Funet CPE

White Box Edge Router

Jani Myyry (CSC/Funet)
GN4-3 WP6 T1

GÉANT Infoshares: White Box in NREN context

8 December 2021

www.geant.org

Agenda

- Introduction and background
- Requirements
- Testbed implementations
- Feature validation
- Performance tests
- Conclusion

Background for “Funet CPE - White Box Edge Router”

- Funet has provided “Campus Edge Router” service since 2012
  - Originally managed edge routers to implement redundant NREN uplinks
  - But some were implemented as an integral part of campus networks
  - Based on Juniper MX80/MX104 routers
  - Same size fits for all?
    - Big university (8000+ staff, 30000+ students)
    - Other 25 customers…
    - Small office of a research infrastructure (10 staff)
- “Funet Kampus” CNaaS evolved service model in 2019
  - Now mostly Juniper MX204/MX10003 based
    - Some Huawei switches for very small environments
  - Configuration automation with Ansible (primarily for Juniper)
  - Juniper MX204 is too future-proof and expensive?
- Currently 27 customers with ~ 100 routers/switches deployed

Common requirements for White Box CPE device

- Scalable solution for smaller campuses and remote offices
  - Redundant and non-redundant edge router(s)
  - 10 Gbps NREN uplink
  - 1/10 Gbps user access
  - 100 Gbps upgrade capability (w/ changing pluggables) preferred
- 5 year support for the NOS and hardware
  - Minimum typical service contract offered
  - NBD or RTF support (redundant setup)
  - Support from a single channel (NOS provider)
- TCO competitive
  - Lower than Juniper MX204
  - Close to low-end switch alternatives
- Enough active users within the R&E community

Other requirements for White Box CPE device

- Critical features
  - Dynamic routing protocols (BGP, OSPFv2/OSPFv3)
  - Stateless traffic and route filtering
  - Management and monitoring cababilities
- Performance
  - Line-rate packet forwarding for 10 Gbps (and 100 Gbps)
  - Hardware with deep buffers (to support interface speed mismatch properly)
  - RIB/FIB scalability for few 1000s routes
  - Support for directly connected research applications (“Science DMZ”)
- Redundant and how-swappable components
  - Power supplies, fans
- High-end complex campus environments out-of-scope
  - MPLS-based L3VPNs, L2VPNs, VPLS, EVPN, …

