



OPEN SOURCE NETWORKING DAYS


Productize programmable network infrastructure

Yi Tseng

MTS, Open Networking Foundation



Scan me to join live Q&A



Or join from the link:
<https://bit.ly/osn-days-qa>



- An overview of Aether project
- Aether edge P4-based disaggregated UPF
- Productize programmable network infrastructure

Yi Tseng

Member of Technical Staff
Open Networking Foundation

2017: Intern

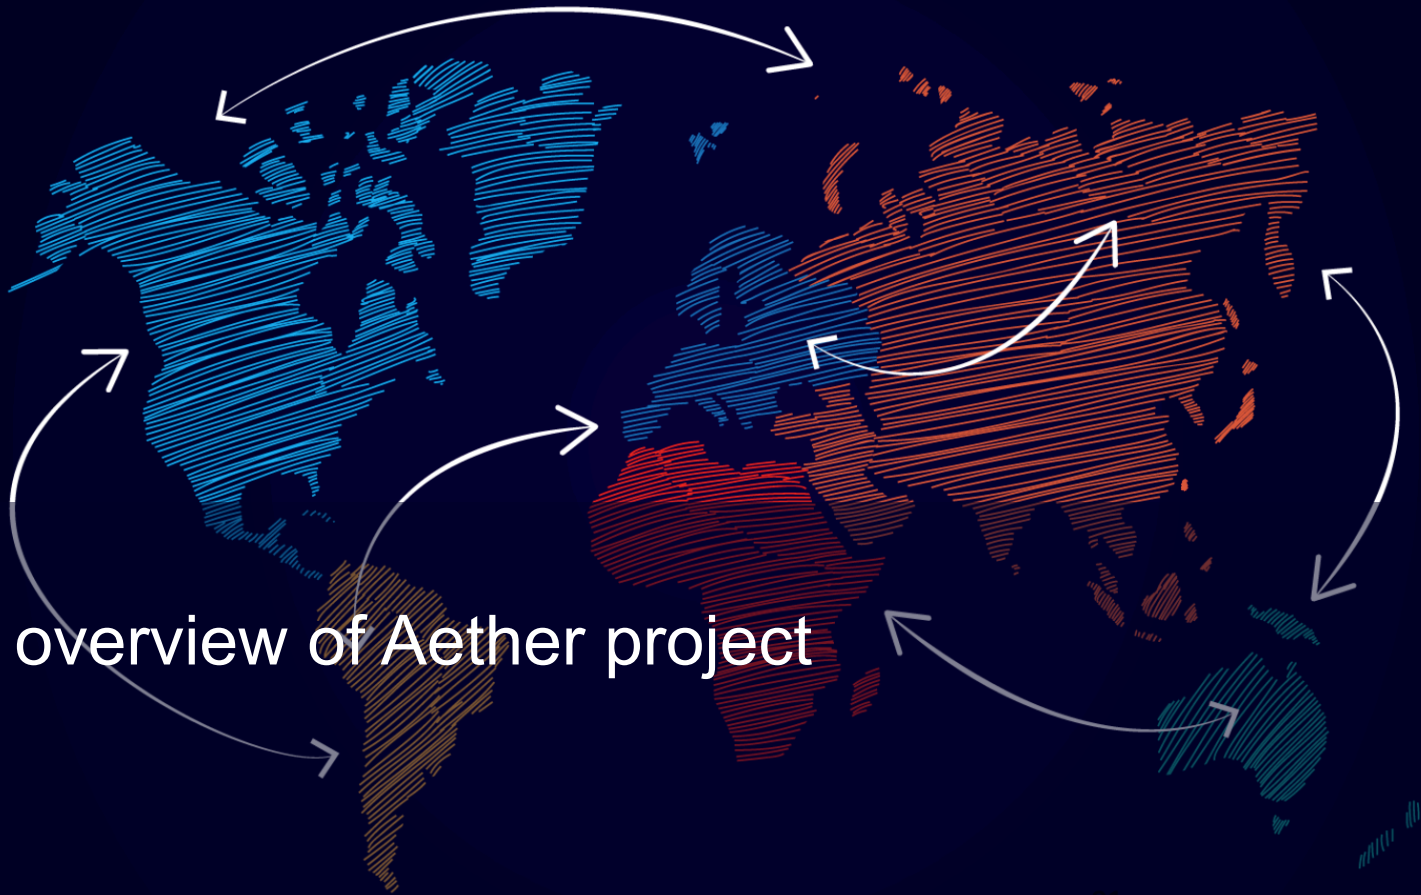
- ONOS
- fabric.p4
- M-CORD

2018-now: MTS (PDP Team)

- Stratum
- Fabric.p4



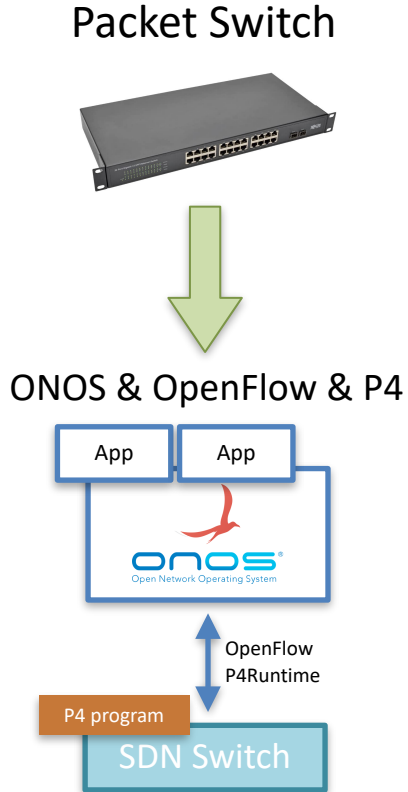
An overview of Aether project



A world map with arrows indicating the flow of goods and services between continents. The map is stylized with a dark background and white outlines for continents. Arrows show a clockwise flow: from North America to Europe, from Europe to Asia, from Asia to Australia, from Australia to South America, and from South America back to North America.



ONF has history of successfully driving disaggregation and SDN




ONF has history of successfully driving disaggregation and SDN



Packet Switch



ONOS & OpenFlow & P4



OpenFlow
P4Runtime

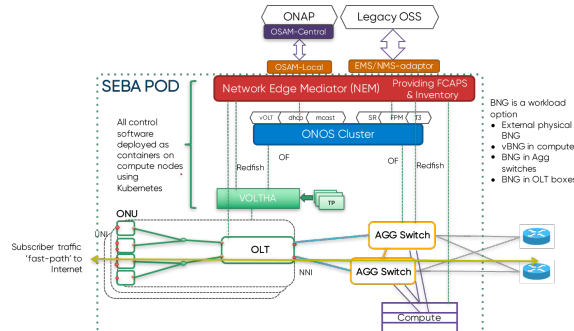
P4 program

SDN Switch

Broadband / PON



SEBA




ONF has history of successfully driving disaggregation and SDN



Packet Switch



ONOS & OpenFlow & P4



OpenFlow
P4Runtime

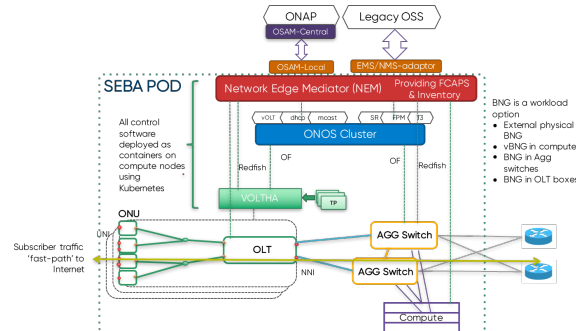
P4 program

SDN Switch

Broadband / PON



SEBA




Optical Transport



ODTN


Open Source
Network Controller



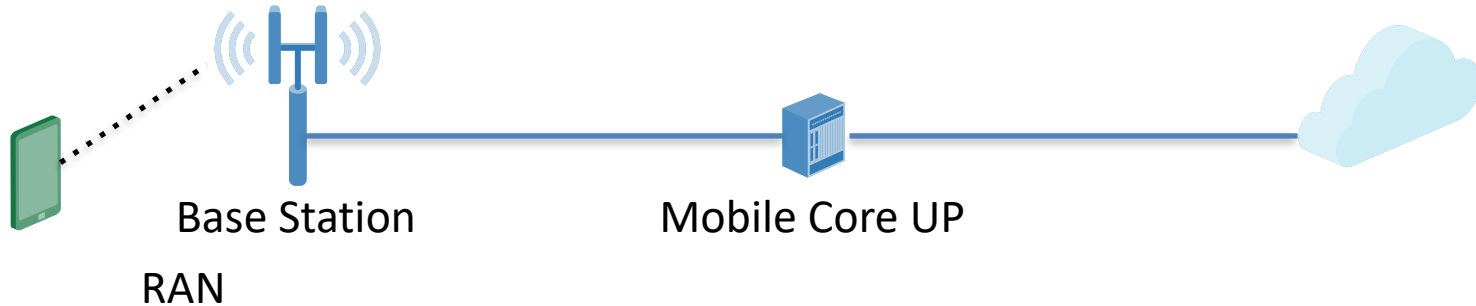
- Multi vendor
- Disaggregated

- Open and standard API
- Common data models

Disaggregation and virtualization for mobile networks




Disaggregation and virtualization for mobile networks




Mobile core control-user plane separation(CUPS)

Disaggregation and virtualization for mobile networks




Mobile core control-user plane separation(CUPS)


