

<https://faelix.link/uknof48>

**HALF-WAY TO
AUTOMATION
ALL THE WAY IN ORKNEY**

About Marek

- ✘ CTO @FAELIX – <https://faelix.net/>
- ✘ PC @uknof – <https://uknof.uk/>
- ✘ Crew @net_mcr – <https://www.netmcr.uk/>

- ✘ Me — @maznu – @NetworkMoose

L2 TO L3 MIGRATIONS

Orcadian Inception

- ✘ Quite common in a lot of organically-grown alt-nets:
 - ✘ "Accidental WISP" resells a DIA via some PtP radio
 - ✘ Victims of own success keep growing L2 domain
 - ✘ Equipment can be mix of "best at the time" tech
 - ✘ No time to learn how to, plan for, and rearchitect
- ✘ WISP in Orkney desired rearchitect from L2 to L3
 - ✘ Also to support multiple 1Gbit/sec backhails
 - ✘ Also because of a 5G New Thinking project

Challenges

- ✘ 5GNT testbed is short-term project
 - ✘ Timescales, reporting, multi-stakeholder
 - ✘ WISP is local but also "incumbent" telco
- ✘ 5GNT project has \$large_vendor involved
 - ✘ \$large_vendor has preference of NOS/hardware
 - ✘ WISP bears most of cap-ex, all long-term op-ex

Opportunities

- ✘ Chance to get involved with operational 5G network
 - ✘ Project partner is leading on the RAN/core
- ✘ Deliver some real benefits to rural communities
 - ✘ Will be building "real infrastructure"
- ✘ Team at FAELIX lives for these kinds of L3 projects
 - ✘ Also might be nice to get out of the data-centre?

L2 TO L3 PLANNING


L2 to L3 Migration

- ✘ Architected / HLD in May 2021 (pen and paper!)
 - ✘ Some existing maps
 - ✘ Several video calls
 - ✘ Lots of chatting back and forth
 - ✘ Starting to appreciate the scale


Not All Choices are Equal

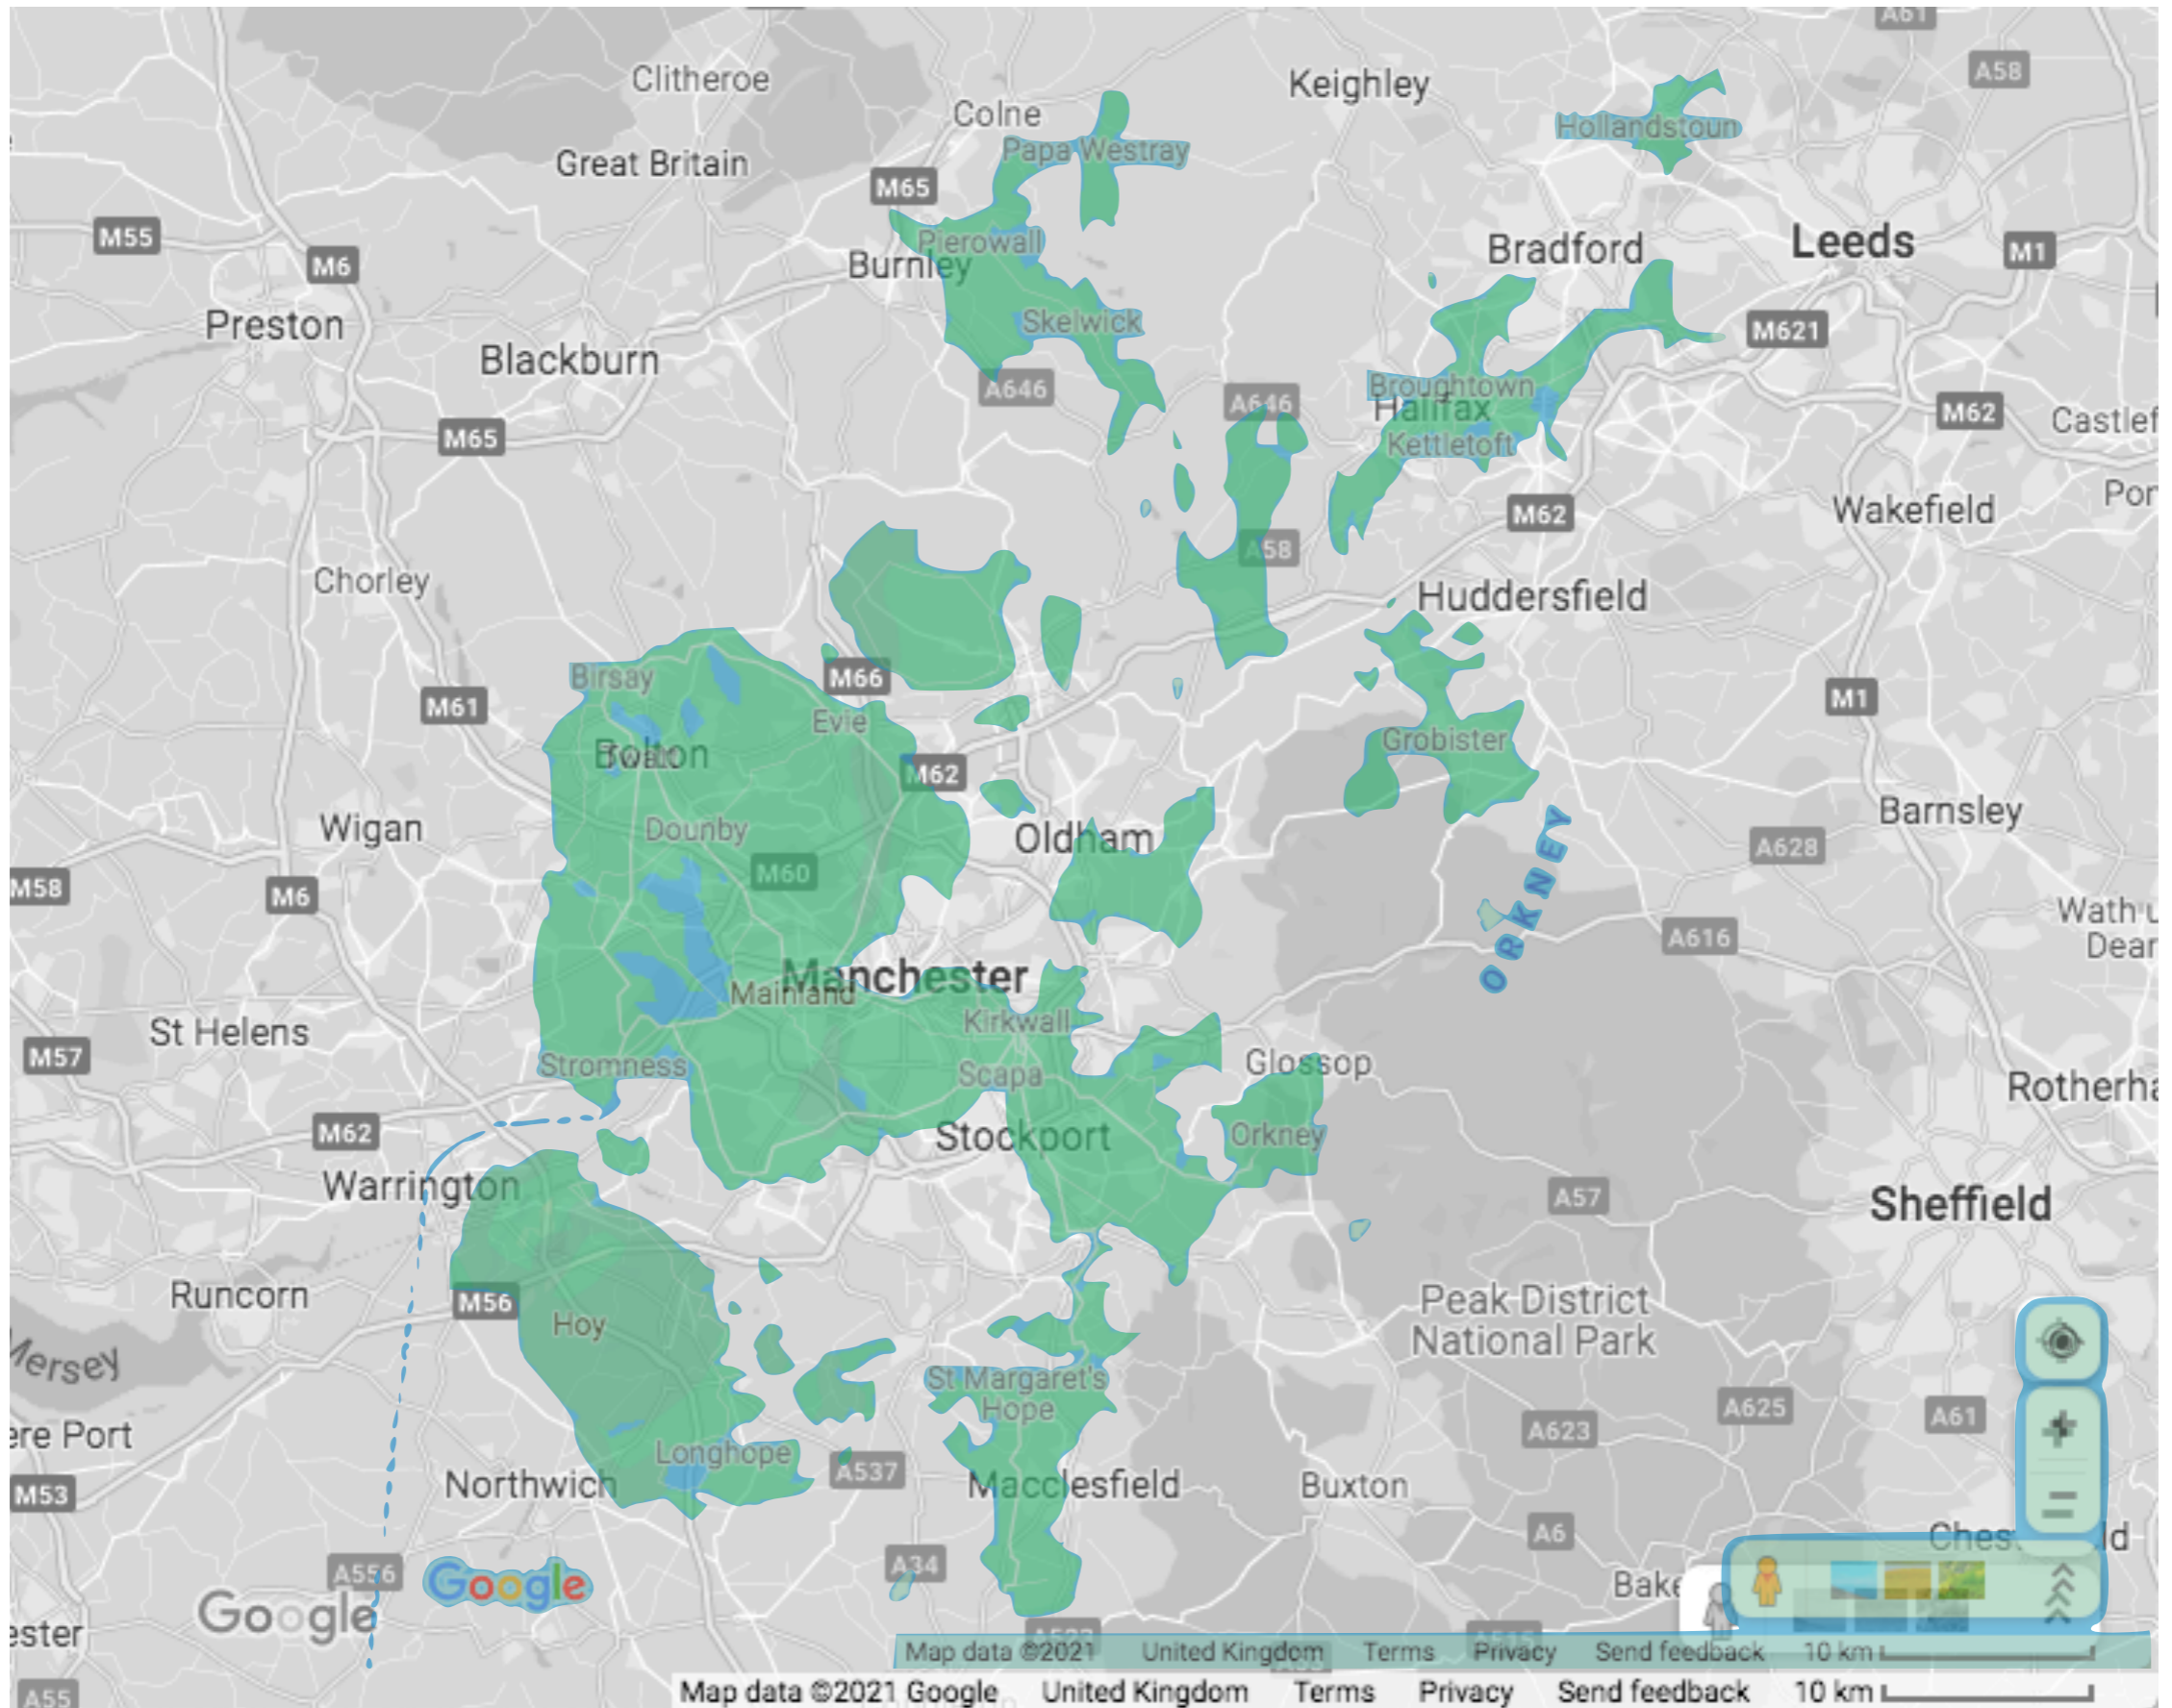
- ✘ Pretty typical \$large_vendor L2/L3 switch:
 - ✘ ££££ to buy (unless \$large_vendor is generous)
 - ✘ ££££ to hold a spare
- ✘ How many spares? (also: field engineers?)
 - ✘ How far can you truck-roll in a SLO period?
 - ✘ How multiples of that area is your total coverage?
 - ✘ Ratio approximates how many spares you need

