

White-Box for NREN

GN4ph3: Objective and strategy

JEANNIN Xavier
WP6T1 task leader

STF meeting, Bucarest, 26/02/2019

Public

www.geant.org

White Box: What is it?

- A Buzzzzzzzz Word



- A pragmatic approach driven by NREN use cases
- Put in « **production** » if operation constraints are satisfied

White Box: What is it?

- A switch/router that is able to run different Network Operating System (NOS)



- **The key idea is to be independent from traditional hardware vendors. Disaggregation trend between Network Operating System and hardware**

→ 2 levels of independence :

- independence from the hardware
 - you can change the hardware vendor and keep the software
- independence from the NOS
 - you change the NOS and keep the hardware

Proprietary design



- Business model
 1. Hardware design Proprietary
 2. Proprietary NOS (embedded)
 3. Hardware maintenance
 4. NOS maintenance
- Dependence from one vendor

White-box Design

- Same ASICs as in many well-known vendors which use off-the-self/commodity ASICs
 - Juniper/Cisco use the same chipset
- Performances and features depend on forwarding chipset (switching ASIC) and the NOS
 - Trident, Trident 2, Tomahawk, Qumran, Jericho
 - Commercial and open source NOS
- New business model
 1. Hardware design Proprietary and Hardware maintenance
 2. NOS (embedded) and NOS maintenance



Current White Box

Edgecore Wedge100BF-32X, 32x100GbE QSFP28 ports, Barefoot Tofino 3.2T, Intel Broadwell DE, Dual AC PSU

- Issue mainly from the data-center world
- Switch with very powerful forwarding capacity and limited number of features
- Cost effective
- New product for network provider
- The architecture of data-center, regional network and telecom carrier become very similar



Routing landscape is changing fast

- Router vendors use ASICs as in many well-known vendors which use off-the-self/commodity ASICs
- **The market is driven by data-center**
 - Architecture of data-center, regional network and telecom carrier become very similar
- **Traditional data-center switch vendor want to enter into the market**
- **Very cost effective**
- Classical vendors will not stay without any reactions
 - Price of legacy vendor box is already decreasing a lot
 - They have already a white offer



White Box: a trend

- ACCTON, Edge-Core, DELL, ...
 - Cloud, Large Data-Centers

- But also traditional router vendor start ...

- JUNOS OS FOR WHITE BOX DATA CENTER SWITCHES

- <https://www.juniper.net/assets/us/en/local/pdf/datasheets/1000641-en.pdf>

- Enabling IOS-XR on Third-Party Network Hardware

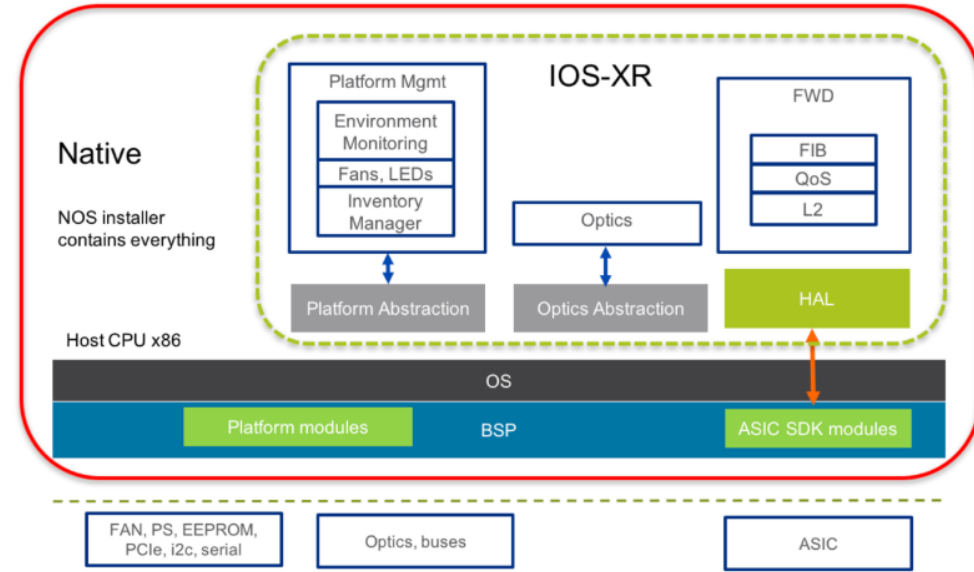
- <https://xrdocs.io/cloud-scale-networking/blogs/2018-03-08-enabling-ios-xr-on-third-party-network->



JUNOS OS FOR WHITE BOX DATA CENTER SWITCHES

Product Description

As cloud computing becomes more pervasive, service providers, cloud operators, and enterprises alike are seeking to deploy more scalable, agile, and automated data center architectures that employ standards-based network protocols and standards-compliant hardware.



