



October 20th, 2020

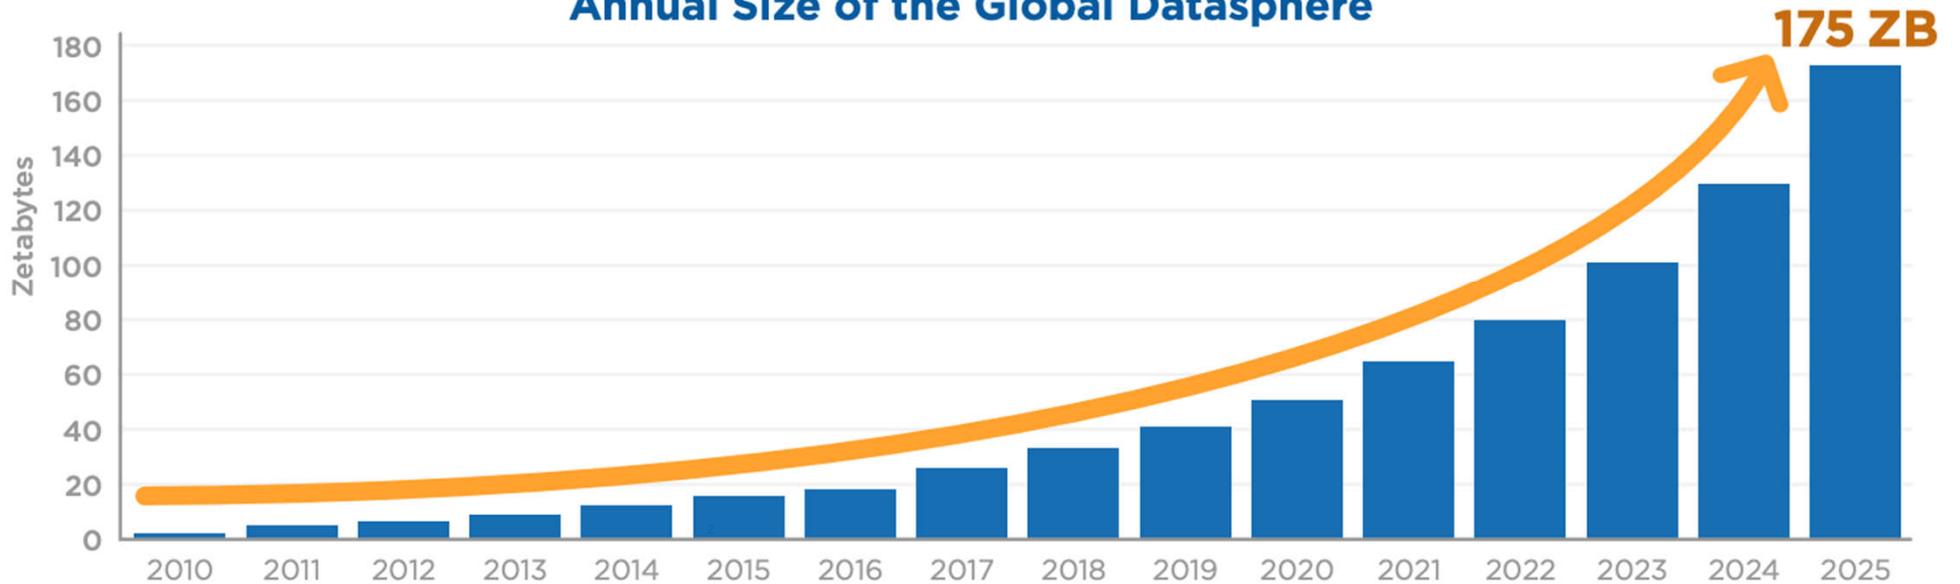
IEEE CVT Briefing

John DeMay
john.demay@intel.com



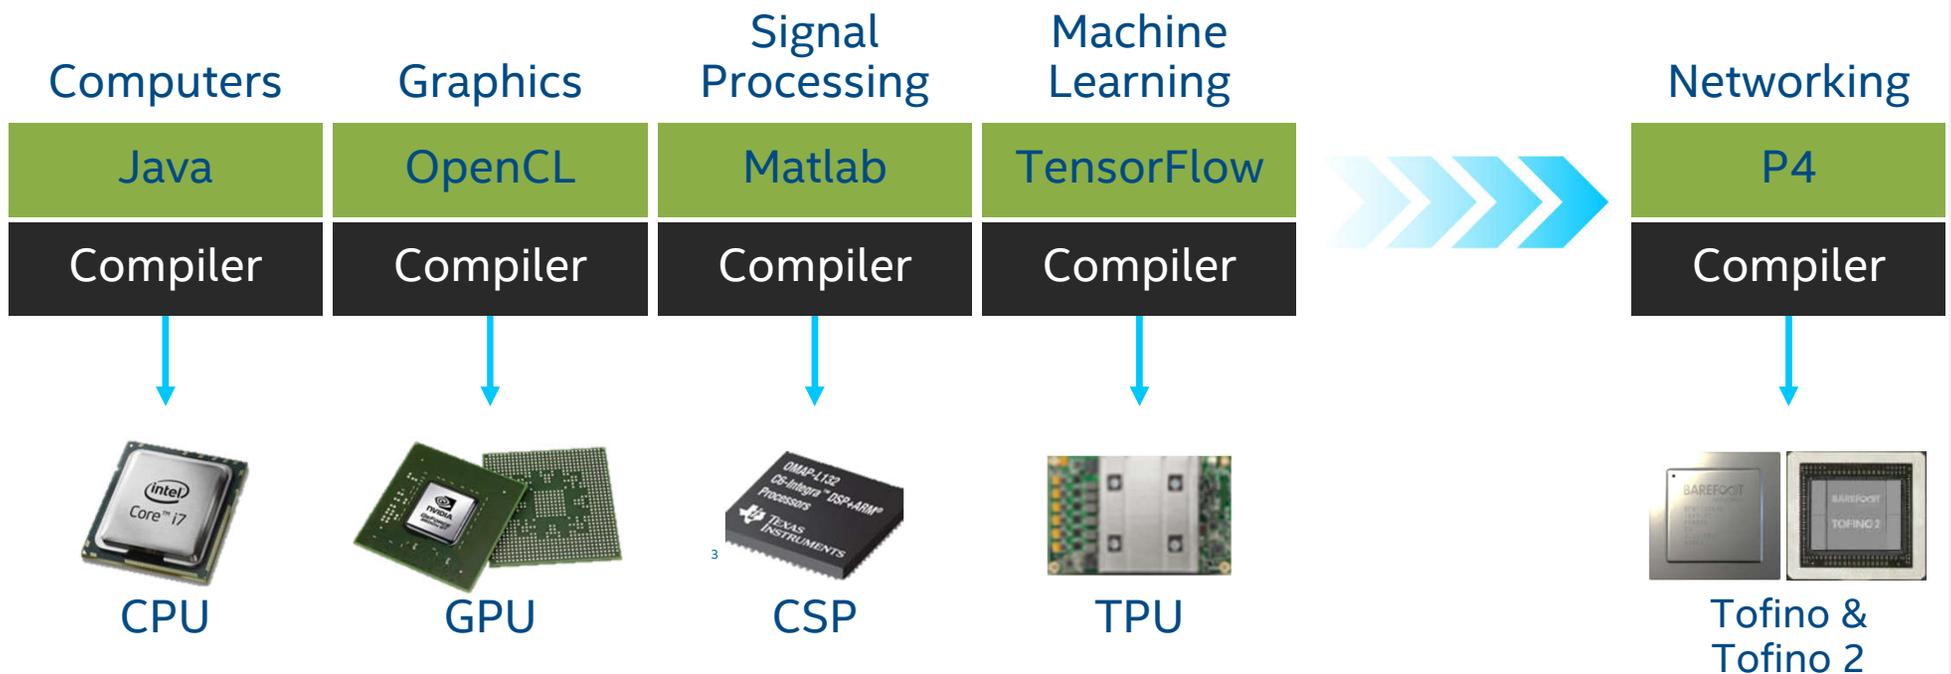
