POLITÉCNICO DE LISBOA

IpLNet www.net.ipl.pt

Mikrotik at Lisbon Polytechnics
MUM Lisbon 2019

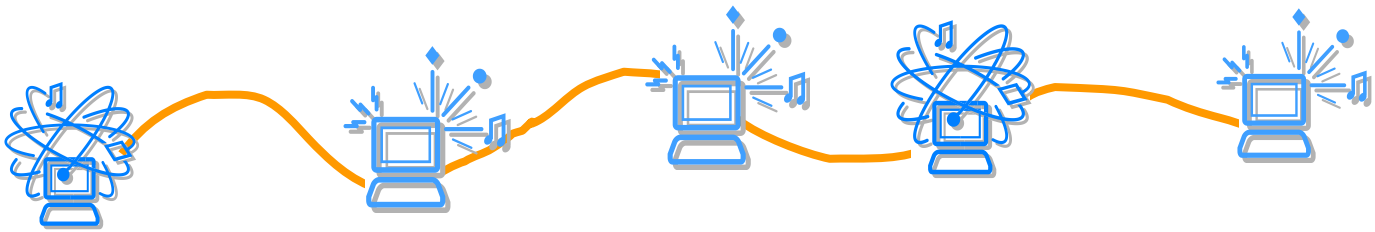

Pedro Ribeiro

Instituto Politécnico de Lisboa (IPL)

Instituto Superior de Engenharia de Lisboa (ISEL)

pribeiro@net.ipl.pt

Who am I?

- Professor at ISEL since 1997
  - Also informal IT coordinator at IPL
- After 2013 leading the IT (DSIC) at IPL
- Skills & experience
  - Linux *kernel*, systems and applications
    - Specialist title obtained with the project “Firewall system based on Netfilter”
  - Cisco routing, switching, WiFi and appliances (ex. WLC)
  - Alcatel/Nokia MPLS Service Routers
  - Mikrotik routing, switching e WiFi
    - MTCNA, MTCRE, MTCWE, ACTR
- Hobbies: ham radio, electronics, microprocessors

IPL Context

- Community of about ~16500 users
- 8 Schools – Arts, Engineering & Health
- 5 Campus – Lisbon metropolitan area

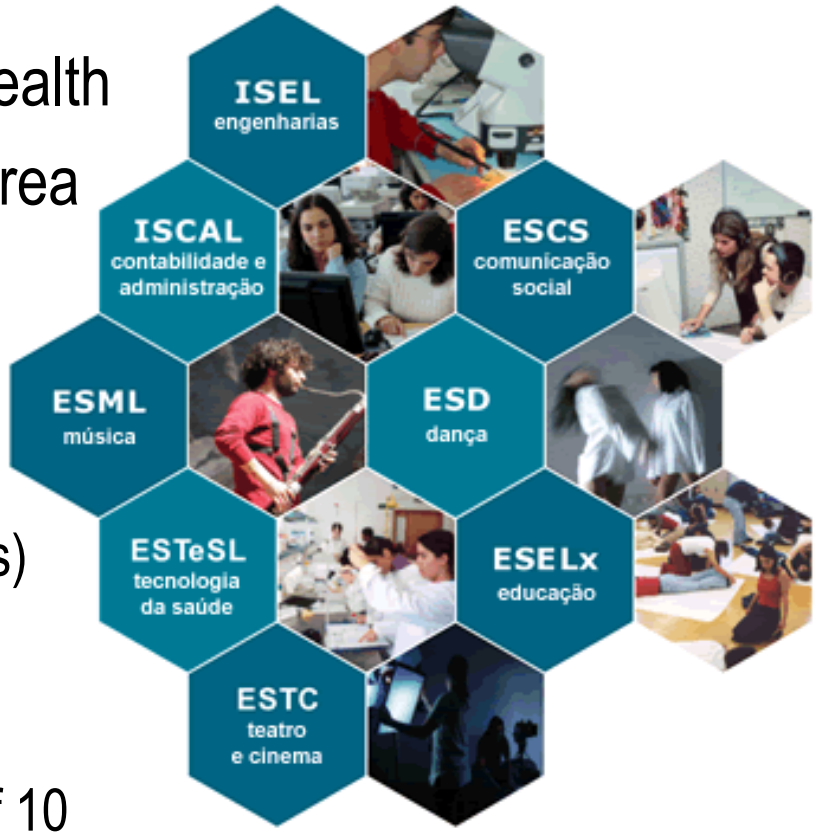

- Networking
  - 8500 access ports
  - 200 eduroam APs (peaks of 3k users)
  - 10+10Gbit/s Internet uplink
  - 350 virtualized servers
  - DSIC (common IPL IT) with a staff of 10

