ONF Open Source Solutions Enable Data Center Interconnect and Broadband Access

Brian O'Connor, Carmelo Cascone, Abhilash Endurthi,
Max Pudelko, Yi Tseng, You Wang
ONF

Webinar on March 17 and 18, 2020

ONF's Project Portfolio Areas

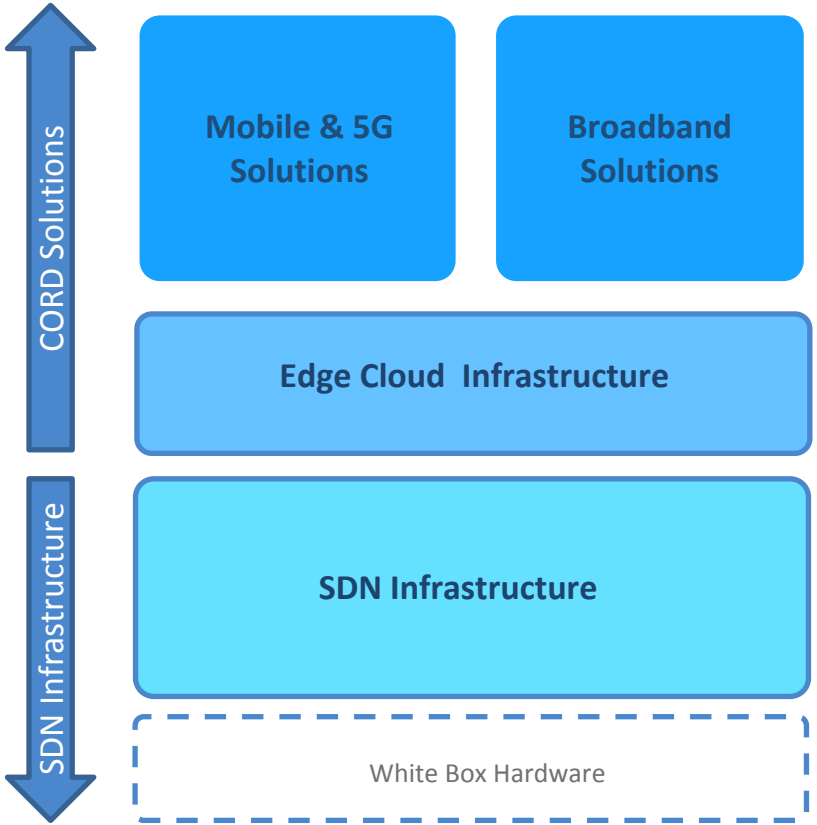

ONF's Projects

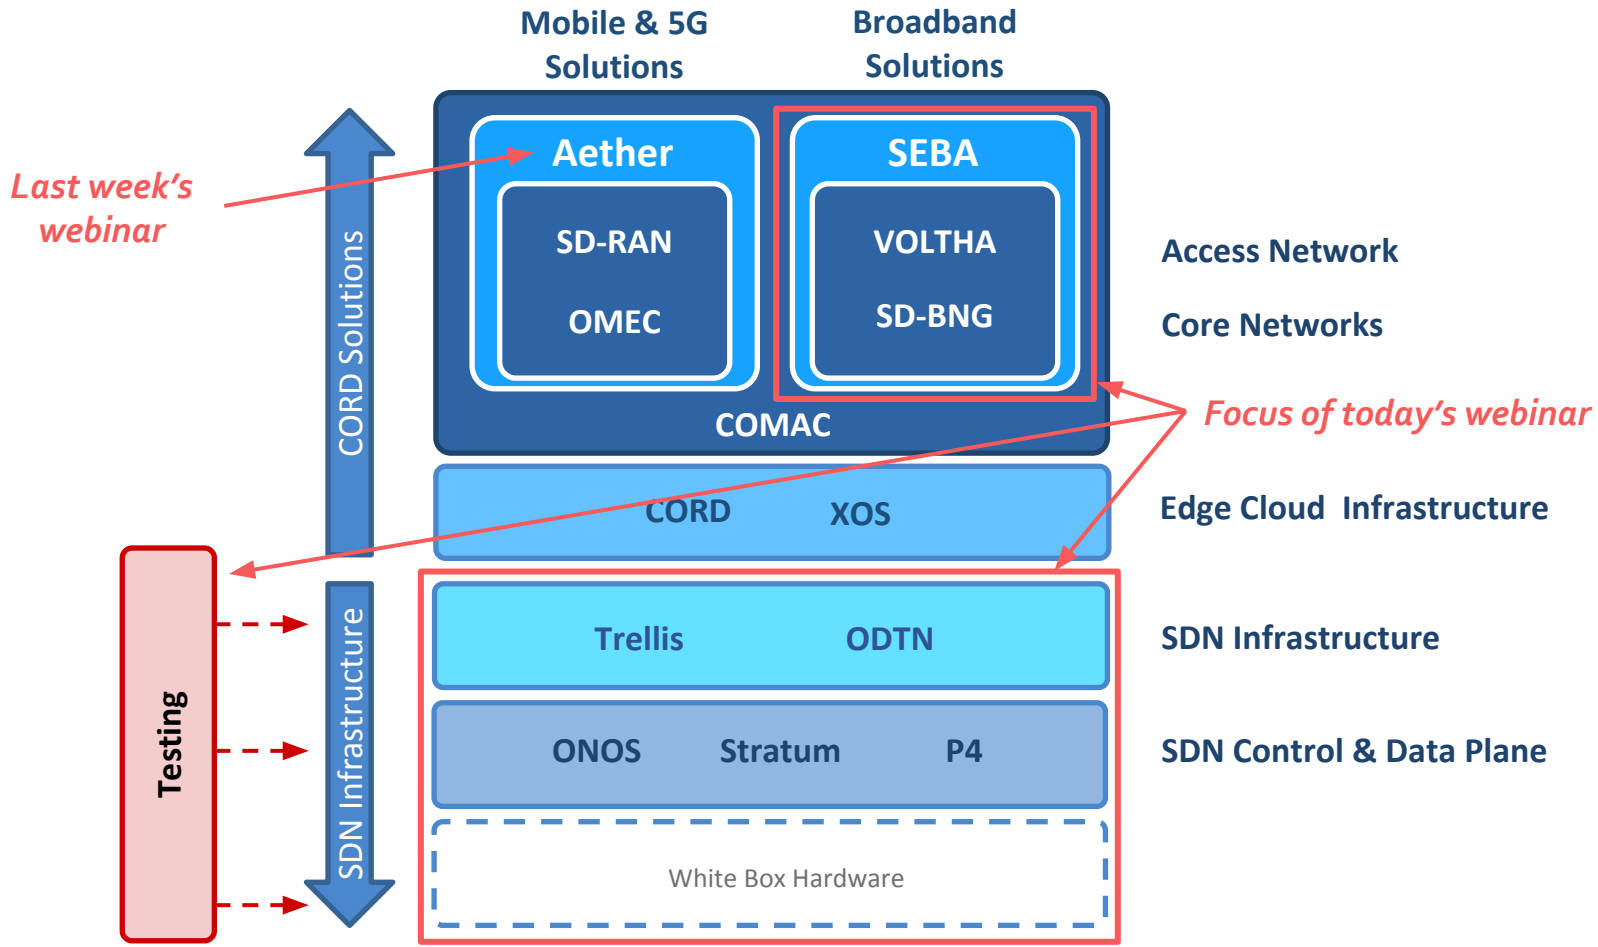

Webinar Overview

1. Production grade, P4 Programmable, SDN Fabric
   - **Trellis** + **Stratum** + **ODTN**
   - Stratum support for **Cassini**, a packet optical transponder
   - Support for Broadcom **OpenNSA** SDK
2. New Use Case enabled by the P4 fabric
   - **SEBA** + Stratum
   - Software-defined BNG (**SD-BNG**) for the access network
3. Testing, Deployment, and Certification
   - **TestVectors**, a black-box, system test framework for Stratum devices
   - ONF's CI/CD pipeline for Stratum
   - ONF's new **continuous certification** (CC) program

Webinar Summary

- Pathway to convergence of Trellis, ODTN, SEBA, and Stratum
  - Today, we demonstrated (Trellis, ODTN, ONOS, Stratum) and (SEBA, ONOS, Stratum)
  - Aether is another logical convergence opportunity for the future
- P4 is enabling new use cases on commodity hardware
  - SD-BNG data plane is built with P4
- ONF testing will enable stable and rapid solution delivery
  - TestVectors enables the Continuous Certification Program for Stratum

Future Work:

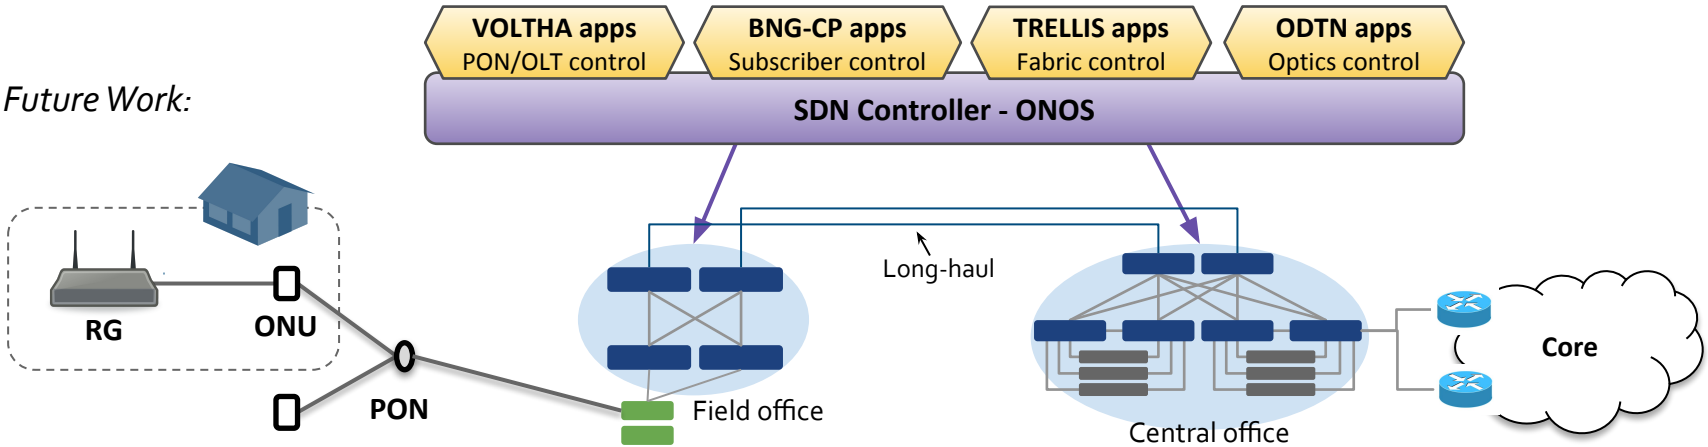

Learning More

ODTN: https://www.opennetworking.org/odtn/

P4: https://p4.org/

Stratum: https://github.com/stratum/stratum

SEBA: https://wiki.opencord.org/display/CORD/SEBA

TestVectors: https://github.com/stratum/testvectors

TestVectors Runner: https://github.com/stratum/testvectors-runner

Trellis: https://docs.trellisfabric.org/

Next steps

- Download the code (and send a pull request)
- Join the mailing lists, Slack workspaces, or TST calls
- Try a tutorial

STRATUM

Production Grade, P4 Programmable, SDN Fabric

OCP Global Summit 2020

Production Grade, P4 Programmable, SDN Fabric

SDN Applications

Network OS
(Control Plane)

Switch OS
(Data Plane)

New Features and Capabilities

1. Stratum's Broadcom implementation has been extended to support **Trellis**
2. Stratum now supports Edgecore's **Cassini** packet optical transponder
3. Trellis and **ODTN** are being run on the same ONOS cluster for the first time
4. Stratum now offers users a choice between Broadcom's **OpenNSA** and SDKLT for Tomahawk switches

Stratum on Cassini

- First open source operating system for Cassini, a packet optical transponder with a merchant silicon switching ASIC
- Optical capabilities managed through **OpenConfig models** and **gNMI**
- Integrated as a native **Trellis spine** (L2 - L4) using Tomahawk+ ASIC

TELECOM INFRA PROJECT

Adding Support for CASSINI to Stratum

- Support for additional OpenConfig models to gNMI
- Support for transceivers via TAI integration into Platform Manager