Exponential growth of data

Annual Size of the Global Datasphere

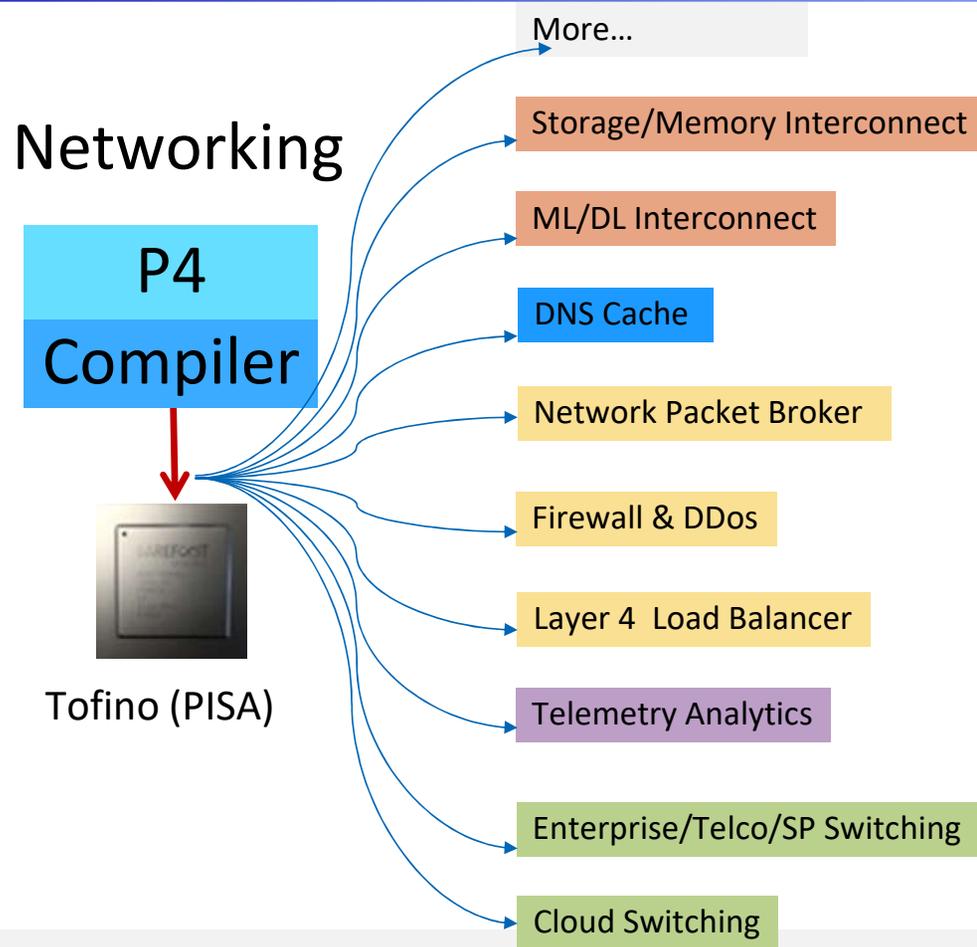


Source: Data Age 2025, sponsored by Seagate with data from IDC Global DataSphere, Nov 2018

