

An abstract graphic on the left side of the image, composed of numerous thin, light green lines that curve and swirl together to form a complex, organic shape resembling a stylized flower or a tunnel. The lines are set against a solid dark blue background.


# Open. Together.



**OCP**  
SUMMIT

# Minipack and F16, Software



Alex Eckert, Software Engineer  
Facebook



Open. Together.



# Minipack and F16, Software



# Software everywhere

Management plane

Control plane

Data plane

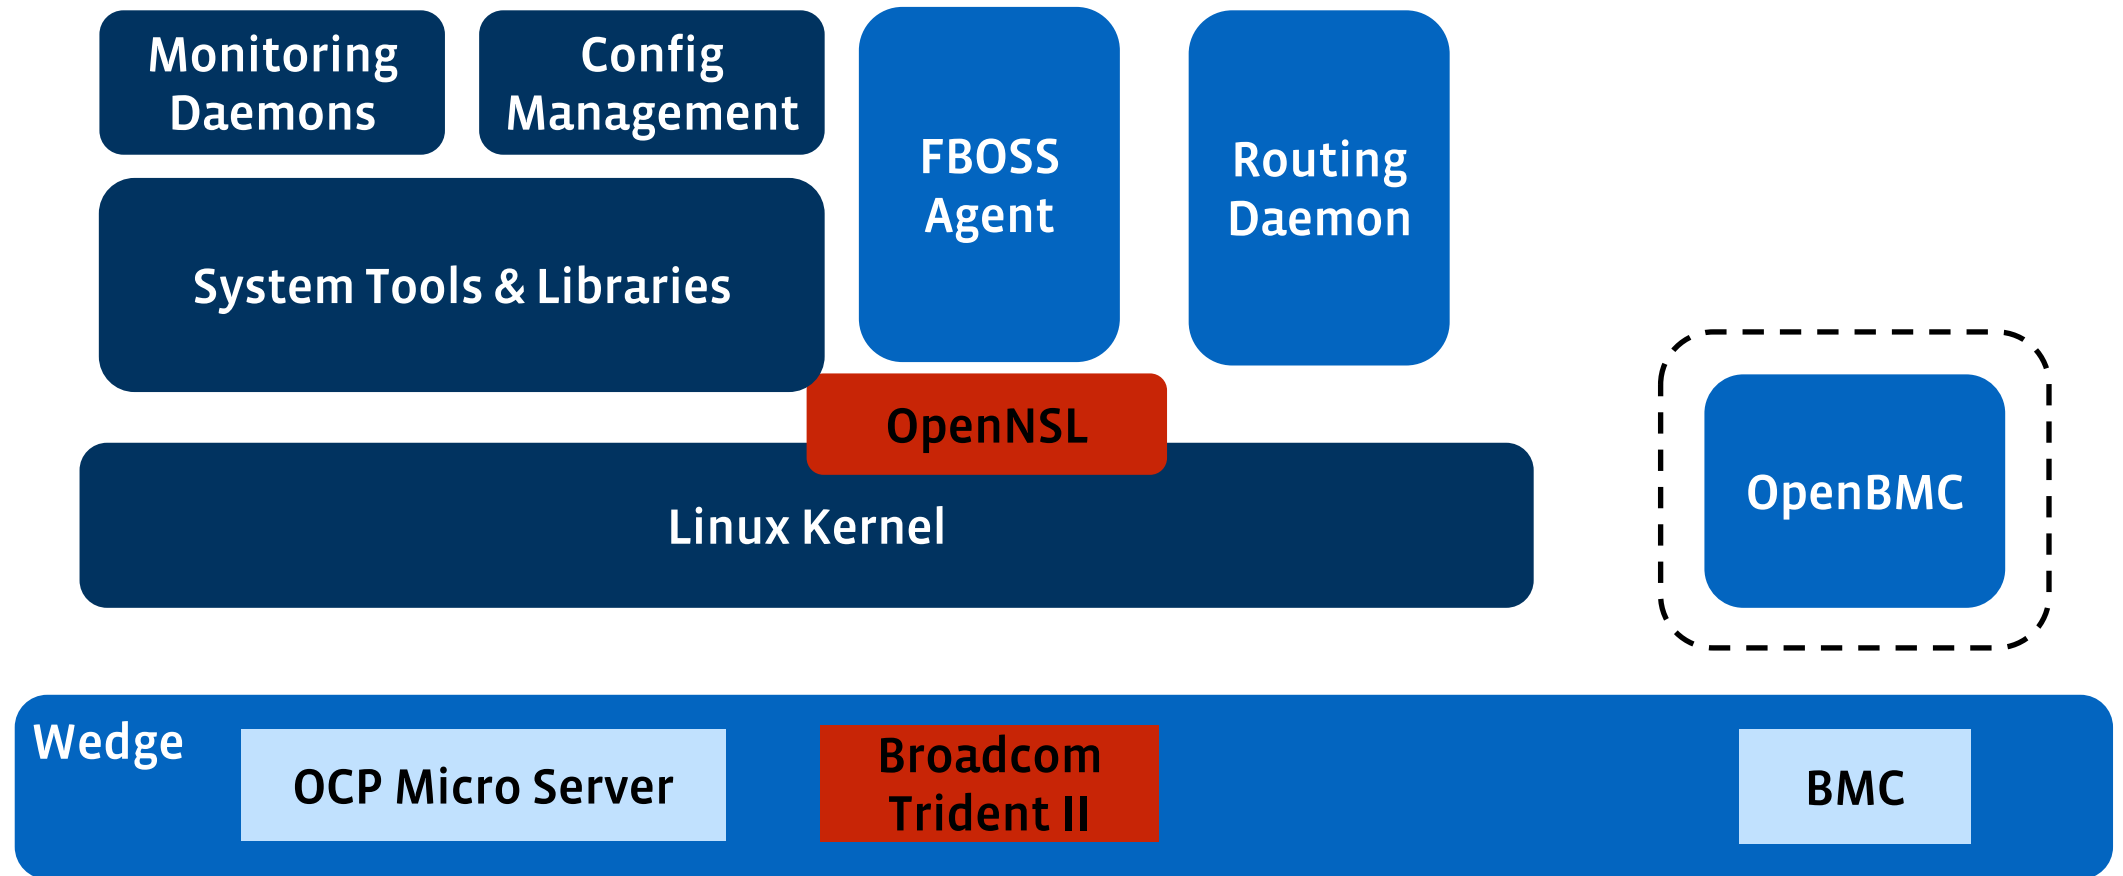
Platform

# Wedge



**16 x 40G - RSW**

# FBOSS Overview





# Sixpack



**128 x 40G - FSW**

# Wedge 100



**32 x 100G - RSW**

# Backpack



**128 x 100G - FSW**


# Wedge 100S



**32 x 100G - RSW**



# Fabric Aggregator



**N x 100G - FA**

# Minipack



**128 x 100G – FSW, SSW, FA**

# New challenges (Minipack)

- Modular switch, single control plane
- Different PIM types, variety of port speeds
- External phy support (Broadcom)
- I2C to 128 optics, MDIO, FPGA
- New microserver (MiniLake), New ASIC (Broadcom TH3)
- Cooling 128 optics in only 4 RU



**128 x 100G – FSW, SSW, FA**

# Arista 7368X4



**128 x 100G – FSW, SSW, FA**



# New challenges (FBOSS on Arista)

- First time operating FBOSS on non-FB hardware
- Different external phy (Credo), Different FPGA
- Unfamiliar microserver, no off-the-shelf UEFI BIOS
- Must manage as either FBOSS or EOS in production
- Complex conversion process b/w EOS / FBOSS