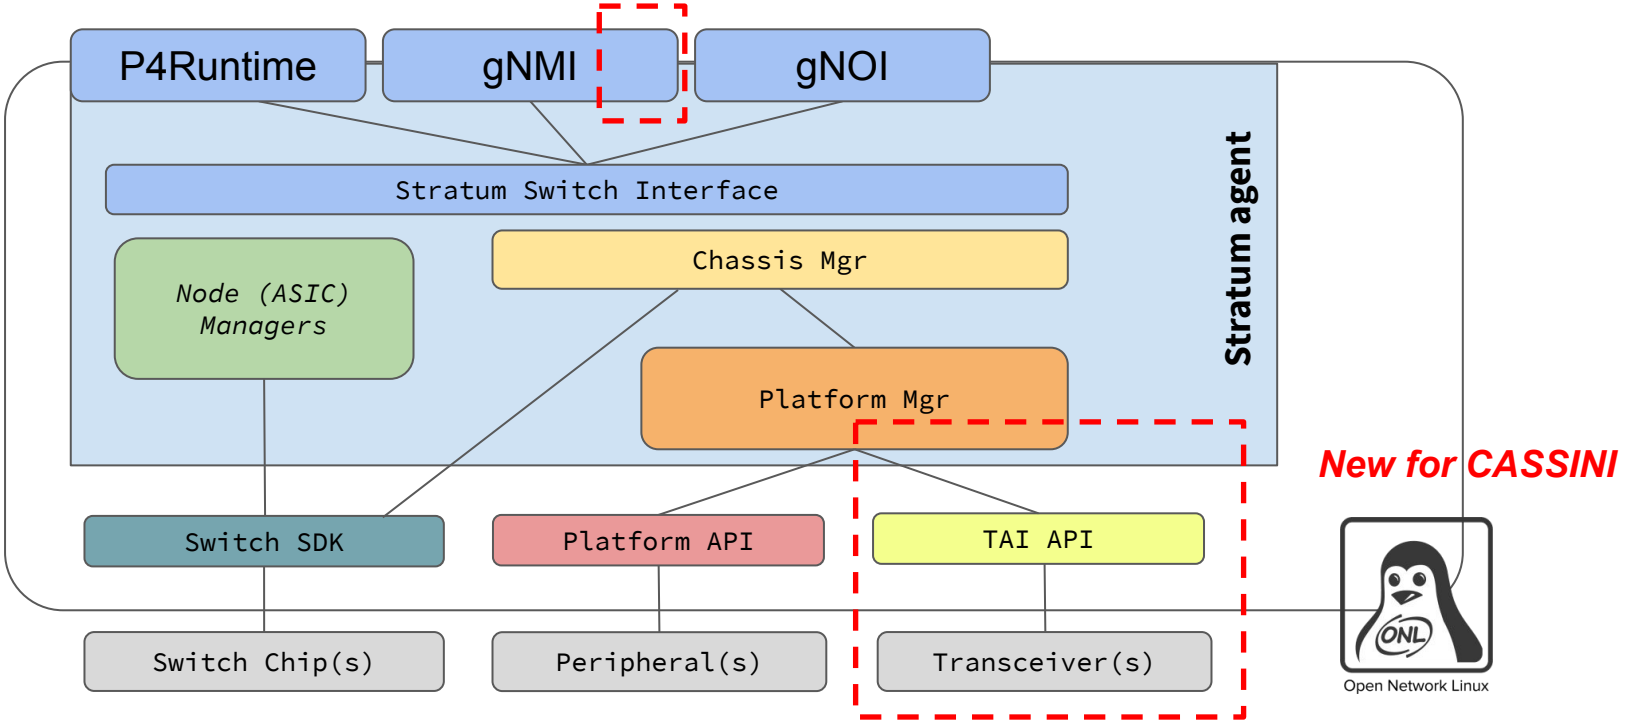

Trellis + ODTN on ONOS

- Added new Pipeconf driver for Broadcom to support Trellis
- Added new Optical behavior to support OpenConfig models over gNMI for Stratum Cassini
- First time that Trellis and ODTN apps have run together on ONOS

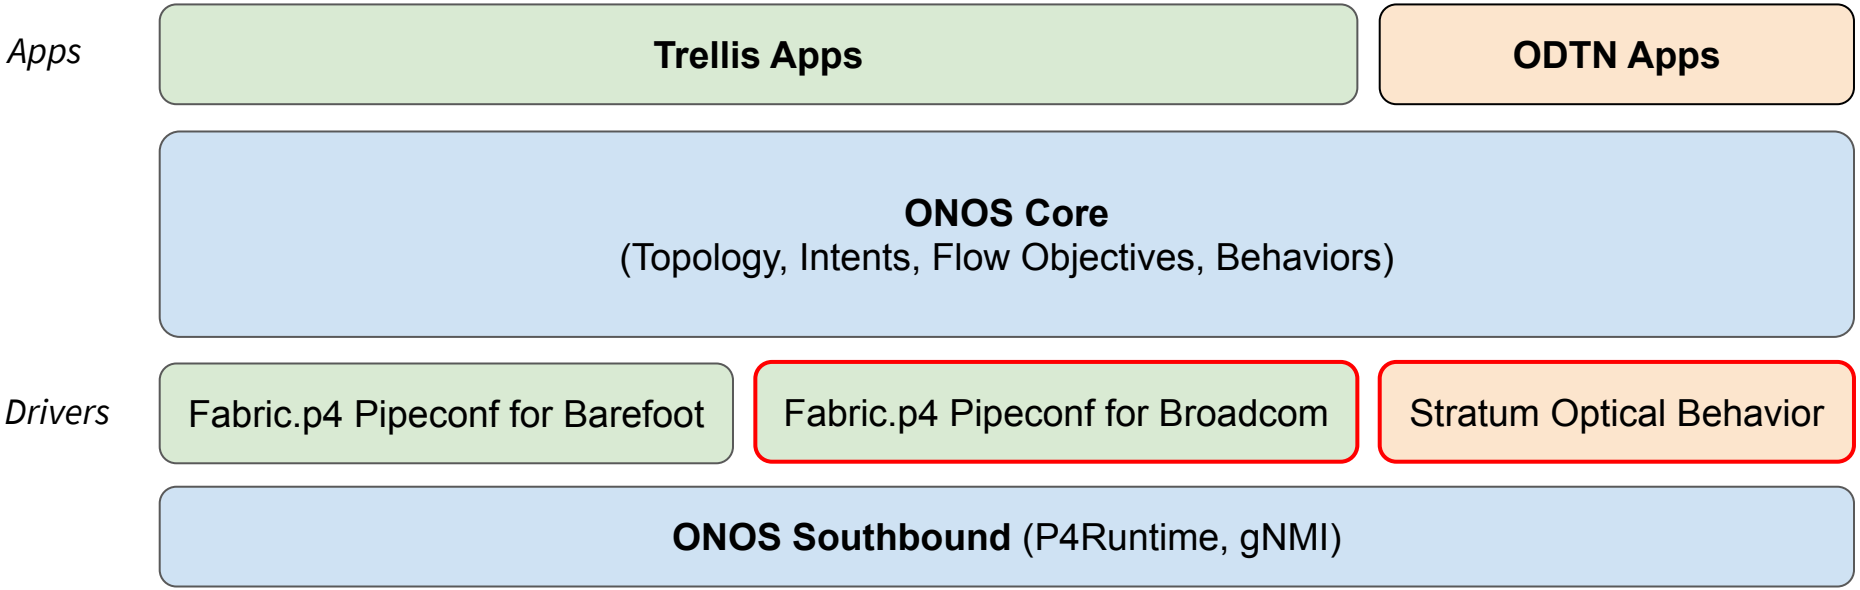

Stratum for Broadcom

- Both SDKs are open
- Choose between **OpenNSA** and **SDKLT** at build time

Bazel

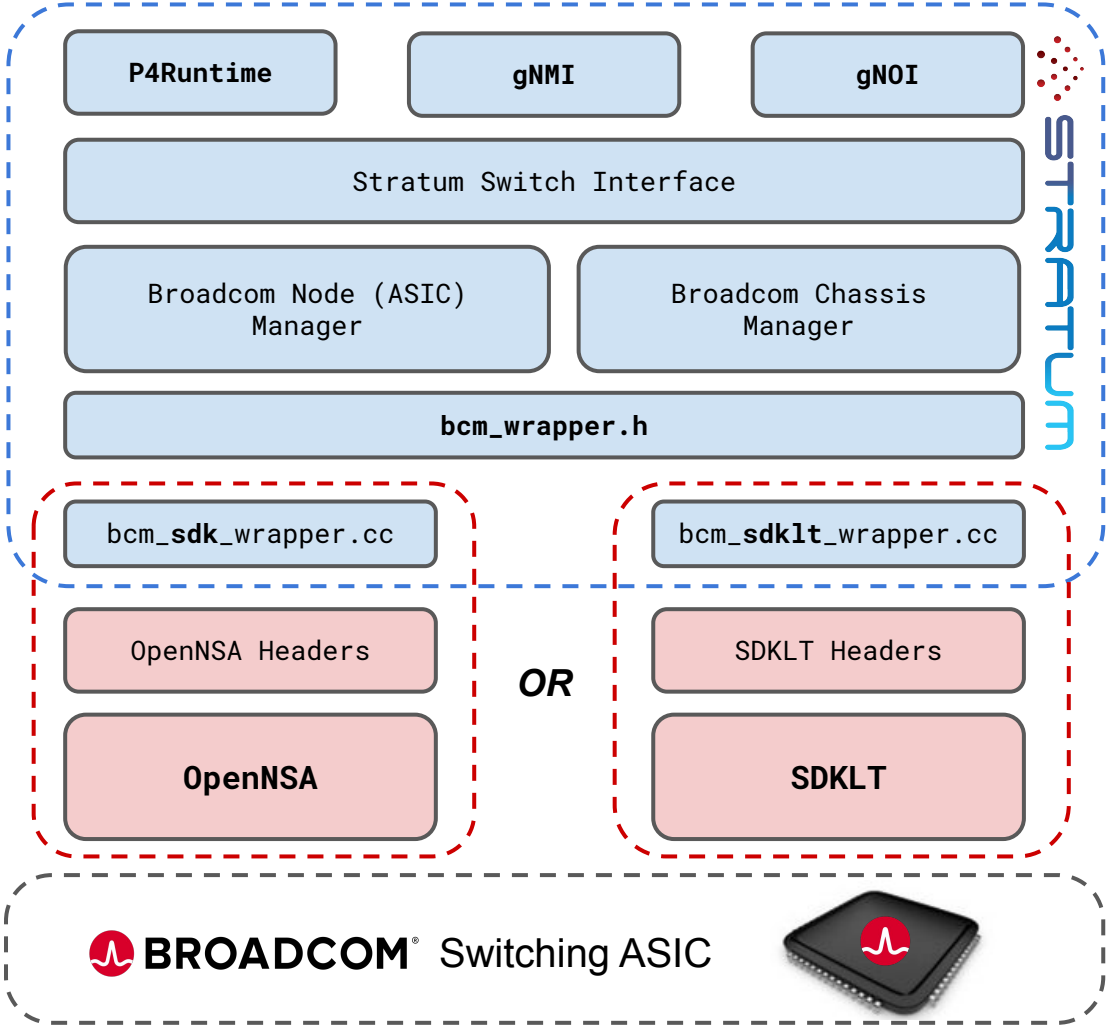

OpenNSA

- Supports:
  - Trident 2
  - Tomahawk
  - Tomahawk 2
  - Tomahawk 3
  - Qumran AX
- Familiar SDK for current users

SDKLT

- Supports:
  - Tomahawk
  - Trident 4
- Pairs with NPL
- Next-gen SDK

Stratum Support Today

Switch Vendor / Switching ASIC	DELL	DELTA	Edge-corE NETWORKS	Inventec	QCT	STORDIS
BAREFOOT NETWORKS **Tofino** *Up to 6.5 Tbps*		**AG9064v1** *64 x 100 Gbps*	**Wedge100BF-32X** *32 x 100 Gbps* **Wedge100BF-65X** *65 x 100 Gbps*	**D5054** *6 x 100 Gbps + 48 x 25 Gbps*		**BF6064X** *64 x 100 Gbps* **BF2556X** *8 x 100 Gbps + 48 x 25 Gbps*
BROADCOM **Tomahawk** *Up to 3.2 Tbps*	**Z9100** *32 x 100 Gbps*		**AS7712** *32 x 100 Gbps* **Cassini** *16 x 100 Gbps + 8 x 200 Gbps (optics)*	**D7032** *32 x 100 Gbps*	**T7032-IX1** *32 x 100 Gbps*	*New*
BROADCOM **Trident II** *Up to 1.28 Tbps*			**AS6712** *32 x 40 Gbps*	*New*		

Near-term future platforms:

- Additional Broadcom StrataXGS platforms (Trident 3, Tomahawk 2, Tomahawk 3) via OpenNSA