Disaggregation and virtualization for mobile networks



Disaggregation and virtualization for mobile networks

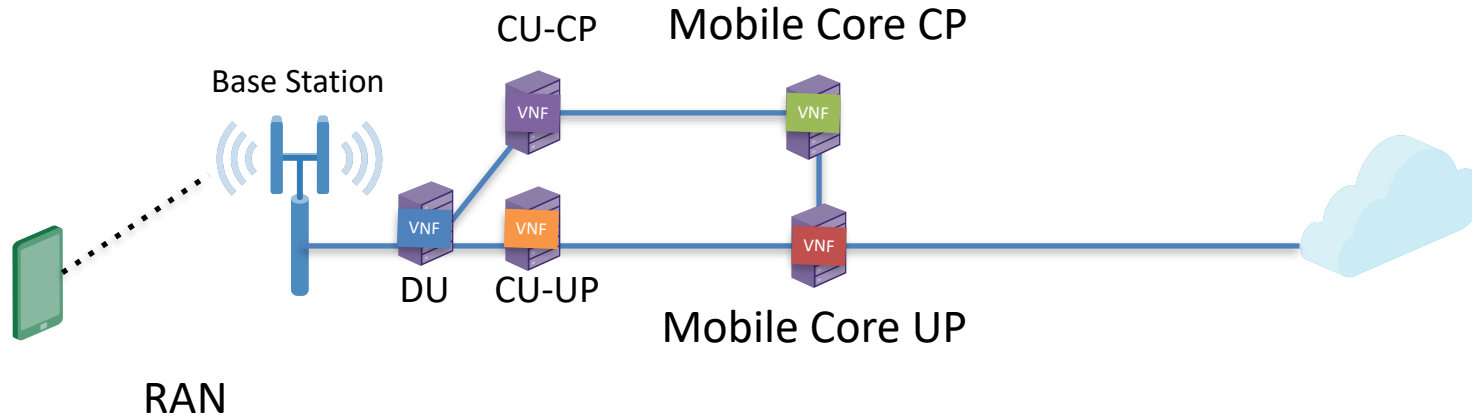


Disaggregation and virtualization for mobile networks



Virtualization

Disaggregation and virtualization for mobile networks



Virtualization


Software-defined networking for mobile networks



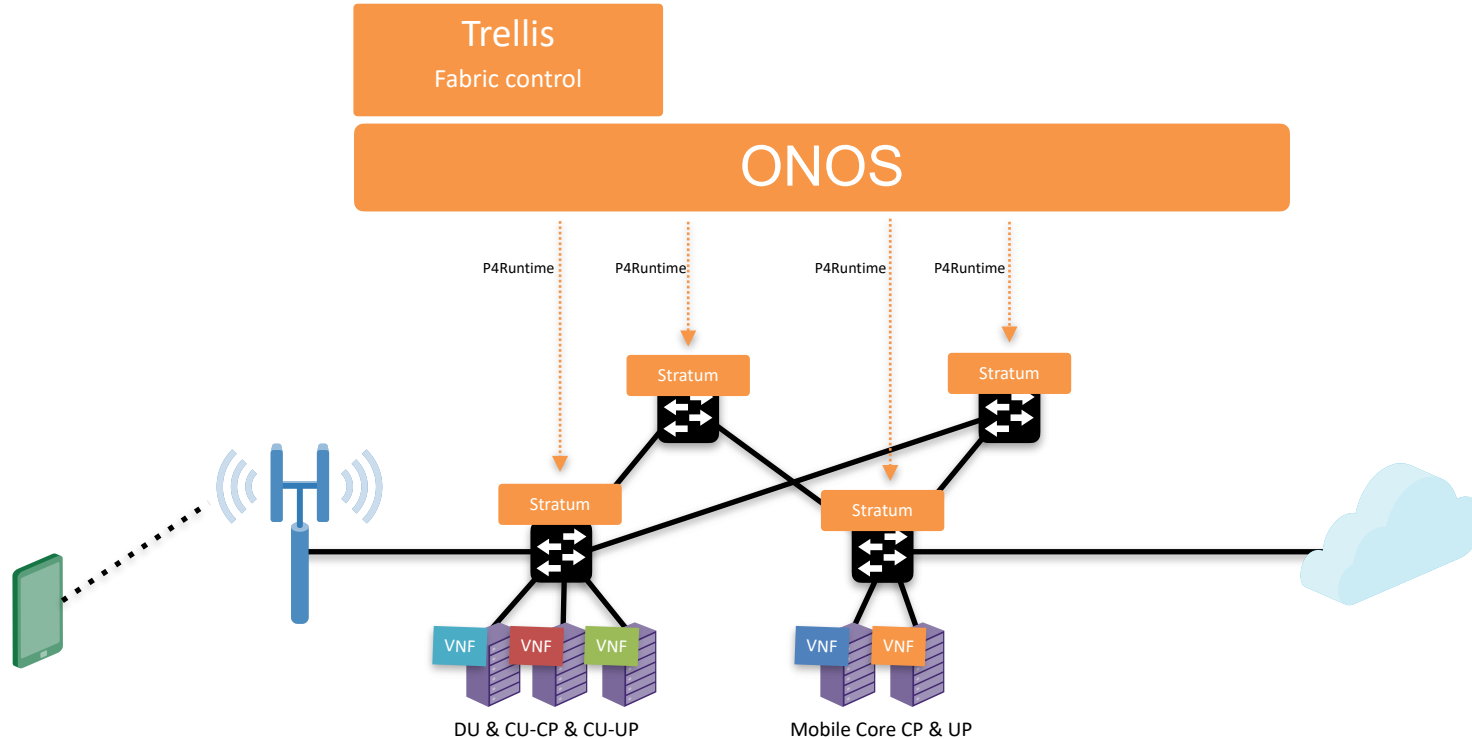
Trellis

Fabric control

ONOS

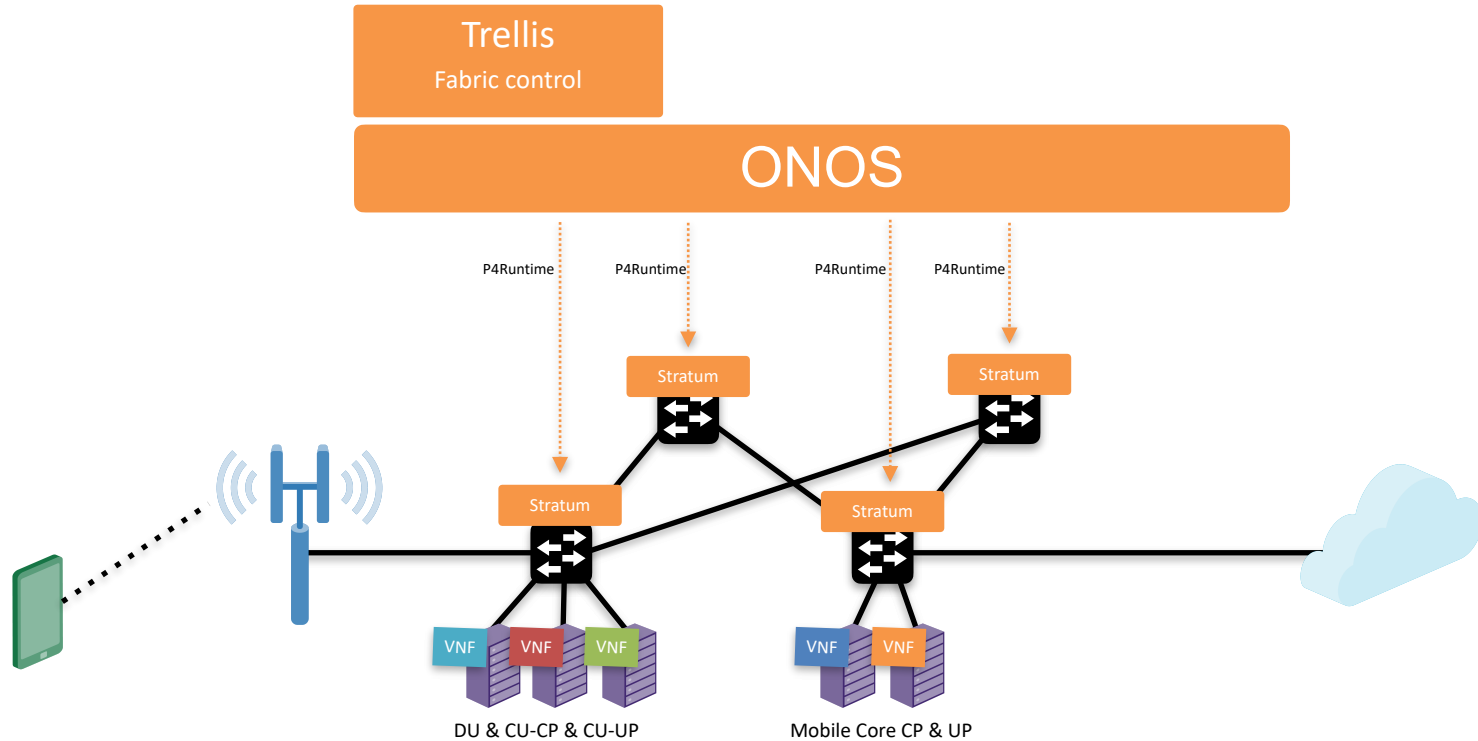


Software-defined networking for mobile networks

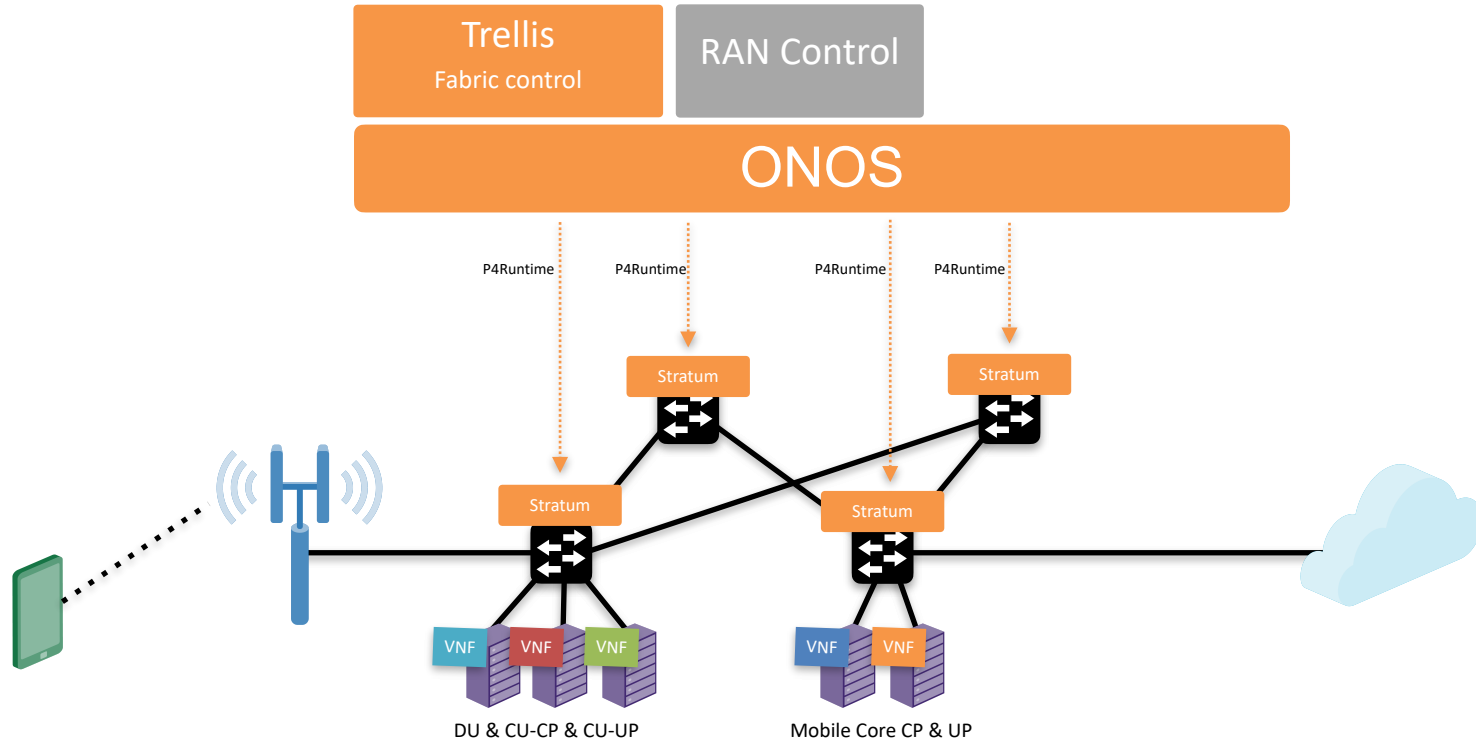


SDN&VNF infrastructure

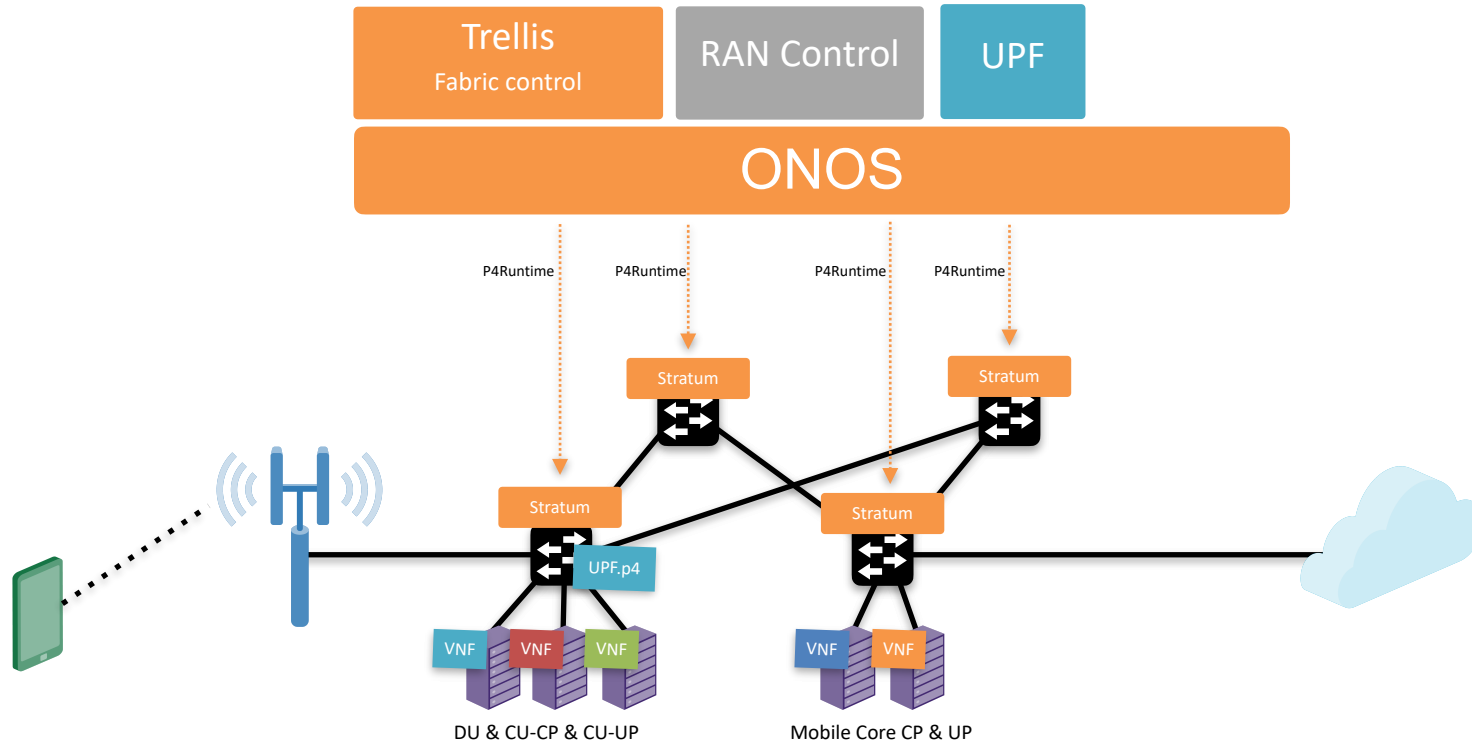
Software-defined networking for mobile networks



Software-defined networking for mobile networks

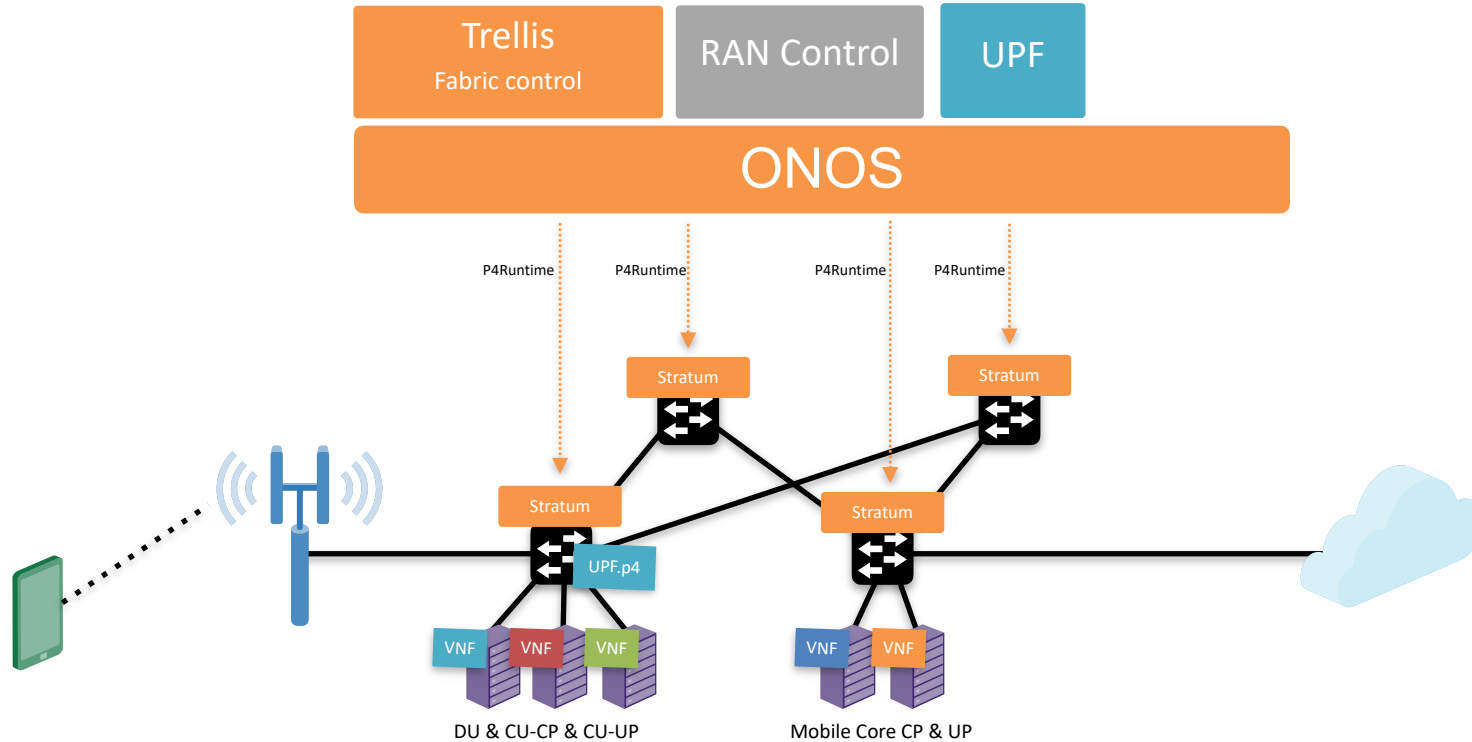


Software-defined networking for mobile networks

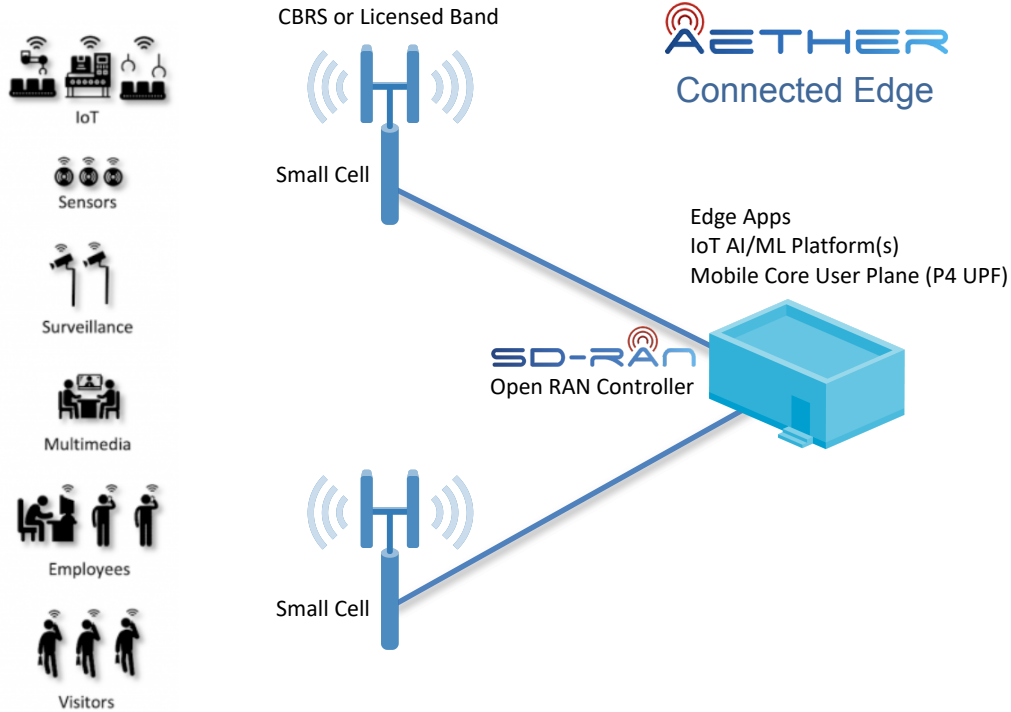


P4-based mobile RAN and core user-planes

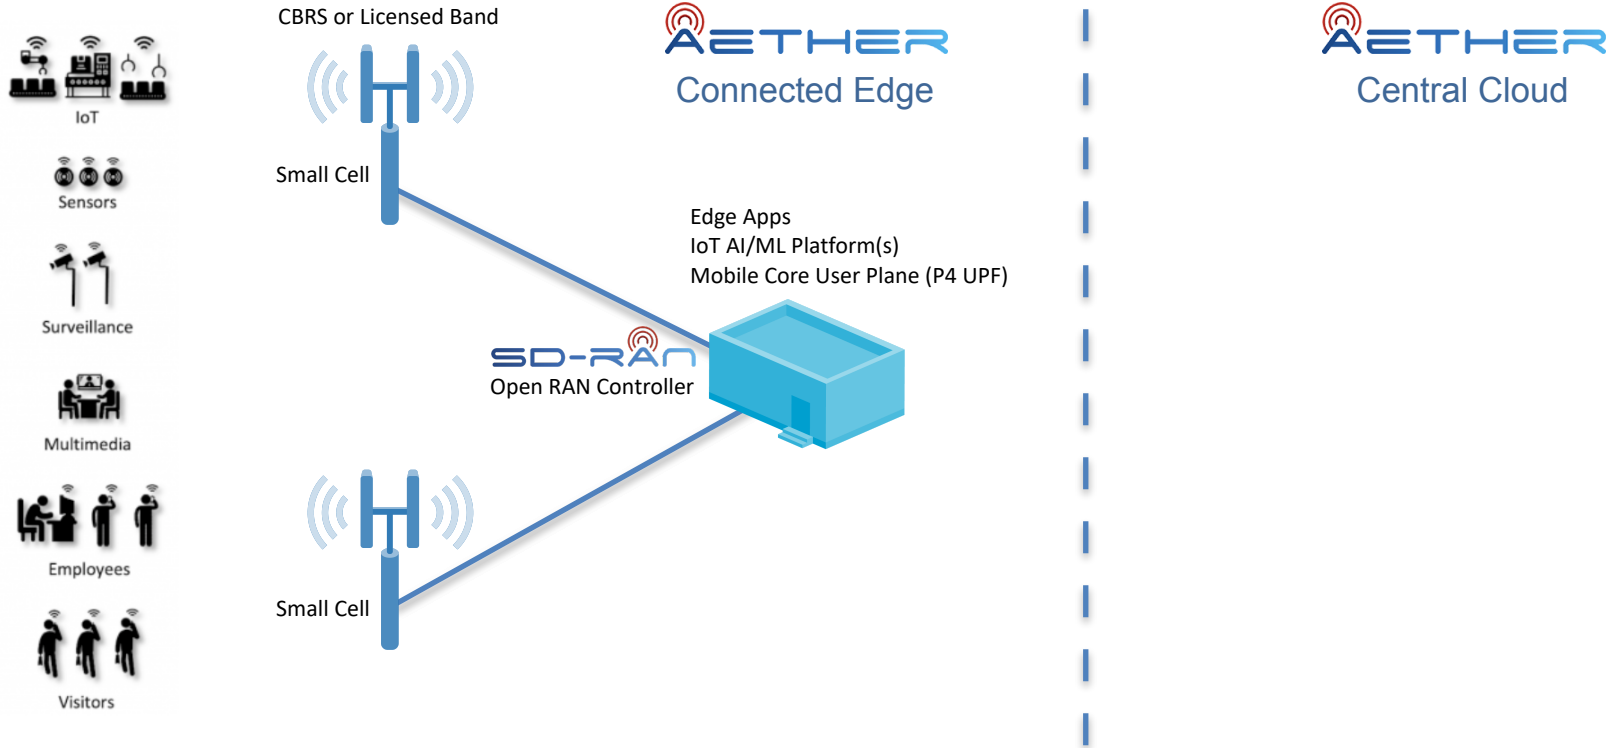
Software-defined networking for mobile networks



