

An aerial photograph of Athens, Greece, showing a dense urban landscape. The Olympic Stadium is visible in the center, surrounded by roads and green spaces. The text "Free IX Greece" is overlaid in the center.

# Free IX Greece

darknob.net

freeix.gr

Free IX Greece

text:darknob.net



**Free IX Greece**

**Free IX Greece**

Free IX Greece



**Free IX Greece**



# What is Free IX GR?

- Physical IXP in Thessaloniki and Athens
- One single Peering LAN between the two cities
- No setup or recurring fees for 10G/25G/100G\* ports
- Non-profit pro-bono service to improve connectivity in the country



# Where is Free IX GR?

- Balkan Gate I (Kalochoi, Thessaloniki)
- TI Sparkle I & II (Metamorfosi, Athens)



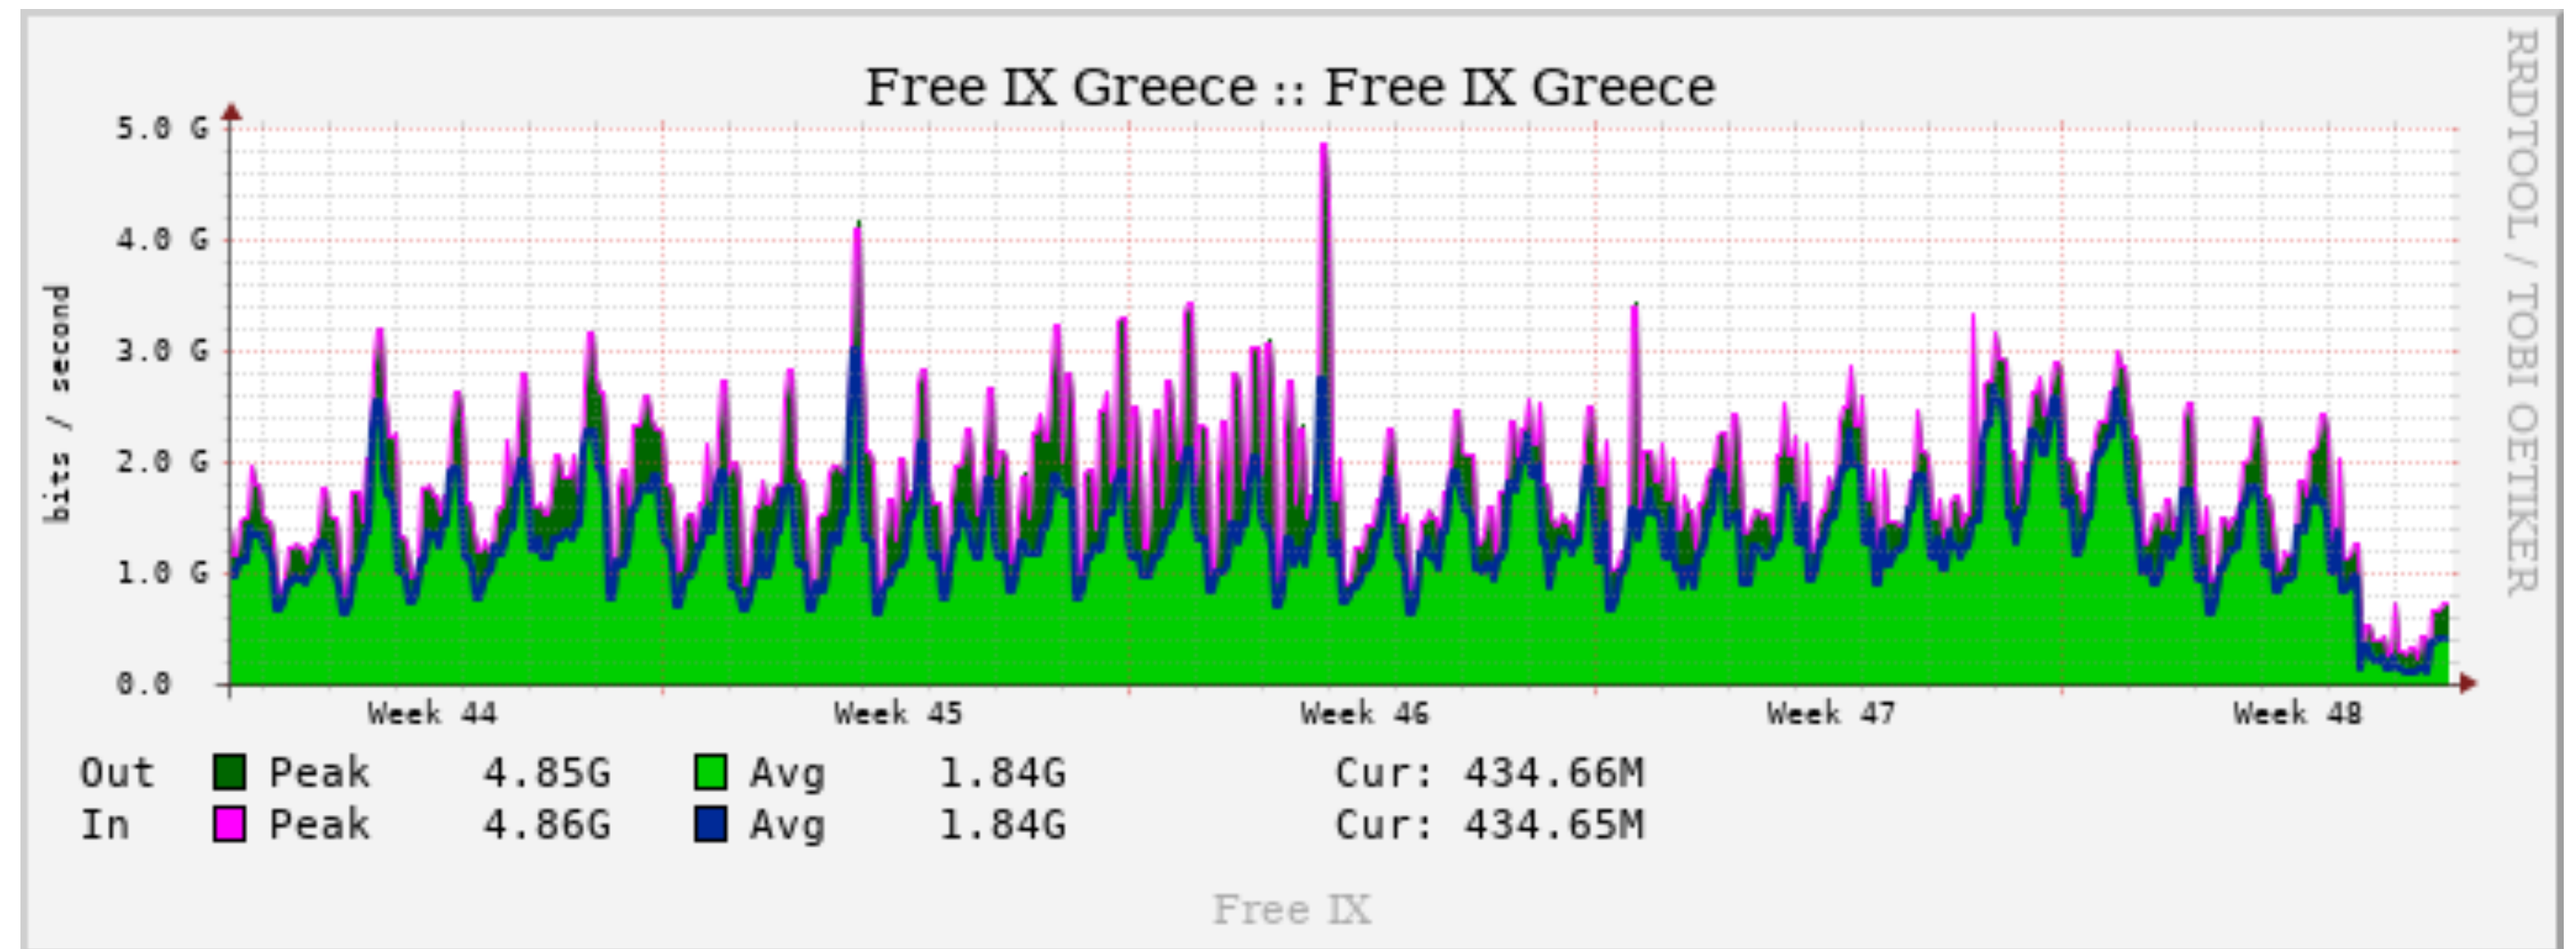
# Where is Free IX GR?