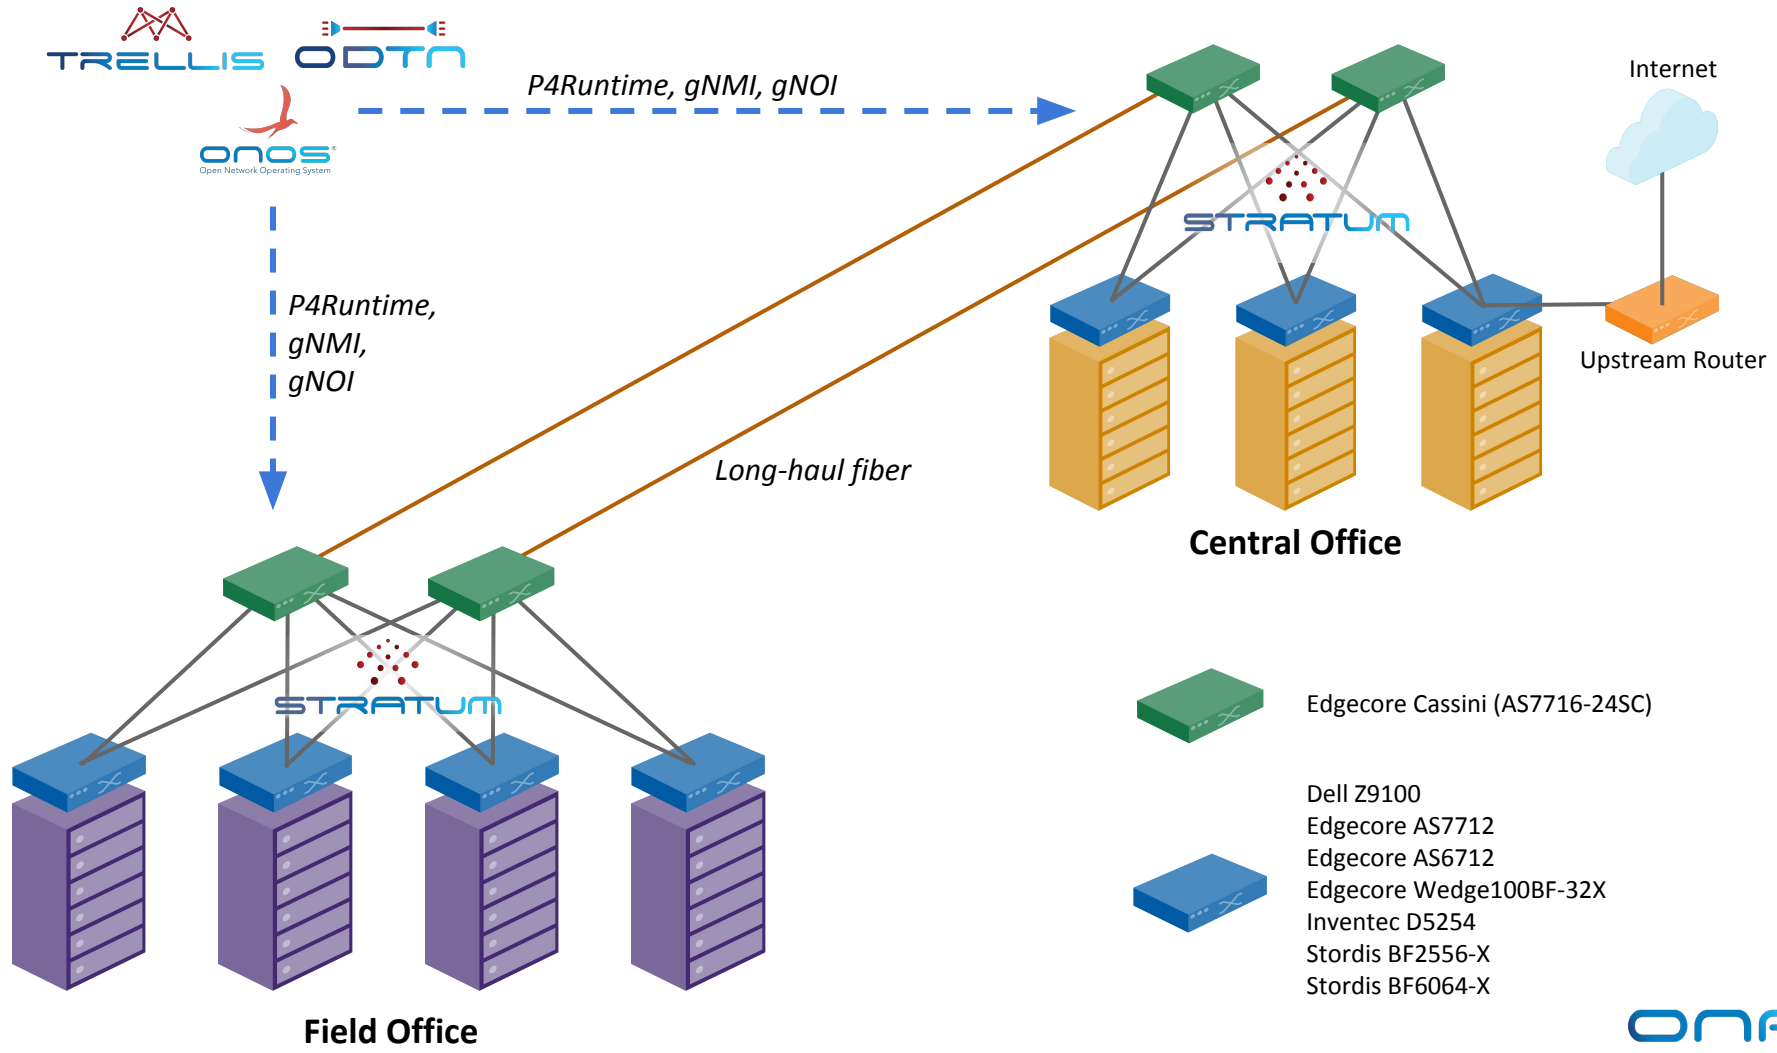
TRELLIS
ODTN
ONOS
Open Network Operating System
P4Runtime, gNMI, gNOI
P4Runtime,
gNMI,
gNOI
Long-haul fiber
STRATUM
STRATUM
Internet
Upstream Router
Central Office
Field Office
Edgecore Cassini (AS7716-24SC)
Dell Z9100
Edgecore AS7712
Edgecore AS6712
Edgecore Wedge100BF-32X
Inventec D5254
Stordis BF2556-X
Stordis BF6064-X

Demo Topology

Two racks (*each its own DC*)

- *Field office*
- *Central office*

Connected by two long-haul DWDM fiber links

- *Example of DCI & backhaul*

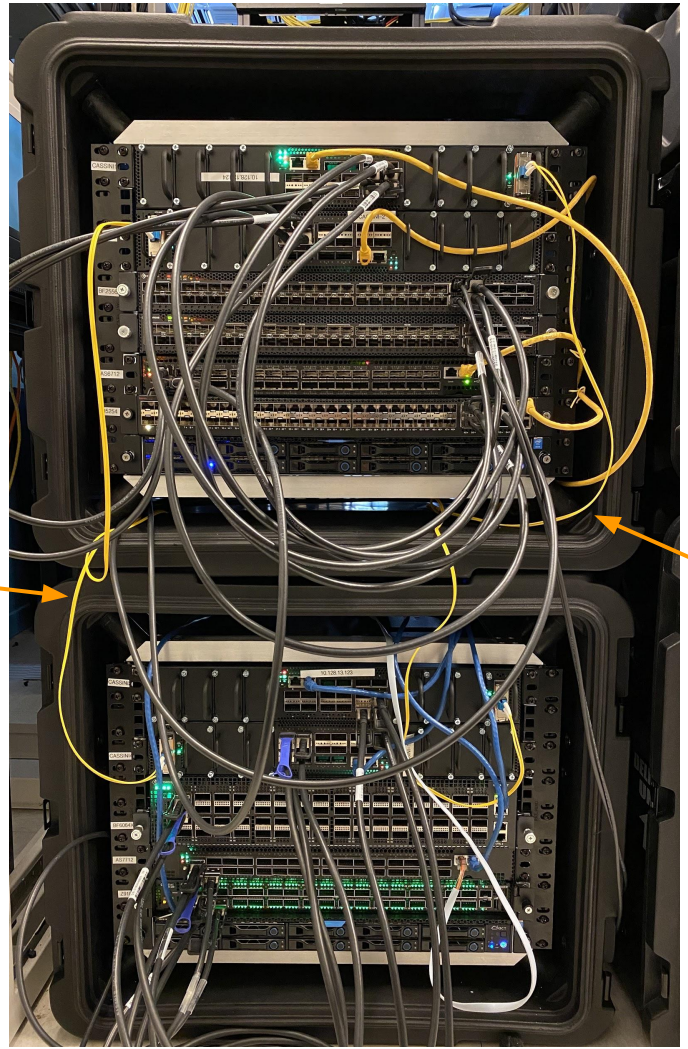

Field Office

Cassini (spine)

Cassini (spine)

BF-2556X (leaf)
BF-2556X (leaf)
AS6712 (leaf)
D5254 (leaf)
Server (hosts)

Long-haul Fiber

Long-haul Fiber

Central Office

Cassini (spine)

Cassini (spine)

BF-6064X (leaf)
AS7712 (leaf)
Z9100 (leaf)
Server (hosts)

SEBA: SDN-Enabled Broadband Access with disaggregated OLT and BNG

Stratum Webinar, March 2020

SEBA Community

SEBA: SDN Enabled Broadband Access

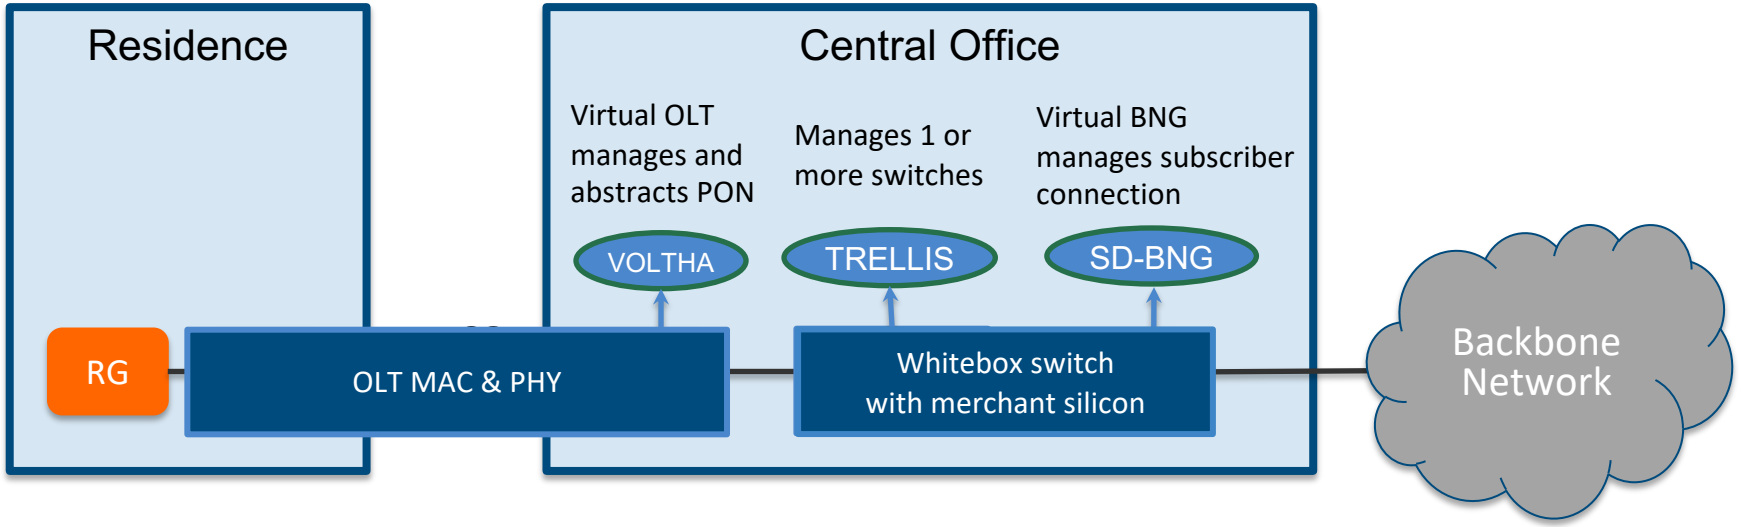

RG – Residential Gateway
ONU - Optical Network Unit
PON - Passive Optical Network
OLT – Optical Line Termination
BNG – Broadband Network Gateway

OLT disaggregation → VOLTHA

On path to production in 2020 at
AT&T, Deutsche Telekom, Türk Telekom

EMS

VLANs, IGMP, 802.1x, Mcast ...

PON MAC blades

Comput e blades

Backplane

Traditional Chassis-based Vendor OLT for PONs (Passive Optical Networks)

SDNize

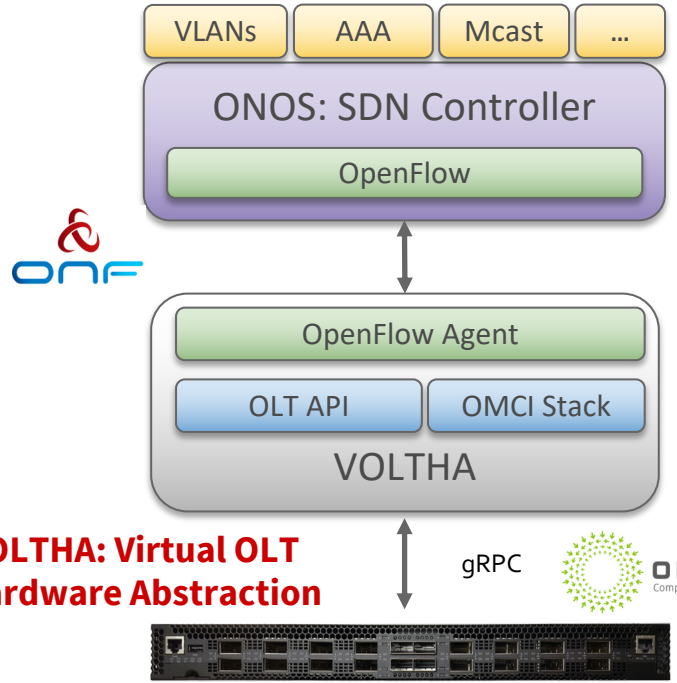

EdgeCore ASXvOLT16 (and others)

