




Stratum Techinar

January 26, 2021 | 9am PST



Brian O'Connor
Stratum Technical Lead



Maximilian Pudelko
Stratum TST Member

Outline



Quick Stratum Intro (1 slide)

State of Stratum: Use cases and Users

Stratum Release Highlights

Stratum Roadmap

What is Stratum?



Open source, production targeted, **thin** switch OS

Thin means bring your own *data plane pipeline* and *control plane apps*

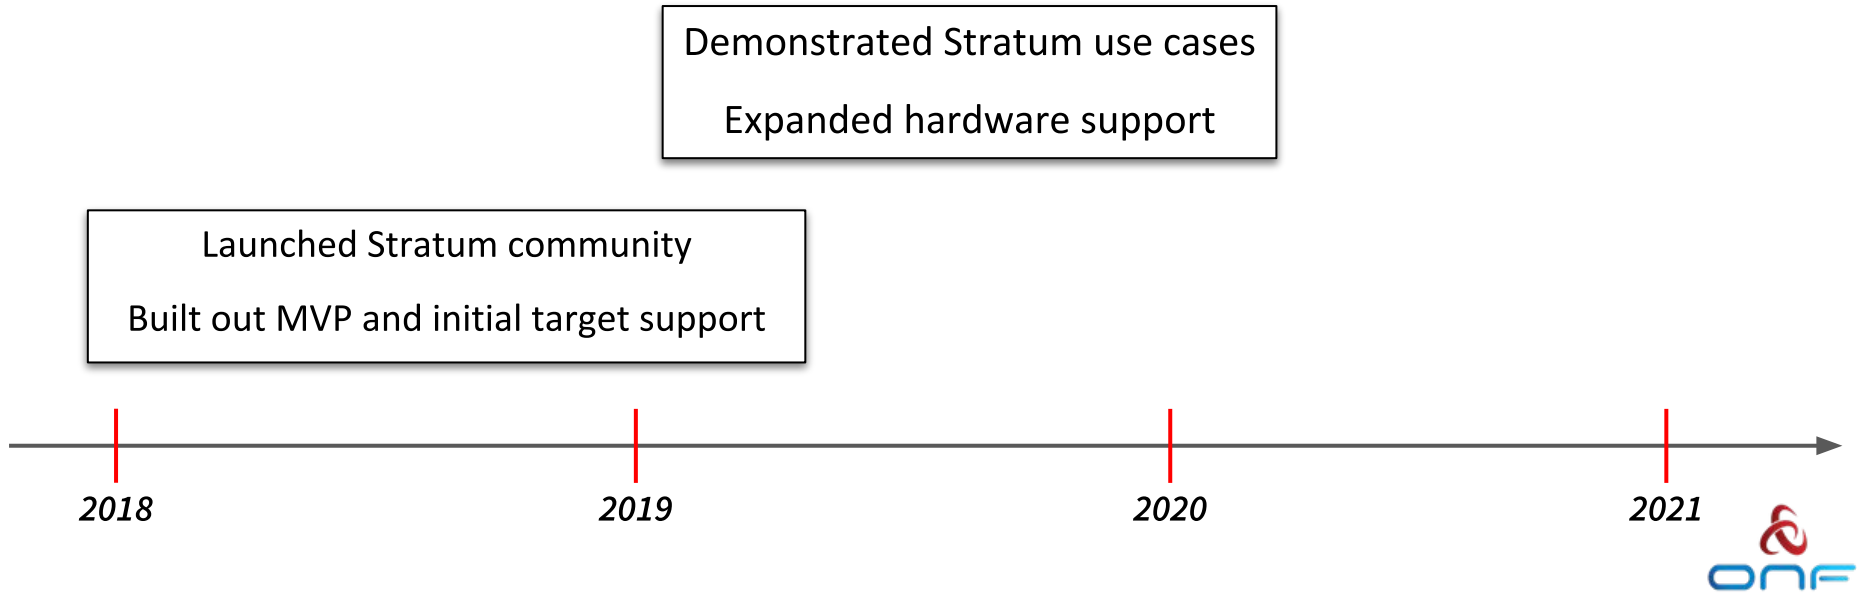
- There's no routing (BGP), trunking (LACP), etc. by default
- Pairs with *fabric.p4*, *ONOS*, and *Trellis* for a complete networking stack

Stratum provides rich, programmatic interfaces

- **P4**: specify the data plane pipeline
 - Works with fixed-function and programmable targets
- **P4Runtime**: program the data plane at runtime
 - Works on the box (distributed) or off the box (SDN)
- **OpenConfig**: model ASIC and platform configuration
 - Provided, but extensible
- **gNMI**: configure and monitor the data plane
 - Enables subscriptions and streaming telemetry