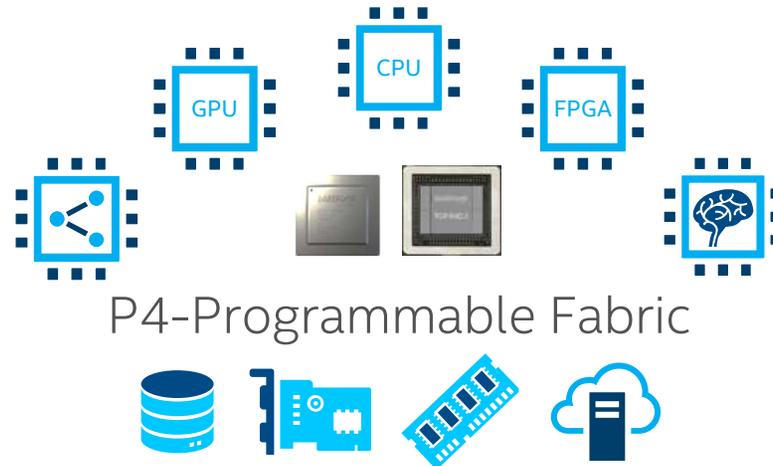
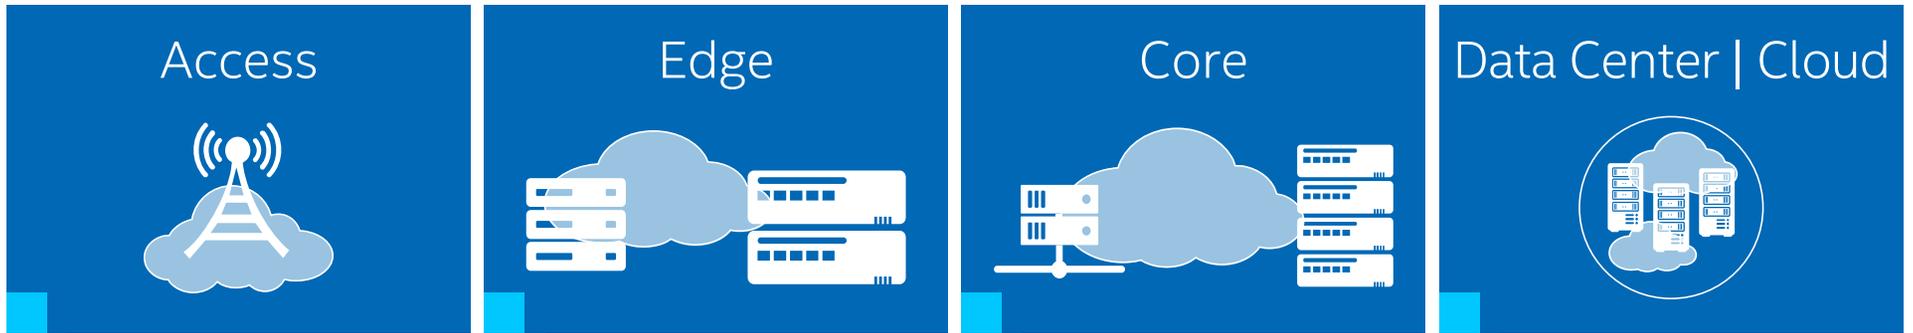
General Industry Trend: Rise of the Domain-Specific Architectures (DSAs)