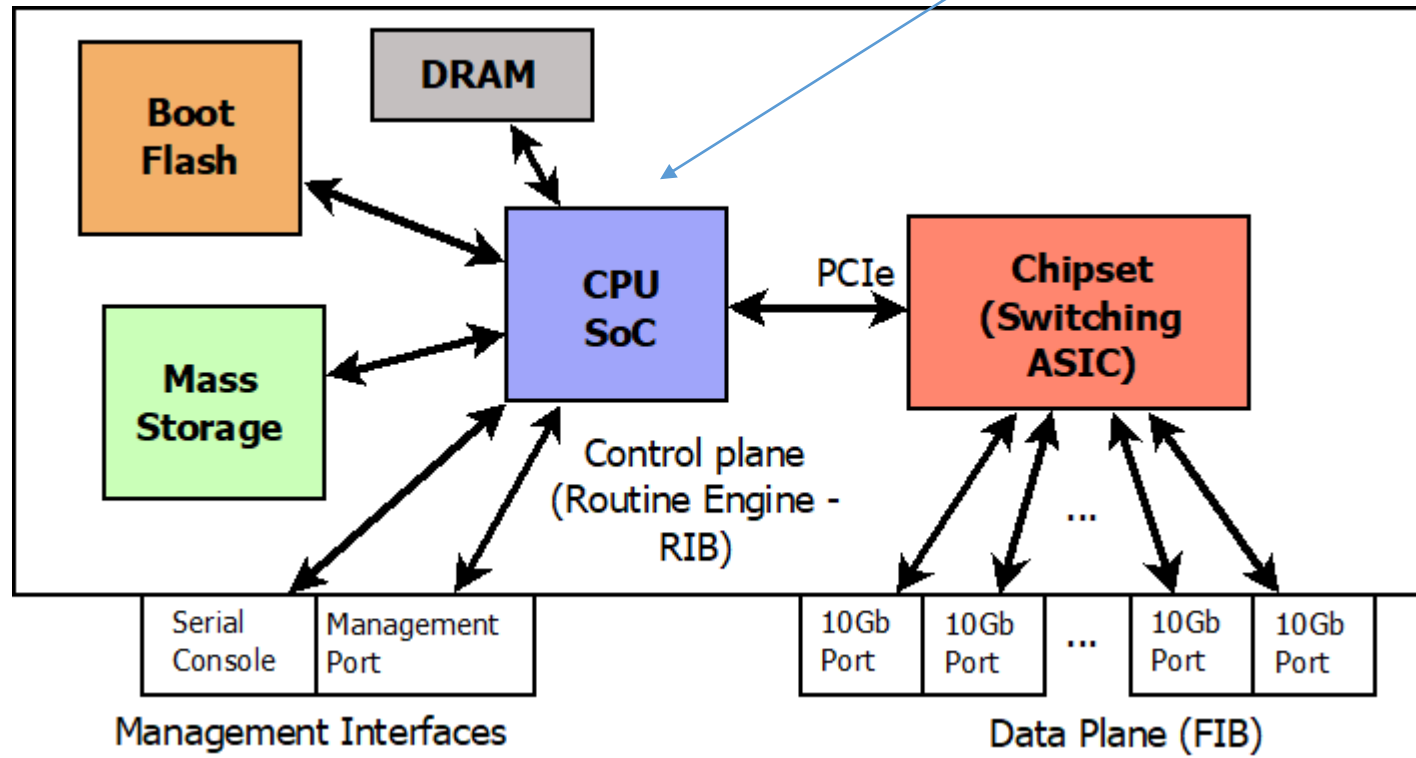
Network Operating System

NOS		commentaires
DELL OS9 ou OS10 (Free BSD)	Commercial	- Linux + Quagga + BGP EVPN (VXLAN ?)
Cumulus Networks	Commercial	- Unix + CLI →Data center
IP Infusion OcNOS	Commercial	- NOS for network operator - Still limited feature (see GN4ph2 report)
Pluribus Networks NetVisor	Commercial	- SDN solution but controller is embedded in each box - Specific new features: data analytics, service chaining and VNF functions
Open Network Linux	Open / Free	- https://opennetlinux.org/
Barefoot networks Sur chipset Tofino	Open et peut être commercial	- The future ? - Based on P4 chipset : Telemetry, DDOS mitigatos, load balancer, ...
Pica8 PicOS	Commercial	Hybrid Networking: OpenFlow agent, with native L2 and L3 features
Big Switch Networks SwitchLight	Commercial	SDN Solution
Canonical Snappy Ubuntu Core	Commercial	
Software for Open Networking in the Cloud SONIC	Open	Microsoft and co-contributors to OCP Orienté cloud pour AZUR
OpenSwitch	Open	Few features but very fast
SnapRoute FlexSwitch	Open	
Open Network Foundation Atrium SDN Distribution	Open	- CORD use case
Open Compute Project http://www.opencompute.org/projects/networking/	Open	Mothership project : - Campus Branch Wireless (CBW) - Open Network Install Environment (ONIE) - Open Network Linux (ONL) - SONIC - Switch Abstraction Interface (SAI)

White-box architecture



In general, 1 routing engine



Chipset characteristics

- Buffer and memory
 - Packet queue length
- Route number limited in FIB
 - For instance, Trident2 is limited to around 200 000 FIB routes
- The same limitation can be seen traditional vendor device that use Trident2 (FIB, TCAM ACL rules, ECMP on VXLAN)