ORKNEY: QUITE BIG



Google






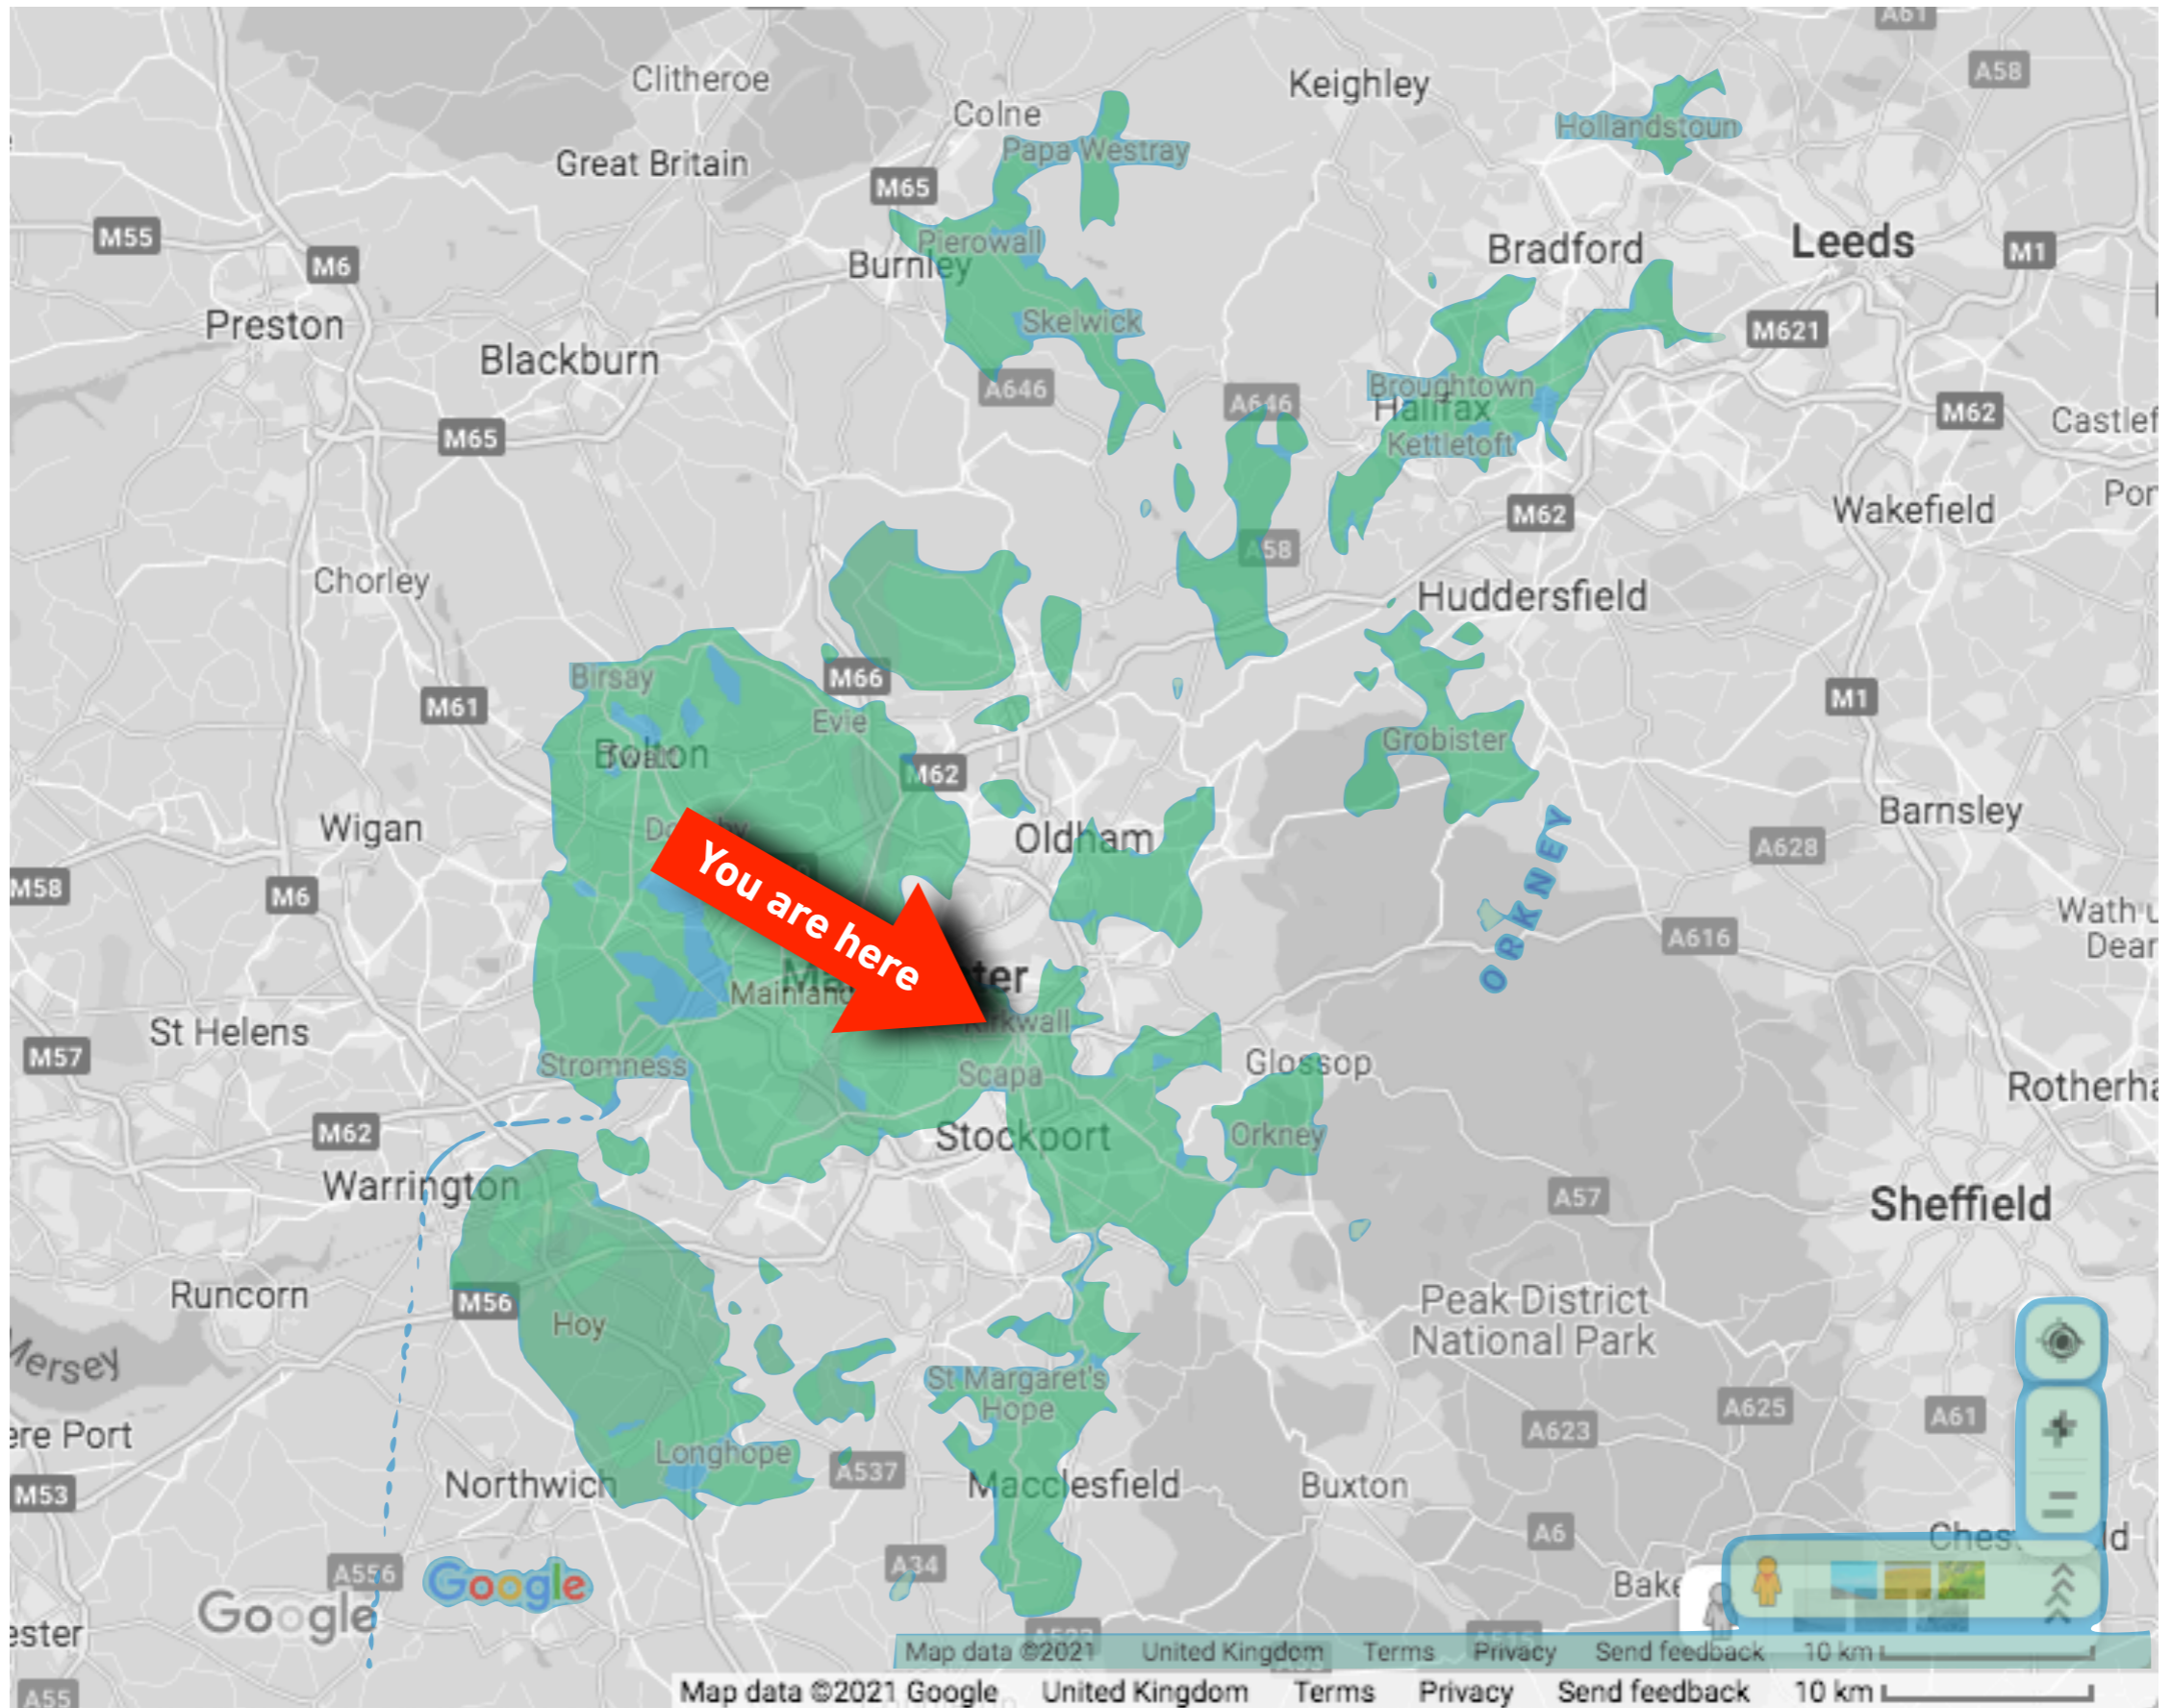
Google

Map data ©2021 United Kingdom Terms Privacy Send feedback 10 km

Map data ©2021 Google United Kingdom Terms Privacy Send feedback 10 km

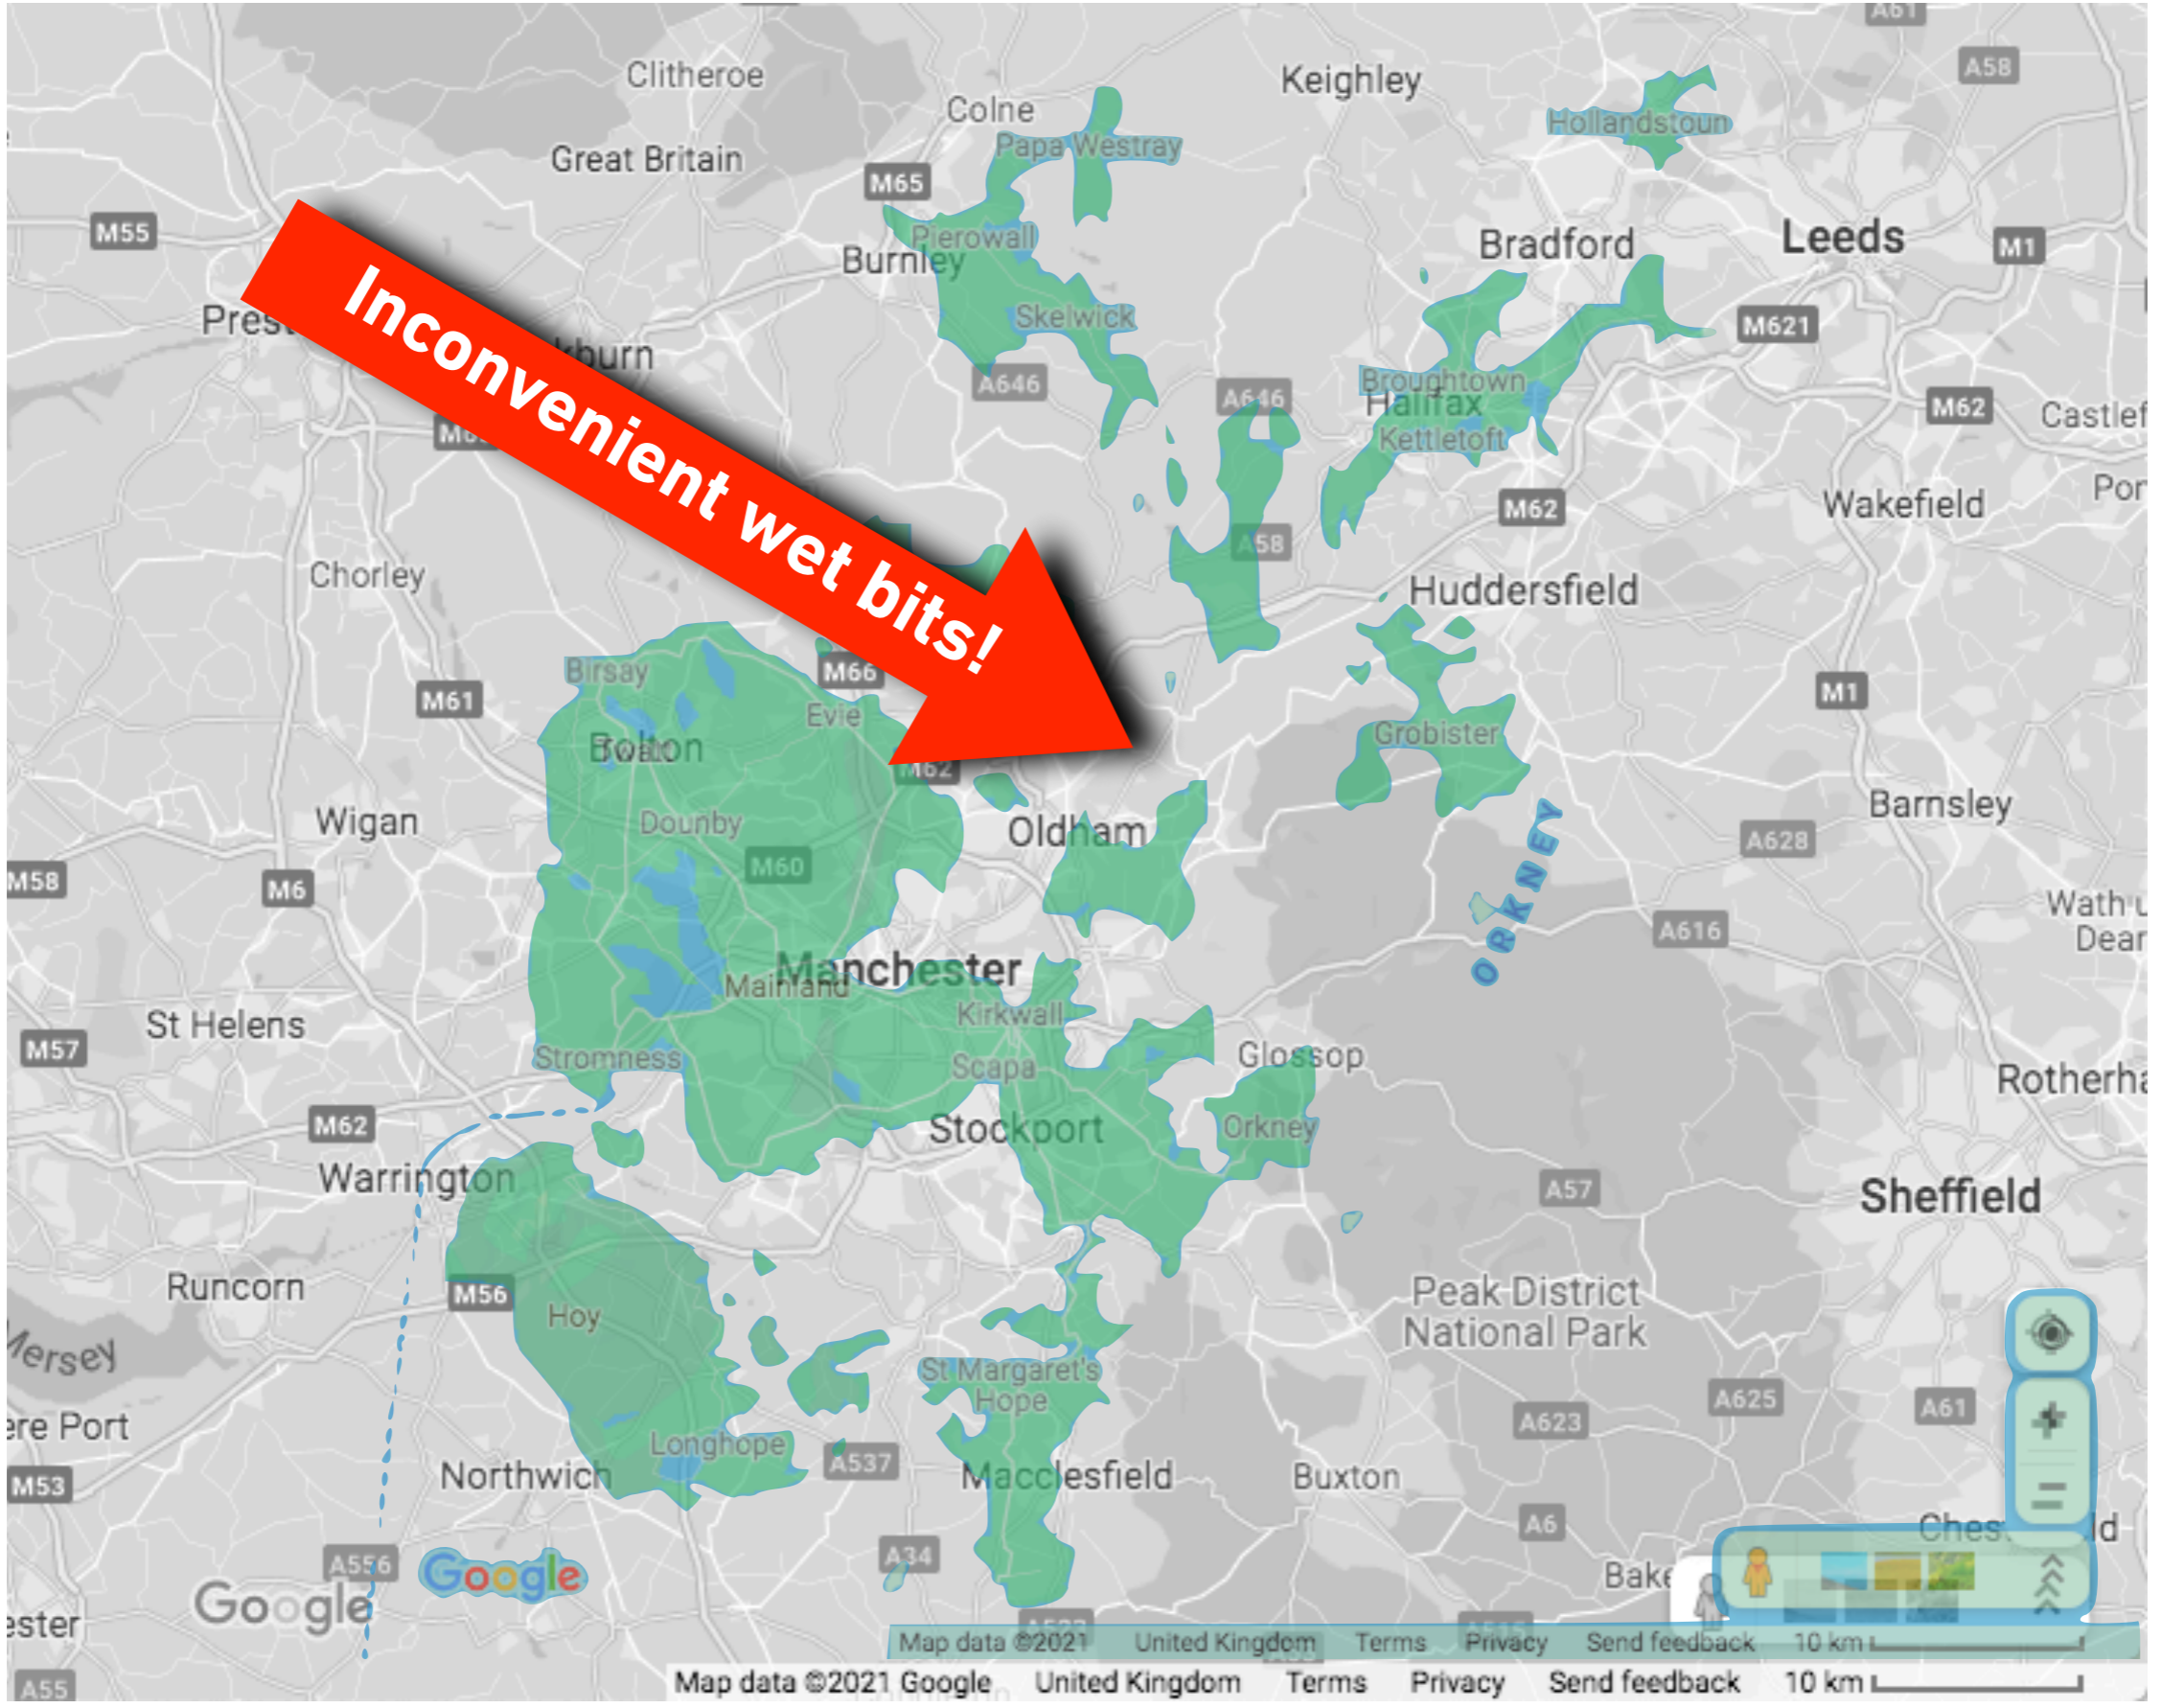
Leeds to Northwich
Bolton to Huddersfield
Burnley to Macclesfield





You are here





Inconvenient wet bits!

