Stratum Techinar

July 19, 2022 | 9am PDT

Brian O'Connor
Intel
Stratum TST Lead

Maximilian Pudelko
Intel
Stratum TST Member

Outline

- Recent Stratum Enhancements and Releases
- Extending Stratum Support to the Server
- Stratum Roadmap

Stratum Releases

- ✓ 20.06
- ✓ 20.09
- ✓ 20.12
- ✓ 21.03
- ✓ 21.06
- ✓ 21.10
- ✓ 21.12
- ✓ 22.03
- ✓ 22.06 (latest)
- ➢ 22.09 (planned)

- Stratum is 9 releases old
- Quarterly release cadence
- “Live at Head”
- Status by Target
  - Intel® Tofino™ programmable Ethernet switch ASIC – stable, active development
  - Broadcom StrataXGS switch series – stable, some maintenance
  - BMv2 – stable, some maintenance
  - np4intel – PoC quality, no active development
  - TDI (DPDK, IPU) – incubating

Stratum Fixes and Enhancements

Improved support for Intel® Tofino™ programmable Ethernet switch ASIC

- Removed Stratum-bf
  - Was based on unmaintained, inflexible PI node backend
- Stratum-bfrt for Intel Tofino has reached maturity / production hardened
  - Used by SD-Fabric, Aether, and Pronto (as well as some other vendors/users)
- Support for Intel P4Studio 9.7.x, 9.8.0 and 9.9.0
- Experimental support for table entry idle timeout notifications on Intel Tofino

Stratum Fixes and Enhancements

P4Runtime Enhancements

- Stratum now uses P4Runtime canonical byte strings by default
  - The old behavior can be enabled with the `enable_bfrt_legacy_bytestring_responses` flag
- Spec compliant P4Runtime mastership arbitration (v1.4.0-rc.1)
- P4Runtime metadata translation is now experimentally supported for stratum_bfrt and is enabled with the `experimental_enable_p4runtime_translation` flag
- New OpenConfig path (`/interfaces/interface[name=*]/state/id`) that can be used for P4Runtime port translation
- Minor fixes: P4 MeterConfig resets, port ID values in the OpenConfig tree