Systems and Infrastructure – DC/COM

- At the ISEL campus
  - Servers, storage, routing e switching inside 16 racks

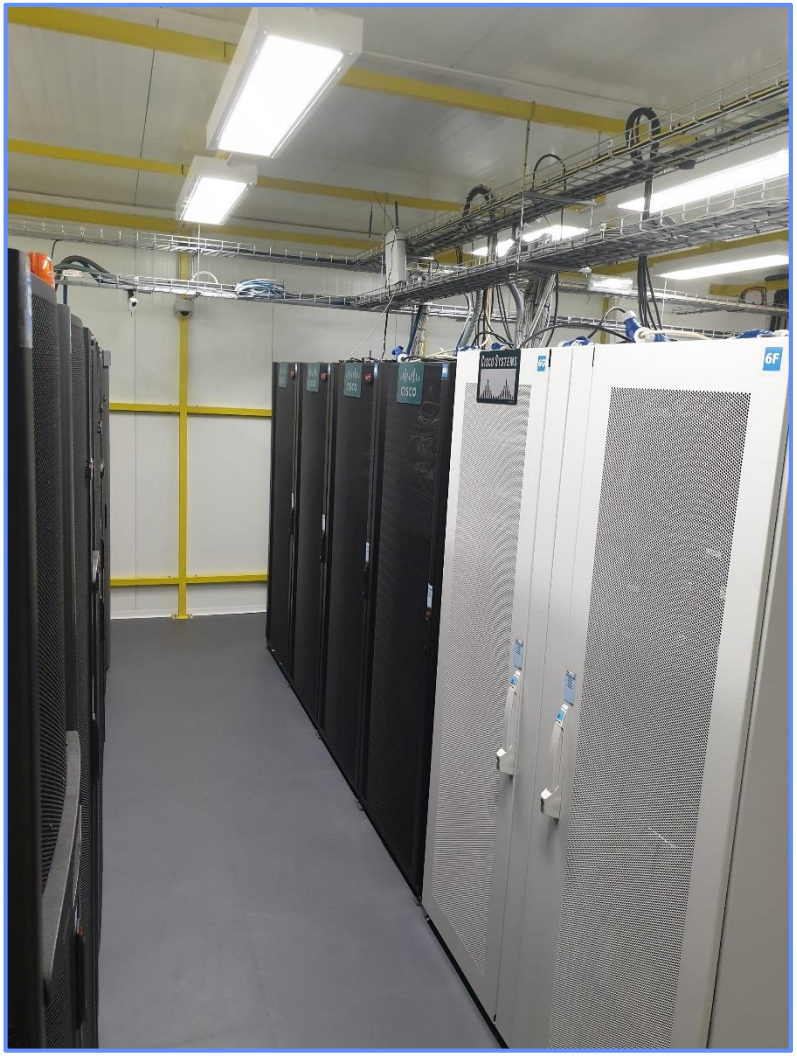

Systems and Infrastructure – DC/COB

- At Benfica campus
  - 3 racks at the room T7 (temporary)
    - backups for the connectivity services
    - Two servers with essential network services
  - Main building for the IT team
- Data/comm center
  - Under project
  - Space for:
    - 10 server racks
    - 6 net racks
    - 2 UPS racks