Google

Map data ©2021 United Kingdom Terms Privacy Send feedback 10 km

Inconvenient Wet Bits



Inconvenient Wet Bits



Not All Choices are Equal

- ✘ Warranty:
 - ✘ Is this hardware made for these environments?
 - ✘ How long would replacement take to arrive?



ICELAND

Labrador
Sea

Thousands of Miles of Sea

UNITED
KINGDOM

IRELAND

Birmingham

Lon

Celtic
Sea

Bay of
Biscay

Quite Windy



**"On Orkney the rain can be vertical.
Vertical upwards."**

Typical High Site



Keeping Dry Inside



Will Warranty Cover This?



Vendor Selection


- ✘ Sometimes you have to accept (and work around) sharp edges in the chosen NOS/hardware because "the customer is always right"
- ✘ Knows the area and their challenges
- ✘ Knows their budget
- ✘ Will be responsible for operations long-term

Architecture Decision

- ❌ IPv6/IPv4 dualstack from the beginning
- ❌ NOS has no IS-IS support
- ❌ Vendor's OSPF is a hot mess
- ❌ Vendor's iBGP with IPv6 is terrible
- ❌ Vendor's BGP next-hop ignores IGP metric

- ❌ Customer has preference — can we make it work?

Architecture Decision

- ❌ IPv6/IPv4 dualstack from the beginning 
- ❌ NOS has no IS-IS support
- ❌ Vendor's OSPF is a hot mess
- ❌ Vendor's iBGP with IPv6 is terrible
- ❌ Vendor's BGP next-hop ignores IGP metric

- ❌ Customer has preference — can we make it work?

Architecture Decision

- ❌ IPv6/IPv4 dualstack from the beginning ✓
- ❌ NOS has no IS-IS support
- ❌ Vendor's OSPF is a hot mess
- ❌ Vendor's iBGP with IPv6 is terrible
- ❌ Vendor's BGP next-hop ignores IGP metric
- ❌ Customer has preference — can we make it work?



don't use IS-IS

Architecture Decision

- ❌ IPv6/IPv4 dualstack from the beginning ✓
- ❌ NOS has no IS-IS support
- ❌ Vendor's OSPF is a hot mess → don't use OSPF
- ❌ Vendor's iBGP with IPv6 is terrible
- ❌ Vendor's BGP next-hop ignores IGP metric
- ❌ Customer has preference — can we make it work?

Architecture Decision

- ❌ IPv6/IPv4 dualstack from the beginning ✓
- ❌ NOS has no IS-IS support → don't use IS-IS
- ❌ Vendor's OSPF is a hot mess → don't use OSPF
- ❌ Vendor's iBGP with IPv6 is terrible → don't use iBGP
- ❌ Vendor's BGP next-hop ignores IGP metric

- ❌ Customer has preference — can we make it work?

Architecture Decision

❖ IPv6/IPv4 dualstack from the beginning ✓

❖ NOS has no IS-IS support

❖ Vendor's OSPF is a hot mess

❖ Vendor's iBGP with IPv6 is terrible

❖ Vendor's BGP next-hop ignores IGP metric

don't use IS-IS

don't use OSPF

don't use iBGP

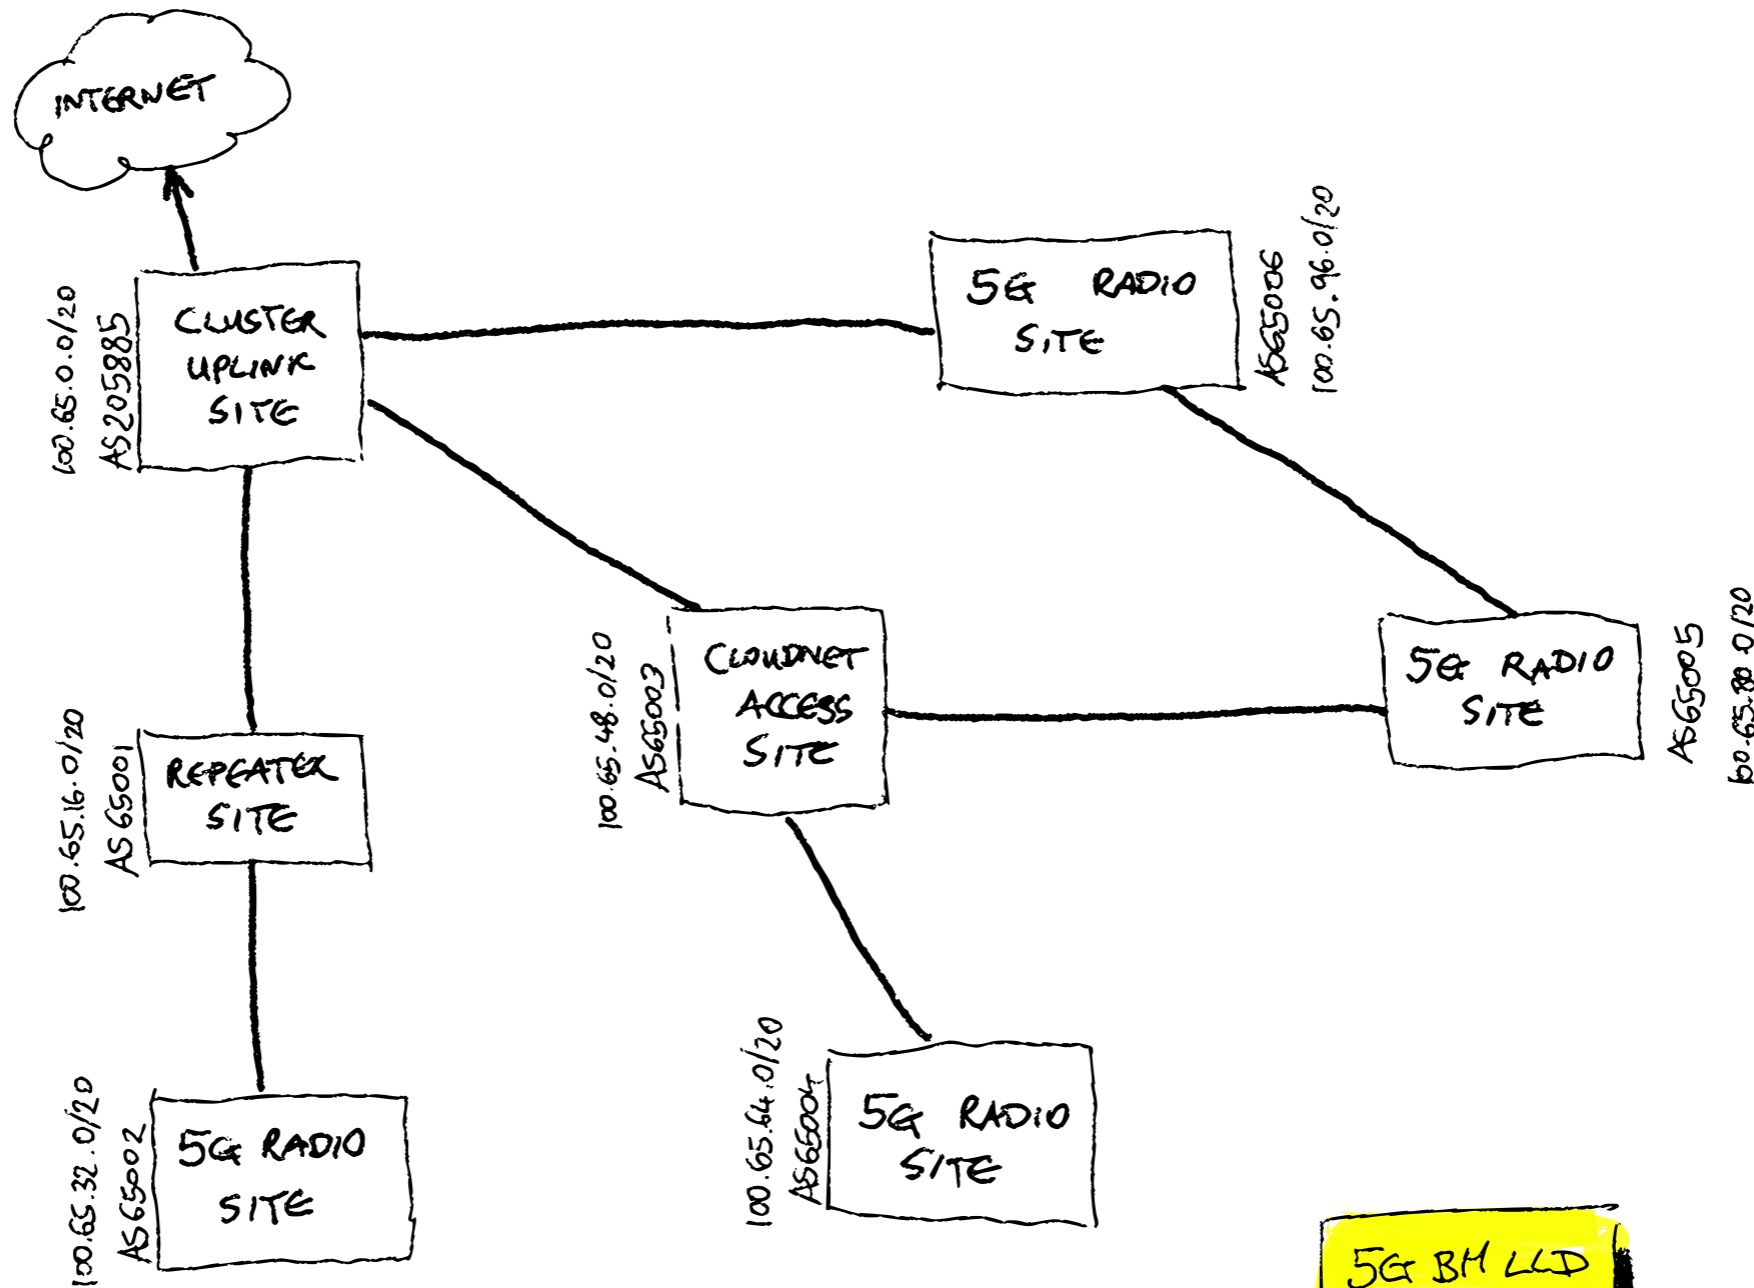
don't have IGP

❖ Customer has preference — can we make it work?

Architecture Decision


- ❌ IPv6/IPv4 dualstack from
 - ❌ NOS has no IS-IS supp
 - ❌ Vendor's OSPF is a hot
 - ❌ Vendor's iBGP with in v
 - ❌ Vendor's BGP next-hop ig
- 
- ❌ Customer has preference — can we make it work?