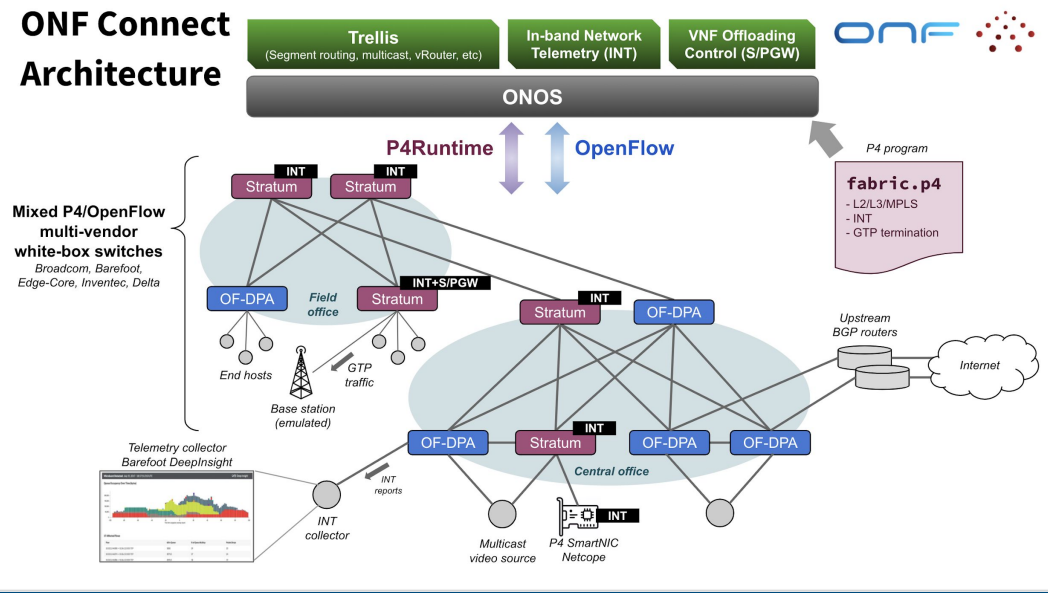
State of Stratum



Demo @ OCP - March 2019



ONF Connect Architecture



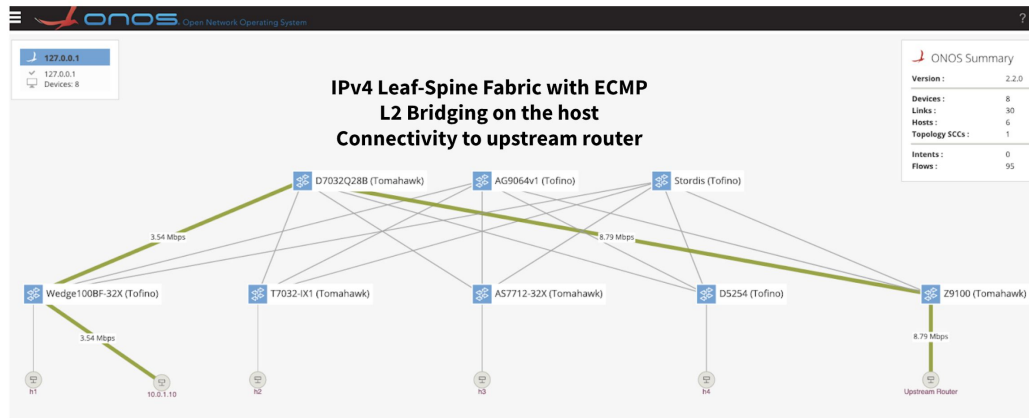
- Multi-stage leaf-spine fabric of Broadcom OF-DPA (OpenFlow) and Tofino Stratum (P4Runtime) switches
- Unified control plane using ONOS and Trellis
- Initial version of SPGW offload and INT using Stratum