Stratum's Use in SD-Fabric

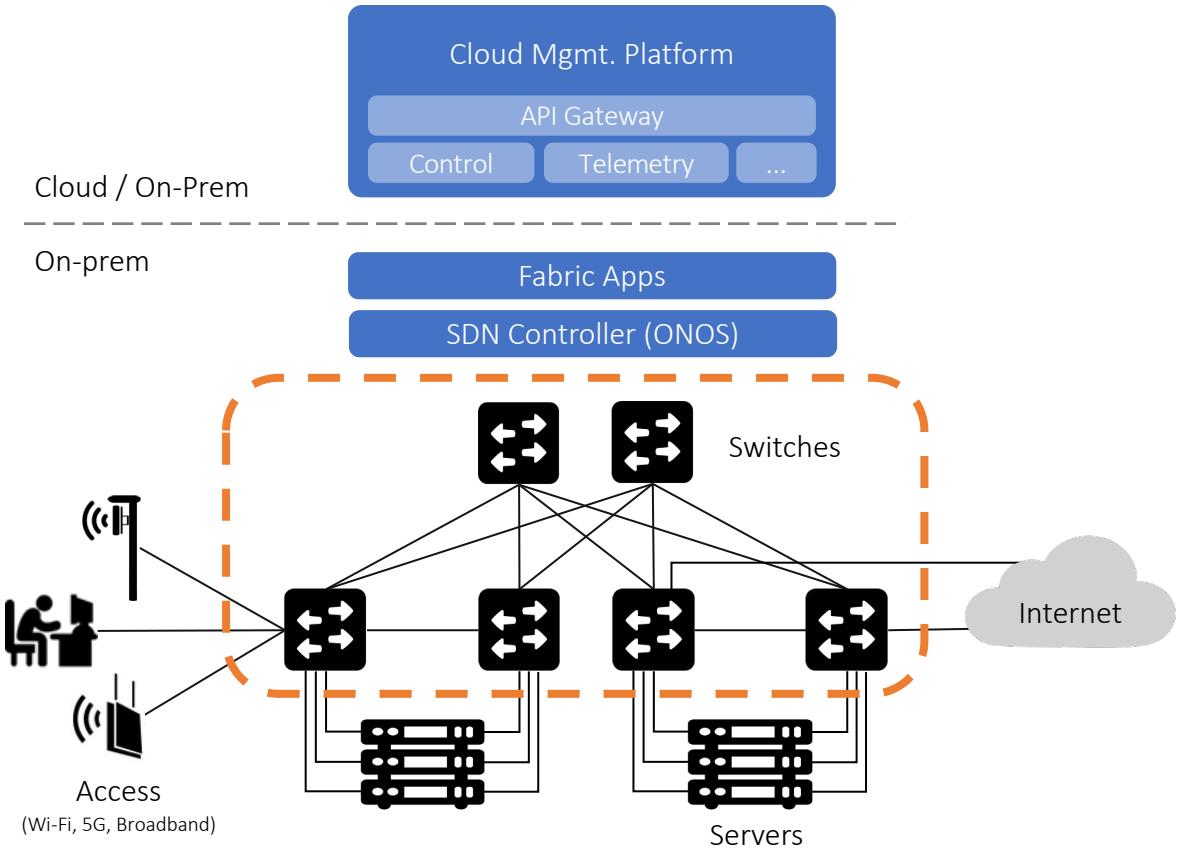

- Stratum runs on the switching infrastructure
- Stratum provides:
  - Precise forwarding control on a P4 defined pipeline
  - Increased visibility via INT
  - Network function offloading
  - P4RT / gNMI as SDN interfaces
- *What about on the server?*

