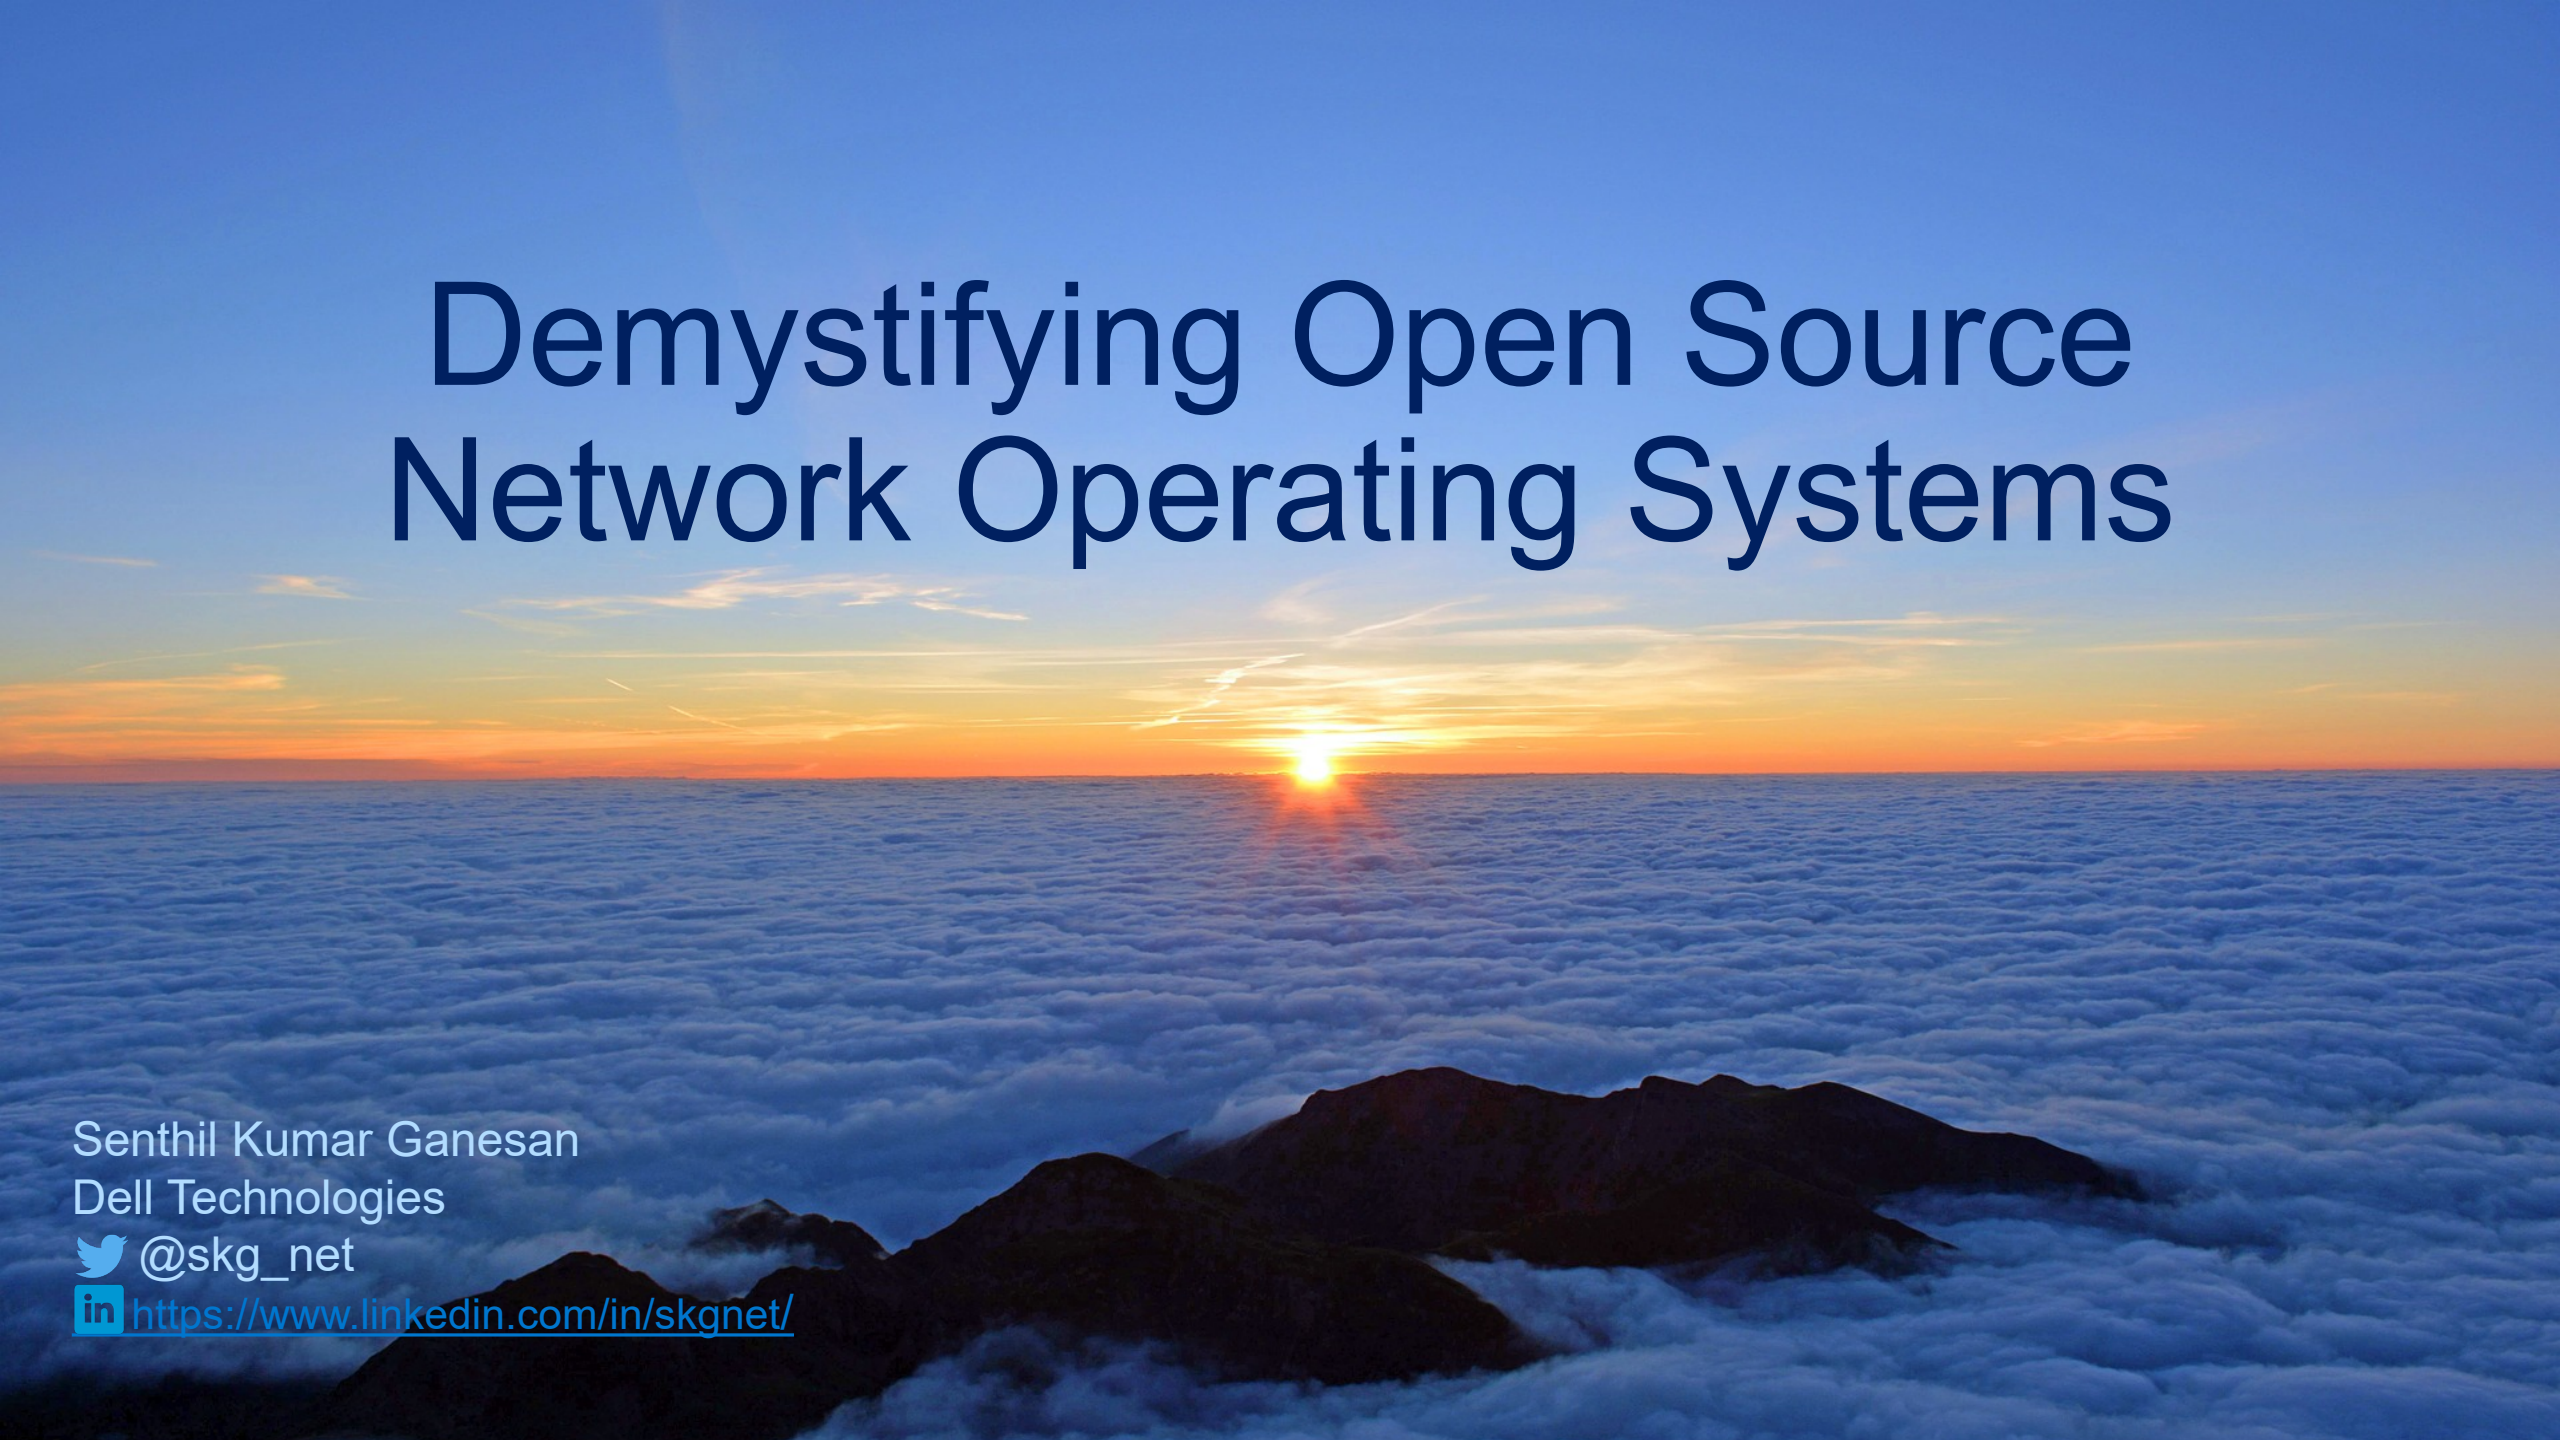

Demystifying Open Source Network Operating Systems

Senthil Kumar Ganesan
Dell Technologies
@skg_net
https://www.linkedin.com/in/skgnet/