**128 x 100G – FSW, SSW, FA**

# Combinations

- Two switch models
- Modularity
- Three layers of F16 network
- Same binary everywhere

# Common abstractions

## Facebook Minipack

OCP MiniLake



Brcm TH3

Brcm Gearbox

Edgecore Board  
(FPGA/PIMs/PSUs/Fans)

## Arista 7368X4

Arista uServer

Brcm TH3

Credo Gearbox

Arista Board  
(FPGA/PIMs/PSUs/Fans)

# CentOS



**FBOSS**




OpenBMC



# New microservers!

## MiniLake

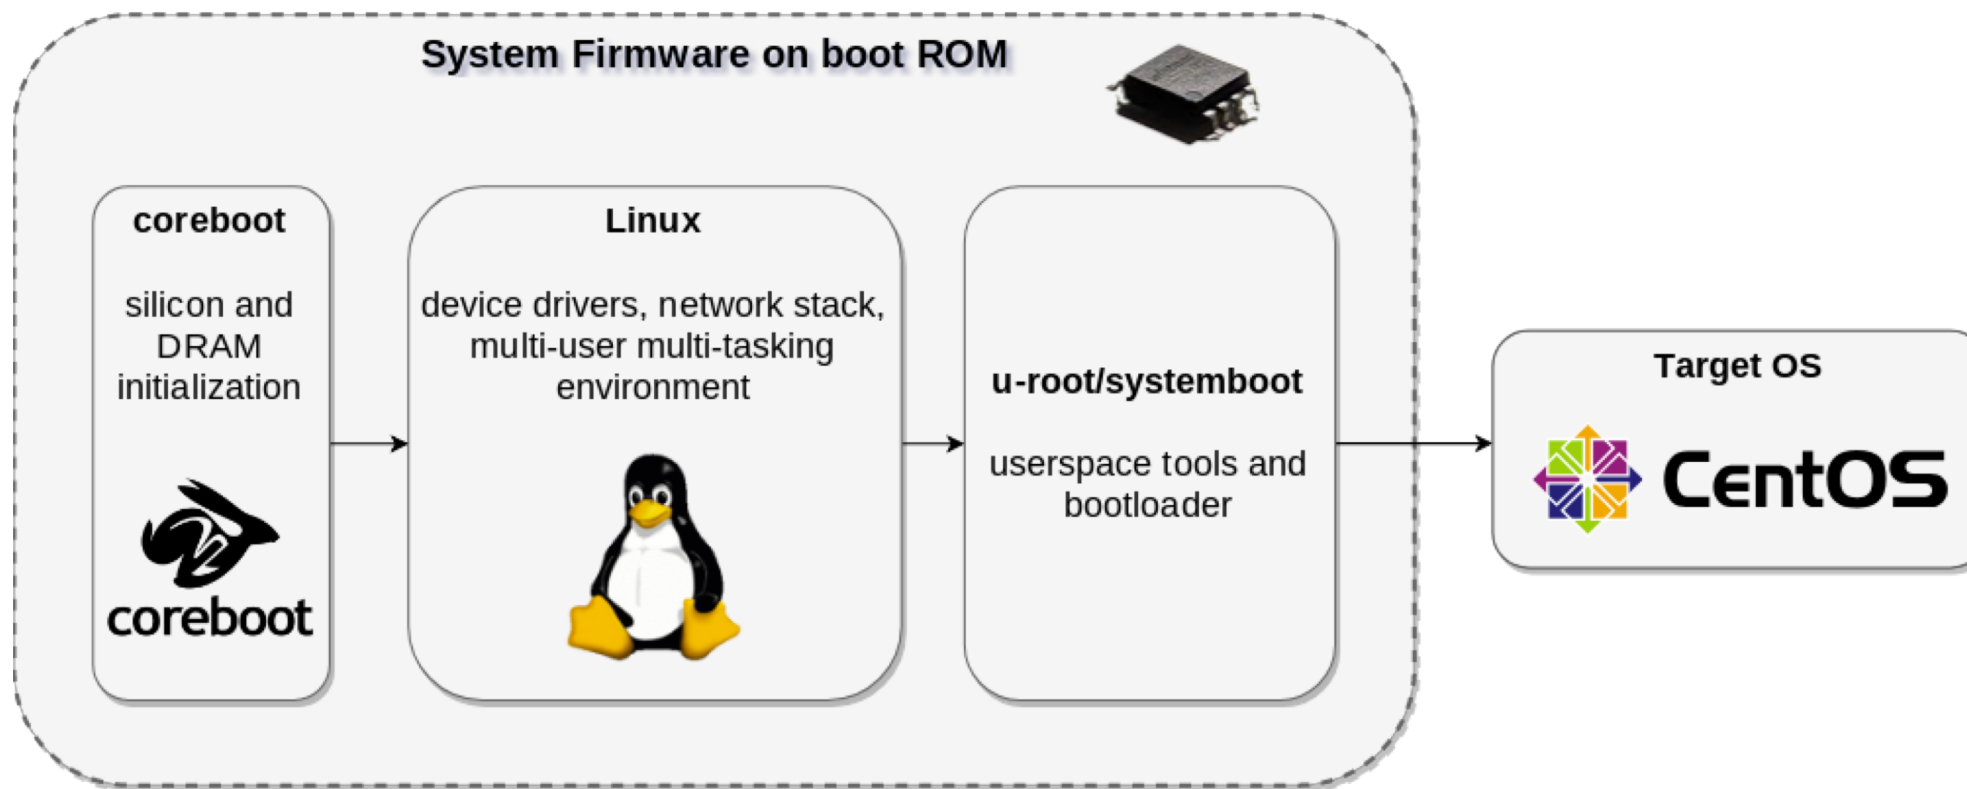