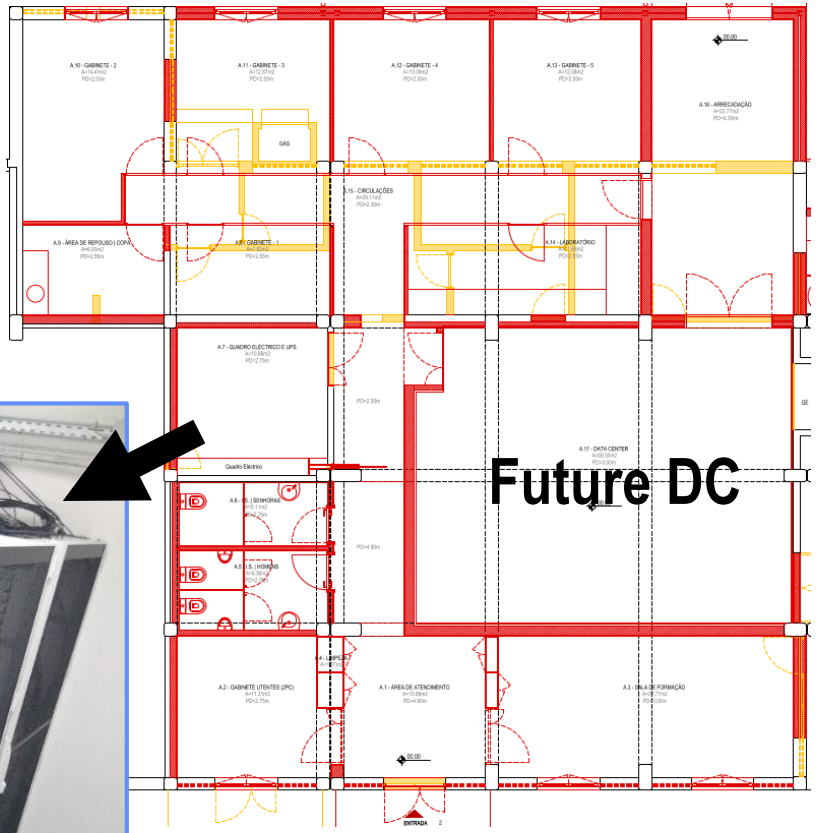

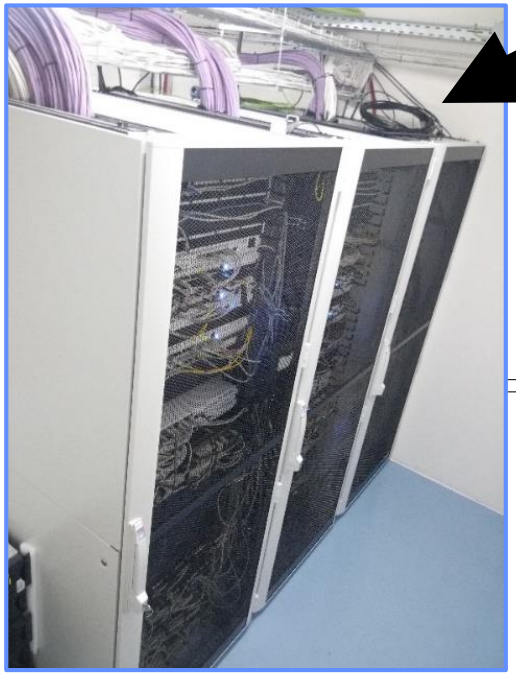

Systems and Infrastructure – DC/SP & other sites

- Presidency building
  - 5 racks with servers supporting the administrative and academic applications

- At each school/site
  - Local IP/IPv6 distribution
  - Switching and aggregation
  - Local servers
  - Servers with offsite backups from elsewere

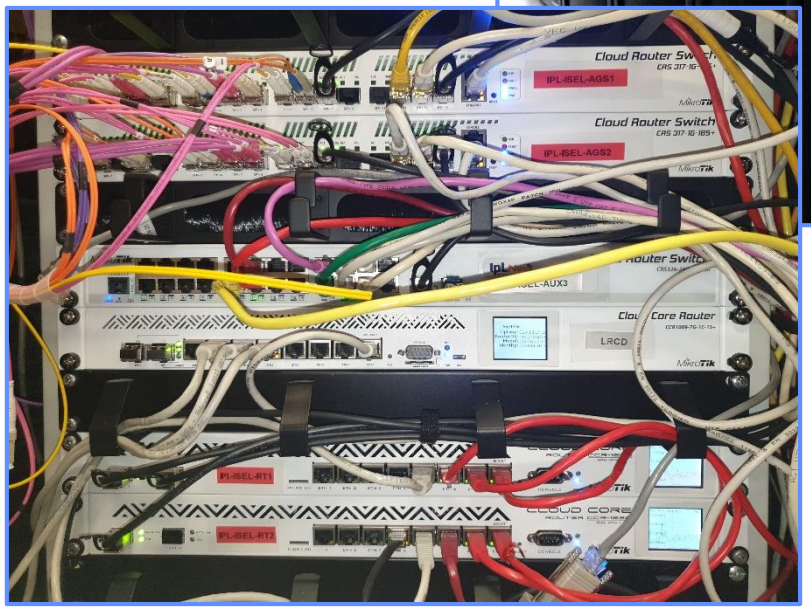

Systems and Infrastructure – Connectivity

- Internet via FCCN/RCTS (NREN) at 10Gbit/s
  - Uplinks at COM & COB sites

- Metro MPLS network interconnects sites at 10Gbit/s+
  - Several transport rings with common legs (~100km)
  - At least 2 uplinks per drop/access site
  - Local IP/IPv6 services routed onsite
    - Interconnection to the core routers with VPLS/PW (or direct 10G Eth when available)
  - Services common to all sites distributed as L2 (switching)
    - VPLSs interconnects routers at the central sites with the remote LANs
      - Printing, network management, eduroam, access control & attendance, etc.