Fabric disaggregation → TRELLIS

Carrier-grade SDN fabric for edge/NFV applications

Based on whitebox switches and merchant silicon

In production at Comcast

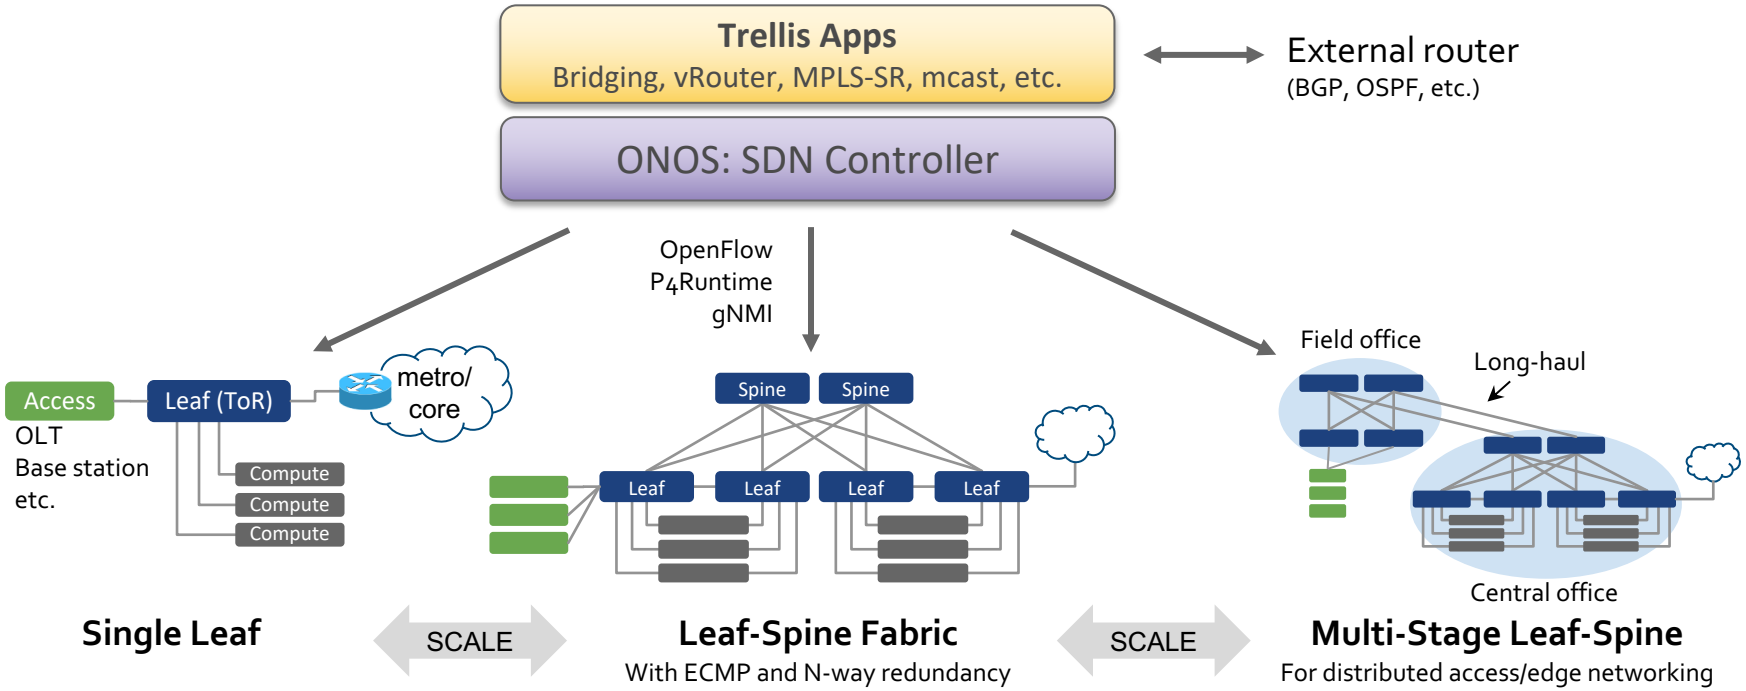

BNG disaggregation → SD-BNG

Traditional Chassis-based Vendor Broadband Network Gateway (BNG)

Specialized router:

- Subscriber termination (QinQ, PPPoE)
- Accounting
- Hierarchical QoS
- Lawful intercept
- Wholesale tunnel relay (L2TP)
- Multicast
- Routing
- Etc.

SDNize

Control Plane Apps (BNG-CP)
AAA, address assignment, mcast, LI, etc.
(or integrates with external CP via CUPS API)

ONOS: SDN Controller

User Plane (BNG-UP)

Use merchant silicon available in fabric and OLT*

OLT

EdgeCore ASXvOLT16
Broadcom Qumran AX
300 Gbit/s, Deep Buffers, HQoS (WIP)

FABRIC SWITCH

EdgeCore Wedge 100BF-32X
Intel Barefoot Tofino
3.2 Tbit/s, P4-defined routing, PPPoE term, accounting, etc.

*Design based on DT Access 4.0 project
Other silicon options are possible
(Qumran 2C, SmartNICs, etc.)

SEBA 2.0 Exemplar Architecture

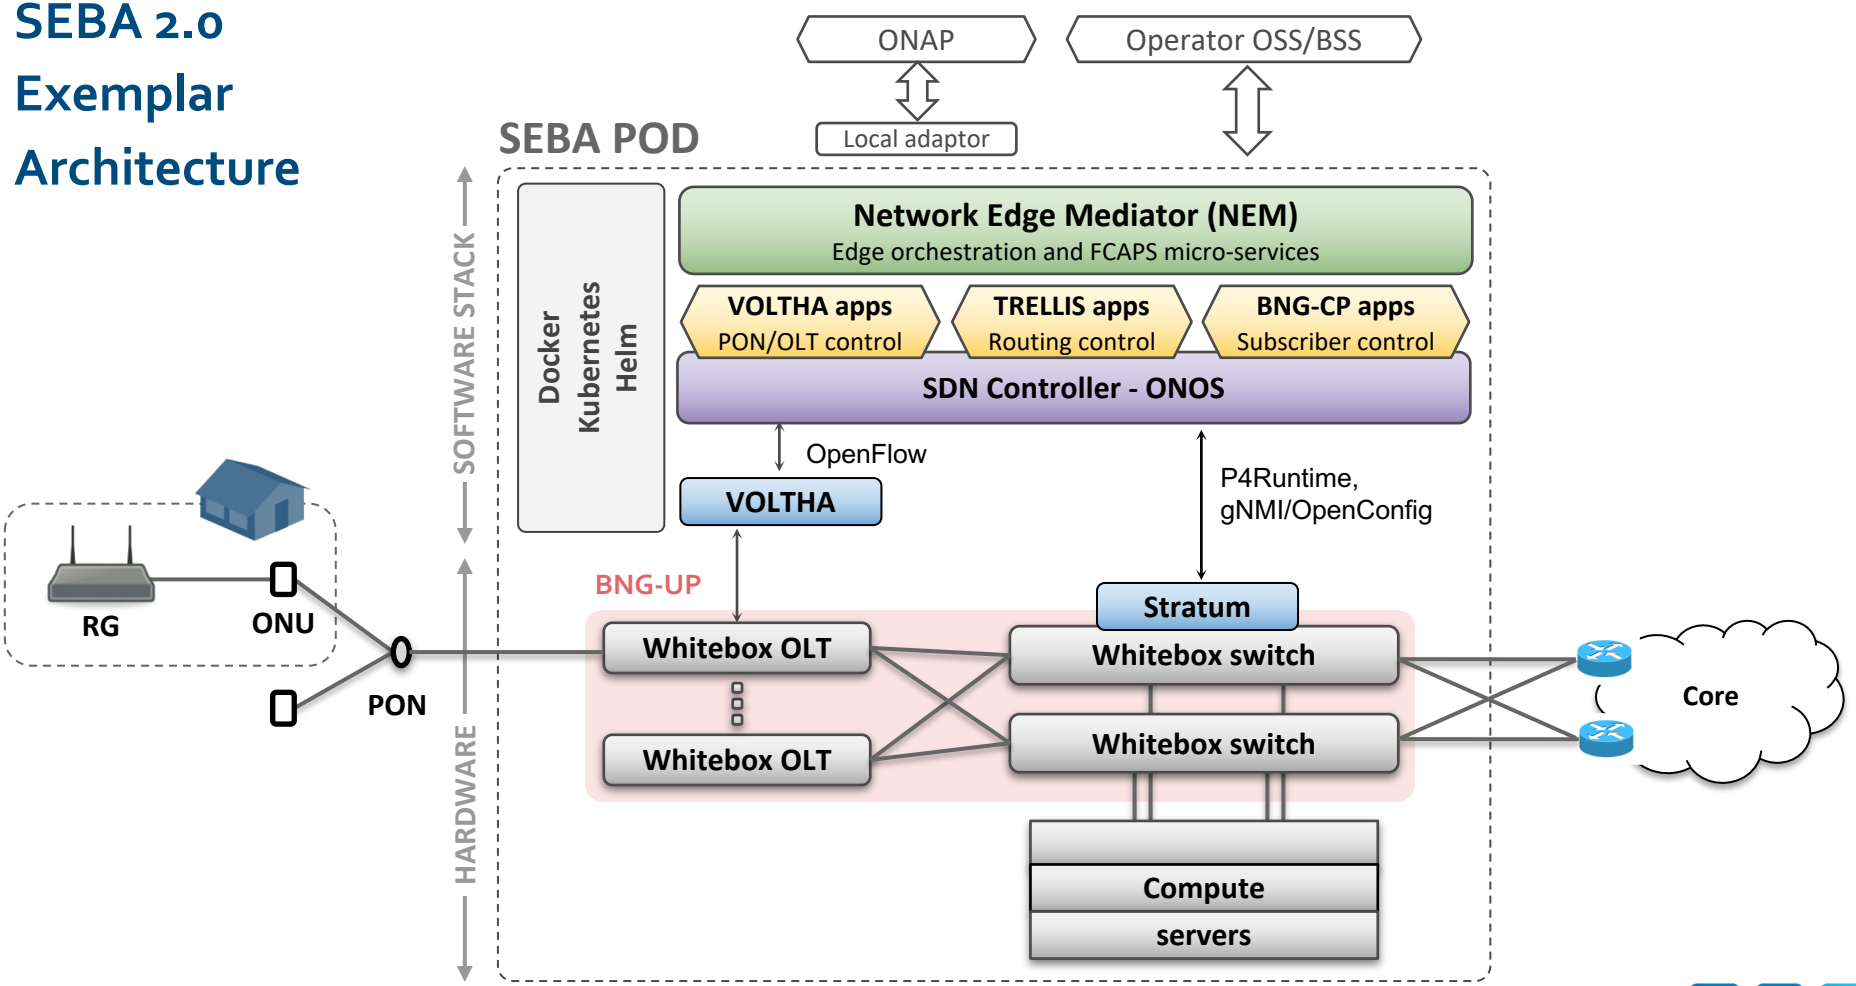

Demo setup

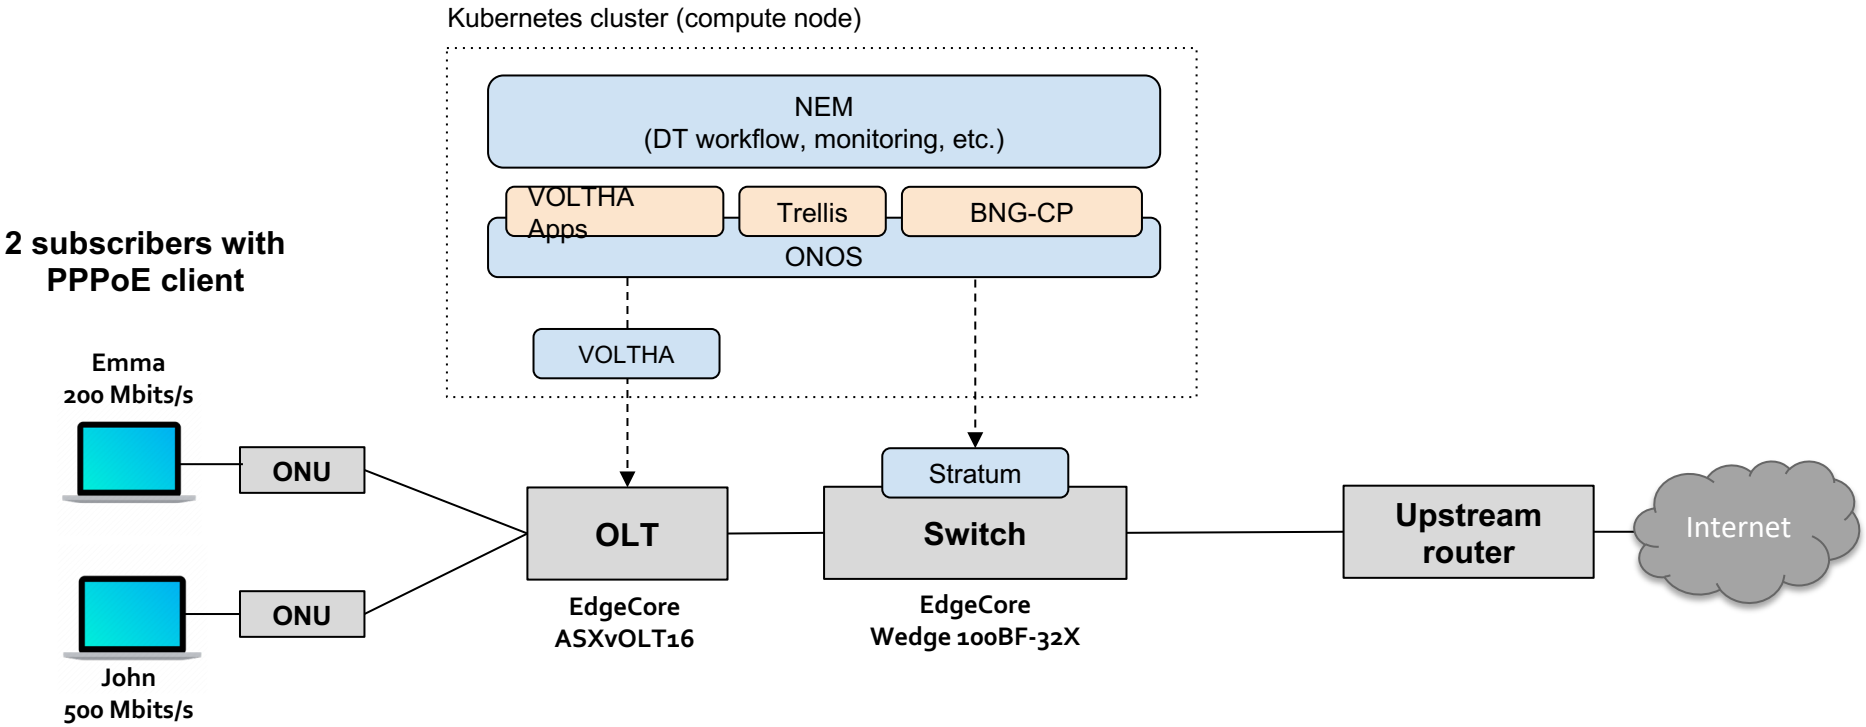
Kubernetes cluster (compute node)
NEM
(DT workflow, monitoring, etc.)
VOLTHA Apps
Trellis
BNG-CP
ONOS
VOLTHA
2 subscribers with PPPoE client
Emma
200 Mbits/s
ONU
John
500 Mbits/s
ONU
OLT
EdgeCore
ASXvOLT16
Stratum
Switch
EdgeCore
Wedge 100BF-32X
Upstream router
Internet

Demo setup

1 - Provision new subscriber

Operator pushes configuration using NEM northbound APIs

ONU whitelist
Subscriber info (c-tag, s-tag, BW profiles, …)
etc.