High Level Design

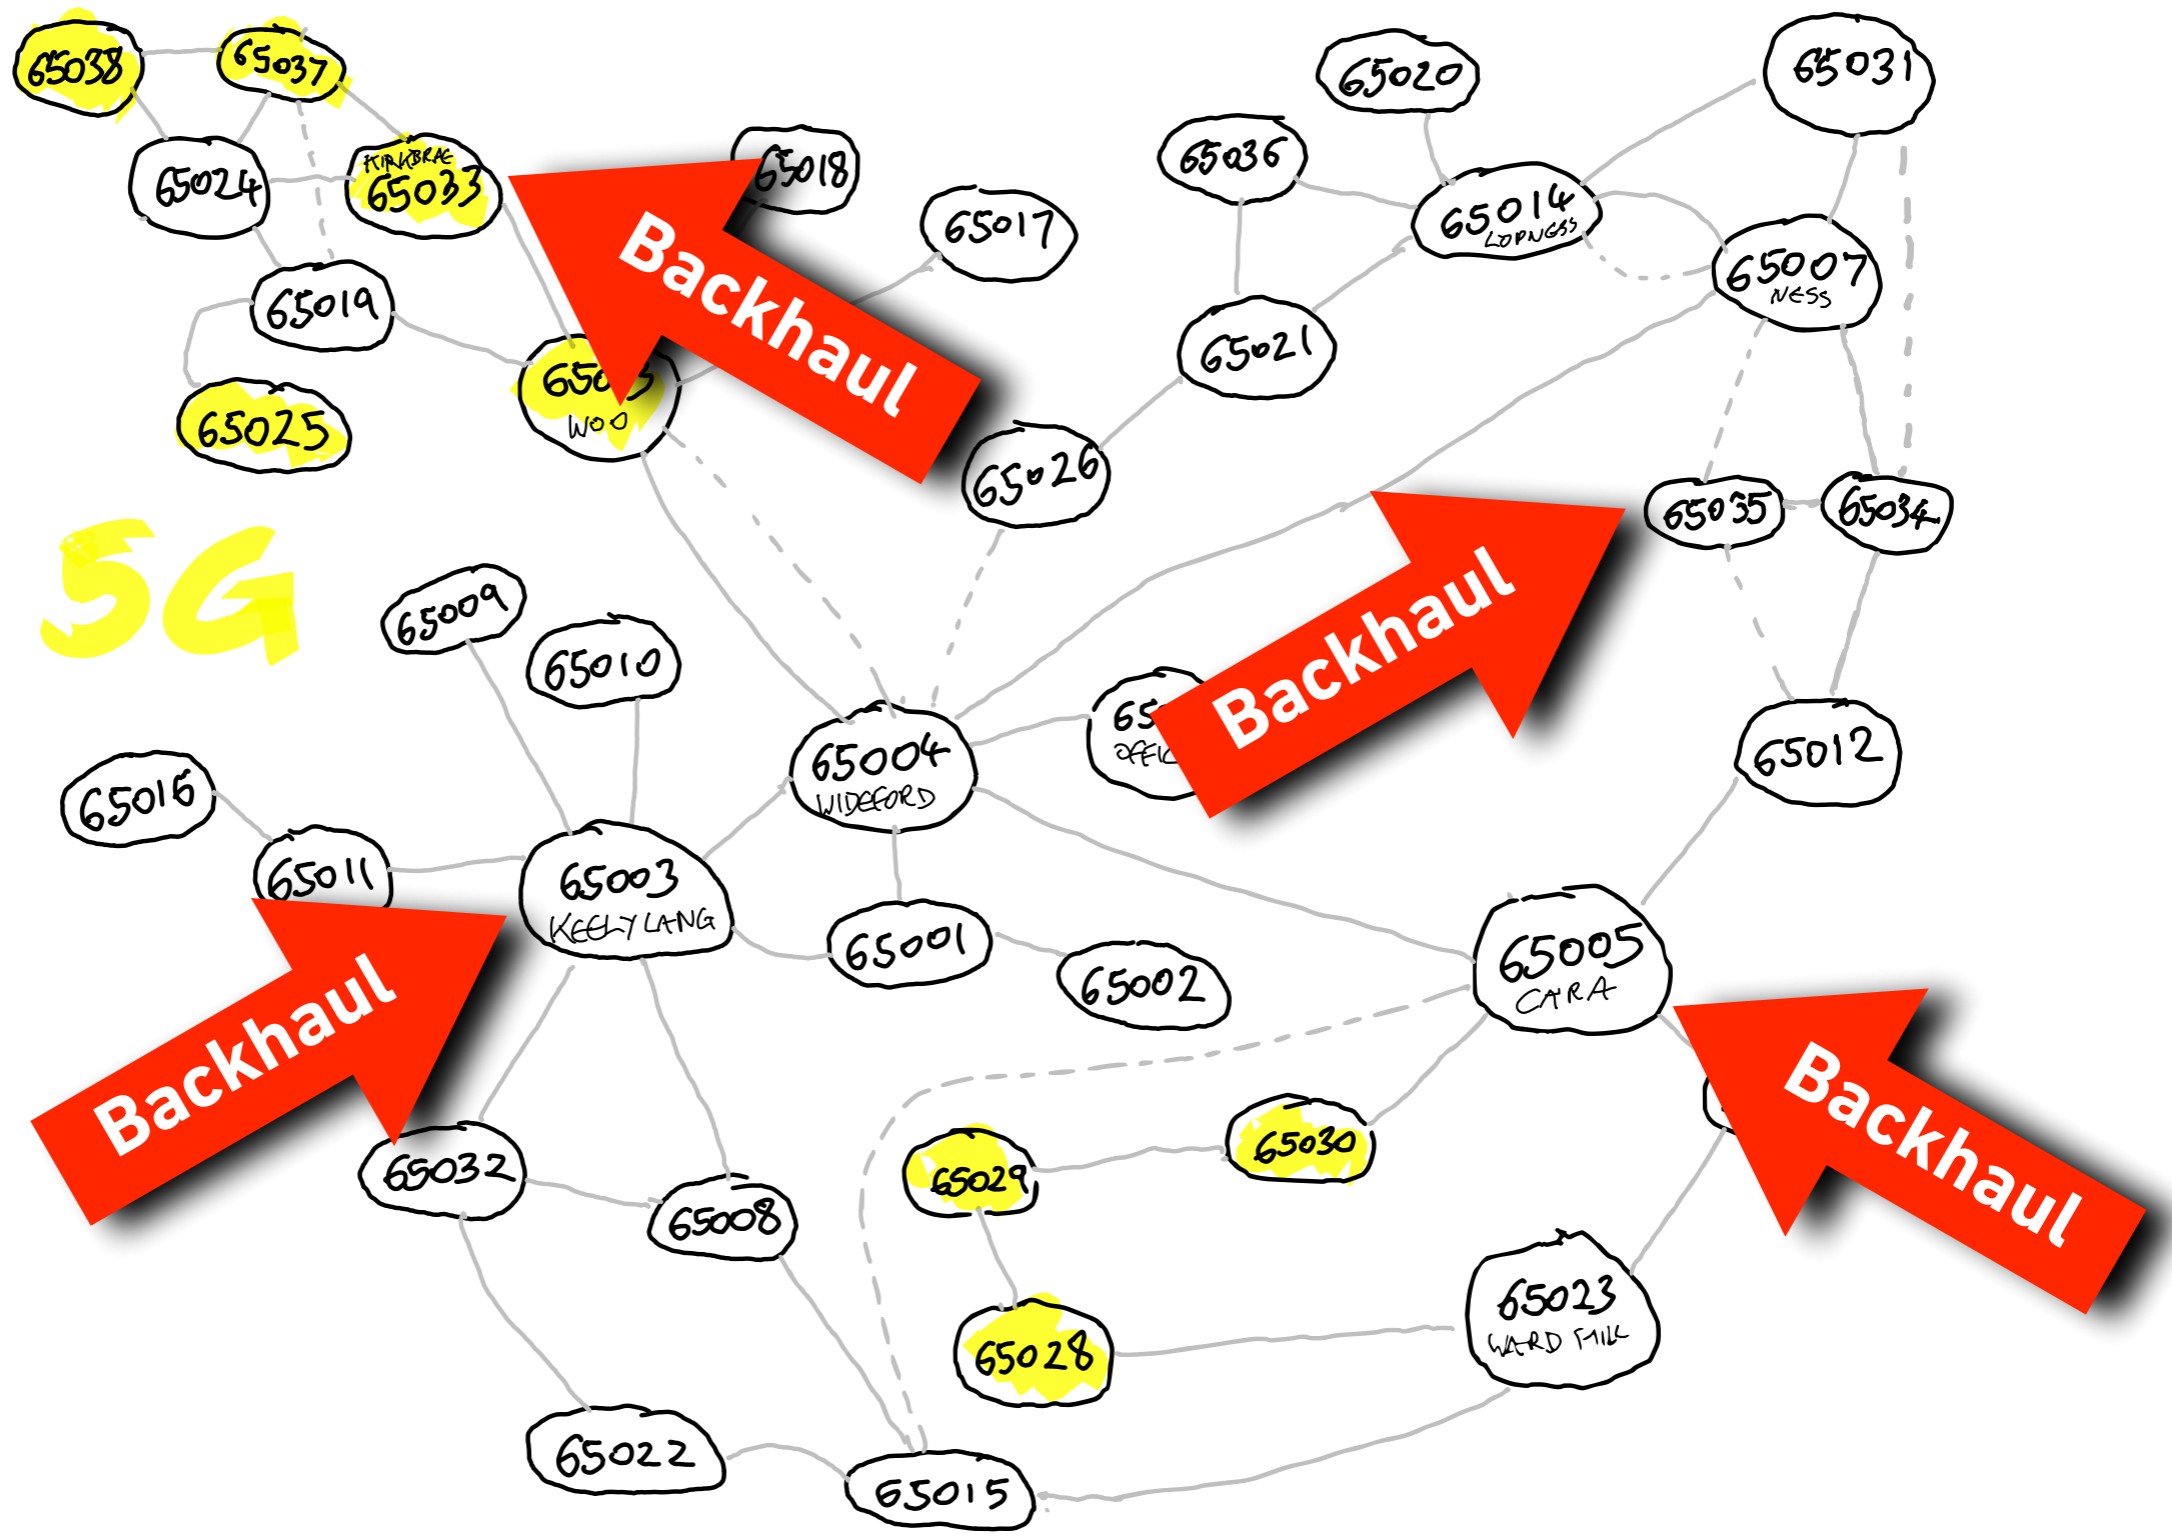


5G BH LLD
SCHEMATIC

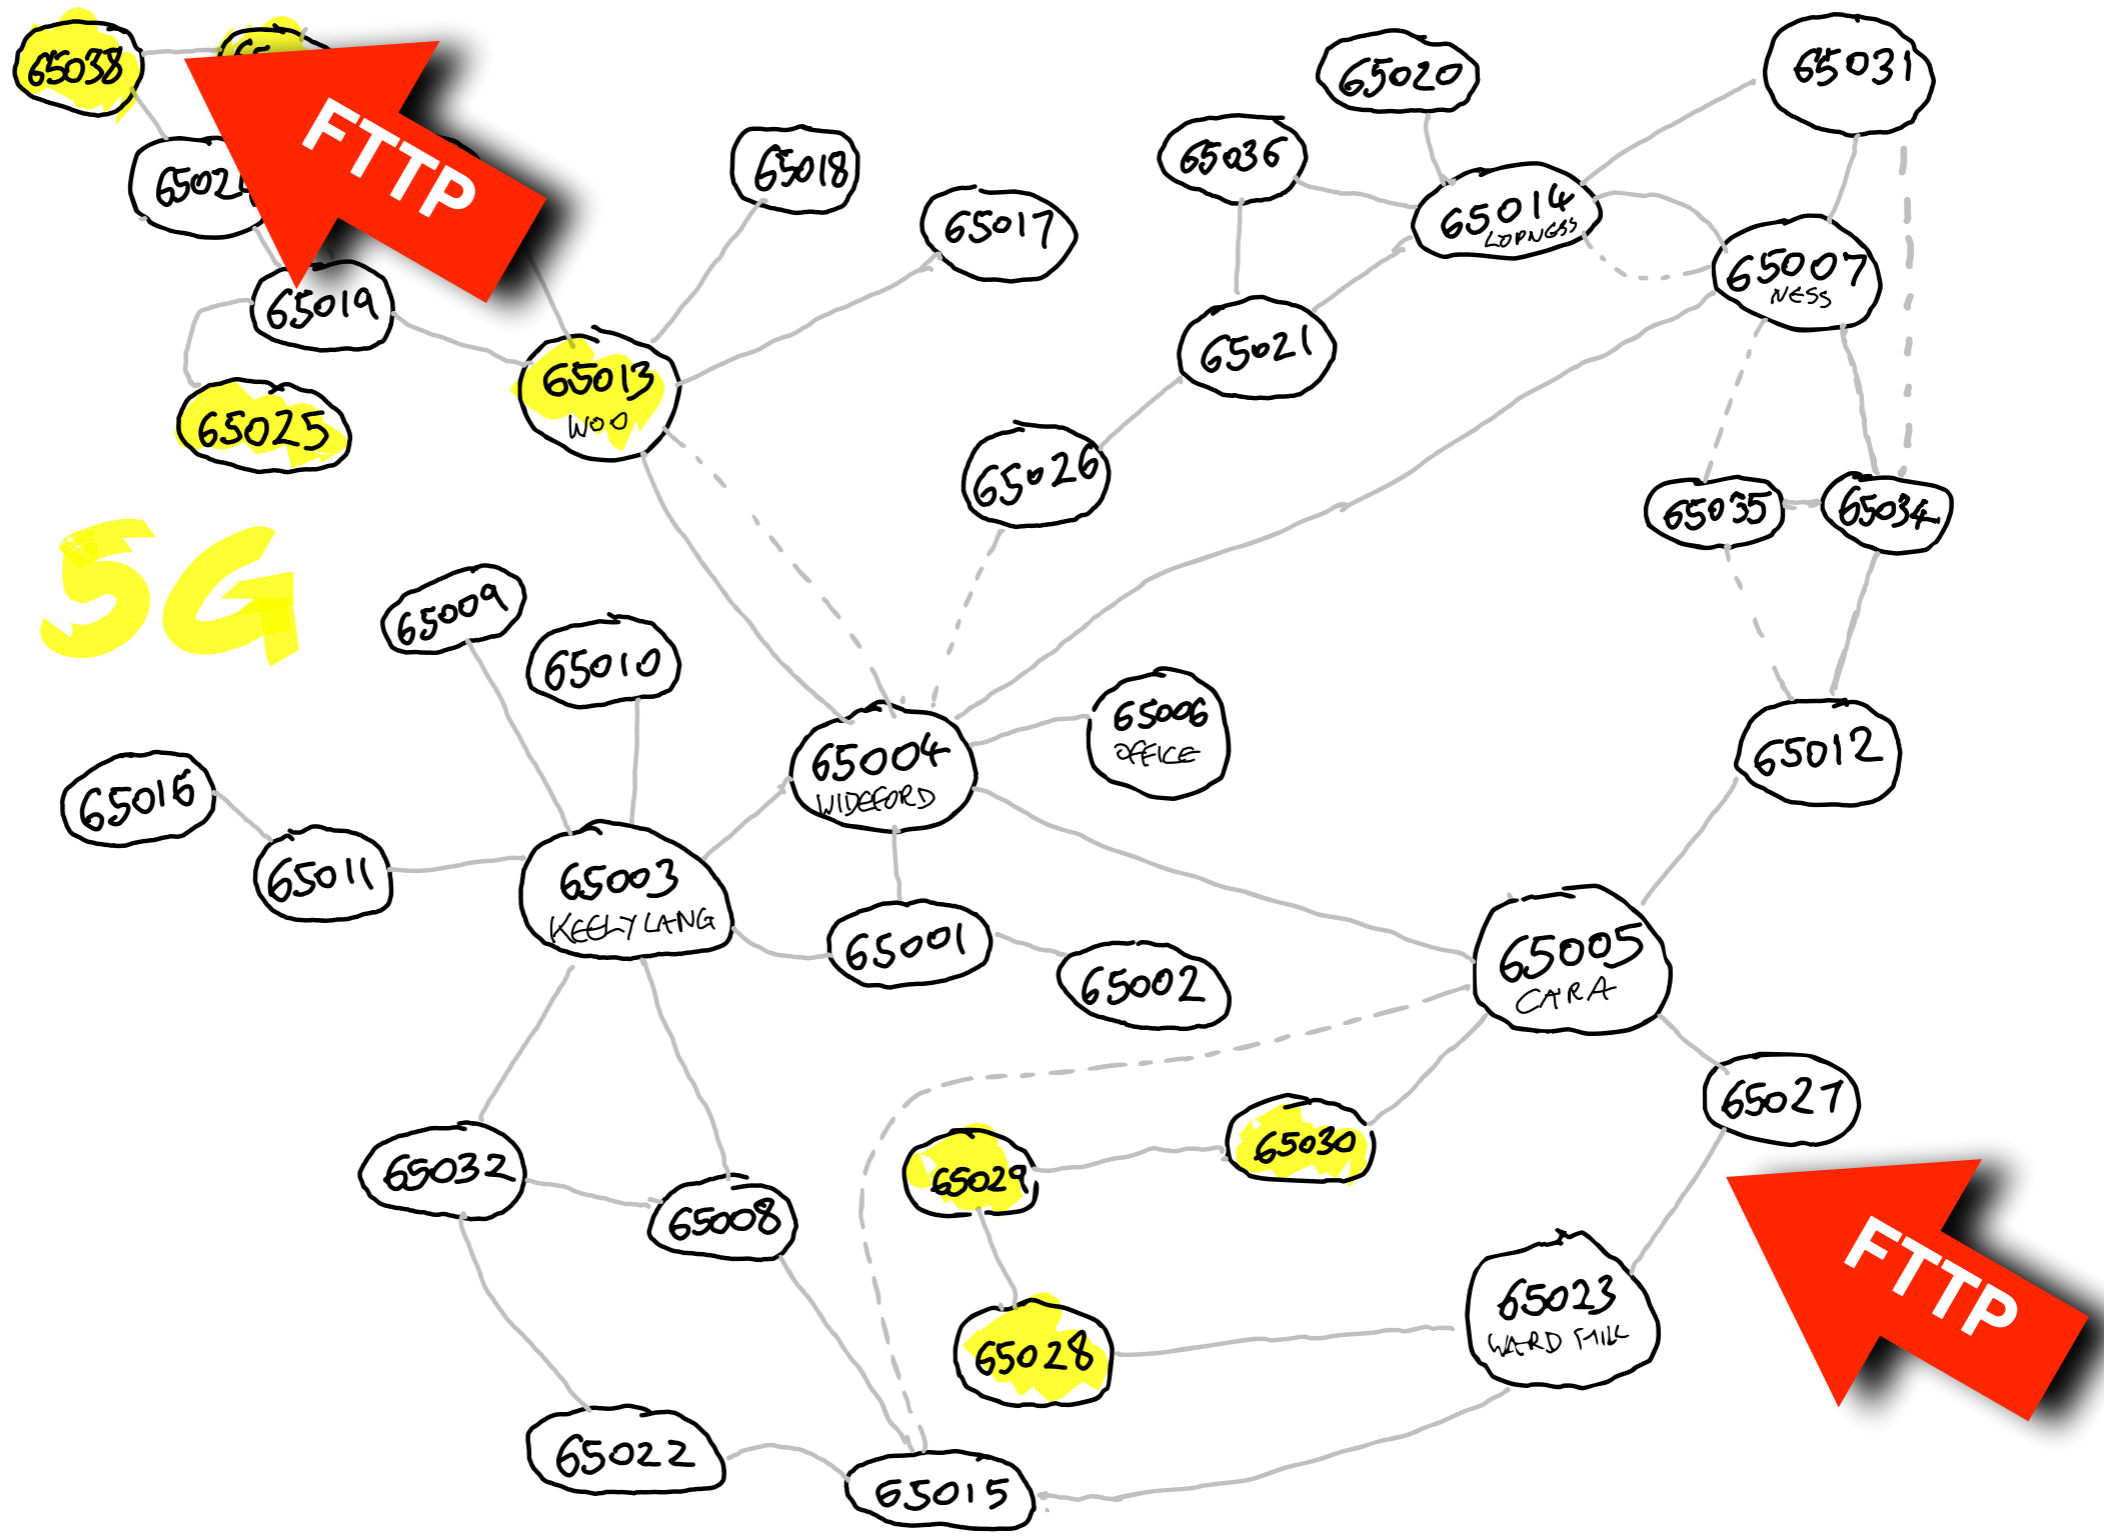
Mid-Level Design



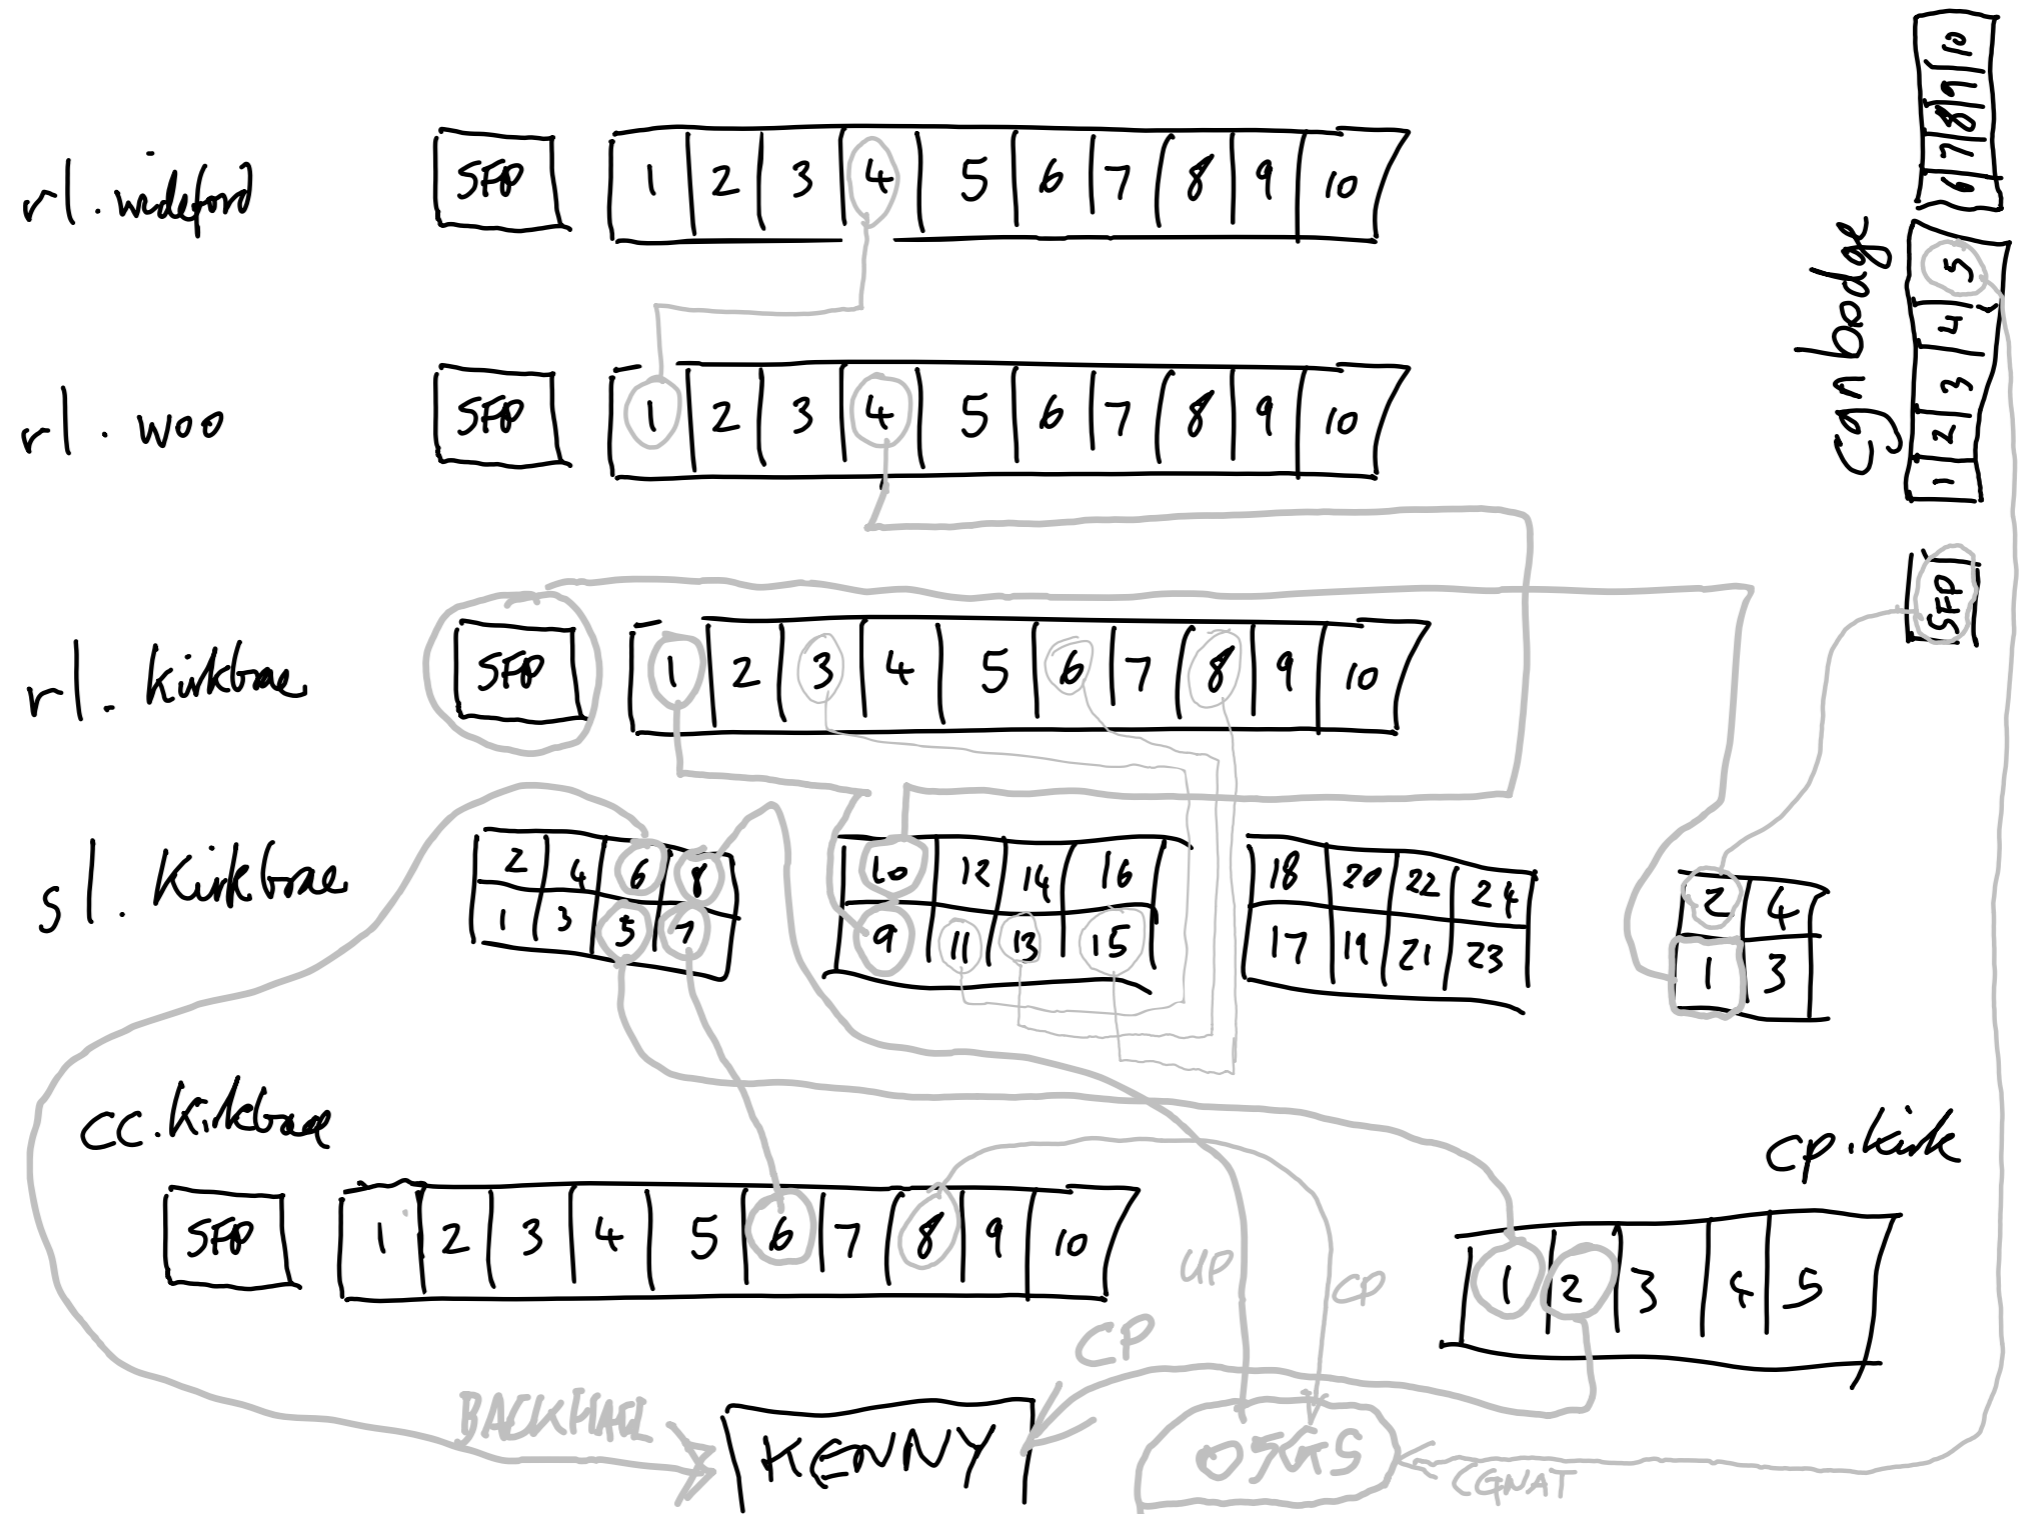
Details



Details



Low-Level Design



Model in Netbox

- ✘ Architected / HLD in May 2021 (pen and paper!)
- ✘ Low-Level Design in June 2021 (Netbox)

- ✘ Entire future-network was modelled in Netbox
 - ✘ Every device, port, cable, VLAN, IP
 - ✘ (including some we thought might be possible?)

Entire Network in Netbox

Organization

| | |
|---------|----|
| Sites | 38 |
| Tenants | 0 |

Inventory

| | |
|--------------|-----|
| Racks | 4 |
| Device Types | 22 |
| Devices | 769 |

Power

| | |
|--------------|---|
| Power Panels | 0 |
| Power Feeds | 0 |

IPAM

| | |
|--------------|------|
| VRFs | 2 |
| Aggregates | 6 |
| Prefixes | 1044 |
| IP Ranges | 0 |
| IP Addresses | 2322 |
| VLANs | 734 |

Circuits

| | |
|-----------|-----|
| Providers | 3 |
| Circuits | 388 |

Virtualization

| | |
|------------------|---|
| Clusters | 2 |
| Virtual Machines | 3 |

Connections

| | |
|-------------------|------|