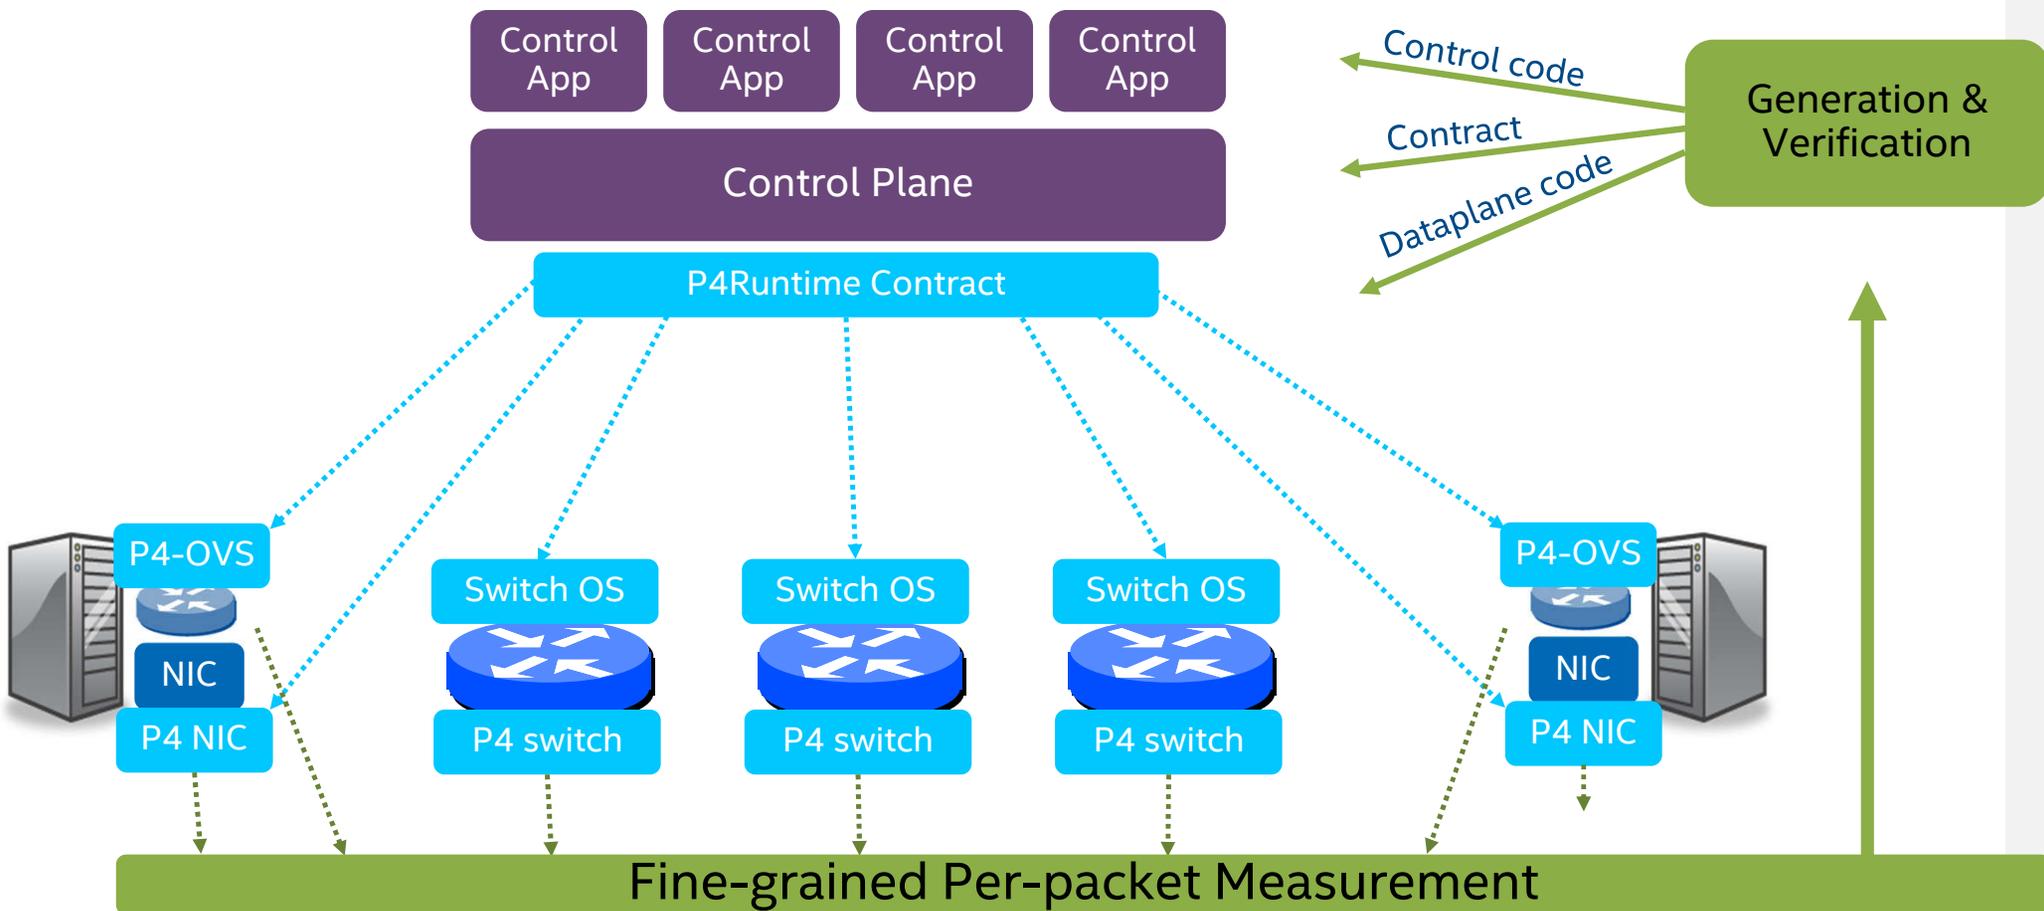
New Use Cases Unleashed Through SW Defined HW