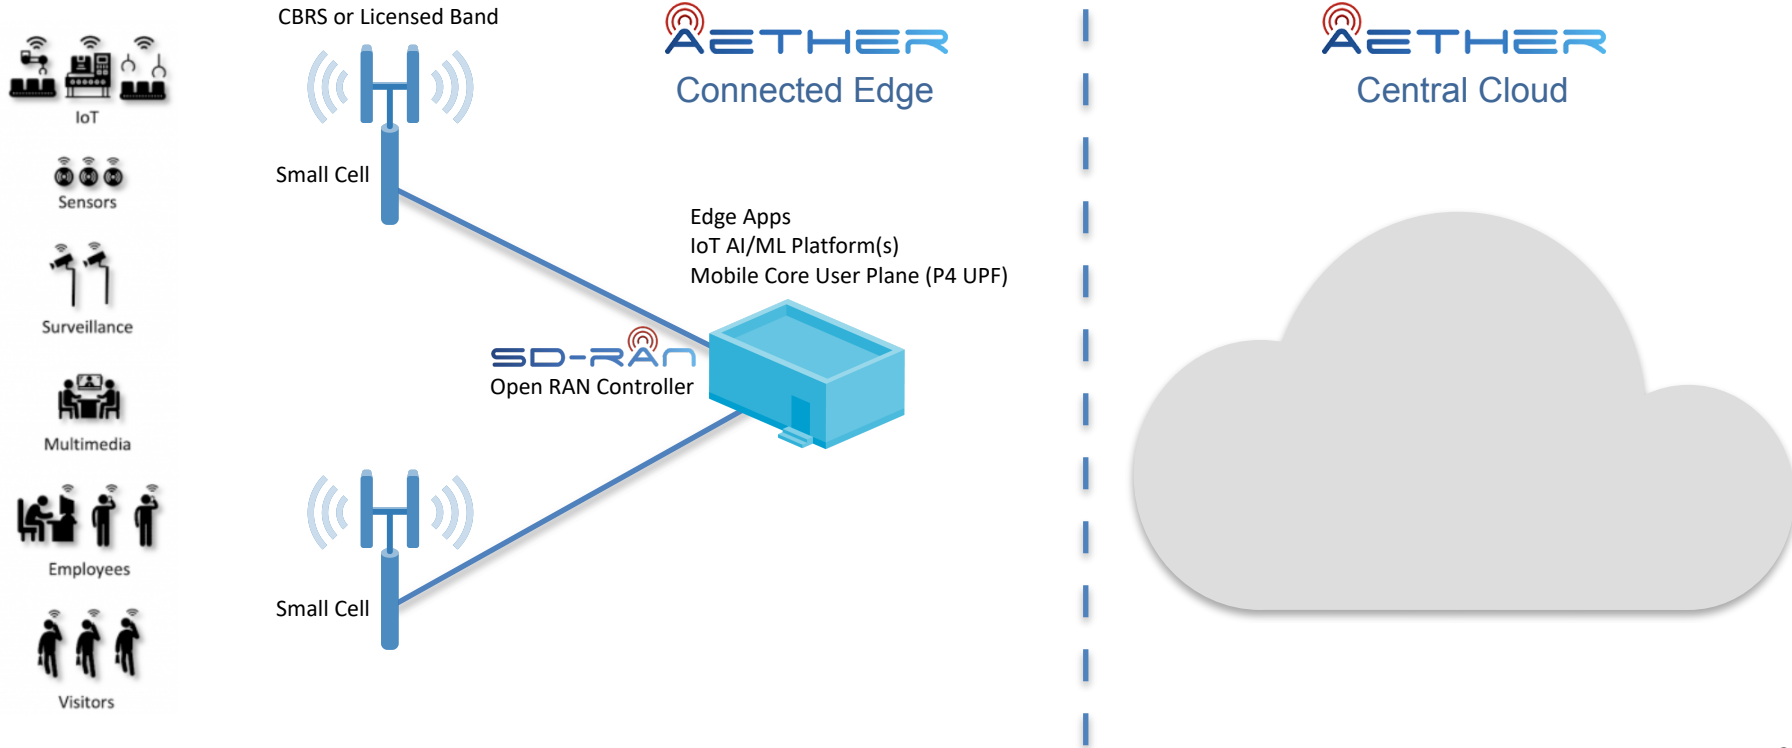
Distributed cloud for mobile networks



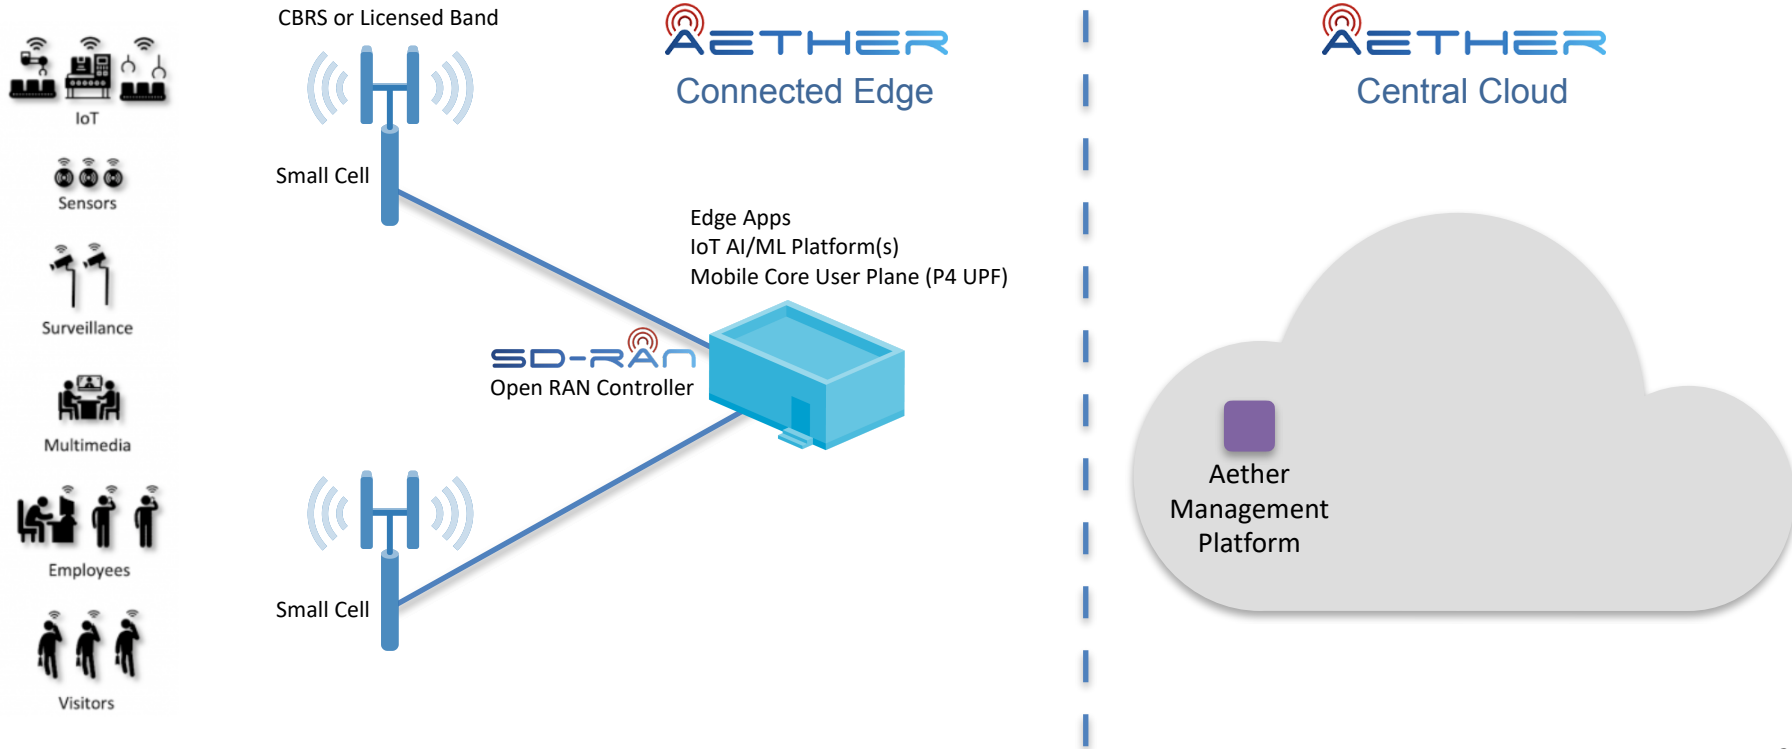
Distributed cloud for mobile networks



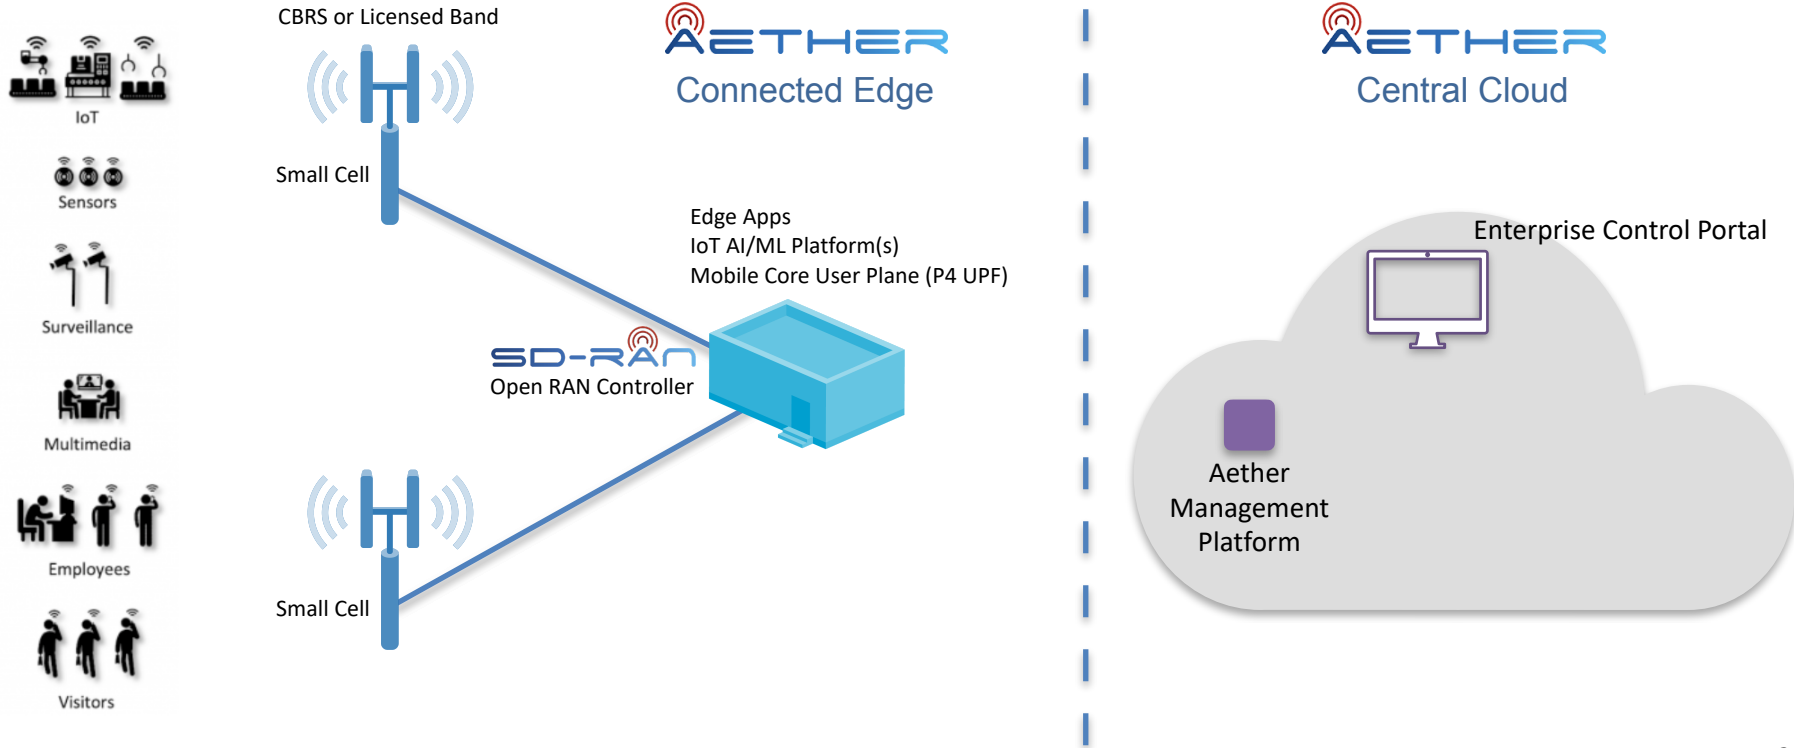
Distributed cloud for mobile networks



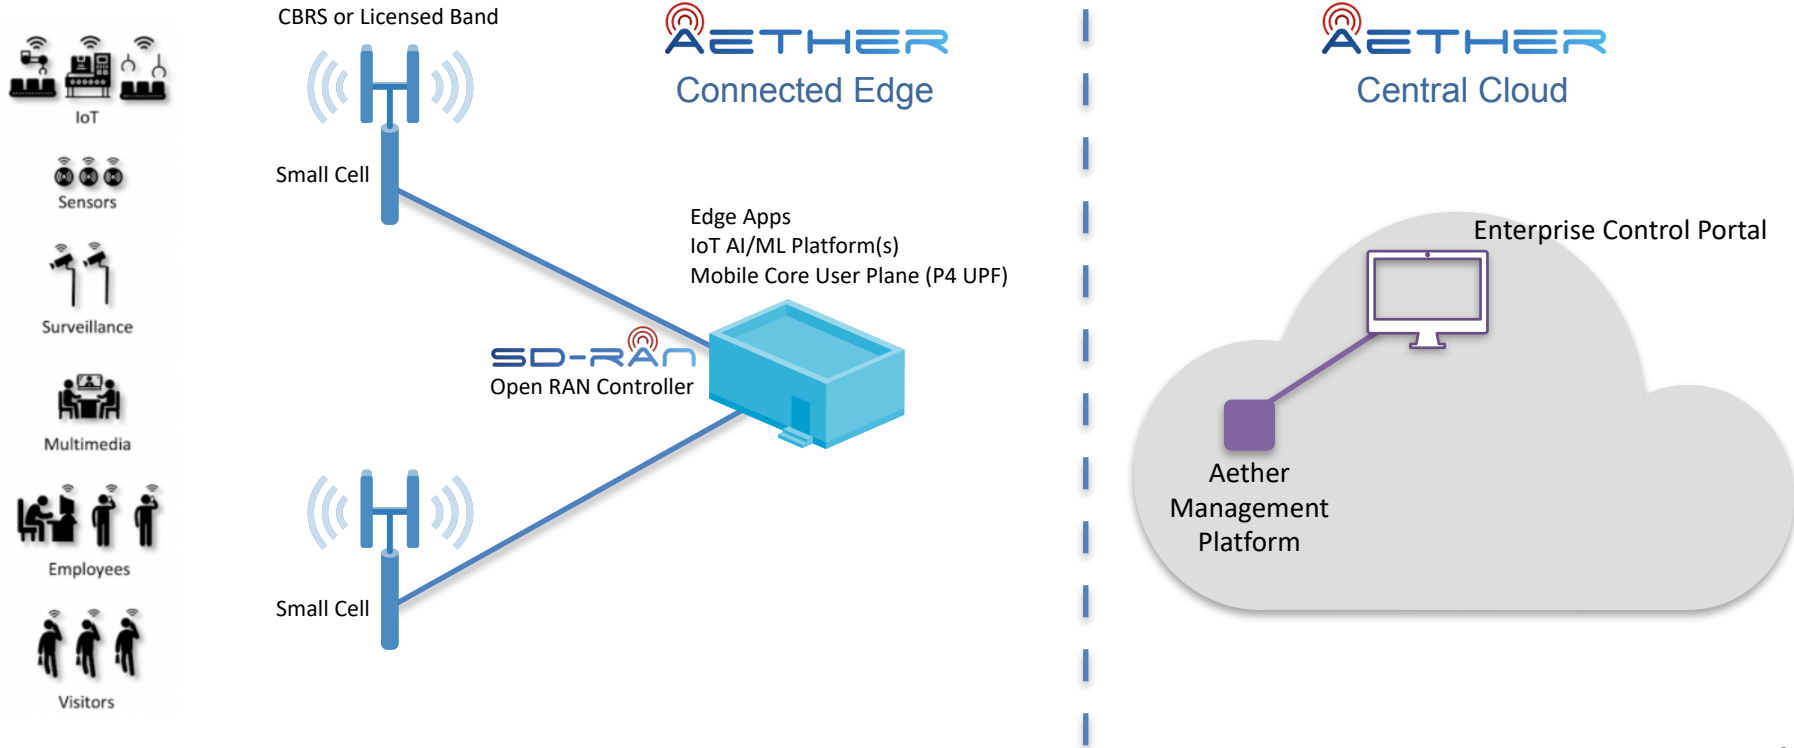
Distributed cloud for mobile networks




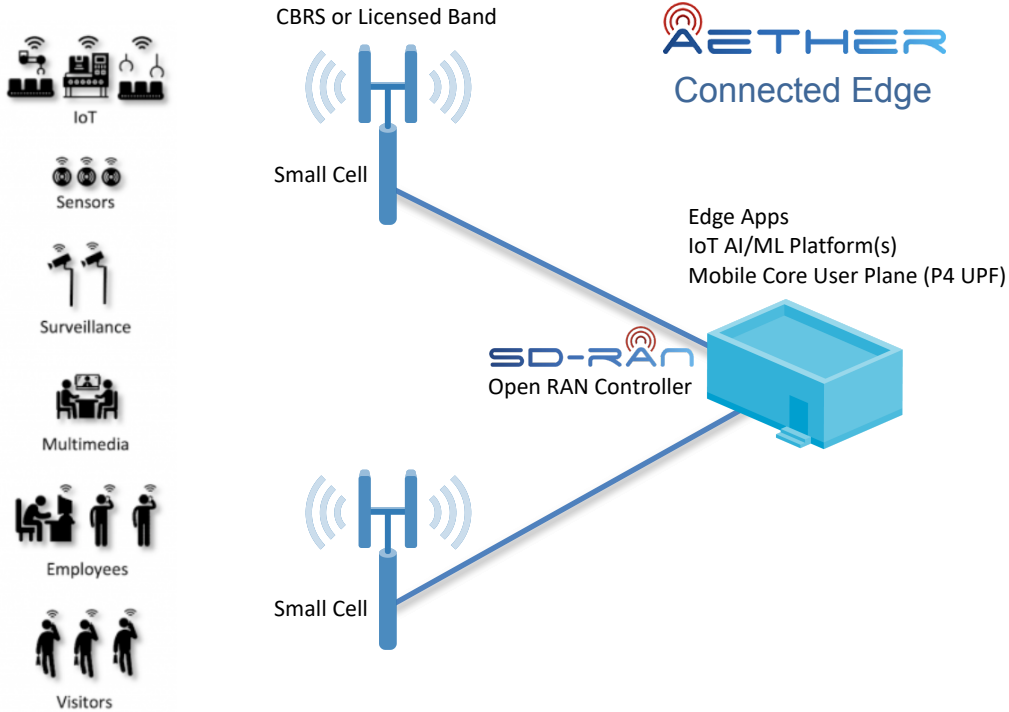
Distributed cloud for mobile networks





Distributed cloud for mobile networks



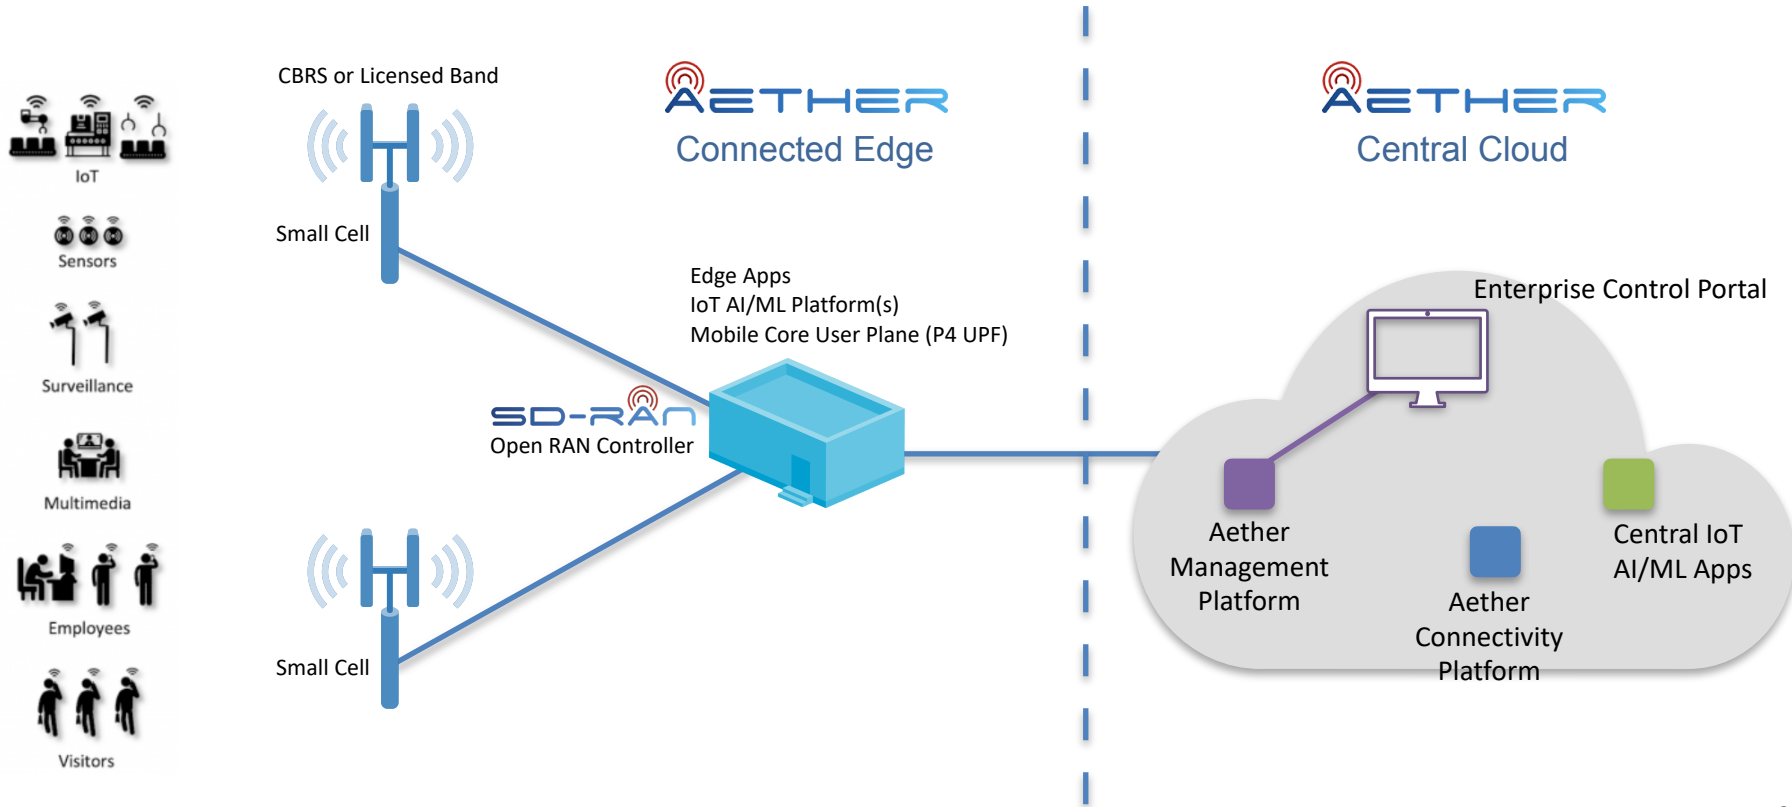
Distributed cloud for mobile networks



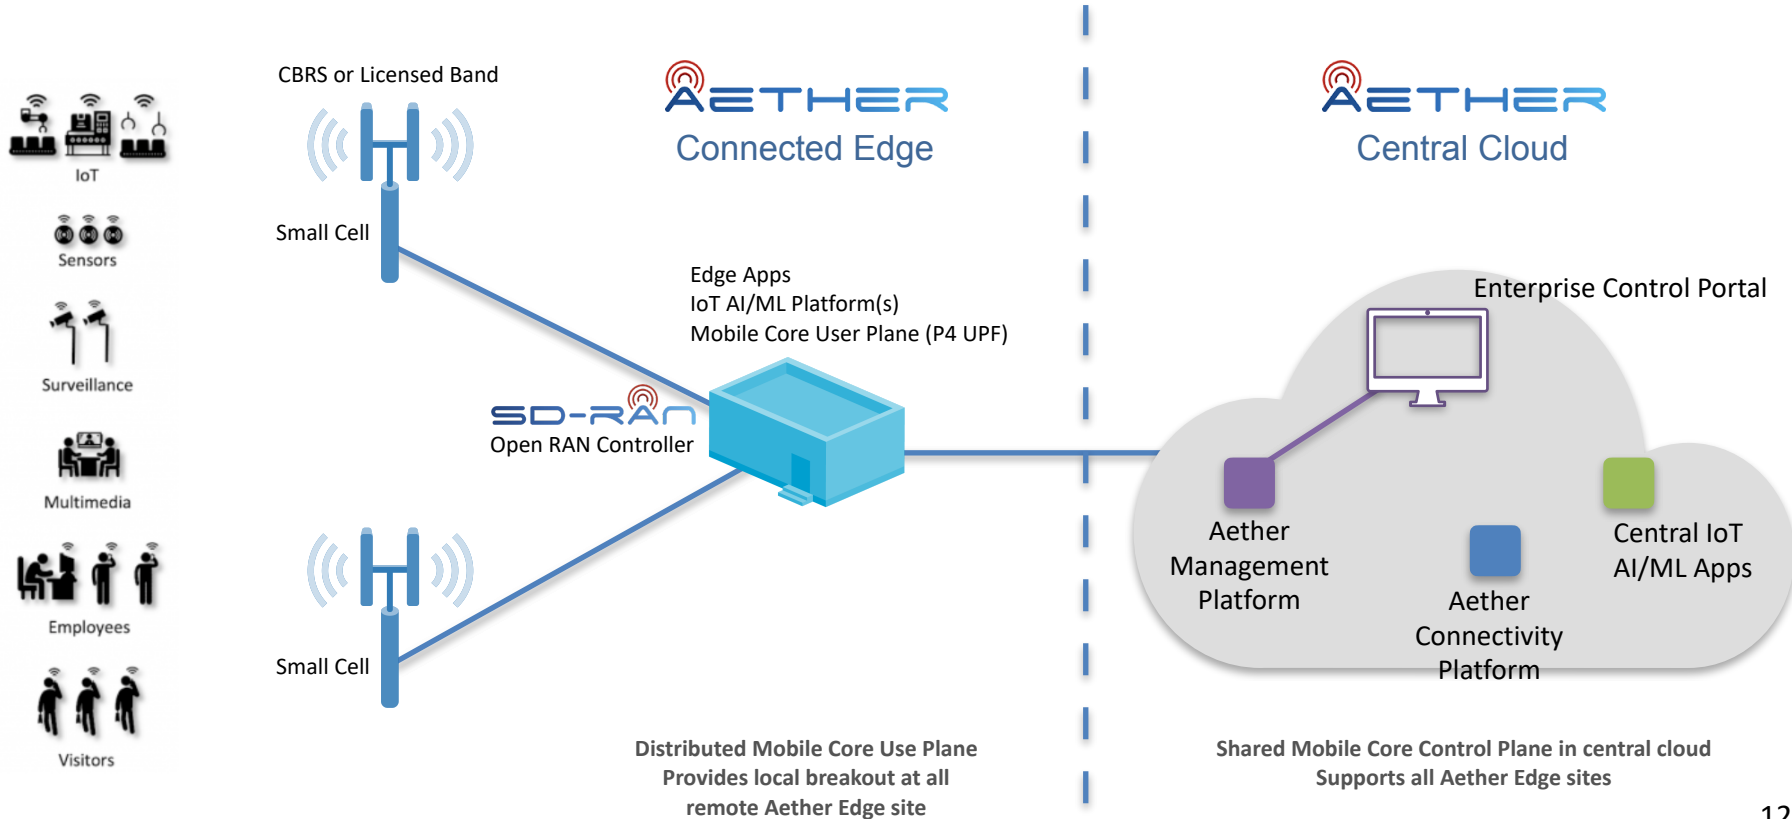
Distributed cloud for mobile networks



Distributed cloud for mobile networks