Disaggregation

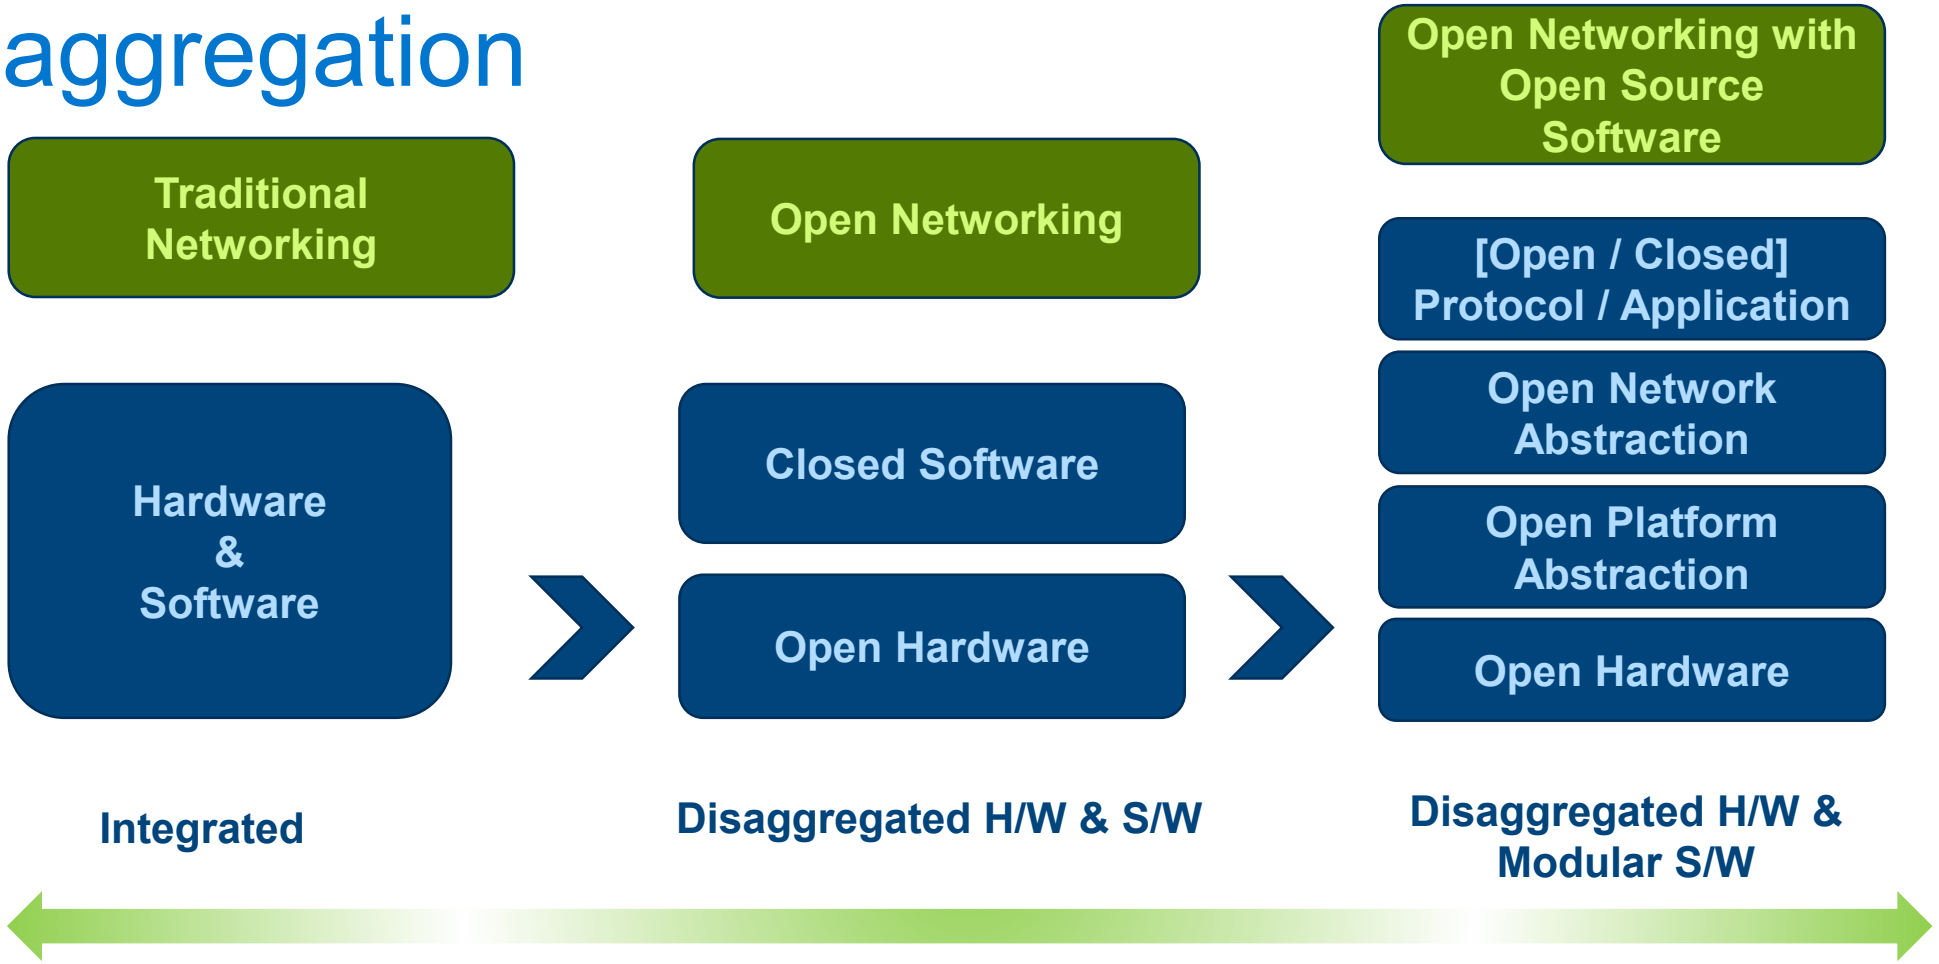

- Rich Feature Set
- Pre Integrated & Tested
- Single Point Support

- H/W Independence
- Power of choice
- Cost Efficiency

- Innovation
- Development Speed
- Cost Efficiency
- Modern Operating Paradigms

Open Networking – Network Operating Systems

Closed Software

Open Hardware

* All product names logos and brands are property of their respective owners

Anatomy of a Network Switch and OS

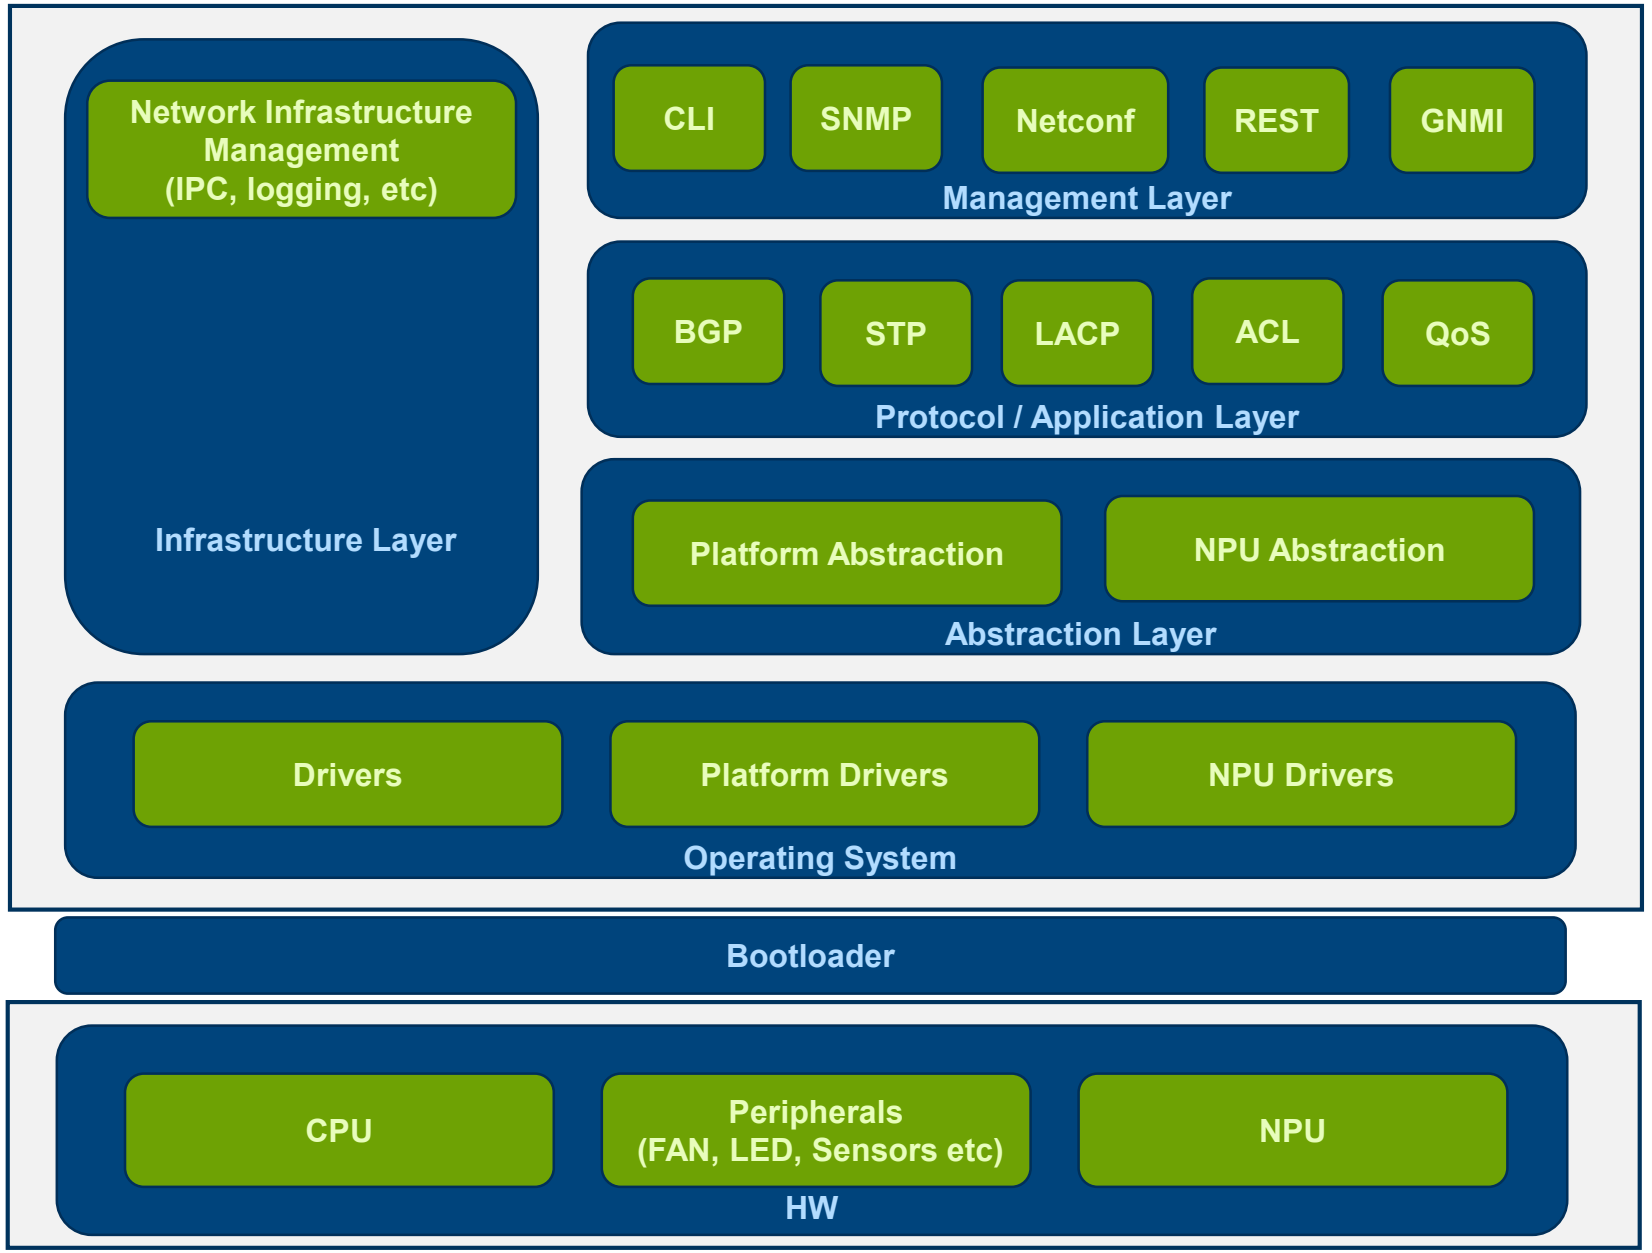

Open Networking Components

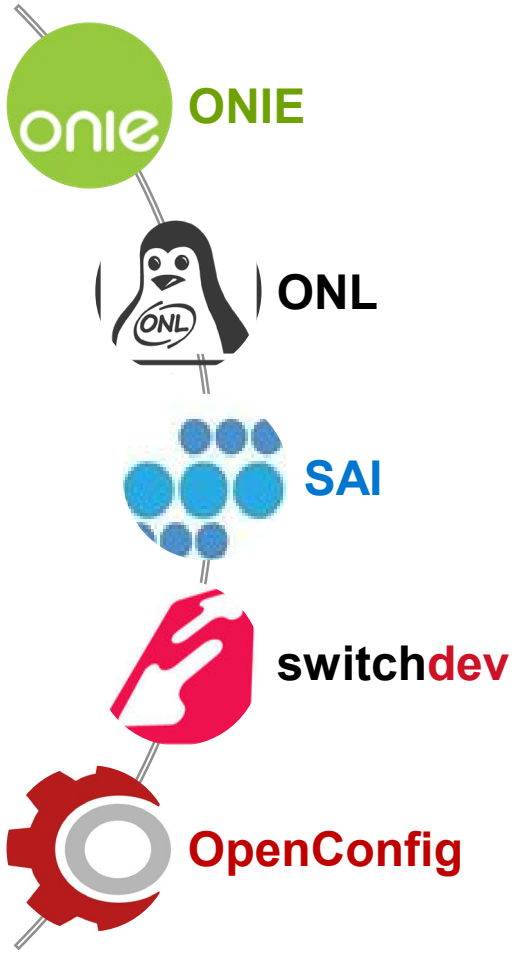

Open Networking Components

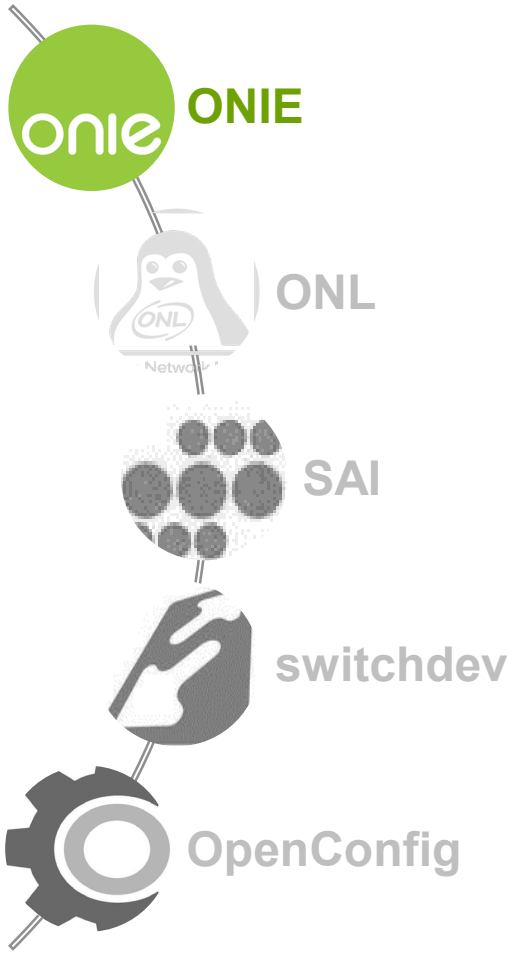

Open Network Install Environment

- Modern, Efficient Network Installer
- Open Compute Subproject, initially contributed by Cumulus
- Defacto standard boot loader, its a small operating system, pre-installed as firmware on bare metal network switches.
- Provides an environment for automated operating system provisioning
- ~161 devices currently supported.

Open Networking Components

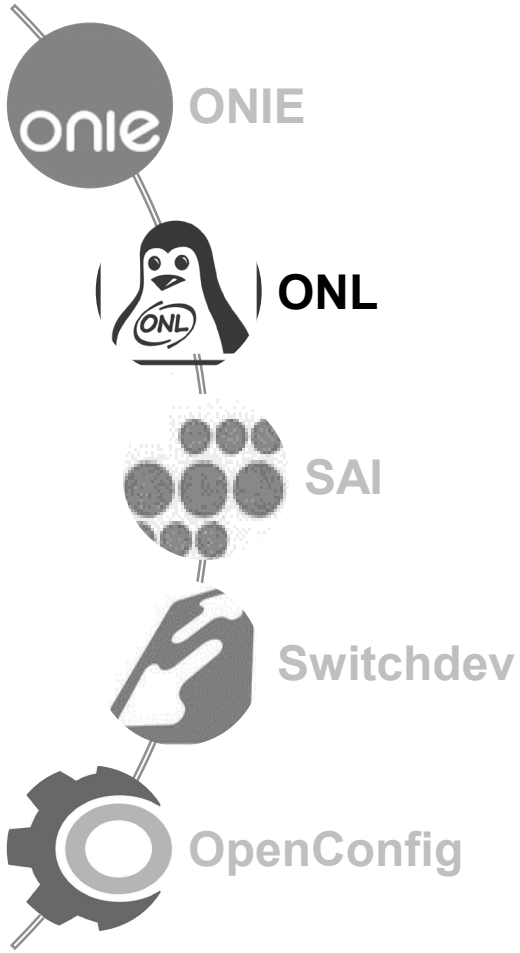

Open Networking Linux

- A Linux Distribution for Open Networking Switches
- Reference NOS for the Open Compute Project (OCP)
- ~120 H/W Platforms supported, becoming the defacto standard.
- Uses Debian and stock LTS Linux kernel
- Provides Platform Abstraction, via ONLP API

Open Networking Components

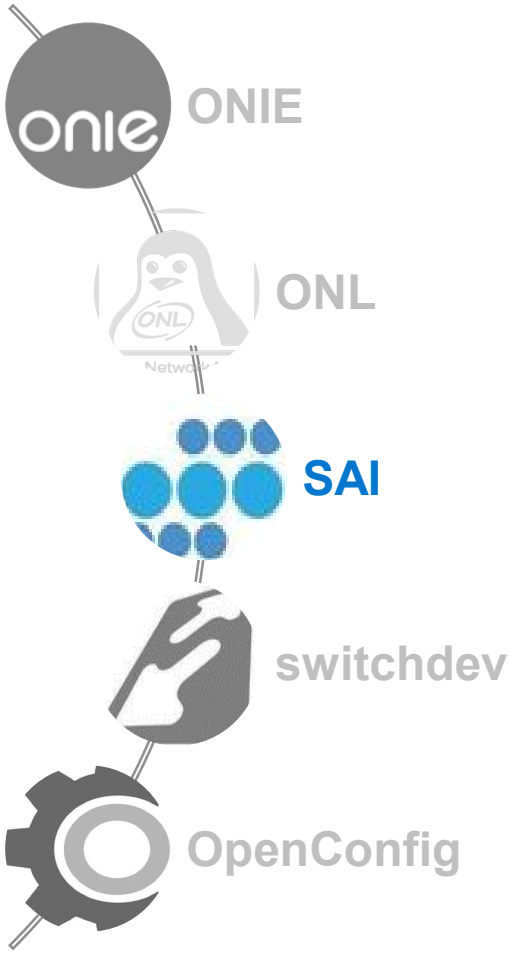

Switch Abstraction Interface

- SAI defines API to provide a vendor-independent way of controlling forwarding elements, such as a switching ASIC, an NPU or a software switch in a uniform manner.
- SAI helps easily port new ASIC by running the same application stack on all the hardware, enabled by a simple, consistent programming interface.
- CRUD operation over extensible Entity/Attribute/Value data model
- Provides Network Abstraction for all major ASIC vendors (Barefoot, Boardcom, Cavium, Innovium, Mellanox)

Open Networking Components

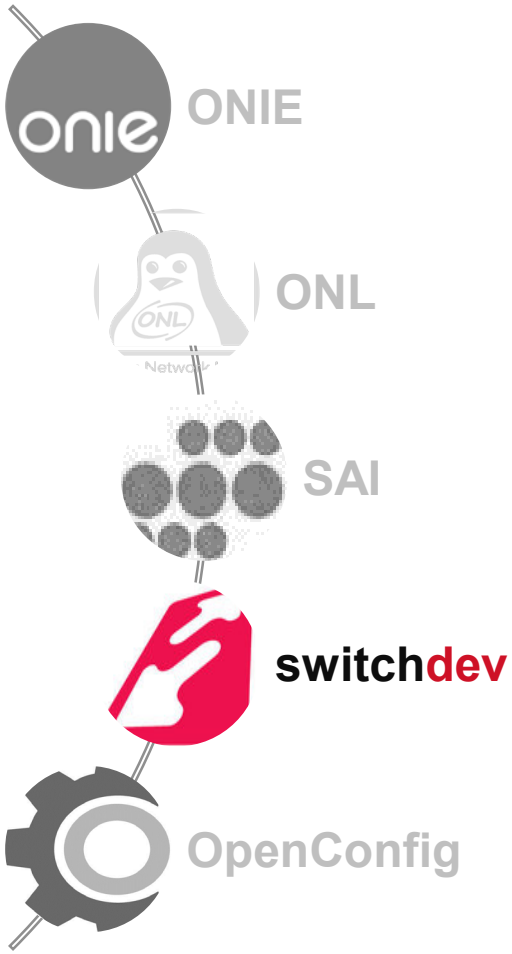

Switchdev

- The Ethernet switch device driver model (switchdev) is an in-kernel driver model for switch devices which offload the forwarding (data) plane from the kernel.
- Provides Network Abstraction – currently only Mellanox is supported
- Offloads L2 & L3 from Linux Kernel and aims to re-use the same set of linux network tool set for switches

Open Networking Components

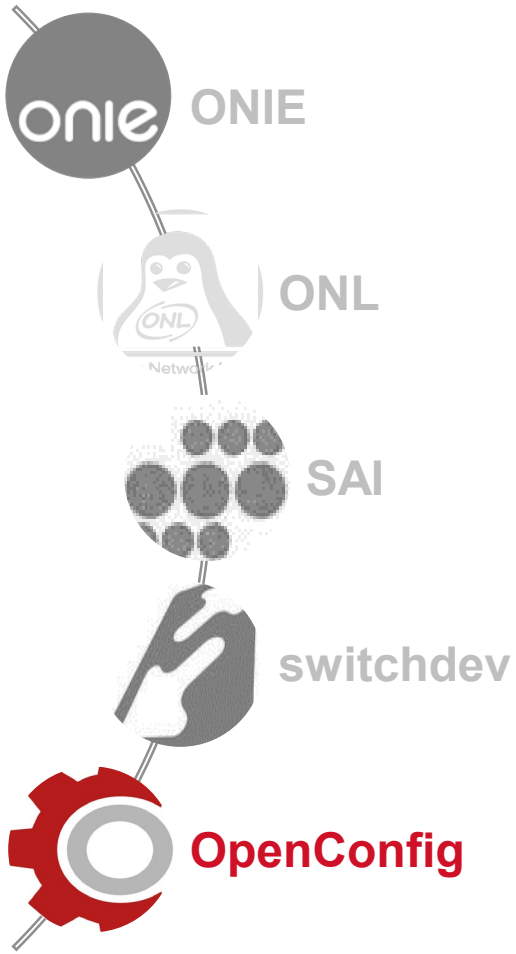

OpenConfig

- OpenConfig provides a consistent set of vendor-neutral data models (written in YANG) based on actual operational needs from use cases and requirements from multiple network operators.
- Openconfig is supported by most of the major networking vendors.
- Provides management layer abstraction – thus enabling common management or controller application to be written.
- Uses modern RPC - NETCONF, RESTCONF, GNMI - GRPC