Dark fiber connecting sites

The network before the evolution to Mikrotik

- Bottlenecks at routers with over 10~20 years
  - Some only had 100Mbit/s ports
  - Others with a couple of 1Gbit/s ports that needed to be shared by N VLANs
  - No software updates or bugfixes
  - Outdated functionalities
  - Limited by excessive use of hardware acceleration chips
- Aggregation and access switching with EOL equipment
  - Security and functional bugs without solution
  - Several limitations dealing with IPv6, multicast and other recent protocols
  - Only a few had PoE, none with the PoE+ power needed by recent devices
- WiFi APs with about 16 years (54Mbit/s were the top at the time)
  - Very slow and with a lot of unsolved bugs (out of support)

Possible evolution

- Several evolution products were proposed by the manufacturer of the installed equipment
  - ▼ All of them lead to strong product dependency (not again!)
    - Features appealing but requiring same manufacturer/product line everywhere
    - Very dependent on central management devices
    - Operation and management with higher abstraction
    - Current products of the manufacturer are highly dependent on cloud & “phone home”
    - Some products require annual licensing (delays at renewals will stop network???)
    - Support will require costlier maintenance contracts
  - ▼ Some devices/functionalities without backup on another chassis
  - ▲ Solution with guarantees from a established manufacturer
  - ▼ First numbers add to around 800keur!!!

Pilot phase

- Several positive references to Mikrotik equipment
  - We had using MT gear before in our homes and some WiFi links
    - Backups links between ESTC/ESCS and ESML/SP

- We acquired 20 CHR licences, 1 CCR1009 and 5 wAP ac
  - ▲ CHR was the solution to our VPNs (and as a WLC)
  - ▲ CCR1009 revealed a good performance and functionalities despite being the little brother of the series
  - ▲ wAPac was able to do all the functionalities needed by eduroam
    - The multicast helper feature solved the IPv6 & Apple problems we had before
    - The dynamic VLAN assignment from RADIUS worked in autonomous and WLC modes
  - ▲ CLI and WEBFIG management were easy explored by the staff

Team training

- Several training sessions done in collaboration with Truenet
  - ISEL/ADEETC provided the lab space
  - Trainers Jorge and Raul (YaTuAprendes)
  - 3 sessions done
  - At the moment our tem has:
    - 3 MTCNA, MTCRE, MTCWE (ACTR)
    - 2 MTCNA, MTCRE
    - 1 MTCNA
- As some staff members are also professors at ISEL
  - We activated the first portuguese Mikrotik Academy
  - One student group already finished the MTCNA (100% approval)
  - Other sessions scheduled for the next months

Products in use

- Acquisitions in several phases (and suppliers)
  - CCRs 1009, 1036 e 1072 (37)
  - CRSs 112, 317, 326, 328 (60)
  - RB 1100AHx4, 260GS e 750Gr3
  - SXT G5HnD e RB911G
  - wAPac (38)
  - SFPs compatible 1000BaseSX/LX e BX* (32)
  - SFP+ compatible 10GBaseSR/LR e BR* (136)
  - SFP+ patch of 1 & 3m (62)
  - SFP/SFP+ RJ45 (39)

* Standards using just one fibre with two wavelengths as implemented by MT S+2332LC10D & S-3553LC20D

CHR for virtualized networking services

- At the moment 17 in production (a few more used for training/testing)
  - Near unnoticeable resource consumption (RAM, CPU, HDD)
- VPNs
  - For systems under contact telemaintenance (SAP, SGA, VoIP, PRINT, SSO …)
  - Intranet accesses (lecturer, students, staff)
  - Liberalized/privileged Internet access (jumping the firewall and getting a public IP)
  - Secure access to administrative applications (+SIBS ATM backend)
  - RADIUS checks user access and return profile attributes to apply (usage limits)
- Routers “intra-vm-cluster”
  - Routing between virtualized systems inside the VM cluster

▲ Reduction of maintenance tasks and time, of the hardware in use, the electrical consumption and end of cipher/protocol limits.