P4-programmable Connectivity Across the Network



Vision: Network as a Programmable Platform



Barefoot Technology & Products



1.2 <-> 6.4 Tb/s

4.0 <-> 12.8 Tb/s

Tofino Tofino 2

Family of P4-programmable Ethernet switch ASICs with P4 compiler and SDK



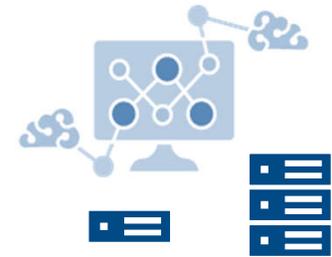
P4 Studio

State-of-the-art compiler for Tofino and Tofino 2 and Development Environment



P4 Insight

P4 Code Debug Tool and P4 Compiler



Deep Insight

Network performance monitoring and analytics software

Tofino Programmable Switching ASIC Portfolio

Tofino



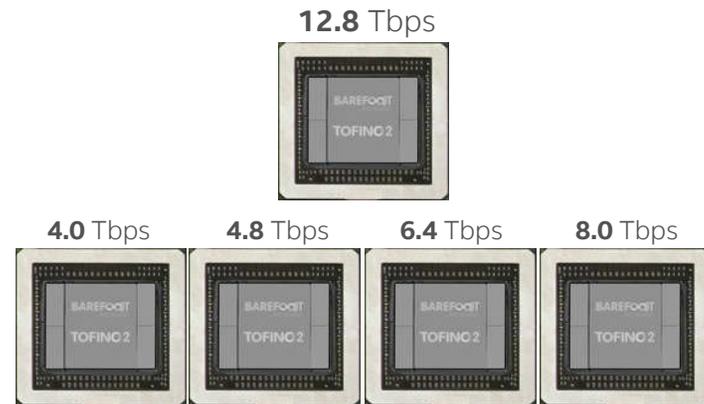
 Programmable

Ethernet Switch ASIC w/ **25Gbps** SerDes

16nm Process Node

Optimized for
100GbE / 50GbE / 40GbE / 25GbE / 10GbE
compute connectivity

Tofino 2



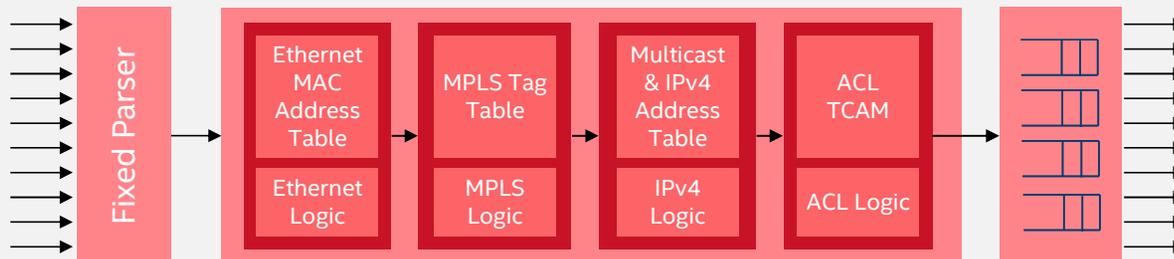
 Programmable

Ethernet Switch ASIC w/ **56Gbps** SerDes

7nm Process Node

Optimized for
400GbE / 200GbE / 100GbE / 50GbE / 25GbE / 10GbE
compute connectivity

Why Programmable Pipeline?