| Cables | 1114 |
| Console | 0 |
| Interfaces | 979 |
| Power Connections | 0 |

Half-Way to Automation

- ✘ Generally backhaul networks are fairly static
 - ✘ Want POP "cores" to be a foundation
 - ✘ Allow POP "edges" to grow (e.g. new access tech)
- ✘ L2-to-L3 journey is quite long
 - ✘ The logical migration will be a lot of planning
 - ✘ The physical migration will be a lot of work
 - ✘ Adapting to new network will be process change


Half-Way to Automation

- ✘ Building a fully-automated network beyond scope
 - ✘ Yet another project dependency for roll-out
- ✘ Try something a bit more traditional
 - ✘ But with a view to automation in future
 - ✘ Cherry-pick important features of automated
 - ✘ Get the deployment benefits of automation without turning small WISP into NetDevOps Corp

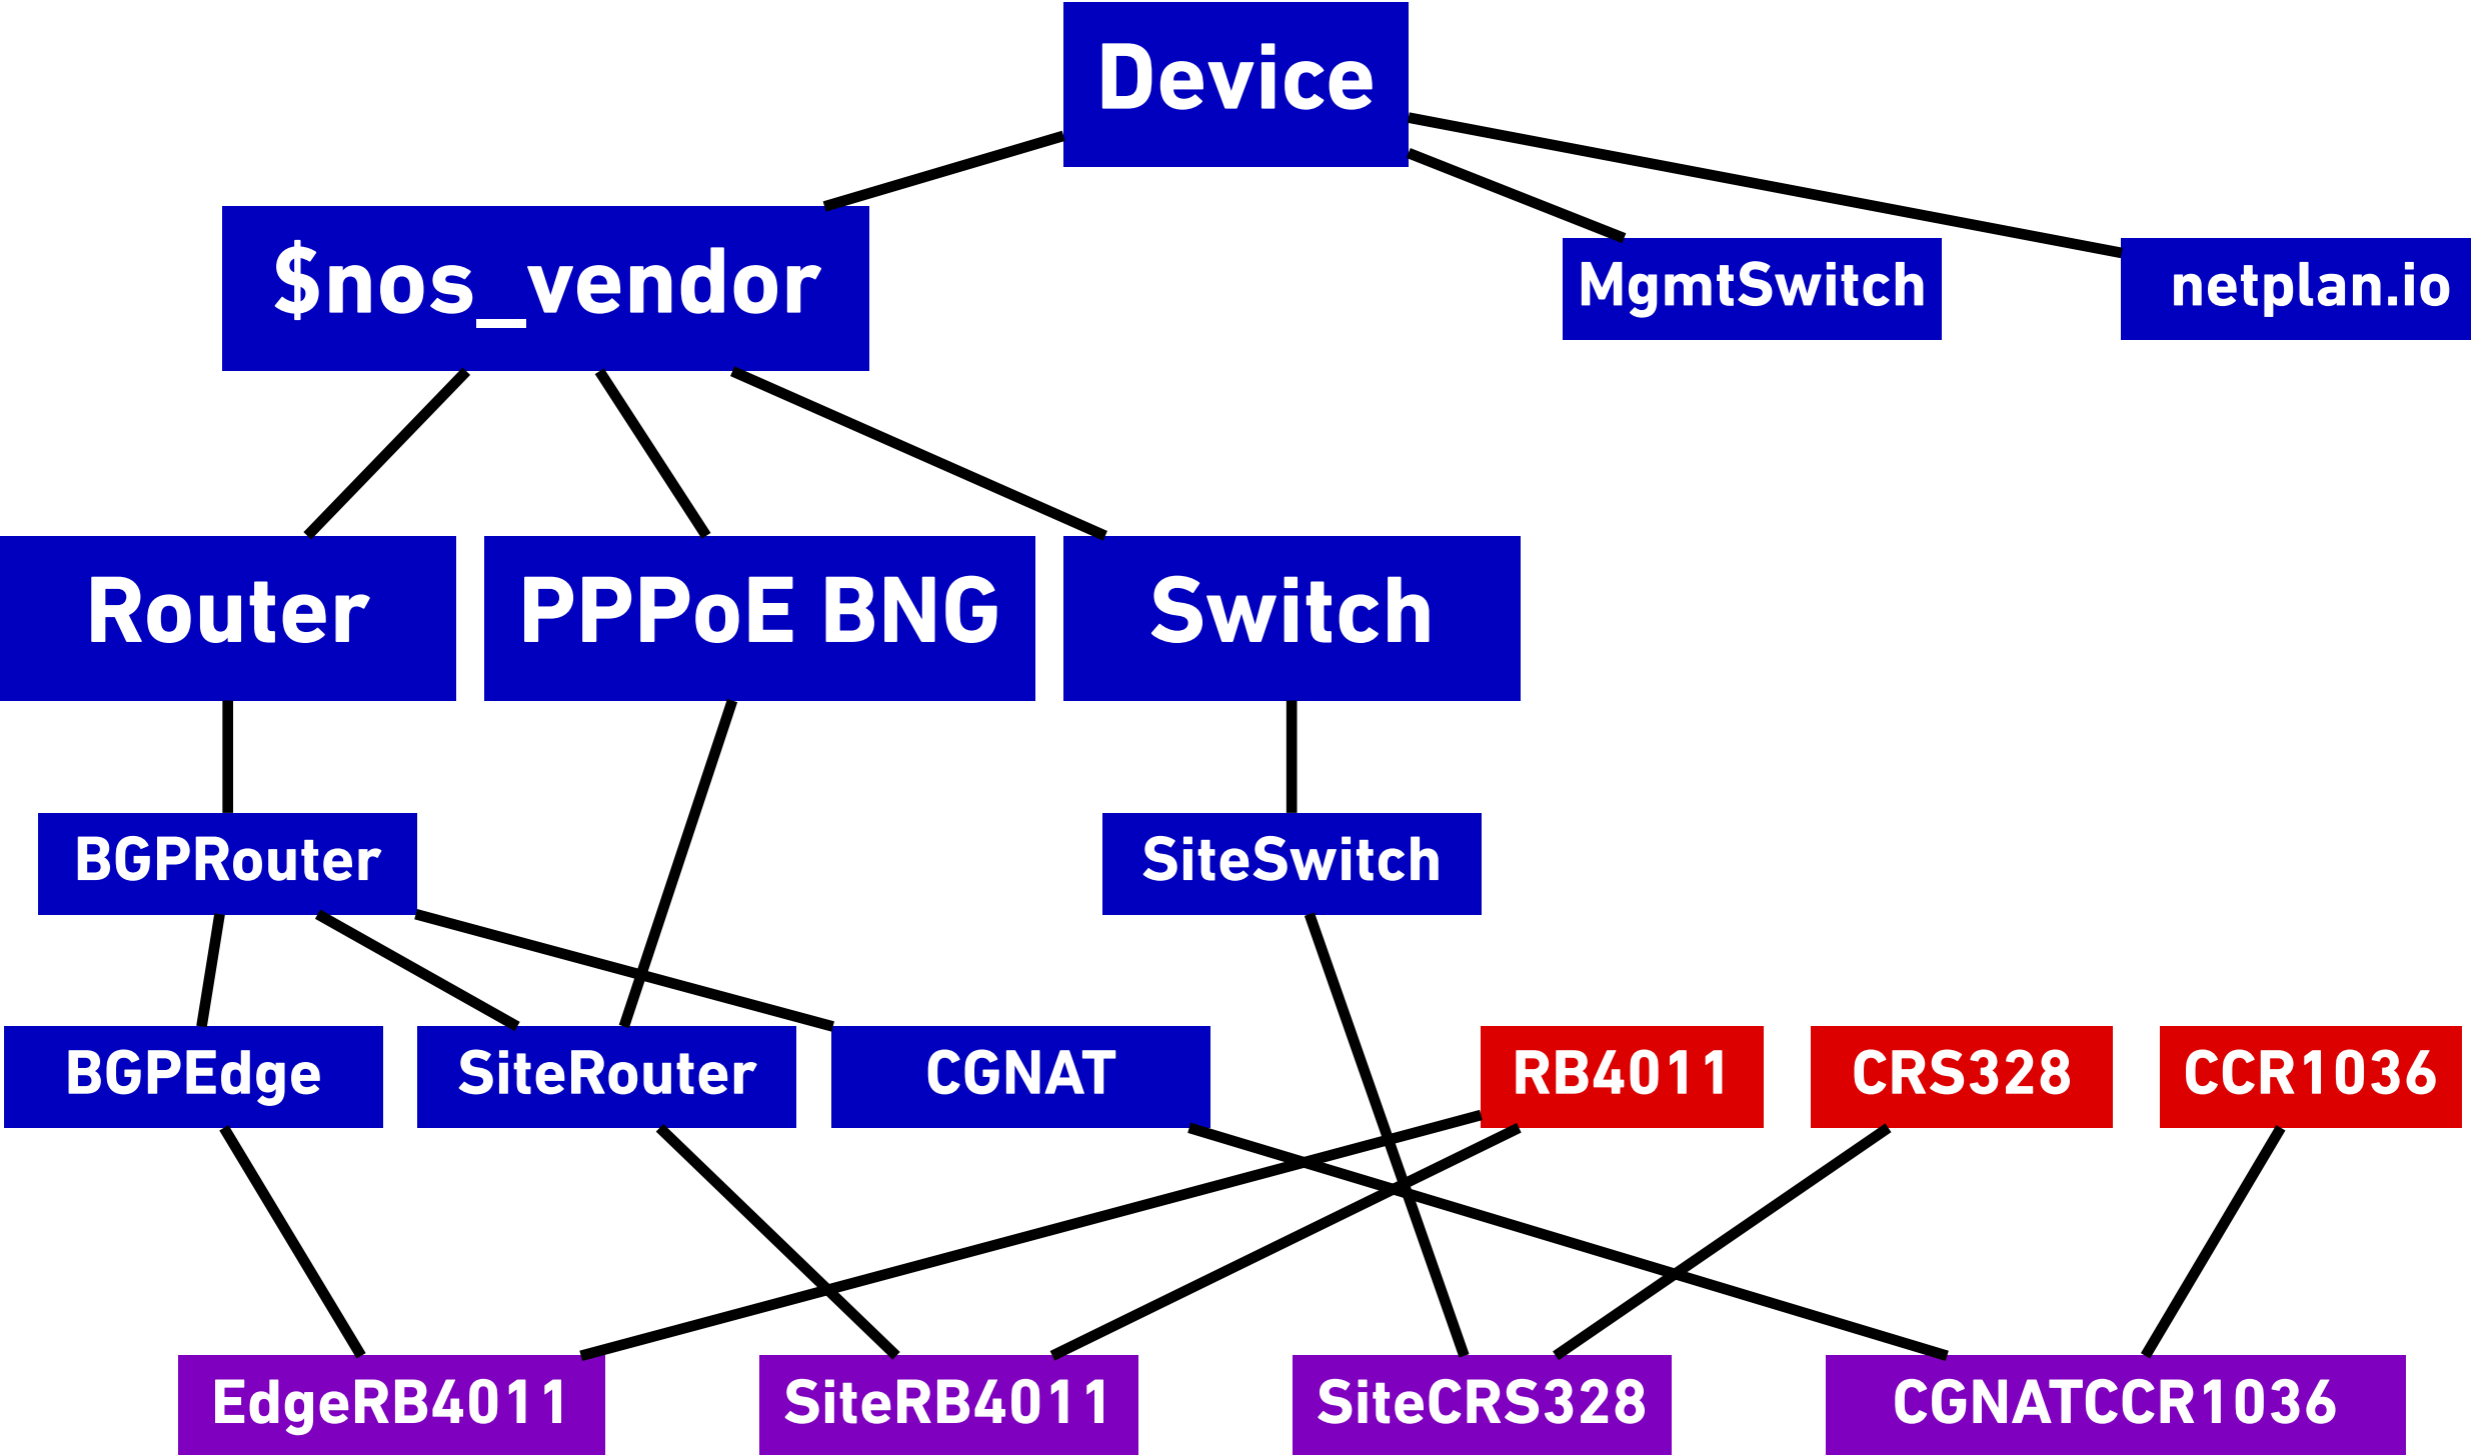
Semi-Automatic Approach

- ✘ Standardise hardware and network software
 - ✘ One device has one function
 - ✘ Eliminate special cases
- ✘ Build configs for all devices with software
 - ✘ Template output for everything
- ✘ /* TODO: how to amend all deployed configs? */