▲ Faster communications when connecting and transferring data

Integration of CCRs in the CORE and DC

- All functions using pairs of equipment's
- Heavier tasks done by CCR1072, the other by CCR1036
- Main equipment mirrored in our sites COB (Benfica) & COM (Marvila)
  - Active/backup roles scattered for each service, none of them is idle
- VRRP used for redundancy of gateway in the terminal networks
- Routes propagated and optimized using OSPF/OSPFv3
- Notes/tips:
  - Use IPv6 to simplified setup of VRRP (even IPv6 isn't used elsewhere)
  - A bridge interface without ports is like a "loopback" needed for some protocols
  - *mac-telnet* gives us out of the box (and recovery) access to devices
  - Prepare scripts to do the common configuration, certificates loading and schedule the regular configuration and inventory backup
  - Disable SFP firewall for optimal performance and stability

CCR at the core and border

- Symmetric topology
  - Simplified and resilient
- Links via:
  - MPLS
  - CRS317 “CC” (no STP)
- BGP/BFD with FCCN/RCTS
  - No problems identified
  - RCTS side using Juniper/Cisco
- OSPF/OSPFv3 as IGP
  - BFD in indirect links (SW/MPLS)

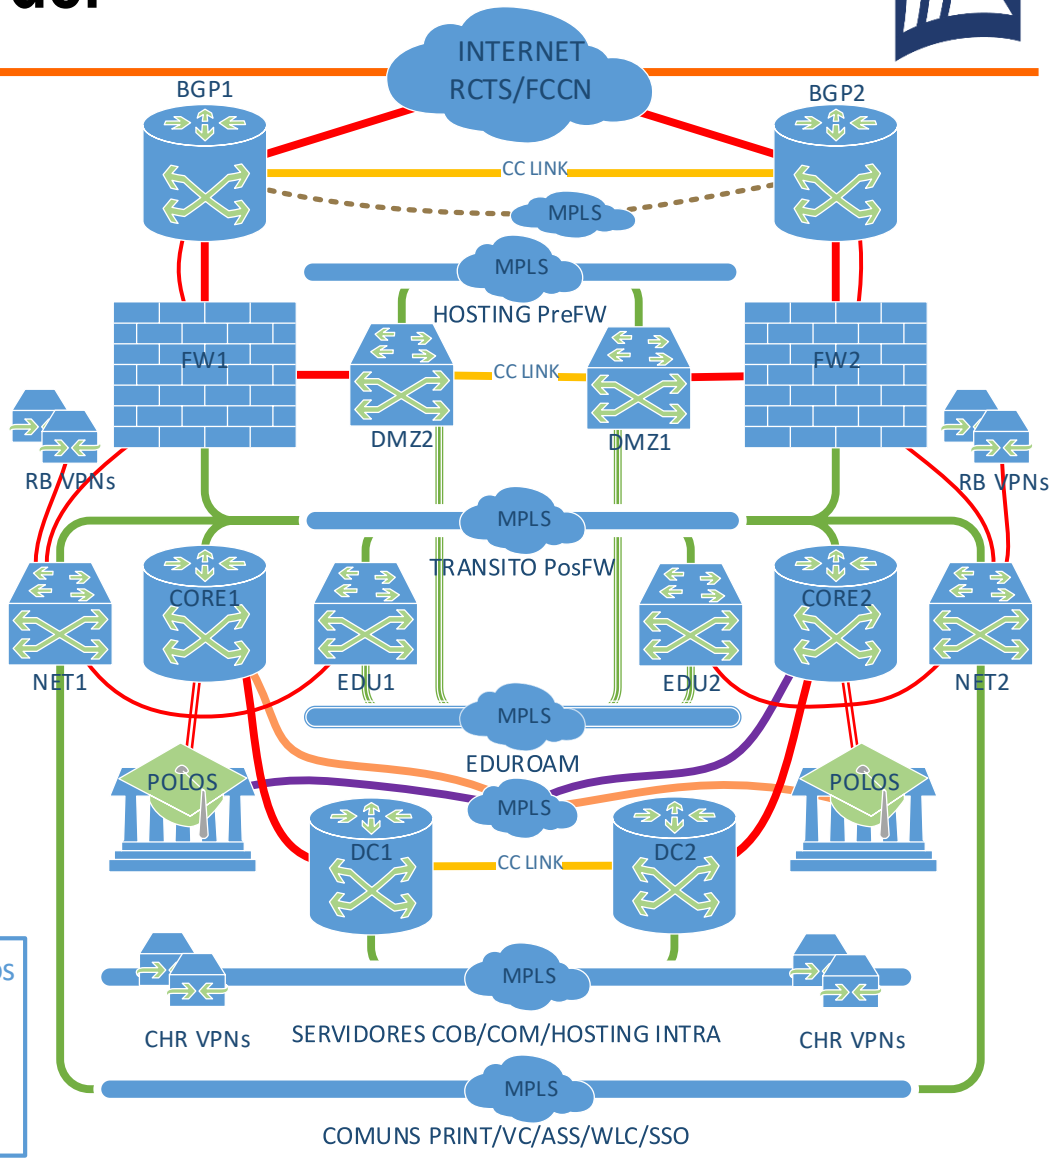

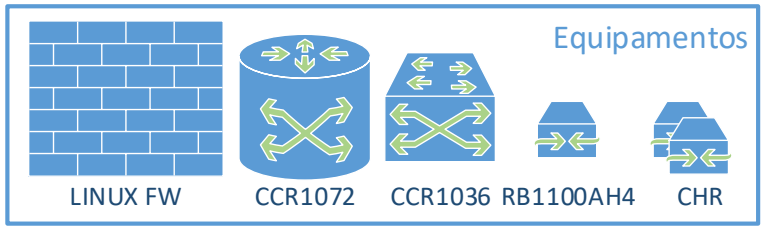

CCR in the campuses and schools

- Using equipment pairs
  - At the moment co-located
- CCR1036, in each:
  - 1 SFP+ port as uplink (main/backup)
  - 2 ether ports do ECMP routing between them
  - 1 SFP+ port serves the more demanding networks
  - Remaining 6 ether ports used for extra local services
  - Active VRRP preference scattered over the two routers

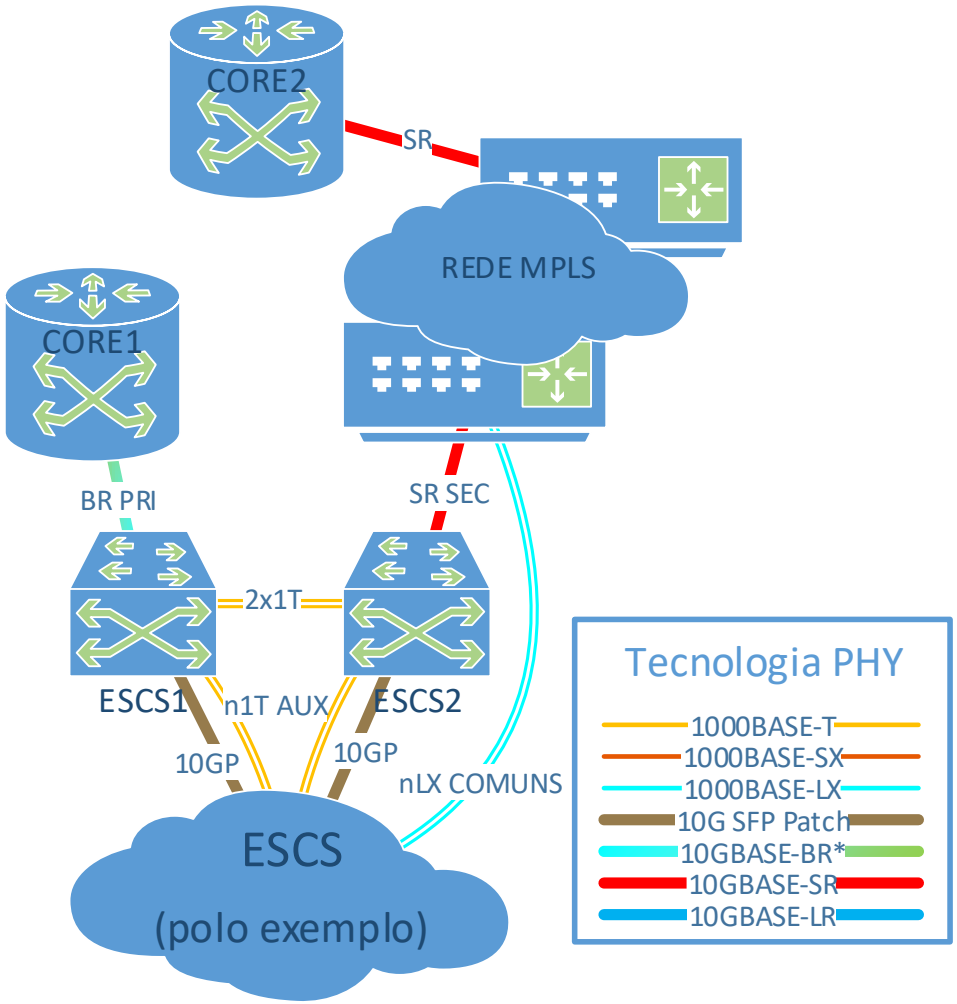

wAPac on “eduroam”