Distributed cloud for mobile networks



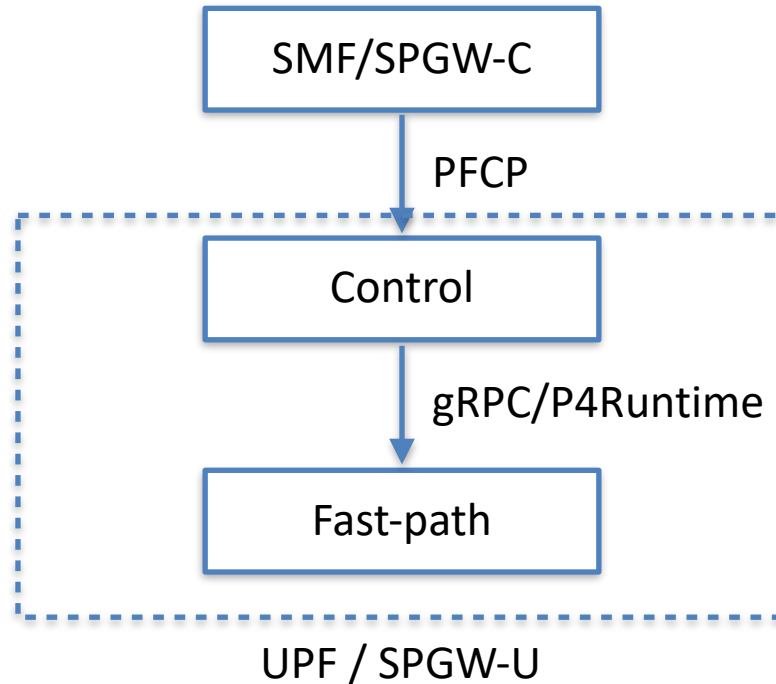
Aether has been operational since December'19



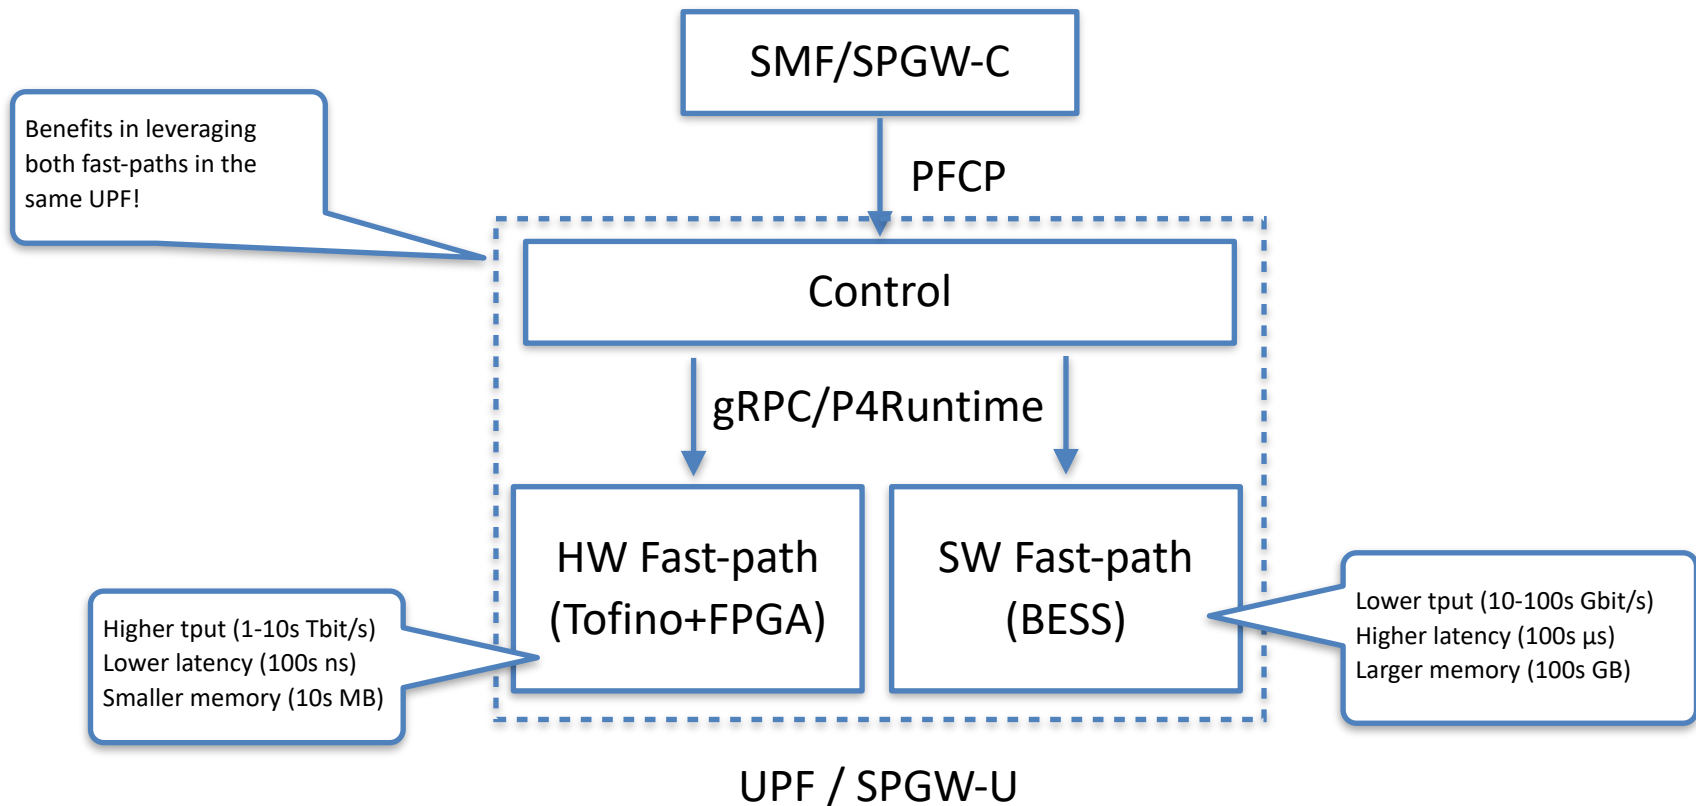


Aether Edge P4-based Disaggregated UPF

A disaggregated UPF



Combine Fast-paths

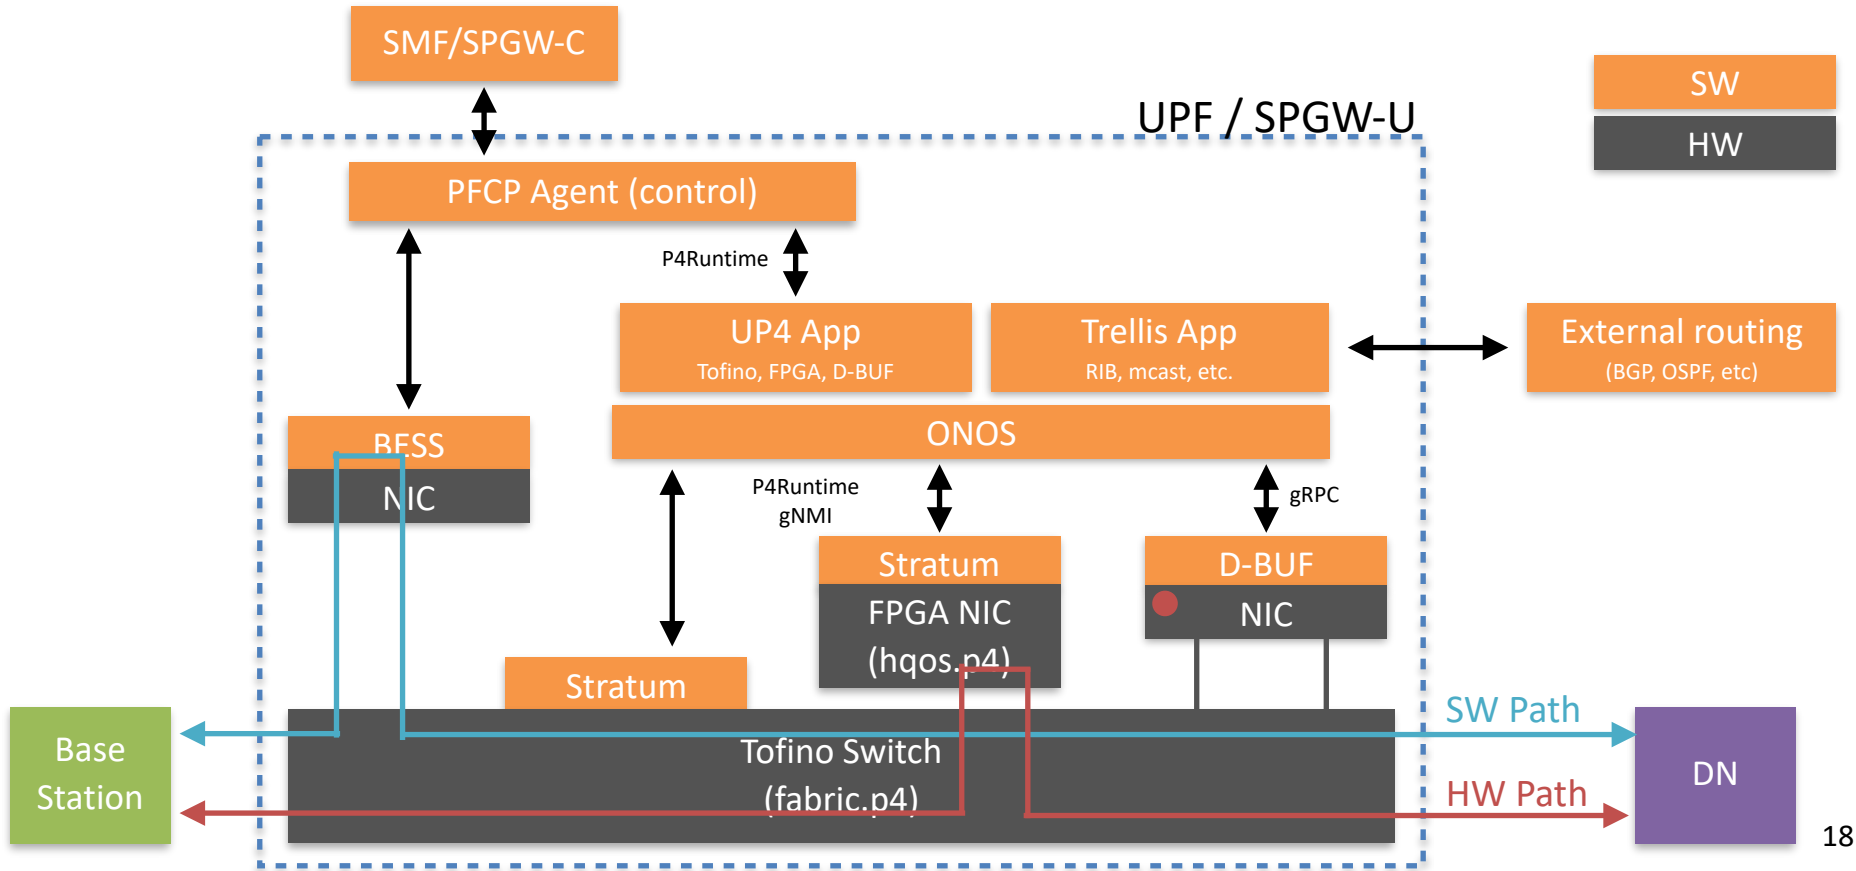


Example: Tesla factory

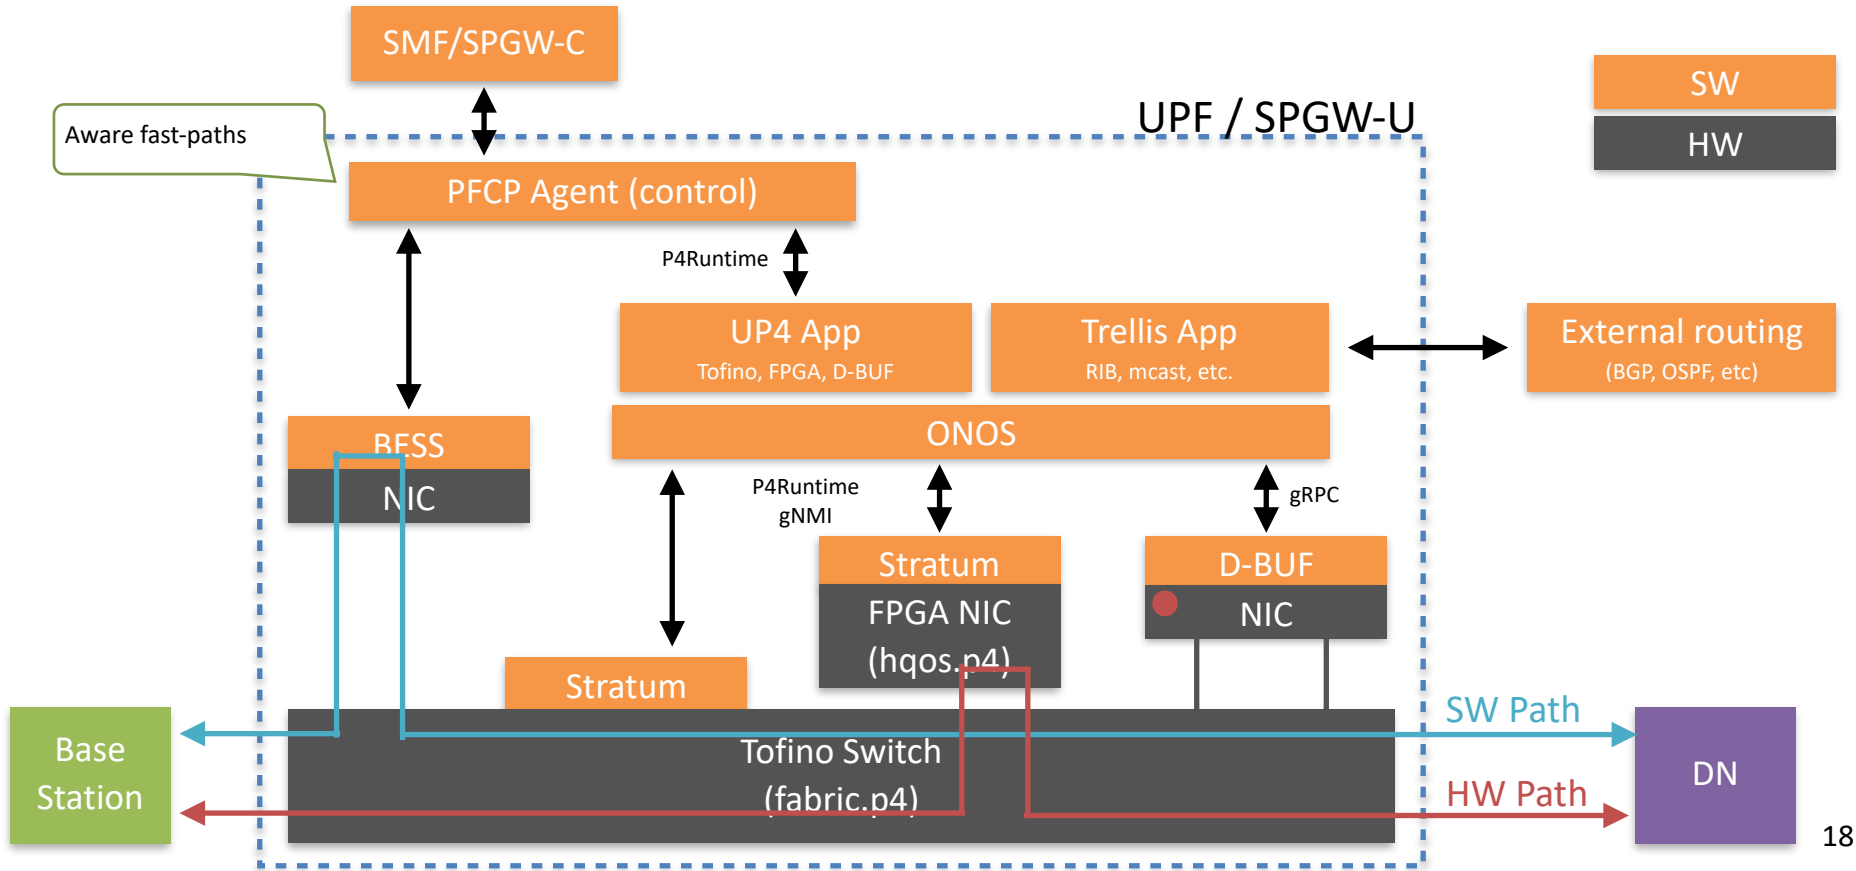


- Requirement: 1M UEs
 - 5% smart phone
 - 10% wideband IoT devices (e.g., HD cameras)
 - 85% narrowband IoT devices (e.g., low data sensors)
- Solution
 - HW fast-path
 - Smartphone + wideband IoT: 150K sessions
 - SW fast-path
 - Narrowband IoT: 850K sessions

