

OPEN POSSIBILITIES.

FBOSS Experience in Onboarding a Second Silicon Vendor



NOVEMBER 9-10, 2021

FBOSS Experience in Onboarding a Second Silicon Vendor

George Kurio , Hardware Engineer, Meta
Srikrishna Gopu, Software Engineer, Meta
Wei-Jen Huang, Distinguished Engineer, Cisco

OPEN POSSIBILITIES.



FBOSS

- **F**acebook **O**pen **S**witching **S**ystem (FBOSS)
- Meta's software stack for controlling/managing network switches deployed in Meta's Datacenters
- One of the largest services in Meta in terms of the number of instances deployed



NETWORKING

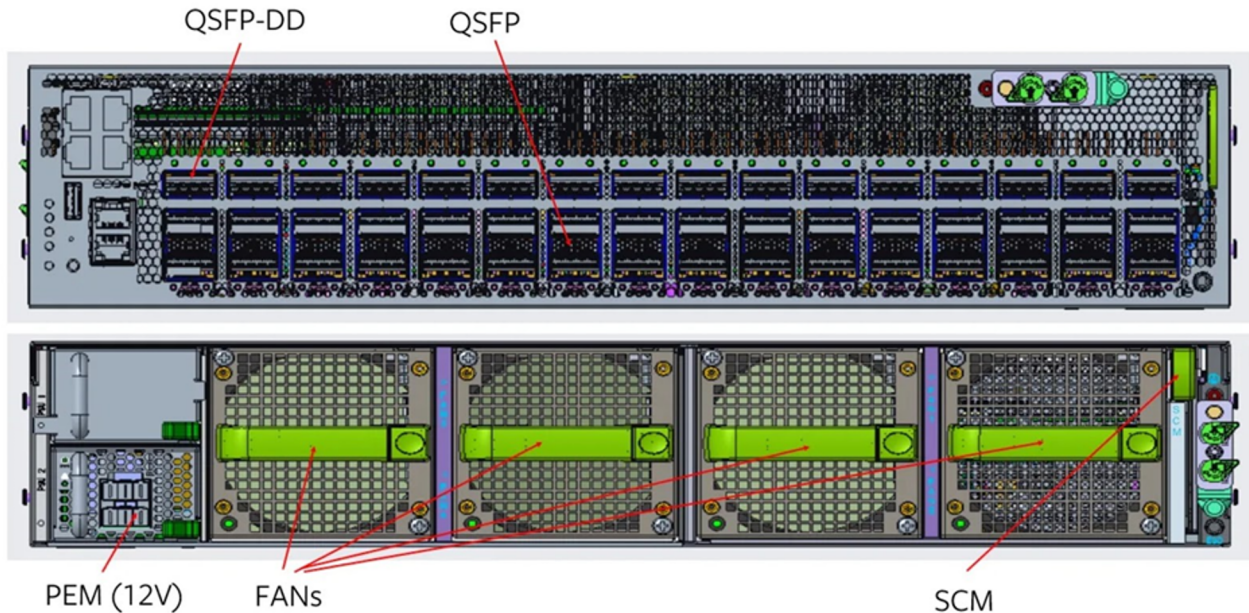
OPEN POSSIBILITIES.



Hardware





NETWORKING



OPEN POSSIBILITIES.