Extending the Control Domain

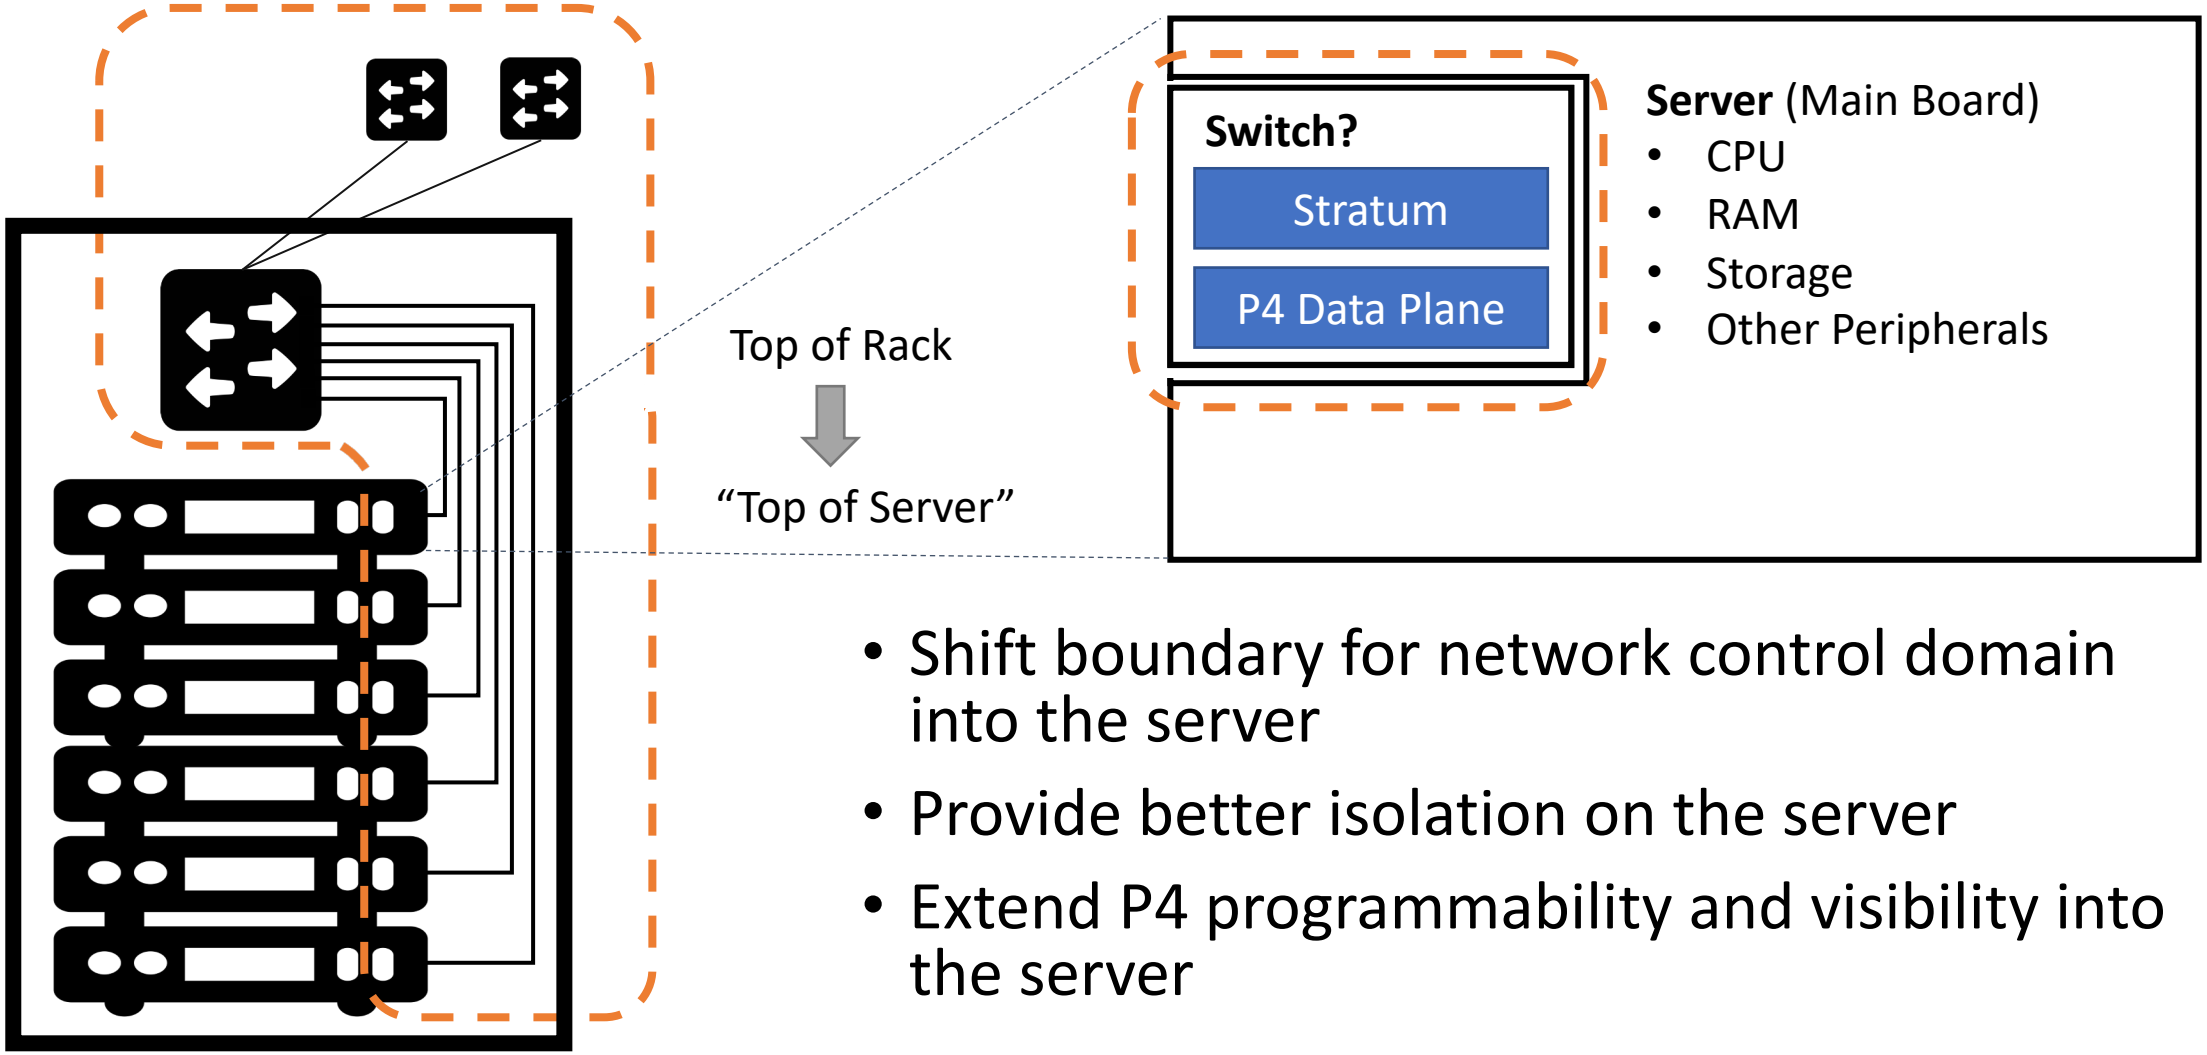

- Shift boundary for network control domain into the server
- Provide better isolation on the server
- Extend P4 programmability and visibility into the server

What is an IPU or DPU?

- **I**nfrastructure **P**rocessing **U**nit *or* **D**ata **P**rocessing **U**nit
- Can be deployed on a PCI card:
  - High speed Ethernet ports
  - Embedded micro-server
  - Programmable network pipeline
  - RDMA, NVMe, DPDK, IPDK, etc.
  - Offload functions (e.g. crypto)
- Standalone operation and management
- Physically isolated from main CPU