CPE deployment scenarios

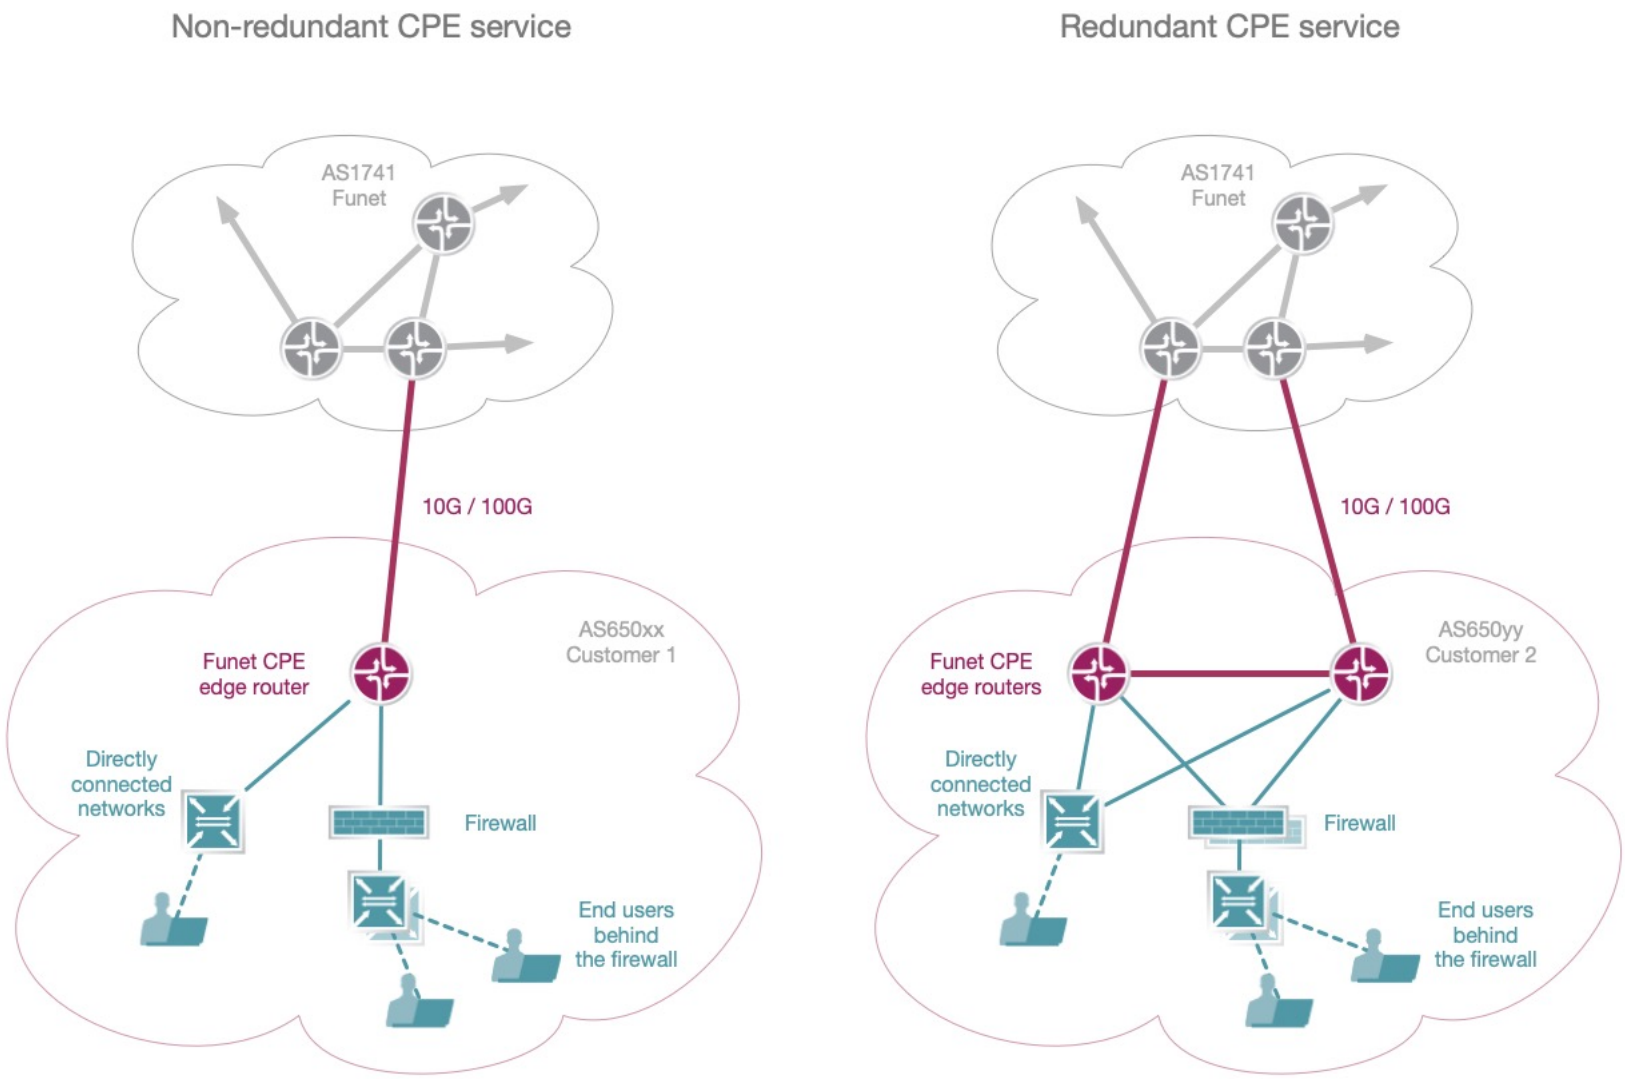

Virtual NOS testbed for White Box CPE

- Cumulus Linux virtual NOS was chosen to test control and management plane features (on VMware)
  - Features complete enough for specified use cases
  - Test environment stable for months
- Hardware selection for CPE not as straightforward as would have been for data centers
  - Optimally having similar features like in traditional routers
    - Deep buffers (Broadcom Jericho/Qumran)
    - Small footprint and support for extended environmental conditions
  - But Cumulus did not support Broadcom Jericho/Qumran
    - And on later date Broadcom at all?

Hardware NOS testbed for White Box CPE

- EdgeCore AS7315-27X became available late 2020
  - Broadcom Qumran-AX based hardware
  - Implements Telecom Infra Project's "Disaggregated Cell Site Gateway" definitions
  - Very good design for smaller sites
    - Native optic support from 1 to 100 GbE
    - Operate in extended temperature range
    - Front-side access only
- Originally tested Cumulus Linux NOS did not support Edgecore hardware
  - All NOS options were from niche players and apparently focused for mobile access
  - ADVA Ensemble Activator was chosen as NOS
  - Hardware (two routers) were sourced from ADVA as well