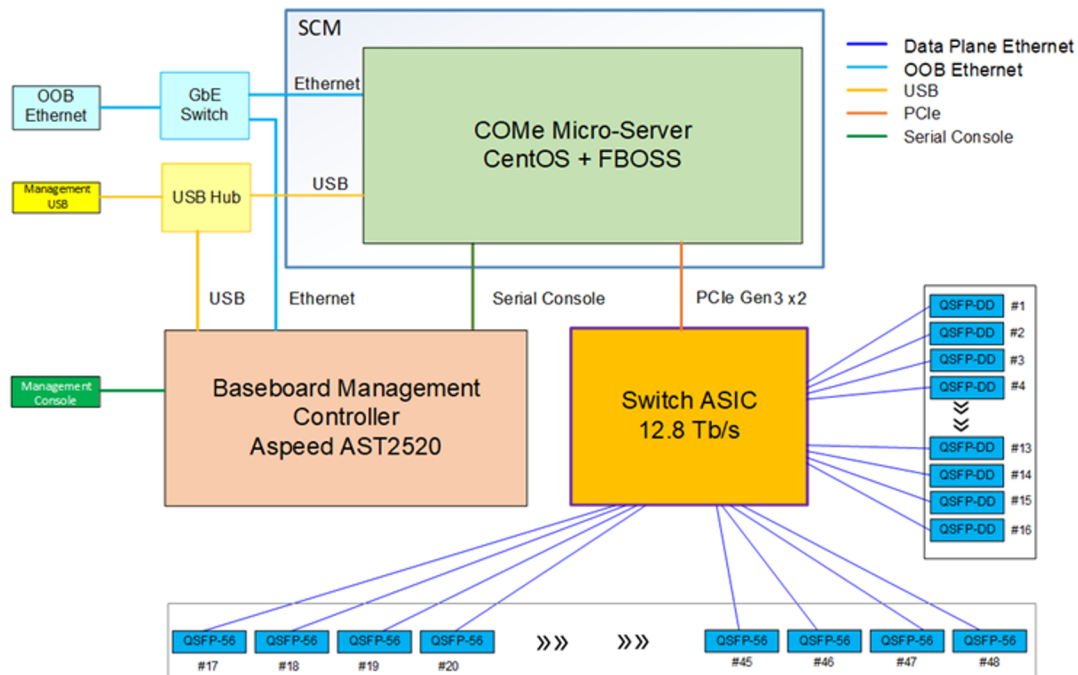
Hardware



NETWORKING

OPEN POSSIBILITIES.

Hardware



NETWORKING

OPEN POSSIBILITIES.

Hardware

- 16 x QSFP-DD: 400G/200G/100G ports
- 32 x QSFP56: 200GE/100GE ports or 2*50GE/4*25GE/4*10GE breakout ports
- RS232 console port to BMC
- OOB GE management port
- USB 2.0 compatible: Support for OCP debug card
- System Controller Module (SCM)
- Power Plane: AC/DC or DC/DC PEM
- Fans: Four 80mm x 80mm x 80mm counter-rotating fan-trays



NETWORKING

OPEN POSSIBILITIES.



Hardware: The Journey



NETWORKING

- Goal: Design and build a switch that is 100% interchangeable with Wedge400
- Cisco gave a series of instructional classes
- ODM engagement: Celestica has been an outstanding partner
- Hardware Development:
 - EVT phase proceeded rapidly and then Covid lockdowns occurred mid-DVT
 - Critical for onboarding a second vendor, Cisco was (and is) a true partner and provided excellent training and support throughout the life cycle of the program development
 - Cisco participated in initial on site hardware bringup and debug
 - There were relatively few hardware related issues. No post DVT HW spin
 - The project transitioned from PVT to MP in June, 2021