- *bridge* join at L2 the interfaces *wlan1*, *wlan2* and *ether1*
- *ether1* receives the “*trunk*” with all VLANs, the management as native
- ▲ RADIUS returns the VLAN of the client on successful authentication
  - Using the RADIUS attribute “Mikrotik_Wireless_VLANID”
- ▲ AP get management IP and the available CAPWAP WLCs from DHCP
- ▼ Users report AP blackouts from time to time (investigating…)
- ▼ Some Apple devices don’t connect if beacon lacks .11d attributes
- Some advices to better service:
  - Disable lower rates from the older standards
  - Correctly select the country and don’t cheat to get more power
  - Use “*multicast helper*” if multicast services needed (IPv6, Apple)
  - On 2,4GHz use 20MHz channels, and *20/40mhz-XX* on 5GHz
    - Priority on the quality of the signal over bandwidth

CRS/ROS at central aggregation – CRS317

- ▲ No performance complains (routing unused, only L2)
- ▲ S+2332LC10D allowed monetizing (€) the leased fibre pairs
  - MPLS runs over one fibre and direct IP/Switching on the other (2x10G!)
- ▲ Product (running ROS) as evolved quite since the purchase
  - Some instabilities in the beginning
- ▼ With VLANs, a lot of care must be taken when managing the *bridge ports/vlans* – Our advice is:
  - Manage VLANs one by one, not grouped
  - Activate the **ingress-filtering** and the appropriate ***frame-types*** in the ports
  - In the access ports (untagged) the **PVID** should be selected on the *bridge port* options
  - Without a good reason for it, don't select untagged VLAN on the VLAN management
  - On the bridge, enable **vlan-filtering** and change the PVID to the management VLAN
- ▼ The servers had failover issues with the first batches of S+RJ10

CRS/ROS as an AP aggregator

- L3 functionalities aren’t used at equipment's in this role
- CRS328 at two places terminating a lot of AP
  - PoE, PoE+ e PassivePoE Mikrotik
  - ▲ Supports Mikrotik APs and APs from other manufacturers using PoE standards
  - ▼ The power hungry APs from some manufacturers (non MT) rapidly exhausts the available power/ports
- CRS112 at use in small aggregations
  - PoE e PassivePoE Mikrotik
  - ▲ Equipment usable for micro sites
    - Low cost and connection versatility (8 x 1G PoE UTP + 4 SFP ports)
    - E.g. rooms with temporary or sudden need for networking (until proper cabling is installed)
  - ▼ Performance very degraded when VLANs are enabled due to switch chip limitations (hardware acceleration is disabled)
    - Documented on the MT wiki for the CRS/switch chip model used

CRS/ROS providing end user and VoIP connectivity

- CRS326 and CRS328 working at pilot zones
- L3 functionalities aren’t used at equipment's in this role
- ▲ Performance without user complains
- ▲ The right/best speed was always found
- ▼ Limited control of rogue devices
  - Domestic networking gear added by users
  - Unwanted access to privileged ports/VLANs
  - ▲ IEEE802.1x was added in recent ROS, that could solve some of this problems
- ▼ VoIP central provisioning doesn’t work, phones need manual configuration onsite due to missing VoiceVLAN info not provided by the switch
  - We still need to use switches from other manufacturers to this role

CRS/ROS Compatibility with STP/VTP

- Some work is needed to good coexistence
  - We have a lot of legacy equipment's from other manufacturers like Cisco, Alcatel/Nokia, Dell, HP & F10
- Mikrotik only supports a single instance of (R)STP over the native VLAN
  - MSTP is not an option due to requirements from the standard (eg. root location)
- If we place CRSs inside of redundant rings with PVSTP equipment's
  - *Loops* and VLAN islands occur due to different views of the (spanning)*tree*

- We accomplished compatibility with this STP/VTP flavours when:
  - The PVSTP switches must process and let thru the BPDUs on the VLAN1/native
  - The trunk ports joining the switches along each ring must all let thru the same VLAN set
  - The VTP pruning must be disabled in the VTP enabled devices

Our wish list for ROS/HW evolution (1)