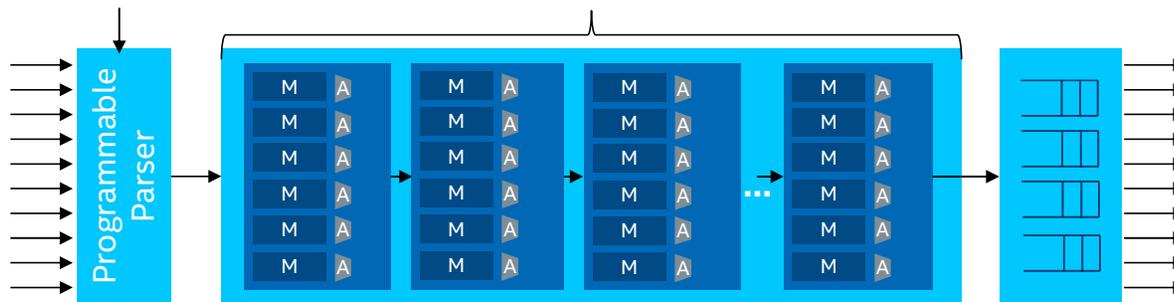


Features and table-sizes are **hard-coded**, NOT optimized and often **unused**

FIXED-FUNCTION

- Behavior cannot be changed
- Poor to little visibility
- New functionality requires hardware upgrade

User specifies headers to be parsed in a P4 program User specifies tables & size needs and packet processing functions in a P4 Program



Pipeline is **fully workload-optimized**

P4-PROGRAMMABLE

- + Optimized to the end-user needs
- + Real-time visibility
- + Software upgradable hardware



P4 Developer Ecosystem



- **Packet Headers**

```
header ethernet_h {
    bit<48>    dstAddr;
    bit<48>    srcAddr;
    bit<16>    etherType;
}
```

- **Match-Action Table**

```
table dmac {
    key = {
        ingress_metadata.bd : exact;
        ethernet.dstAddr    : exact;
    }
    actions = {
        dmac_hit;
        dmac_miss;
        dmac_redirect_to_cpu;
    }
    default_action = dmac_miss;
    size = 131072;
}
```

- **Actions**

```
action dmac_hit(bit<9> egress_port) {
    ig_intr_md_for_tm.ucast_egress_port = egress_port;
    l2_metadata.same_if_check =
        ig_intr_md.ingress_port ^ egress_port;
}
```

P4 Features

- Open Spec, Compiler, Test Frameworks, and more
- Protocol Independent
- Target Independent