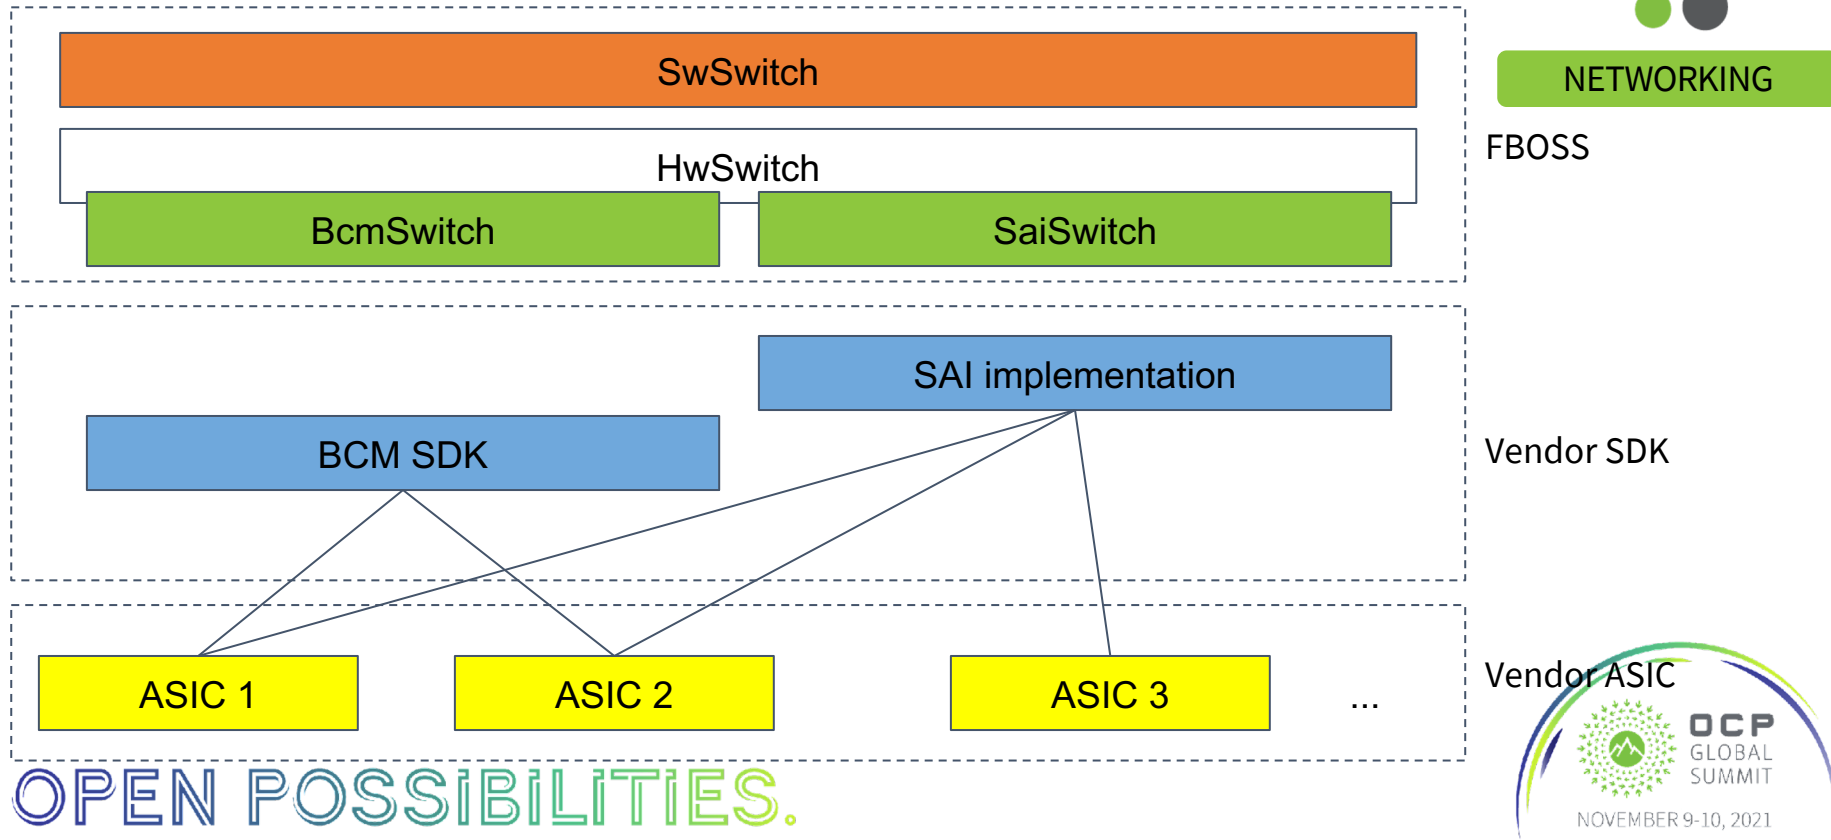
OPEN POSSIBILITIES.