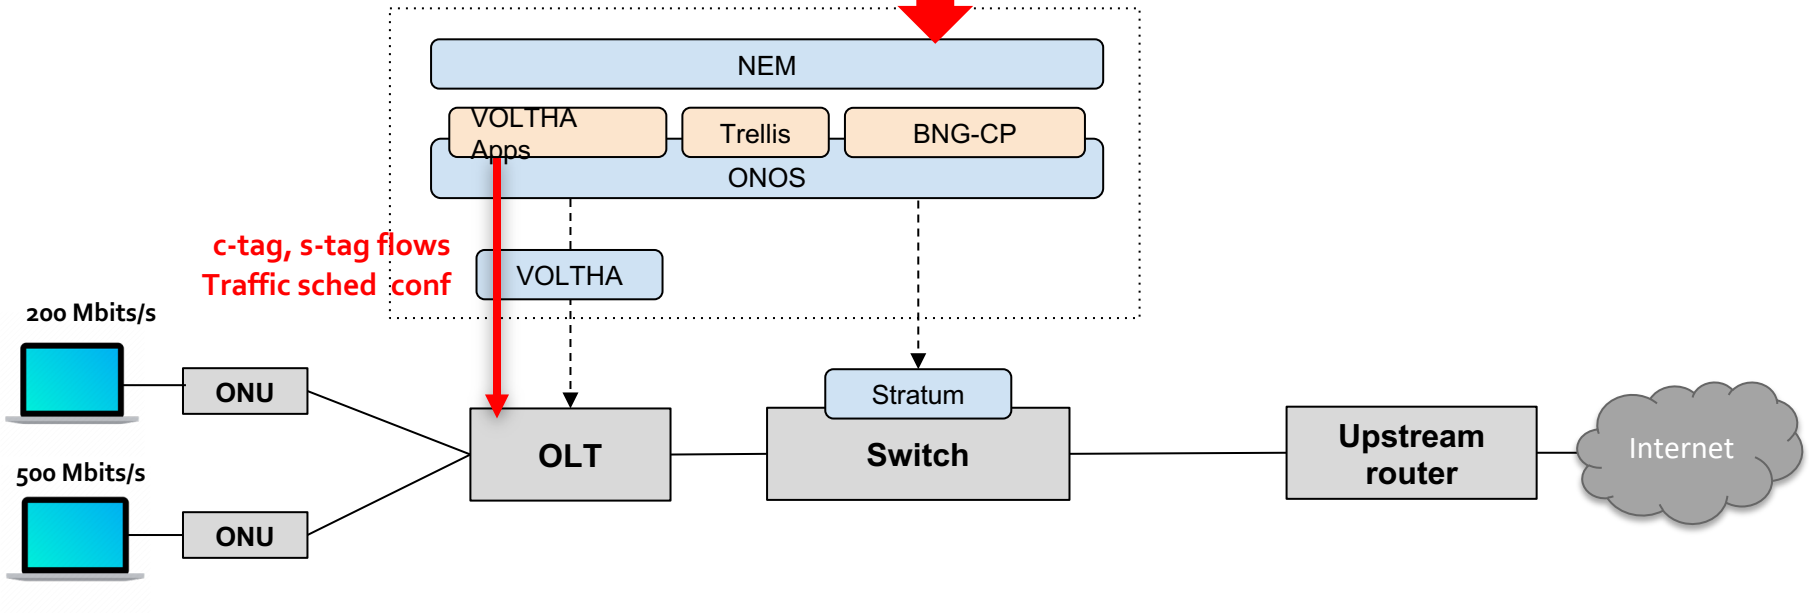

Demo setup

2 - Subscriber authentication

BNG-CP handles packets for PPPoE session discovery, authentication, and address assignment

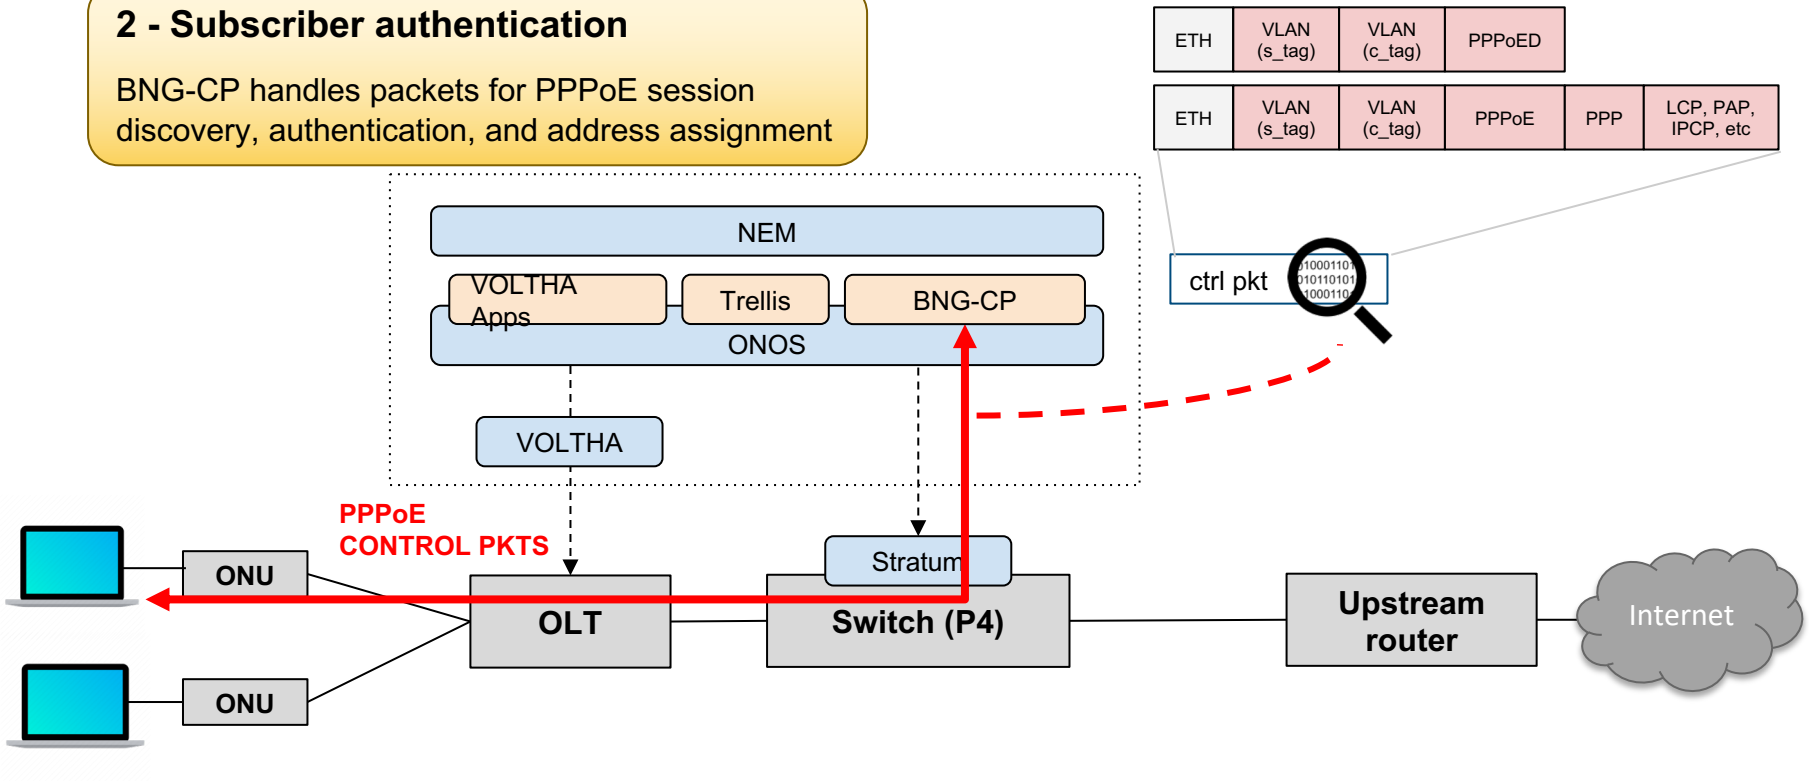

Demo setup

3 - BNG user plane setup

ONOS apps write P4Runtime entries to terminate and route PPPoE data packets to/from the subscriber