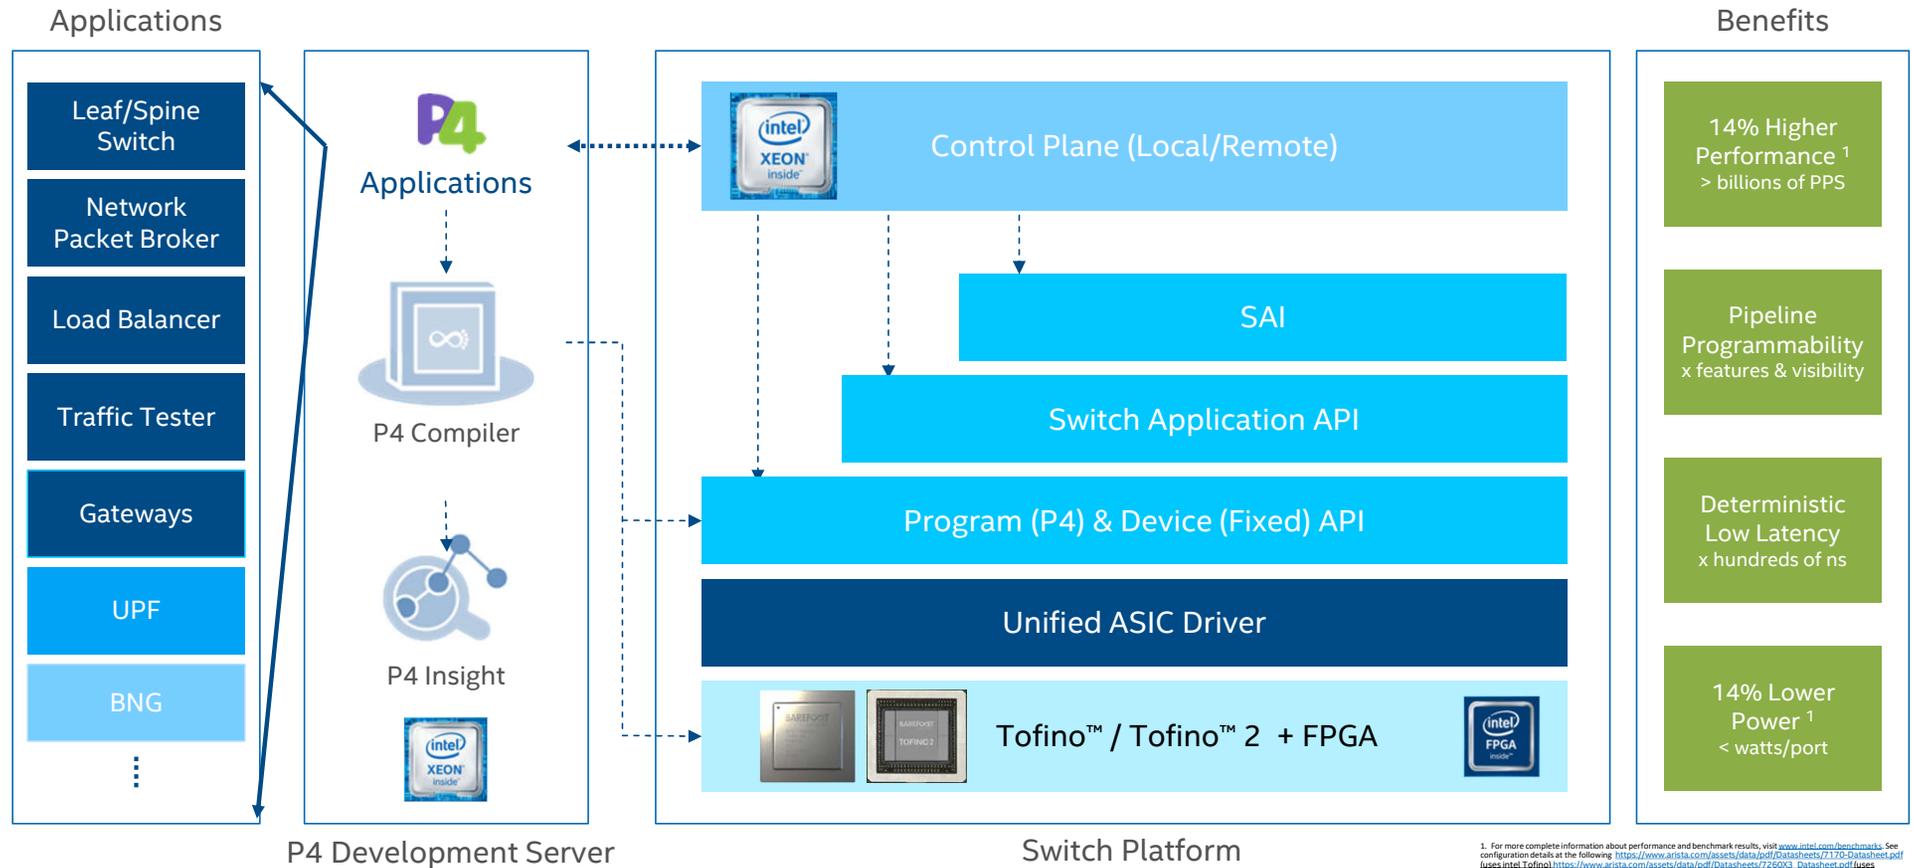
Strong community

- 4000+ developers trained and growing
- 100+ member organizations
- Expanding across the globe

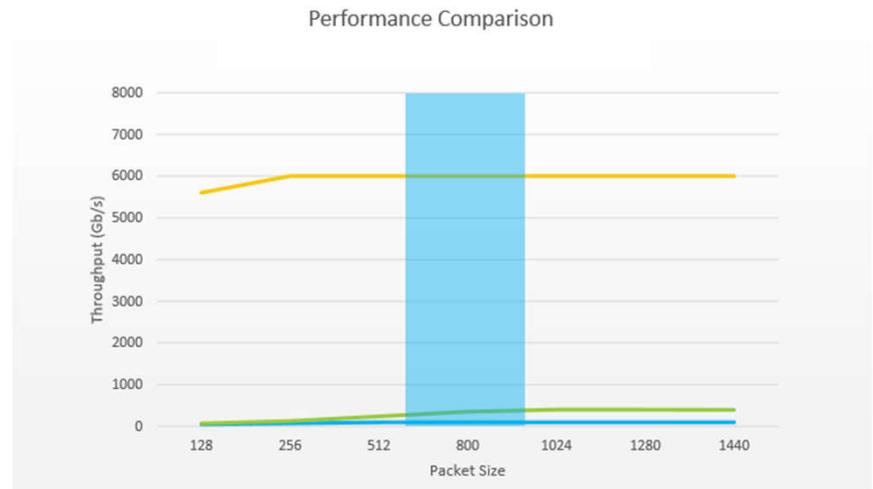
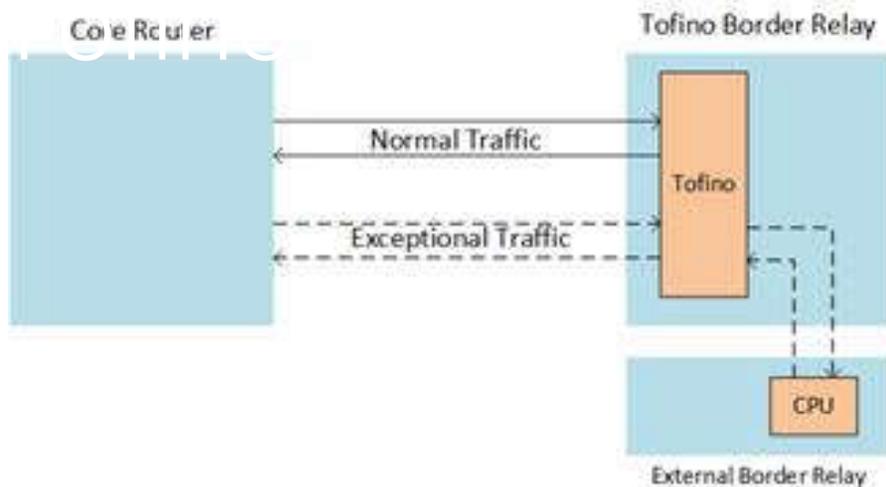
Accelerating adoption