FBOSS Software Architecture



NETWORKING



FBOSS

Vendor SDK

Vendor ASIC



FBOSS + SAI



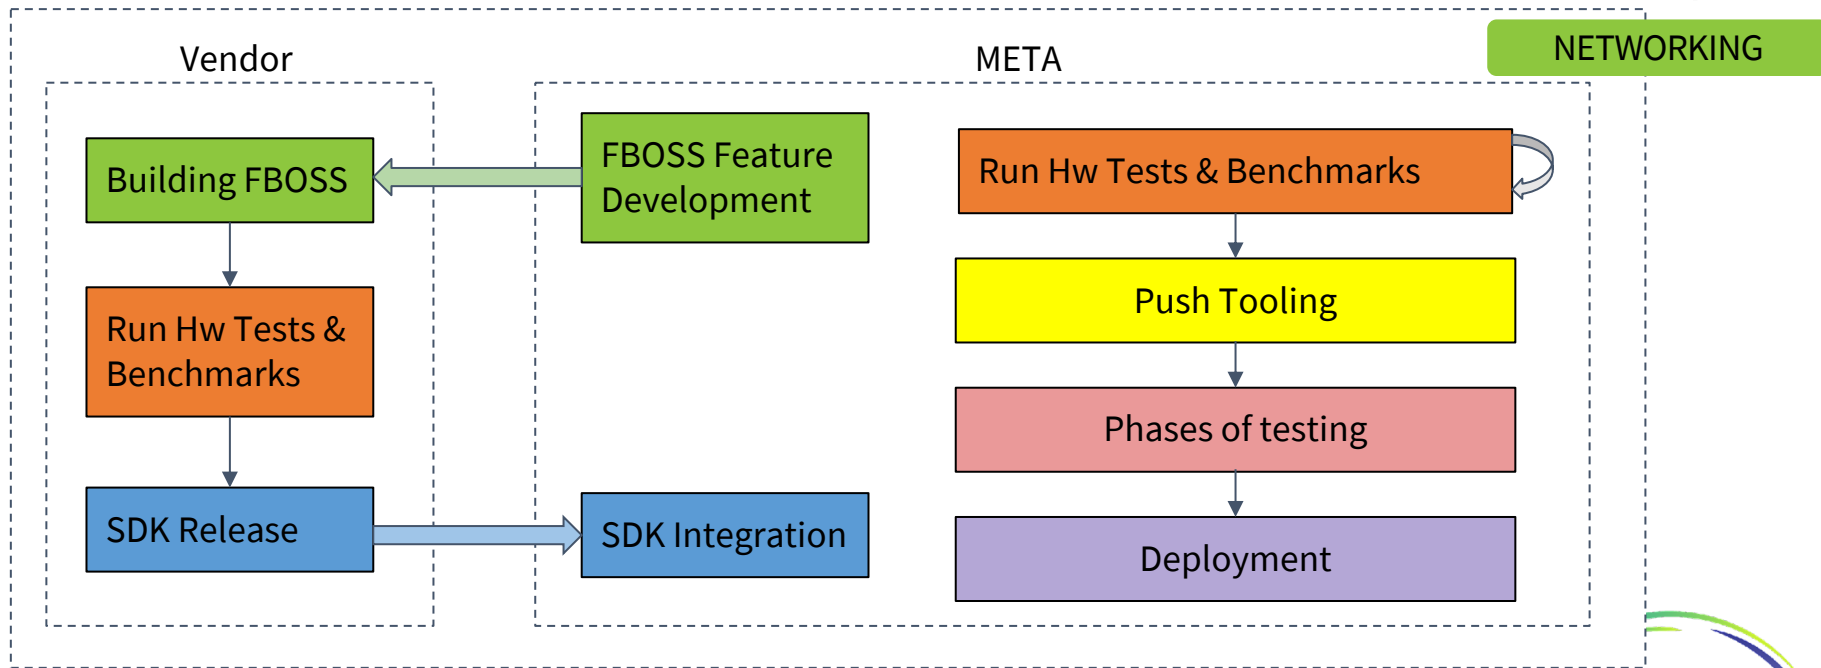
NETWORKING

- SAI
 - **S**witch **A**bstraction **I**nterface
 - Project under **O**pen **C**ompute **P**roject (OCP)
 - Open source API to control forwarding elements
 - Vendor independent
- FBOSS SAI based implementation:
 - HwSwitch: multiple ASICs, ASIC vendors
 - Easy to onboard newer ASICs
 - Open source contributions
 - FBOSS is open source
 - Meta contributes to SAI spec

OPEN POSSIBILITIES.



Software



OPEN POSSIBILITIES.

Software - SDK Integration



NETWORKING

- Difference in tool chains and other external dependencies
- Integrating the SDK into Meta infrastructure is a challenge
 - ~ 3 - 4 weeks during initial phase
 - ~ 2 days to integrate now.
- Several SDK drops to attain feature parity
- SDK release cadence of 2 weeks

OPEN POSSIBILITIES.




Software - Validating functionality



NETWORKING

- Existing tests were vendor dependent
- Ported ~1000 tests to a vendor agnostic framework
 - The same tests runs on sai platforms and non-sai platforms.
 - Automated to run every day on all platforms
 - Performs warmboot
- Validated asic functionalities like ECMP, QOS, Hashing, Port Up/Down, ACL, Buffers, Schedulers
- Worked with vendor to ensure every release enables more tests

Single Box Dataplane testing



OPEN POSSIBILITIES.



Software - FBOSS Open Source



NETWORKING

- ~90% of FBOSS code is open sourced
- Feature Development was done in parallel for FBOSS and Vendor
- Vendor can build and run the hw tests to validate the asic functionality
- Gating factor before making a new SDK release

OPEN POSSIBILITIES.