Anatomy of a Network Switch and OS

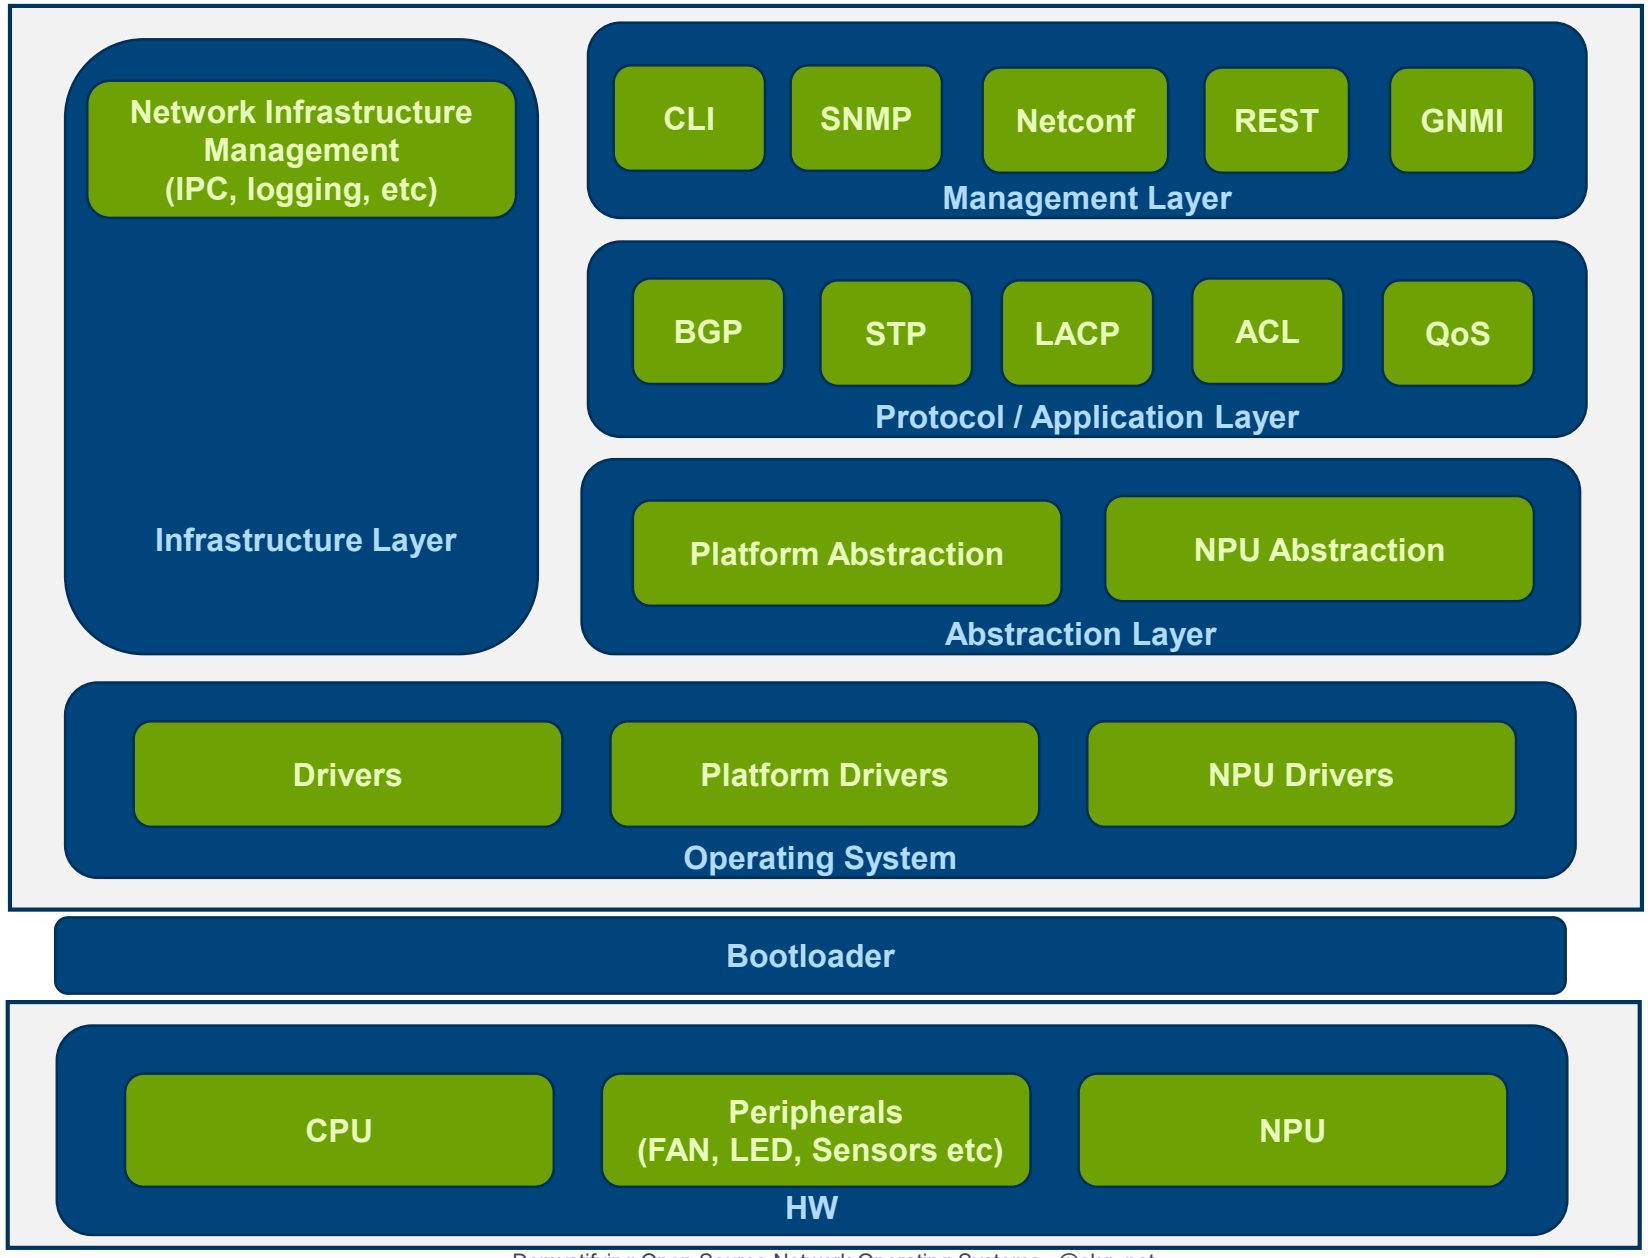

Anatomy of a Network Switch and OS

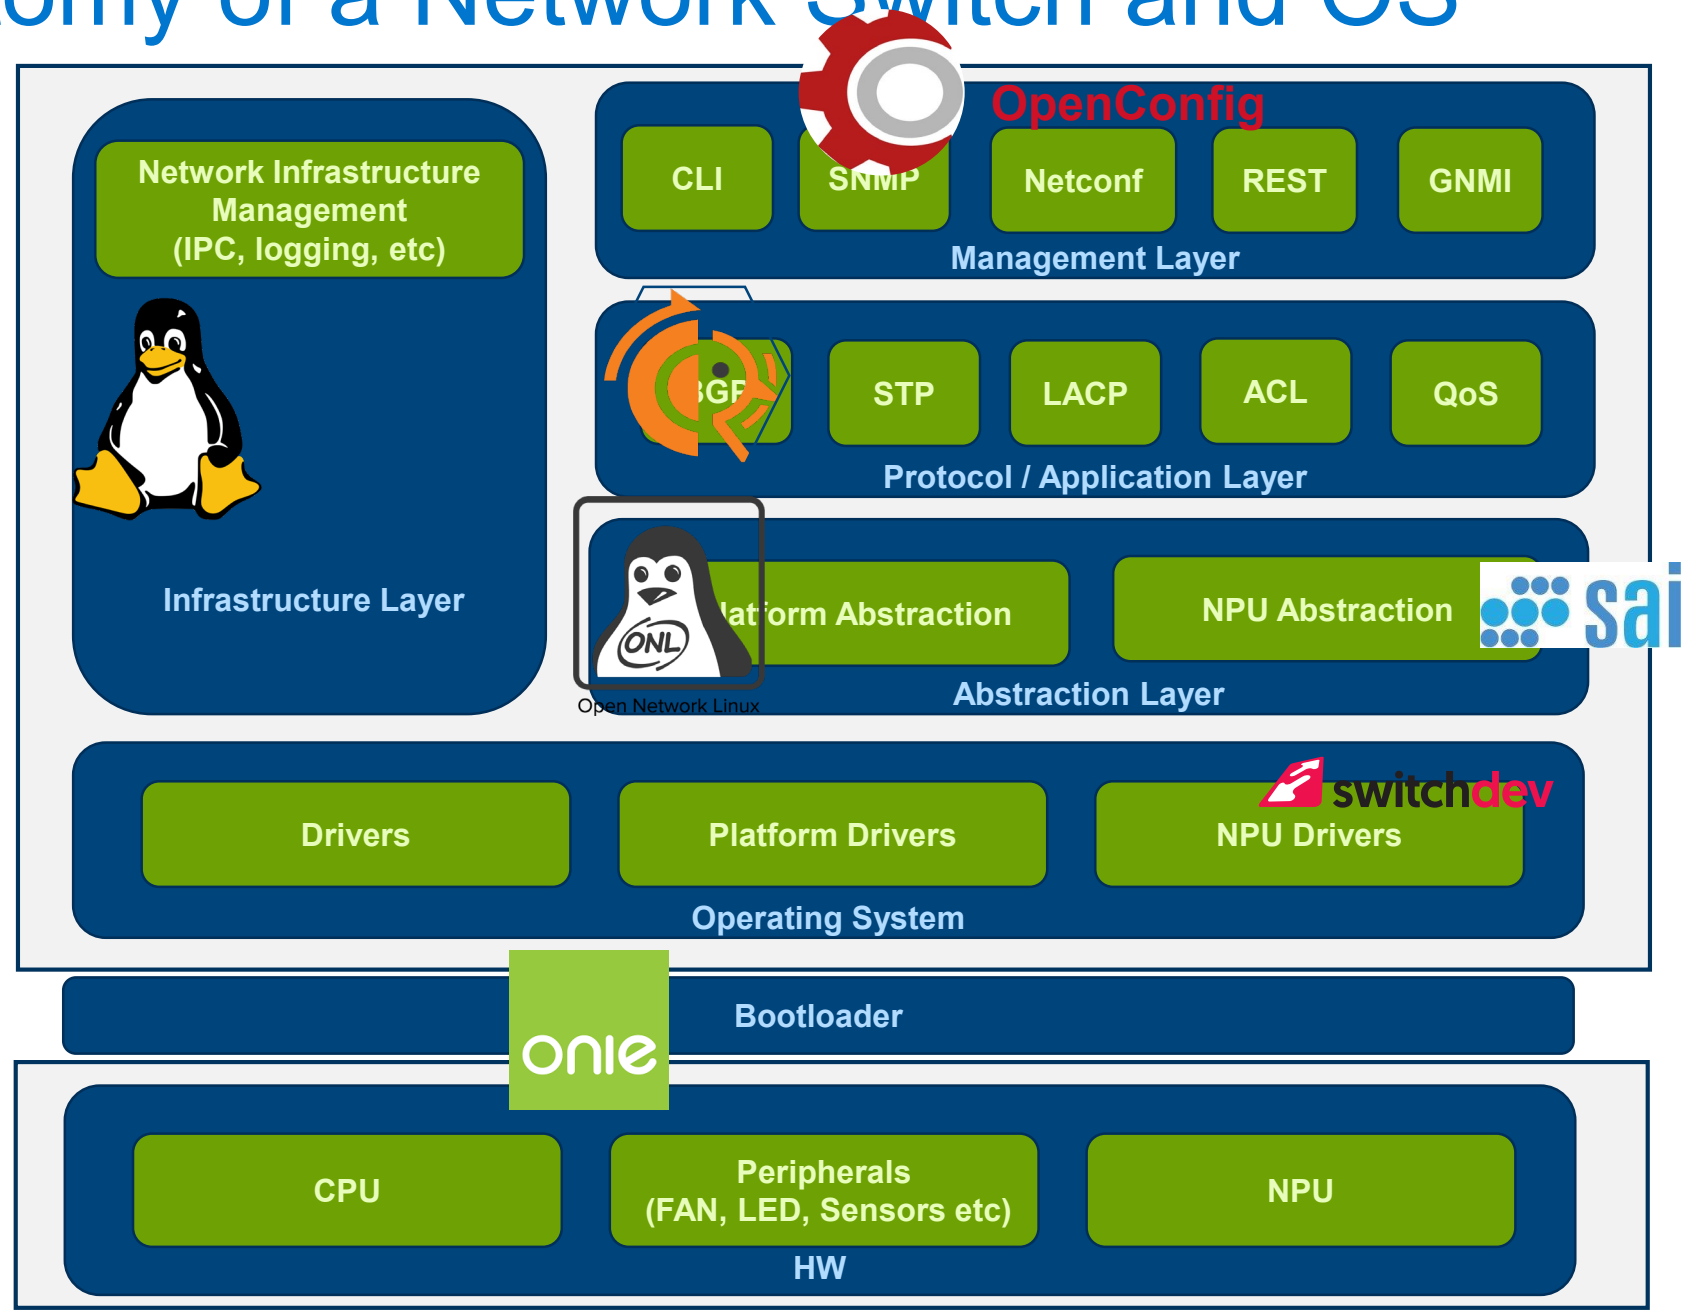

Open Source Network Operating Systems

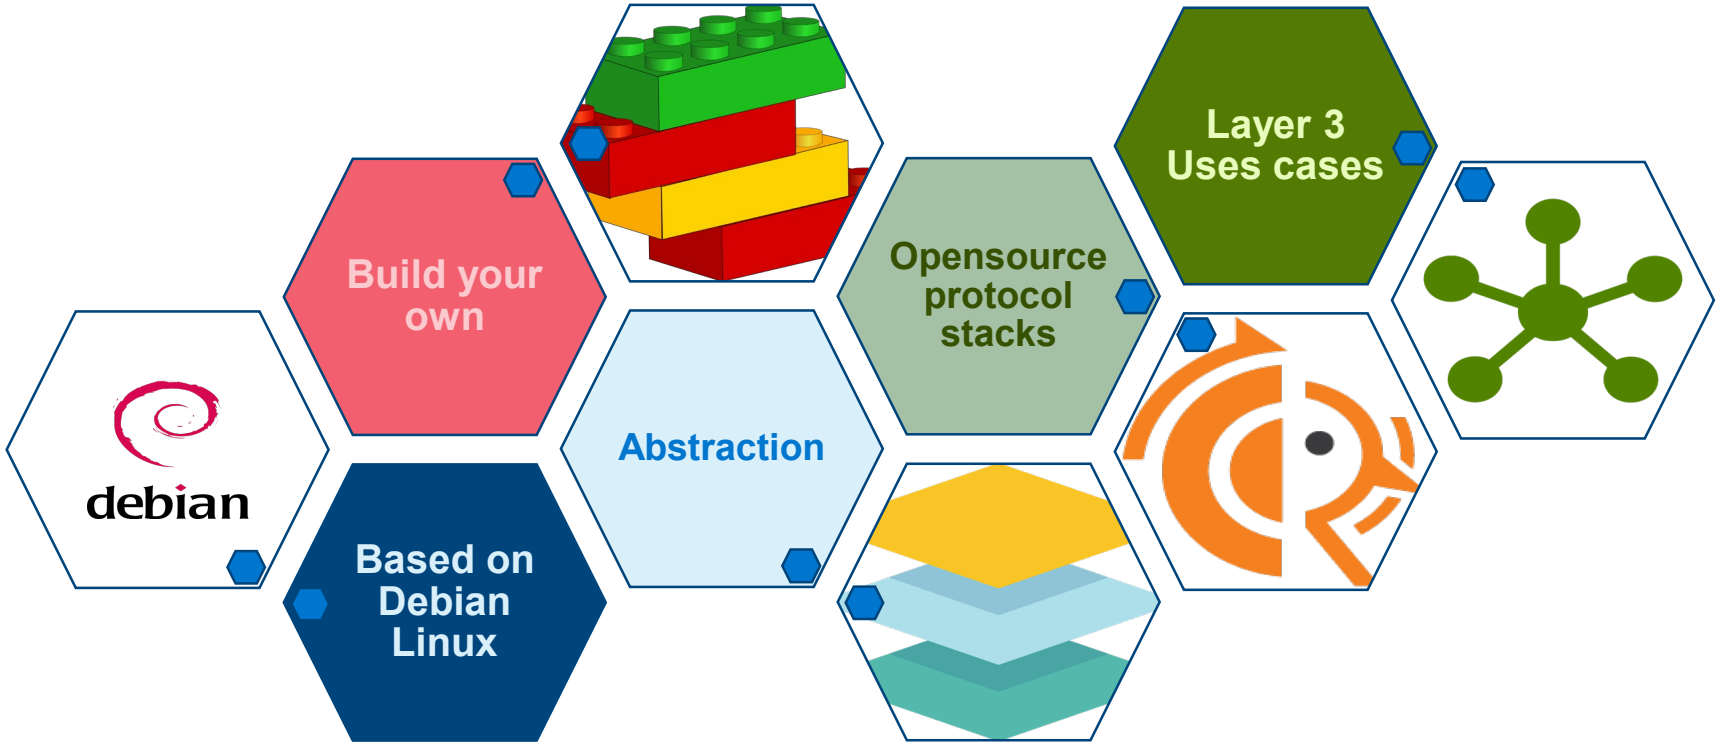

Open Source Network Operating Systems

Winding Down

LibreSwitch

FlexSwitch

Present

Open Network Linux

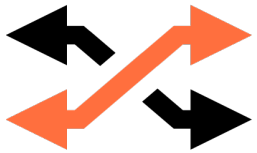

Open Switch

STRATUM

Upcoming

* All product names logos and brands are property of their respective owners

Open Networking Linux (ONL)

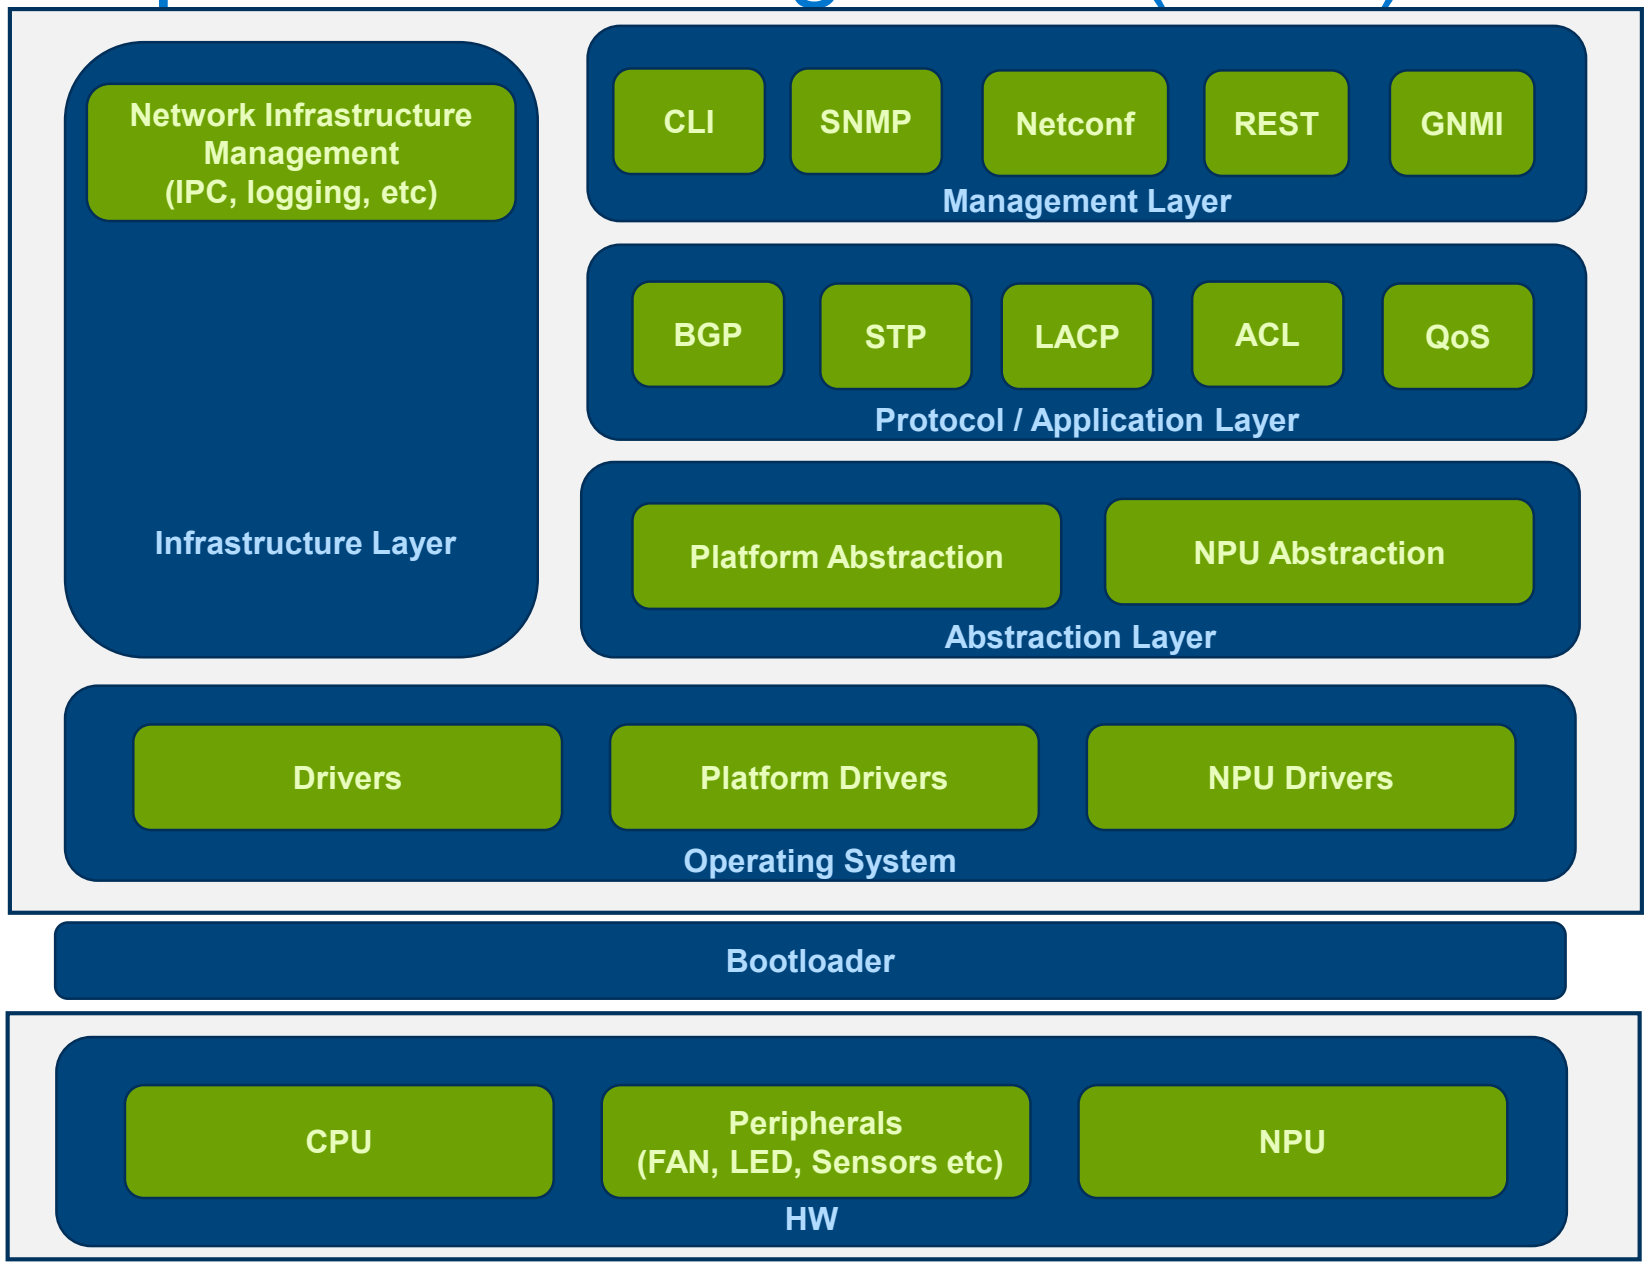

Open Network Linux

- Provides Platform Abstraction
- Support Multiple NPU abstraction. OpenNSL, SAI, OF_DPA.
- Supports Multiple Control Plane.
- Thin OS with Controller (ONOS)
- Thick OS with FRR/ORC
- Becoming the de-facto standard for platform abstraction ~120 platforms (BigSwitch, Startum, Arrcus, SnapRoute etc)
- Use Cases:
  - CORD Leaf Spine Fabric (ONOS / Indigo Agent)
  - EVPN with GoBGP / Zebra / ORC / Open NSL
  - FBOSS Leaf Spine Fabric

Open Networking Linux (ONL)