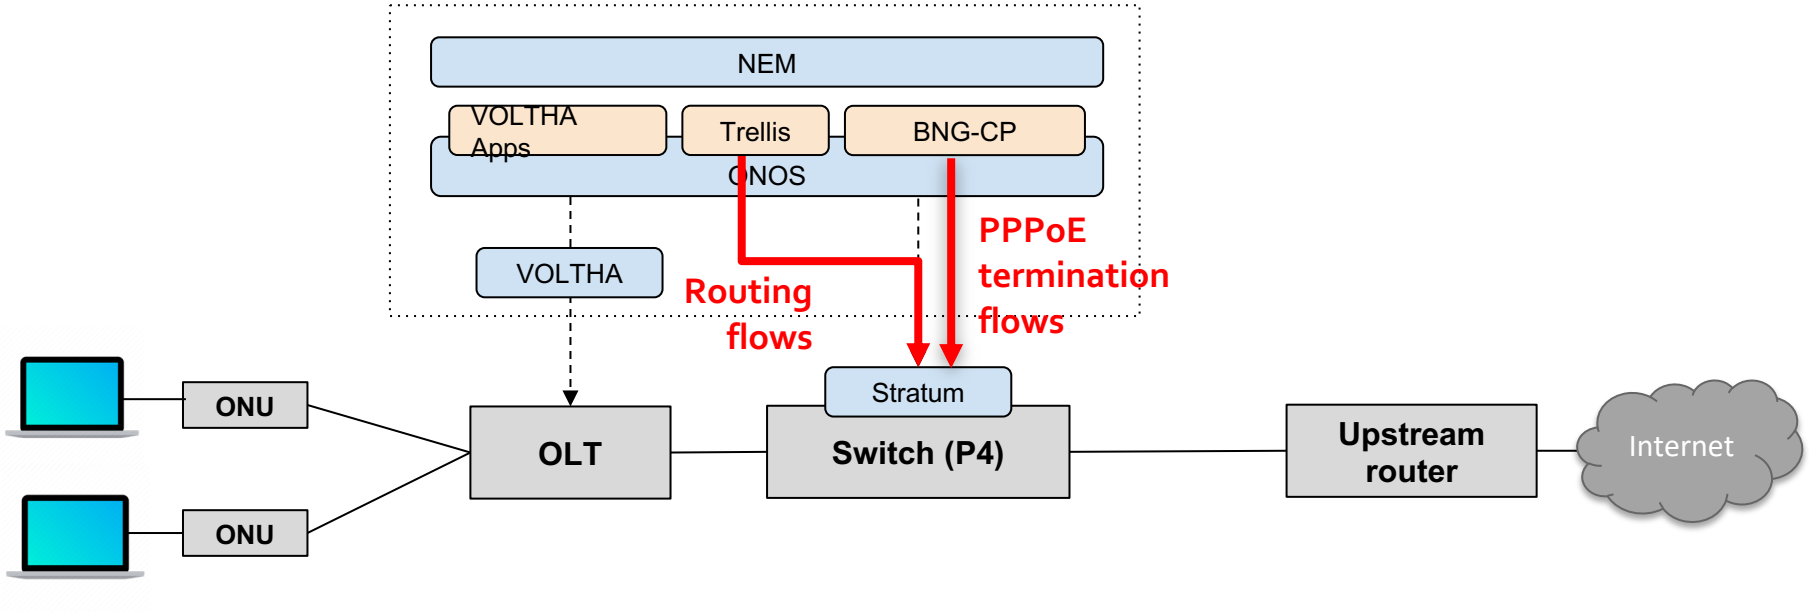

Demo setup

4 - Subscriber is connected

P4 switch performs termination and routing of PPPoE data packets

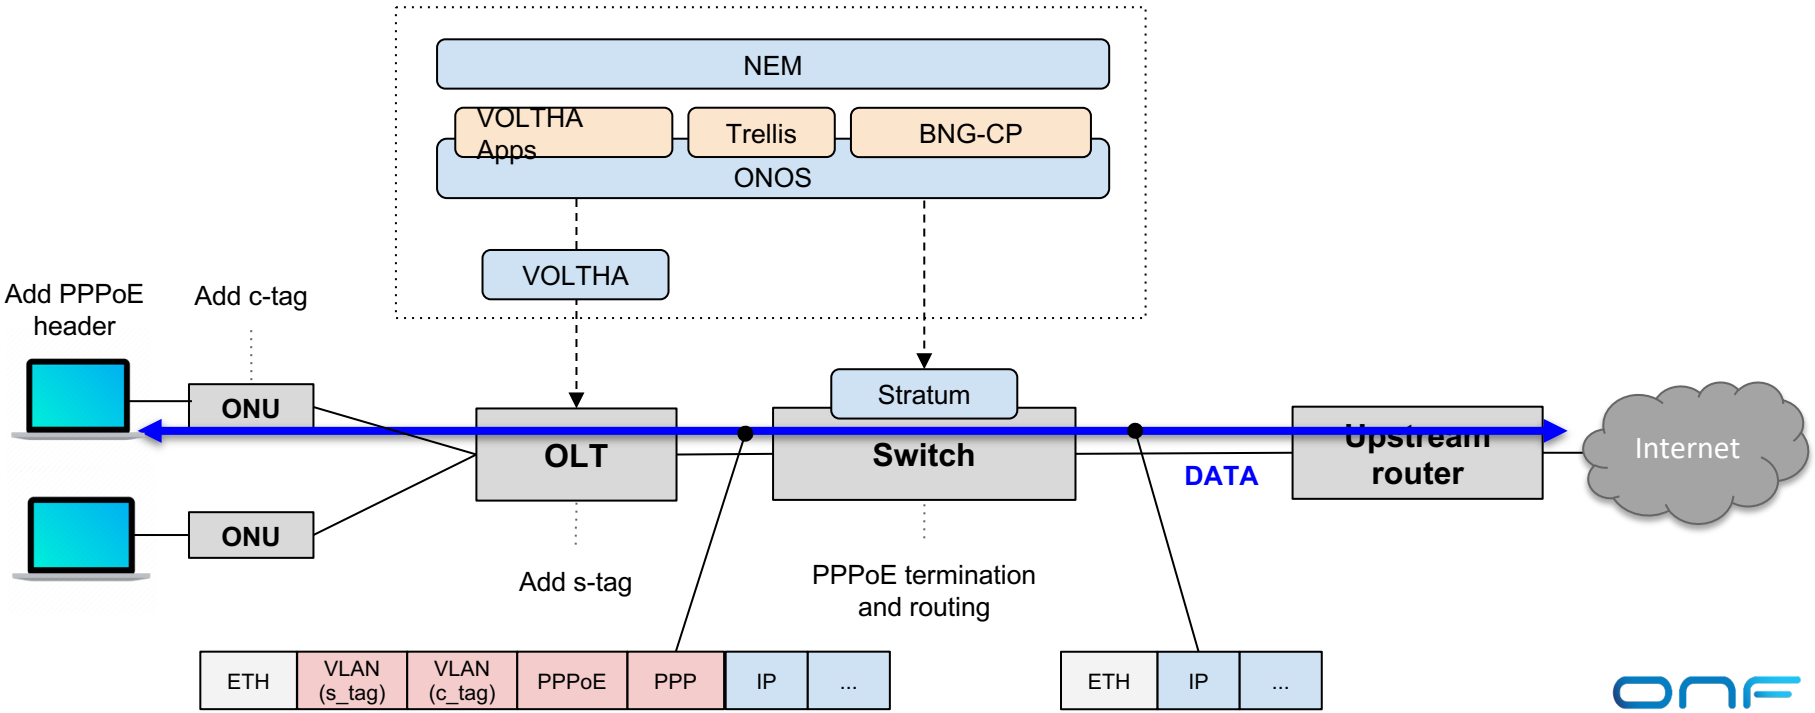

Demo setup

Accounting Dashboard

Stats
Alarms
Events

GRAFANA
KIBANA
PROMETHEUS
ELK STACK
KAFKA
XOS

NEM
VOLTHA Apps
Trellis
BNG-CP
ONOS
VOLTHA

BNG user plane counters
PON counters

ONU
ONU
OLT
Stratum
Switch
Upstream router
Internet

5 - Stats collection

Collect per-subscriber counters via NEM monitoring infrastructure

Summary

- **SEBA and VOLTHA already on path to production in 2020**
  - Productization with 3 Operators - AT&T, DT and Türk Telekom
  - External BNG
  - Community Progress
    - VOLTHA and SEBA have healthy active communities
    - Brigades - a way to effectively manage community resources

- **Looking forward: SEBA 2.0 with SD-BNG**
  - Stratum and P4 enable BNG disaggregation
  - Next steps:
    - Add HQoS, and other missing features
    - Integration with standard BNG CUPS API
    - Support other merchant silicon options

Backup

SD-BNG with user plane functional split

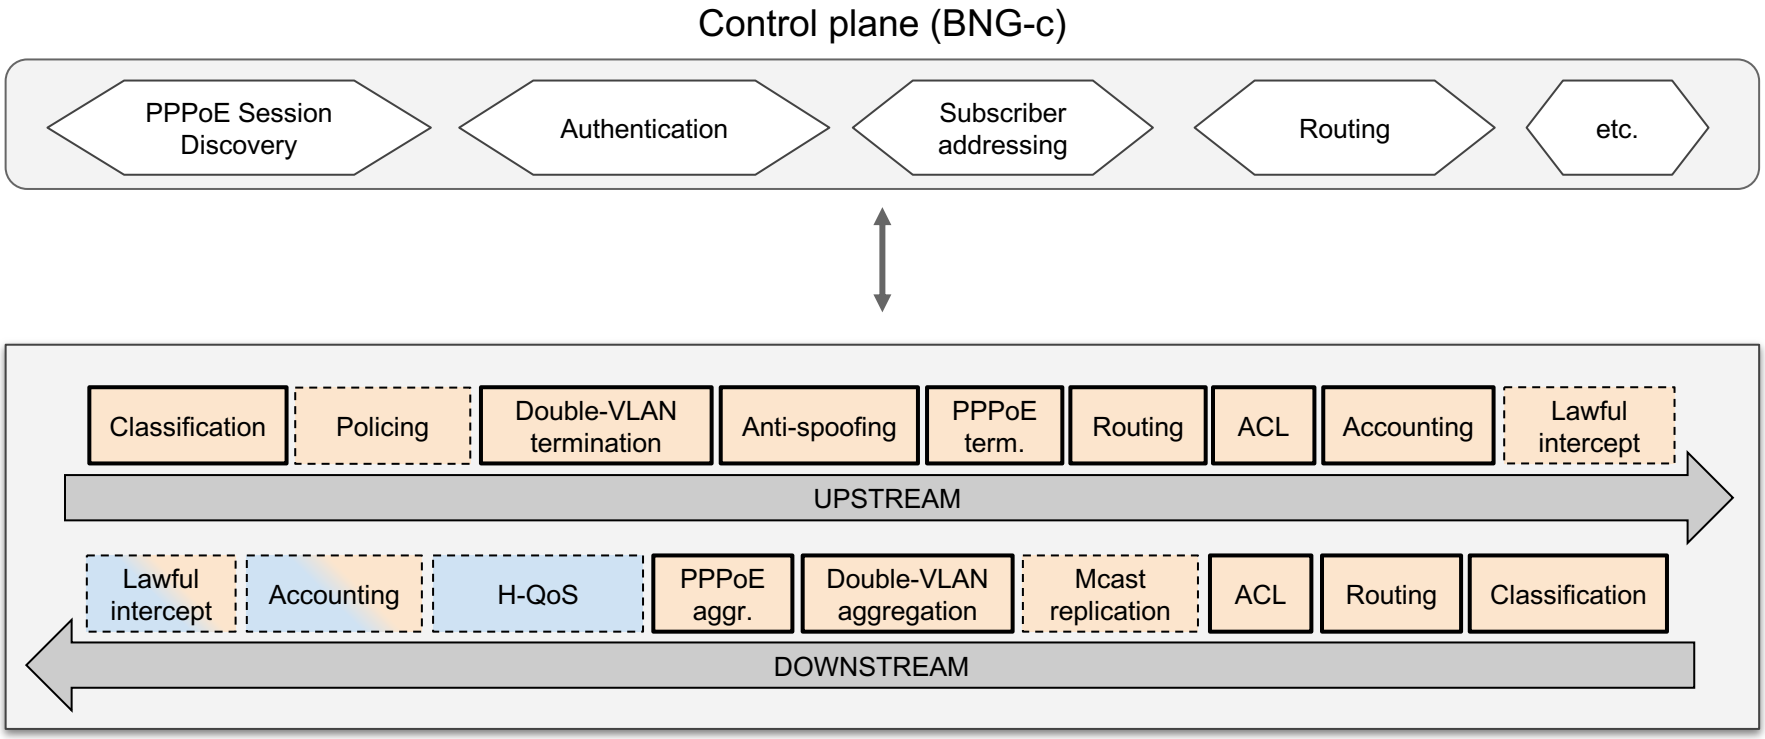

Barefoot Tofino (fabric sw)

Broadcom Qumran AX (OLT)

Demonstrated today

Work in progress

BNG disaggregation → SD-BNG with CUPS API

Traditional Chassis-based Vendor Broadband Network Gateway (BNG)

Specialized router:

- Subscriber termination (QinQ, PPPoE)
- Accounting
- Hierarchical QoS
- Lawful intercept (LI)
- Wholesale tunnel relay (L2TP)
- Multicast
- Routing
- Etc.

SDNize

External Router
BGP, OSPF, etc.

External BNG-CP
Auth, addr assignment, accounting, LI, etc.

Trellis
vRouter, mcast, etc.

BNG CUPS API

ONOS: SDN Controller

BNG-UP

Use merchant silicon available in fabric and OLT*

OLT

Broadcom Qumran (DNX)
300 Gbit/s, Deep Buffers, HQoS

FABRIC SWITCH

Intel Barefoot Tofino
3.2 Tbit/s, P4-defined routing, PPPoE term, accounting, etc.