Software - Tooling

Several tools have to be developed and updated in order to onboard a new asic.



NETWORKING

- *Push tooling*: Tools that are used to deploy the binary in the fleet should be aware of the 2nd asic.
- *Kernel module installer*: Kernel module has to be installed and deployed across the fleet
- *Chip specific Configuration Generation*: Vendor specific configuration has to be generated
- *Test infrastructure*: In order to run hardware tests on sai platforms, test infrastructure needs to be updated.

OPEN POSSIBILITIES.



Software - Challenges



NETWORKING

- **Different vendors using different SAI version**
 - Hard to force a vendor to be on a particular SAI version
 - Currently, FBOSS uses the header files provided by the vendors
- **SAI Extensions:**
 - Functionality diverge between vendors. For eg, LED
 - Generic attributes , propose it to the SAI community.
 - Specific attributes, use SAI extension.

OPEN POSSIBILITIES.



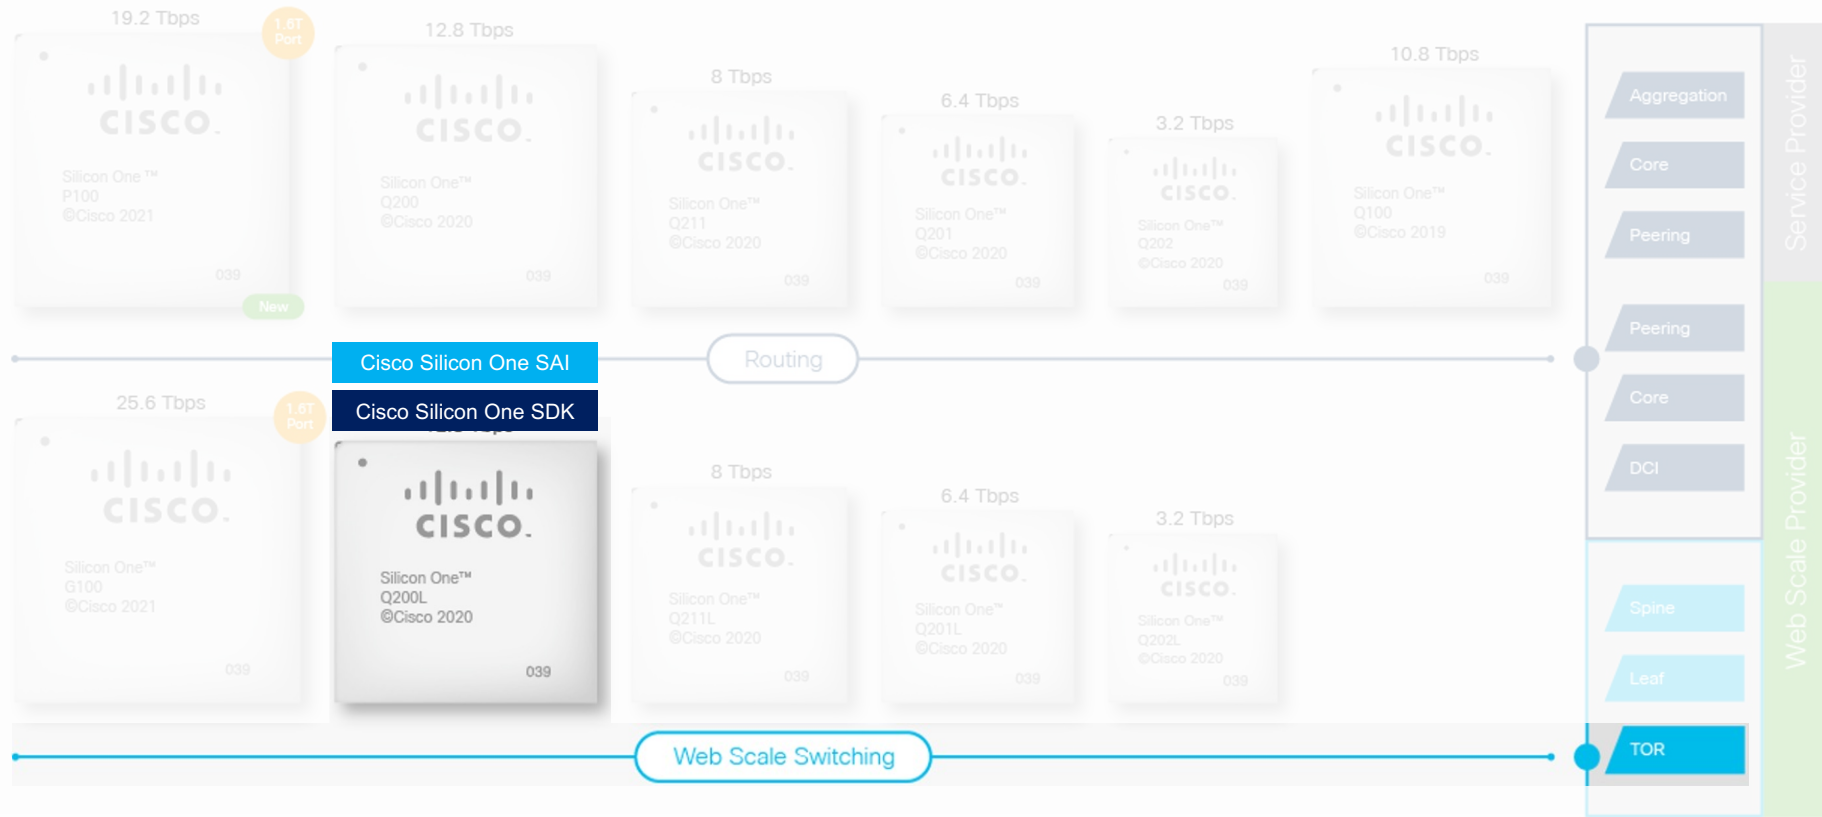
Cisco Silicon One Portfolio

One Architecture. Multiple Devices. No Compromise.



Cisco Silicon One Q200L Powering Wedge 400C

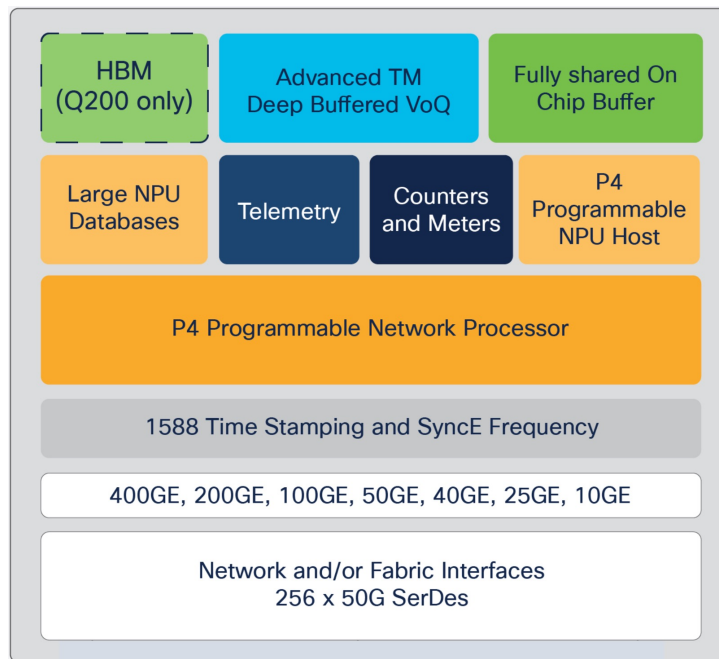
One Architecture. Multiple Devices. No Compromise.



Cisco Silicon One Q200 Family



NETWORKING



Higher bandwidth



P4 Programmable



Optional Deep Buffer

Q200 – DCI, WAN Network
Q200L – Web Scale Data Center Network

OPEN POSSIBILITIES.



