UNBOXING THE WHITE BOXES

ANDREAS POLYRAKIS

CHRISTOS ARGYROPOULOS

TASOS KARALIOTAS

GRNOG 6/7/2018

OUTLINE

- Scope
- Use Case
- Overview
- Merchant Silicon Vendors
- HW vendors
- SW options
- Some details
- What we did

SCOPE

- Give an overview of the whiteboxing trend
- Discuss about possible use cases and implementation fields
- Present some options, bits and bytes and first thoughts
  - We just unboxed few whiteboxes
  - Not a thoroughly tested field
- Discuss what actually changes in the networking world
- Disclaimer: surely not all players mentioned, not intention

USE CASES [GR-IX UPGRADE]

- Now a pure Ethernet fabric
- Loop avoidance by STP
- Goals
  - More bandwidth
  - High speed ports (40G/100G)
  - Loop free topology without disabling BB links => IP fabric with some overlay (VxLan) and finally EVPN control plane
  - Unchain reseller function
- Budget Constrains

GR-IX TARGETED TOPOLOGY

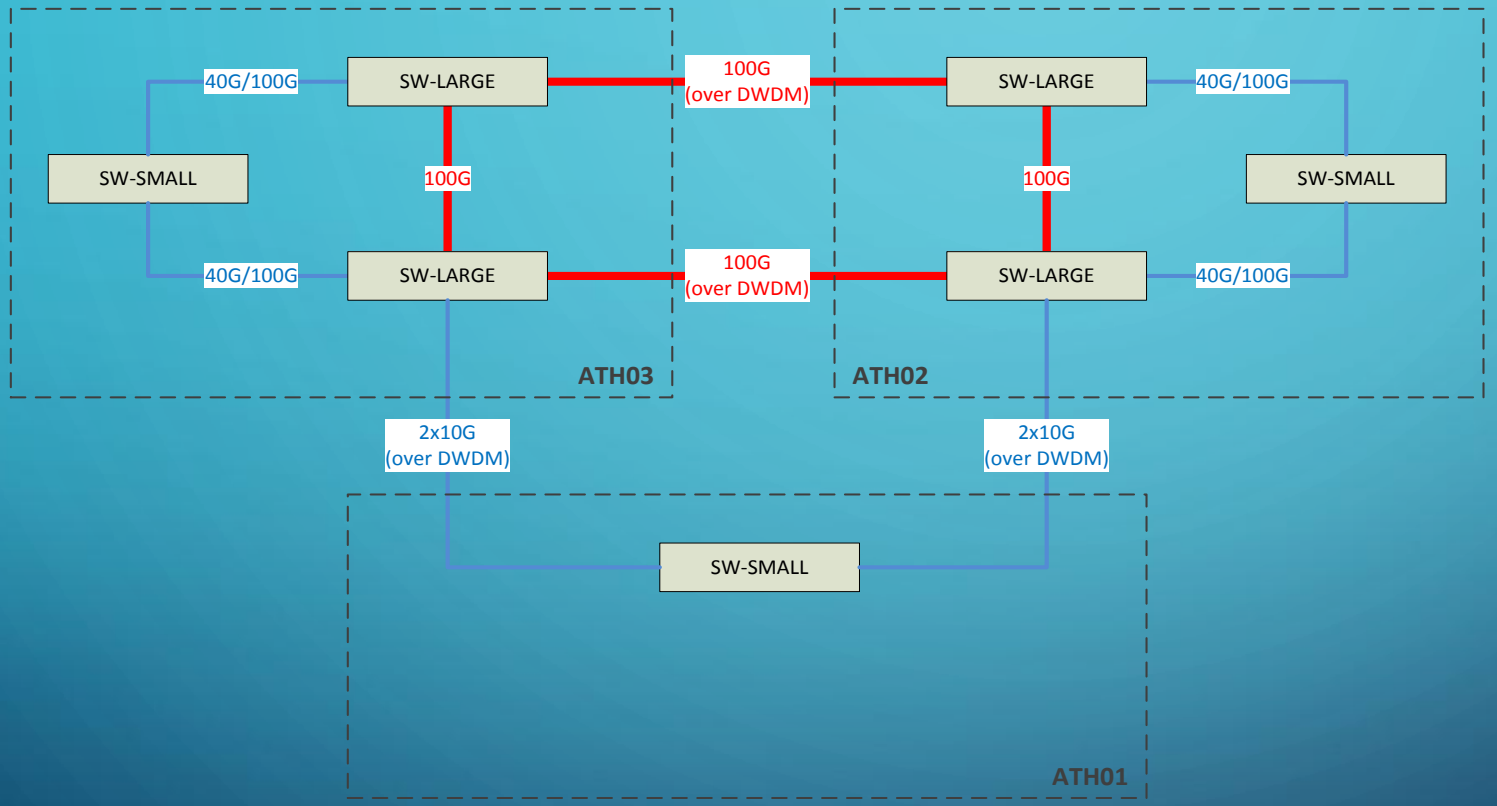

Where:

SW-LARGE: ≥ 12xQSFP28 ports

SW-SMALL: ≥ 2x(QSFP+ or QSFP28), ≥48xSFP+ ports

So sw-[large|small] could be

- Classical HW&SW in a box
- Disaggregated solutions with separate HW &SW vendor.

WHITEBOXING IN A NUTSHELL

- Merchant silicon (ASICs)
- HW vendors - integrate merchant silicon (ASIC) and of-the-self components to switches
- SW vendors that develop OS, protocols and technologies

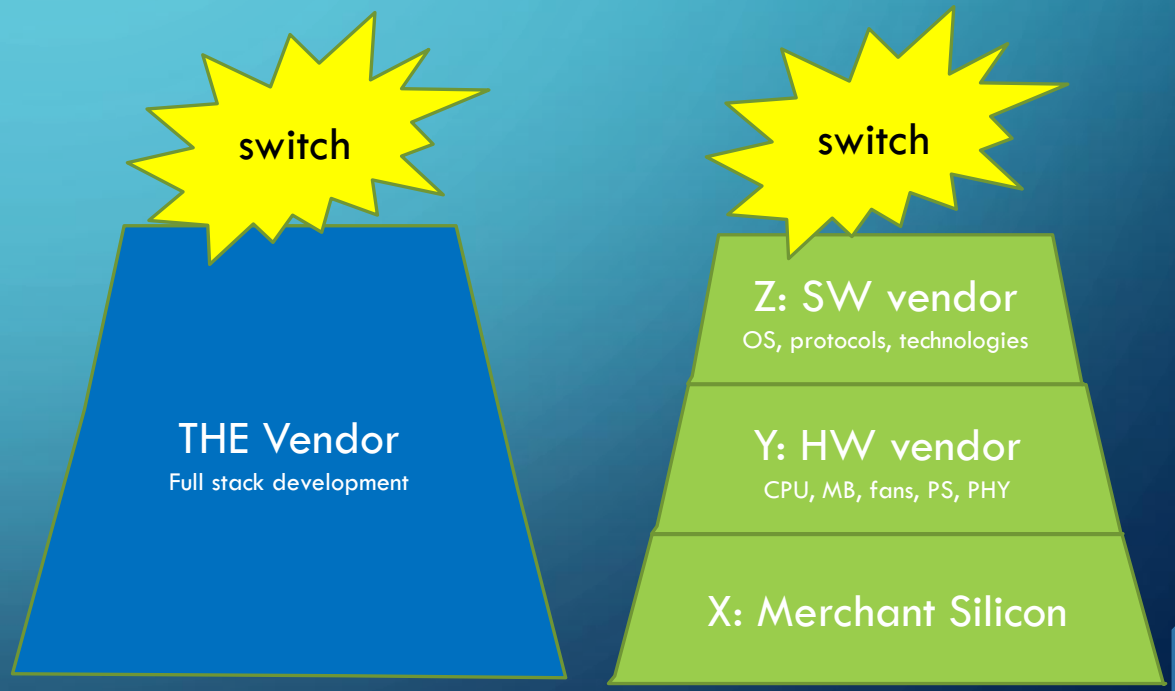

DEFINITIONS

- Proprietary Switch

  Switch that carries a custom made ASIC.