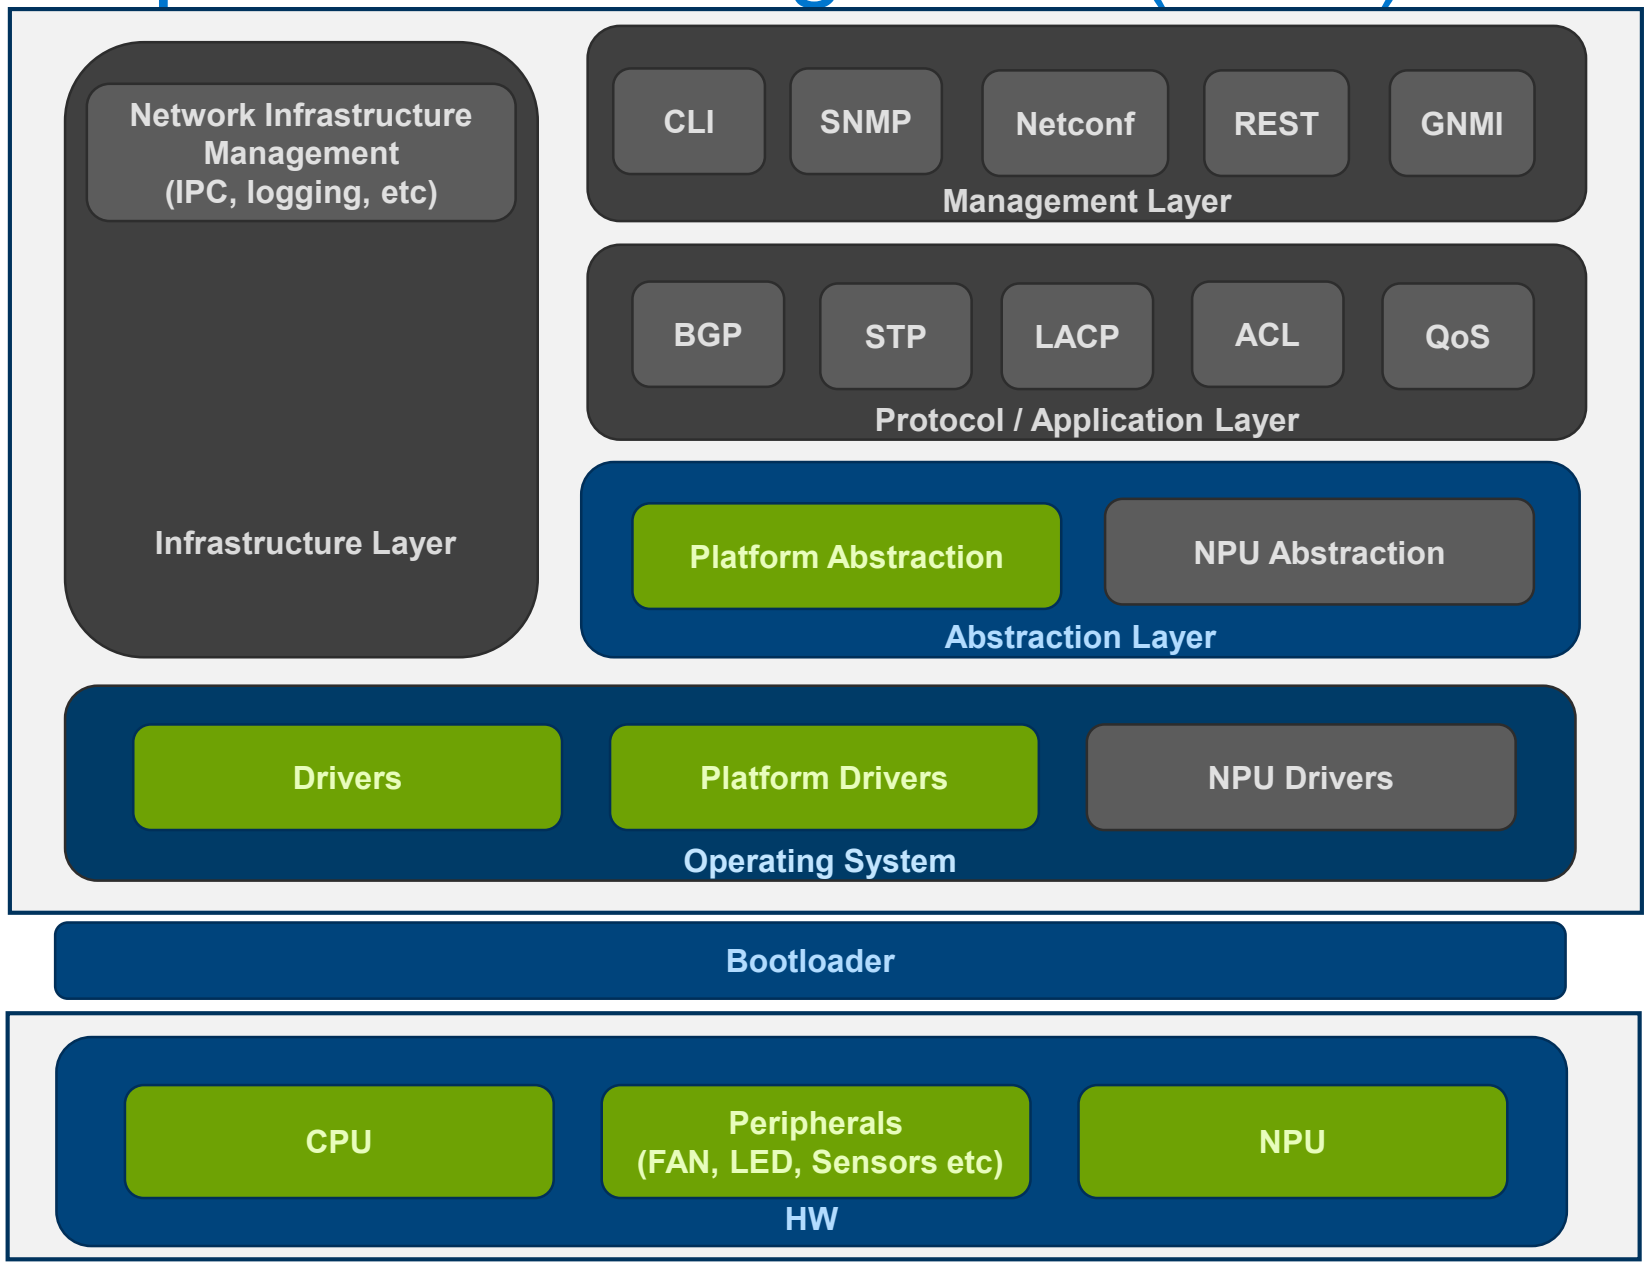

Open Network Linux

- Provides Platform Abstraction
- Support Multiple NPU abstraction. OpenNSL, SAI, OF_DPA.
- Supports Multiple Control Plane.
- Thin OS with Controller (ONOS)
- Thick OS with FRR/ORC
- Becoming the de-facto standard for platform abstraction ~120 platforms (BigSwitch, Startum, Arrcus, SnapRoute etc)
- Use Cases:
  - CORD Leaf Spine Fabric (ONOS / Indigo Agent)
  - EVPN with GoBGP / Zebra / ORC / Open NSL
  - FBOSS Leaf Spine Fabric

Open Networking Linux (ONL)

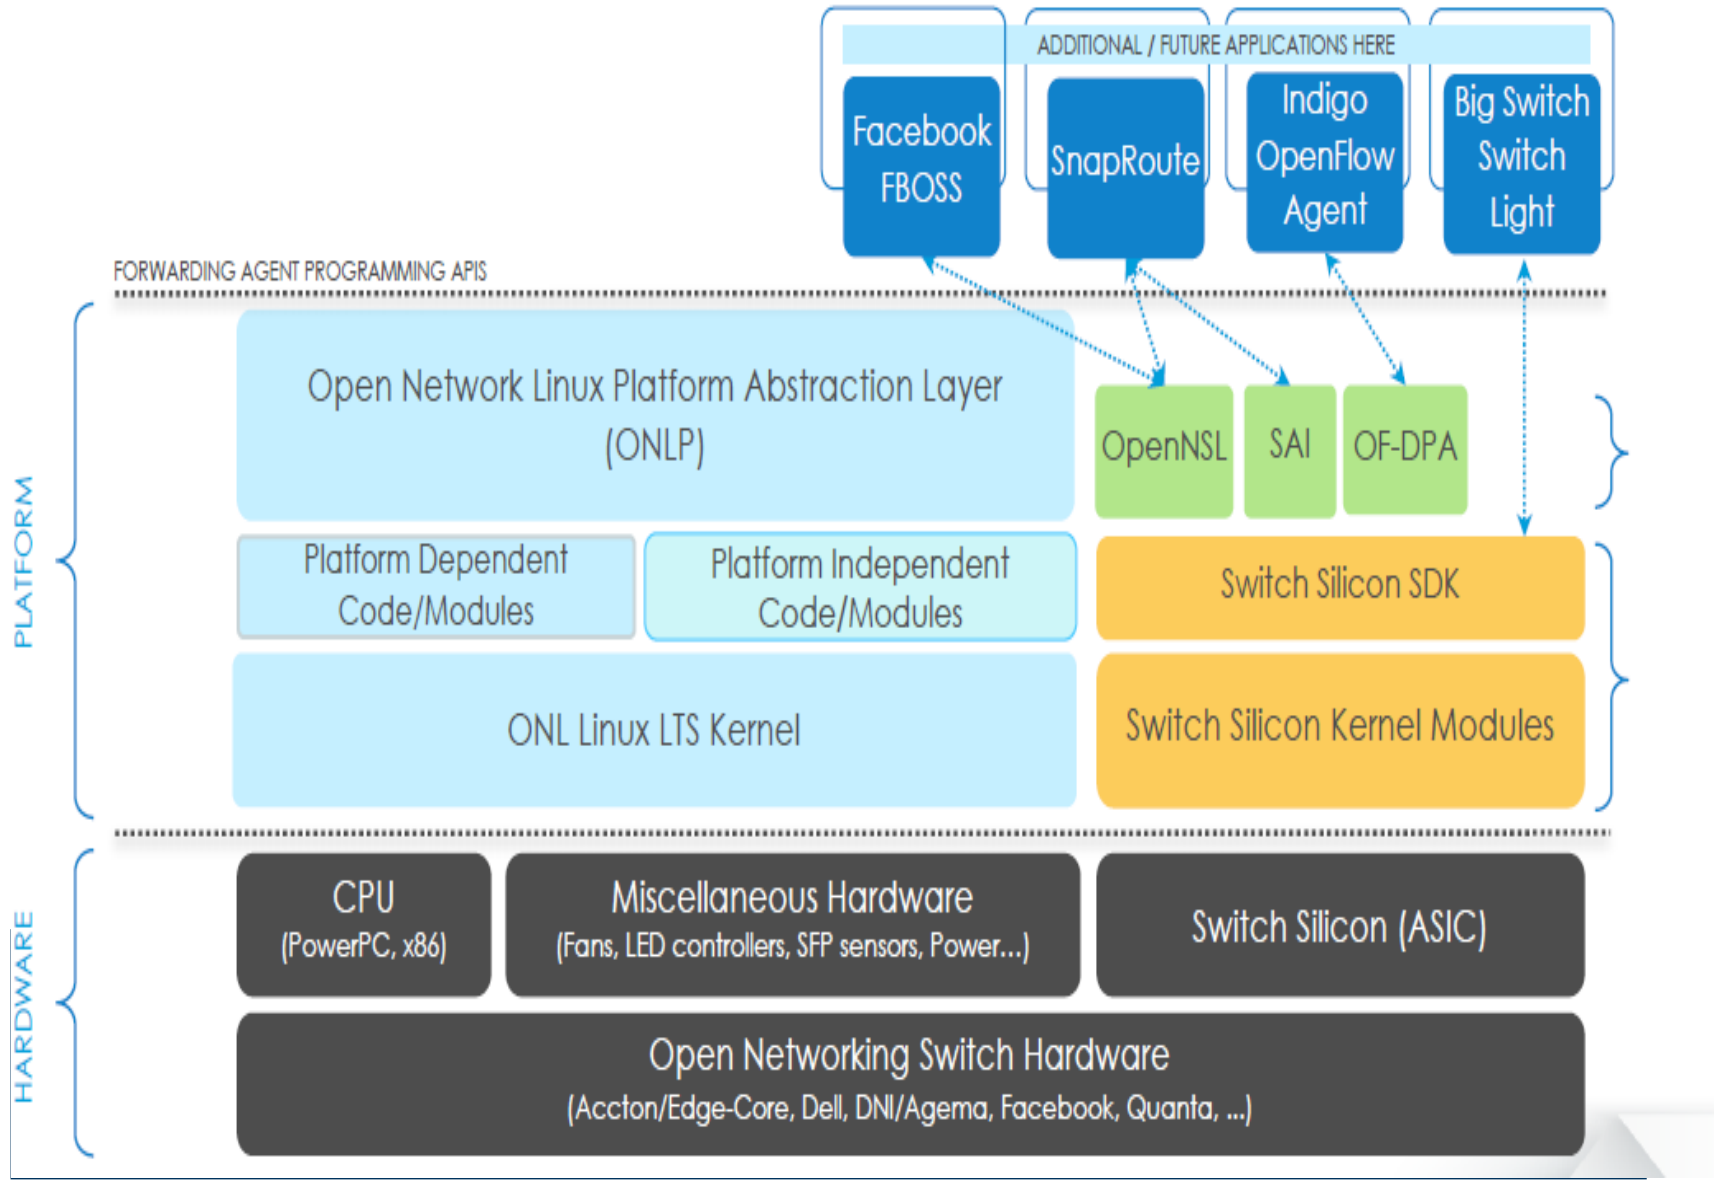

Open Network Linux

- Provides Platform Abstraction
- Support Multiple NPU abstraction. OpenNSL, SAI, OF_DPA.
- Supports Multiple Control Plane.
- Thin OS with Controller (ONOS)
- Thick OS with FRR/ORC
- Becoming the de-facto standard for platform abstraction ~120 platforms (BigSwitch, Startum, Arrcus, SnapRoute etc)
- Use Cases:
  - CORD Leaf Spine Fabric (ONOS / Indigo Agent)
  - EVPN with GoBGP / Zebra / ORC / Open NSL
  - FBOSS Leaf Spine Fabric

OPX Open Switch THE LINUX FOUNDATION

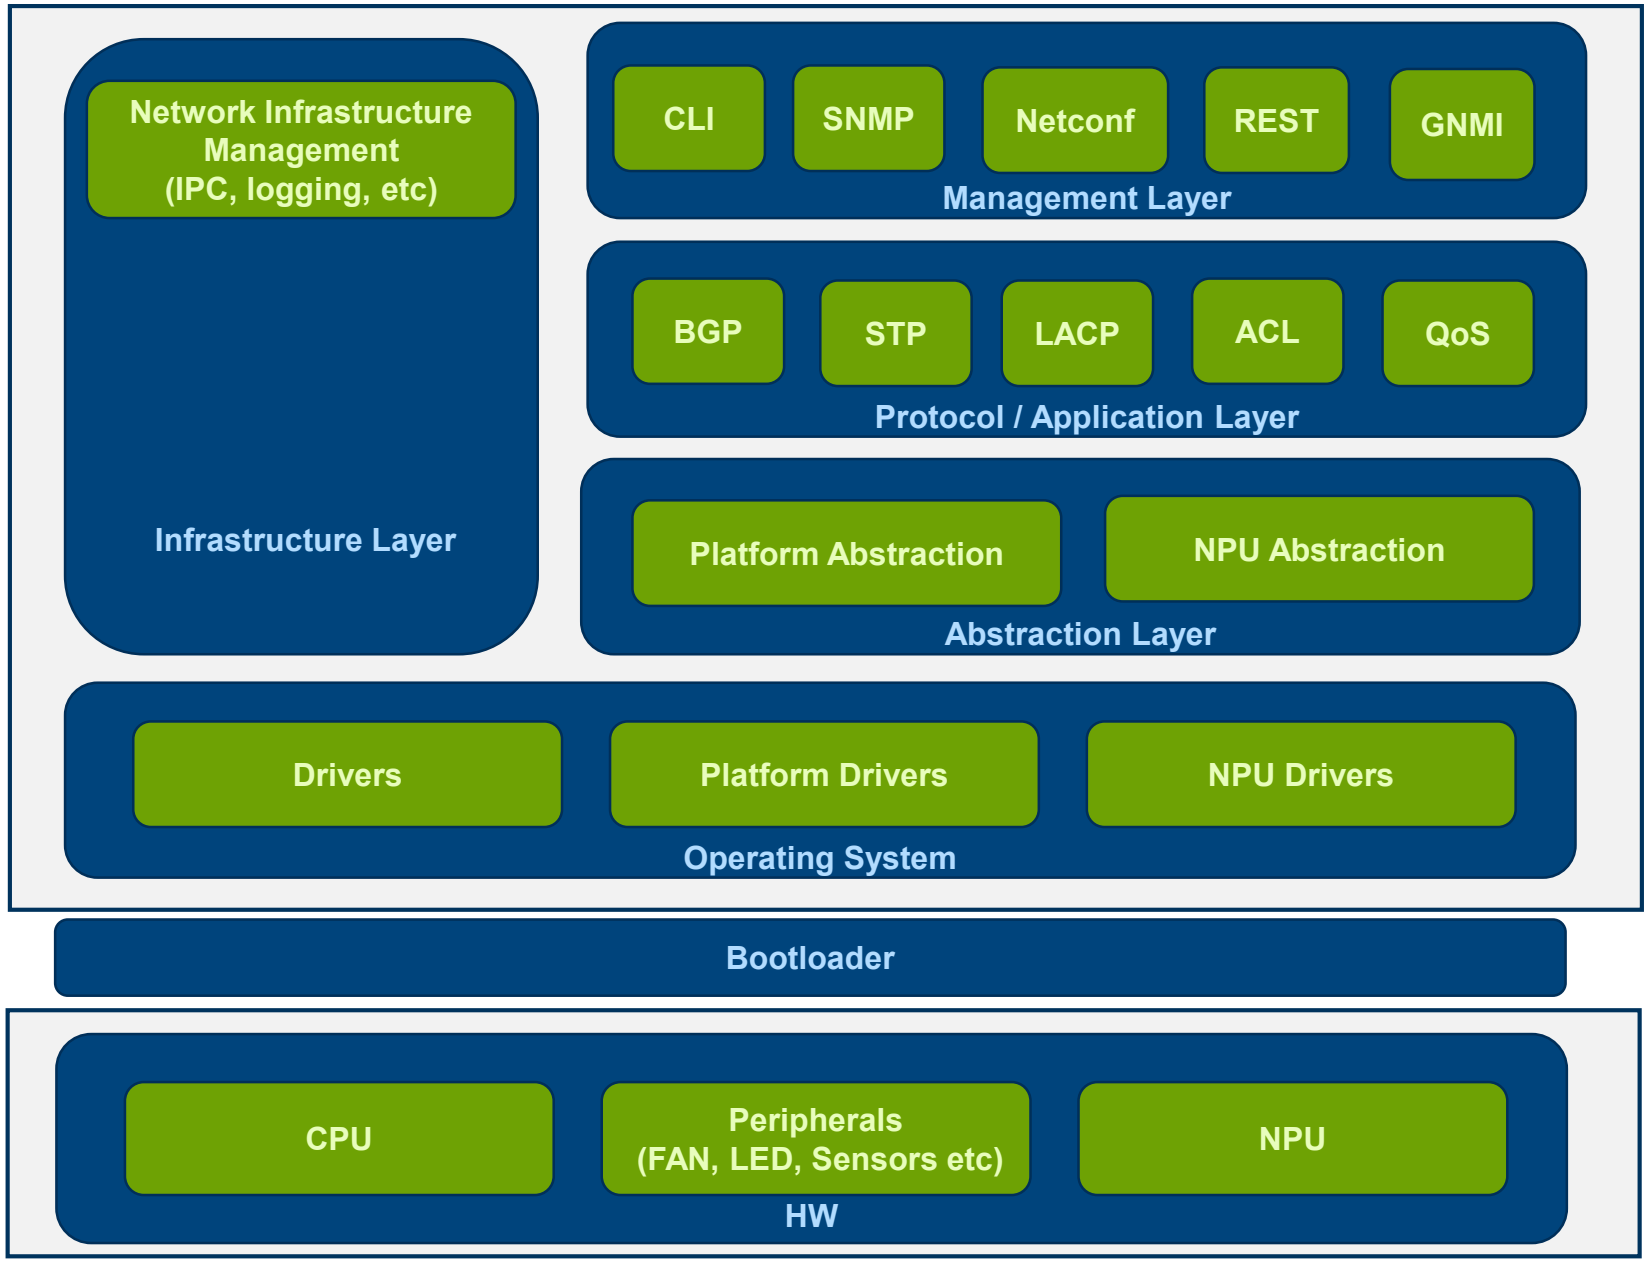

OVS Open vSwitch

- Provides Platform Abstraction and Network Abstraction
- By default supports the Linux Protocol & Application Stacks
- Used as base operating system for Dell EMC Enterprise Operating OS10
- Can be leveraged to write custom application using Control Plane Services API.
- Production Ready, Deployed in multiple Verizon, AWINX etc.
- Part of Linux Foundation