Cisco Silicon One Foundation

IOS-XR

Platform Specific API 1

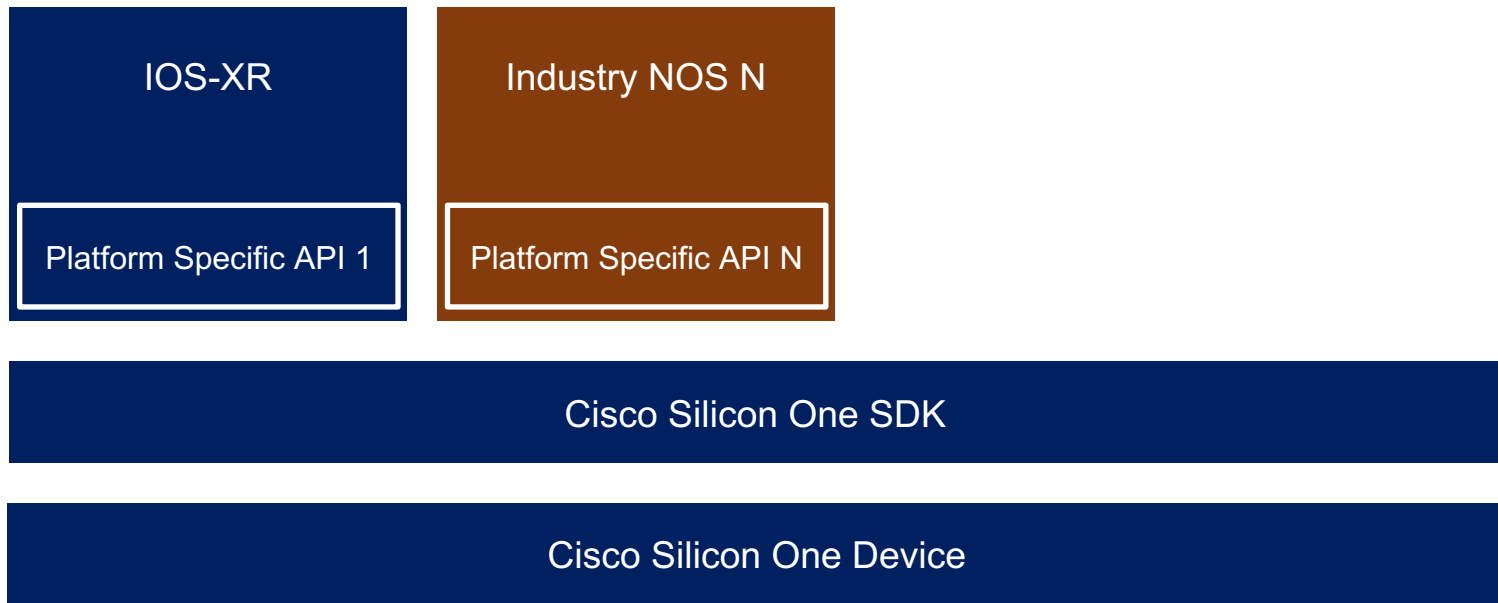
Cisco Silicon One SDK

Cisco Silicon One Device

OPEN POSSIBILITIES.