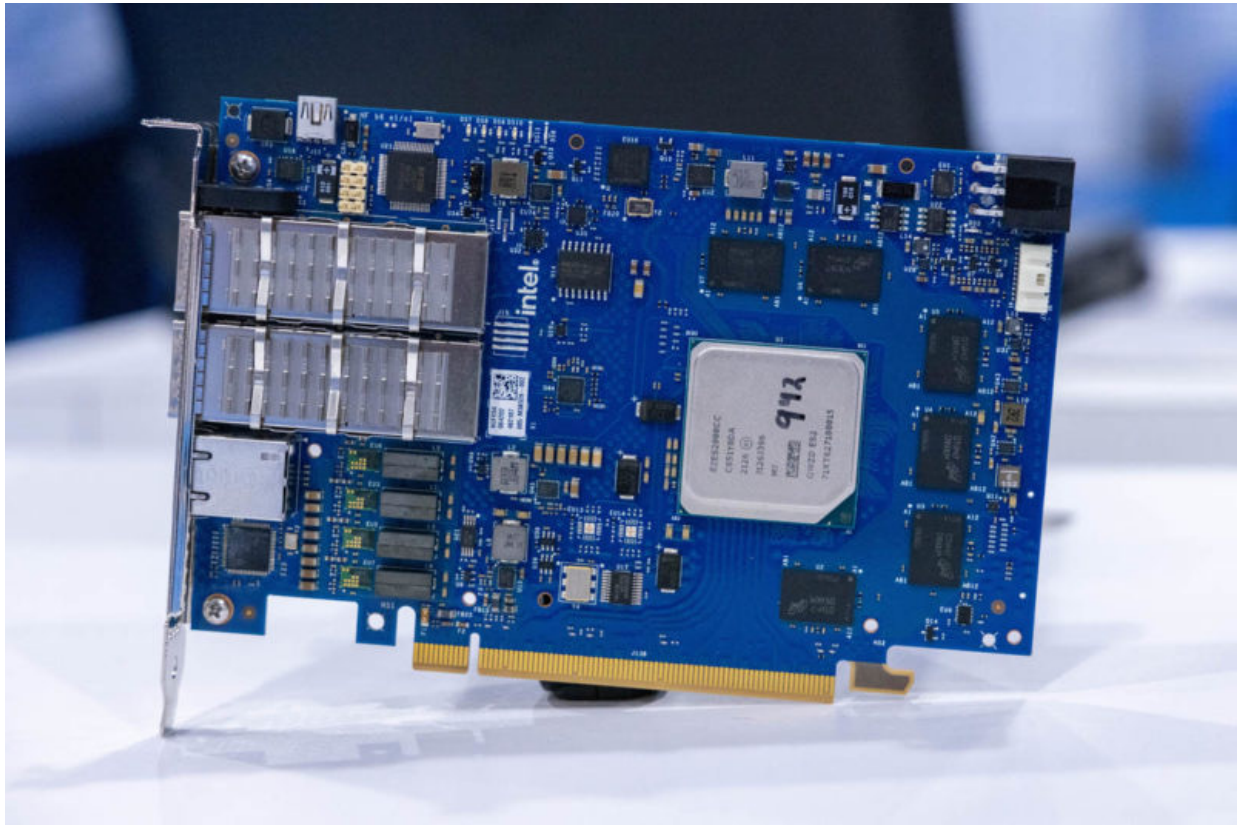

Intel® Mount Evans IPU at Intel Vision 2022

Photo credit: https://www.servethehome.com/new-intel-mount-evans-ipu-asic-dpu-at-intel-vision-2022/

Programming the IPU using TDI

- TDI (Table Driven Interface) provides a common interface for different targets
- TDI is independent of P4 program, architecture, and backend target
- Entry schema (tdi.json) is driven by the P4 program and target-specific fixed functions (e.g. traffic manager)