Demo @ ONF Connect - Sept. 2019



Stratum Interoperability Demo



BAREFOOT
NETWORKS

BROADCOM



DELTA

edge-core
NETWORKS

Inventec

QCT

STORDIS


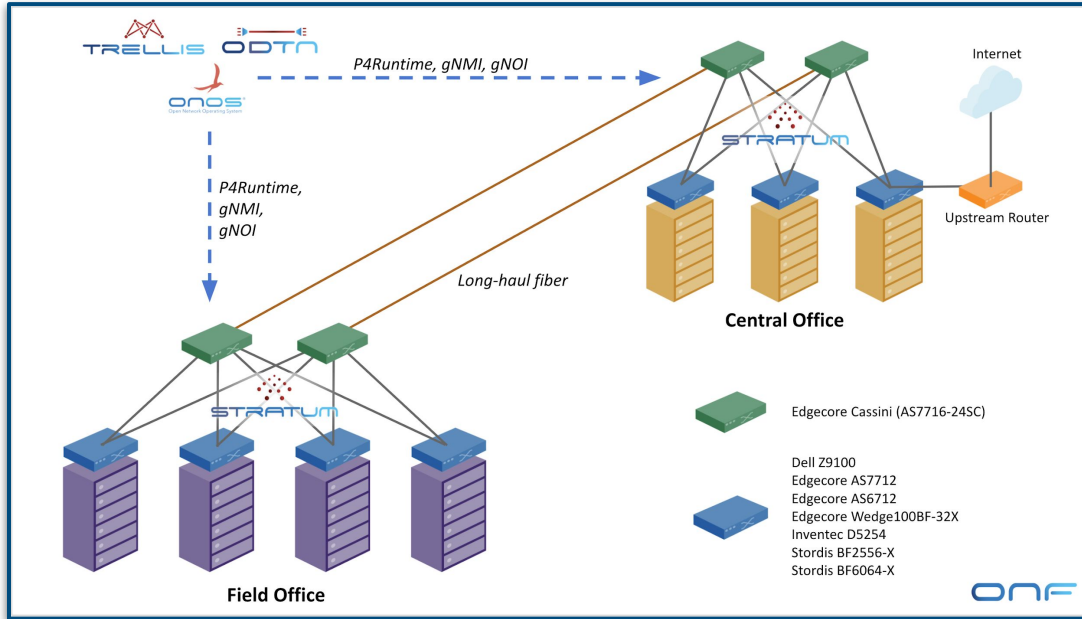
Stratum Open Source release

3 demos

- **SEBA** - Stratum used to offload BNG data plane functions
- Stratum used with ONOS to demonstrated **zero-touch provisioning**
- Demonstrated **Stratum on Tofino and Broadcom** Tomahawk silicon in IP routed, leaf-spine fabric