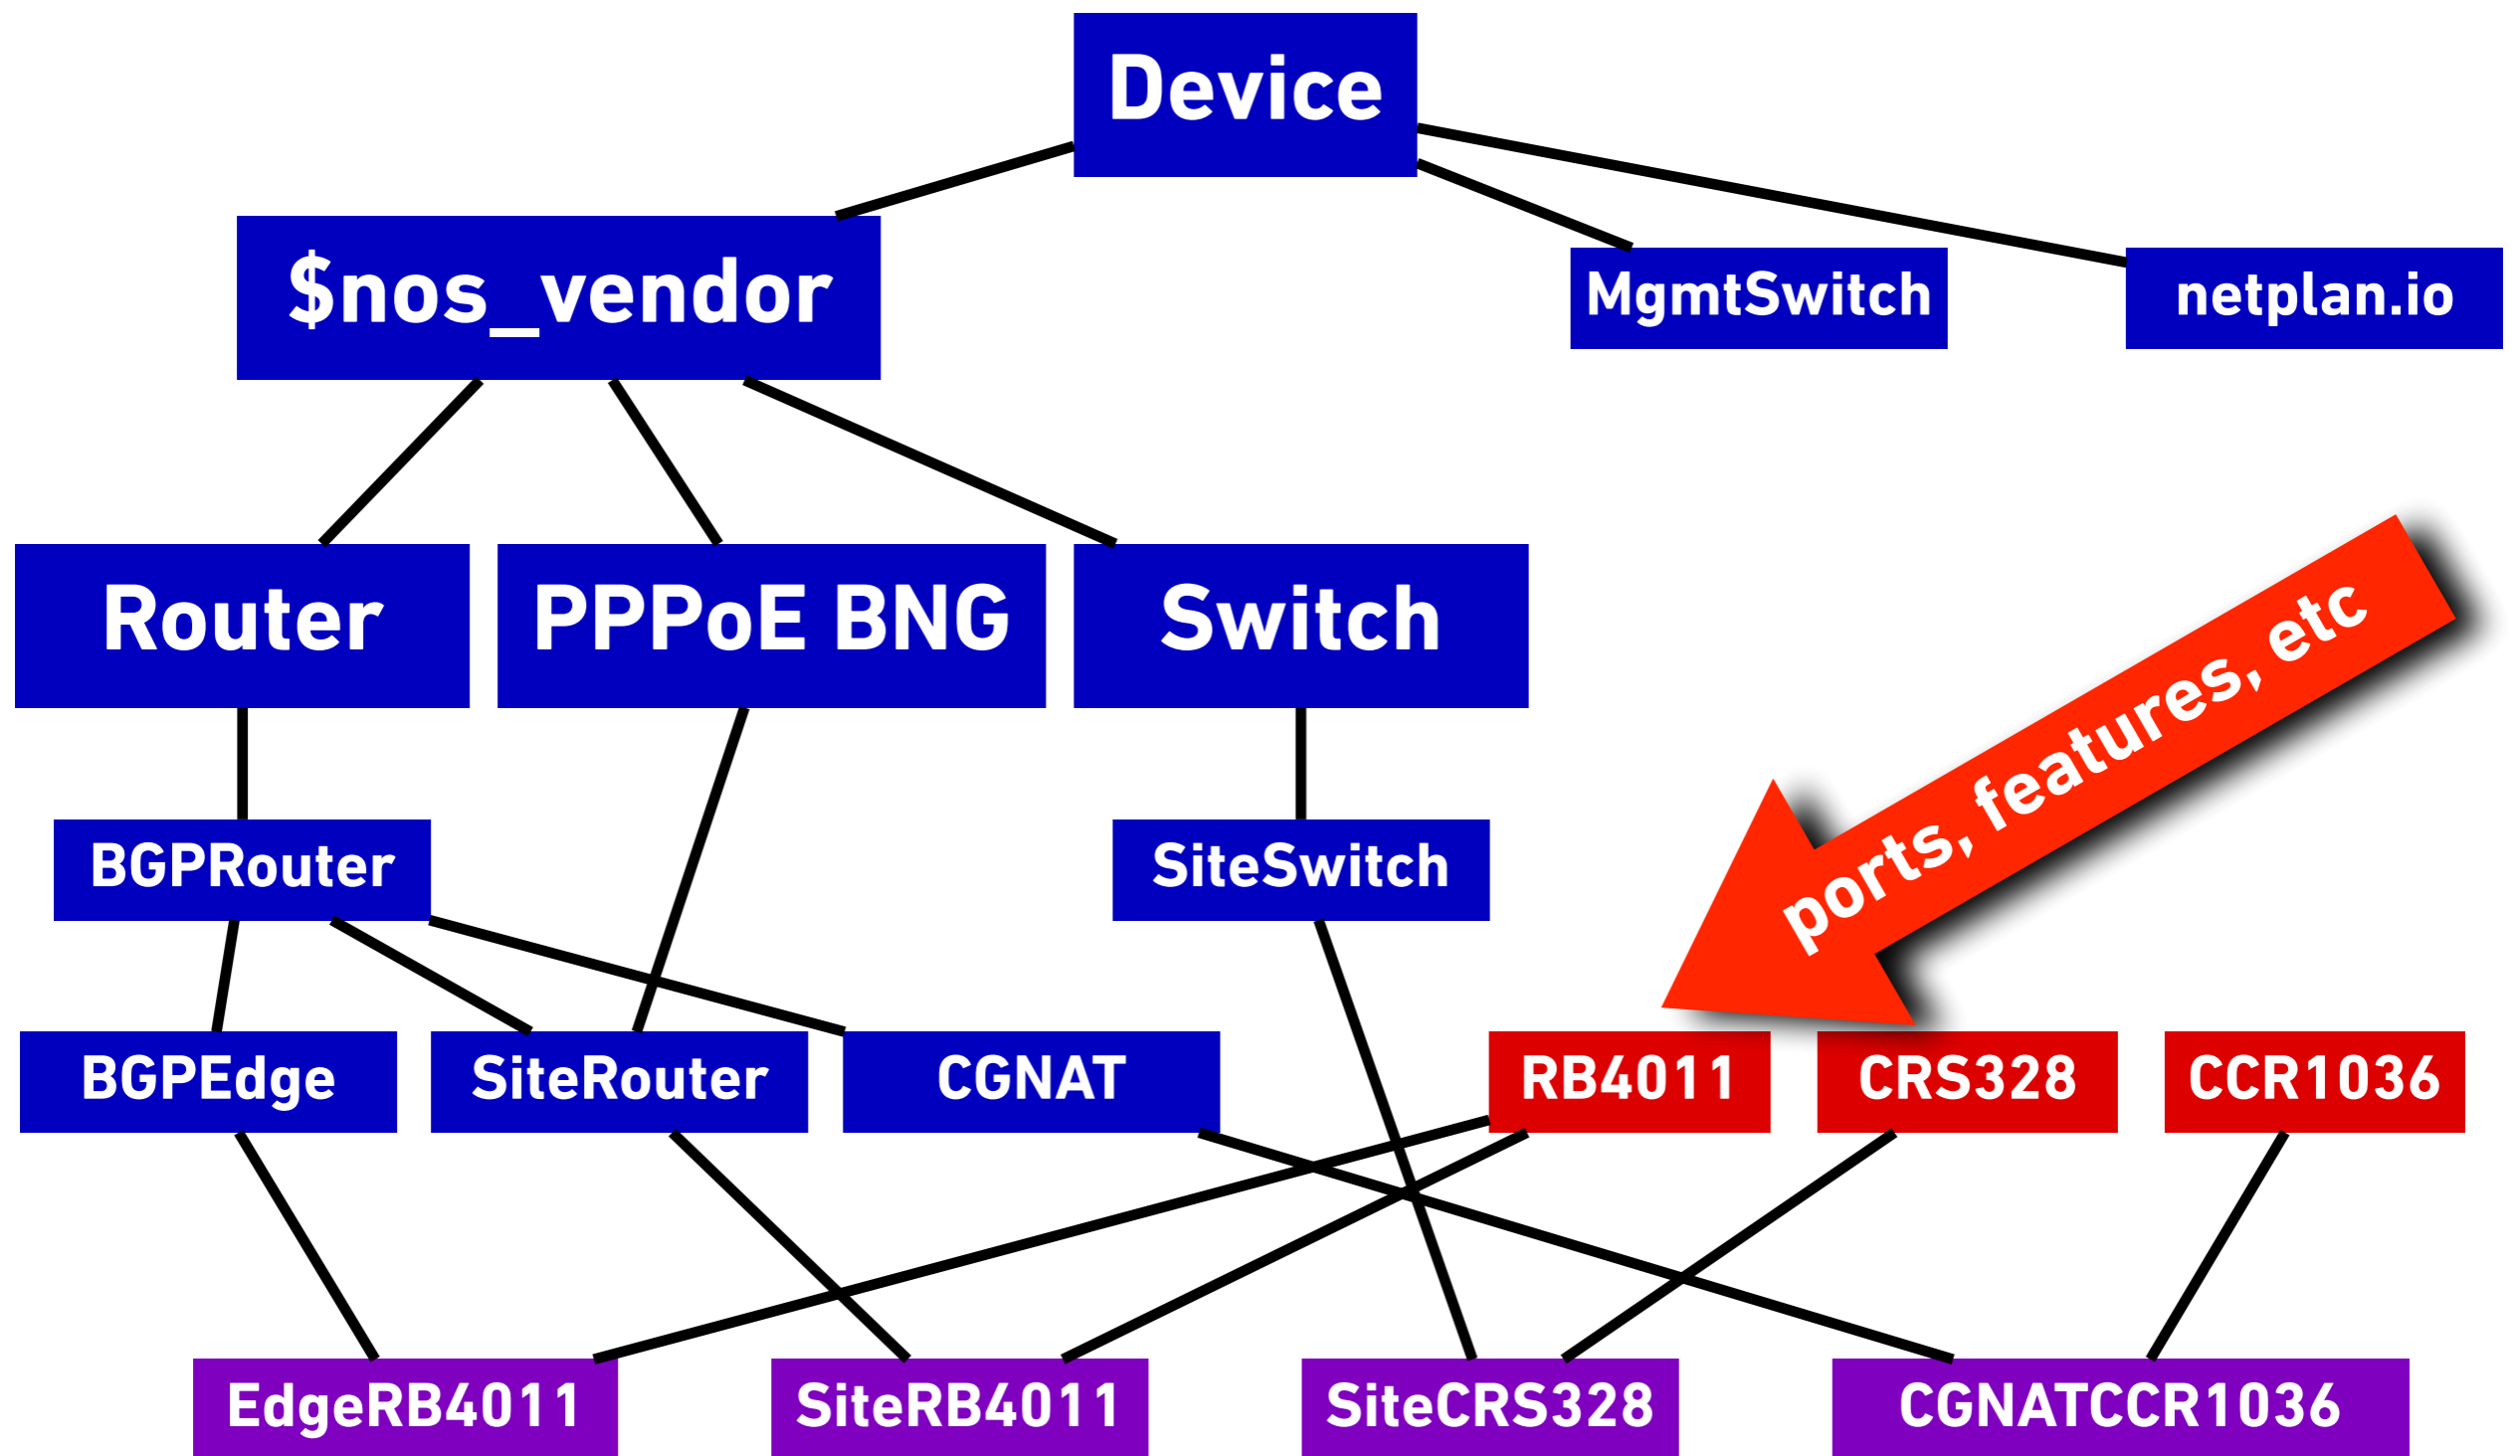
Just Add Code



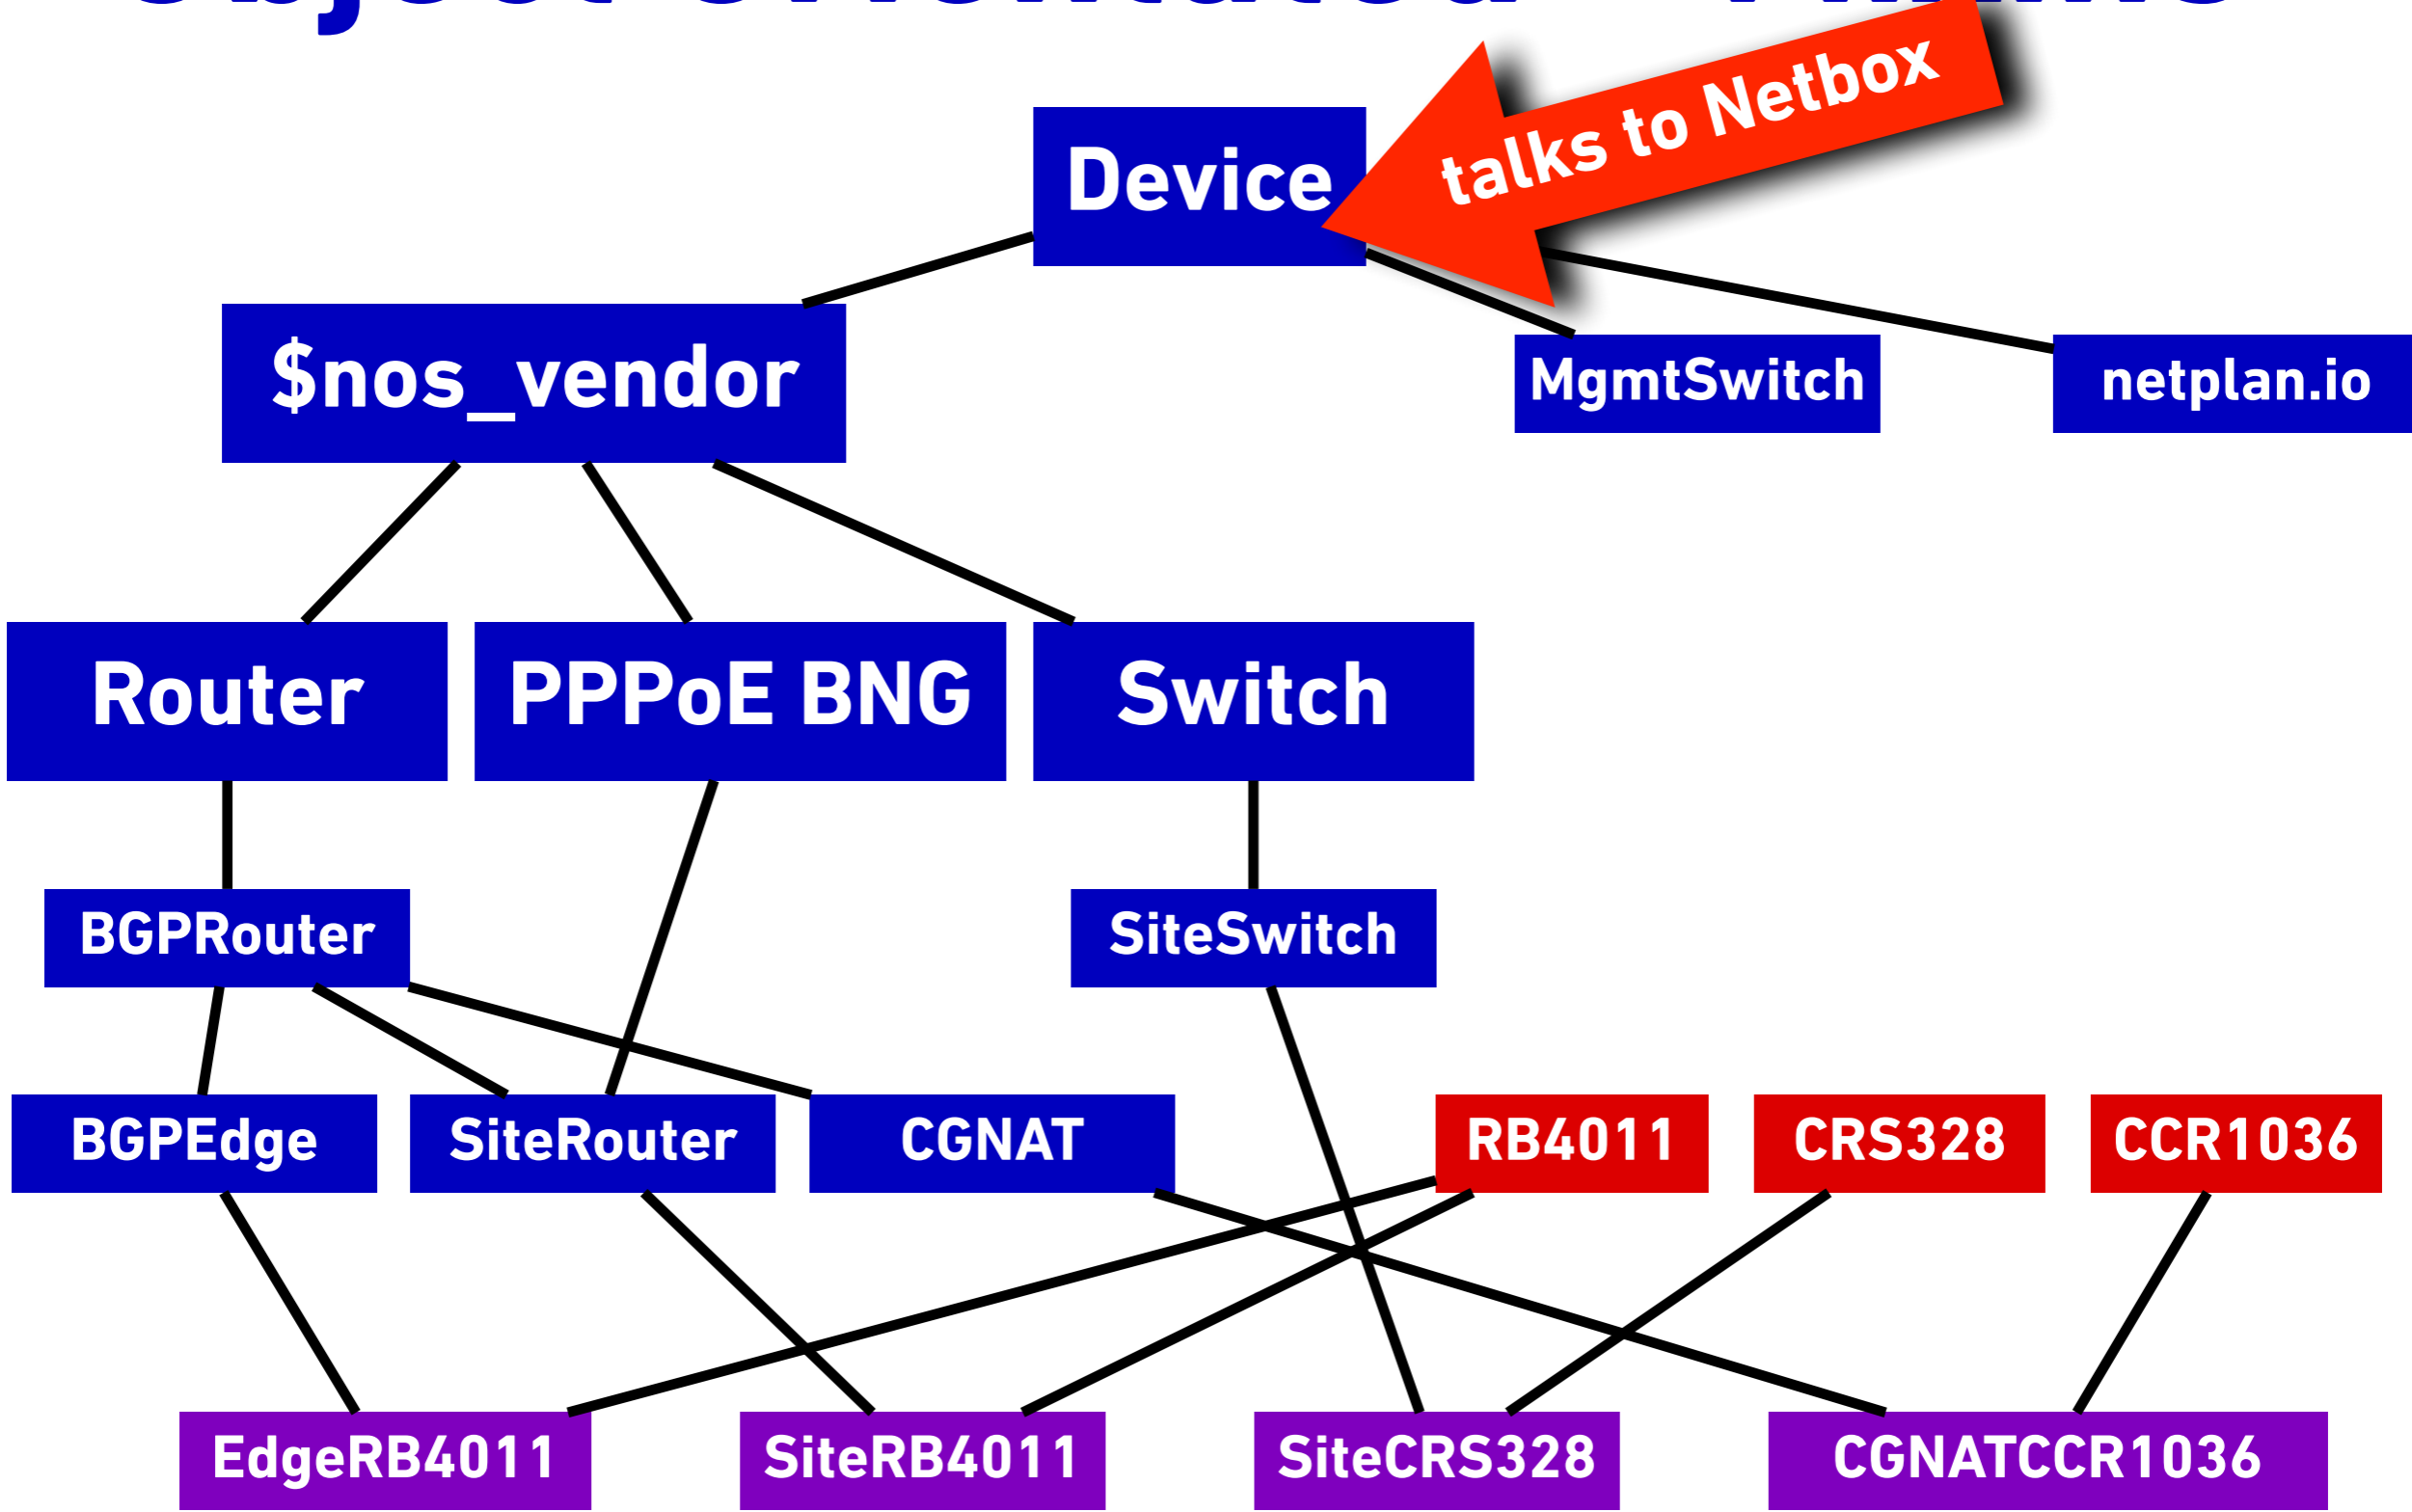
Object Orientated + Mixins