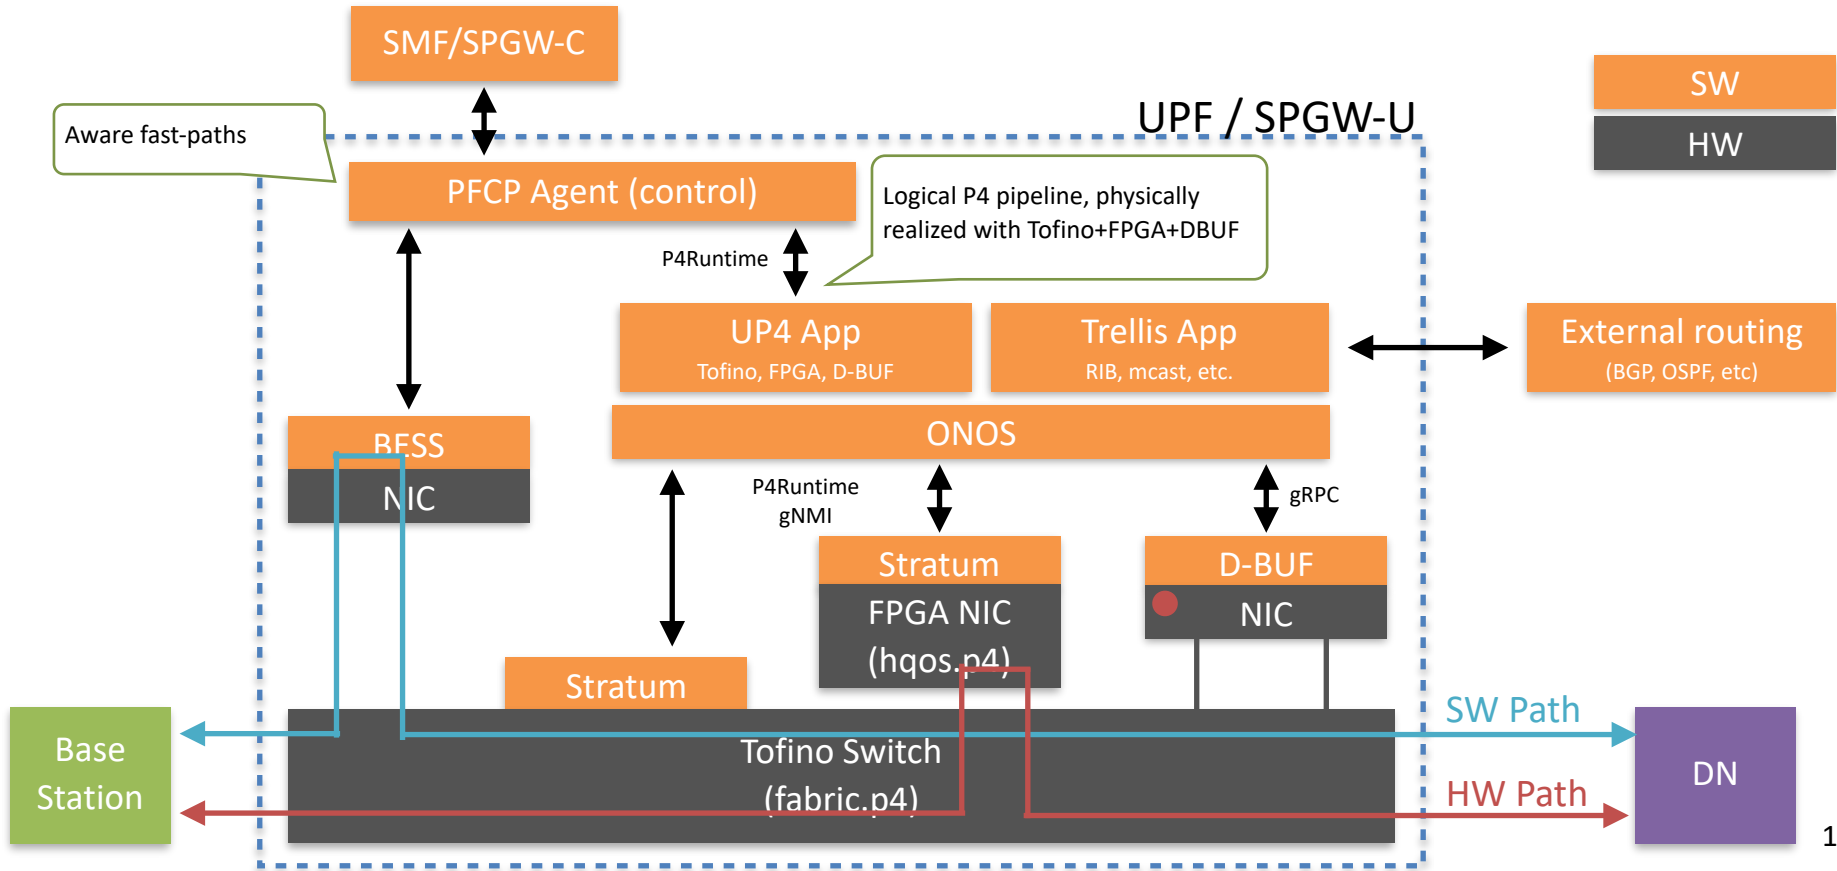
Architecture



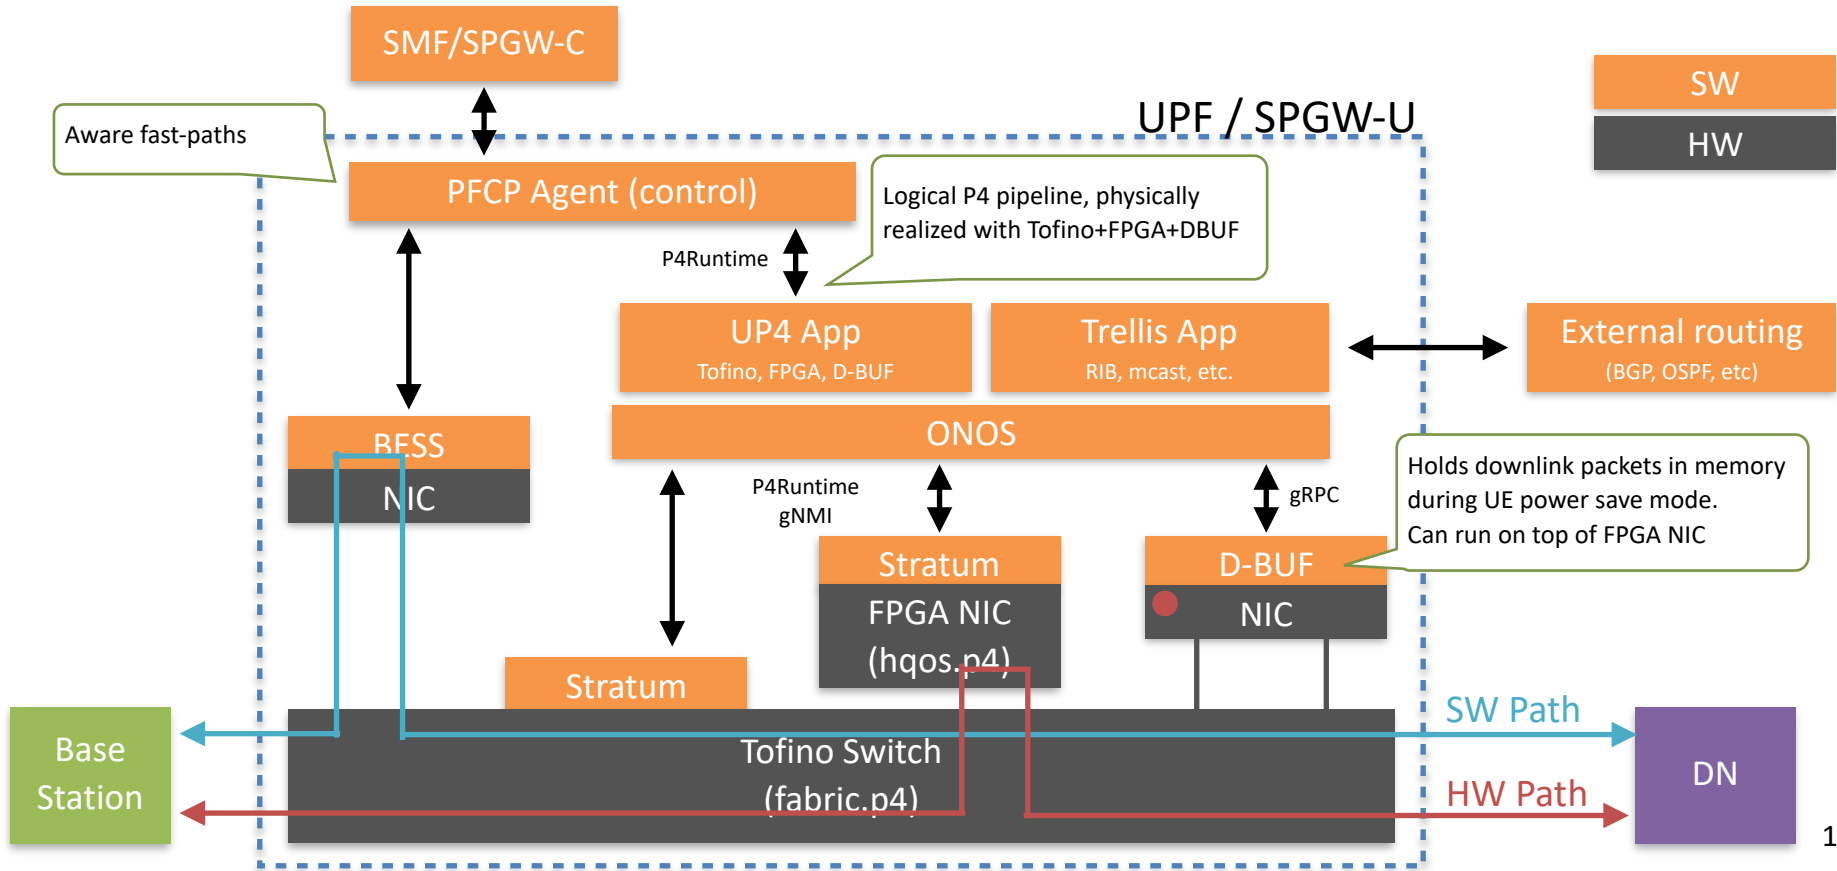
Architecture



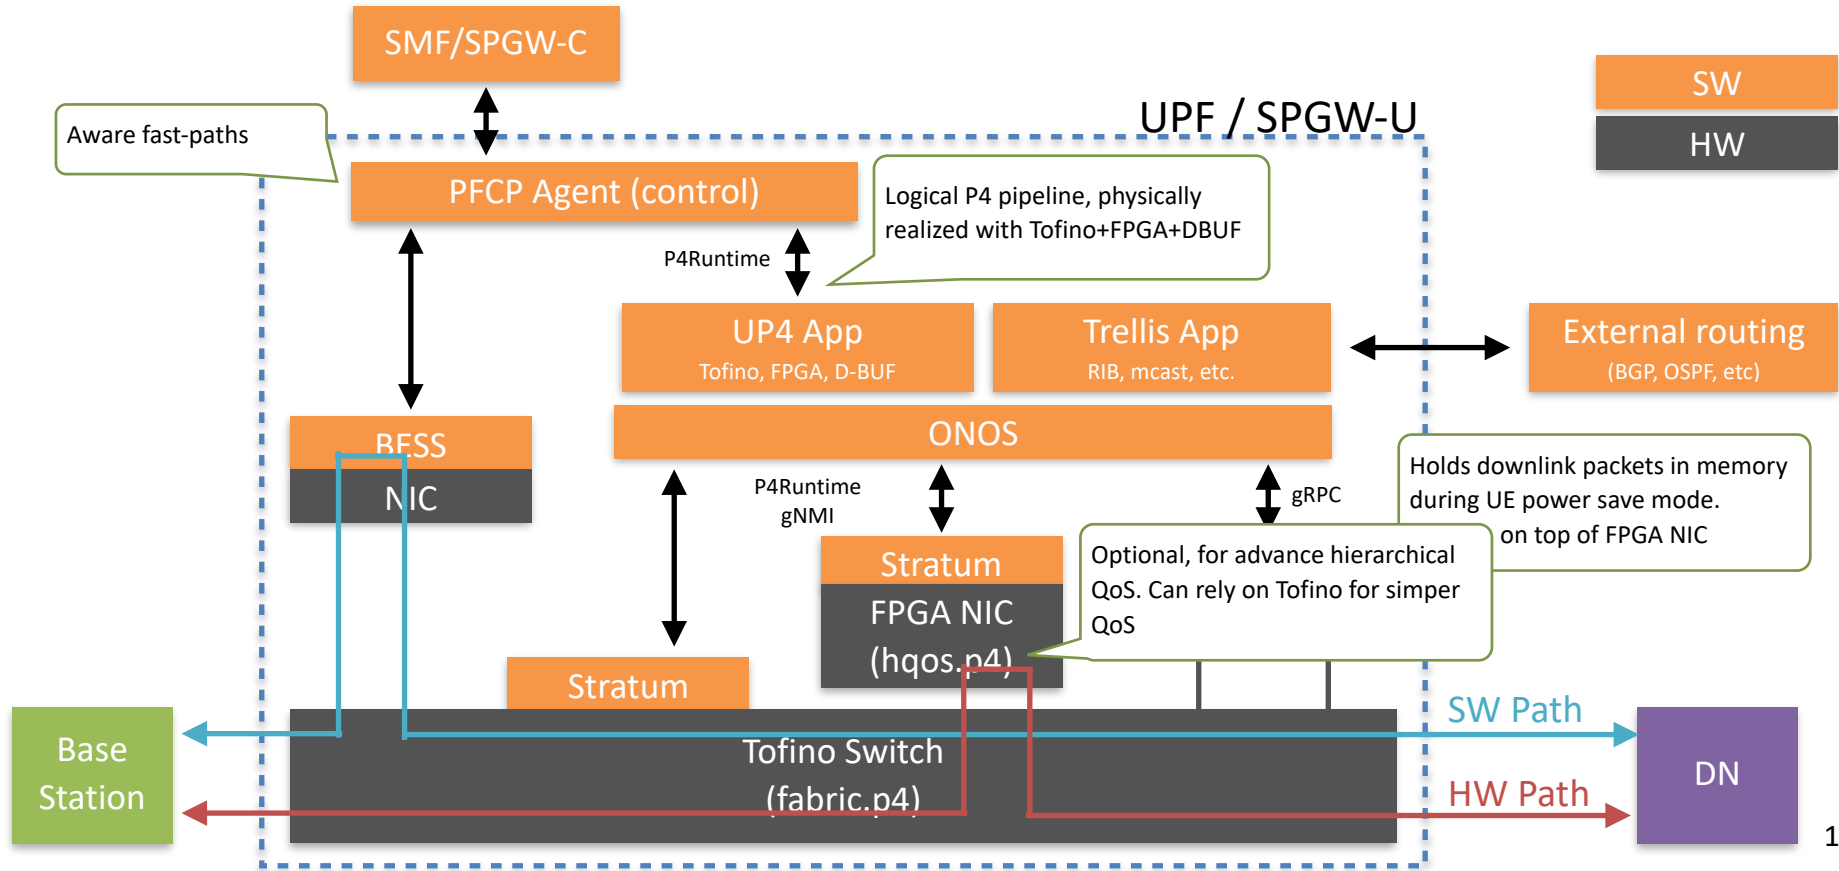
Architecture



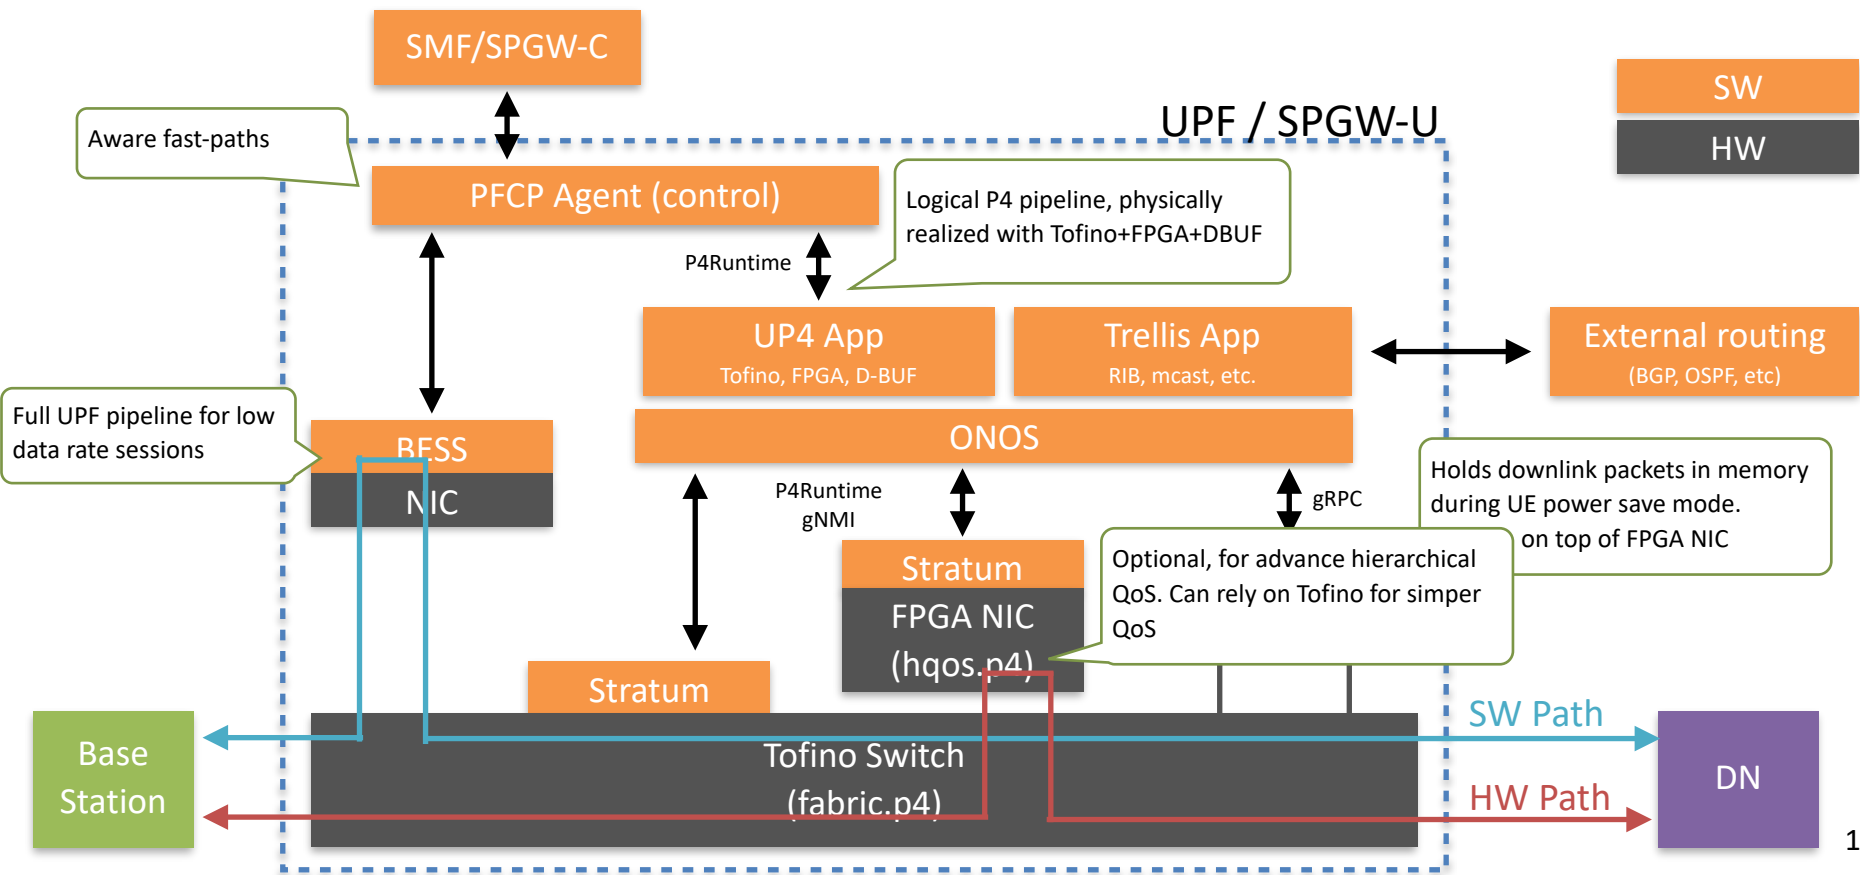
Architecture



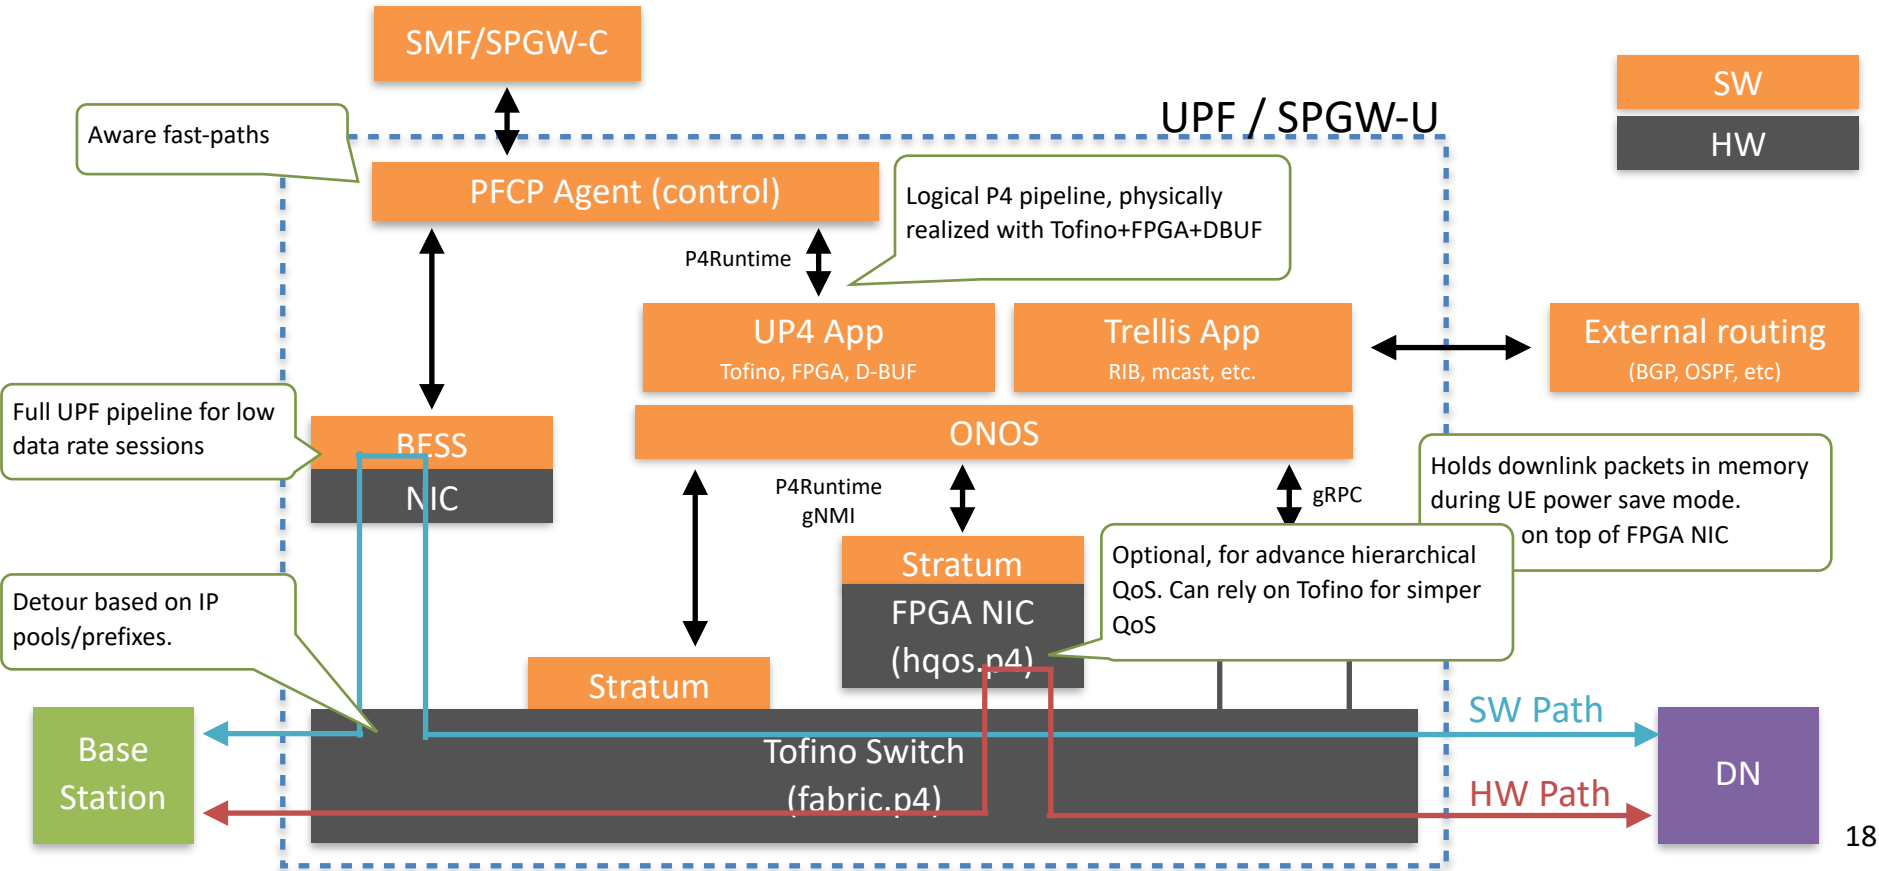
Architecture



Architecture



Architecture



Status (as of December 2020)




- Already started rolling out at several Aether Edge sites:
 - GTP termination and accounting on Tofino, integrated with Trellis/ONOS fabric control
- Q1 2021
 - Downlink buffering via DBUF
 - QoS on Tofino with shared queues
 - Improved scale: 10k UEs
- Q4 2021
 - Integration with FPGA NIC for advance QoS
 - Scale improvements
- Long-term
 - Integration with SW-based fast-path (BESS)

A stylized world map with a dark blue background. The continents are filled with a dense, hand-drawn texture. North America is light blue, South America is light orange, Europe and Africa are dark blue, and Asia and Australia are dark orange. Five white curved arrows form a continuous loop around the globe, starting from the top left, moving clockwise through the top, right, bottom, and back to the top left.

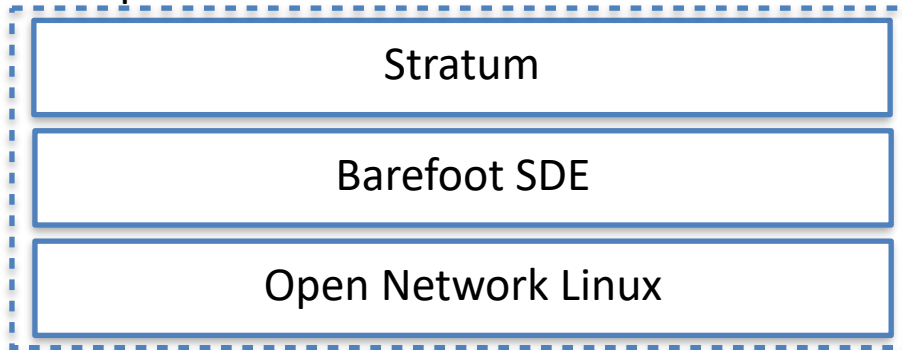
Productize programmable network infrastructure



Control plane software



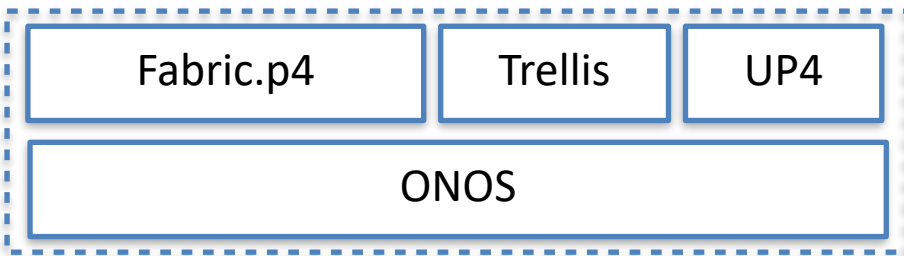
Data plane software



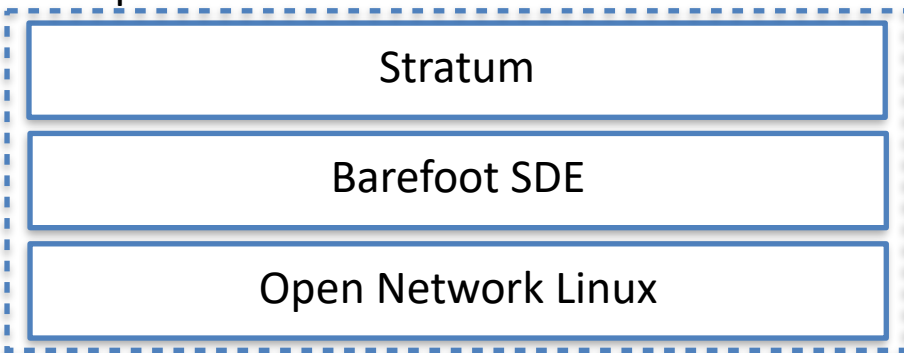
Improved, optimized software stack



Control plane software




Data plane software



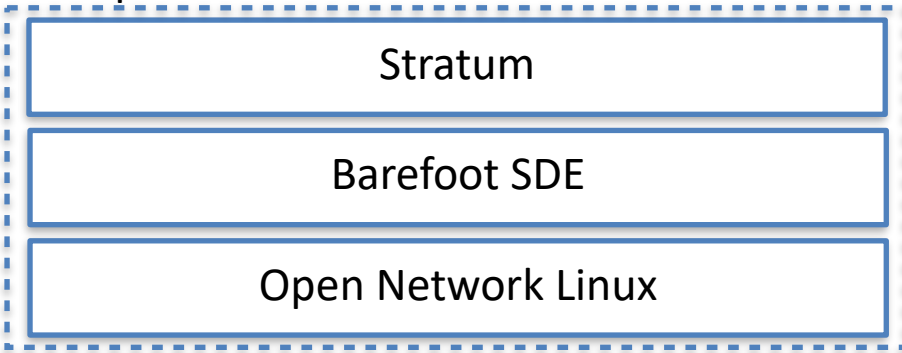
Improved, optimized software stack



Control plane software




Data plane software