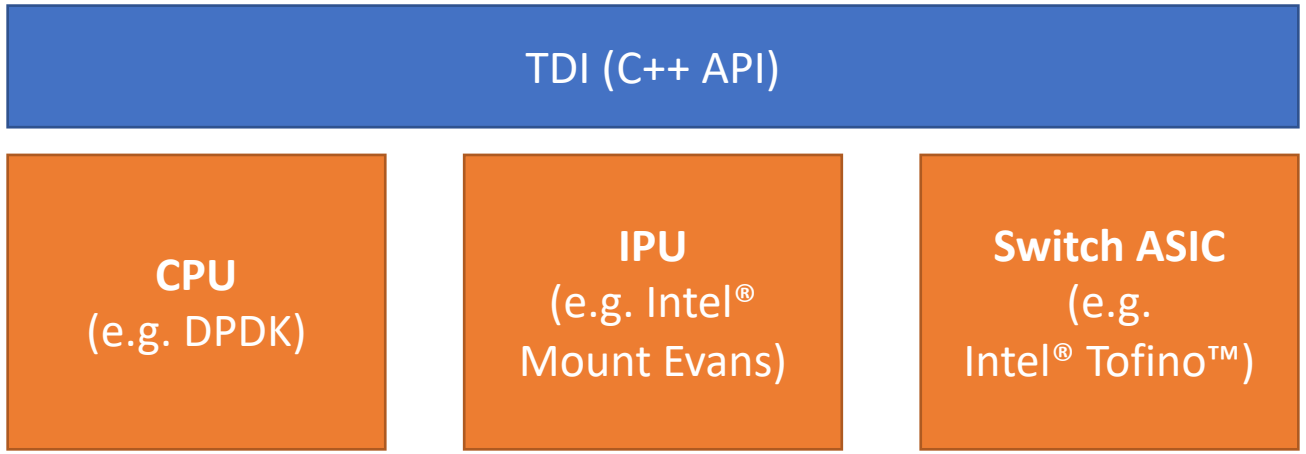

Stratum Roadmap: Stratum_TDI

- TDI is already open sourced: https://github.com/p4lang/tdi
- Internal version of P4-OvS (**`infrap4d`**) is built on TDI
  - Based on **`stratum_bfrt`**
  - BfRt-based code refactored to use TDI
- Later in 2022, plan to upstream new **`stratum_tdi`** target that will support:
  - P4-DPDK, Intel® Mount Evans IPU, Intel® Tofino™ switch ASIC
  - FPGA, additional switch ASICs (future)
- Upstreaming brings many benefits
  - Provide a high(er) performance Stratum software switch target, based on DPDK
  - Allow P4-OvS to more easily consume Stratum fixes and enhancements
  - Path for unified software stack for different P4-enabled targets