# New microservers!

## Arista control module



# OCP Open System Firmware



# Conversion

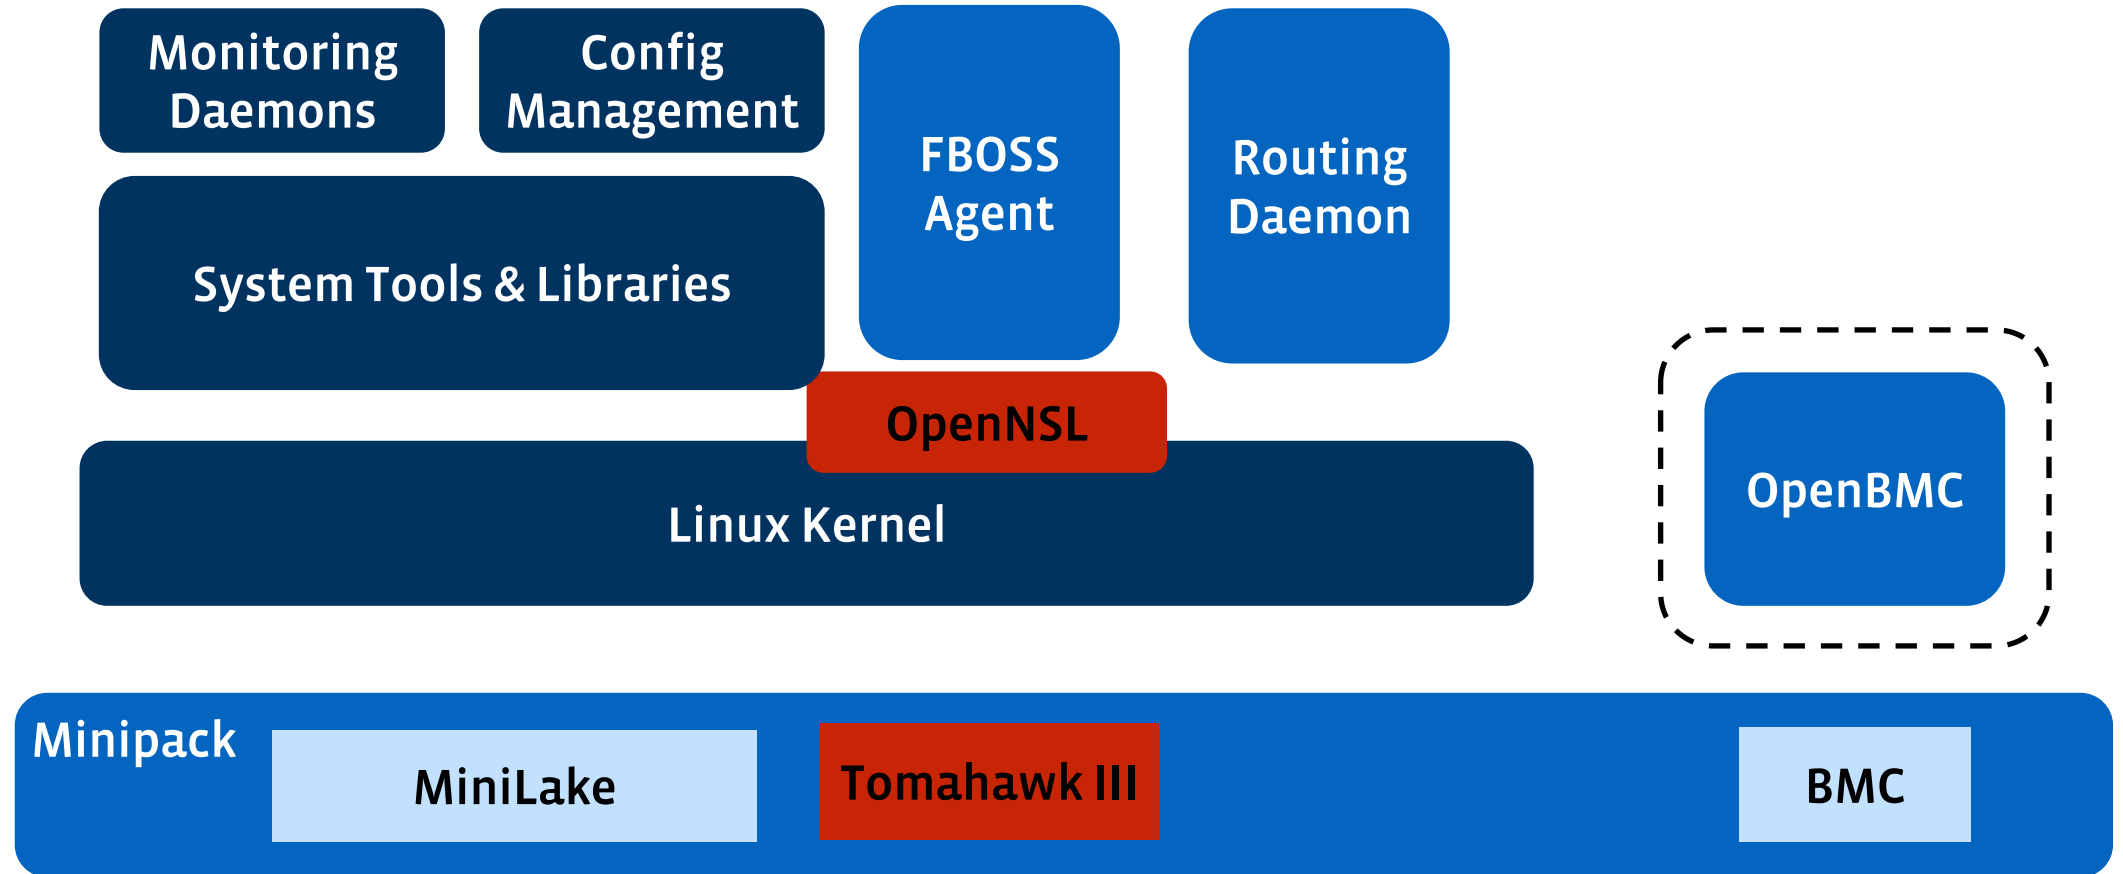
ARISTA  
EOS



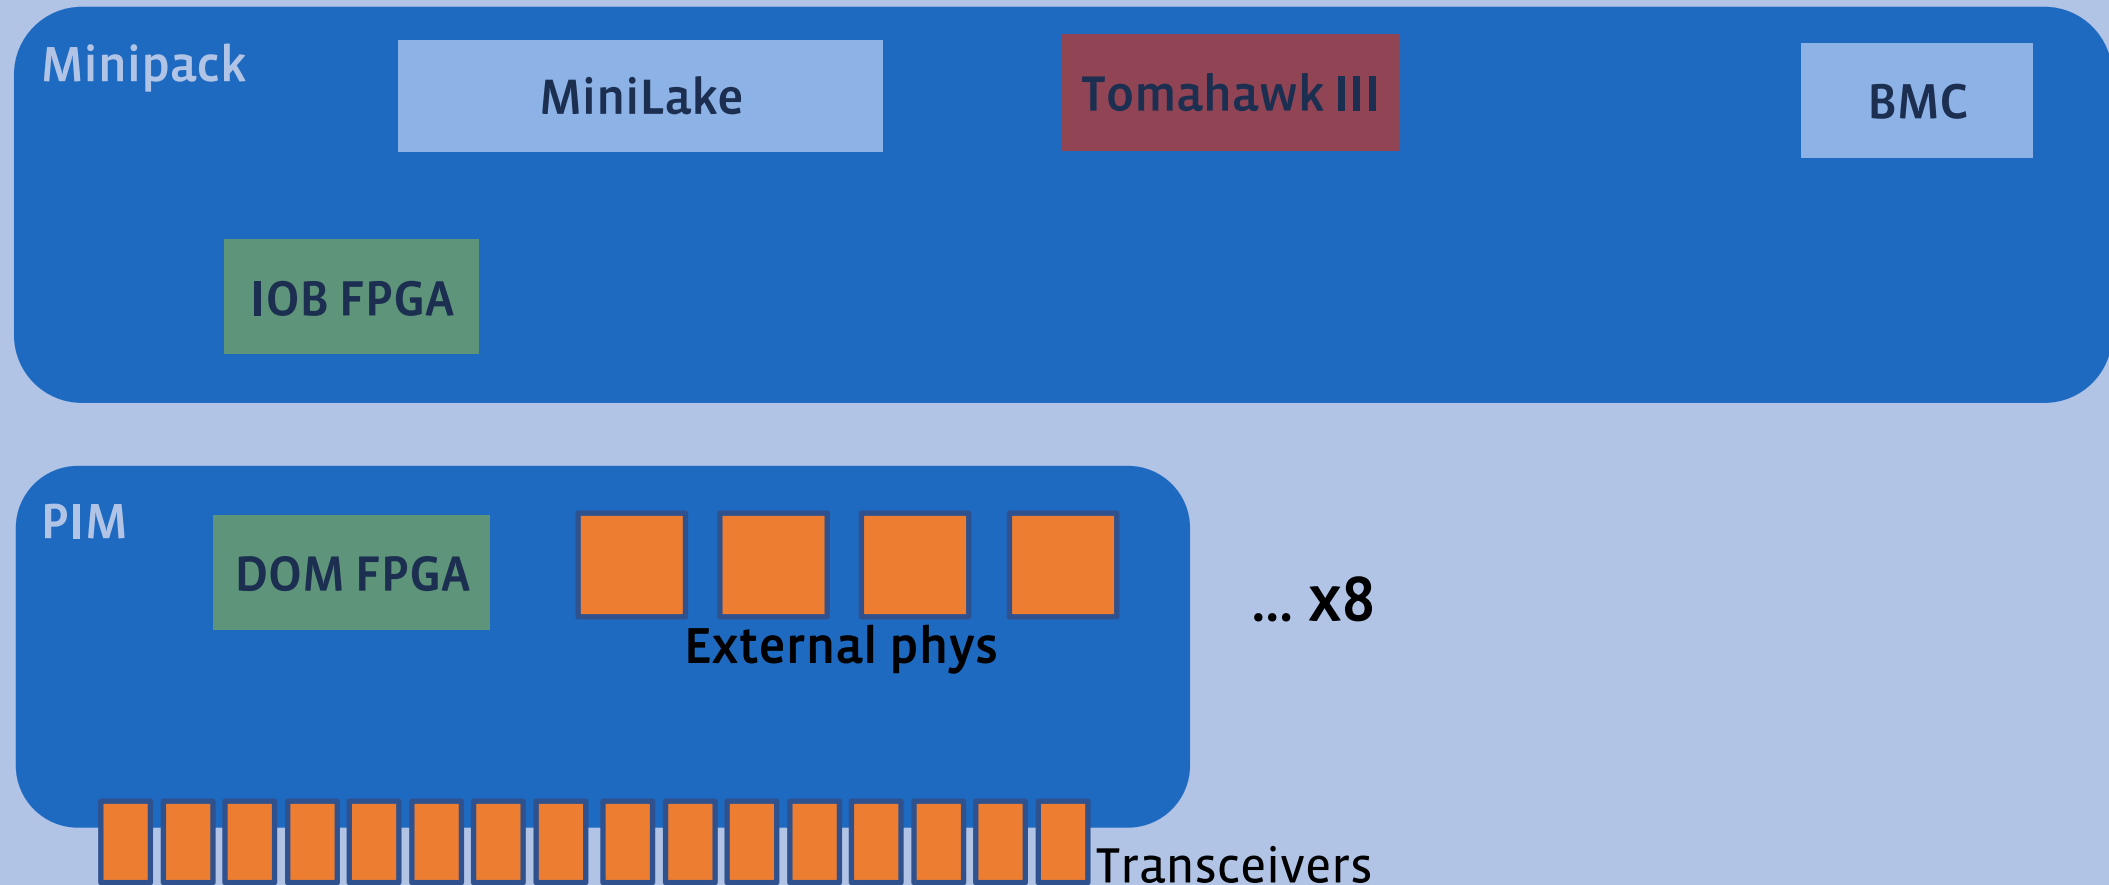