Improved, optimized software stack




Control plane software



Rewrite fabric.p4 from v1model architecture to Tofino Native Architecture(TNA).

Allows us to create more advance and optimized P4 program.


Data plane software




Improved, optimized software stack



Control plane software




Data plane software



Improved, optimized software stack

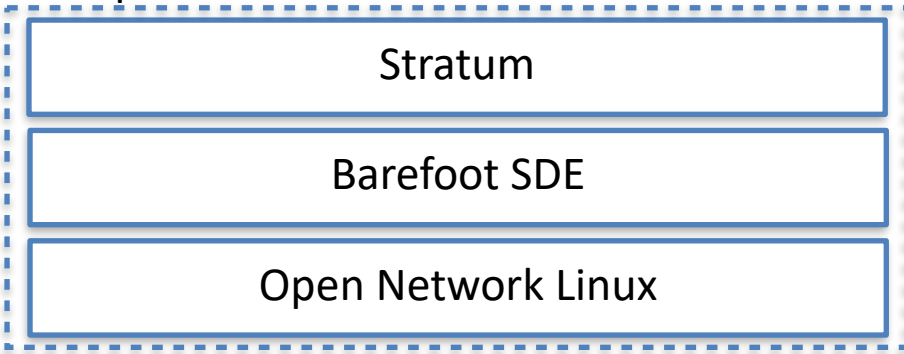


Control plane software



Decoupled from the ONOS code-base with new release cycle.


Data plane software



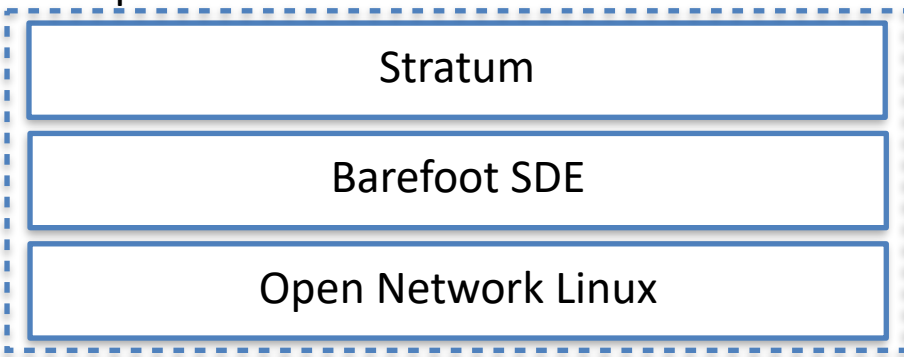
Improved, optimized software stack



Control plane software




Data plane software



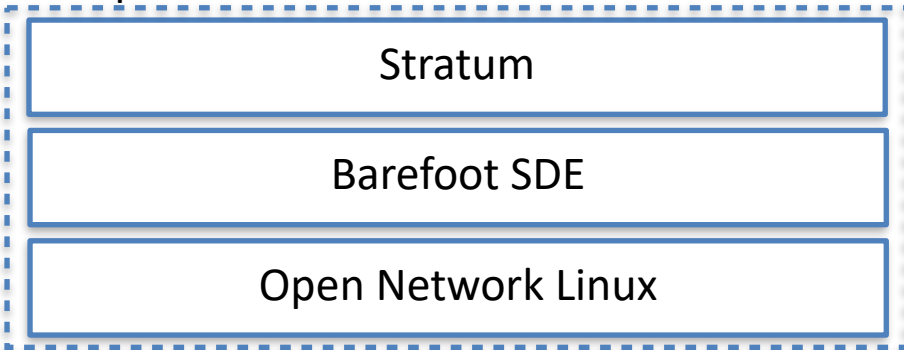
Improved, optimized software stack



Control plane software



Data plane software




Use new ONOS LTS with lots of stability, performance, and availability improvements


Improved, optimized software stack



Control plane software




Data plane software



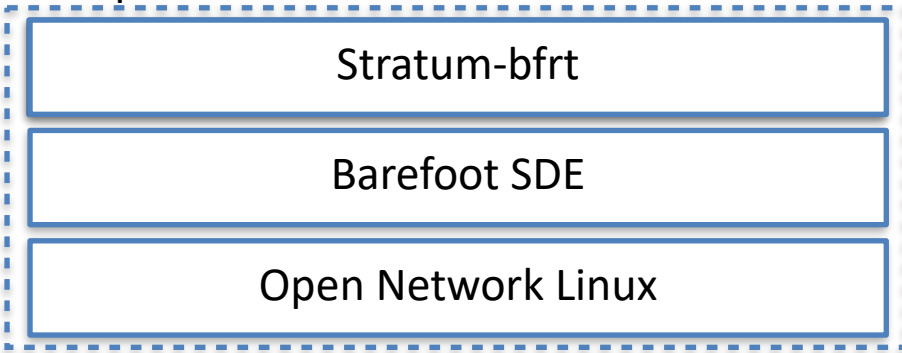
Improved, optimized software stack



Control plane software




Data plane software



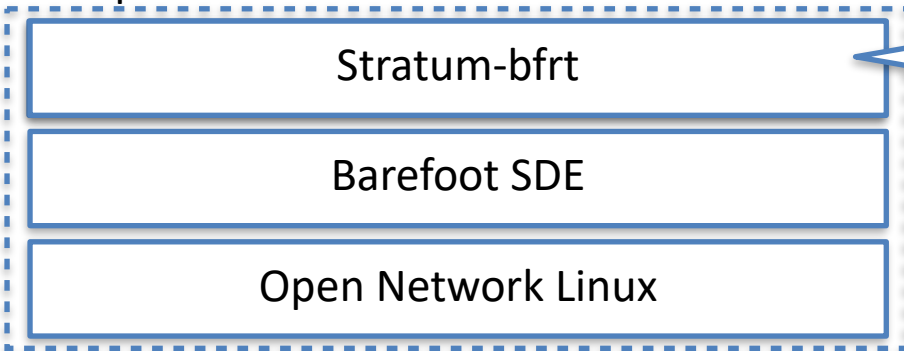
Improved, optimized software stack



Control plane software



Data plane software




New Stratum implementation based on Barefoot native API unlocks more advance ASIC management.


Improved, optimized software stack



Control plane software




Data plane software



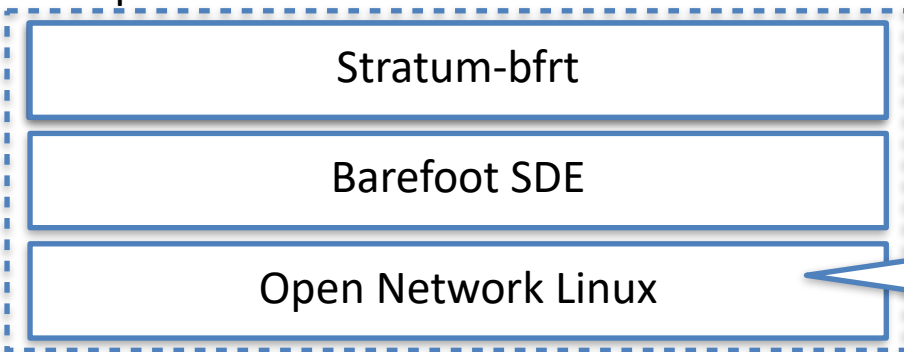
Improved, optimized software stack



Control plane software



Data plane software




Several improvements to support fast deployment and troubleshooting.