OPEN Compute Project

*Design based on DT Access 4.0 project

SEBA 2.0 Exemplar Architecture with BNG CUPS API

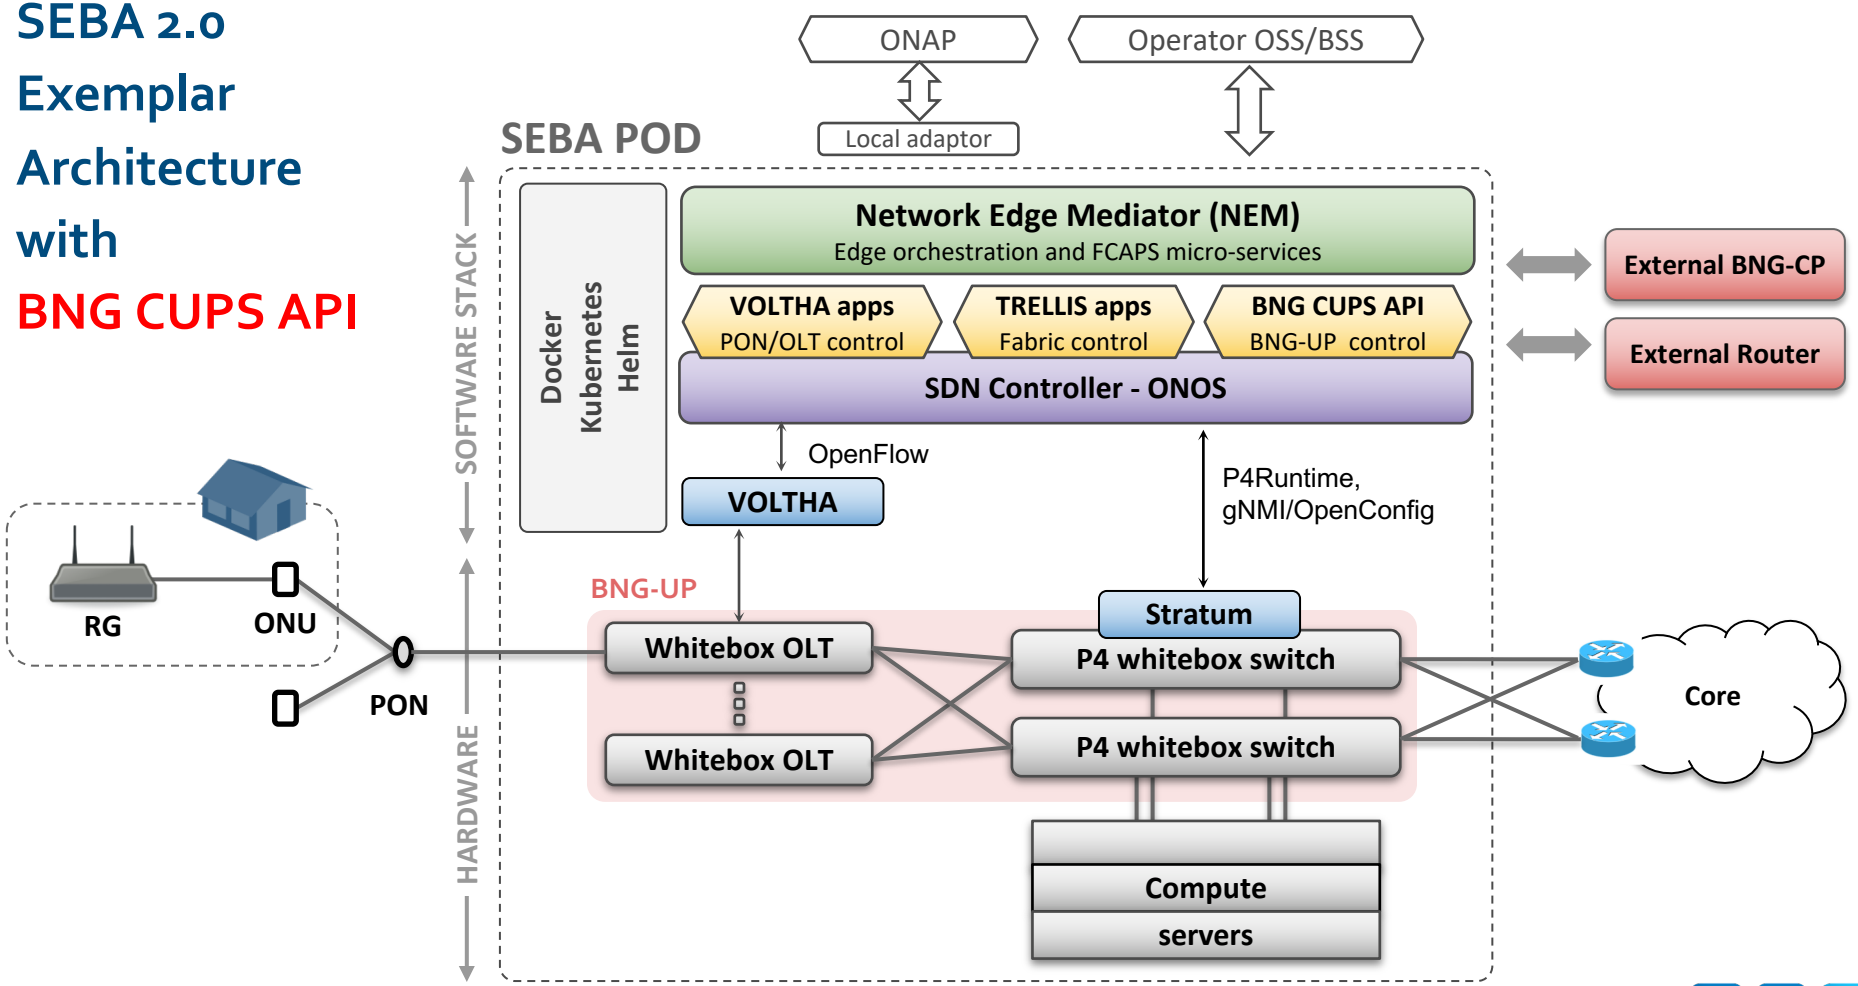

Bringing CI/CD to Switch OS Platforms

OCP Global Summit 2020

Motivation

- Networks have not kept pace with the speed of cloud development
  - They are effectively 'black box' systems that are just expected to deliver packets
  - Functionality is dictated by ASIC switch vendors, and innovation takes years to become avail (e.g VxLAN)
- With P4 and NG-SDN Interfaces, network can now be enhanced at cloud-like speed
  - Becomes a tool for innovation
- Network innovation needs to incorporate CI/CD best practices
  - Accelerate feature development and deployment of Stratum
  - Continuously certify Stratum enabled switches to work with the latest software releases

System Test with Test Vectors

Test Vectors serve as compliance tests for Stratum-based devices
They can be written **manually** or **generated automatically**

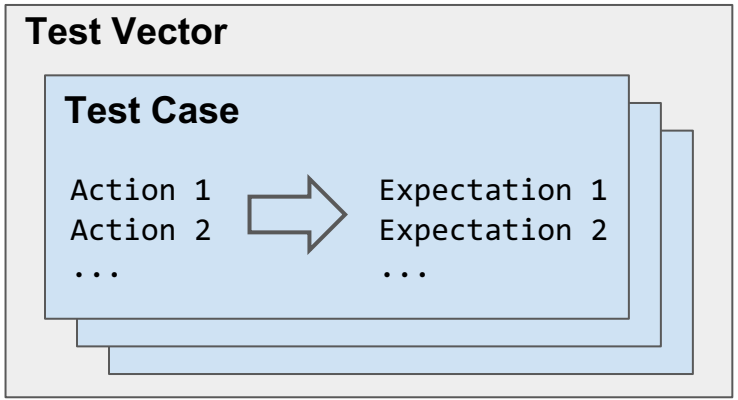

TestVectors Runner is a data-driven framework that enables users to execute Test Vectors

- Reference impl. in **golang**; supports **P4RT**/**gNMI**

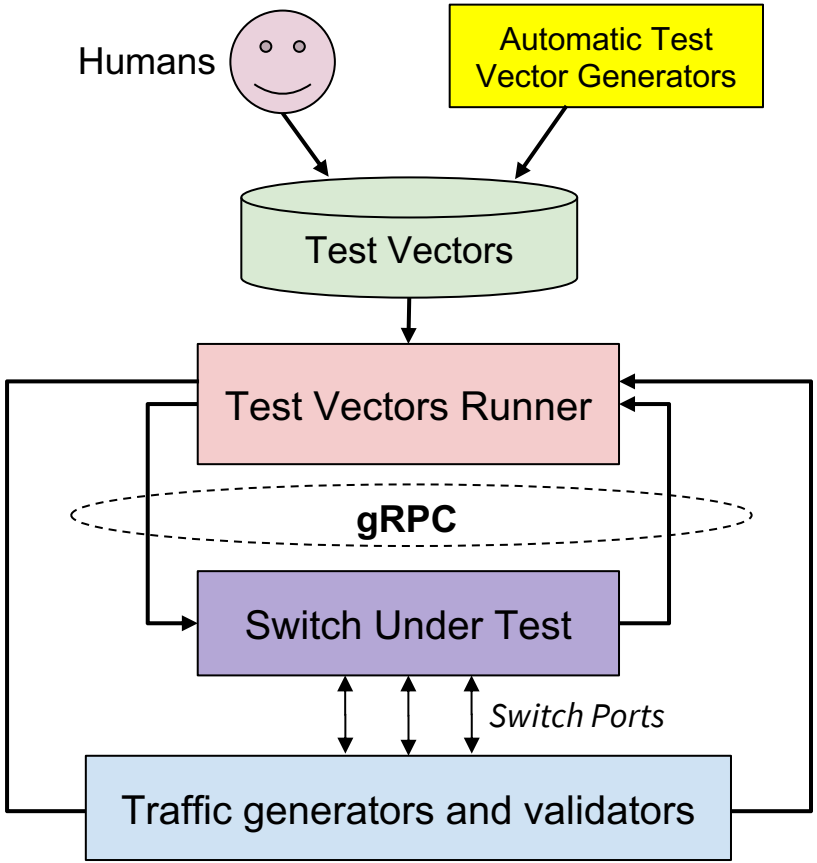

CI Workflow

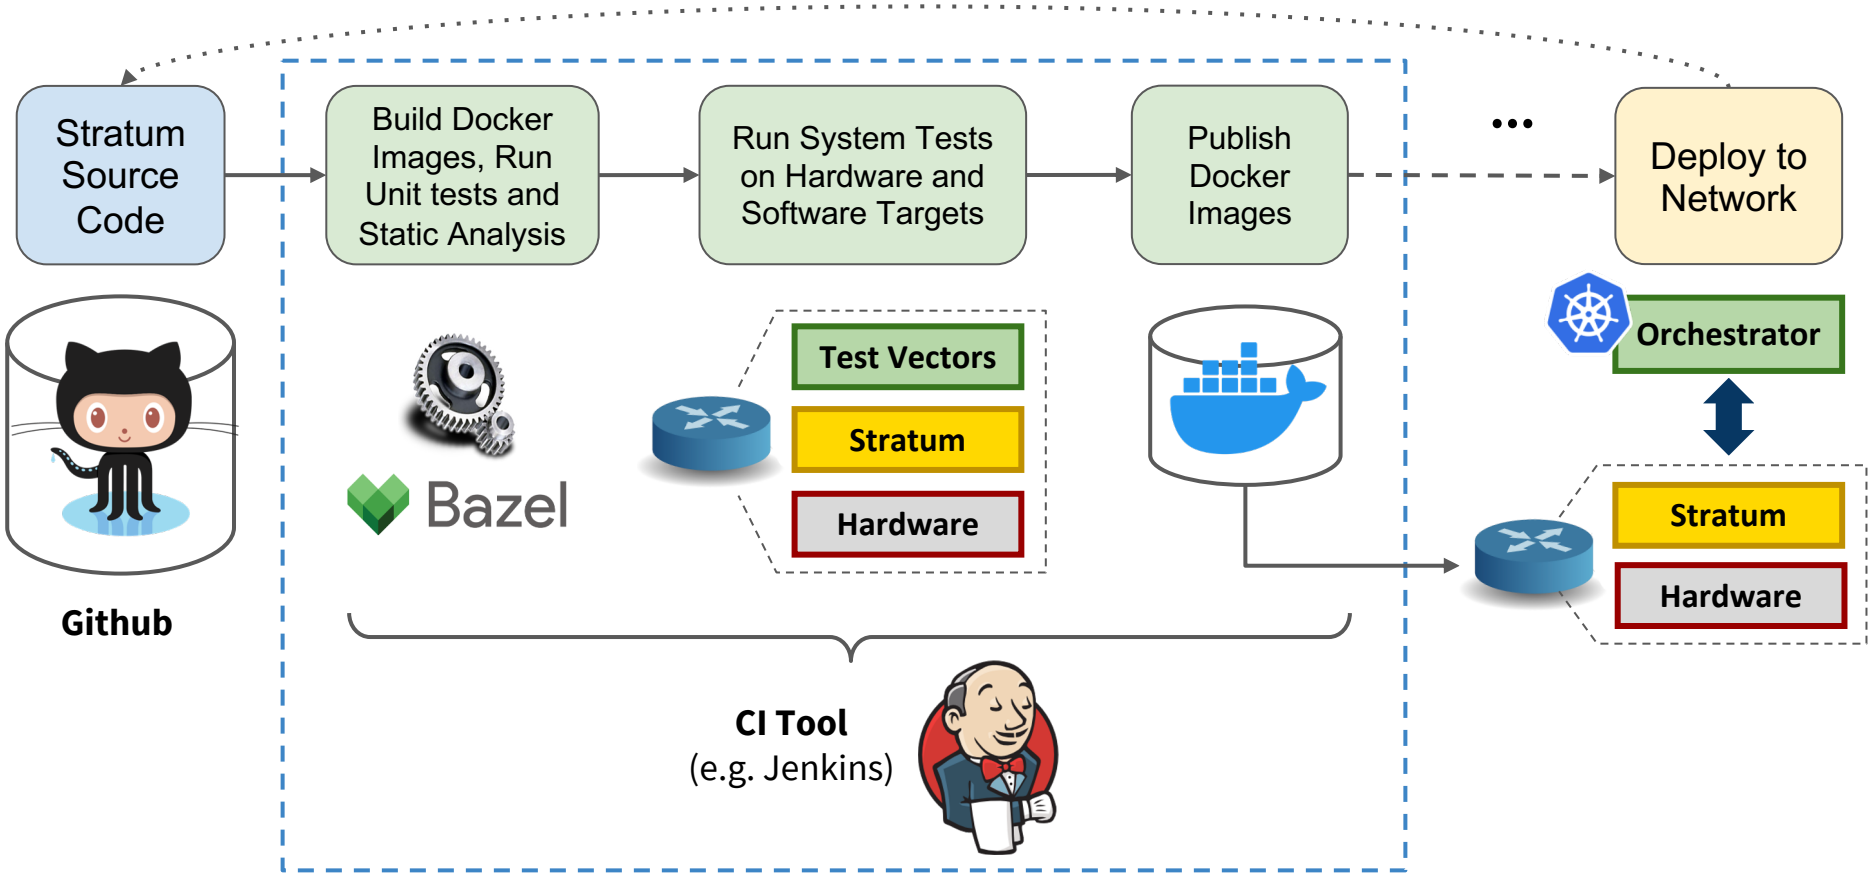
Stratum Source Code
Build Docker Images, Run Unit tests and Static Analysis
Run System Tests on Hardware and Software Targets
Publish Docker Images
...
Deploy to Network
Github
Bazel
Test Vectors
Stratum
Hardware
Orchestrator
Stratum
Hardware
CI Tool
(e.g. Jenkins)

ONF's New Continuous Certification (CC) Program

- Enable supply chain vendors to embed their products for ongoing **test and compliance certification** with ONF's open source software projects
  - Test against the latest master branch and maintain historical test data
  - Publish a list of certified devices for each software release