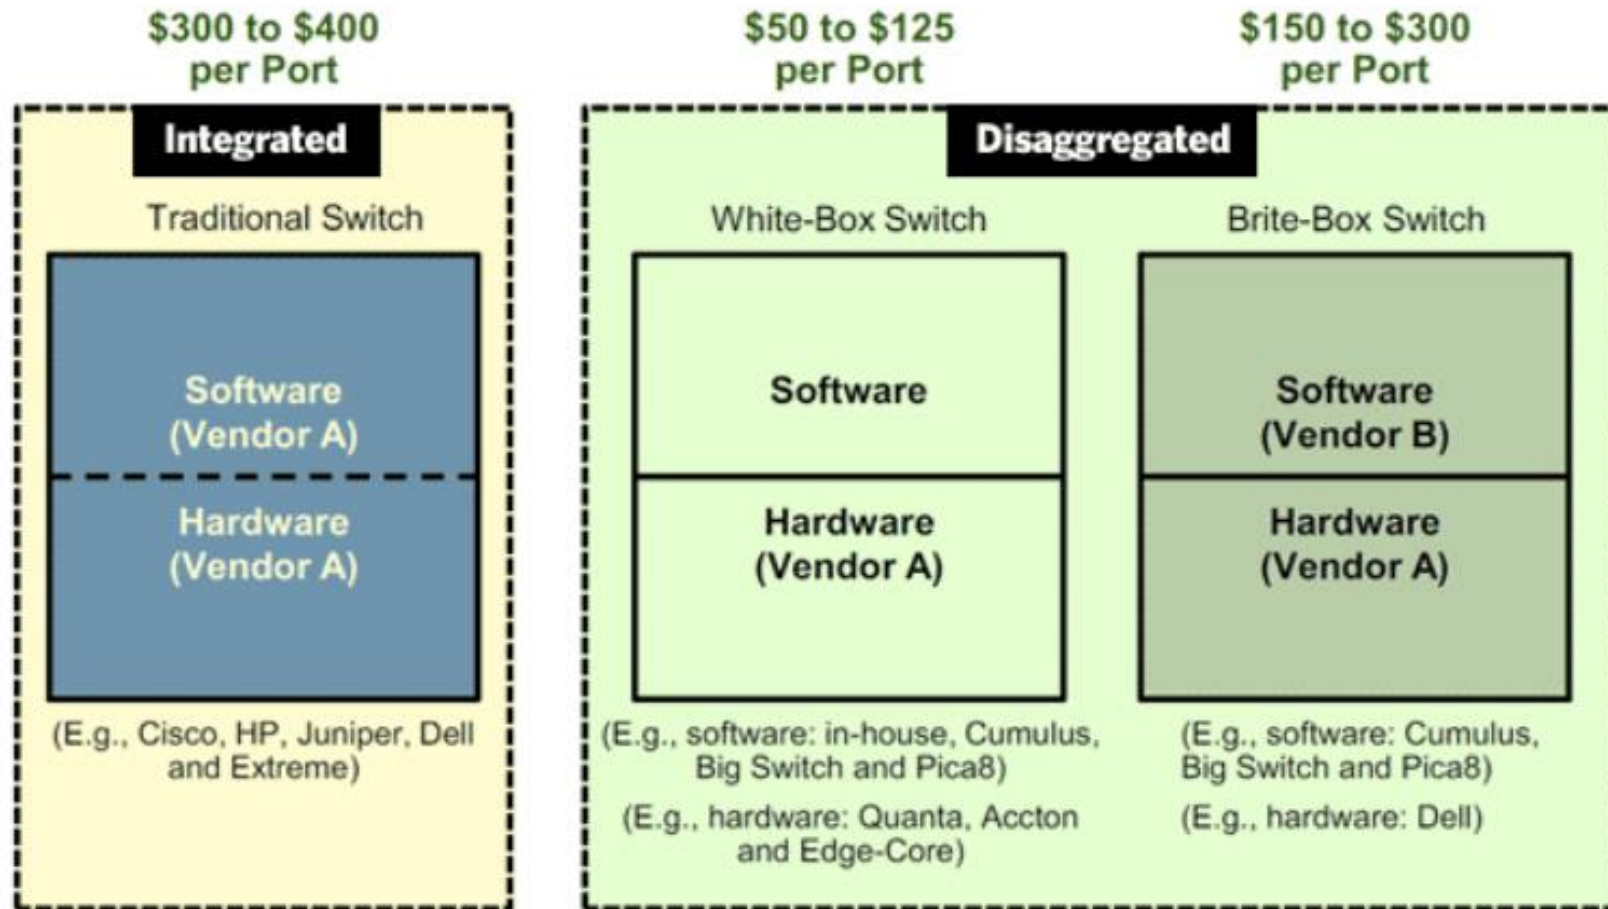
```
S4048-1#show forwarding profile limit
Configured profile : l2-profile-three
Forwarding profile : l2-profile-three(Active in hardware)

-----
|                               Forwarding Profile Table Size                               |
|-----|-----|-----|-----|-----|-----|
| Profile Name | MAC-Table | Host-Table | Prefix-Table | | | |
|---|---|---|---|---|---|---|
|               | MAC ADDR | IPV4 UC | IPV6 UC | IPV4 | IPV6 | IPV6 |
|               |           | (IPV4)  | (IPV6)  | (64) | (64) | (>=65)|
|-----|-----|-----|-----|-----|-----|
| l2-profile-one   | 288k | 16k | 8k | 16k | 8k | 0k |
| l2-profile-two   | 224k | 80k | 40k | 16k | 8k | 0k |
| l2-profile-three | 160k | 144k | 72k | 16k | 8k | 0k |
| l3-profile       | 96k  | 208k | 104k | 16k | 8k | 0k |
| l3-128bit-profile | 96k  | 208k | 104k | 8k  | 4k  | 2k  |
| lpm-profile      | 32k  | 16k  | 8k  | 256k | 128k | 0k  |
| lpm-128bit-profile | 32k  | 16k  | 8k  | 128k | 16k  | 16k |
-----
S4048-1#
```



Example cost analysis gathered on the Web

Traditional, White-Box and Brite-Box switching models



Source: Gartner (December 2014)

Our approach

- Assess white-box use in NREN context → **WB deployment**
 - Explore NREN-relevant use cases: Data-centre, IX router, CPE, P router, ...
 - With the objective to deploy them in production
- White box programmability → **Data Plane Programming**
 - New usage: Monitoring, Security
 - Router for Academic, Research and Education (RARE)
- Concentrate the effort over the next 2 years

Lessons learnt

- Management Plane (Operation, automation, security), documentation, maintenance model are mandatory
- White Box adoption/uptake strategy is a key point
- Do not expect better than you have
 - Same will be a great result
 - New usage/feature will be a breakthrough



Lesson learnt

- White box will not replace our Juniper or Cisco boxes in a first stage
 - Instead move specific services on white box: GIX, LHCONE, ...
- A long term approach is necessary
 - Linux does not replace Solaris in one year
 - The landscape is changing continuously
- Not a reversible path?
 - For White Box
 - For data plane programming