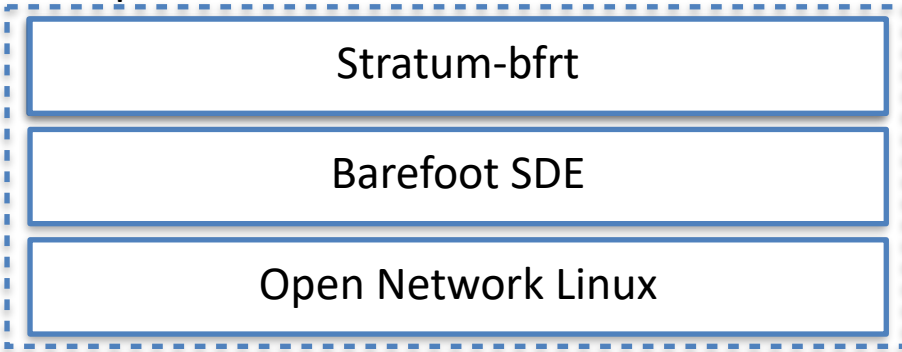
Improved, optimized software stack



Control plane software



Data plane software

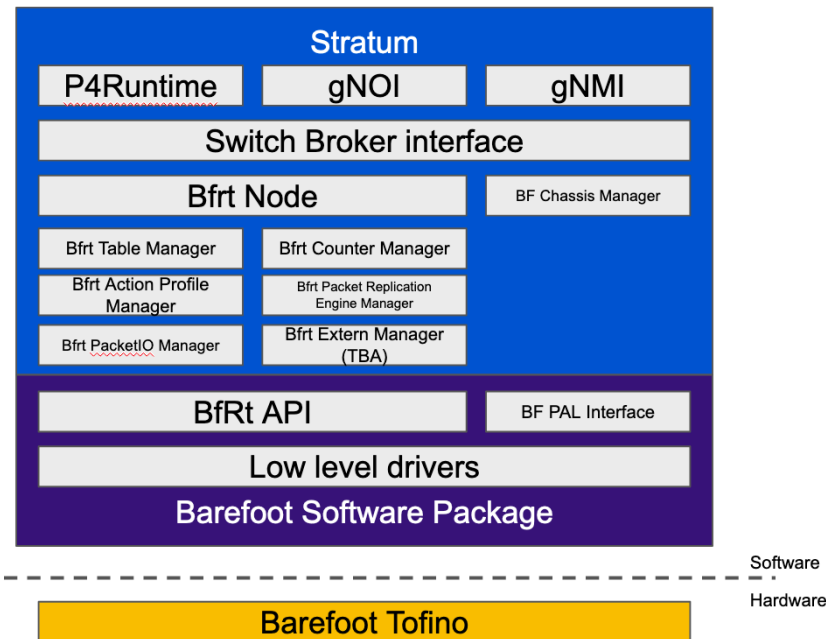




- ONF's fabric.p4 on Tofino Native Architecture(TNA)
- Supports Aether Edge use-cases
 - Trellis (Bridging, Routing, ...)
 - UPF/SPGW-U
 - Simple QoS, accounting
 - Integrate with D-BUF
 - Inband Network Telemetry (INT)
 - Advance telemetry report mechanism




- Stratum implementation with Barefoot BfRt C++ API
- Performance improvement
- Advance ASIC control
 - Batching/Transaction
 - Register
 - Traffic manager
 - Egress mirroring
 - Folded/Multi pipeline
 - ...






TOST: Trellis ONOS Stratum Tofino

TOST container image



Stratum container image



ONIE installer





Tofino Switch

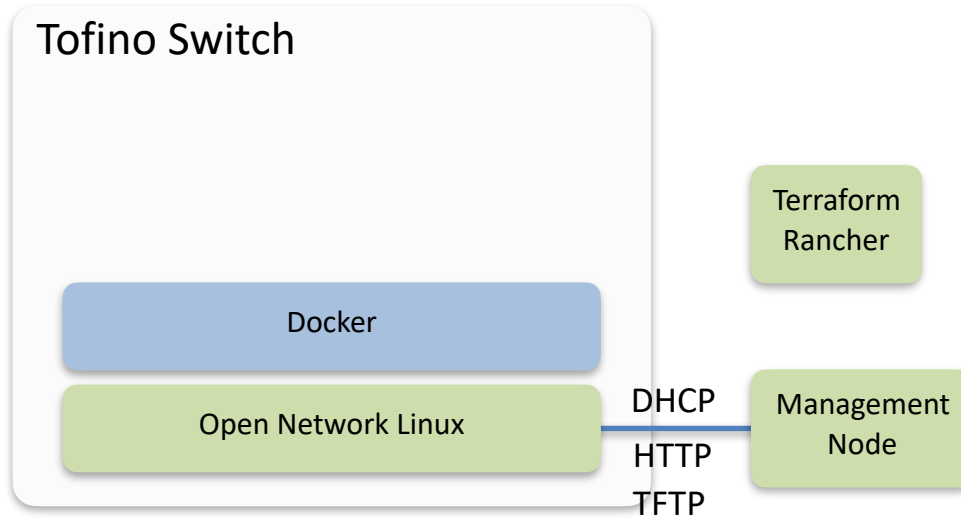
Terraform
Rancher

Management
Node

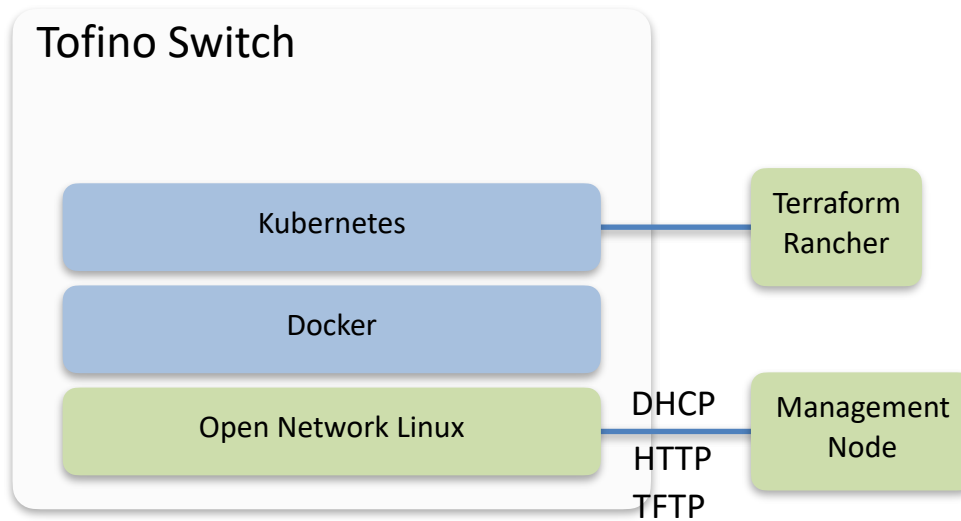
Kubernetes Integration



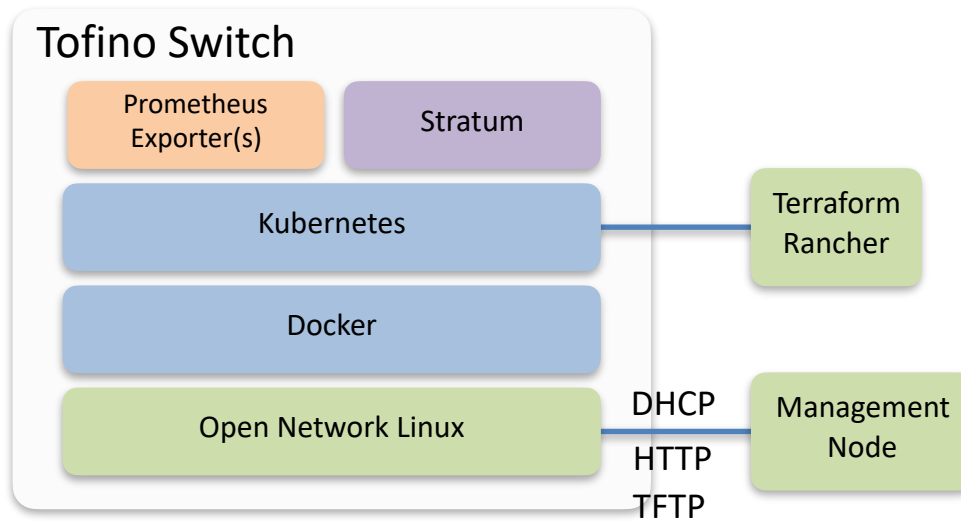
Kubernetes Integration



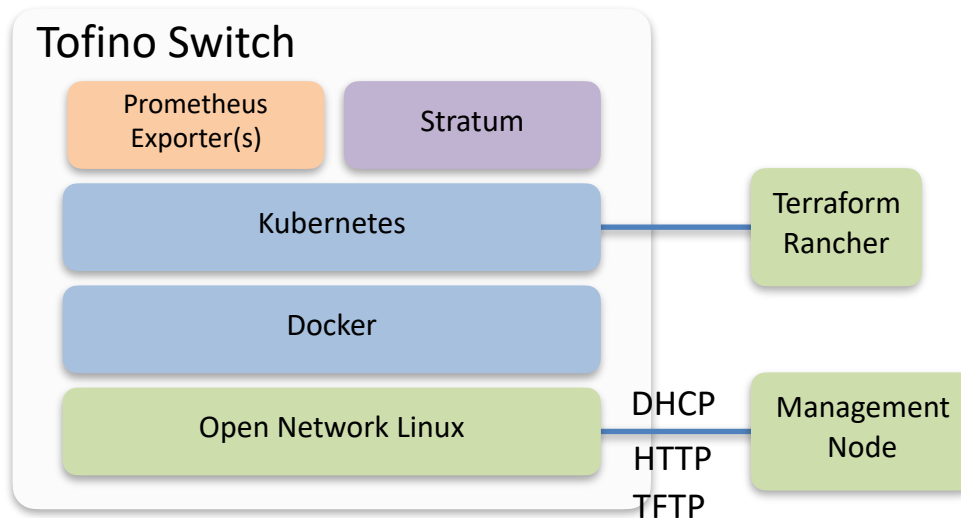
Kubernetes Integration



Kubernetes Integration



Kubernetes Integration




- Tofino switch as a Kubernetes worker node
 - With special taint and label to make sure only Stratum is deployed on it
- Stratum is deployed as Kubernetes service
 - Deployed as DaemonSet. There will be one and only one instance on each switch node
 - P4RT/gNMI exposed via NodePort
 - externalTrafficPolicy=Local so the traffic won't get load-balanced to other switches

Automation - Build and Release




- Git-triggered automated build and release process for Trellis apps and control plane container image
- Build and release Stratum image weekly



Automation - Build and Release



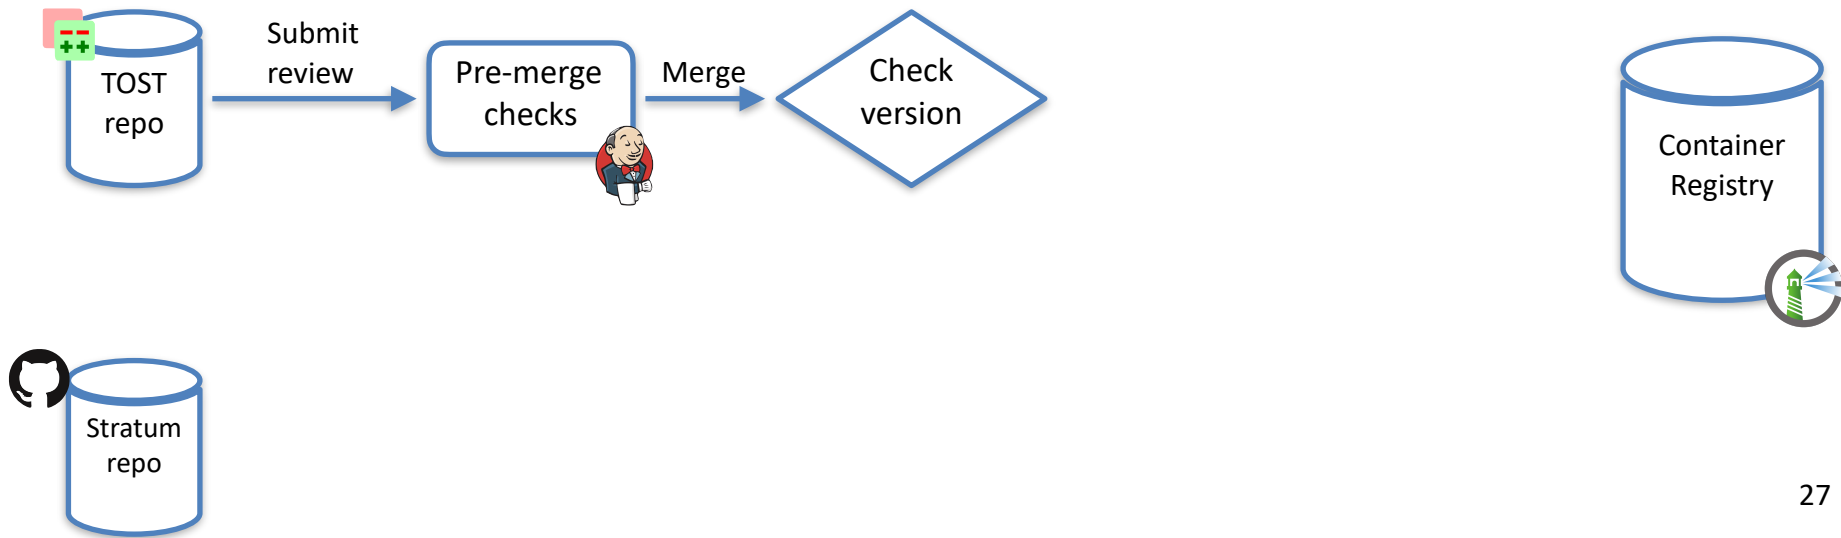
- Git-triggered automated build and release process for Trellis apps and control plane container image
- Build and release Stratum image weekly



Automation - Build and Release



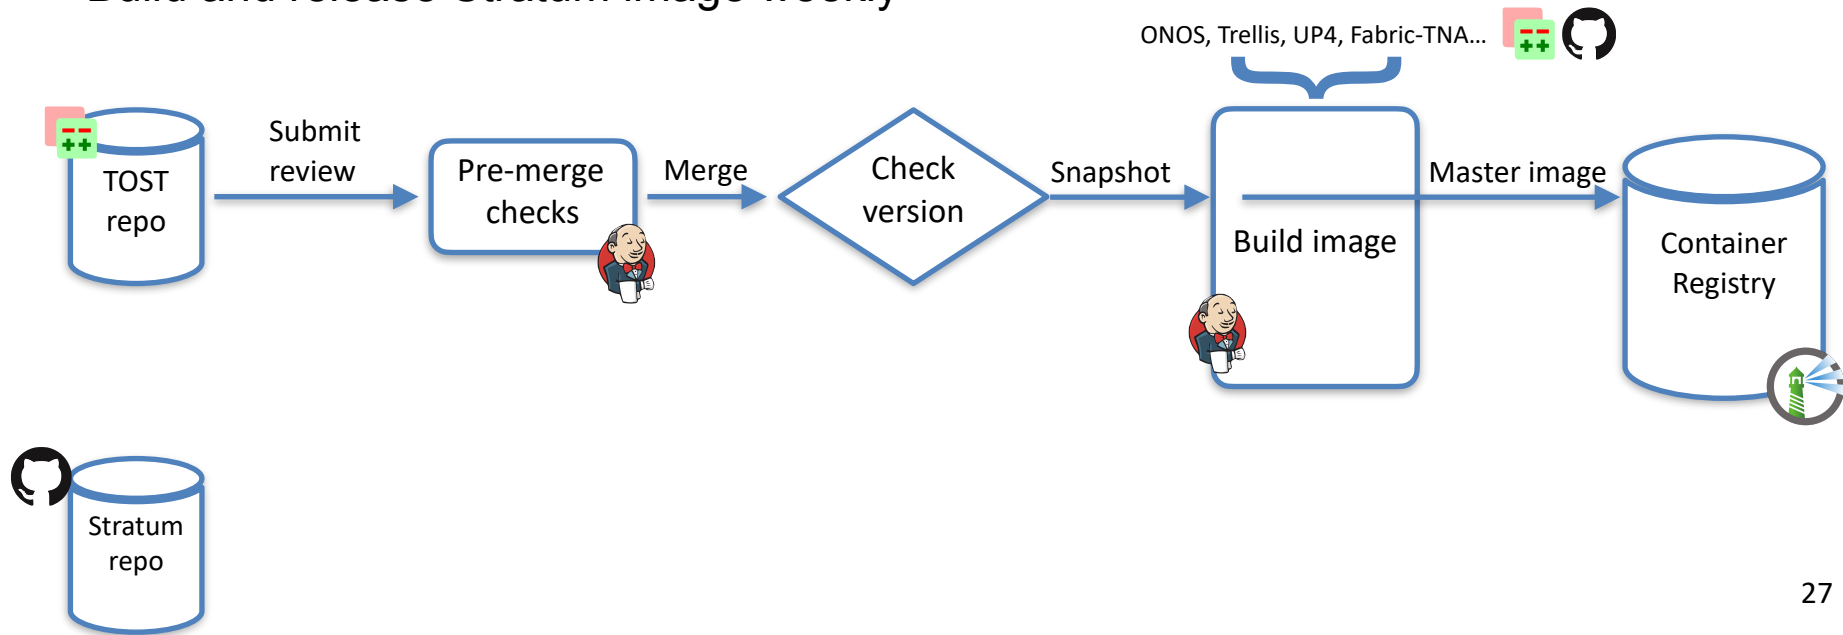
- Git-triggered automated build and release process for Trellis apps and control plane container image
- Build and release Stratum image weekly