  Usually developed under one vendor's realm together with NOS

- Bare Metal Switch

  Switches from Original Device Manufacturers (ODMs) with no network operating system loaded on them. (EdgeCore/Delta/Alpha... etc)- ONIE enabled

- White-Box Switch

  A switch made by of-the-shelf components & merch. silicon with a NOS

- Brite-Box Switch

  Bare metal switch with a brand on it (eg Dell, HP). The Brand promises to offer better support and act as a single point of contact for both HW & NOS

IN A TABLE

	Bare metal switch	Branded bare metal (BBM)	White-box (WB) switch	Proprieratry switch
Definition	Hardware only with basic support from original design manufacturer	Hardware only with orginal equipment manufacturer branding and warranty/support/services	Commodity hardware and NOS preloaded	Propriertary hardware and NOS
Hardware cost	low	low	low	high
Type of hardware components	Off-the-shelf components including ASIC	Off-the-shelf components including ASIC	Off-the-shelf components including ASIC	Proprieratry
Network operating system	None (customer can load PicOS/Cumulus/Big Switch)	None (customer can load PicOS/Cumulus/Big Switch)	Vendor's own or 3rd party already loaded	Vendor's own NOS
Examples	Accton AS5712 (Broadcom) Penguin 4800 (Broadcom) Quanta 3048 (Broadcom)	Dell S4810-ON/S6000-ON (Broadcom) HP 5700/5712/6700 (Broadcom) HP 5700/5712/6700 (Broadcom)	Arista 7250X (Broadcom) Dell S6000 (Broadcom) HP 5930 (Broadcom)	Nexus 7000 HP 3500/5400/8200 (HP ProVIsion) Juniper 9200 (Trio)

https://go.forrester.com/blogs/15-02-23-what_is_the_difference_between_white_box_bare_metal_branded_bare_metal_and_oem_network_switches/

WHITEBOXING

- Not a new idea…completely analogue to IBM compatible PC story
- Open Compute Project
  - a father and great promoter (https://www.opencompute.org/projects/networking)
- What is supposed to offer

  Decoupling layers in networking gear manufacturing and production hopefully will
  - Reduce TCO for equipment & network
  - Create space for smaller players => Competition & Innovation
  - Give providers more control on their assets
- Standardization on how components interact
  - Open Network Install Environment (ONIE) – Cumulus/OCP
    - Standardized bootloader that permits to load NOS on a bare metal switch
  - Switch Abstraction Interface (SAI) - OCP
    - Set of open APIs allowing code written for ASIC A to run on ASIC B as long as both support SAI
- New support model for equipment components (HW & SW)
  - Split model
  - One stop shop is offered but things generally may get messy (to be proved)

SILICON VENDORS

- **BroadCom**
  - **StrataXGS Family – Trident, Trident II and Tomahawk.**
    - Intended for low cost top of rack (ToR) switches, low density, and network edge use cases.
    - Broadcom BCM56960 Tomahawk 5. 25/40/50/100Gbe ToR switch.
  - **StrataDNX Family – Dune.**
    - Used for higher cost devices with more demanding needs/ 800 Gbps of packet processing external packet buffers and advanced packet processing
- **Barefoot Networks**
  - Tofino Chip
- **Mellanox**
  - Spectrum
- **Cavium**
  - XPliant® CNX780XX/CNX680XX Family
- Marvell Technology Group
  - Prestera

BARE METAL VENDORS

- Corsa
- Edge-Core (Accton)
- Penguin
- Quanta (QCT)
- Mellanox
- Dell
- HPE

WHITE BOX VENDORS

- Arista
- Cumulus (Cumulus Express)
- Mellanox
- Dell
- HPE

NOS VENDORS

- Big Switch: Security and DC implementations
- Metaswitch: Tailor made packages
- IP in Fusion: OcNOS
- Cumulus Networks : Cumulus Linux
- Dell: OS10
- Pluribus
- SONiC (Software for Open Networking in the Cloud) - Azure
- PICA8: PICOS

PARADIGM SHIFT

- Purpose made software stacks
  - Pay only for the software you need
  - Supposedly less buggy (proportional bugs/code lines)
- Changing HW Vendor but stick on the same NOS (or less often vice versa)
  - No learning curve
  - Continuity on tools, procedures and apps
  - Less Vendor lock-in (better bargaining position)
- Software Defined Networking

SOME MORE BITS

- Great majority of NOSes are based on Linux
  - Code base
  - Tools and mindset from Linux word, automation / portability
    - Puppet /chef for switch management
    - Unifying the way network and servers are managed
  - Protocol implementation market
- Offerings are mainly focused on DC implementations
  - Cause it is a push from DC operators
  - There is progress towards campus/enterprise and SP environments
  - Good penetration on IX market (at least in Europe)
  - Diversity on the focus and “special characteristics”