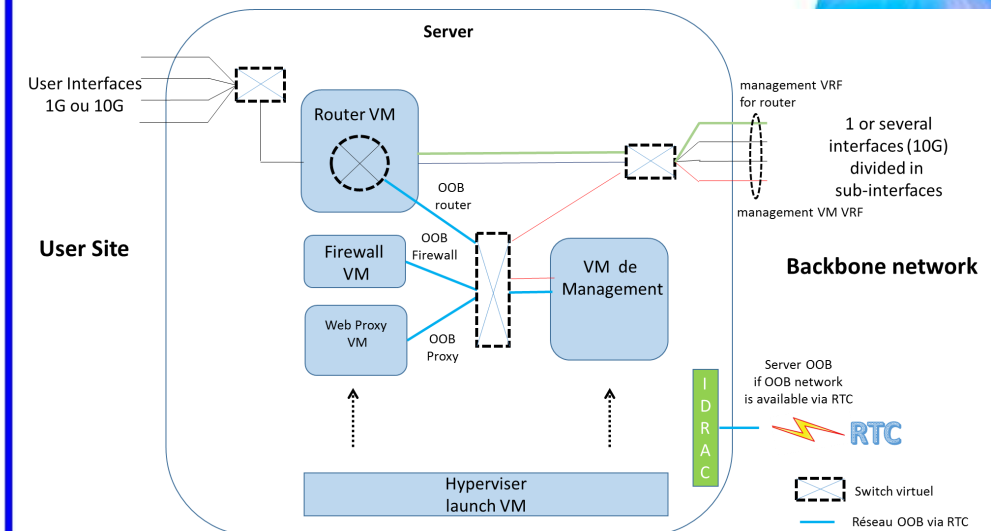


WB deployment: Scope of our work

- Switch with Linux on which you can download your NOS
- Switch with ONIE
 - DELL, Edge Core, ...
- A X86 server than run several NOS

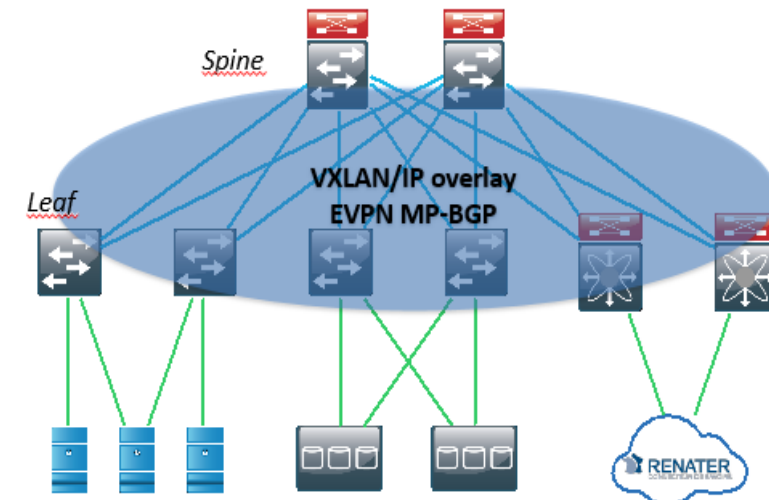

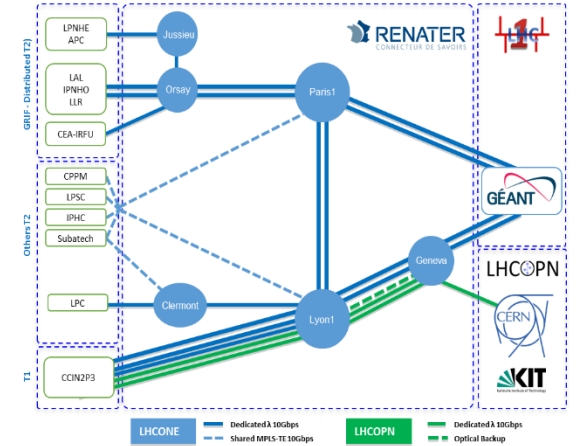


a X86 server with a NOS



R&E use cases

- NREN backbone (PE and P routers)
- Regional network
- Campus network
- Science project
- Global Internet eXchange (GIX)
- Cloud Fabric
 - Monitoring, Telemetry
 - Security