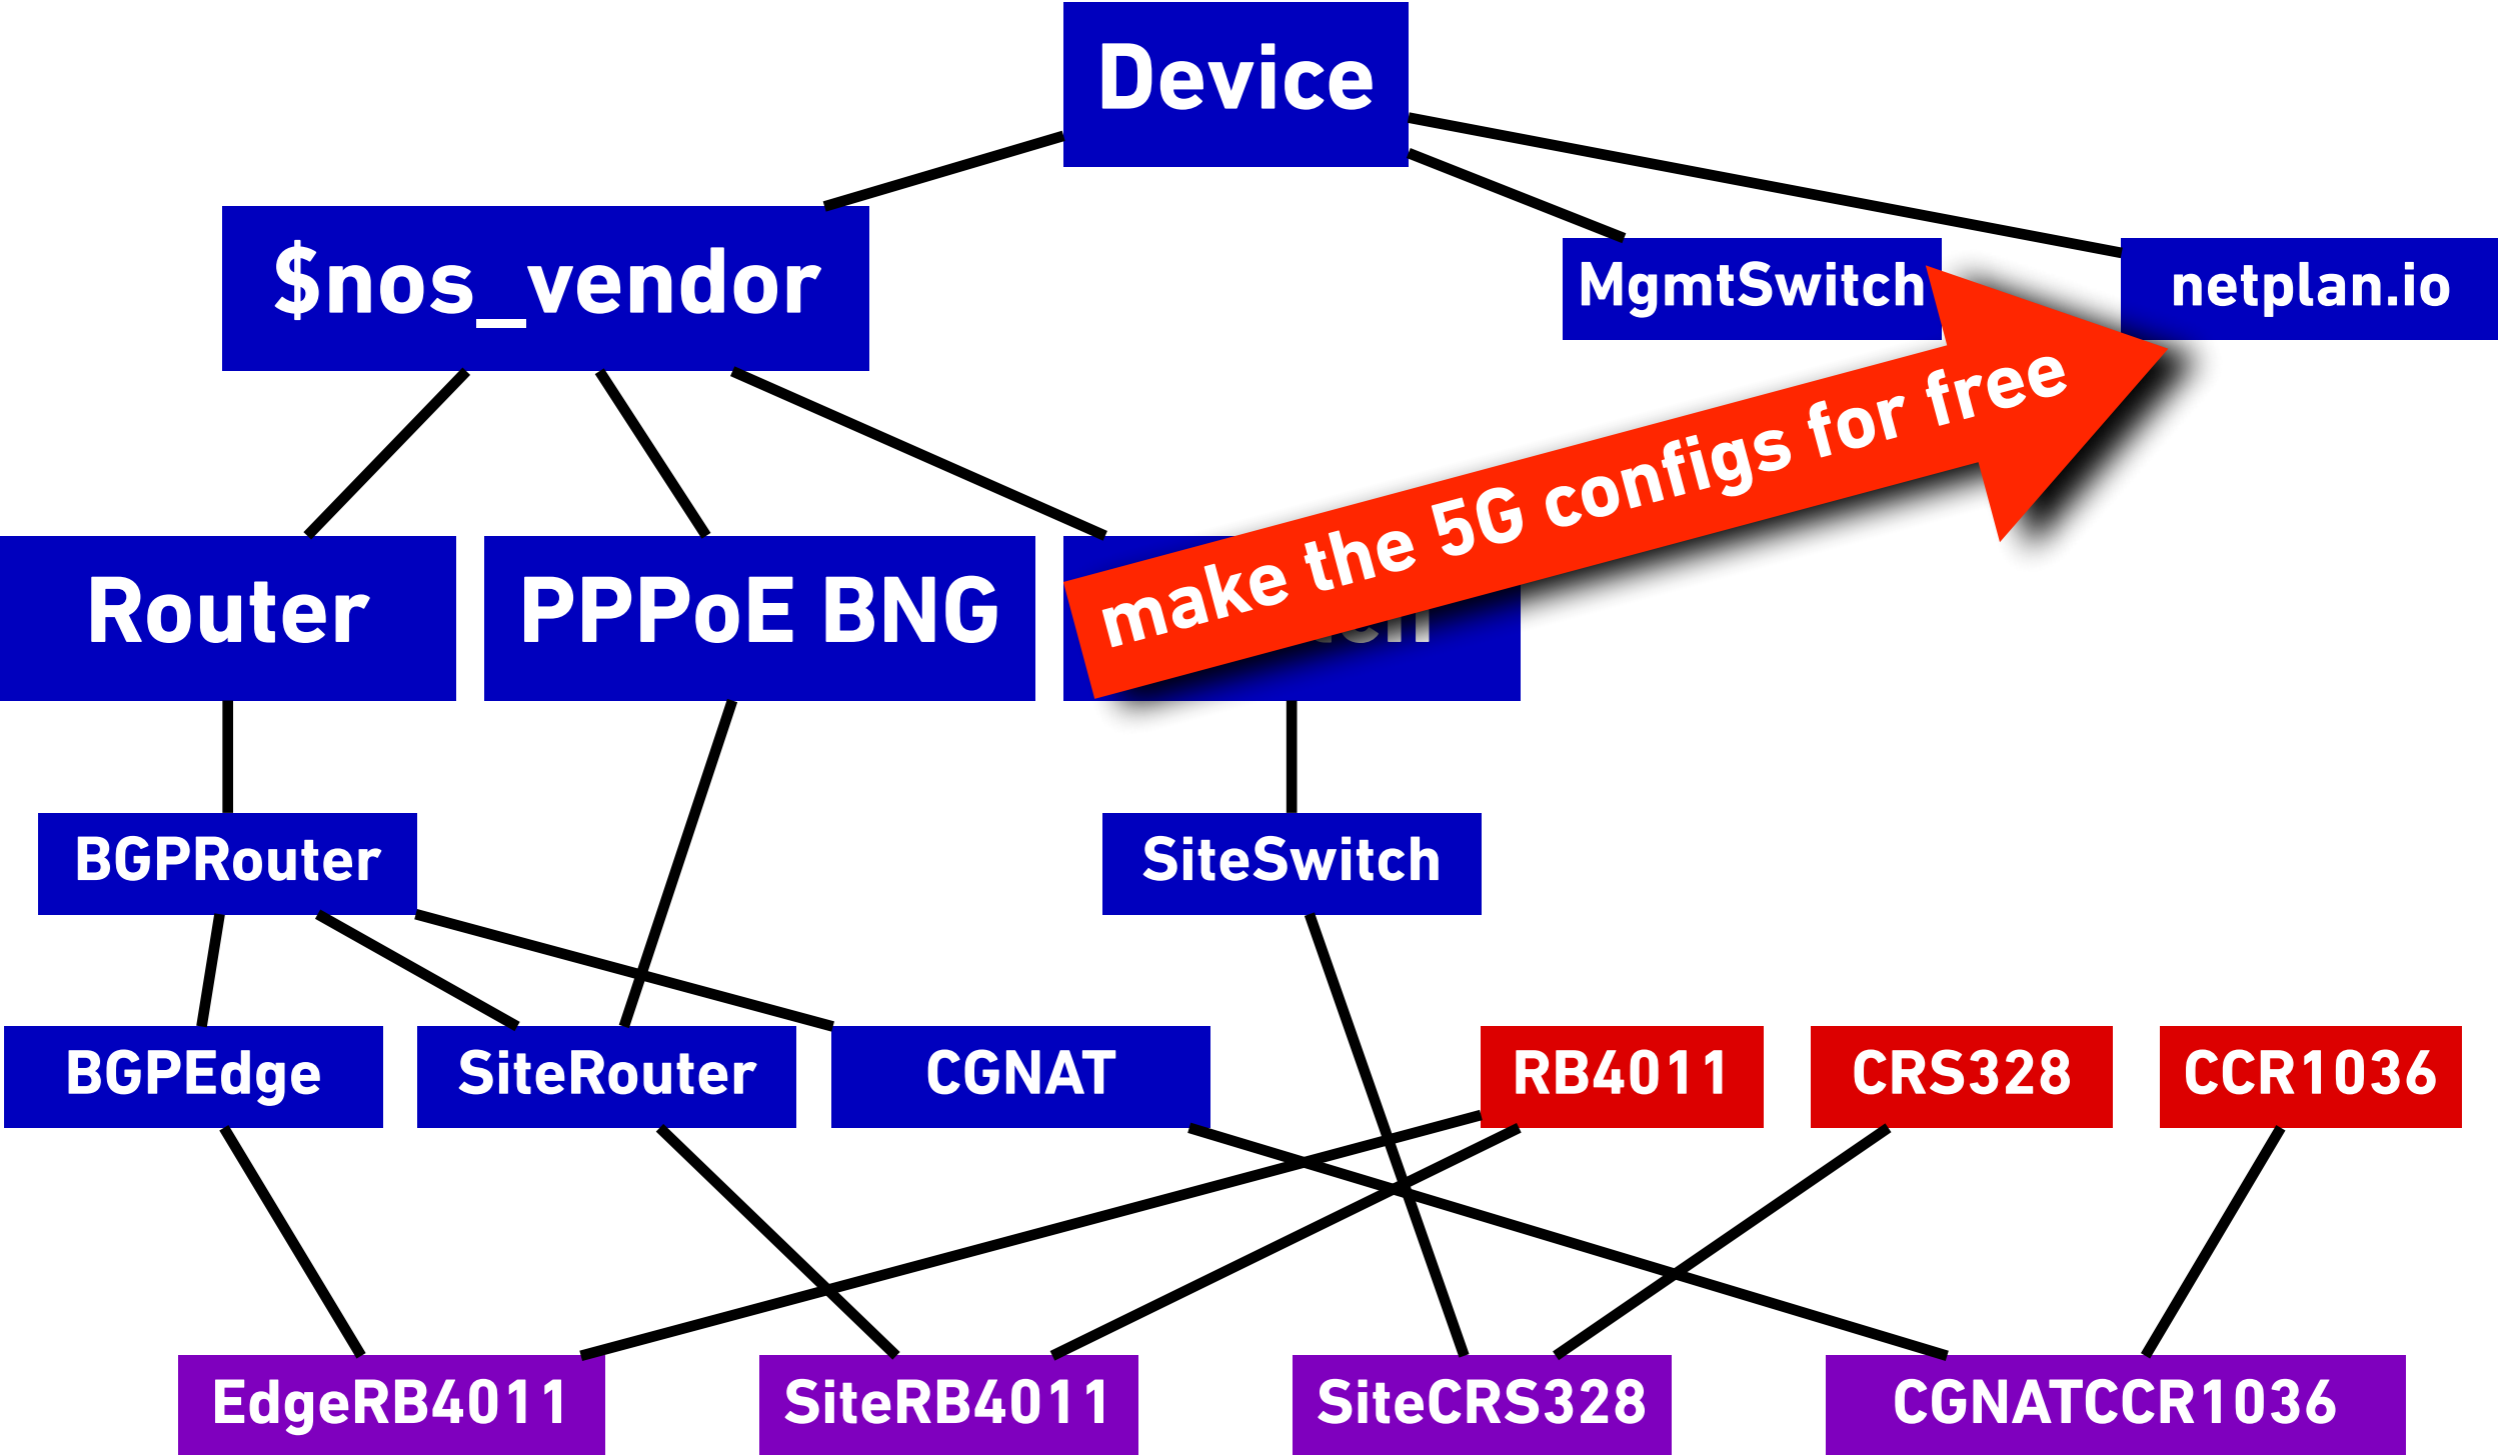
Object Orientated + Mixins




Object Orientated + Mixins