Automation - Build and Release



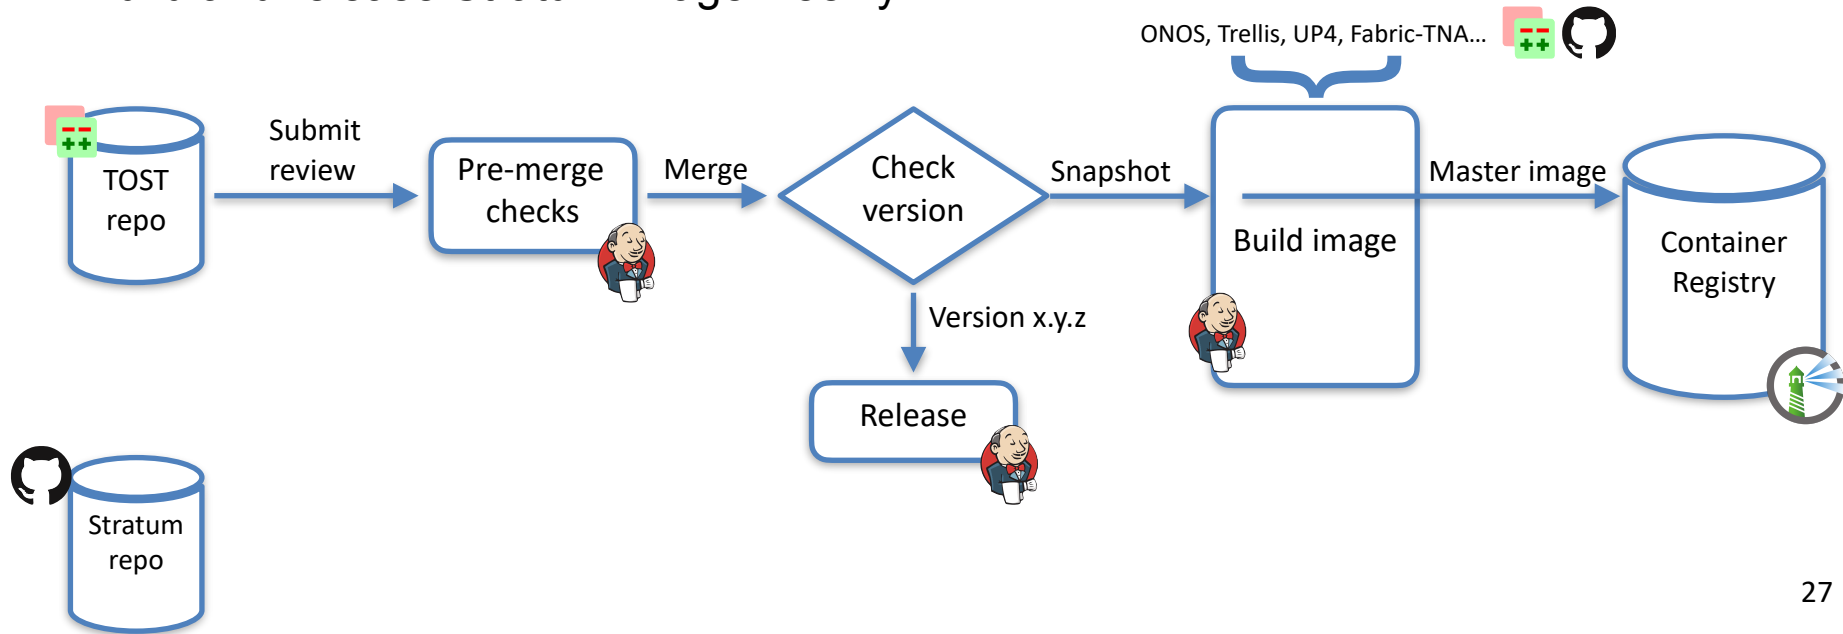
- Git-triggered automated build and release process for Trellis apps and control plane container image
- Build and release Stratum image weekly



Automation - Build and Release



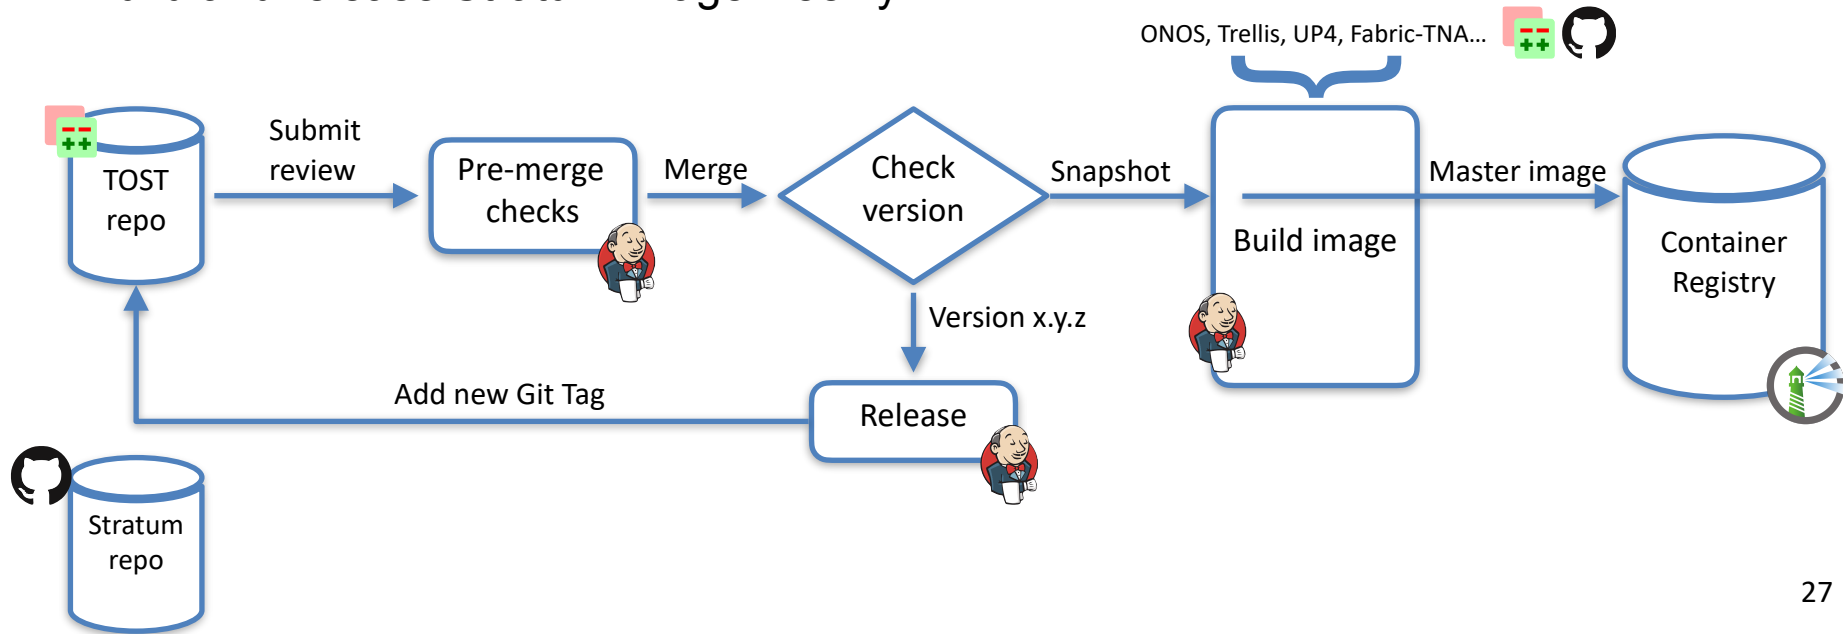
- Git-triggered automated build and release process for Trellis apps and control plane container image
- Build and release Stratum image weekly



Automation - Build and Release



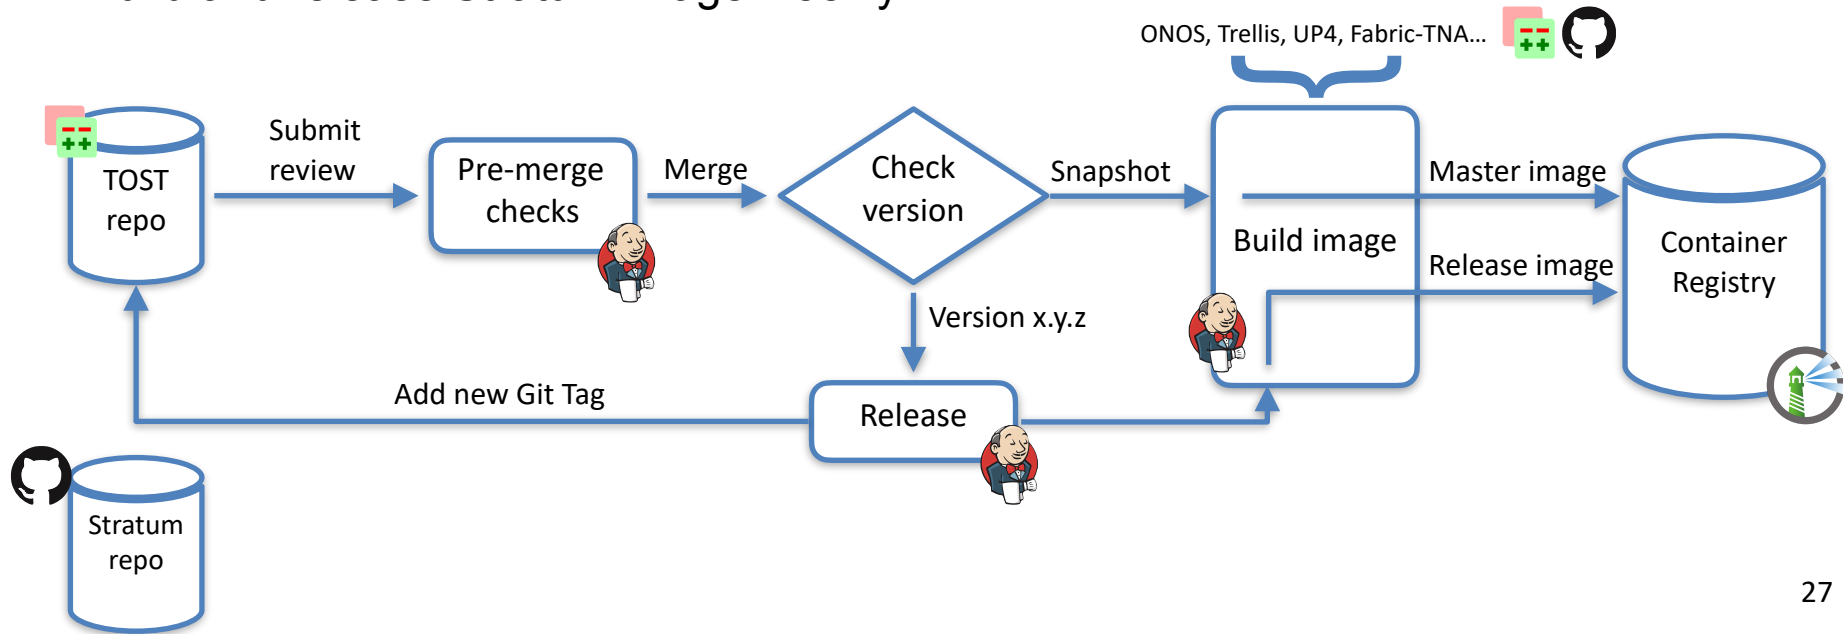
- Git-triggered automated build and release process for Trellis apps and control plane container image
- Build and release Stratum image weekly



Automation - Build and Release



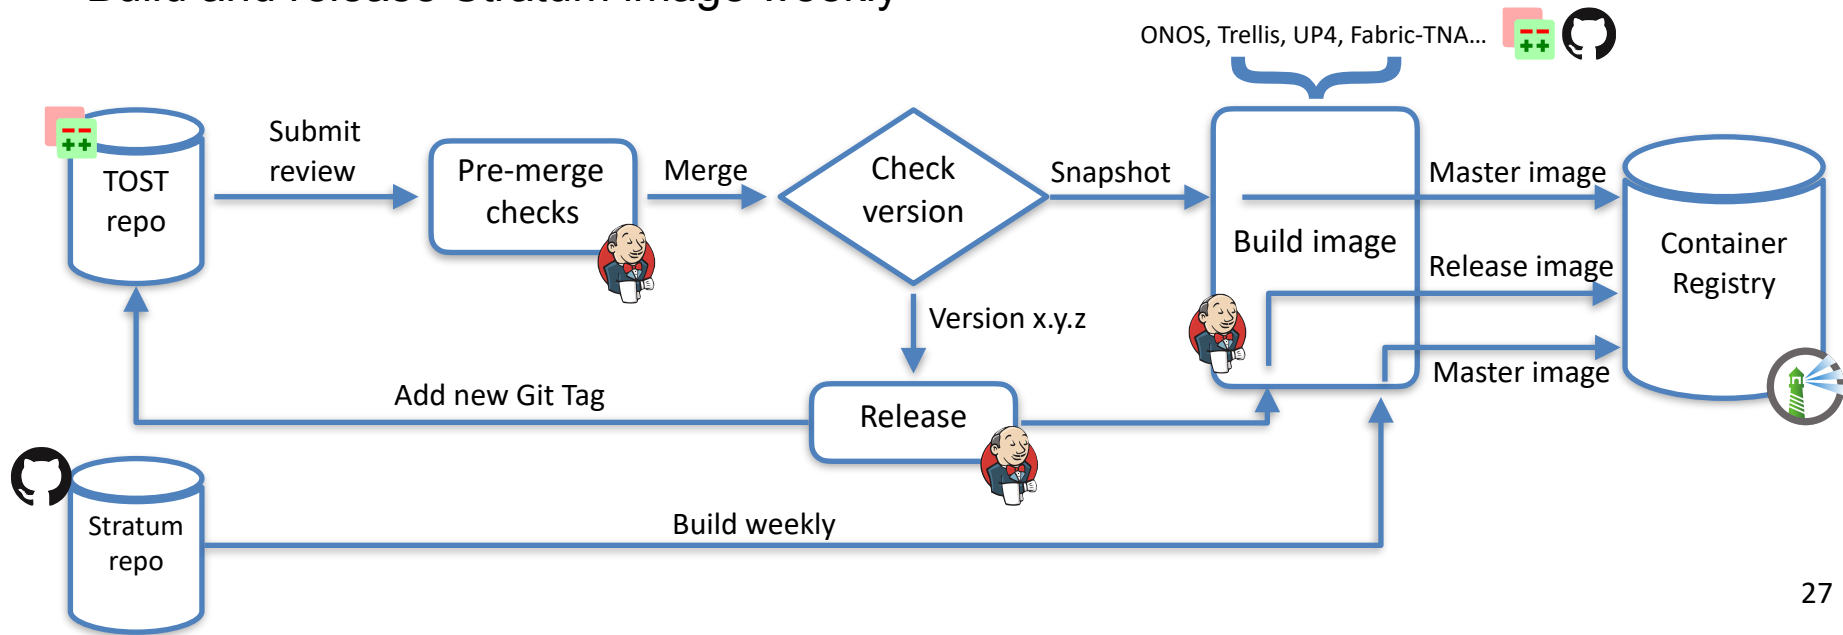
- Git-triggered automated build and release process for Trellis apps and control plane container image
- Build and release Stratum image weekly



Automation - Build and Release



- Git-triggered automated build and release process for Trellis apps and control plane container image
- Build and release Stratum image weekly






- Human-triggered Jenkins pipeline based on Terraform
 - Explicitly-defined helm chart version
 - Get rid of issues seen in Rancher CLI



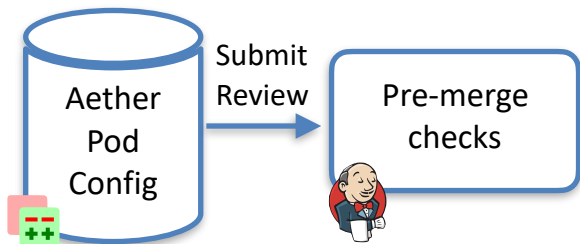


- Human-triggered Jenkins pipeline based on Terraform
 - Explicitly-defined helm chart version
 - Get rid of issues seen in Rancher CLI



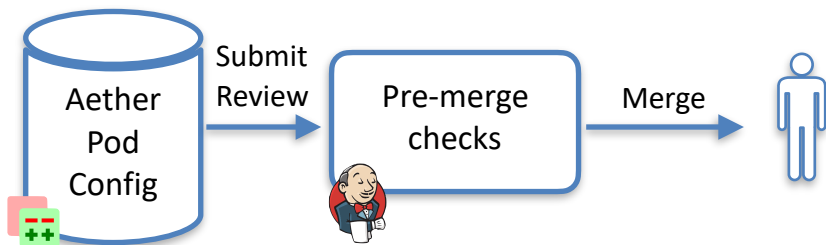


- Human-triggered Jenkins pipeline based on Terraform
 - Explicitly-defined helm chart version
 - Get rid of issues seen in Rancher CLI



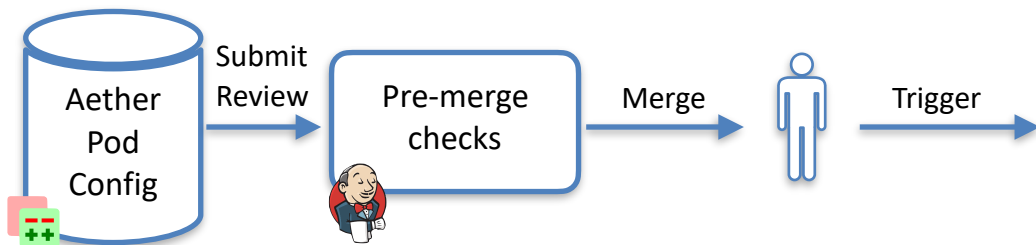


- Human-triggered Jenkins pipeline based on Terraform
 - Explicitly-defined helm chart version
 - Get rid of issues seen in Rancher CLI






- Human-triggered Jenkins pipeline based on Terraform
 - Explicitly-defined helm chart version
 - Get rid of issues seen in Rancher CLI



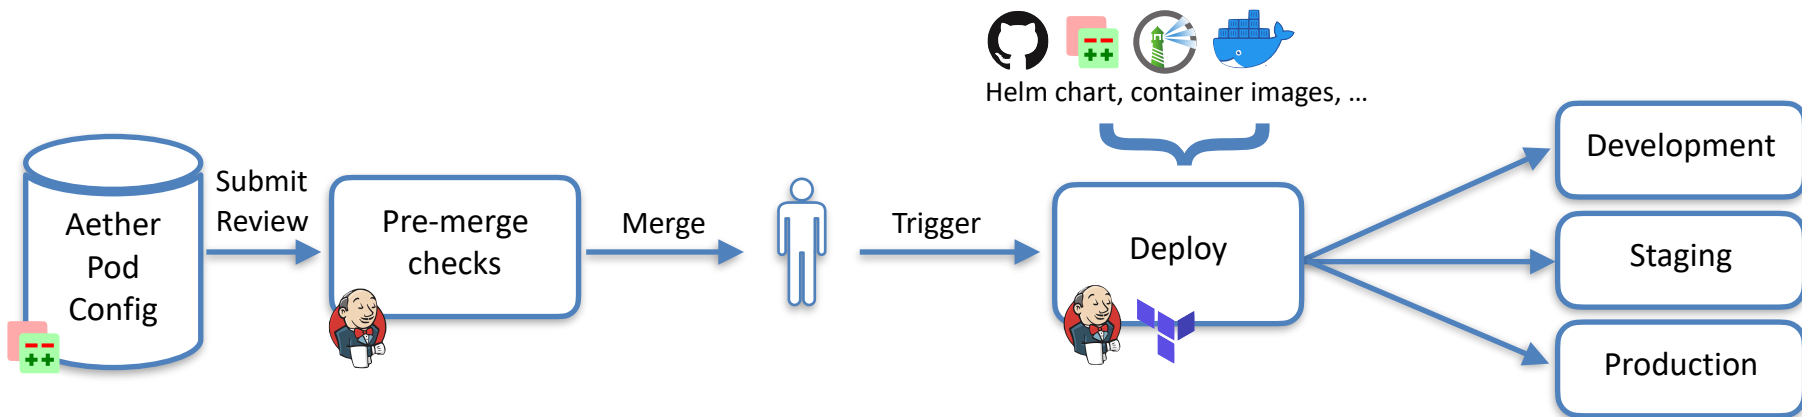


- Human-triggered Jenkins pipeline based on Terraform
 - Explicitly-defined helm chart version
 - Get rid of issues seen in Rancher CLI





- Human-triggered Jenkins pipeline based on Terraform
 - Explicitly-defined helm chart version
 - Get rid of issues seen in Rancher CLI





- Aether - 5G/LTE Enterprise Private Edge Cloud
- P4-based disaggregated UPF
- Highly automated network infrastructure





- Aether
 - 5G/LTE Enterprise Private Edge Cloud
 - <https://aetherproject.org>
- Trellis
 - Leaf-spine SDN fabric for edge
 - <https://opennetworking.org/trellis>
- Stratum
 - Silicon-independent switch operating system for SDN
 - <https://stratumproject.org>
- Slack Channel: onf-community



Thank you