WHAT WE DID ON WB

- Tested a number o NOSes on virtual lab environment (EVE-NG)
  - First impression
  - In parallel with the RFI => Better understanding
- Tested more extensively Cumulus Linux on EVE-NG
  - Based on the RFI results it scored better against our bill of features
  - Created the whole fabric in the LAB
  - Tested specific characteristics and functionality
- Demo Dell Switches with CL

GR-IX PoC - CUMULUS

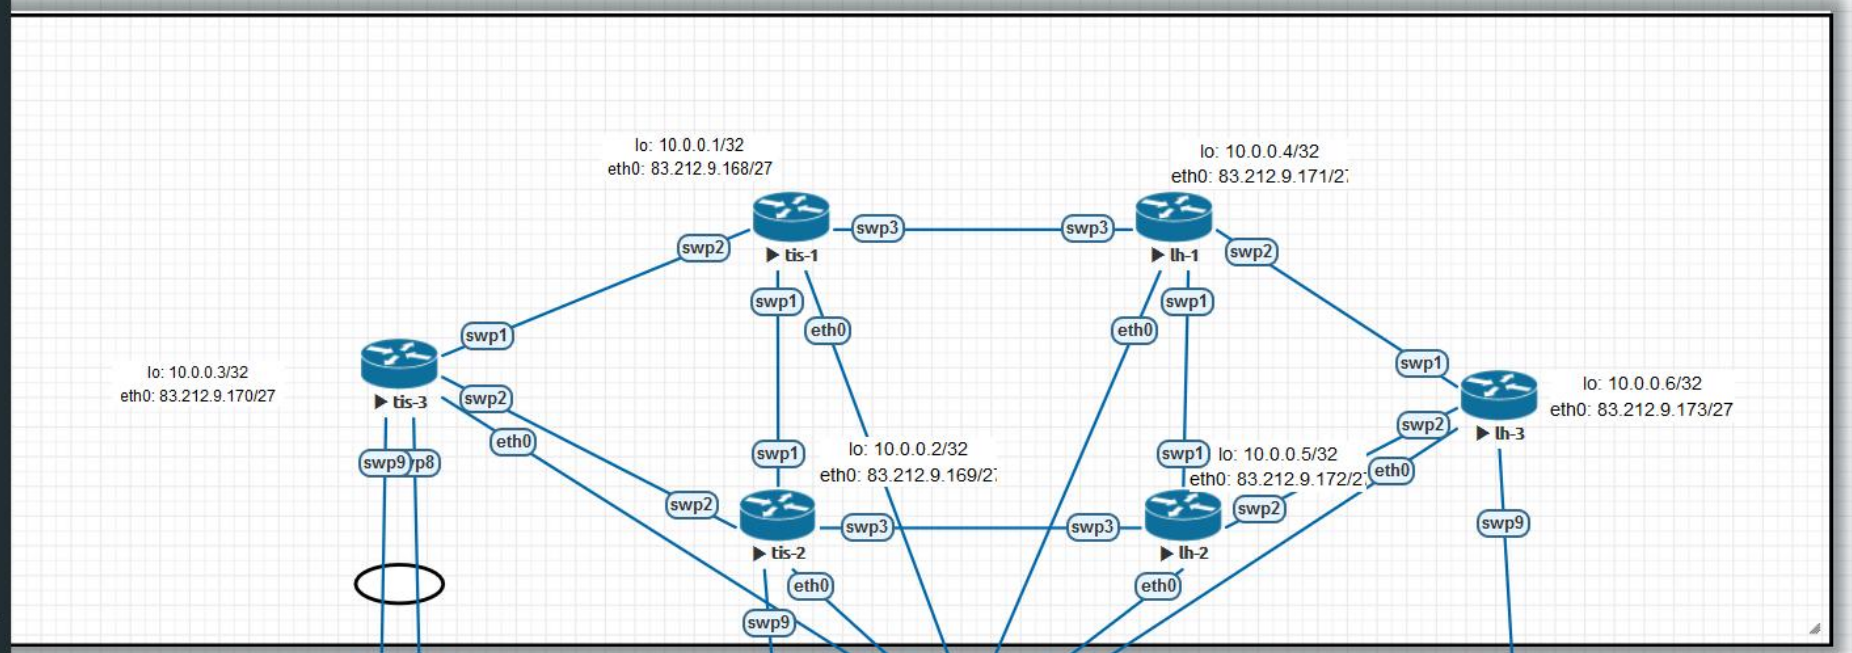

THANK YOU!

QUESTIONS ?