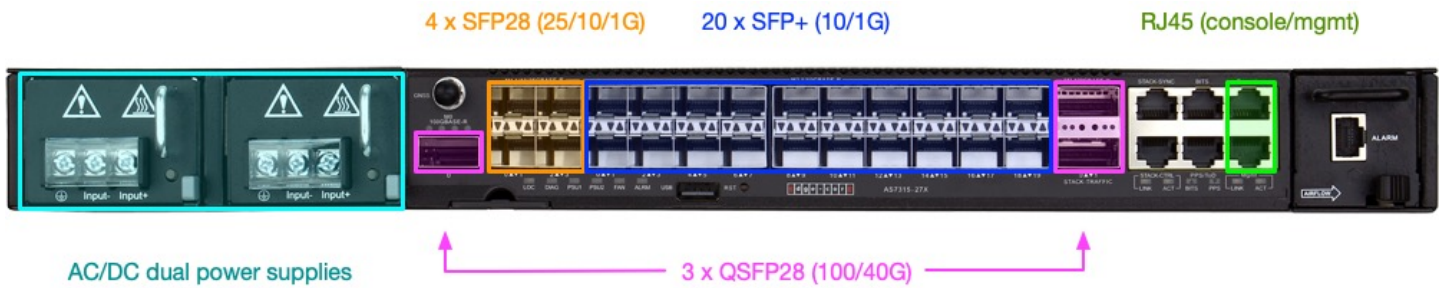

Preparing NOS hardware testbed for White Box CPE

- EdgeCore AS7315-27X supports both AC and DC power supplies
  - AC used everywhere but due to schedules, DC was the only option available (required external rectifiers)
- Hardware arrived without NOS installed
  - NOS installation via network-based bootstrap environment
- Internal storage issues with one of routers
  - Required vendor support to restore into operational mode
  - But still did not behave 100% correctly during the tests

Feature validation (expected results)

- Dynamic routing protocols
  - BGP (*), OSPFv2/OPPFv3
- Access networks
  - SLAAC, VRRP (*)
- Management
  - SSH, OOB
- Monitoring
  - SNMP, syslog
- Performance (10 / 100 Gbps mixed environment)
  - Line-rate forwarding and traffic filtering (10 Gbps)

(*) issues in next slide

Feature validation (issues)

- Dynamic routing protocols
  - BGP (IPv6 route filtering limitations)
- Access networks
  - VRRP (issues w/ MD5 authentication)
  - DHCPv4/DHCPv6 relay (limited configurability and issues)
- Traffic filtering
  - No egress filtering support
  - No L4 port support with IPv6
  - No control plane filtering support
- L2 bridging
  - No real support available

Performance tests

- Performance tests were performed with EXFO FTB-890NGE tester
  - Up to 10 Gbps supported by the tester
  - Tester directly connected to both routers
- Test method 1: RFC 2544
  - To evaluate IPv4 and IPv6 forwarding performance
  - Included throughput, back-to-back frames, frame loss and latency tests
  - Frame size distribution: 70, 128, 256, 512, 1024, 1280 and 1518 bytes
- Test method 2: traffic generation with 7 parallel streams
  - To evalute IPv4 ingress traffic filtering performance
  - Stream bandwidth distribution: 40%, 20%, 10%, 10%, 10%, 5% and 5%
  - Streams number 2 (20%) and number 3 (10%) were filtered
- All tests were successful and gave expected results
  - Line-rate forwarding and traffic filtering

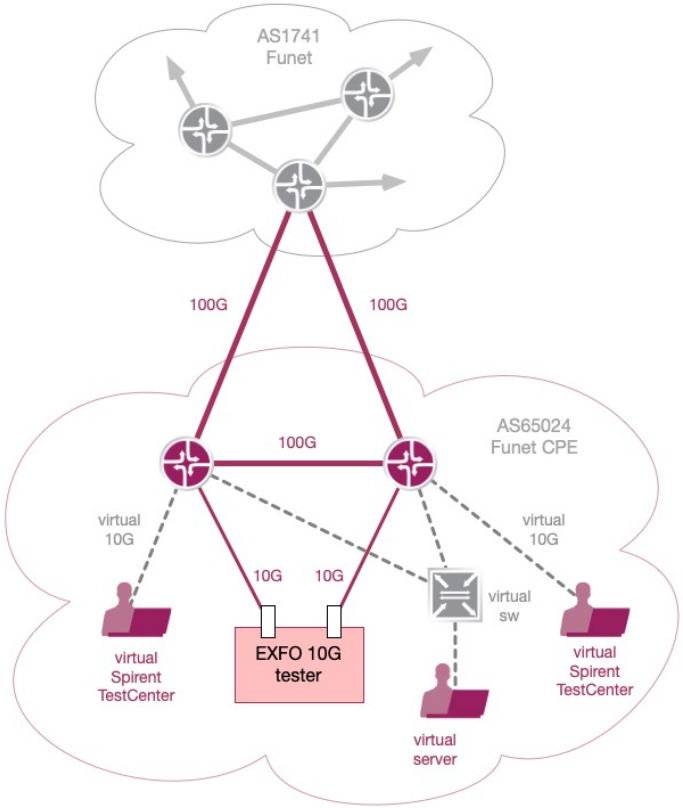

Conclusion

- Edgecore hardware pretty ok for many smaller use cases
  - Open hardware specifications
  - Supports 1/10/25/40/100 GbE from the same device
  - Designed to operate in non-datacenter conditions
  - Limited scalability (300 Gbps without option to over-subscribe)
- Tested NOS recently introduced for the hardware
  - Some of issues found were clearly bugs in the implementation
  - NOS primary focus: mobile or data center networks?
- White Box with both an open NOS and an open hardware?
  - Microsoft SONiC (states support for Edgecore AS7315-27X)?
  - Support would be complicated without an NOS vendor?

GÉANT
Networks • Services • People

Thank you

Any questions?

www.geant.org

The research leading to these results has received funding from the European Union's Horizon 2020 research and innovation programme under Grant Agreement No. 856726 (GN4-3).