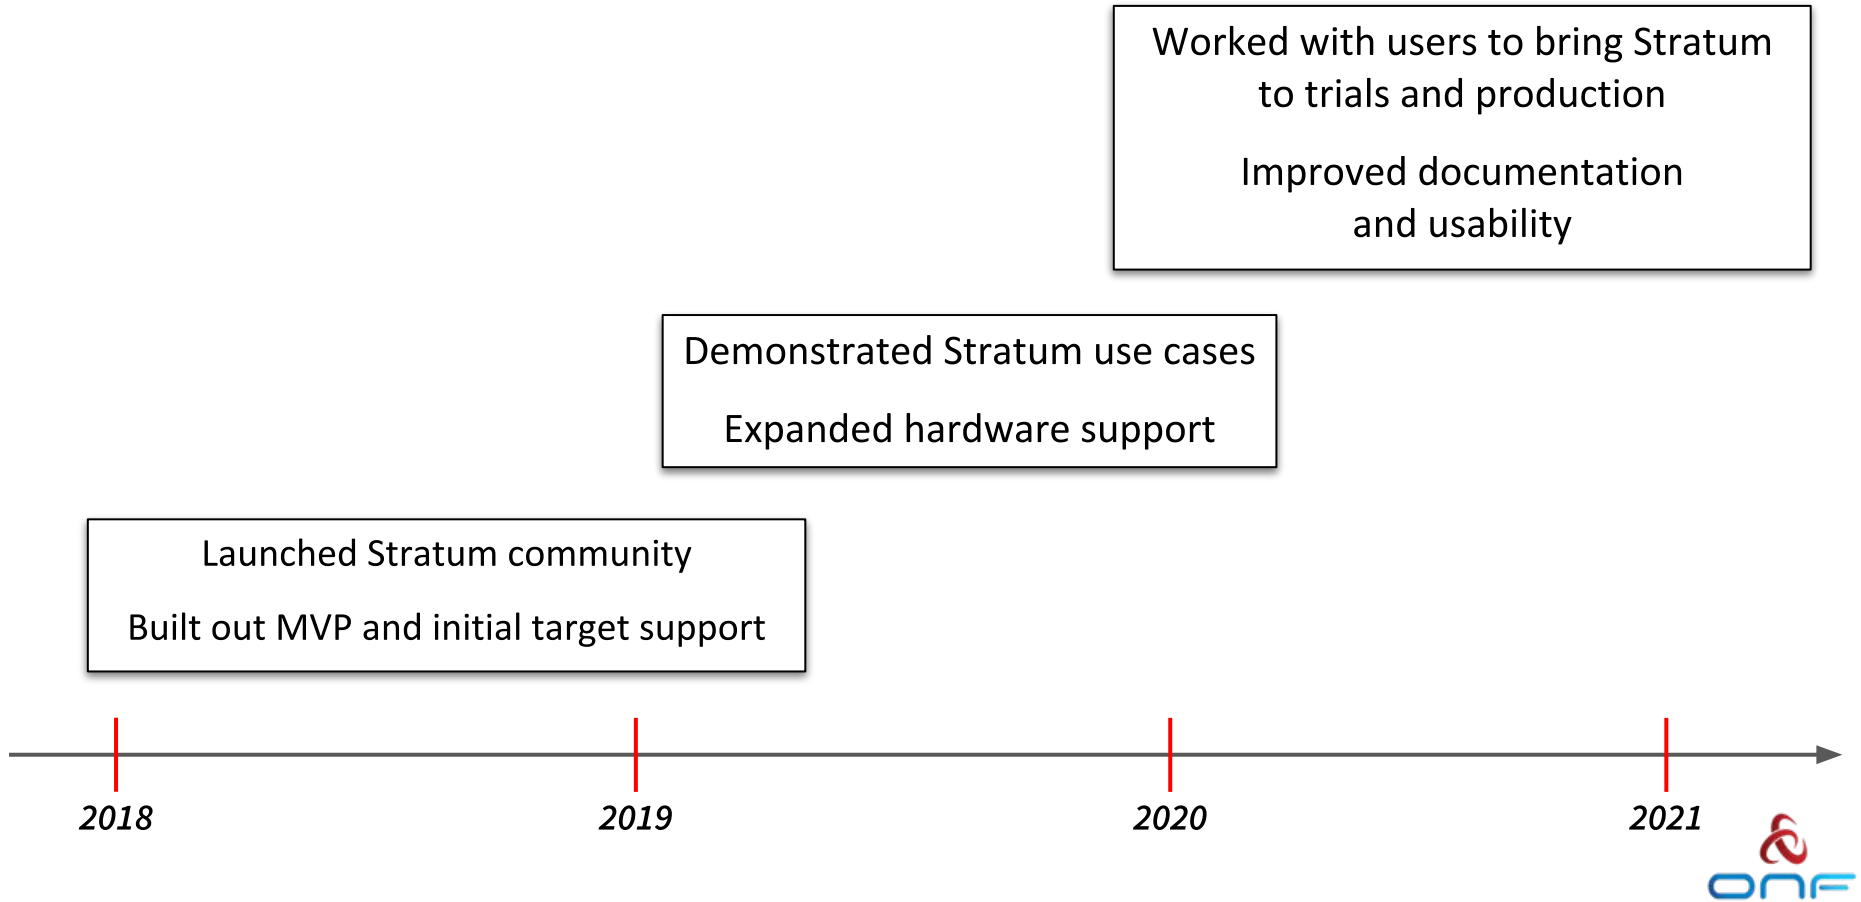


Demo @ OCP - May 2020



- Demonstrated data center fabric topology with remote field office using Trellis, ONOS, and a heterogeneous set of Stratum hardware
- Data center interconnect (DCI) using Stratum on Edgecore Cassini packet-optical transponder

State of Stratum



Production-Grade Stratum



- Stratum switches enter field trials and production
 - Aether Mobile Edge fabric
 - Canopus Flow Classifier
 - ... and others
- FPGA NIC support added to Stratum
- Tassen - Building a BNG product using Stratum
 - FPGA target: Dell, Intel, Benu Networks
 - Tofino target: ONF, Intel
- Stratum selected as switch OS for Pronto platform



What is it?

- Disaggregated BNG born from DT's Access 4.0 platform
- Enables Control and User Plane Separation (CUPS)
- Provides a logical abstraction that works on multiple targets

How does it use Stratum?

- Stratum runs the user plane (including offloaded functionality) and provides the control plane interface

Why does it need Stratum?

- User plane pipeline is defined using P4
- Stratum provides common control plane interfaces across targets (including switching ASIC and FPGA)

Options for BNG Disaggregation

Traditional chassis-based vendor
Broadband Network Gateway (BNG)



*Stratum-powered
user plane*

Specialized router

QinQ, PPPoE, IPoE, accounting,
HQoS, lawful intercept, wholesale
tunnel relay (L2TP), multicast, etc.