- Neighbour discovery protocols under more control
  - Selective optional activation for CDP, LLDP and MNDP
  - CDP/LLDP need to support sending VoiceVLAN attributes (VoIP phones need it!)
  - Filtering to avoid flooding of this packets (some arrive a few switches away!)
    - By default CDP and LLDP packets shouldn't be forwarded and they should be only processed and sent on the native VLAN and with differentiated attributes for each bridge port
- Sorting of /export lines and parameters
  - All the lines that ordering doesn't change the behaviour (e.g. /ip address)
  - /export decoupled from the order of command insertions
    - E.g. firewall, bridge vlans/ports
- API
  - Should provide an inventory of chassis and SFPs
- WEBFIG
  - Should confirm "destructive" actions (slip pointers)

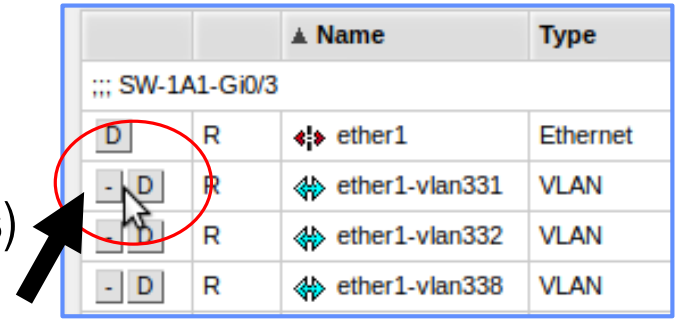

Our wish list for ROS/HW evolution (2)

- OSPF/OSPFv3
  - Easy way for control the static routes redistributed
- DHCP (client on APs)
  - ~~The IP list included in the CAPWAP attribute should be processed in preference order (avoid random CAPSMAN selection) –~~ It's already fixed in latest ROS!
- IPv6 should send the gateway preference on RA (RFC4191)
  - Native failover and route selection scheme for terminal equipment
  - Only 2bit need to be manipulated on sent RAs
- Serial CLI
  - Should be more robust against random bootloader aborts on boot
    - `/system routerboard settings`

      `set enter-setup-on=delete-key` – Isn't enough, noise/coupling cause hangs
    - Don't leave cables connected to the serial console ports

Our wish list for ROS/HW evolution (3)

- Product for PS/RPS role
  - With dual AC input for high availability and flexible maintenance
  - To supply power to groups of MT devices on racks (e.g. CRS326s)
- wAP
  - Beacons should include IEEE802.11d attributes to avoid Apple problems
- Switching ROS
  - Functionality's to secure the users edge
    - Limiting the number of MAC addresses connected to each port and sticky learning
    - Simplified filtering of STP BPDUs (avoid interferences between zones)
      - *bridge filter*? Seems risky and can affect performance on some devices
  - Managing the “bridge port” and “bridge vlan” in a more intuitive way
  - Listing “bridge hosts” should have some kind of filtering on WEBFIG
    - Switches/browser become irresponsive when the list is long

New projects (probably using Mikrotik gear …)

- Remodelling of our computing & virtualization clusters
  - CRS312-4C+8XG-RM seems a good solution with their 10G UTP ports
- Expansion of the VoIP network
  - Needs VoiceVLAN support on CRS328-24P-4S+RM to be an option
- Sharing the SAN networks between datacentres
  - Seems an option using the CWDM-MUX8A multiplexer and SFP of distinct wavelengths to share the current leased fibres between the Fiberchannel and IP/MPLS uses
- Sensor monitoring and domotics control on campus and buildings
  - We already participate in a wide area pilot covering the capital city (with gateways from other manufacturers)
  - New wAP LoRa8 seems an option to consider

Some conclusions …

- ▼ Switching
  - Integration requirements forced us to acquire switches compatible with all the APs and VoIP phones in use
  - The CRS switching is still in pilot phase for edge uses, we are only using it when capacity is a priority and we edge security isn't a must
- ▲ WiFi eduroam
  - wAPac are still a preferred option
  - Some acquisitions to cover new areas
  - ▼ We still need some time to investigate the random connectivity hangs
  - ▼ Some compatibility problems (seem easy to solve!)
- ▲ If Mikrotik has the needed product, it's our preference
  - Due to their ratio of capabilities vs cost
  - Due to the flexible, uniform and simplified configuration
  - We should support the European products

Thank you!

- Acknowledgments
  - To the event organization for invitation
  - To my colleagues at DSIC by their collaboration and dedication to the service
  - To all of you for taking your time to this presentation

- To Mikrotik for the fantastic products they have made!
  - And for giving some time to our bug reports