OPX Open Switch THE LINUX FOUNDATION

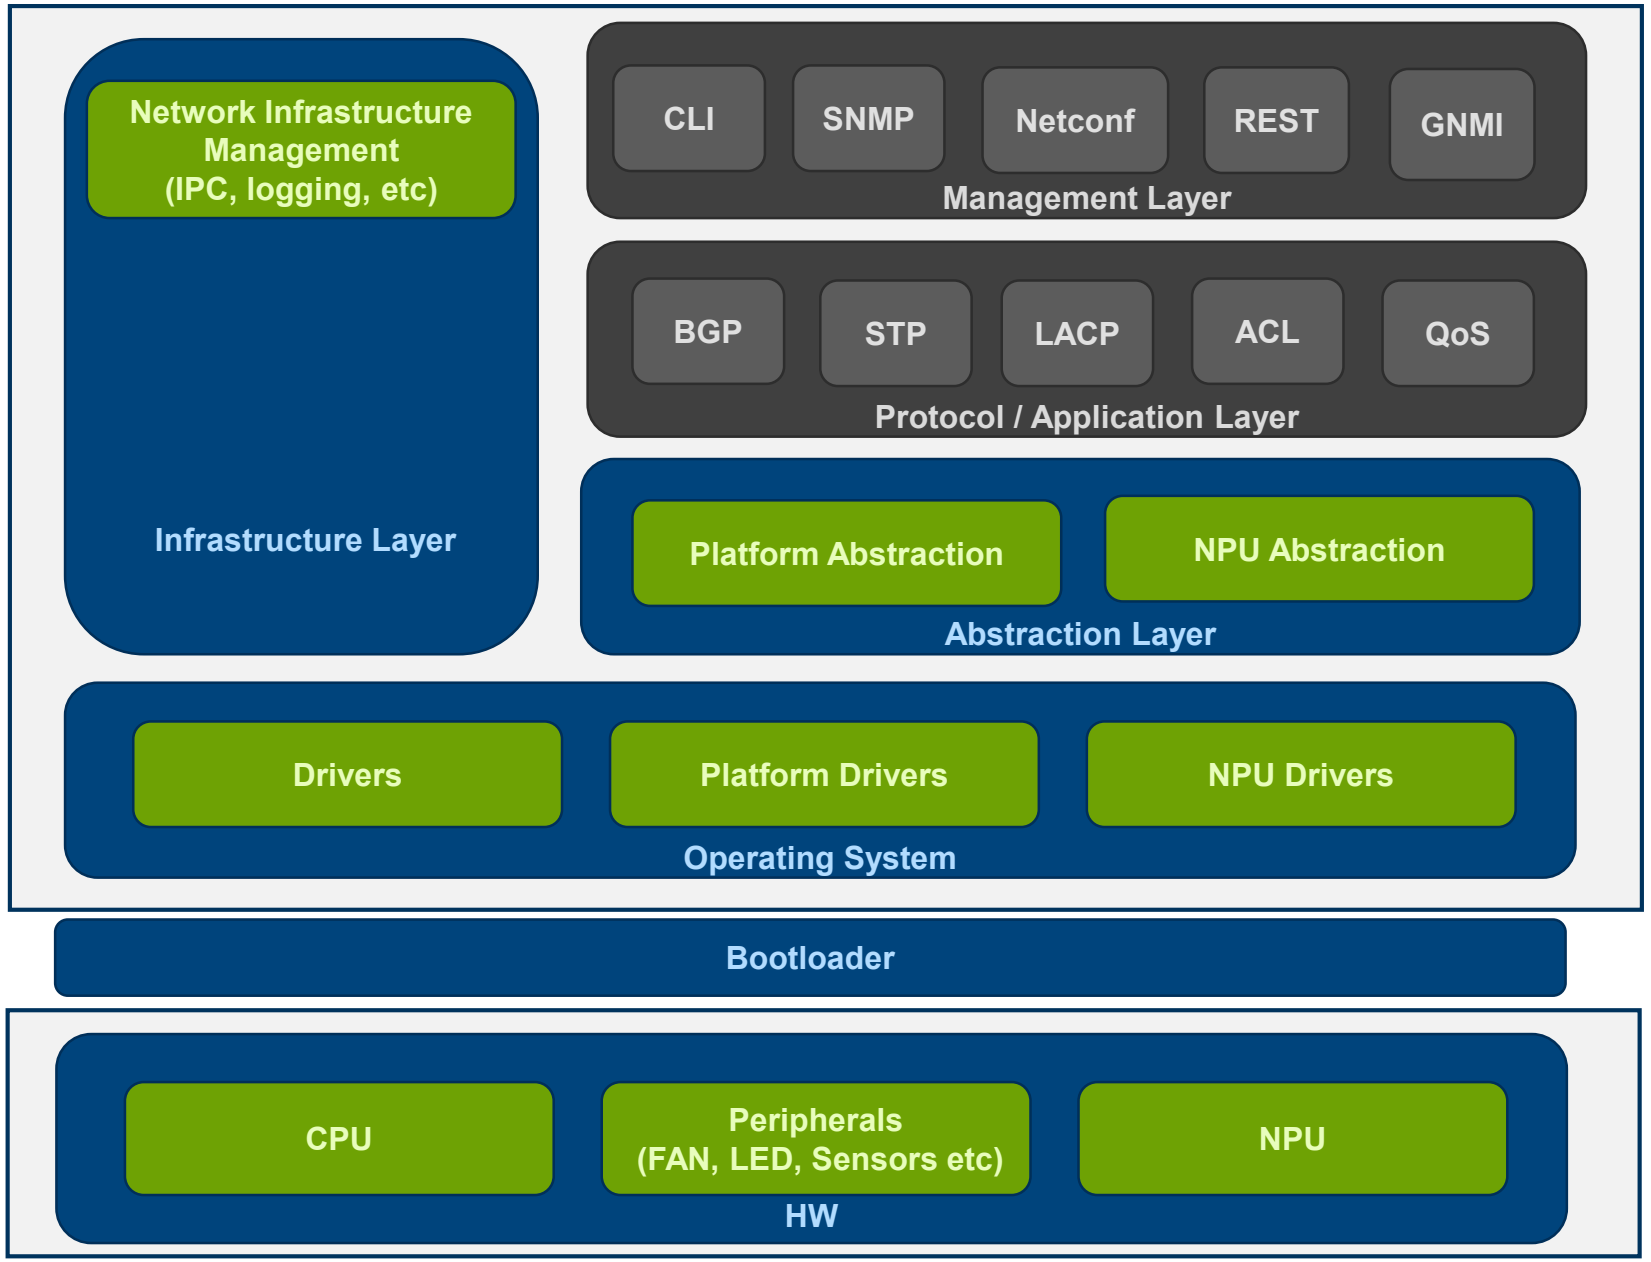

OVS Open vSwitch

- Provides Platform Abstraction and Network Abstraction
- By default supports the Linux Protocol & Application Stacks
- Used as base operating system for Dell EMC Enterprise Operating OS10
- Can be leveraged to write custom application using Control Plane Services API.
- Production Ready, Deployed in multiple Verizon, AWINX etc.
- Part of Linux Foundation

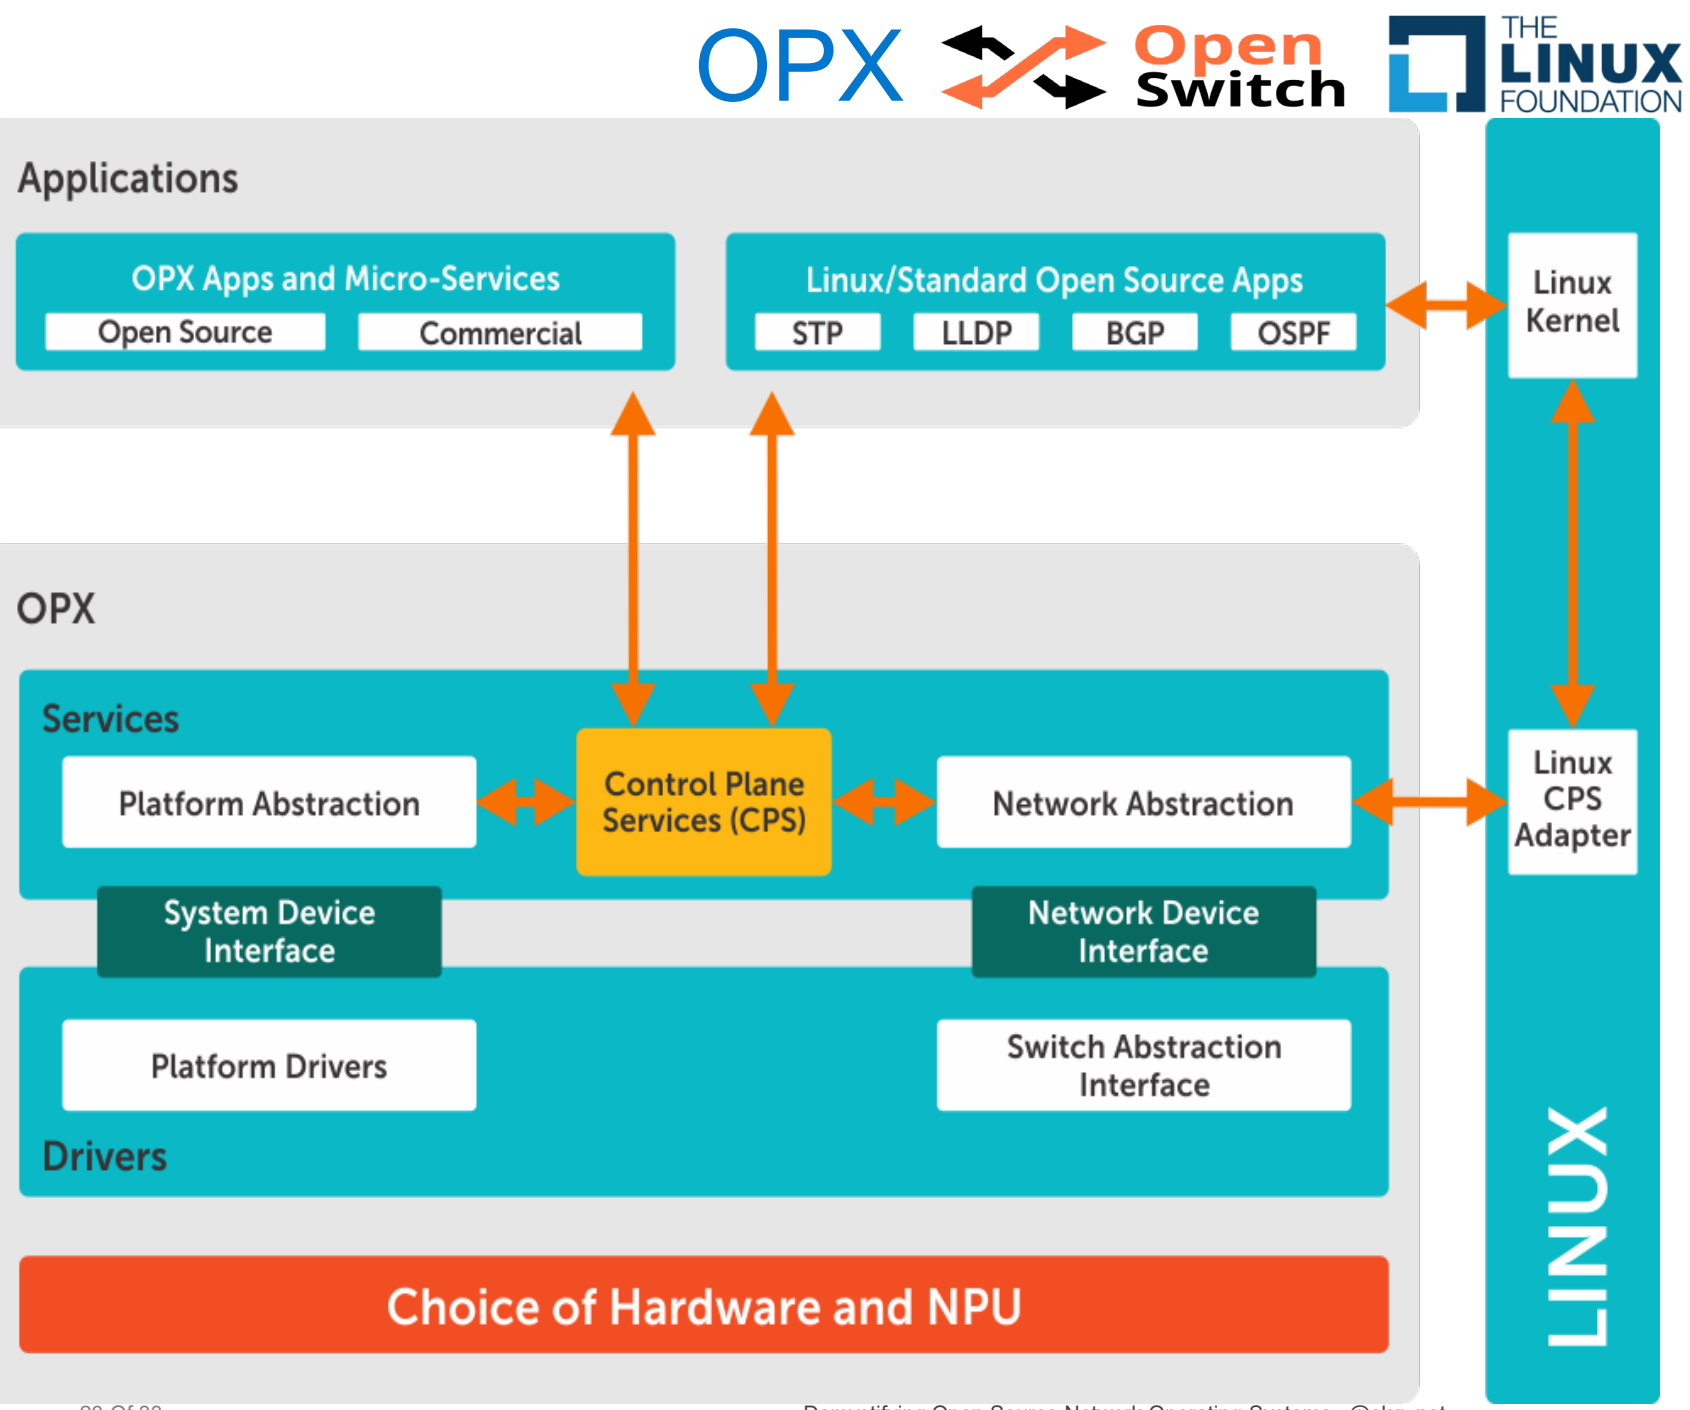

Open vSwitch

- Provides Platform Abstraction and Network Abstraction
- By default supports the Linux Protocol & Application Stacks
- Used as base operating system for Dell EMC Enterprise Operating OS10
- Can be leveraged to write custom application using Control Plane Services API.
- Production Ready, Deployed in multiple Verizon, AWINX etc.
- Part of Linux Foundation