Adapting Stratum for P4-OvS

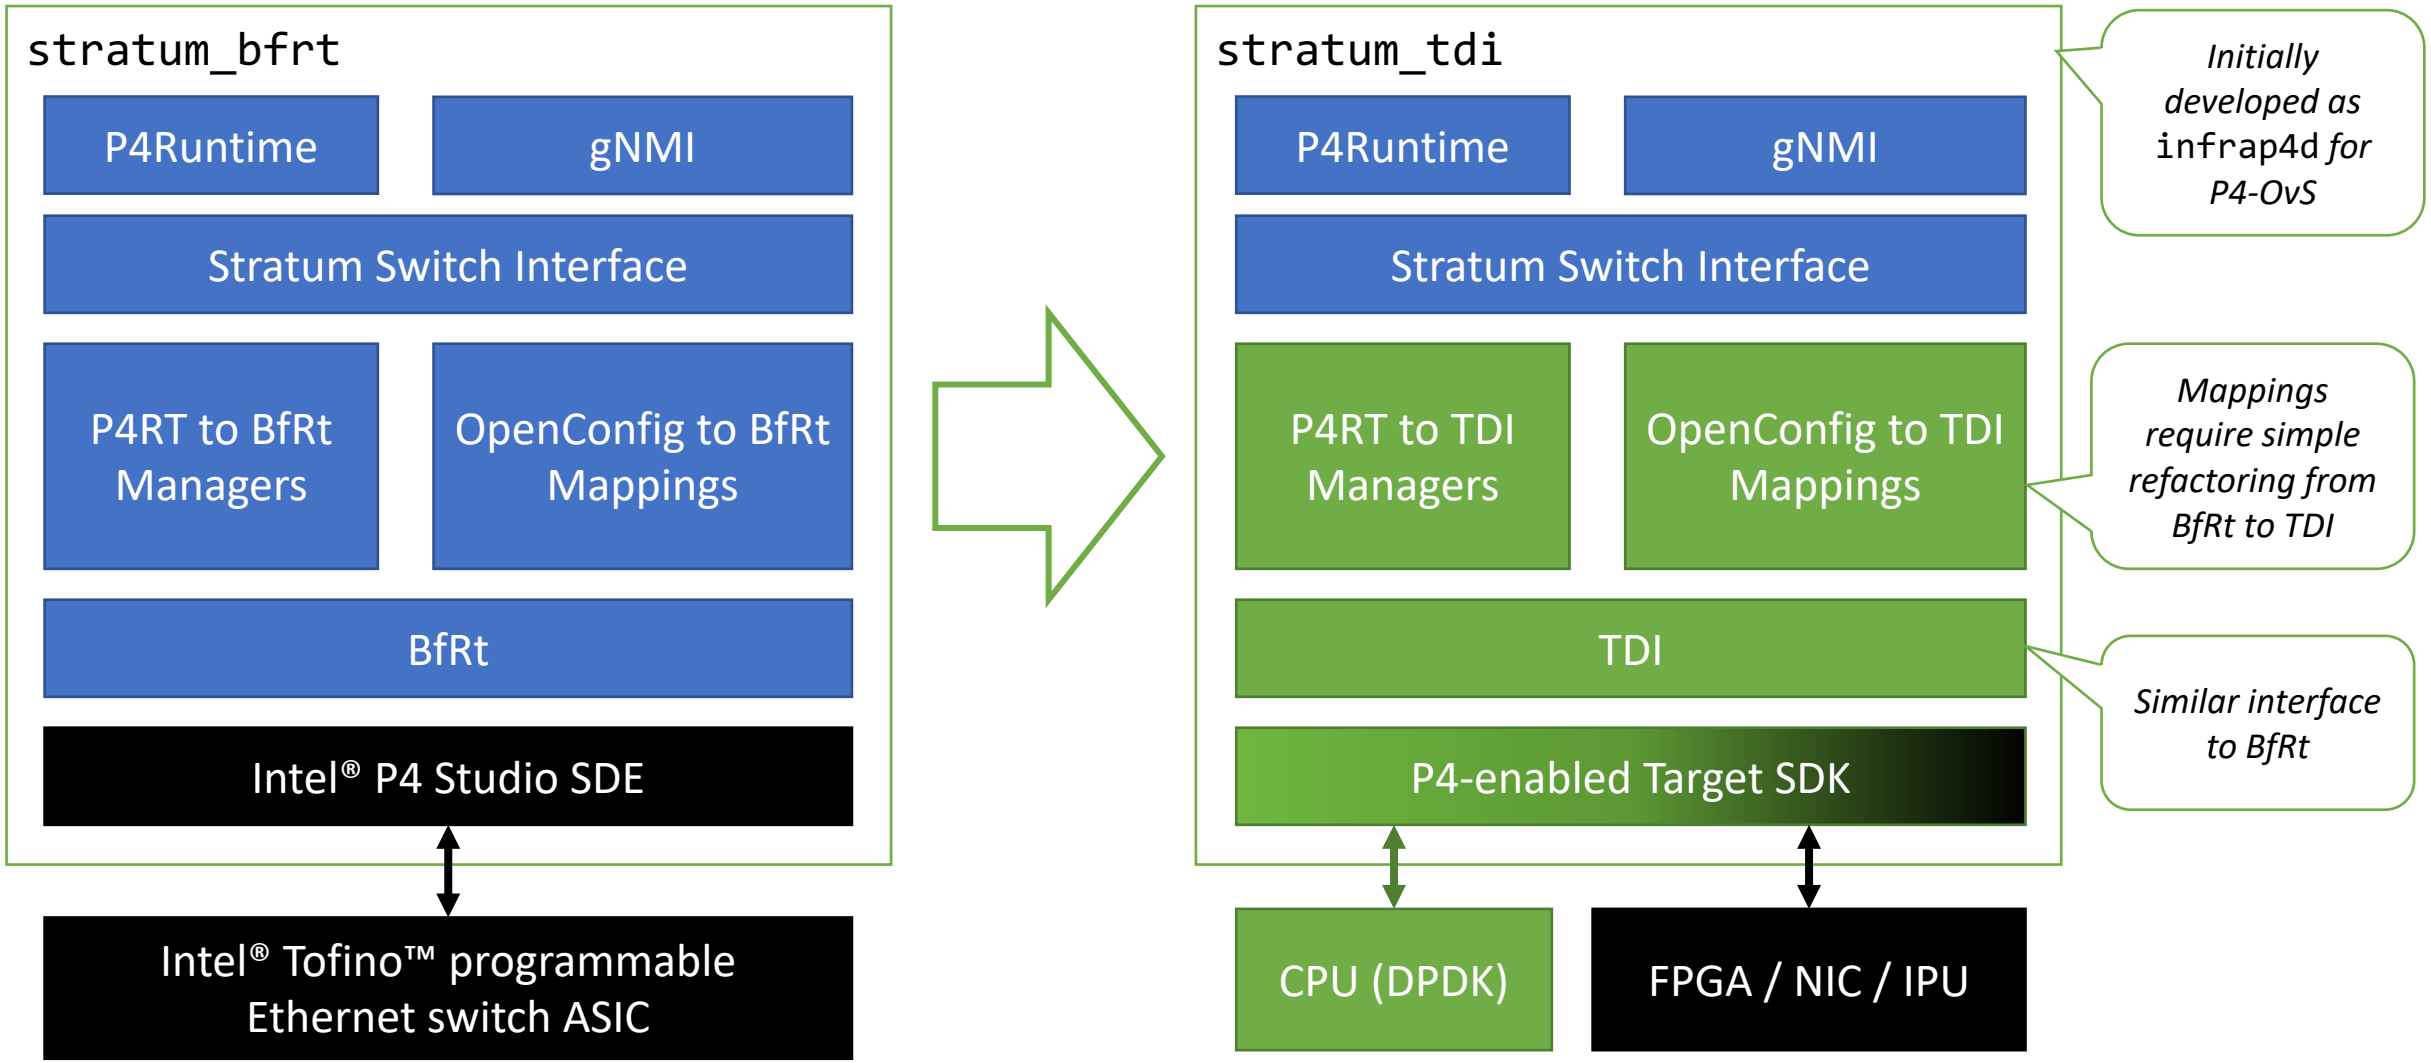

Stratum Roadmap: Base OS Background

- Stratum runs on Linux
  - Not tied to any particular distro, but it's helpful to provide a default
  - Users have a lot of flexibility
- Open Network Linux (ONL) has been the default base OS for many years
  - ONLPv2 branch and platform API
  - Very few commits over past 2 years
  - Based on older distro and kernel
  - Lack of new platform support, bugs, performance issues
- SONiC is a better choice for the Stratum's base OS for switches
  - Also based on Debian
  - Actively maintained by a diverse community
  - Regular releases (twice a year) and updates/fixes

Stratum Roadmap: Base OS Plan

- ONL is the default base OS for Stratum 22.06
  - Users can continue to use ONL (or another distro) if they prefer
- Moving forward, SONiC will be the default base OS
  - Note: the SONiC control plane is disabled
  - We just use the Linux distro, kernel, and platform API
- Stratum releases starting with 22.09 will:
  - Point to a specific SONiC image
  - Contain kernel modules for recently supported SONiC images (and ONL, too)
  - Be tested on the SONiC base OS
- Longer term opportunities
  - Create a Stratum PHAL adapter for SONiC's platform API
  - Plug Stratum into SONiC's CLI (e.g. allow users to configure and monitor ports)

Summary

- Stratum 22.06 recently released in June
  - Stratum_BfRt fully replaces Stratum_Bf
- Stratum is venturing into server via the IPU / DPU
- Stratum 22.09 planned for September
  - Initial version of Stratum_TDI
  - SONiC as the default base operating system

Getting Involved

How do I get involved with Stratum?

- Send an email to stratum-dev or a Slack message to #stratum-dev
- Join the bi-weekly technical steering team (TST) meeting
- Send a pull request with your changes to the Stratum Github repo
- Try the NG-SDN tutorial
- Help review and test Stratum on P4-DPDK

Notices and Disclaimers

- Intel technologies may require enabled hardware, software or service activation.
- No product or component can be absolutely secure.
- Your costs and results may vary.
- Intel does not control or audit third-party data. You should consult other sources to evaluate accuracy.
-

Thank You

Follow Up Links:

https://opennetworking.org/stratum/

https://github.com/stratum/stratum