***FBOSS***



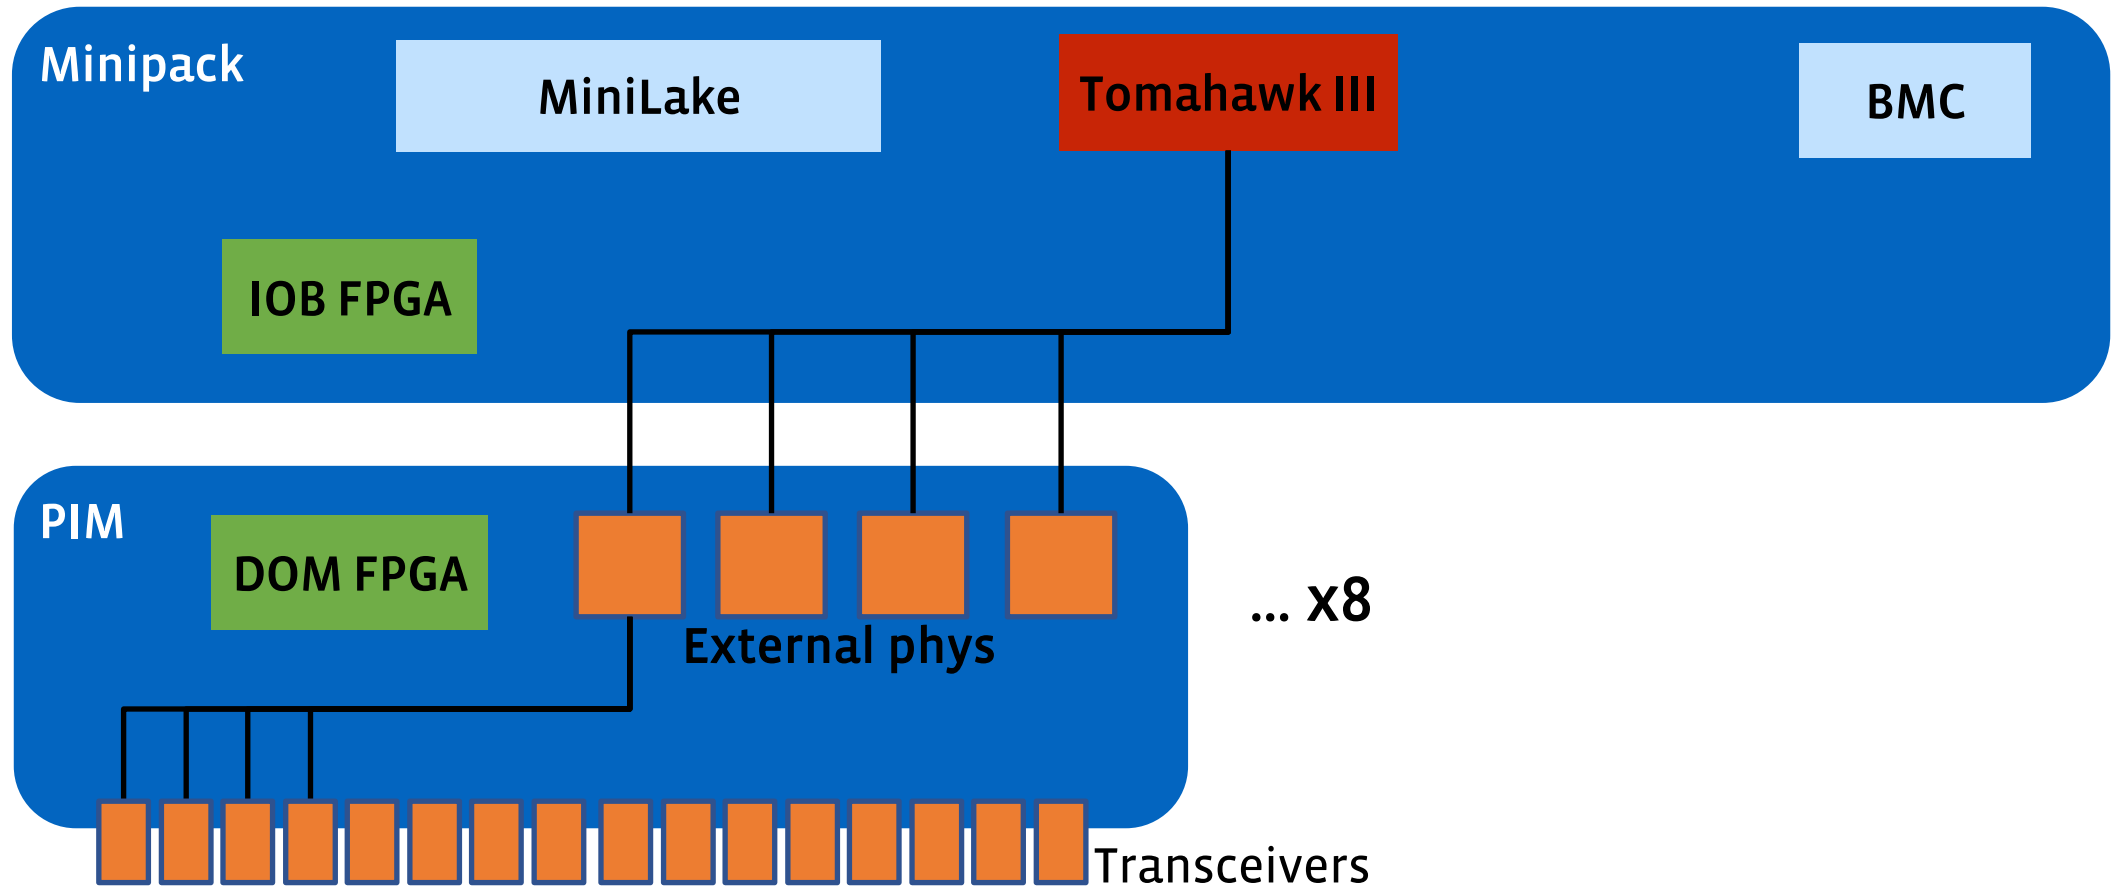
# FBOSS Overview




# Minipack Hardware Components



# Data path

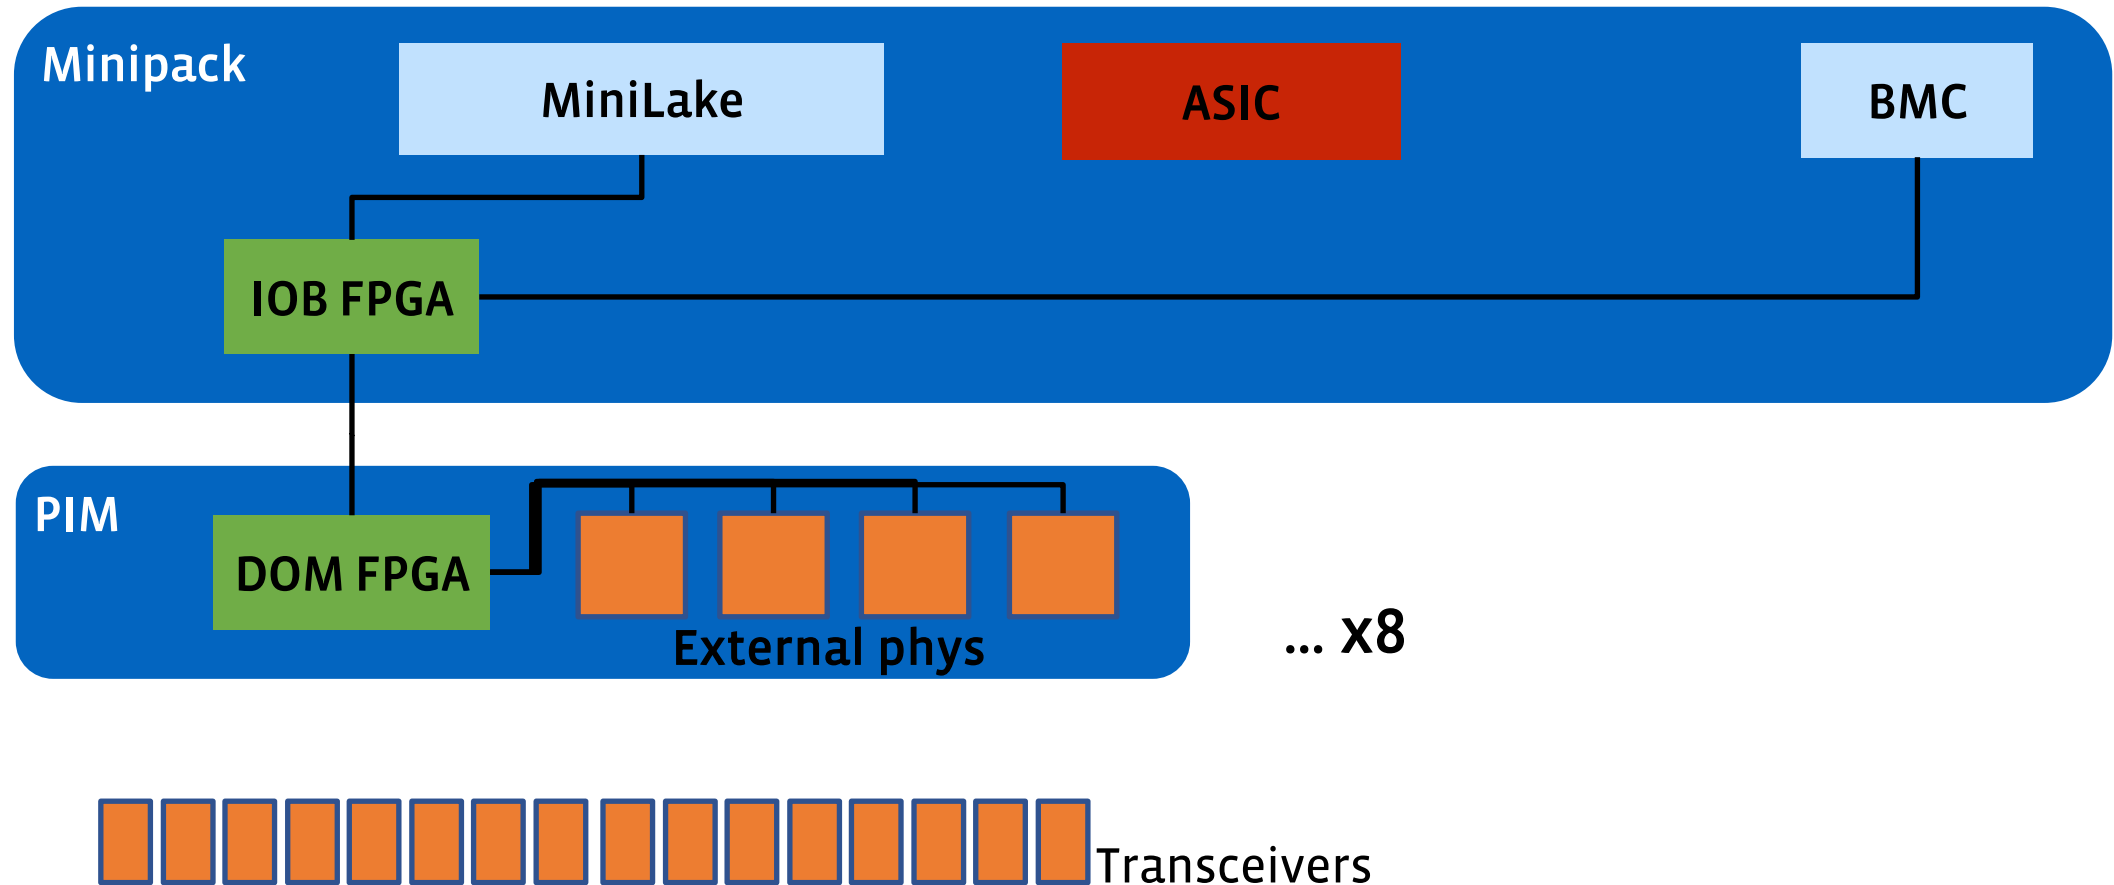


# Management path – I2C







# Management path – MDIO




# Programming External PHYs



# Programming External PHYs



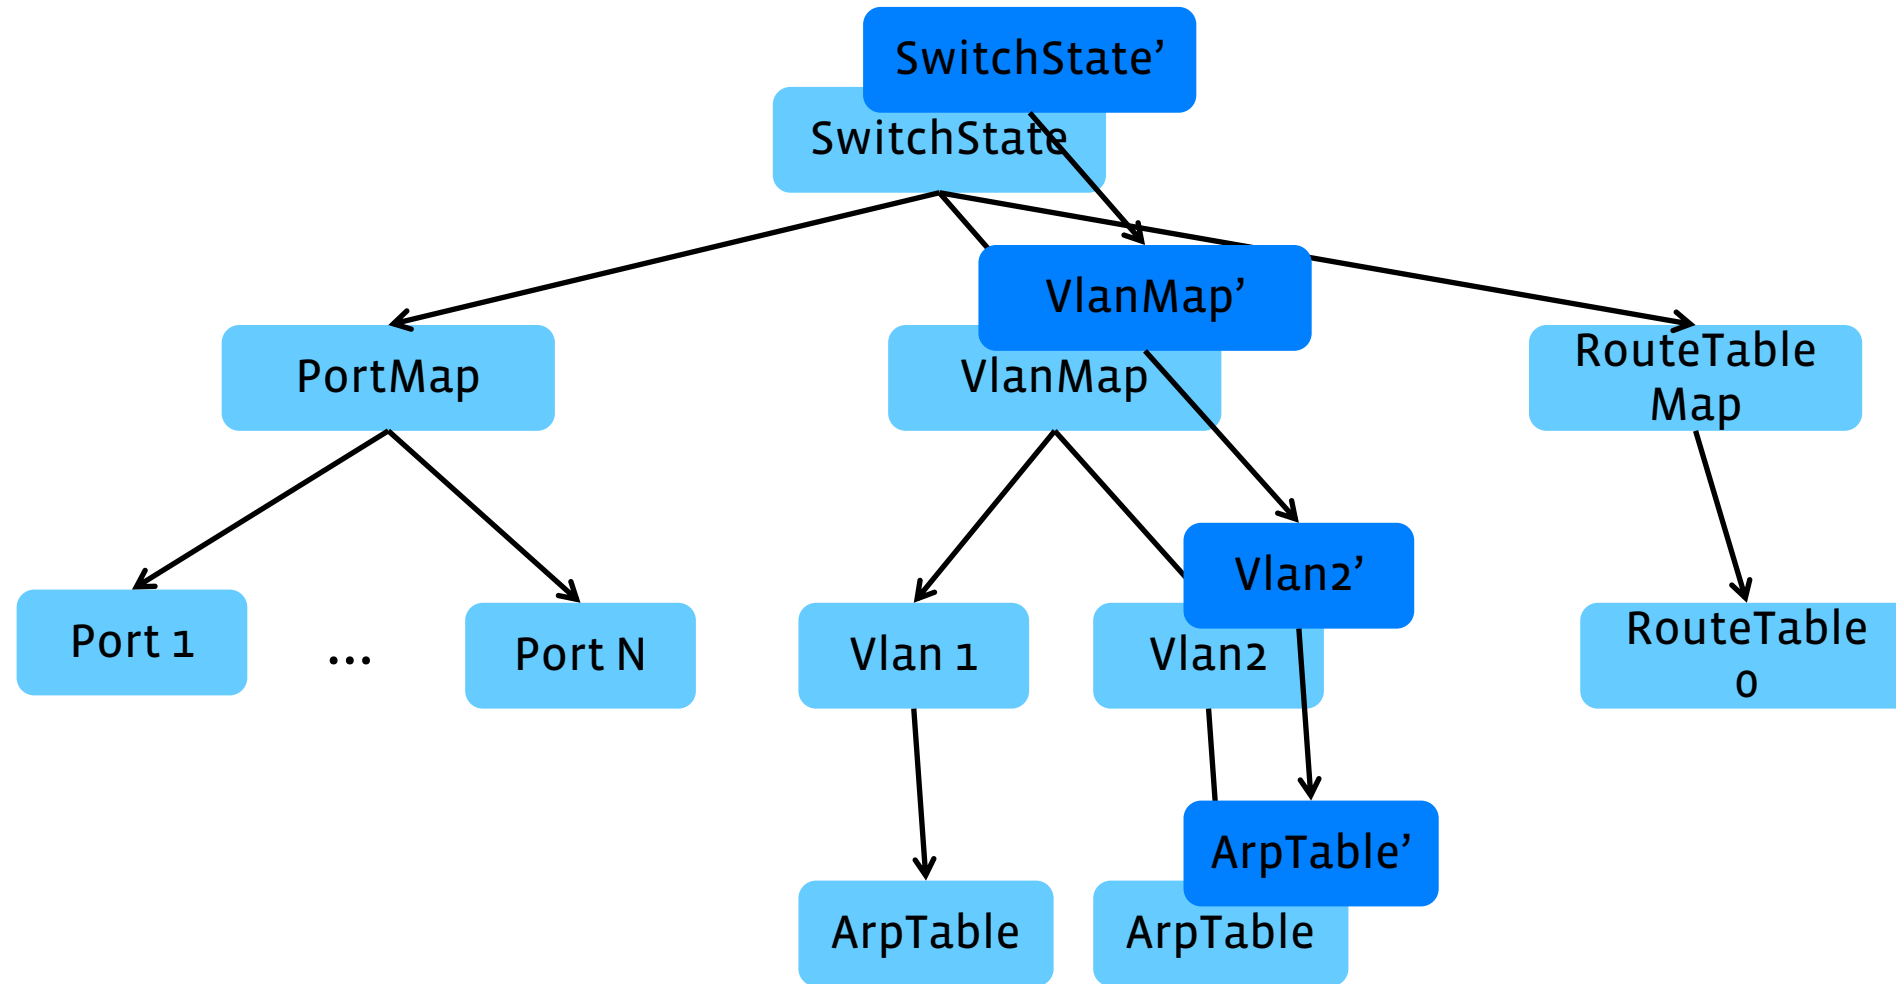
# Programming External PHYs






# State management problem

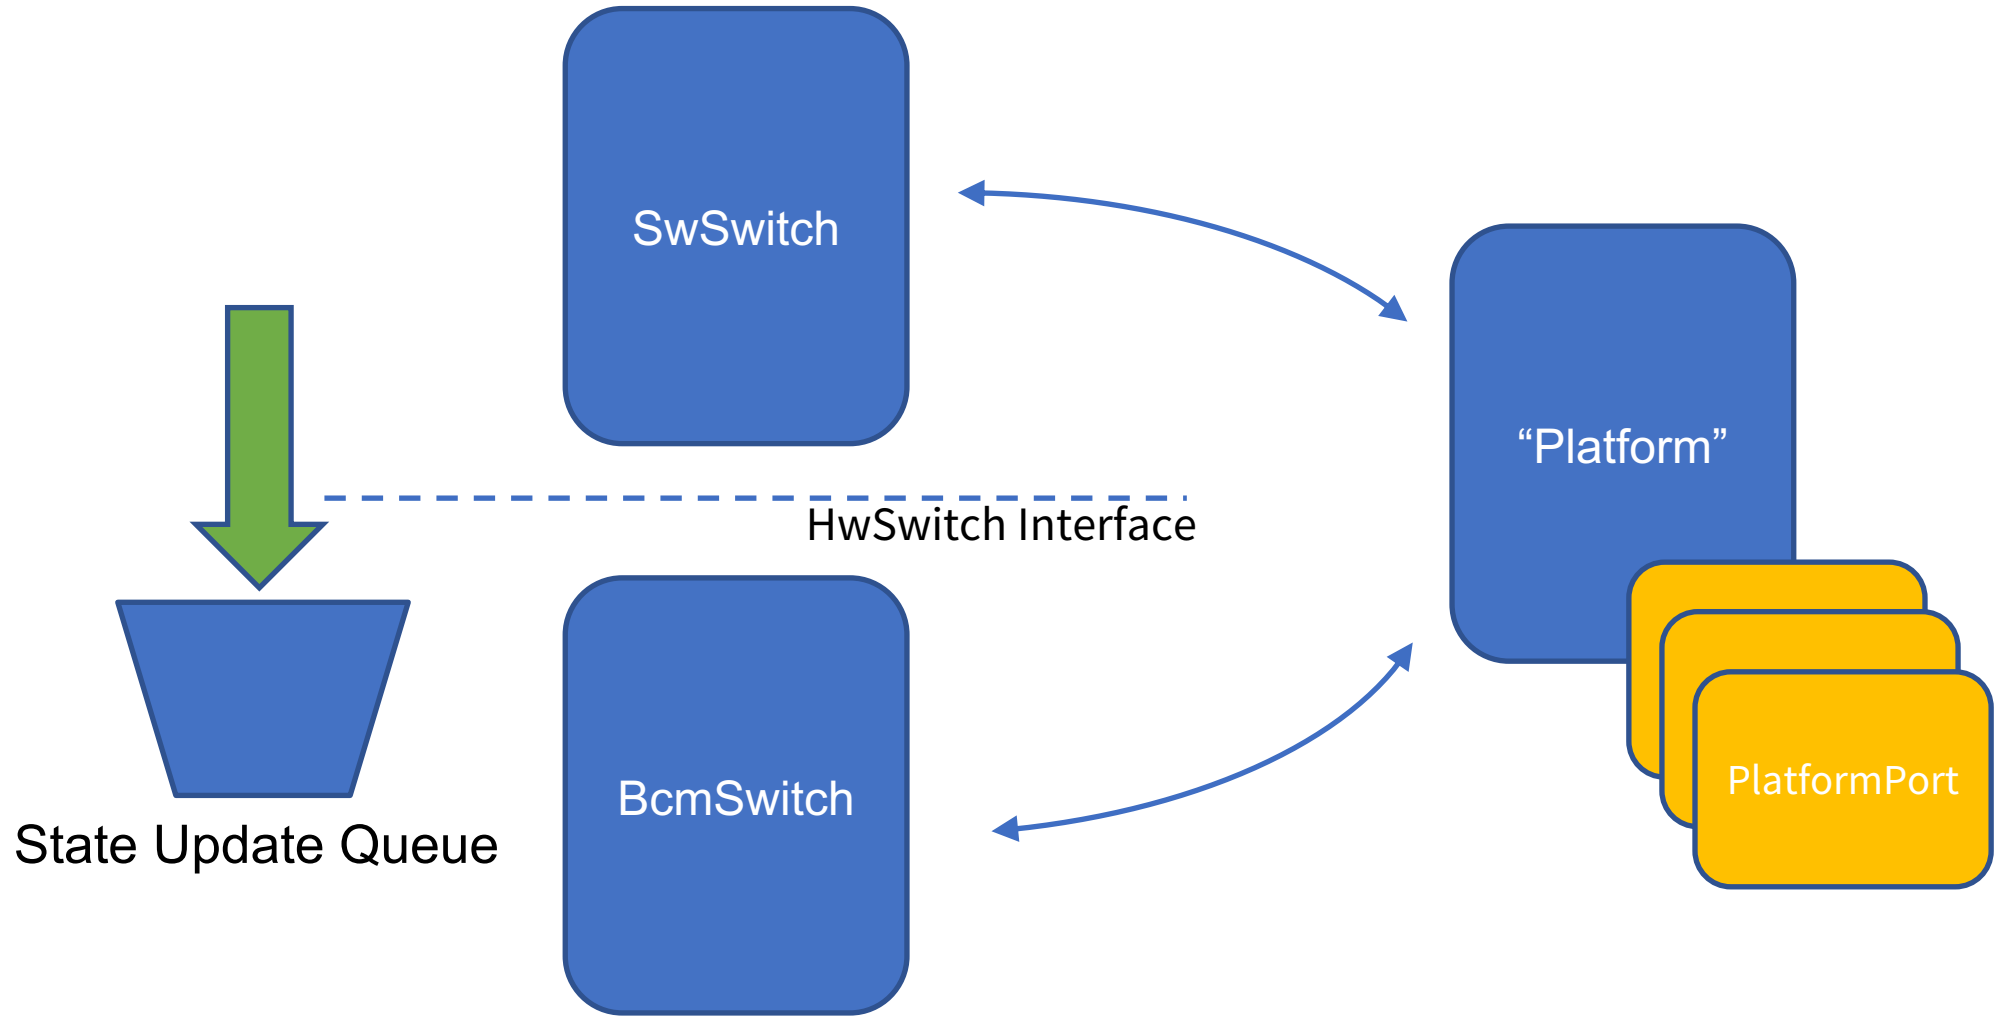
# Switch State



# FBOSS Agent



# FBOSS Agent



# Platform Configuration

- Each port + speed has unique settings (PlatformPortSettings)
- Model config as a map<speed, PlatformPortSettings>
- Hides complexity from SwSwitch