Use cases handled by WP6T1 WB deployment

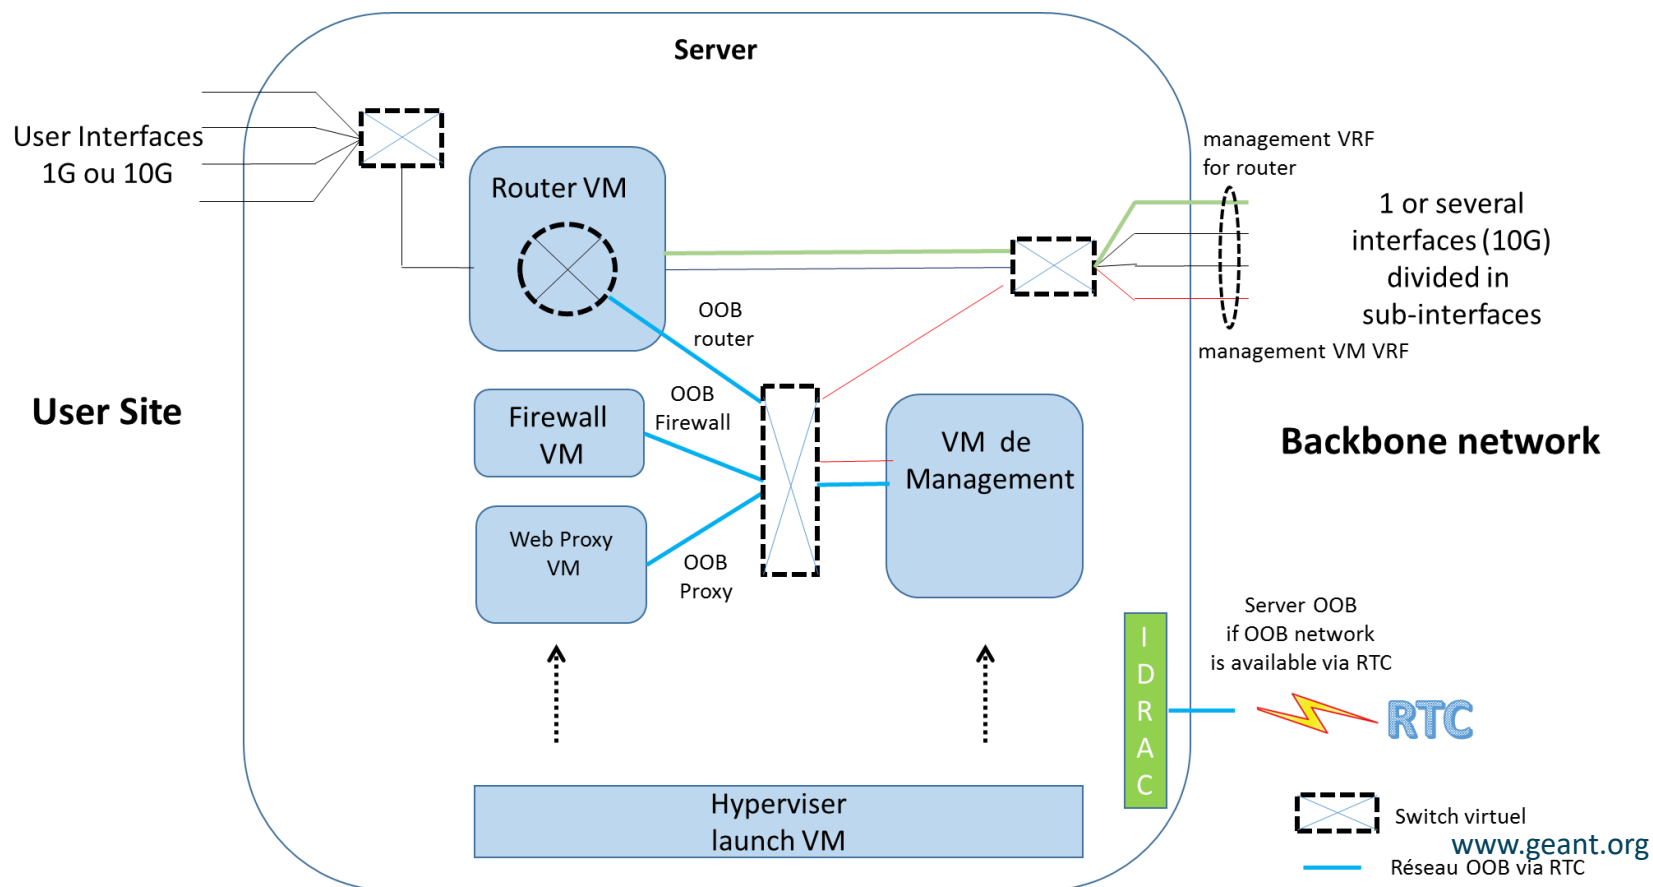
- Internet eXchange point
 - SFINX, RENATER Internet eXchange point (Paris/2 locations)
- CPE
 - High school in Normandy Region (France) - RENATER
 - FUnet
- Data Center
 - Normandy Region (France) – RENATER
 - GRnet
- P router - PSNC