- Expanding adoption by new vendors
- Switches, NICs, FPGA, Software Data Planes

Tofino Applications and Benefits



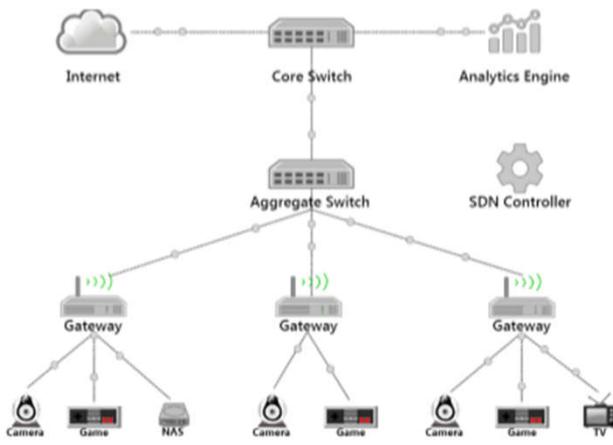
Stateless NAT Implementation



Tofino Border Relay

- Normal TCP and UDP packets are translated in Tofino ~99% of traffic
 - MAP pre-fixes handled in Tofino
 - Tofino performs anti-spoofing checks for v4 & v6
- Exceptional packets are punted to CPU w/o decrementing TTL
- Hybrid approach provide faster deployment & significant cost savings
 - Equivalent solution would require 15ea. appliances with 5KW power in 15RU
 - Tofino solution draws 830W max in 3RU

P4 & INT Implementation For DDoS Mitigation

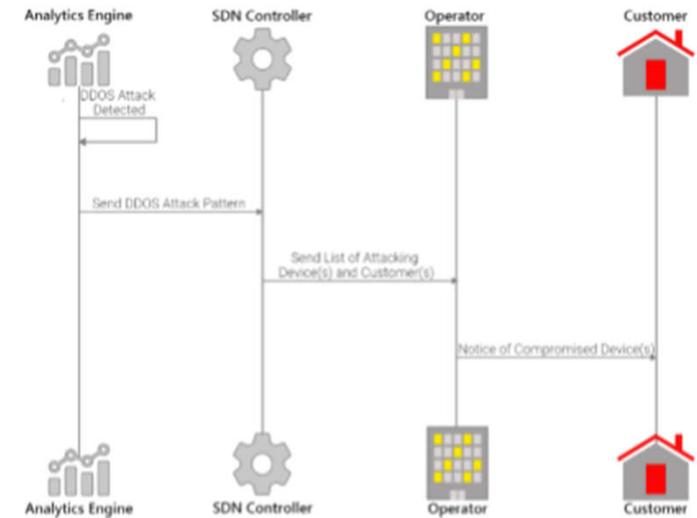
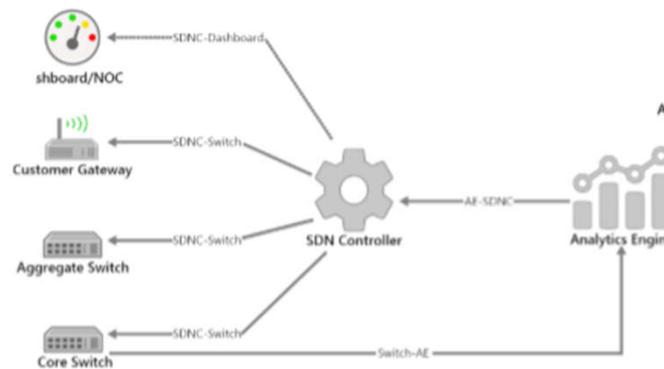


3 customer premises w/ 2 or more devices

- P4 enabled gateways are connected to aggregate switch
- P4 enabled aggregate switch connected to core switch
- P4 enabled core switch connected to the internet and analytics engine for pattern recognition

When malicious patterns are detected:

- SDN controller is notified
- Controller updates P4 enabled devices thru GRPC, Thrift, HTTP or RPC
- P4 devices are instructed how to handle packets based on the pattern signature



When controller receives an alert signature

- Controller alerts the operator
- Gateway attributes enable customer identification
- Operator can contact the customer to notify them which devices are compromised

Keysight UHD100T32

Problem Statement:

- With mass deployment in modern data centers, 100G has matured
- Customers are challenged with density and cost pressures on their test Infrastructure

Solution:

- 1RU, 32x100G, line rate traffic tester based on Tofino
 - Support fan out speeds (50/40/25/10GE)
- Supported by innovative clientless operation via web interface with REST APIs
- Provides extremely cost effective solution



Thank You!



intel®