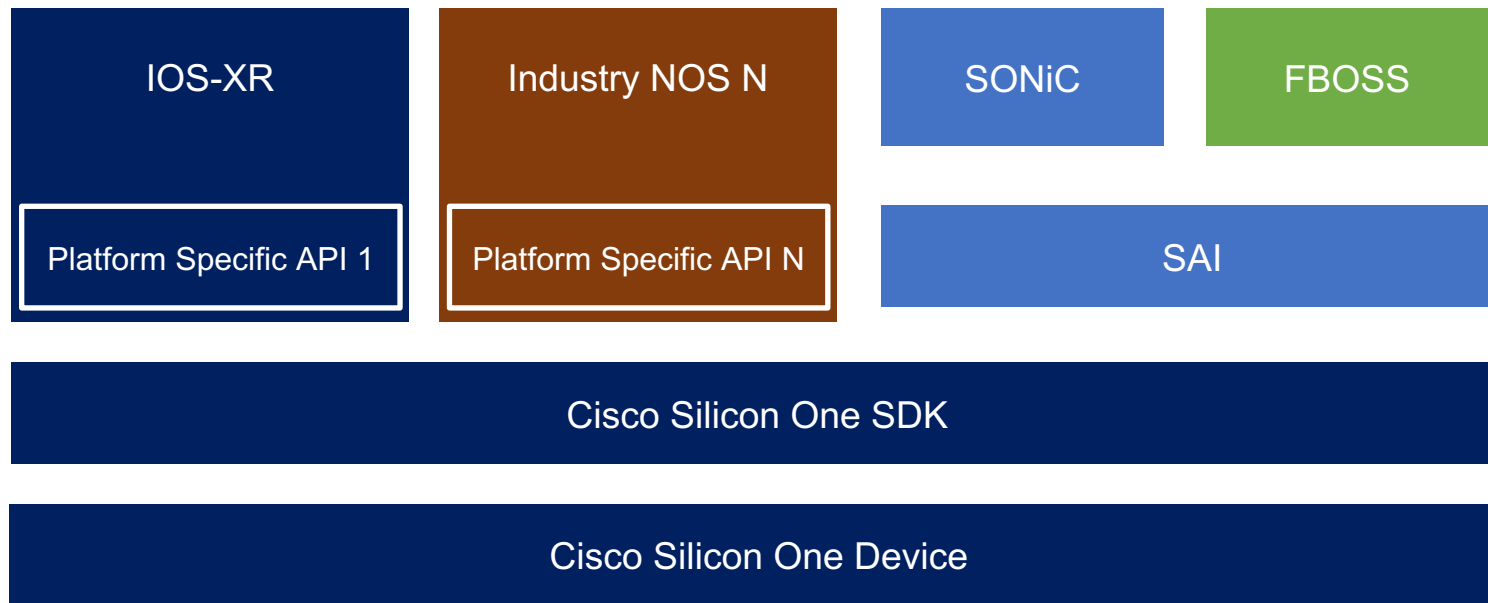
Cisco Silicon One Foundation



OPEN POSSIBILITIES.



Cisco Silicon One Foundation



OPEN POSSIBILITIES.





Joint Plan

Training

Hardware Build Out
Schematic Review
Simulation Review
Bring Up
Validation

CICD

Customer Production

Tests Completion

100%

Time

OPEN POSSIBILITIES.



Cisco Silicon One Success Elements



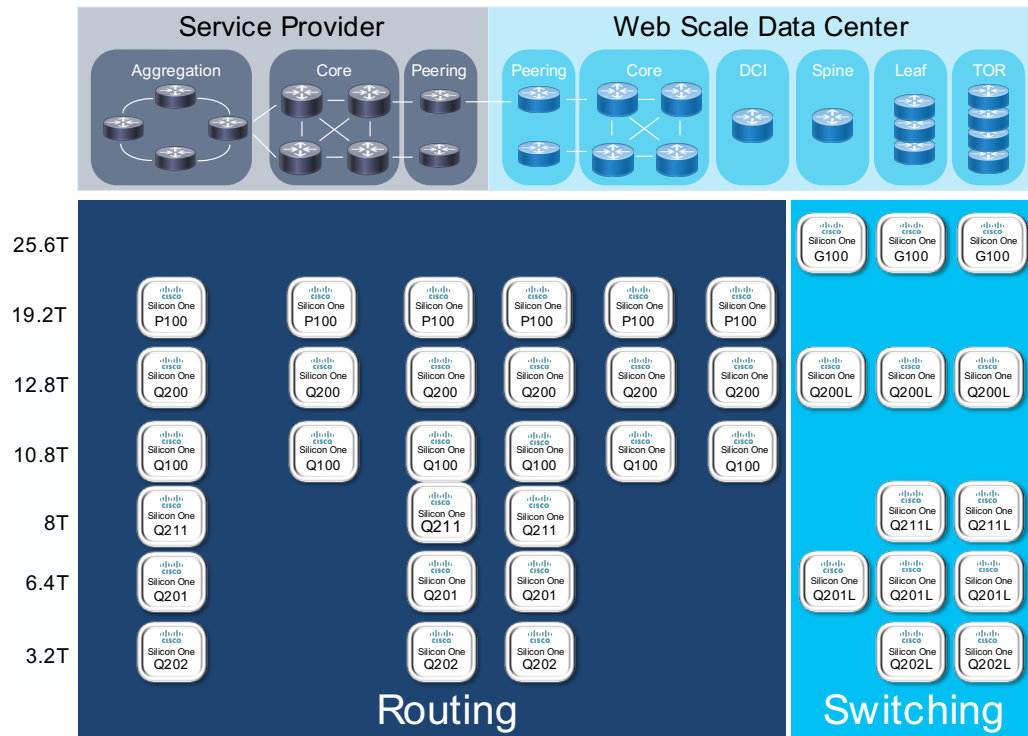
NETWORKING

- ✓ Exceptional teamwork across Facebook, Cisco and ODM vendor
- ✓ Continuous issue tracking with rapid closure
- ✓ Joint world class engineering talent

OPEN POSSIBILITIES.



Cisco Silicon One Across The Network



One Architecture
One SDK
One Forwarding Code
One Form-Factor
One Design

One Network
One Experience

Without Compromise

OPEN POSSIBILITIES.

Call to Action

- SAI Spec revisions should not break warm-boot
 - e.g. enum re numbering has broken warm-boot in the past.
- SAI Spec enhancements
 - Faster turnaround

OPEN POSSIBILITIES.



Thank you!



NOVEMBER 9-10, 2021