OPX Features

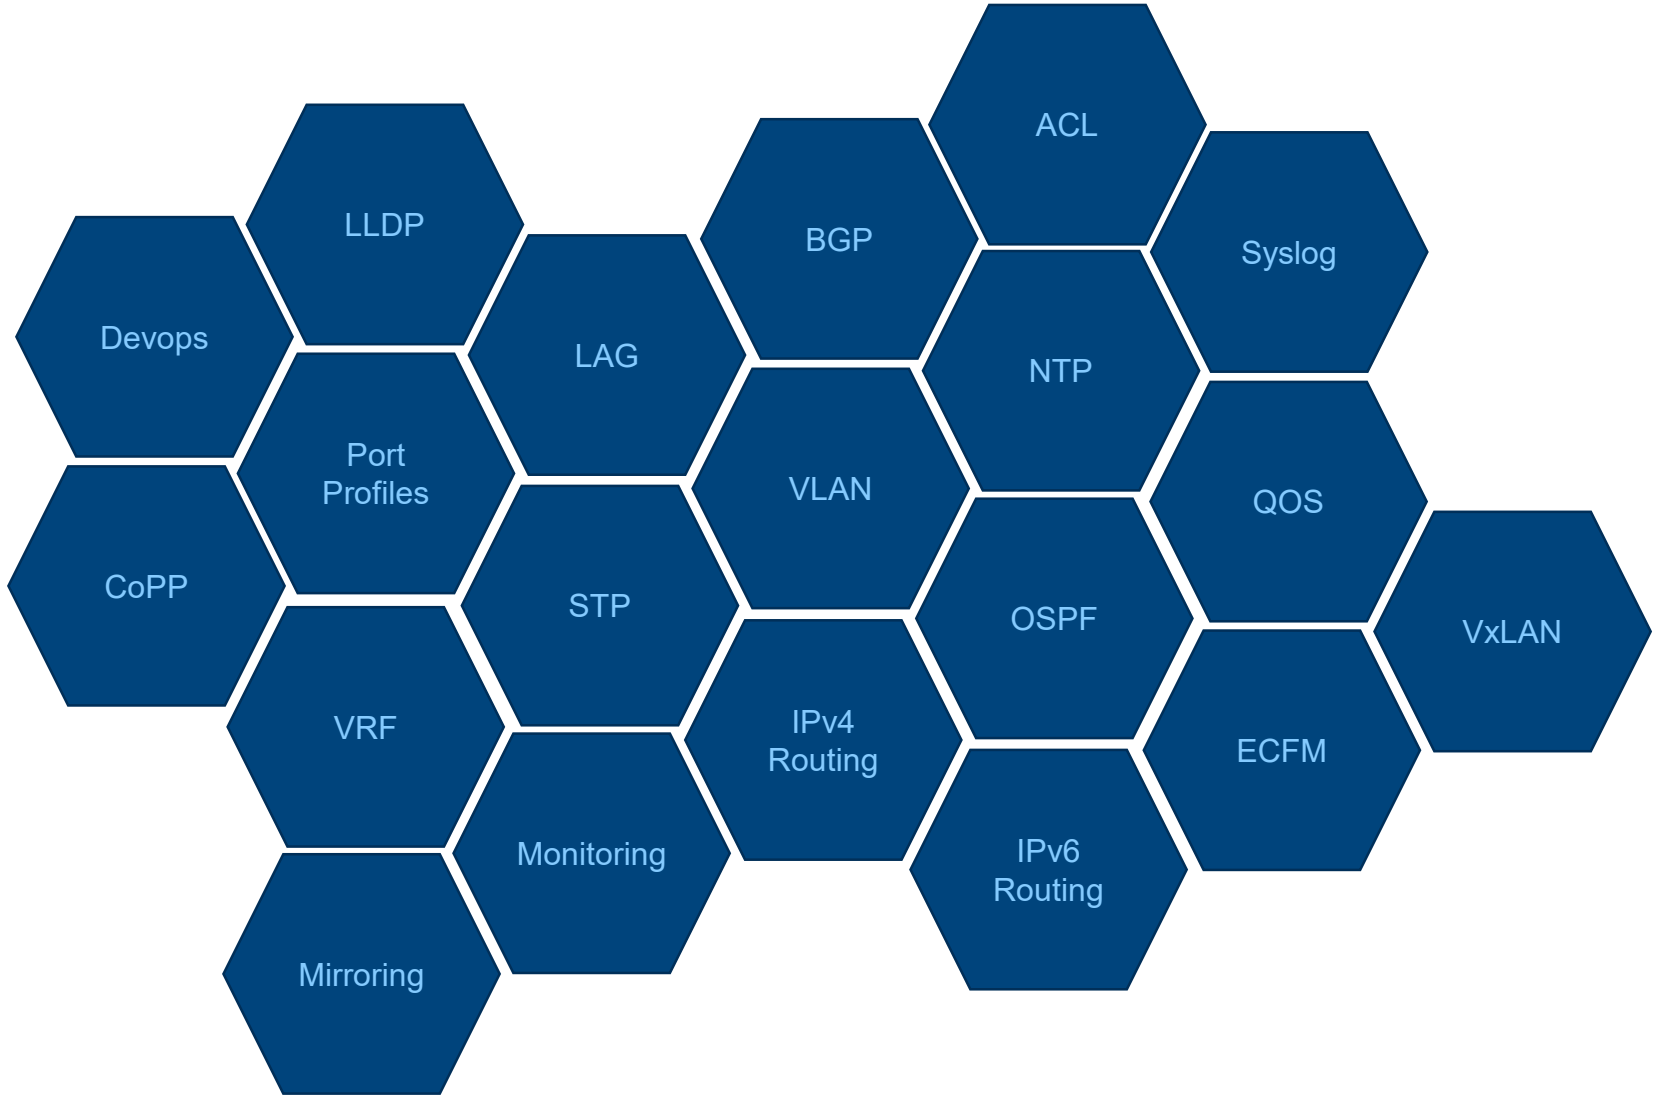

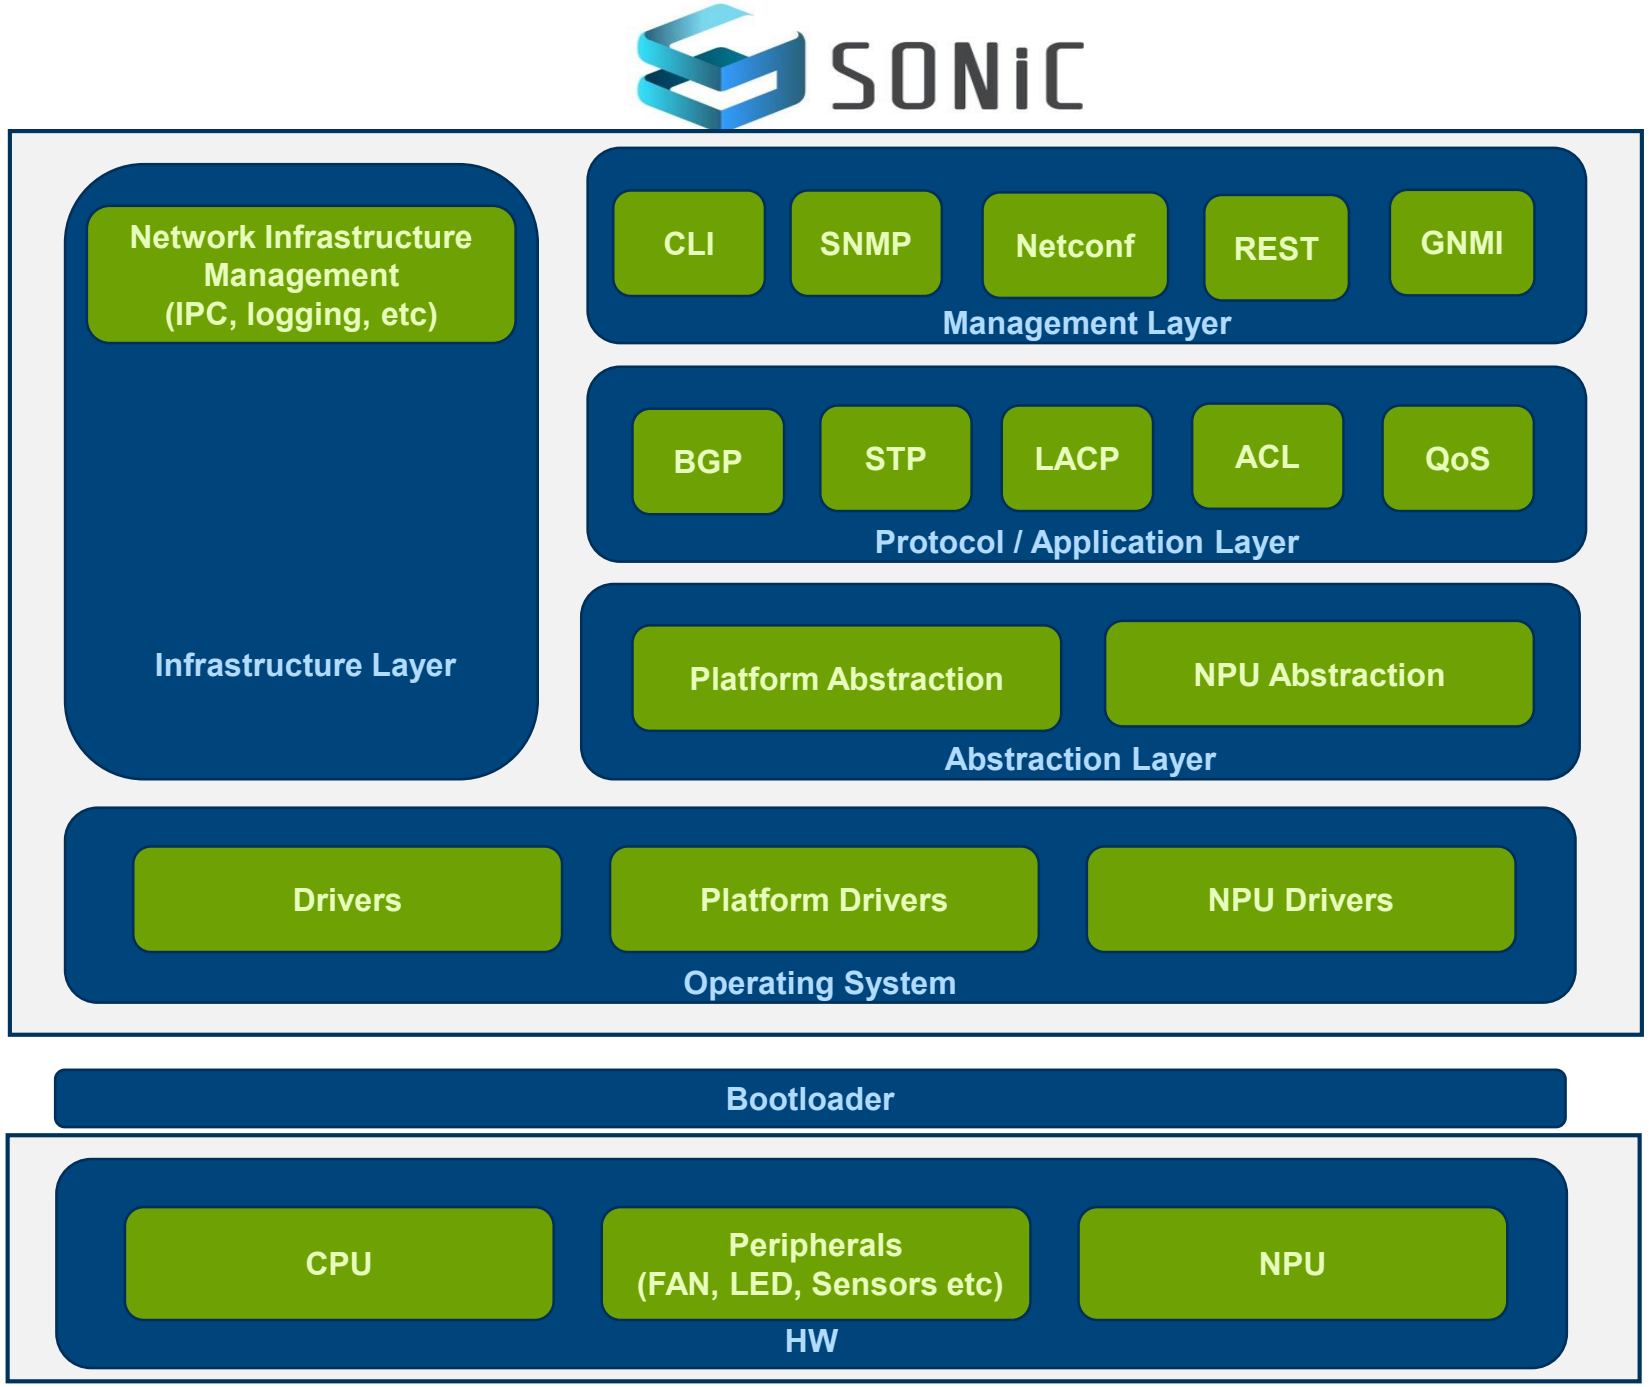

- Provides a complete NOS with the modern database and micro services based architecture.
- By default supports the Linux Protocol & Application Stacks
- Production Ready & Deployed by multiple Hyperscalers and Enterprise (Microsoft, Alibaba, Tencent, LinkedIn, Comcast etc)
- Supported by Dell Technologies and Mellanox
- Deployment velocity and scale
- Use Cases:
  - Data Center Leaf & Spine

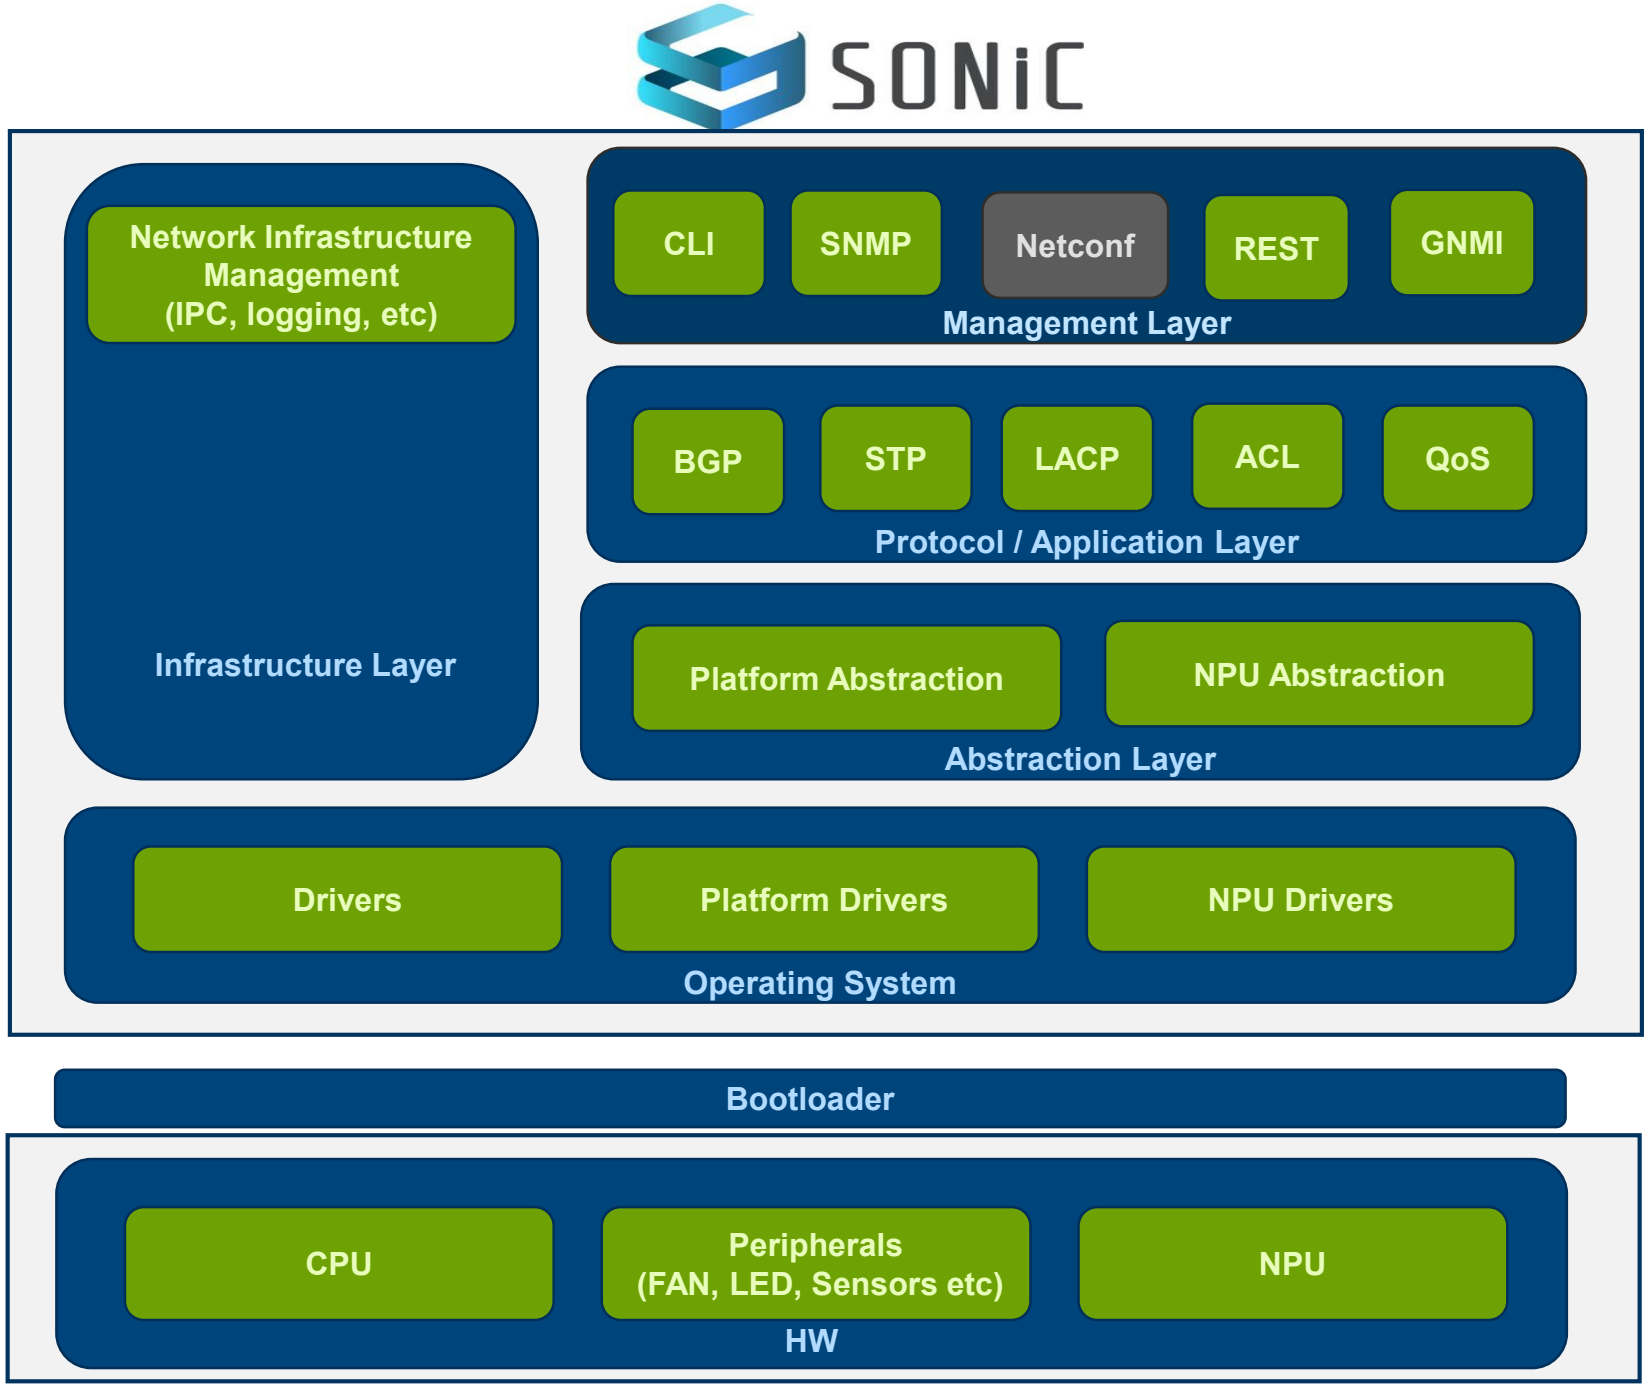

- Provides a complete NOS with the modern database and micro services based architecture.
- By default supports the Linux Protocol & Application Stacks
- Production Ready & Deployed by multiple Hyperscalers and Enterprise (Microsoft, Alibaba, Tencent, LinkedIn, Comcast etc)
- Supported by Dell Technologies and Mellanox
- Deployment velocity and scale
- Use Cases:
  - Data Center Leaf & Spine

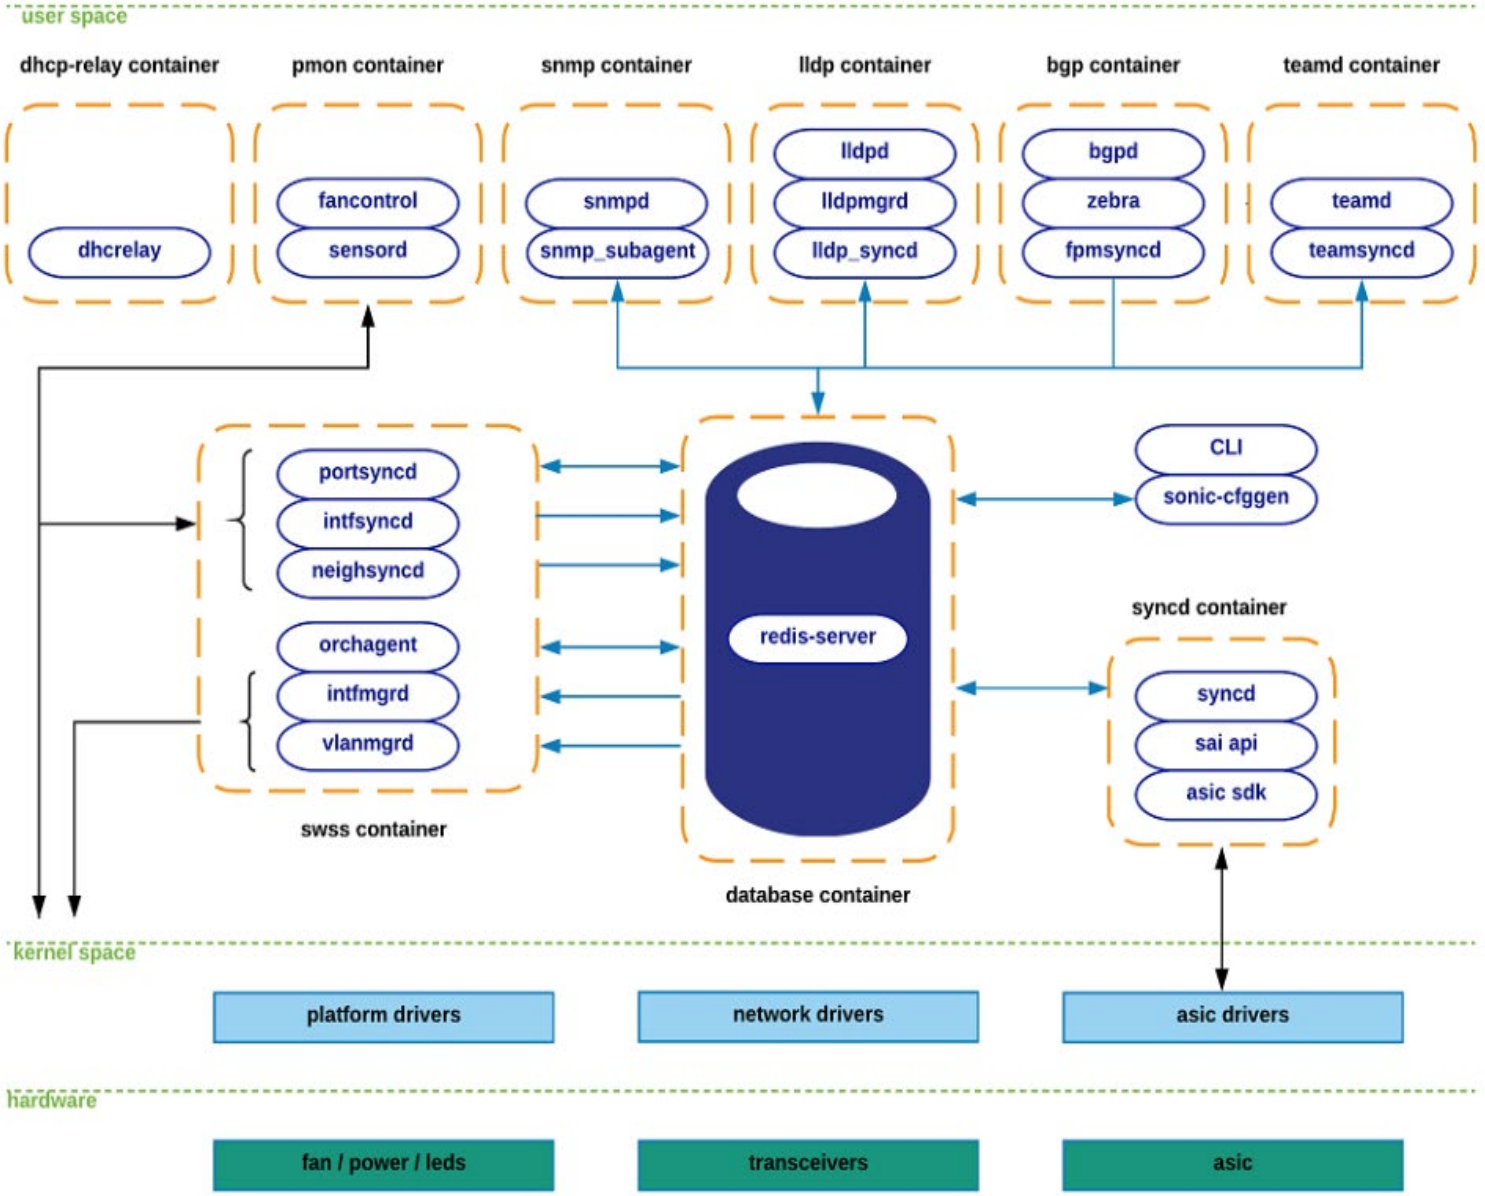

- Provides a complete NOS with the modern database and micro services based architecture.
- By default supports the Linux Protocol & Application Stacks
- Production Ready & Deployed by multiple Hyperscalers and Enterprise (Microsoft, Alibaba, Tencent, LinkedIn, Comcast etc)
- Supported by Dell Technologies and Mellanox
- Deployment velocity and scale
- Use Cases:
  - Data Center Leaf & Spine