- Balkan Gate I (Kalochori, Thessaloniki)
- TI Sparkle I & II (Metamorfosi, Athens)
- Digital Realty ATH\* (Koropi, Athens)
- Digital Realty HER1 (Heraklion, Crete)



# Stats

- 10 Members
- ~ 2 Gbps of traffic
- 2 Switches





# Switches

- Cisco Nexus 93180YC-FX
- 48x25G + 6x100G
- NX-OS 10
- ~€500 on eBay ;)





# Servers

- Dell PowerEdge R640
- IXP Manager
- bird2 with extra Info + Action Communities
- AS112 DNS
- Metrics (Prometheus, LibreNMS, ...)
- ~€800





# Fabric

- 2x100G between the two cities / switches
- IP Underlay with OSPF & ECMP
- EVPN-VXLAN Full Mesh for Peering LAN & Private VLANs
- BUM elimination
- 1'500 MTU for all access ports

**An IXP on a Cisco?**



**Let's try!**

# Problems

- Licensing
- Software
- Hardware
- Proprietariness / Vendor Lock-in



# Software

- NX-OS BGP is... \\_(ツ)\_/
- Options are moved everywhere and deprecated constantly
- Stumbled into some bugs
- Required redundant config (additional route maps for next-hop)
- Troubleshooting (EVPN) was difficult / felt the lack of tooling
- The software seemed to be designed for other use cases (e.g. L3, not L2)

# Hardware



# Hardware

- Some basic necessities were not available in the silicon / ASIC
  - BCP 214 for BGP Session Culling (the ACLs couldn't match src && dst and ingress/egress)

# Hardware

- Some basic necessities were not available in the silicon / ASIC
  - BCP 214 for BGP Session Culling (the ACLs couldn't match src && dst and ingress/egress)
- TCAM space was limited and carving required predicting the future
  - We had to choose between QoS and L2/L3 ACLs
  - Any change to this in the future will require a hard reboot / downtime (if it's even possible)



# Hardware

- Some basic necessities were not available in the silicon / ASIC
  - BCP 214 for BGP Session Culling (the ACLs couldn't match src && dst and ingress/egress)
- TCAM space was limited and carving required predicting the future
  - We had to choose between QoS and L2/L3 ACLs
  - Any change to this in the future will require a hard reboot / downtime (if it's even possible)
- We can't remap VLANs (e.g. in "reseller" ports) so 100,101,102 -> 50 / Peering LAN
  - Port Security / ACLs / troubleshooting is port-granular
  - Problems with resold "ports" have to be dealt at the physical port level (collateral damage)

# Proprietaryness

- EVPN VXLAN needs “Route Targets” & “Route Distinguishers”
  - They can be static or set based on standards
  - NX-OS supports automatically setting them / managing them for you
    - If you use them, you can't ever add another vendor to your IXP
      - Unless you take down everything and re-create it from scratch

# Is this bad?

- No, each product has its uses
- We can work around most limitations with good engineering
- Free IX Greece has been stable for > 1 year with no issues
- We feel confident in our ability to add new PoPs or increase capacity



**free-ix.gr**

[info@free-ix.gr](mailto:info@free-ix.gr)