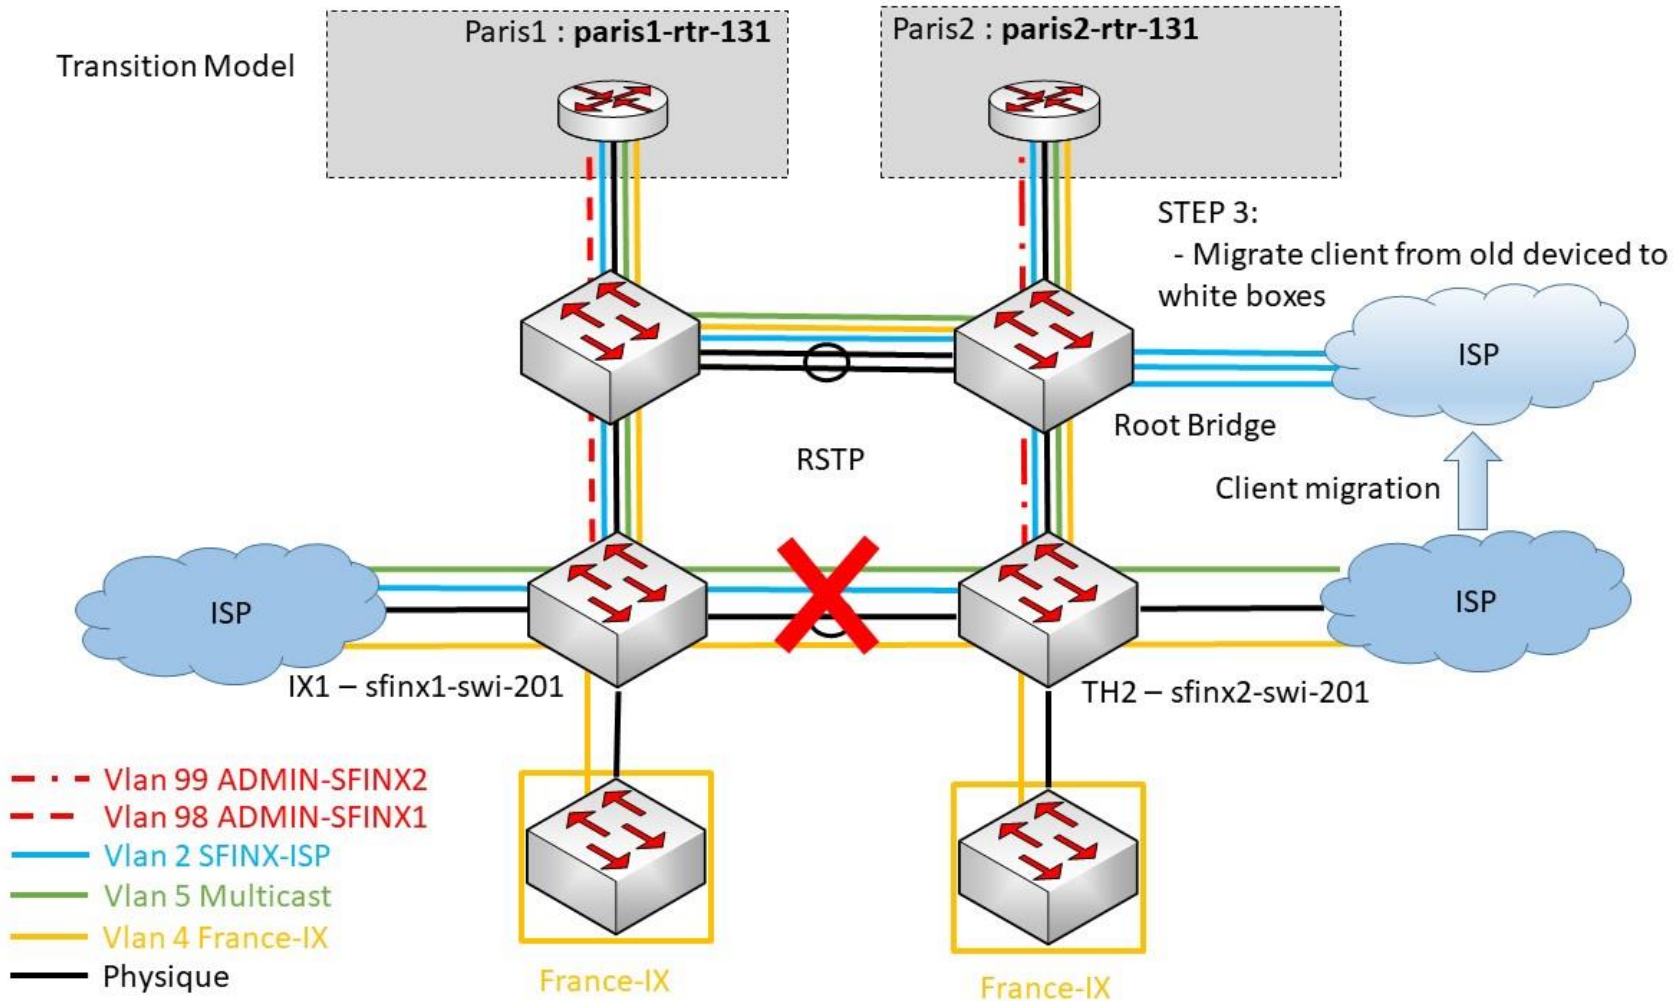
CPE project in France

- A X86 server with a NOS can be considered as White Box
 - Limited number of use-cases due to performance

Example of CPE design



Internet eXchange Point



Data-center Normandy project



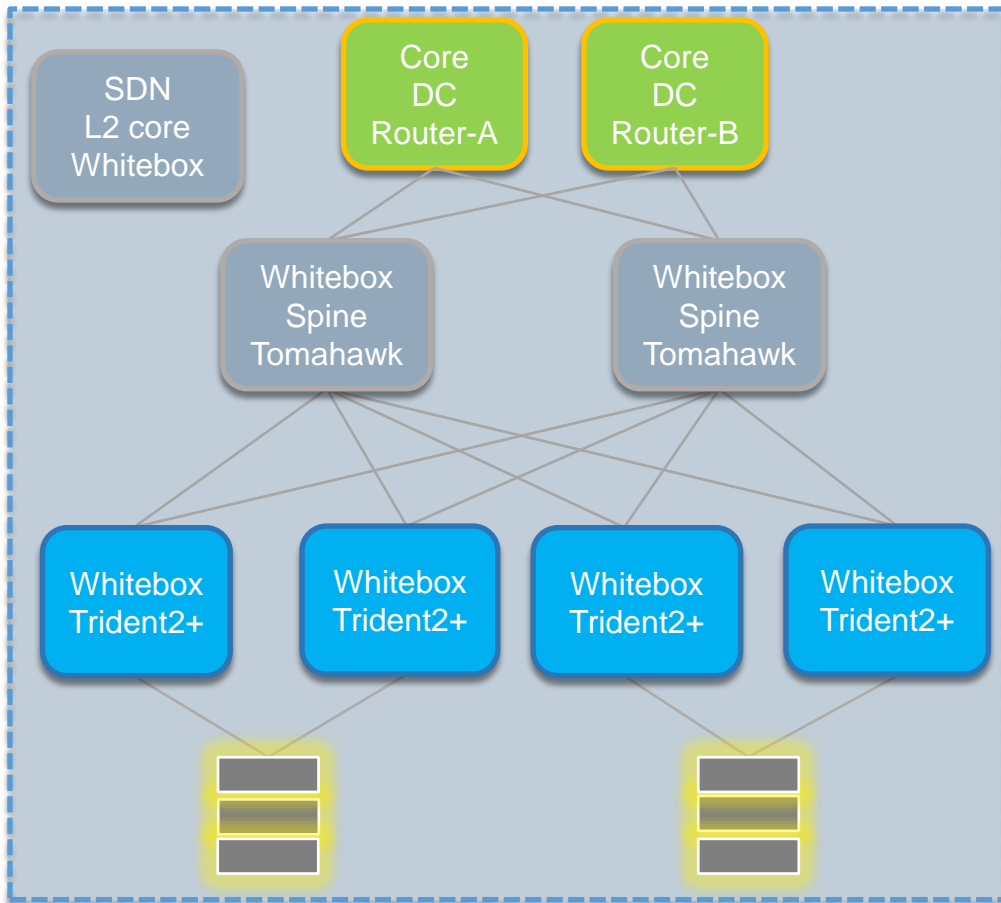
WB : Datacenter switching

Goals

- Cost reduction
 - this is the main objective : the initial vendor solution cost 200k€ for around 1500 1Gbit ports
- Interoperability
 - separate hardware from software is interesting when upgrading or extending part of the infrastructure
 - avoid vendor lock situation especially when it comes to big infrastructure
- Automation and orchestration
 - no more CLI configuration
 - network configuration will be driven by application services