SONiC Features

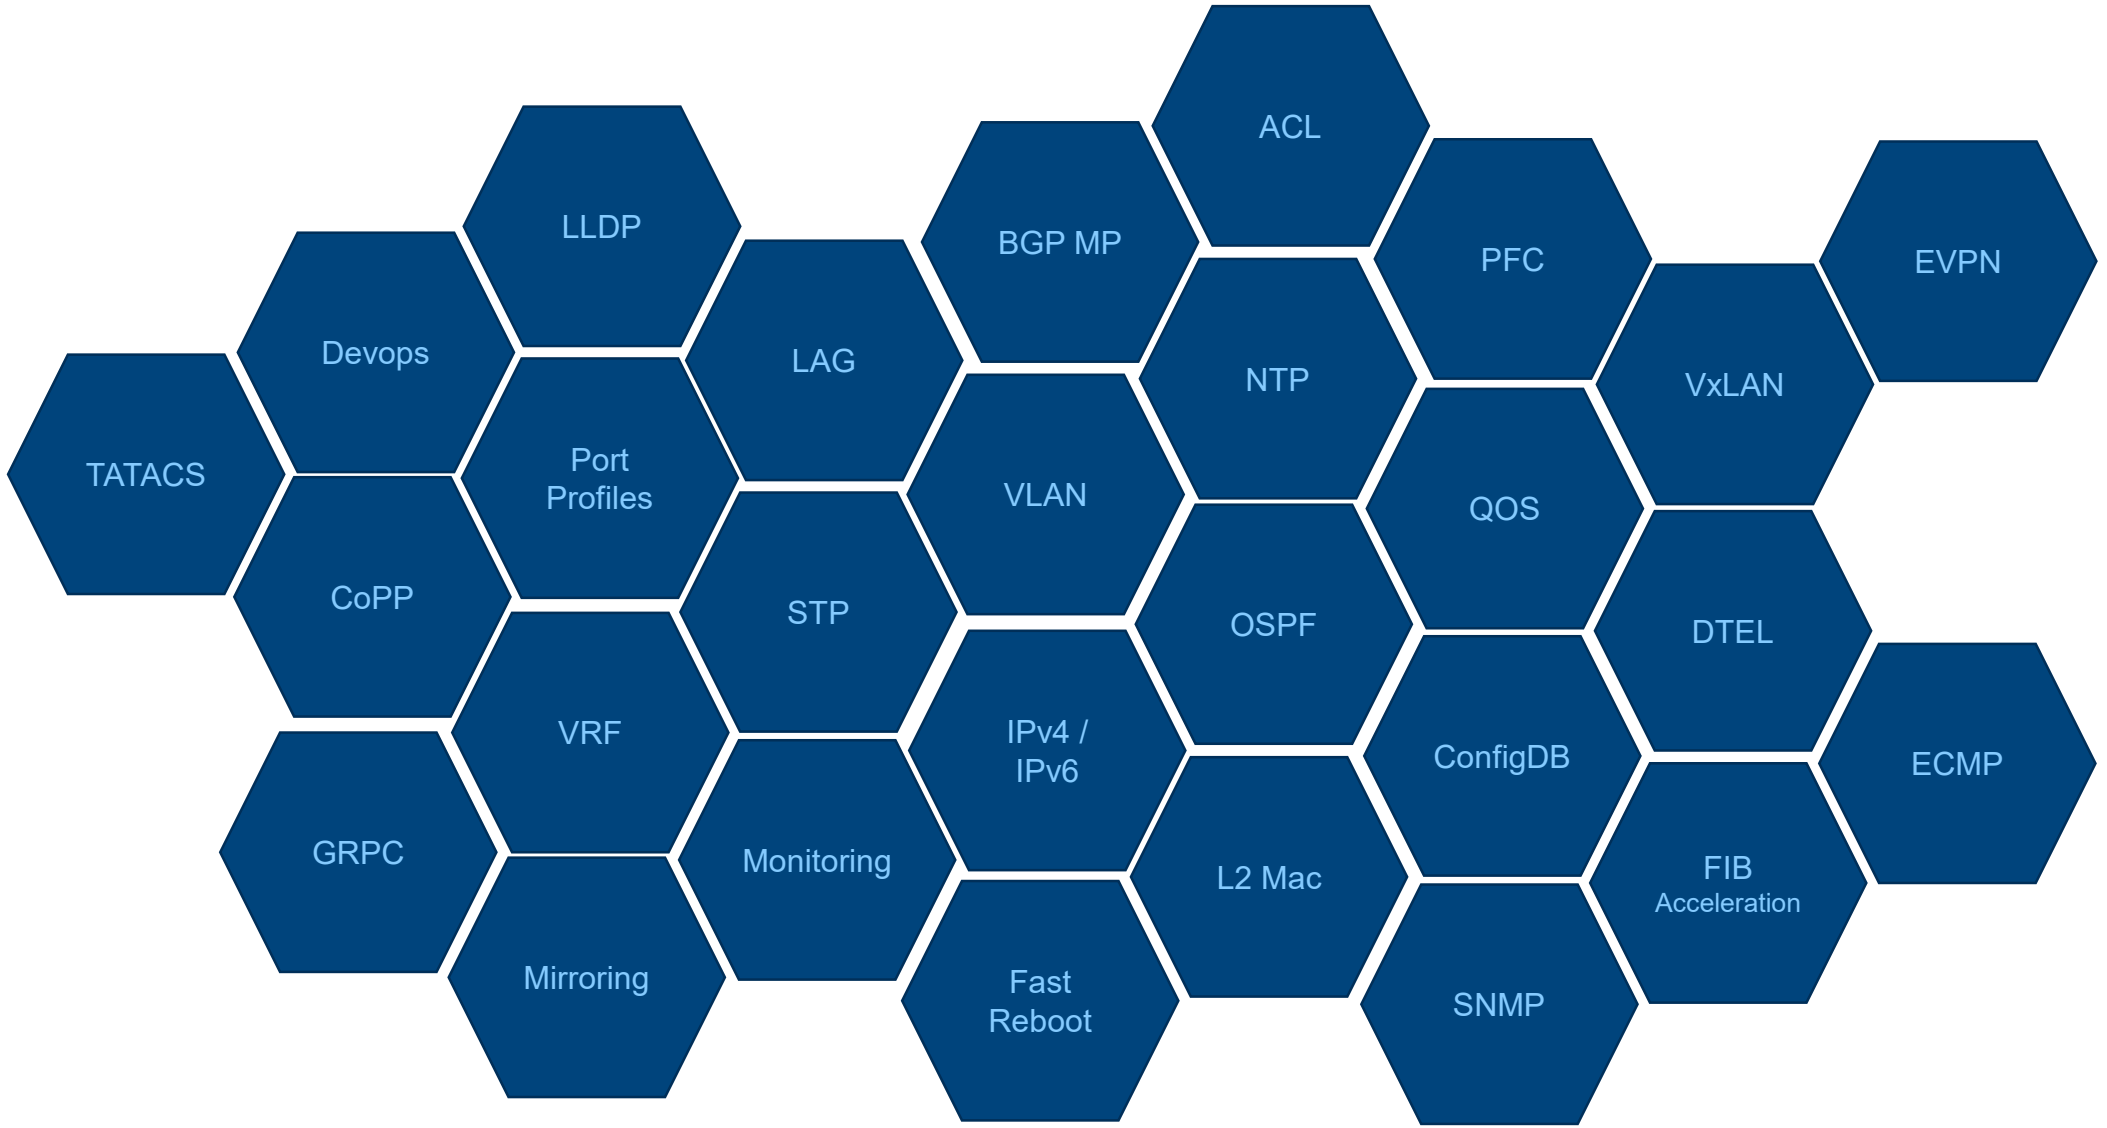

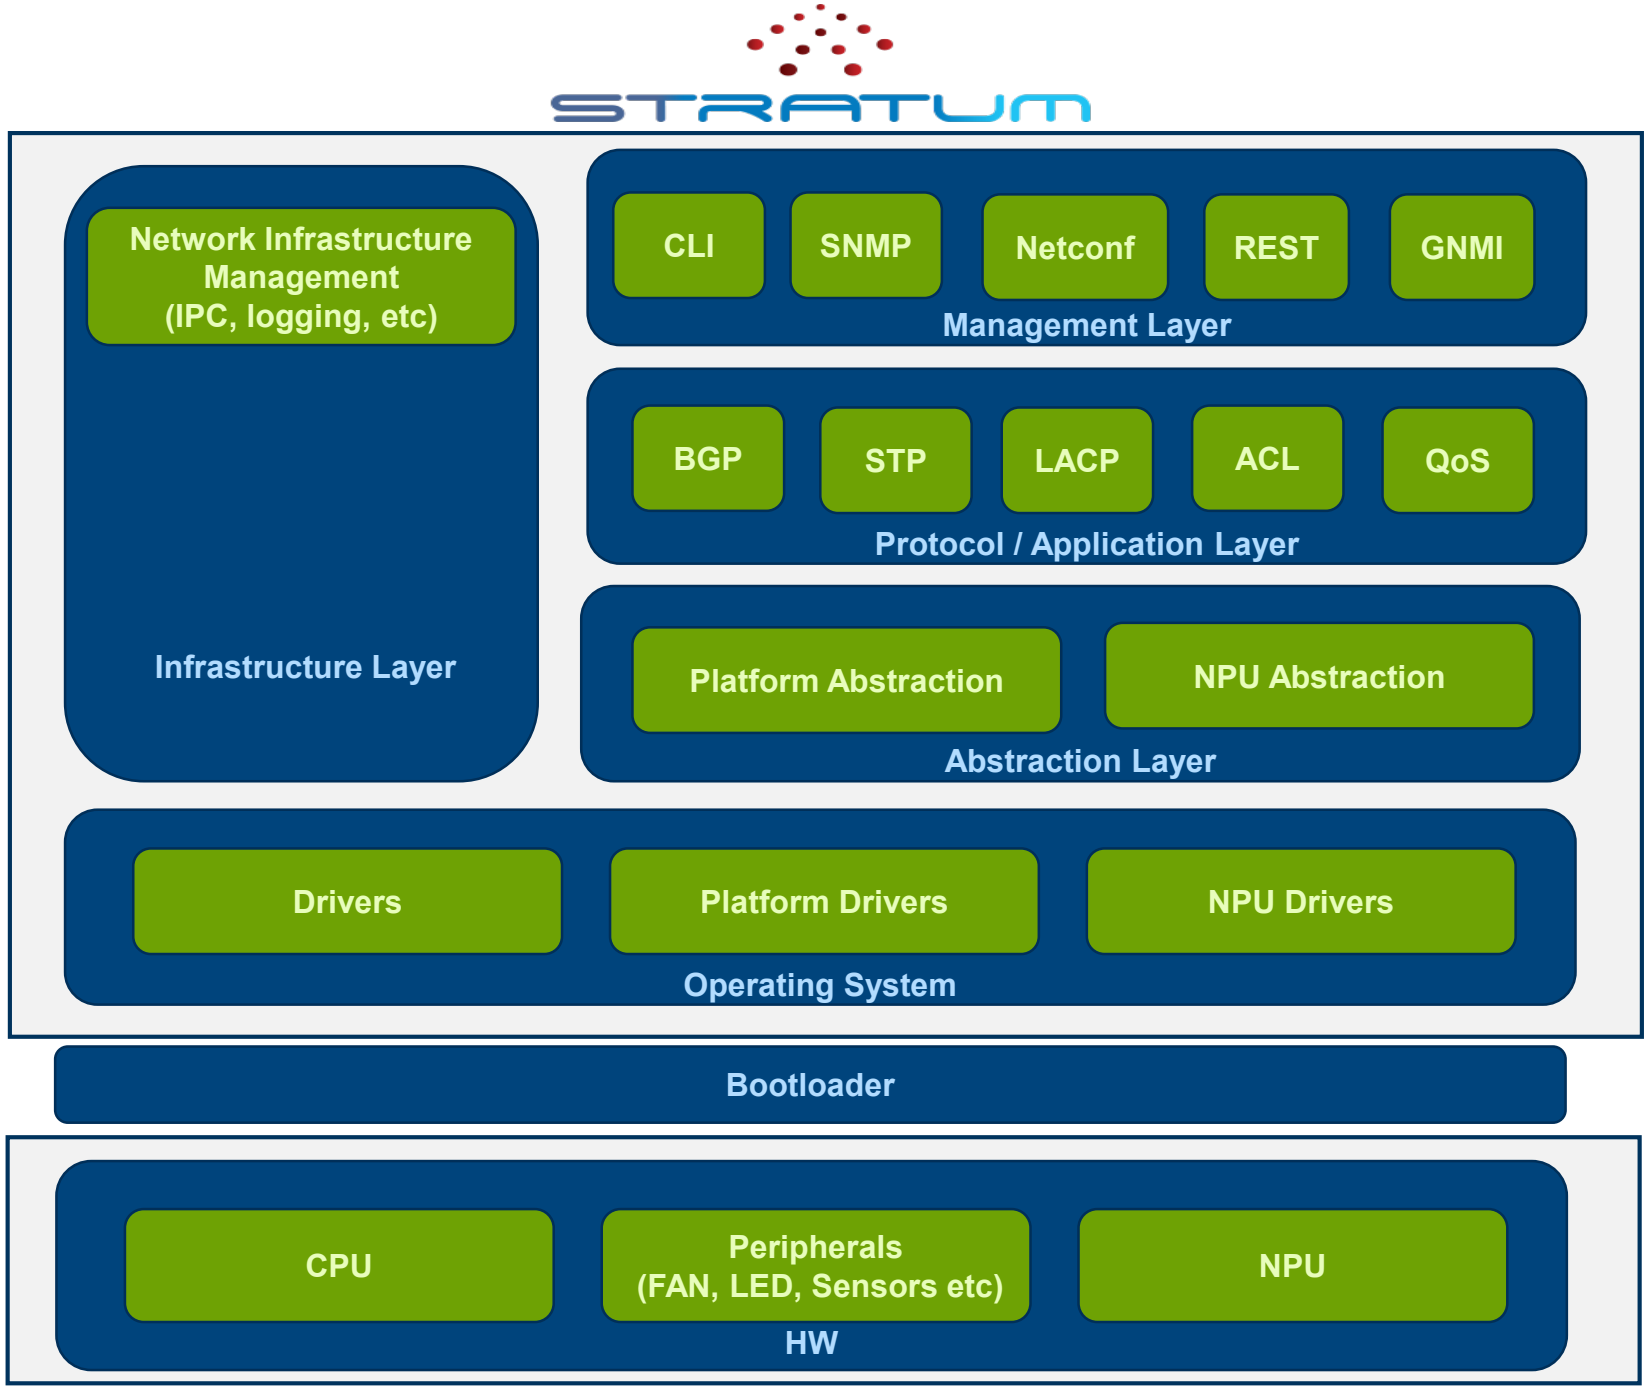

ONF

- New, extensible control interface GNMI, GNOI, P4 Runtime
- Common models and interface for configuration, management & operations
- Common Platform Abstractions
- Vendor-neutral control planes
- Unified device management
- Simplified migration
- Deployment velocity and scale
- Open Source Sep 2019

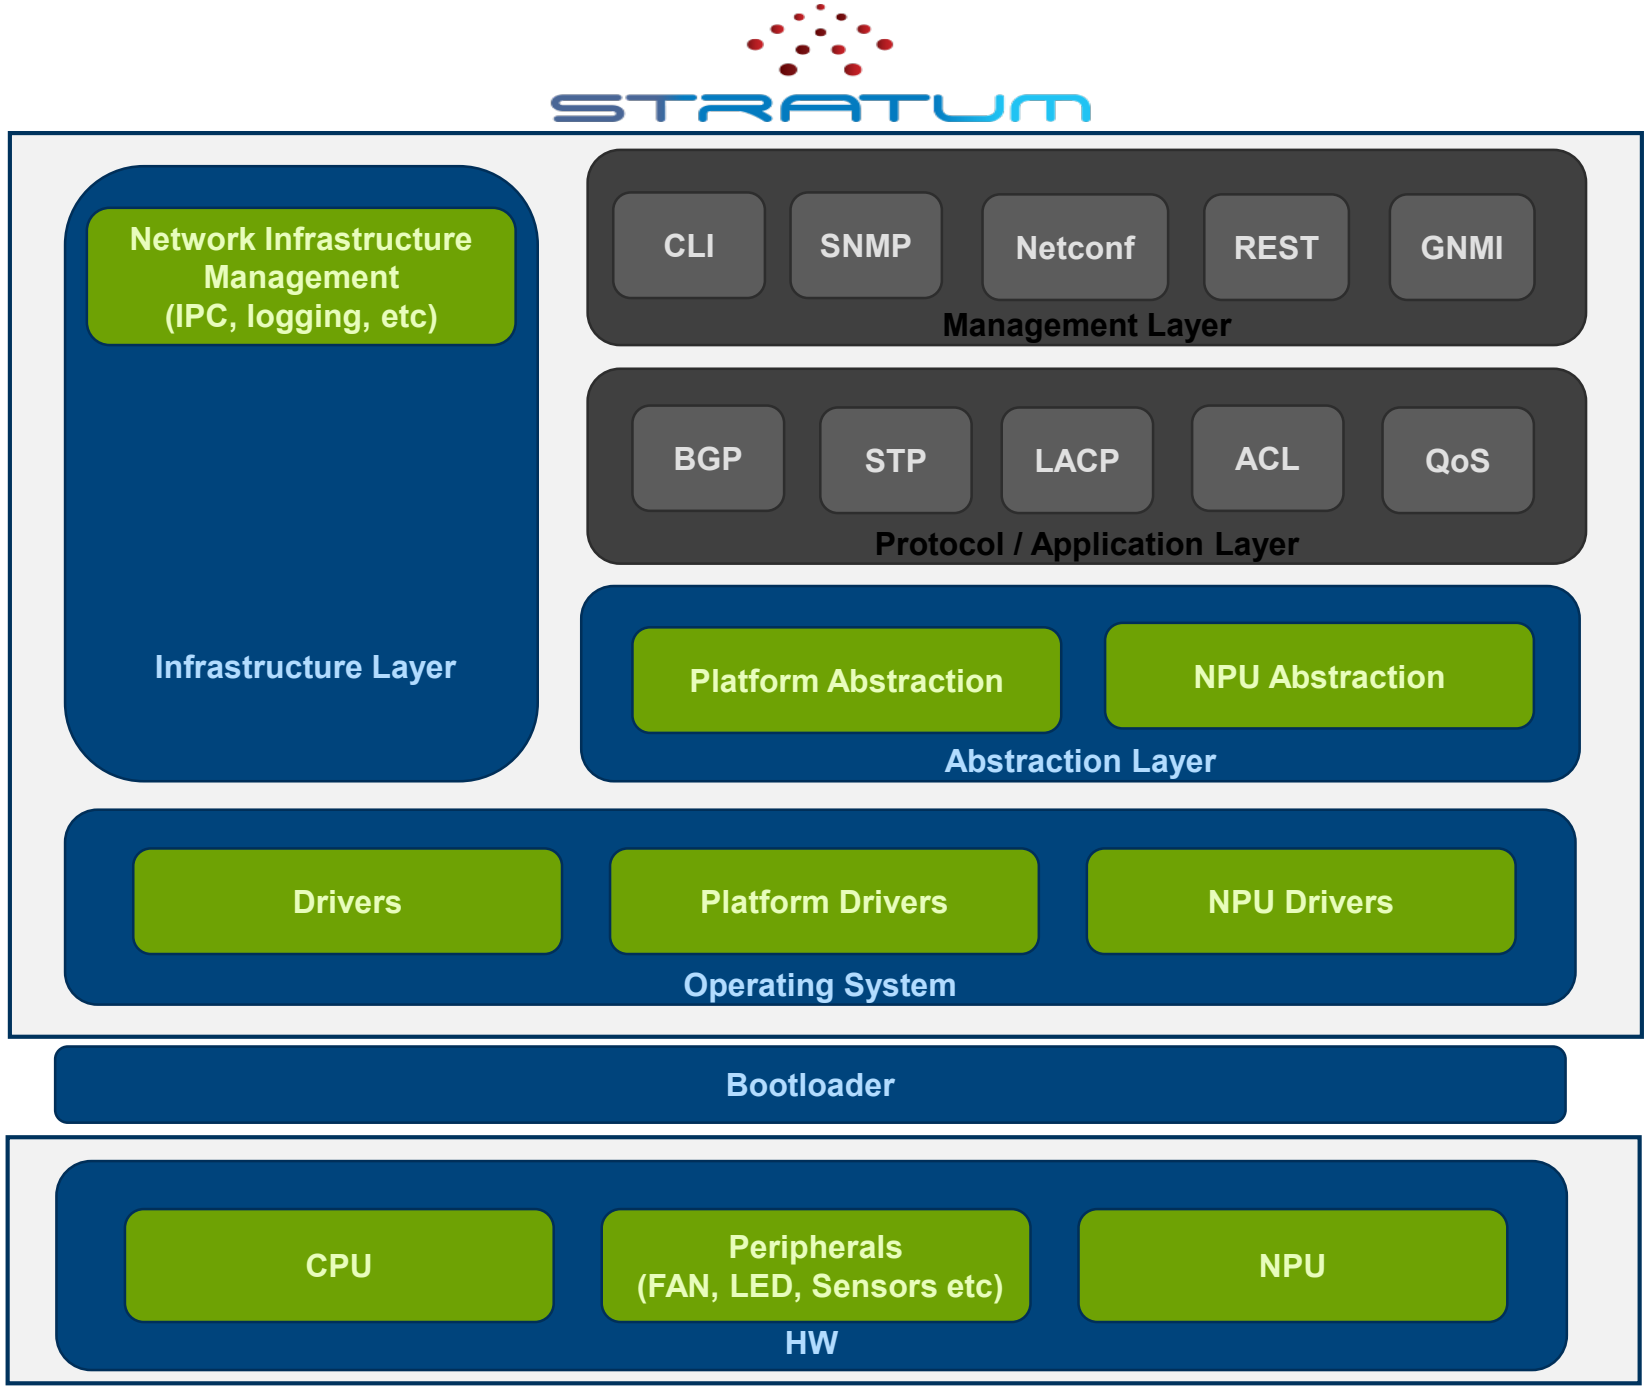

ONF

- New, extensible control interface GNMI, GNOI, P4 Runtime
- Common models and interface for configuration, management & operations
- Common Platform Abstractions
- Vendor-neutral control planes
- Unified device management
- Simplified migration
- Deployment velocity and scale
- Open Source Sep 2019