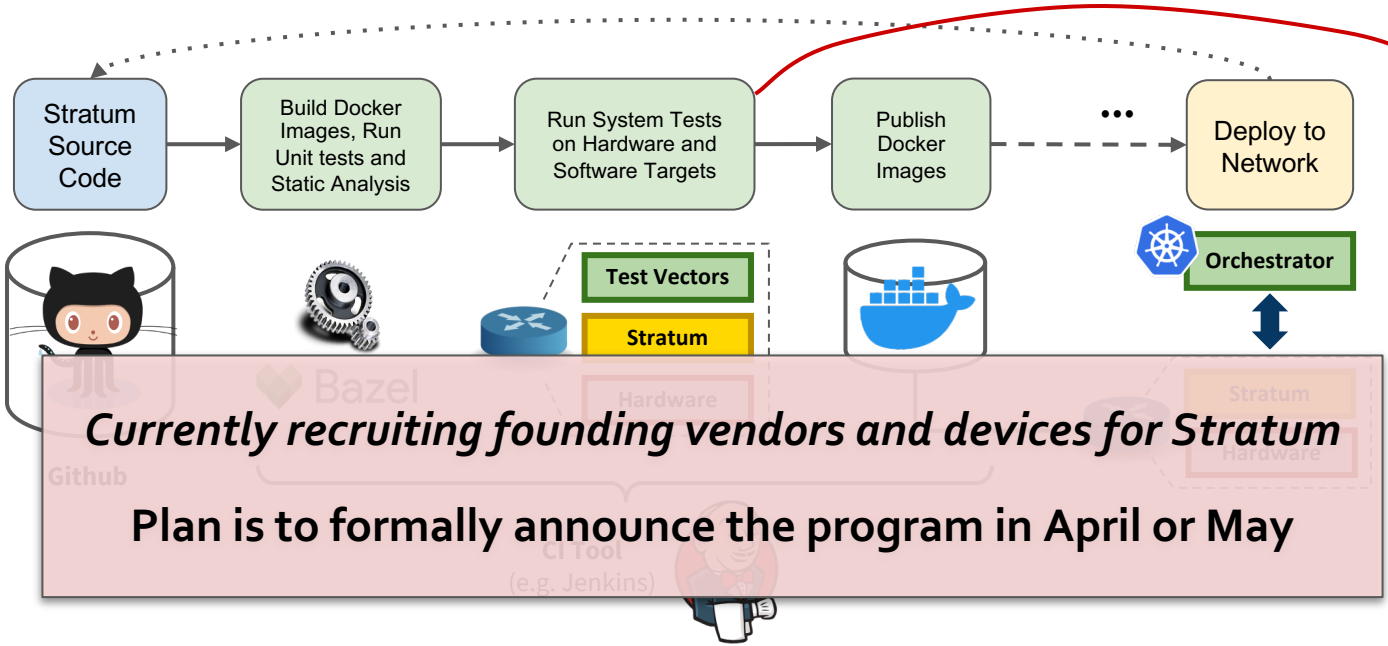

Currently recruiting founding vendors and devices for Stratum

Plan is to formally announce the program in April or May

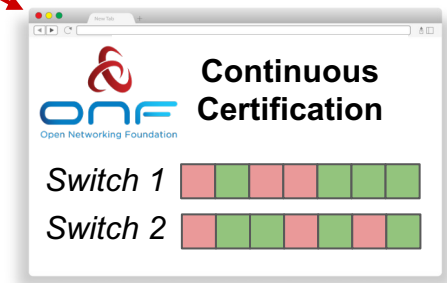

Backup

Test Vectors Implemented

- p4runtime
  - PktIoOutDirectToDataPlaneTest
  - PktIoOutToIngressPipelineAclPuntToCpuTest
  - PktIoOutToIngressPipelineAclRedirectToPortTest
  - PktIoOutToIngressPipelineL3ForwardingTest
  - PacketIoOutDirectLoopbackPortAclTest
  - PacketIoOutDirectLoopbackL3ForwardingTest
  - RedirectDataplaneToCpuACLTest
  - RedirectDataplaneToCpuNextHopTest
  - RedirectDataplaneToDataplaneTest
  - L3ForwardTest
- gnmi
  - Subscribe_Health_Indicator
  - Config_expectation_1
  - Config_expectation_2
  - …
  - Config_expectation_36
- e2e
  - SubRedirectDataplaneToDataplane

- Targets supported: bmv2, Barefoot Tofino, Broadcom Tomahawk

```
test_cases: <
  action_groups: <
    sequential_action_group: <
      actions: <
        control_plane_operation: <
          write_operation: <
            p4_write_request: <
              device_id: 1
              election_id: <
                low: 4
              >
              updates: <
                type: INSERT
                entity: <
                  table_entry: <
                    table_id: 33573106
                    match: <
                      field_id: 1
                      ternary: <
                        value: "\000\000\000\252\252\252"
                        mask: "\377\377\377\377\377\377"
                      >
                    >
                    action: <
                      action: <
                        action_id: 16832439
                      >
                    >
                    priority: 10
                  >
                >
              >
            >
          >
        >
      >
      actions: <…
      >
      actions: <…
      >
    >
    action_group_id: "ag1"
  >
  test_case_id: "insert_write"
>
```

```
test_cases: <
  expectations: <
    telemetry_expectation: <
      gnmi_subscribe_request: <
        subscribe: <
          subscription: <
            path: <
              elem: <
                name: "interfaces"
              >
              elem: <
                name: "interface"
                key: <
                  key: "name"
                  value: "veth3"
                >
              >
              elem: <
                name: "state"
              >
              elem: <
                name: "counters"
              >
              elem: <
                name: "out-unicast-pkts"
              >
            >
            mode: SAMPLE
            sample_interval: 3000
          >
          updates_only: true
        >
      >
      action_group: <
        sequential_action_group: <
          actions: <…
          >
        >
        action_group_id: "ag1"
      >
      gnmi_subscribe_response: <…
      >
      gnmi_subscribe_response: <…
      >
    >
    expectation_id: "e1"
  >
  expectations: <
    data_plane_expectation: <…
    >
    expectation_id: "e2"
  >
  test_case_id: "subscribe"
>
```

```
test_cases: <
  action_groups: <
    sequential_action_group: <
      actions: <
        control_plane_operation: <
          write_operation: <
            p4_write_request: <
              device_id: 1
              election_id: <
                low: 4
              >
              updates: <
                type: DELETE
                entity: <
                  table_entry: <
                    table_id: 33572104
                    match: <
                      field_id: 1
                      exact: <
                        value: "\000\000"
                      >
                    >
                    match: <
                      field_id: 2
                      lpm: <
                        value: "\n\002\000\000"
                        prefix_len: 16
                      >
                    >
                    action: <
                      action_profile_member_id: 1
                    >
                  >
                >
              >
            >
          >
        >
      >
      actions: <…
      >
      actions: <…
      >
    >
    action_group_id: "ag2"
  >
  test_case_id: "delete_write"
>
```

Test Vector Runner Architecture

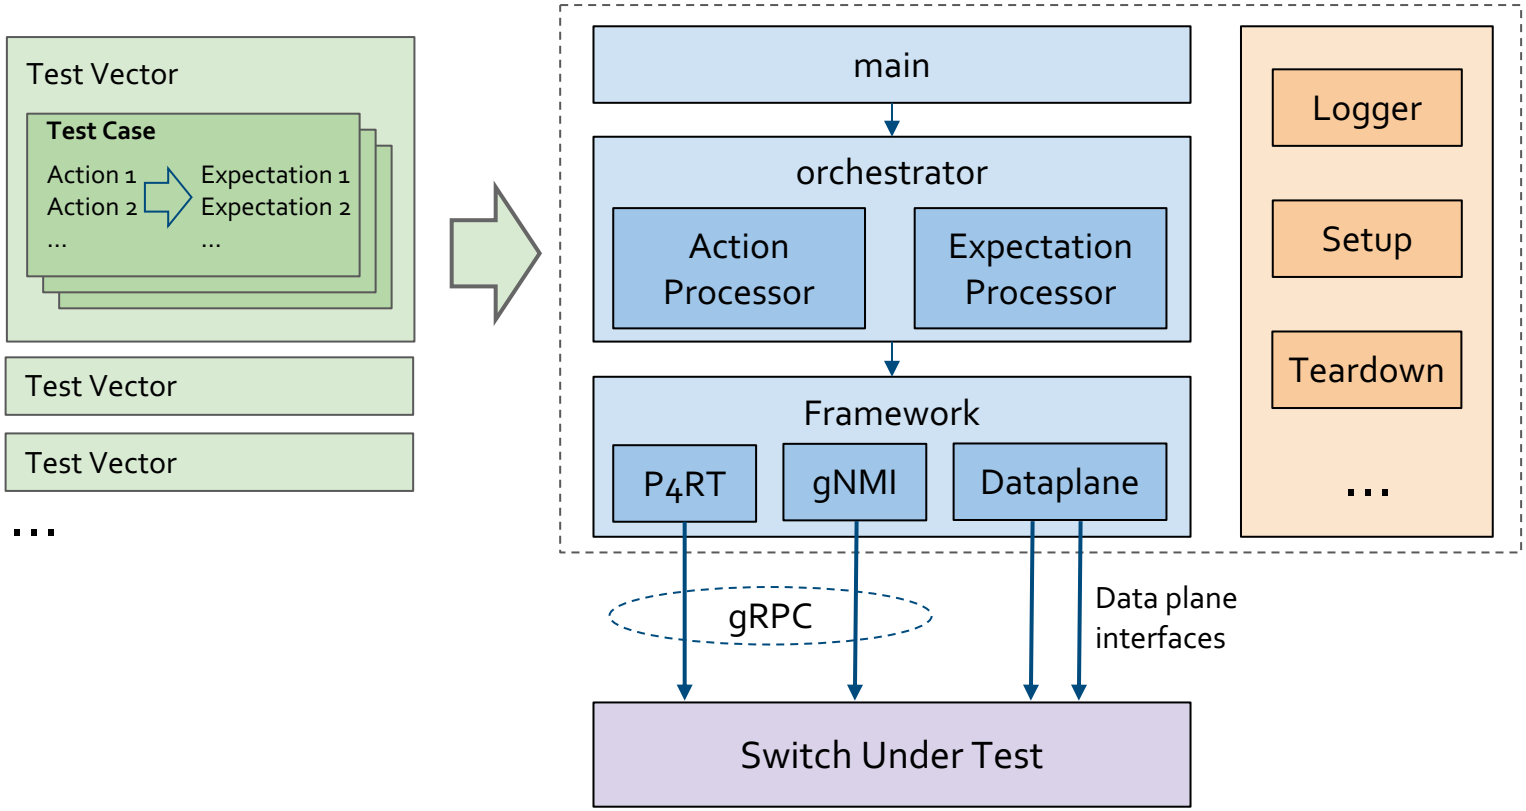

Deployment Scenarios

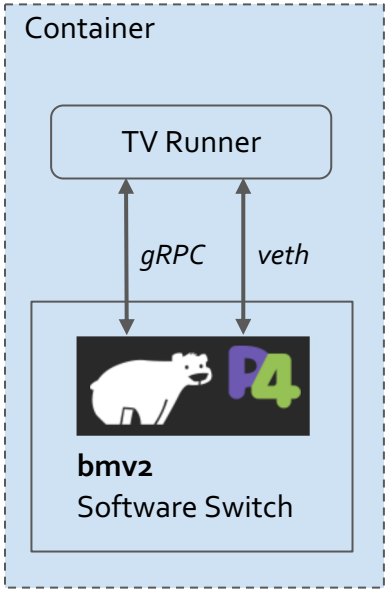

Container Mode

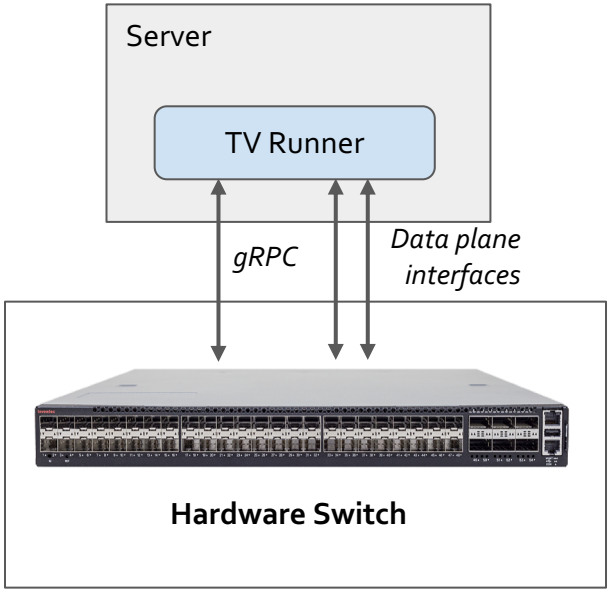

Direct Mode

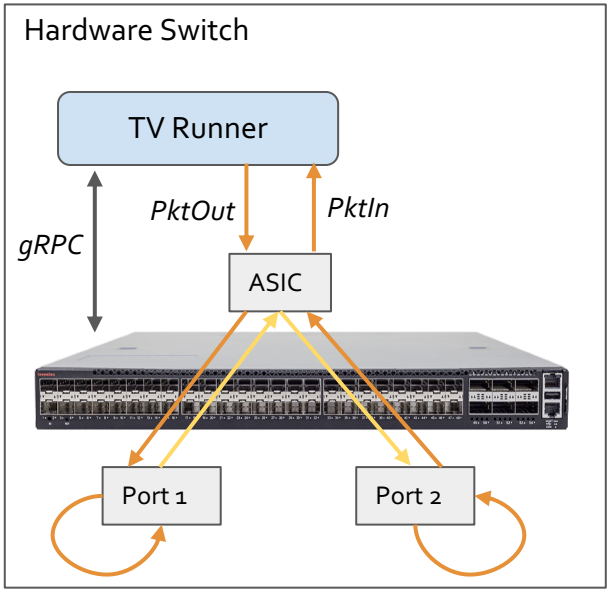

Loopback Mode

10