GRNET data center projet - IP Fabric topology

GRNET whiteboxing testbed



- Spine & Leaf topology
- SPINE: Tomahawk ASIC
 - ▣ nx40G uplink to GRNET DC routers
 - ▣ nx100G uplink to GRNET L2 core SDN
- LEAF: Trident2+ or Maverick
 - ▣ 2x40G uplink
- Server:
 - ▣ 2x1/10G UTP
 - **Multihoming: In pairs of racks**
 - ▣ LACP or Active-Backup

Condition to adopt white-box model?

- Identify a first use case
 - **The same services** that NRENs provide to their end-users?
 - OR a **new service**?
- Use Case validation
 - Feature?
 - Performance?
 - Maturity and Reliability?
 - What is missing in an open network operating system before going into production?
 - Manageability, Open network operating system security, Documentation, Maintenance model
 - ...
- **Transition model** that could be put in place in order to go in production?
- What is the **total cost ownership** for such technology?



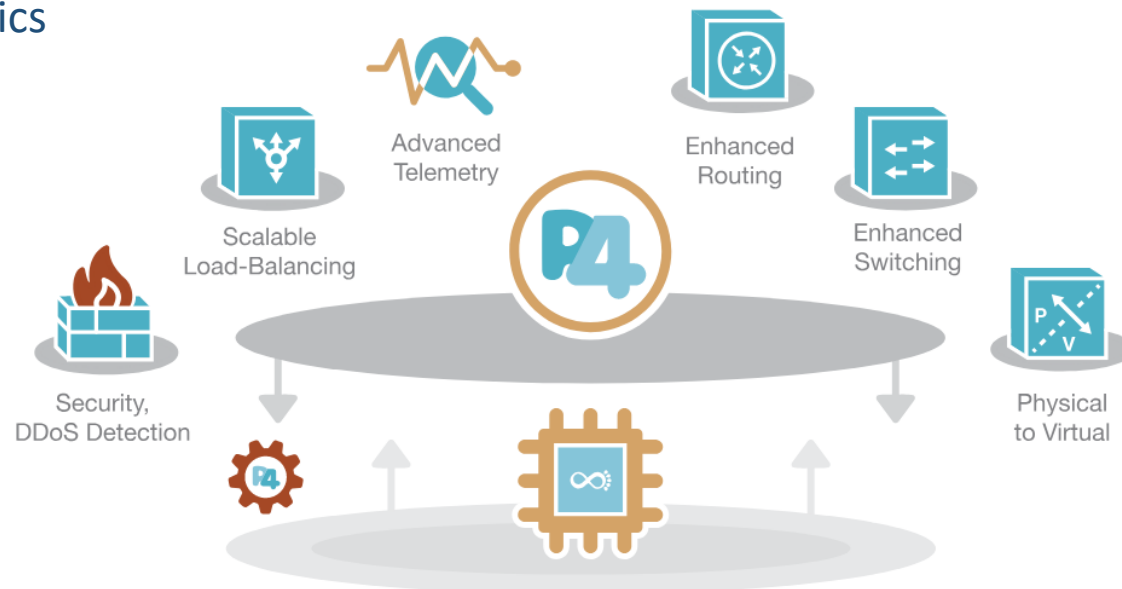
WB deployment strategy

- Identify appropriate use cases that will be put in production
 - White Box adoption/Uptake strategy
- Technical validation
 - Methodology
- Business model
 - Licence model and TCO
- Deployment plan
- Qualification for production
- Production