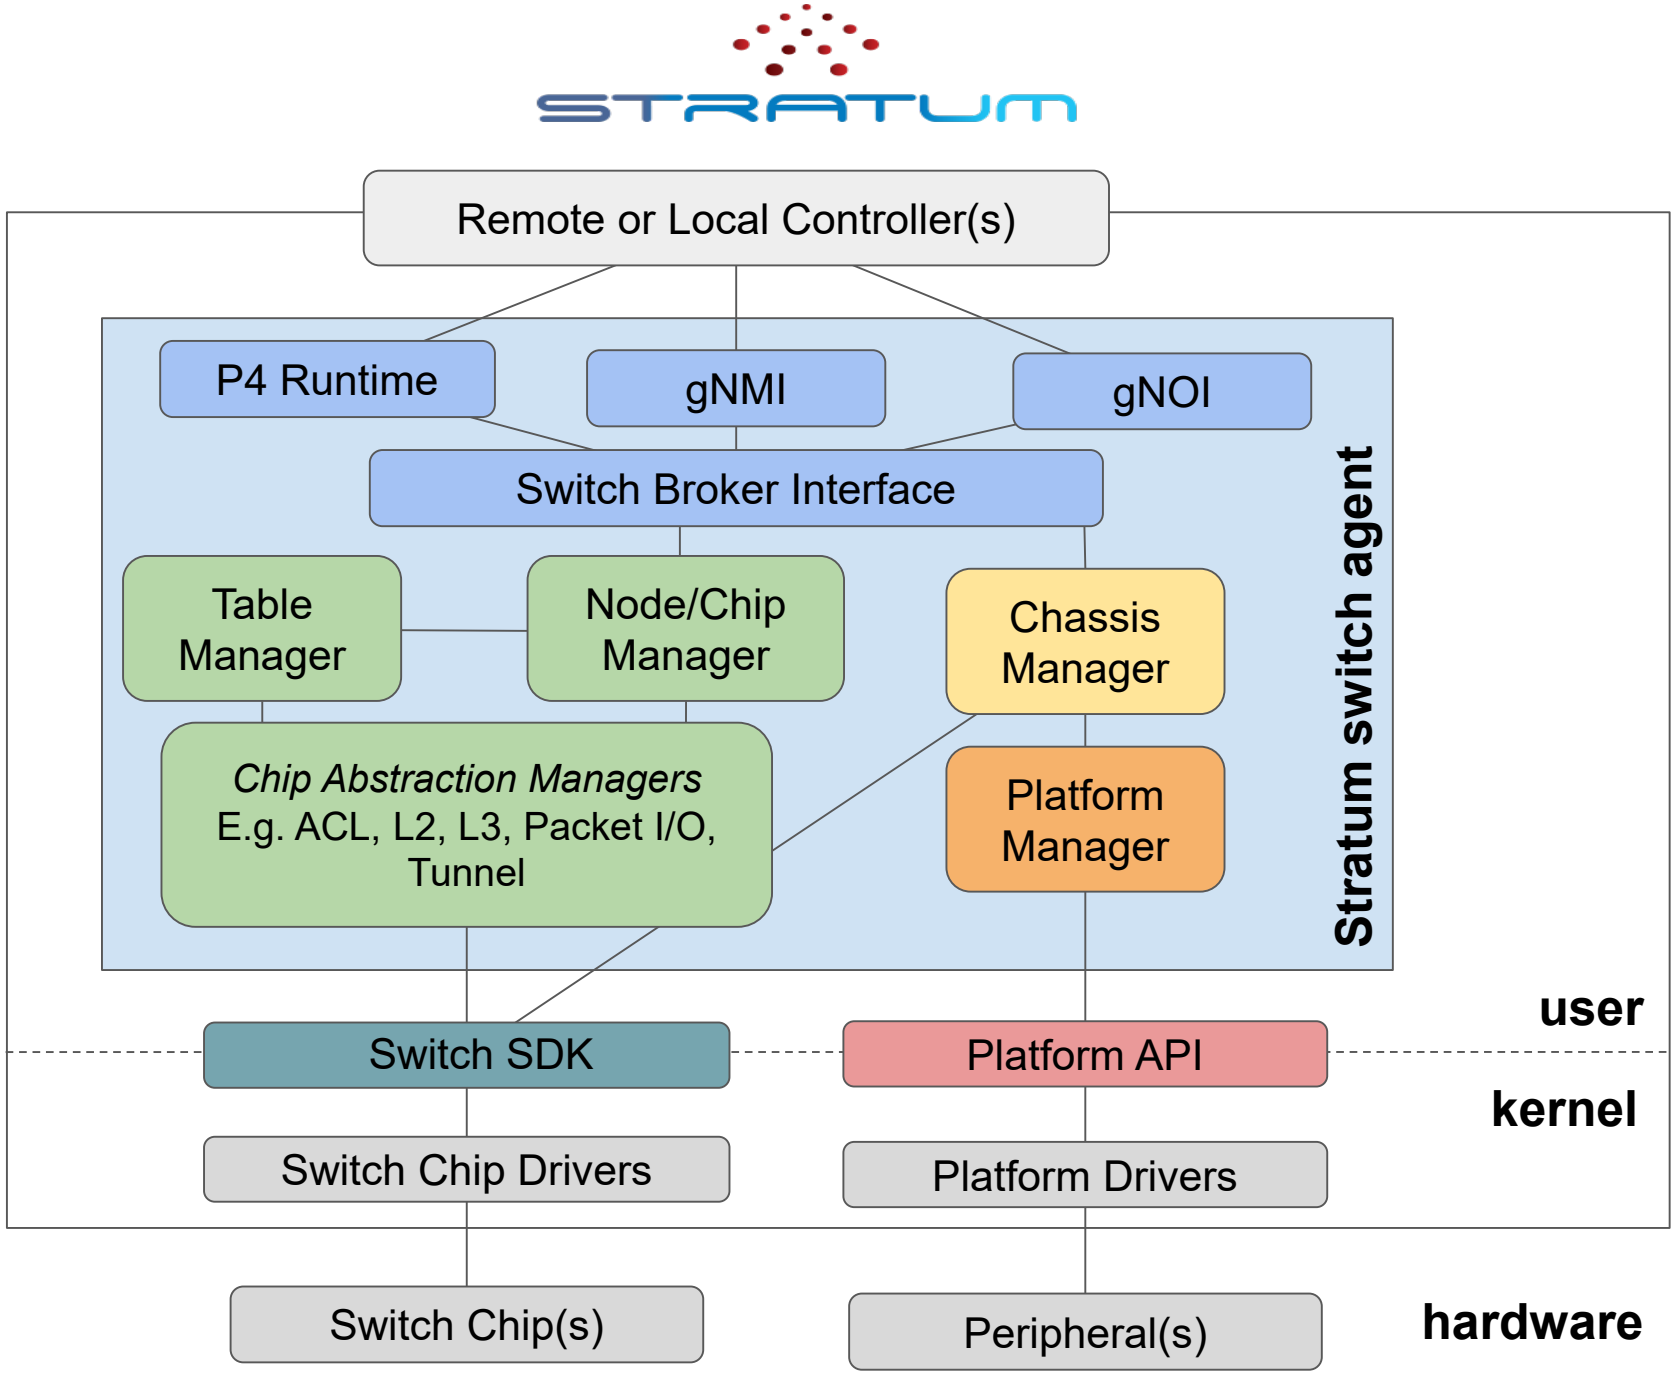

ONF

- New, extensible control interface GNMI, GNOI, P4 Runtime
- Common models and interface for configuration, management & operations
- Common Platform Abstractions
- Vendor-neutral control planes
- Unified device management
- Simplified migration
- Deployment velocity and scale
- Open Source Sep 2019

Shared (HW agnostic)
Chip specific
Platform specific
Chip and Platform specific

Stratum Use Cases

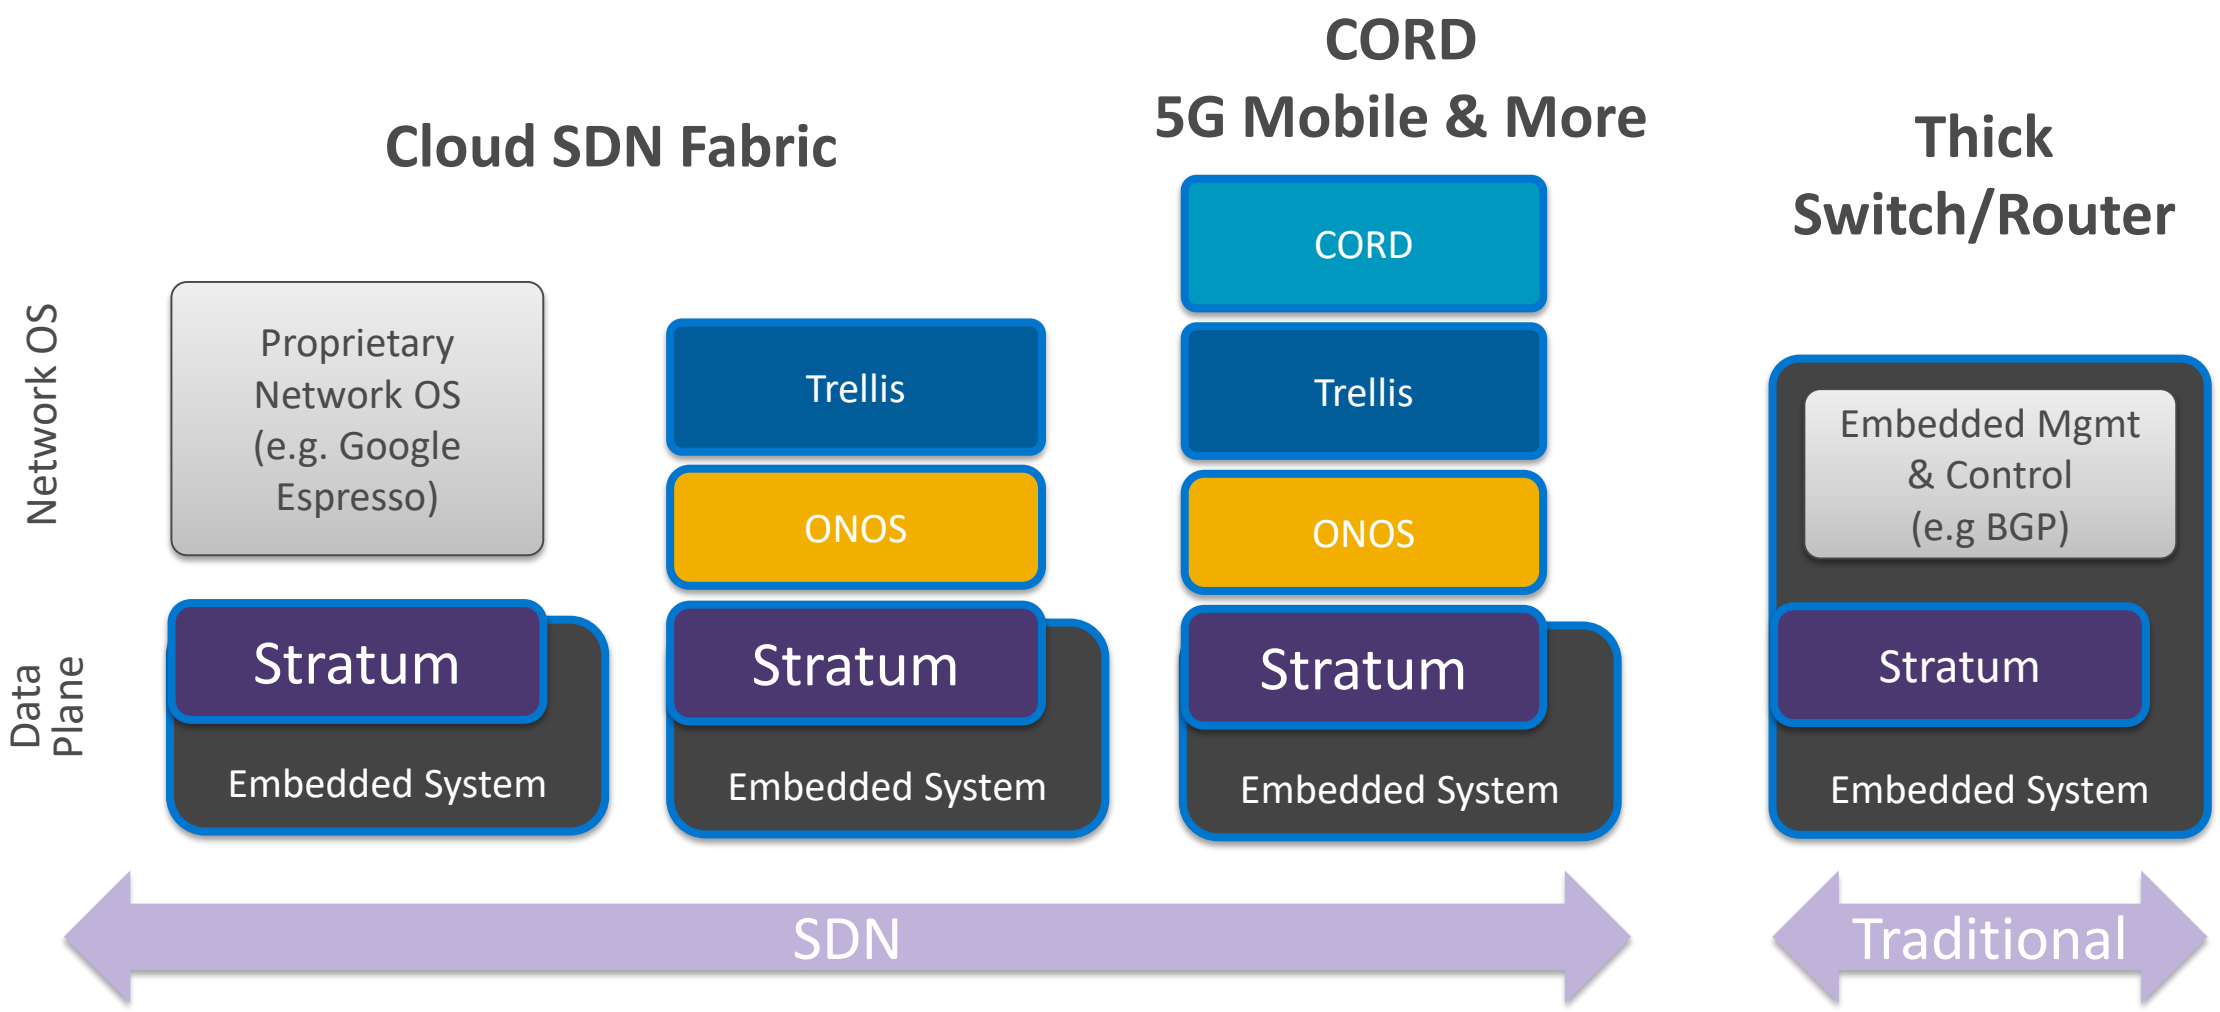

Open Networking & Open Source NOS Use cases

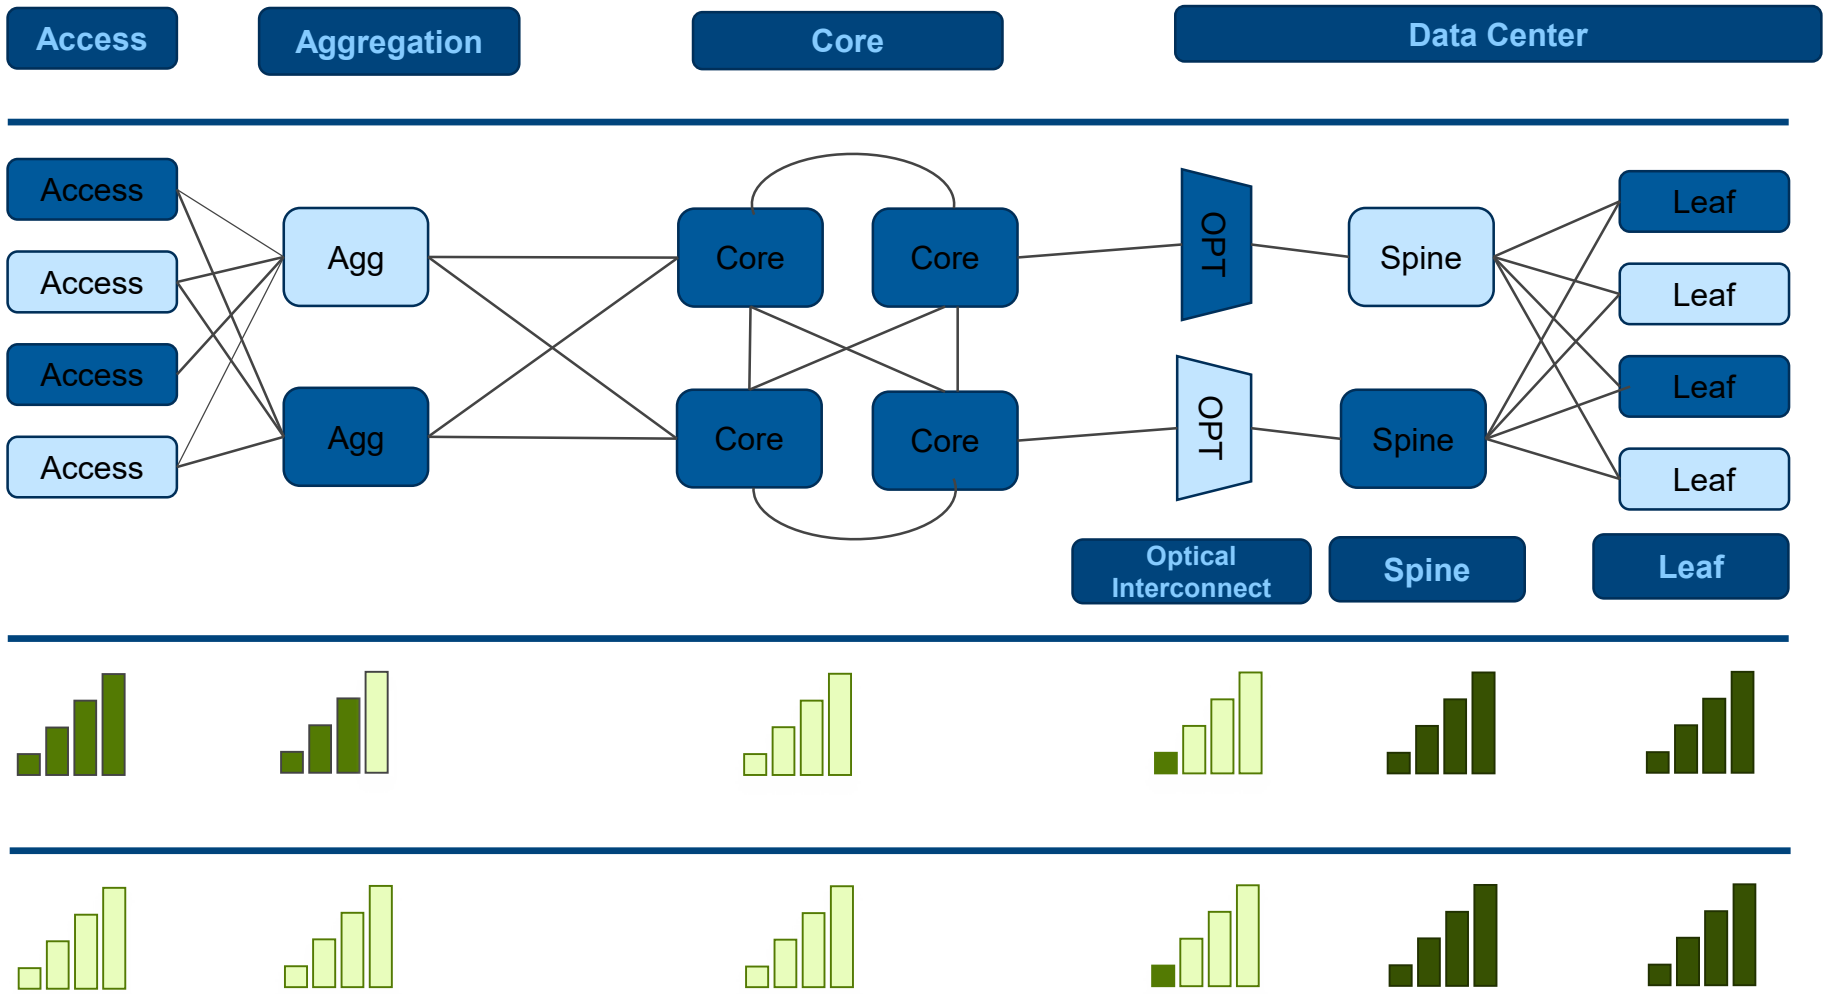

Journey so far…

Hyperscalers → Service Providers → Enterprise

Challenges & Call for Action?

Image by [Peggy und Marco Lachmann-Anke]

Support

Quality

Challenges

Call for Action

Explore & Understand

Test Drive

Deploy

@skg_net

https://www.linkedin.com/in/skgnet