<https://github.com/opennetworkinglab/tassen>

Slide courtesy of Carmelo Cascone, Craig Stevens, and Mario Kind

BNG-CP → Intelligence

Choice of programming language and architecture

Open CUPS API



*Provides silicon-independence
& facilitates interoperability*

BNG-UP → High-speed forwarding

Choice of merchant silicon and port configuration

Magenta Switch

Single-chip Qumran 2C



Deutsche Telekom

SmartNIC

Single-chip FPGA



Dell, Intel, Benu

Tofino + Qumran AX

Dual-chip, HQoS in OLT (WIP)



ONF

OLT (QAX)
Leaf switch
(Tofino)

Tofino + FPGA

Dual-chip, HQoS in FPGA



TU Darmstadt

*Example of HW-based implementations existing today.
Others are possible, including SW-based processing.*





What is it?

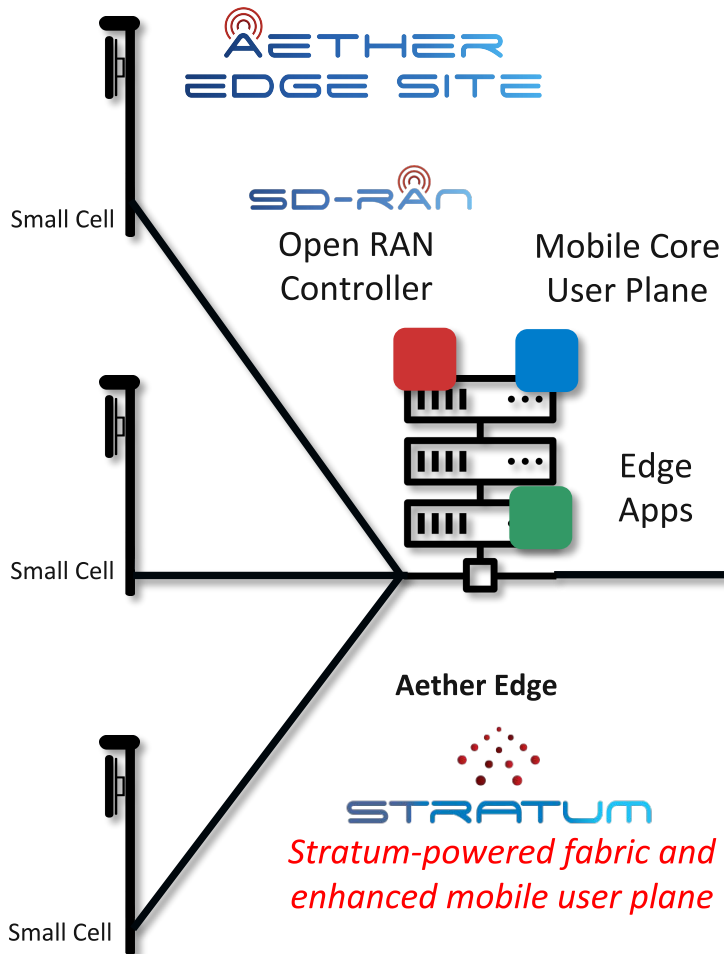

- Private 5G/LTE edge access and compute platform targeted for enterprise
- Works with licensed and unlicensed (e.g. CBRs) spectrum
- Connects to public cloud for subscriber management and configuration
- Enables Industry 4.0 applications (real-time analytics, AI/ML offload, IoT)
- ***Already operational at more than 10 ONF and partner sites across the globe***

How does it use Stratum?


- Stratum runs as:
 - the switch OS in the Aether Edge pod
 - As the “NIC OS” for select FPGAs that are part of the user plane
- Performs routing in the edge pod using Trellis and offloads of some user plane functions (e.g. GTP encap/decap)

Why does it need Stratum?

- 5G/LTE user plane brings unique data plane requirements
- Need fine grained control of traffic for user plane services (forwarding, QoS, metering, accounting, etc.)
- Stratum offers exceptional visibility and control via P4Runtime and gNMI



**AETHER
CORE**



<https://aetherproject.org>

Slide courtesy of Oğuz Sunay

ONF



What is it?

- \$30M DARPA program to create a 5G network with verifiable, closed-loop control
- Control plane applications and data plane extensions built on Aether platform in partnership with Cornell, Princeton, and Stanford

How does it use Stratum?

- Stratum runs as the switch OS in the Pronto edge fabric as with Aether
- Performs fabric / mobile user plane functions along with exporting additional visibility via INT

Why does it need Stratum?