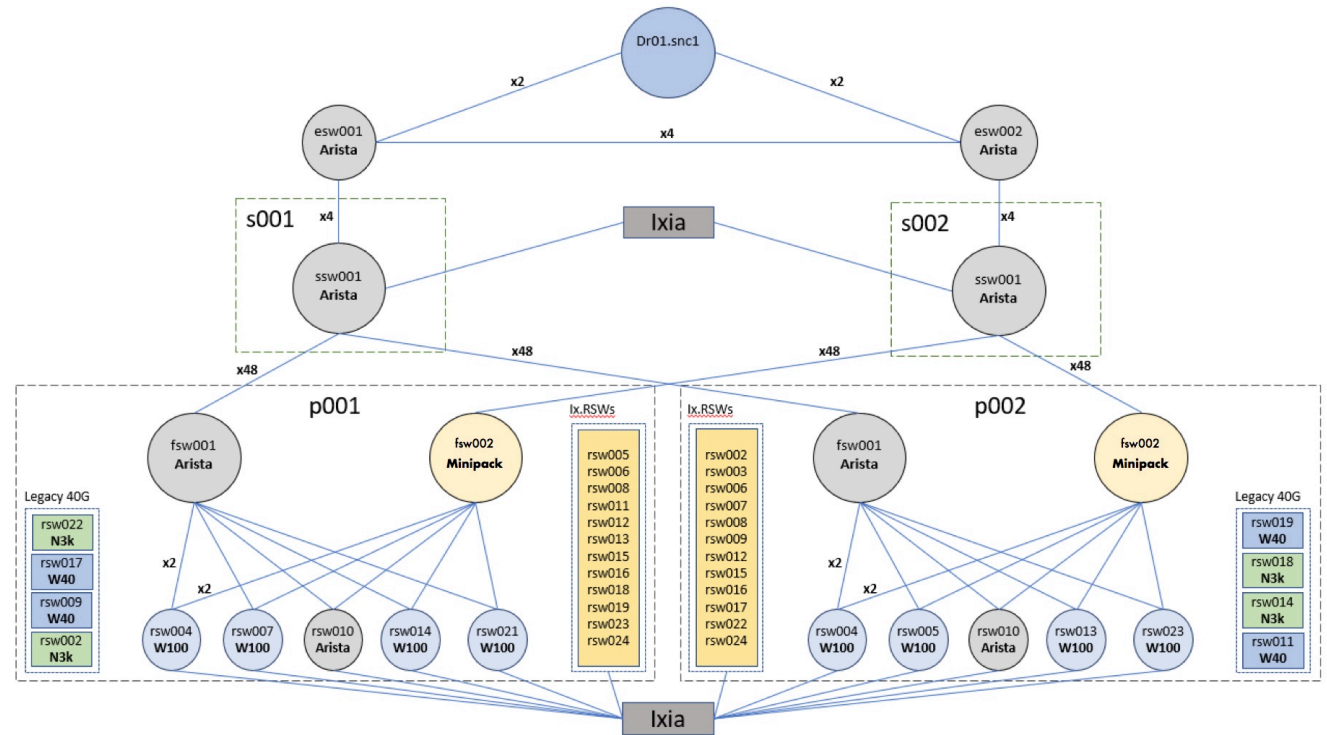




# Testing

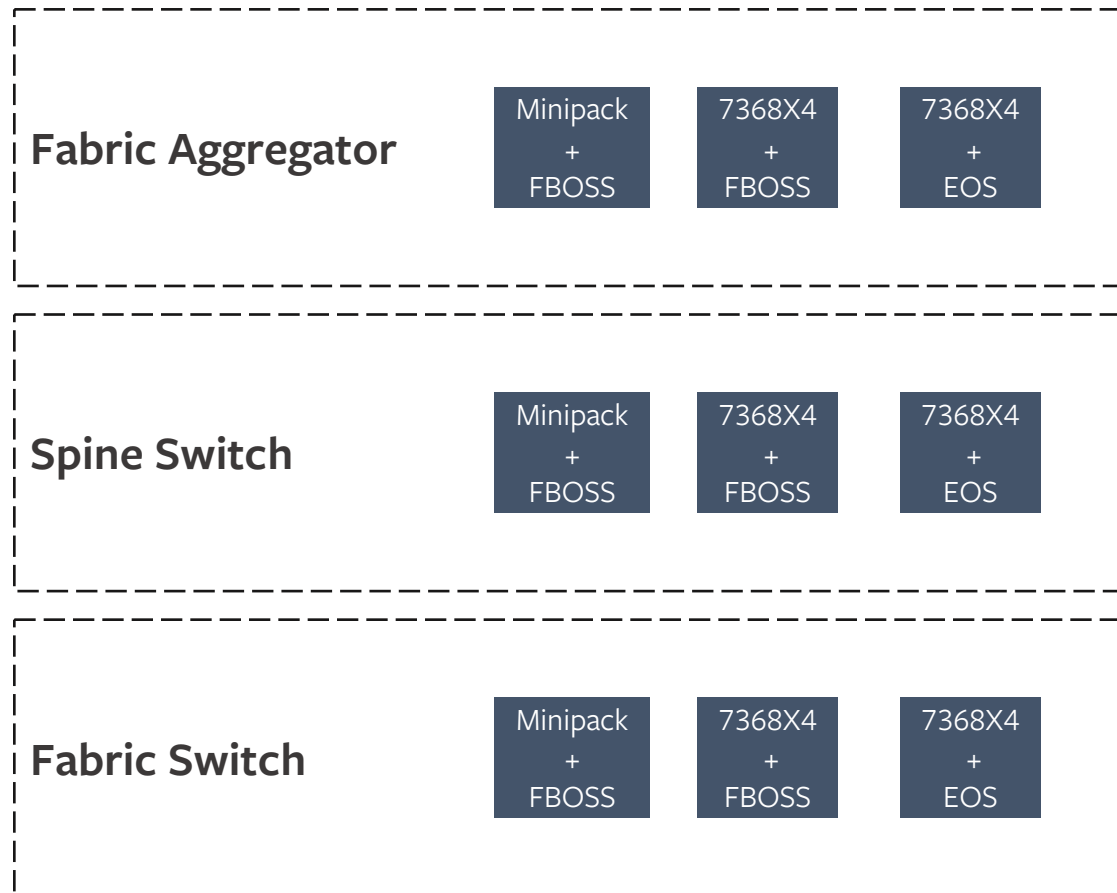
# Testing strategy

- Targeted tests in realistic lab environments
- Deploy early
- Invest in automation

# Scale challenge



# Test all combinations




Deploy early

# Invest in automation

- Automate as much as possible of our lab testing
- Develop targeted ASIC data plane tests

Test specific asic functionality, verify, warm boot





# Utilize common FB testing infra

● Good • fboss_test_infra • 12.81 sec	<div>✓ Good Tests 1,913 100%</div> <div>⚠ Broken Tests 0</div> <div>⚠ Flaky Tests 0</div> <div>🌱 New Tests 0</div>
netcastle_test/fboss_bcm/tomahawk_alpm/6.4.10-4_6.5.13-1 - AclEmptyCodeIcmp (warm_boot.BcmAclQualifierTest) architecture: x86_64, buildsystem: buck, compiler: gcc, sanitizer: none ● Good • fboss_test_infra • 416.86 sec	
netcastle_test/fboss_bcm/tomahawk_alpm/6.4.10-4_6.5.13-1 - AclIcmp6Qualifiers (cold_boot.BcmAclQualifierTest) architecture: x86_64, buildsystem: buck, compiler: gcc, sanitizer: none ● Good • fboss_test_infra • 12.89 sec	
netcastle_test/fboss_bcm/tomahawk_alpm/6.4.10-4_6.5.13-1 - AclIp4TcpQualifiers (warm_boot.BcmAclQualifierTest) architecture: x86_64, buildsystem: buck, compiler: gcc, sanitizer: none ● Good • fboss_test_infra • 71.08 sec	
netcastle_test/fboss_bcm/tomahawk_alpm/6.4.10-4_6.5.13-1 - AclMirror (cold_boot.BcmMirrorTest/1) architecture: x86_64, buildsystem: buck, compiler: gcc, sanitizer: none ● Good • fboss_test_infra • 8.55 sec	
netcastle_test/fboss_bcm/tomahawk_alpm/6.4.10-4_6.5.13-1 - AclMirror (warm_boot.BcmMirrorTest/1) architecture: x86_64, buildsystem: buck, compiler: gcc, sanitizer: none ● Good • fboss_test_infra • 1182.56 sec	

# Takeaways

- Minipack is a powerful modular building block for our networks.
- Software support for modularity has its challenges.
- We were able to overcome these challenges through improved hardware layer abstractions and investing heavily in early deployment and automated testing.

Thank you



# Open. Together.

OCP Global Summit | March 14–15, 2019