White Box Programming – Data Plane Programming

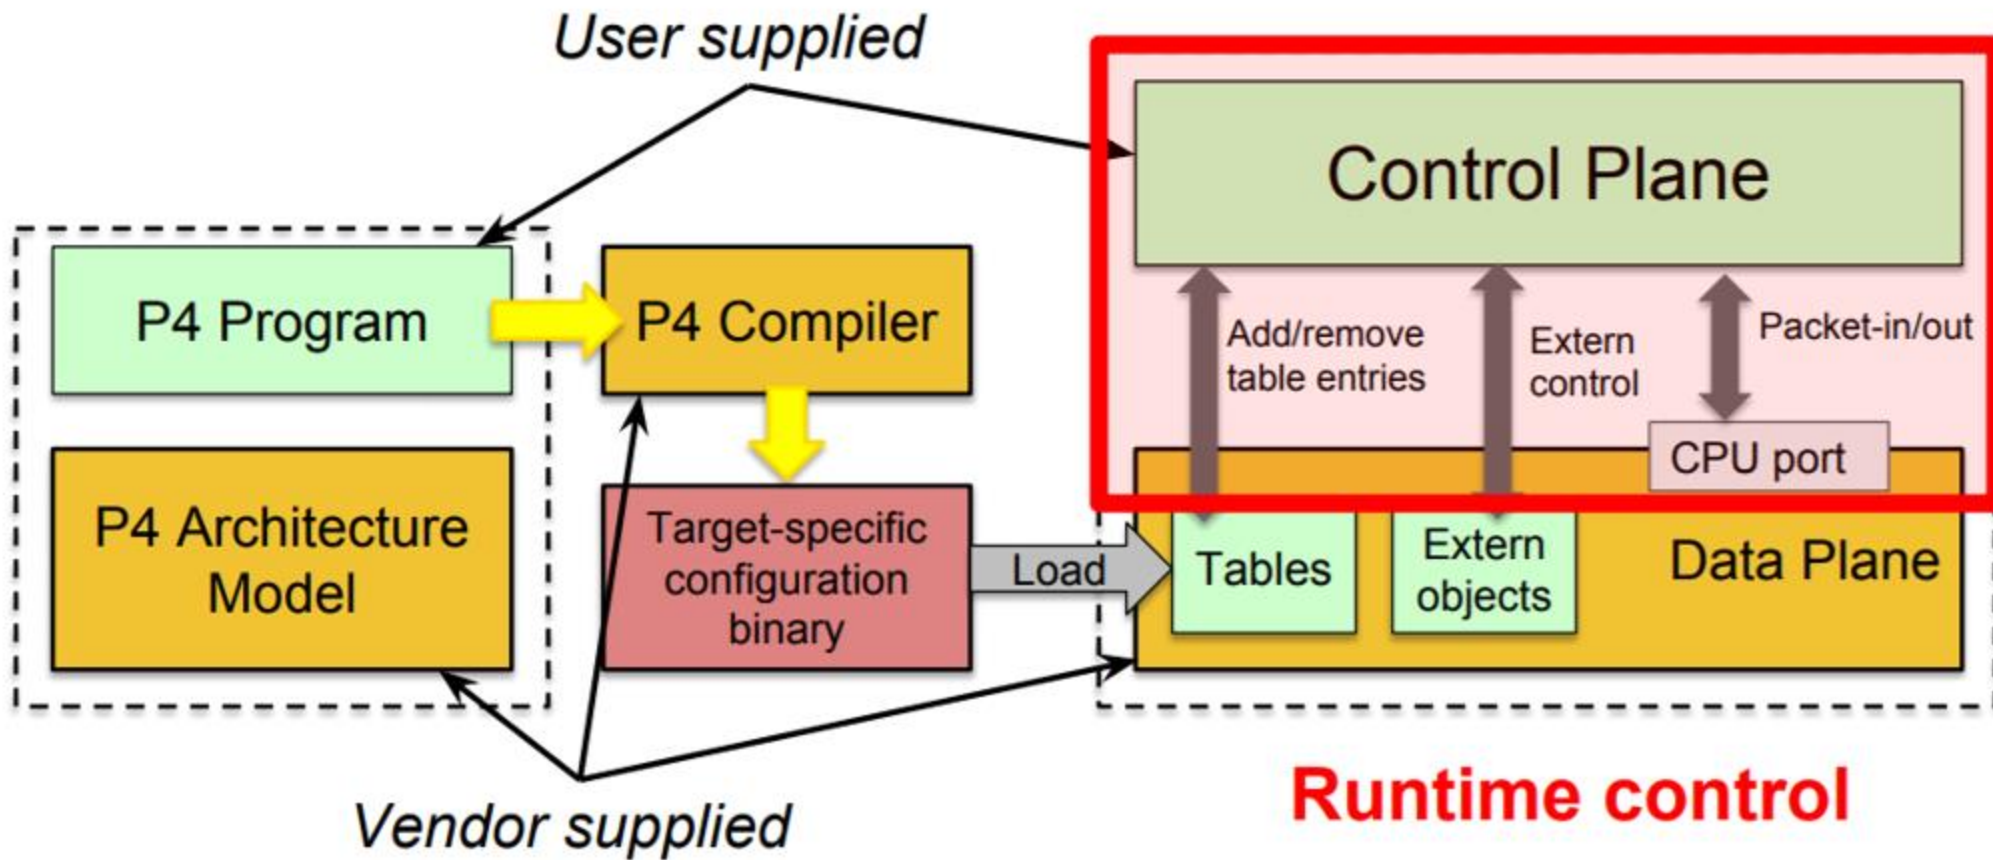
- Data Plane Programmable – P4 language
 - Based on PISA architecture [Protocol-Independent Switch Architecture], FPGA, Open VSwitch Tofino chipset barefoot network (<https://www.barefootnetworks.com>)
- Application
 - Monitoring: Advanced Network Monitoring/Telemetry
 - Security: In-Network DDoSDetection
 - Performance: Layer 4 Load Balancer
 - Analytics



World of P4 Advanced Apps

<https://p4.org/> www.geant.org

P4 Workflow



Data Plane Programming (DPP) and P4 language

- **DPP advanced feature**
 - Coordinator: Mauro Campanella
 - Objectives: Explore new feature provided by programmable white box
 - Monitoring
 - Security (DDOS)
- **Router for Academic, Research and Education (RARE)**
 - Coordinator: Frédéric Loui
 - Objectives: Demonstrate the feasibility of a router for academic community
 - Data plane: DPP
 - Interaction control plane and Data plane

Summary

- **The routing landscape is changing fast**
 - Especially in data-center ...
 - Business model – Cost - New NRENs (African countries,)
- **WB deployment**
 - Test and put in production WB in NREN use cases context
- **WB programmability**
 - Investigate new usage: first Monitoring and security
 - Open source router for academic and research
 - Universities are very interested for its own research in developing P4 usage
- **Build a community around white-boxing**

Building a community around white boxing

- **Workshop and dissemination:**

- White-Boxing: 4th April 2019, SUNET, Stockholm, (before SIG-NGN)
- P4 Programming: 20th June TNC19, BoF - exploring R&E community use cases

- **Feedback from our community**

- **NREN**
- **Education and Research community**



Thank you

Any questions?

www.geant.org