- Top-down, full control of the data plane pipeline through P4
- Fine grained visibility and control of network state
- Enhanced visibility of packet-level behavior using INT

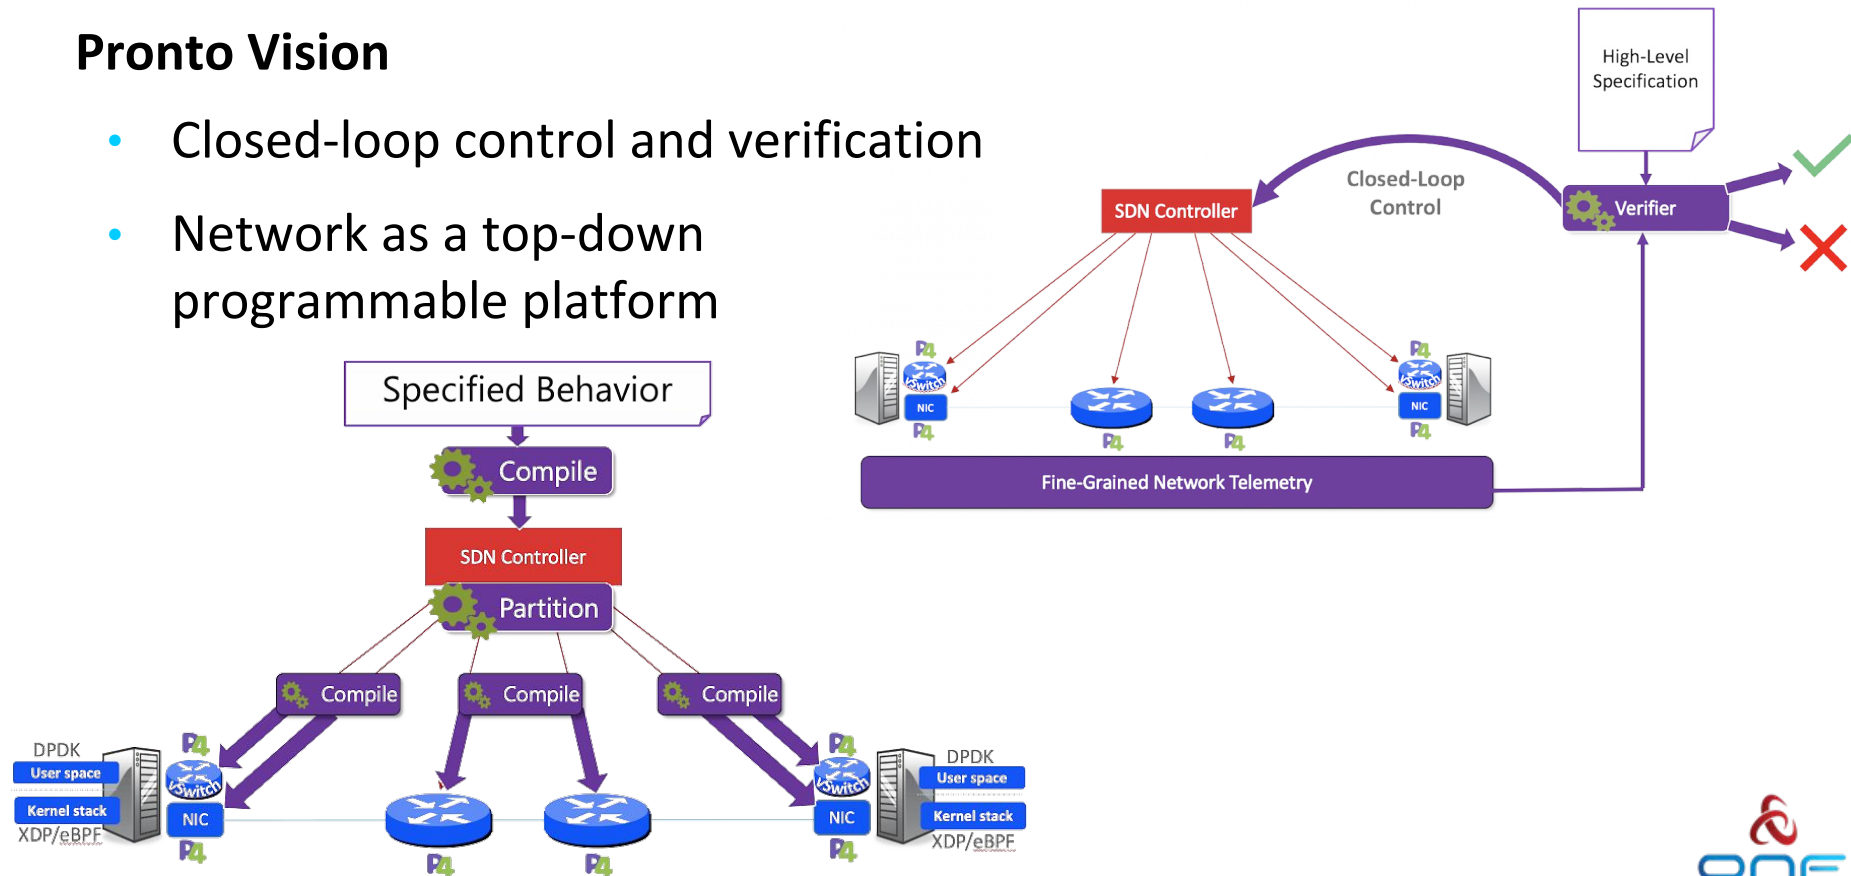
Pronto

<https://prontoproject.org>



Pronto Vision

- Closed-loop control and verification
- Network as a top-down programmable platform



