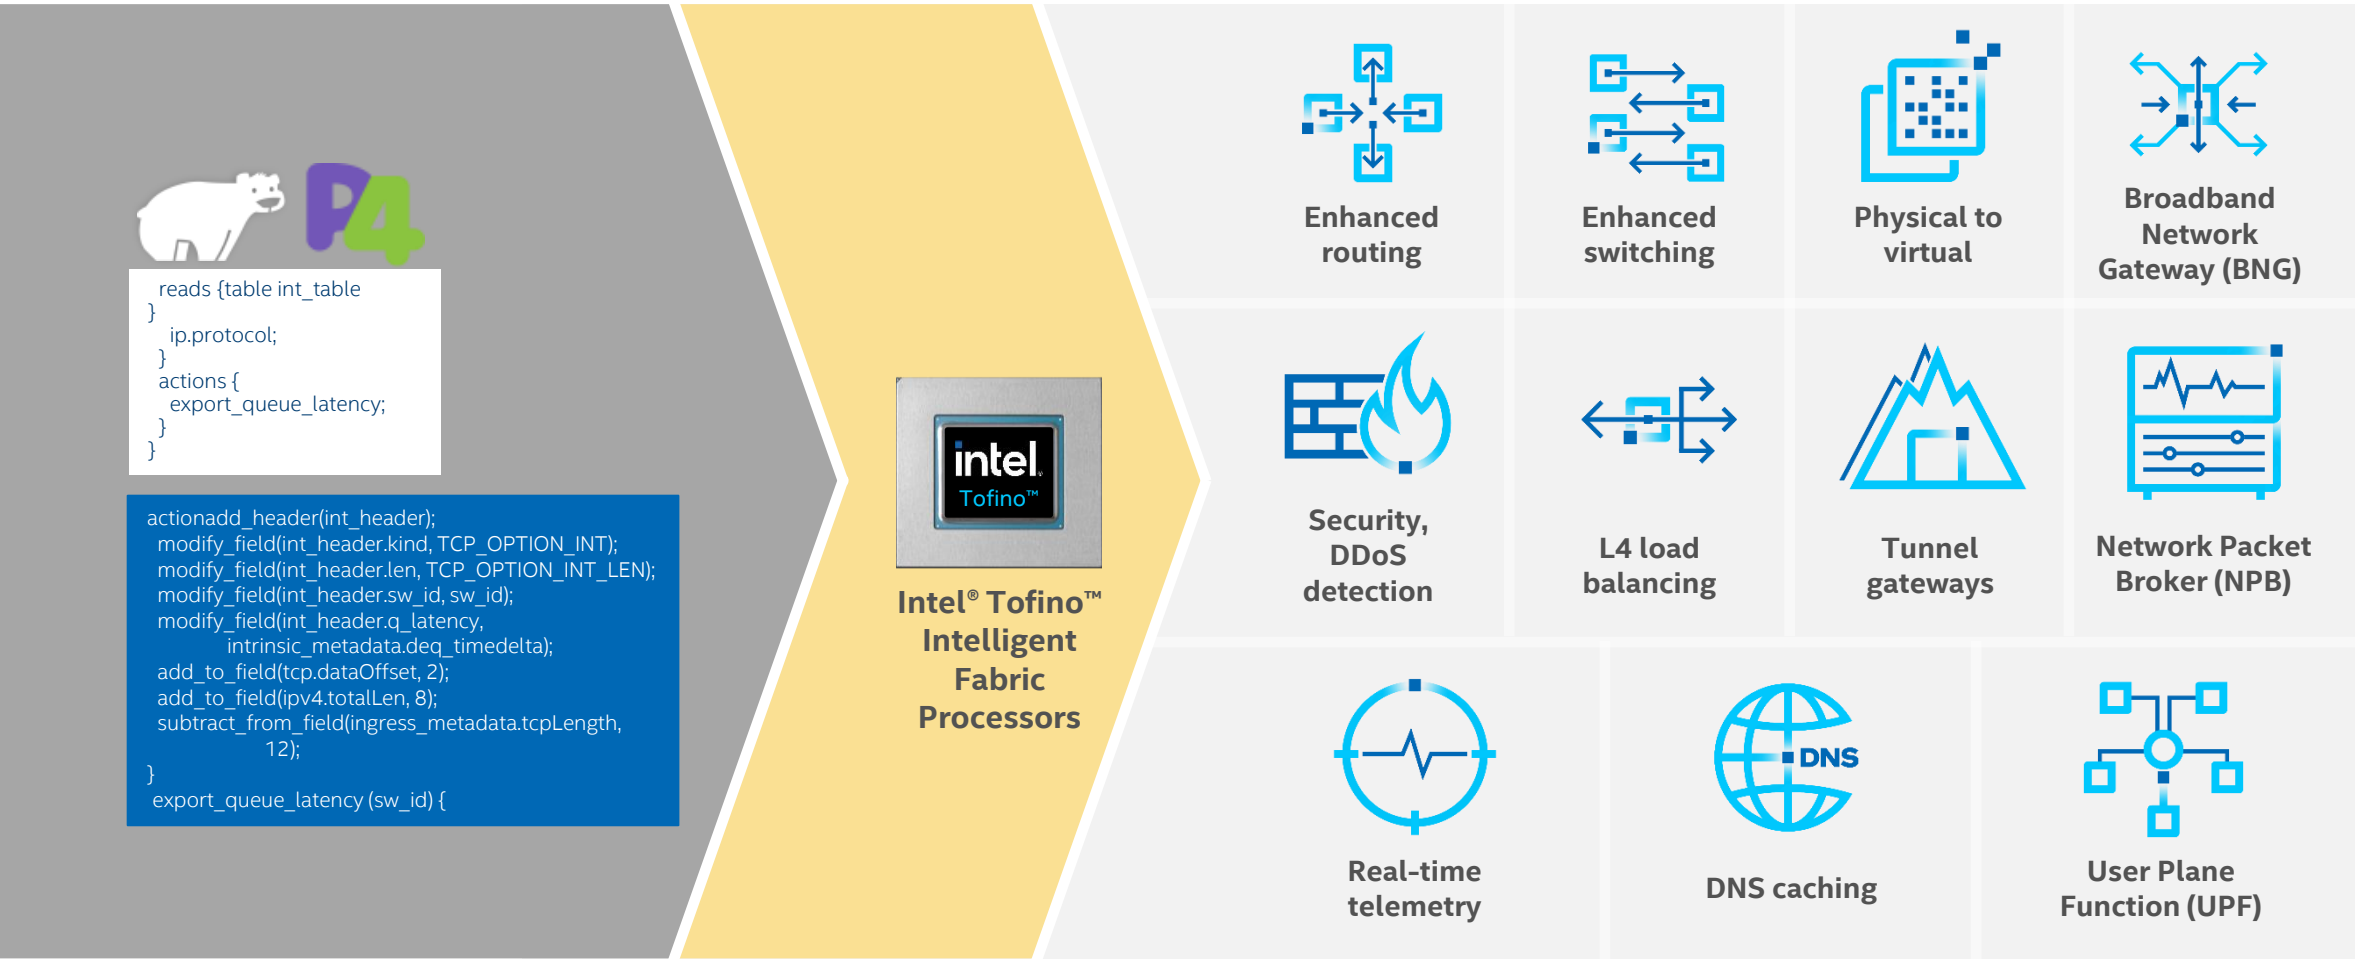

University Grant Opportunity:

Intel Connectivity Research Program (ICRP) Fast Forward Initiative

- Goal to accelerate the best network programmability research of tomorrow
- We are looking for the most ground-breaking, industry-advancing experiments involving the P4 programming language
- Grants will assist ICRP members in acquiring hardware for their projects
- Participating vendors: APS Networks, Edgecore Networks, Netberg, UFI
- Special consideration given to projects which advance the state of the art in the following areas: Intelligence, Performance, Visibility & Control
- Applications due December 31, 2021 with selection by January 14, 2022

Link to application (https://bit.ly/3F3qAAC)

Summary/Call to Action – Ways to Get Involved

- Join the P4.org community
- Contribute to the OCP SONiC project
- Learn at the Intel® Connectivity Academy and Contribute to the Intel® Connectivity Research Program

- Visit intel.com/IFP for more info on Intel® Intelligent Fabric Processors (Intel® IFPs)

Notices and Disclaimers

- Performance varies by use, configuration and other factors. Learn more at www.Intel.com/PerformanceIndex.
- Performance results are based on testing as of dates shown in configurations and may not reflect all publicly available updates. See backup for configuration details. No product or component can be absolutely secure.
- Your costs and results may vary.
- Intel technologies may require enabled hardware, software or service activation.
-


intel