State of Stratum



**Stratum is starting to be used
by the community in real world,
production scenarios.**

Worked with users to bring Stratum
to trials and production
Improved documentation
and usability

Demonstrated Stratum use cases
Expanded hardware support

Launched Stratum community
Built out MVP and initial target support

2018

2019

2020

2021



Part 2: What's new since open source launch?



- Road to Production Worthiness
- Testing Story
- Usability
- Release Strategy
- New Platforms and ASICs

What's new? - Road to Production Worthiness



Production use uncovered platform and interface shortcomings

- Missing features - Stratum not taking advantage of the whole API
- Poor platform support - Reliance on incomplete or unmaintained APIs

Replaced with better implementations to accommodate production uses

- Intel: PI / v1model → BfRt and Tofino Native Architecture (TNA)
- Broadcom: SDKLT → OpenNSA
- *ONLPv2 → Tofino BSP, SONiC platform API (in progress)*

What's new? - Road to Production Worthiness



Customers demand pipeline flexibility, e.g. support for Tofino Native Architecture (TNA) and BfRt

Stratum accommodates this need with a new target: Stratum-bfrt

- Maximize portability with v1model
- Or, unlock the full capabilities of the Tofino hardware:
 - Register access
 - Full control over metadata bridging
 - Resource optimizations, like pipeline folding
 - Access to traffic manager
 - Multi-pipe programs

What's new? - Road to Production Worthiness



Started departure from ONL as sole choice of base OS and platform API

- ONL Pv2 support often incomplete, buggy or plain unavailable
- Unclear future of ONL project
- Conflicts with heterogeneous deployment reality (Ubuntu, Debian 10, BYO)

Stratum's new platform plugin model allows choosing the best available implementation

- ONLP
- Transponder abstraction interface (TAI)
- Tofino board support package (BSP) *[preliminary]*
- SONiC platform API *[planned]*

What's new? - Testing Story



A comprehensive testing, certification and deployment lifecycle:

- Added unit testing guidelines and improved test coverage
- Enabled integration testing on software switches (beyond bmv2)
 - Tofino model now available for TNA pipelines
 - “Test what you deploy”
- Started Continuous Certification Program (certification for the CI/CD era)
 - Nightly hardware tests with newest code using Jenkins
 - Powered by Test Vectors (Stratum’s black-box test suite)
- Created Stratum-replay
 - Capture and replay production traces in a development environment
 - Prevent regressions by adding traces to integration tests

What's new? - Usability



When open sourced, Stratum had a usability problem:

- Lots of tribal knowledge required
- Manual system setup
- Inconsistencies between Tofino and Broadcom targets

Make Stratum deployable by non-developers:

- Provide universal binaries, not an explosion of options to choose from
 - Complete packages (Debian) and images (Docker)
 - Stamping and labels
- Enhance documentation
 - Add configuration, runtime, and troubleshooting sections
 - Remove redundant sections
- Improve user experience
 - Do the right thing by default and remove extra flags
 - Automate whenever feasible (git checkout to .deb file in a single command)
- Streamline pipeline packing format on Tofino
 - From implicit, arcane binary packing to a standardized protobuf
 - Includes tooling that bridges the gap between p4c and Stratum
- Simplify configuration files
 - Hide irrelevant system internals, like SDK port numbers

Today, Stratum is deployed at ONF by a team separate from development



What's new? - Release Strategy



First open source release in Sept 2019

Continue “live at head” philosophy

- No LTS baggage, feature backports or Semantic Versioning
- Lightweight and fast moving

... but for deployment, we've have quarterly, qualified releases

- June 2020
- Sept 2020
- Dec 2020
- *March 2021 (planned)*

What's new? - Release Strategy



“No LTS/SemVer?! How are breaking changes handled?”

Non-atomic refactoring with continuous deployments (cf. [Abseil](#), [Bazel](#), Protobuf):

0. Release **N**: Contains old feature
1. New feature added at head
 - Discuss at TST meetings and on pull requests
 - Ensure no changes to the old feature (sometimes this means introducing a new target)
 - Include tools and docs for assisted migration
 - Mark old functionality as deprecated
2. Release **N+1**: New feature is included
 - Old and new feature is available; old one is deprecated (expect logs)
3. Release **N+2**: Old feature is removed
 - Clients must use the new feature

New Platforms & ASICs



Tofino

- APS Networks BF6064X
- APS Networks BF2556X
- Delta AG9064v1
- **Edgecore Wedge100BF-32QS**
- Edgecore Wedge100BF-32X
- Edgecore Wedge100BF-65X
- Inventec D5254
- **Inventec D5264Q28B**
- Inventec D10056
- Inventec D10064
- **Netberg Aurora 610**
- **Netberg Aurora 710**
- **Netberg Aurora 750**

Broadcom

- Tomahawk
 - Edgecore AS7712-32X
 - **Edgecore AS7716-24XC Cassini**
 - Dell Z9100
 - Inventec D7032
 - QCT QuantaMesh T7032-IX1
- **Tomahawk+**
 - **Edgecore AS7716-24SC Cassini**
 - **Inventec D7054**
- **Tomahawk2**
- **Tomahawk3**
- **Trident2**
 - **Edgecore AS6712**
 - **Inventec D6254**

Smart NIC

- **Intel PAC N3000**

Software Switches

- bmv2
- dummy
- **tofino-model**

** New since 2019*

For an up-to-date listing, see <https://github.com/stratum/stratum#supported-devices>



2021 Technical Roadmap for Stratum



1. Support production users

- QoS model
- Tofino features (scheduler, traffic manager)
- External API enhancements (P4Runtime, gNMI)
- Testing (pre-merge hardware validation, CI system)

2. Portability

- Standardized config (move from Chassis Config to OpenConfig)
- Platform API integration / plugins (focus on transceivers and ASIC)
- Run / test on other switch OSes

3. Usability

- Documentation
- Tofino model usability (Mininet integration, P4 program linking)
- Build and release automation



What about SONiC?



Comparing Stratum and SONiC's dataplane (SAI)



- ✓ Vendor-independent switch interface
- ✓ Provides common, fixed logical switch pipeline that is well suited for L2/L3 and other data center uses
- ✓ Common API and pipeline allows development of common control plane components (or easy integration of control plane applications; e.g. FRR)
- ✗ Changes require community consensus, which takes time and required vendor support
- ✗ Vendor-specific functionality exposed through private extensions
- ✗ Pipeline control plane lives on the box

- ✓ Vendor-independent switch interface
- ✓ Fully user programmable switch pipeline using P4
- ✓ Vendor-specific functionality easily exposed in P4 to users
- ✓ Supports on the box and SDN-style control planes
- ✗ Requires target support for P4
- ✗ No common pipeline means control plane components are often pipeline-specific (and may need to be changed when the pipeline does)

What about SONiC?



Exploring new SONiC features: PINS (P4 Integrated Network Stack)



- ✓ Vendor-independent switch interface
- ✓ Provides common, fixed logical switch pipeline that is well suited for L2/L3 and other data center uses
- ✓ Common API and pipeline allows development of common control plane components (or easy integration of control plane applications; e.g. FRR)
- ~~✗ Changes require community consensus, which takes time and required vendor support~~
- ~~✗ Vendor-specific functionality exposed through private extensions~~
- ~~✗ Pipeline control plane lives on the box~~

- ✓ Vendor-independent switch interface
- ✓ Fully user programmable switch pipeline using P4 (*extensible for supported targets*)
- ✓ Vendor-specific functionality easily exposed in P4 to users (*for supported targets*)
- ✓ Supports on the box and SDN-style control planes
- ~~✗ Requires target support for P4~~
- ~~✗ No common pipeline means control plane components are often pipeline-specific (and may need to be changed when the pipeline does)~~

PINS has extensible SAI vs. Stratum has fully programmable



Stratum Big Picture Roadmap



**Stratum is great for users that want/need
full customizability of the data plane.**

(All of the use cases we've talked about today fall into this bucket.)

1. Support production deployments (new features, bug fixes)
2. Play to Stratum's strengths by enhancing programmability
3. Make Stratum a complementary option for any switch OS by improving interoperability, portability, and platform integrations
 - *Think of Stratum as an application that provides a rich, programmable interface to the data plane that runs anywhere*

Action Items



Take the [2021 Stratum User survey](#)

- Help shape our priorities and future of Stratum

How to get involved?

- Join the TST calls, *#stratum-dev* Slack, and mailing lists
 - Details here: <https://github.com/stratum/stratum/wiki/Discuss>
- Try the [NG-SDN tutorial](#)
- Send PRs



Thank You

Follow Up Links:

<https://opennetworking.org/stratum/>

<https://github.com/stratum/stratum>