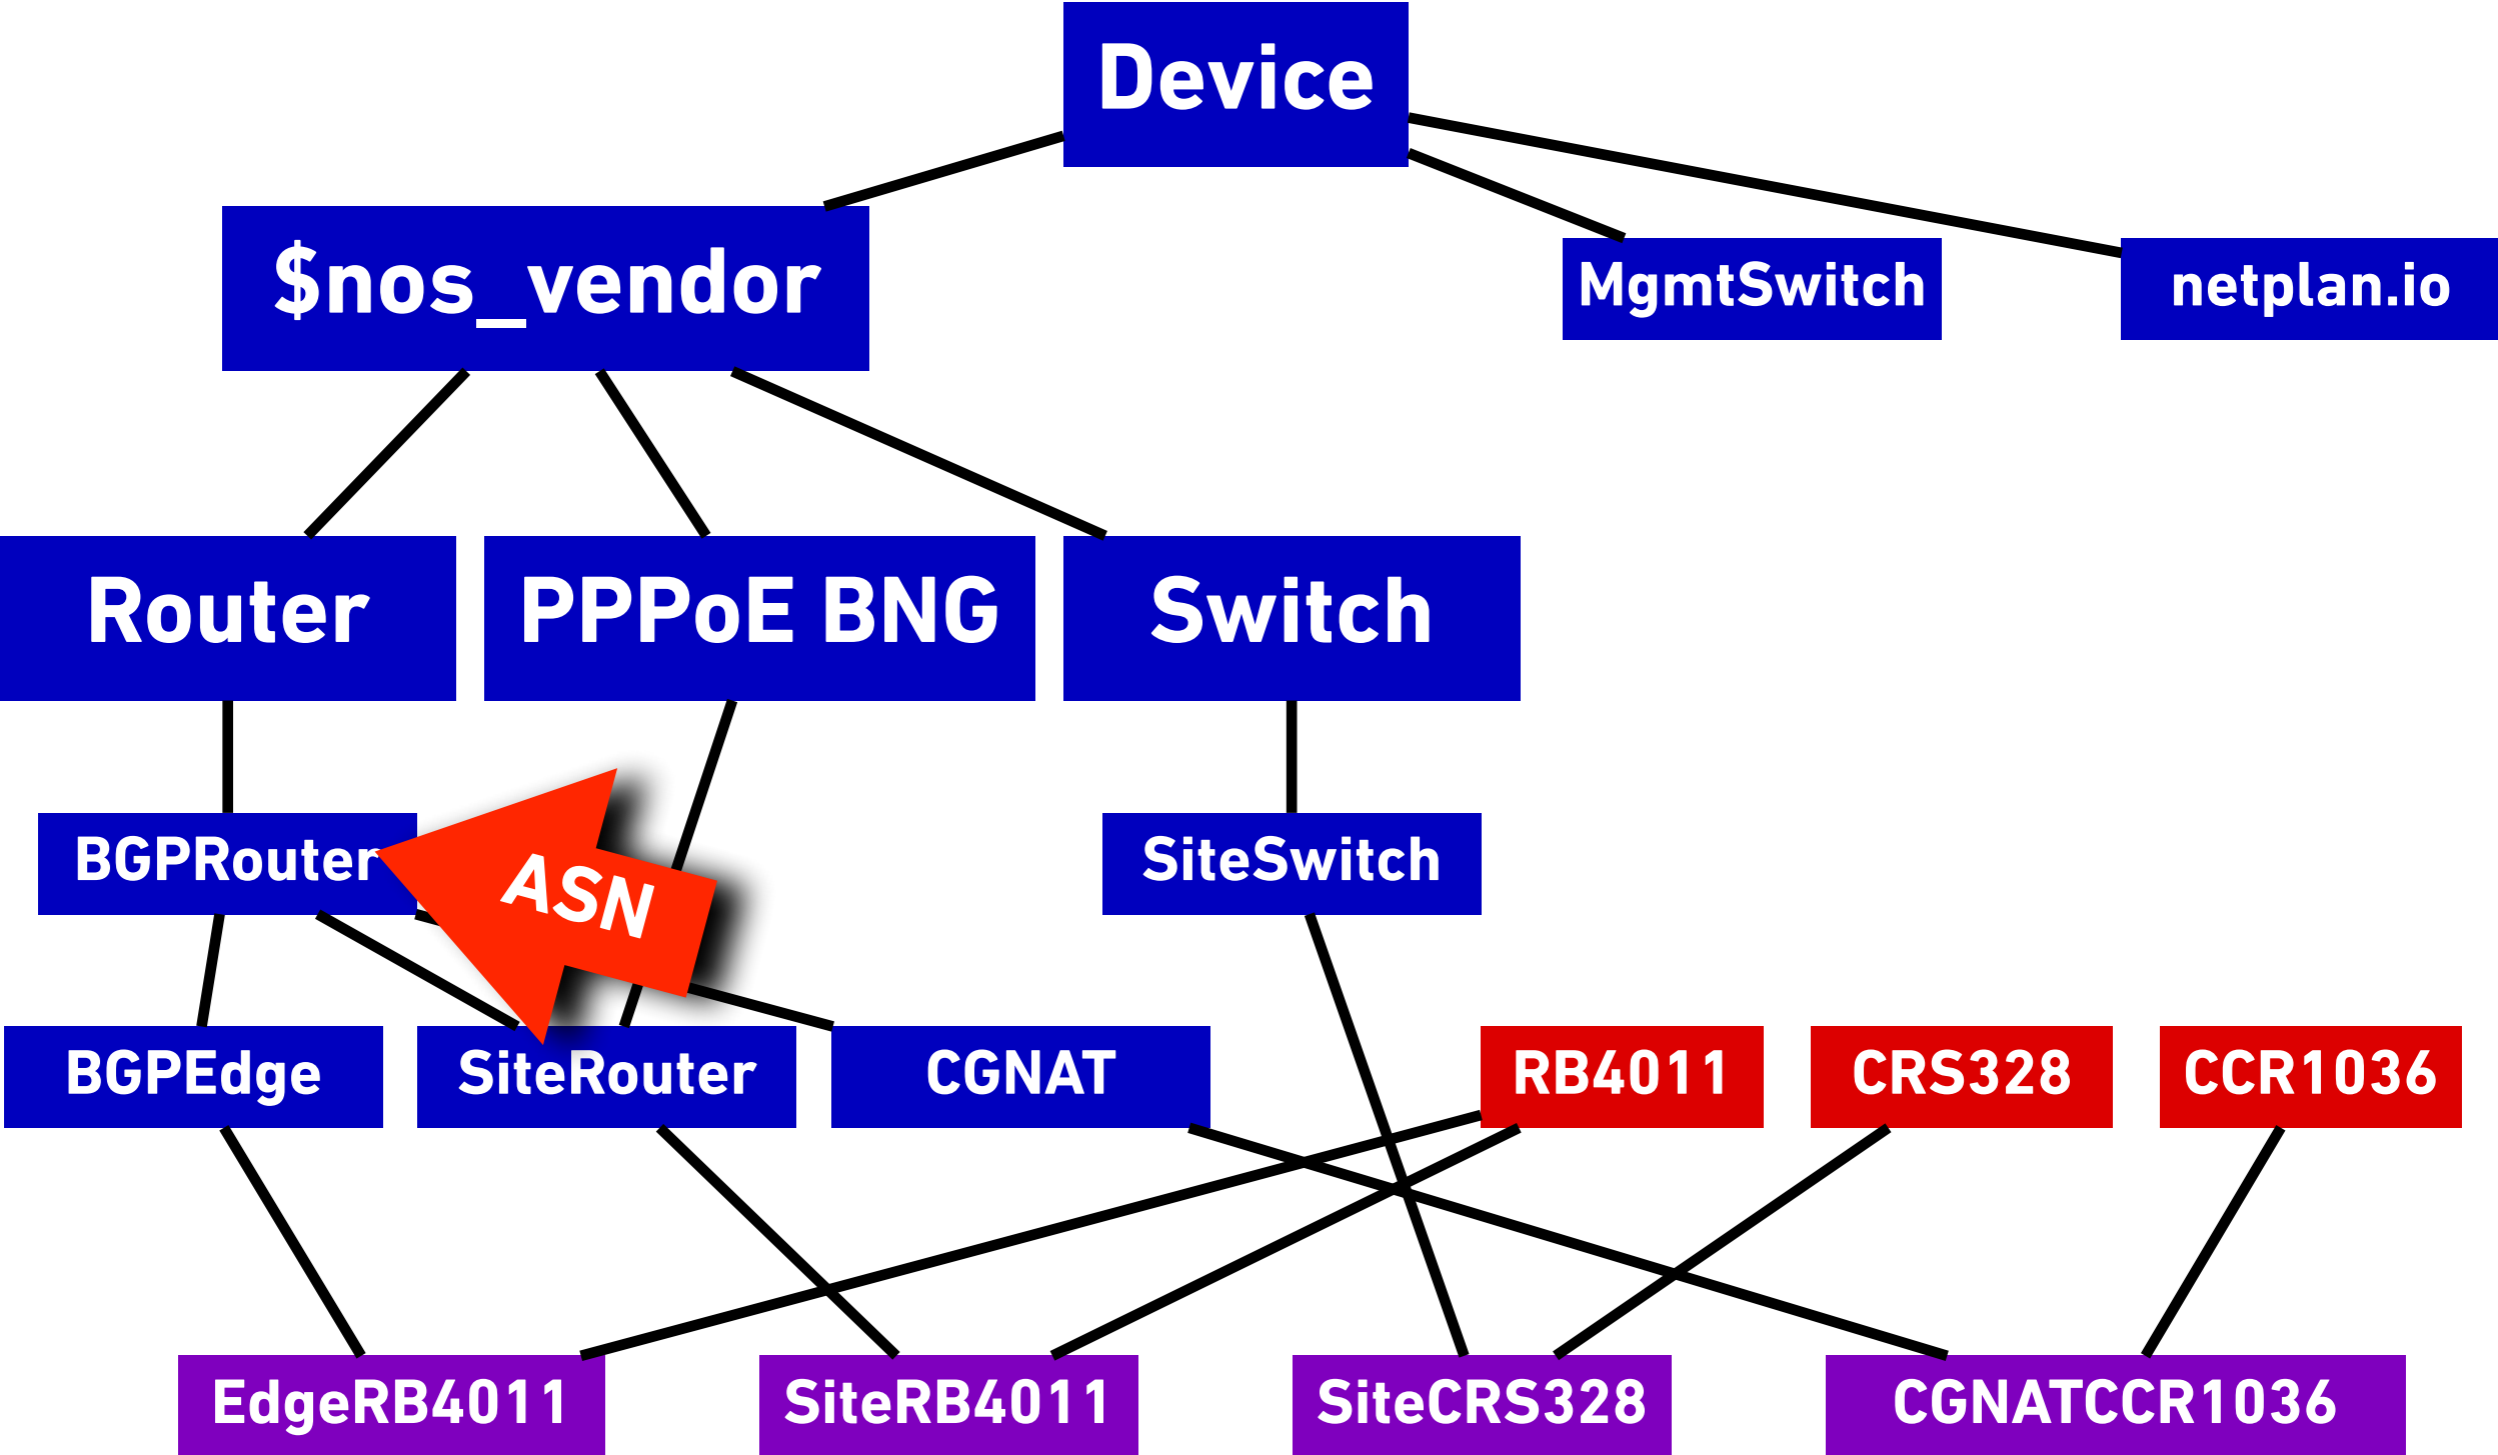
Object Orientated + Mixins



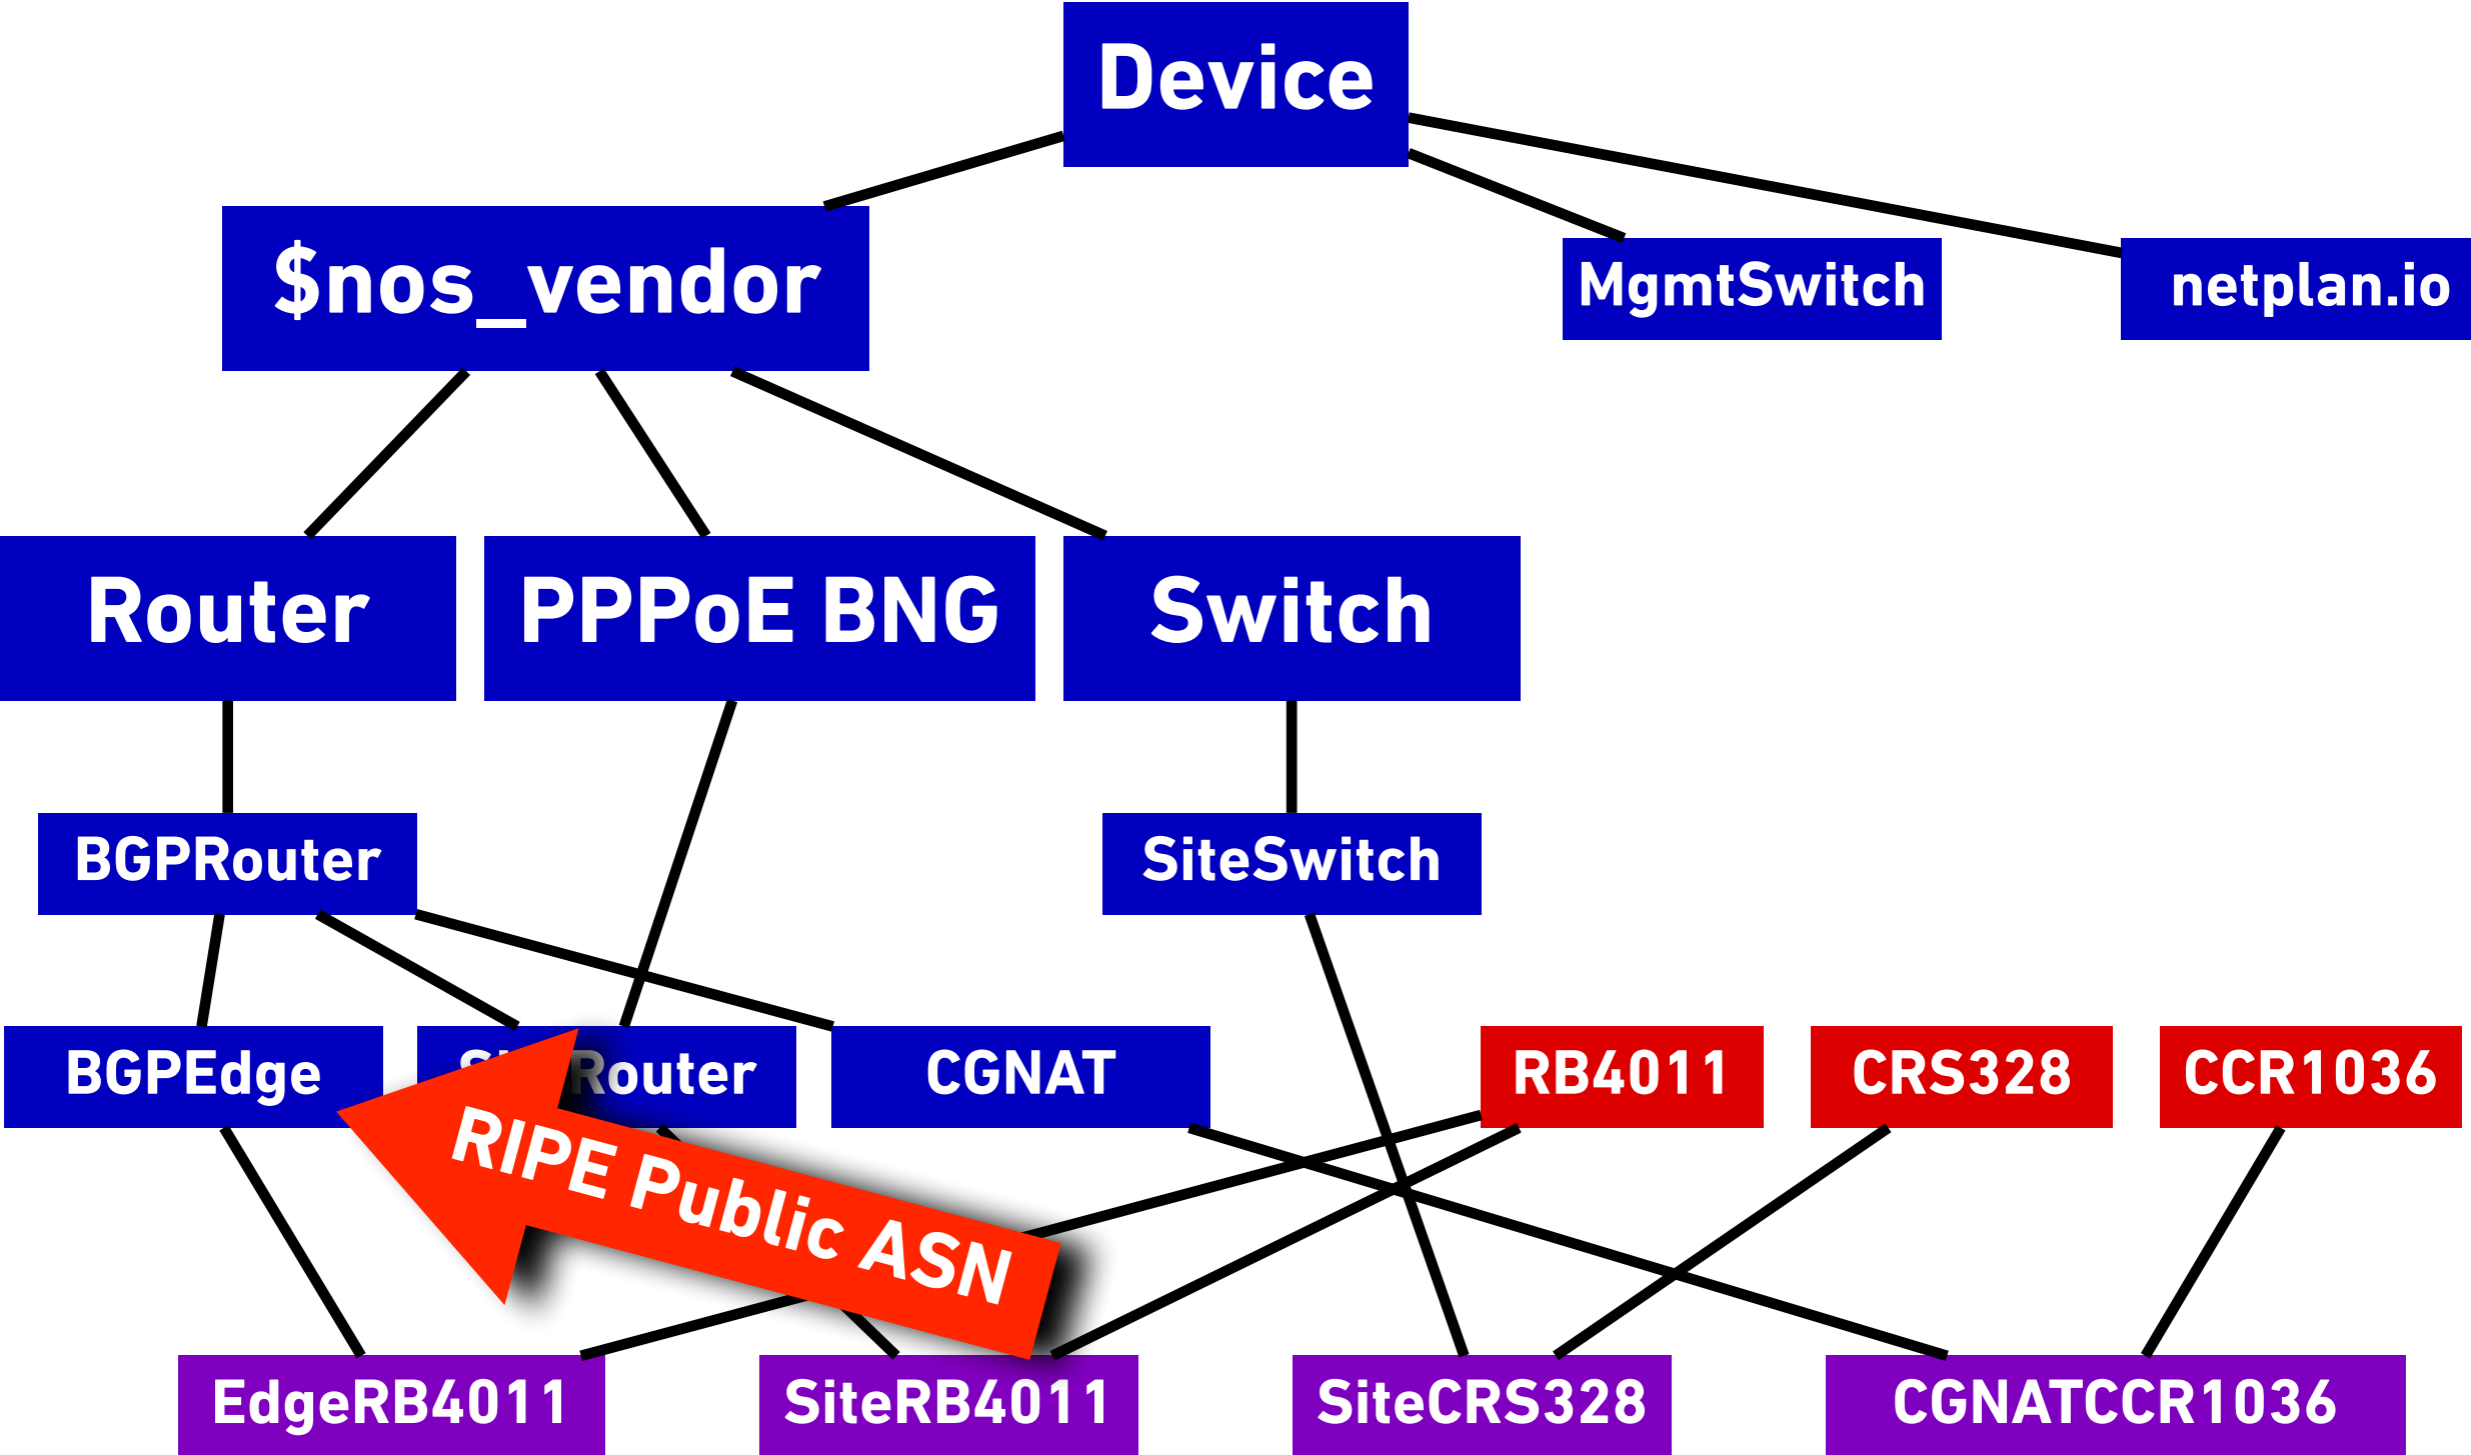
Object Orientated + Mixins



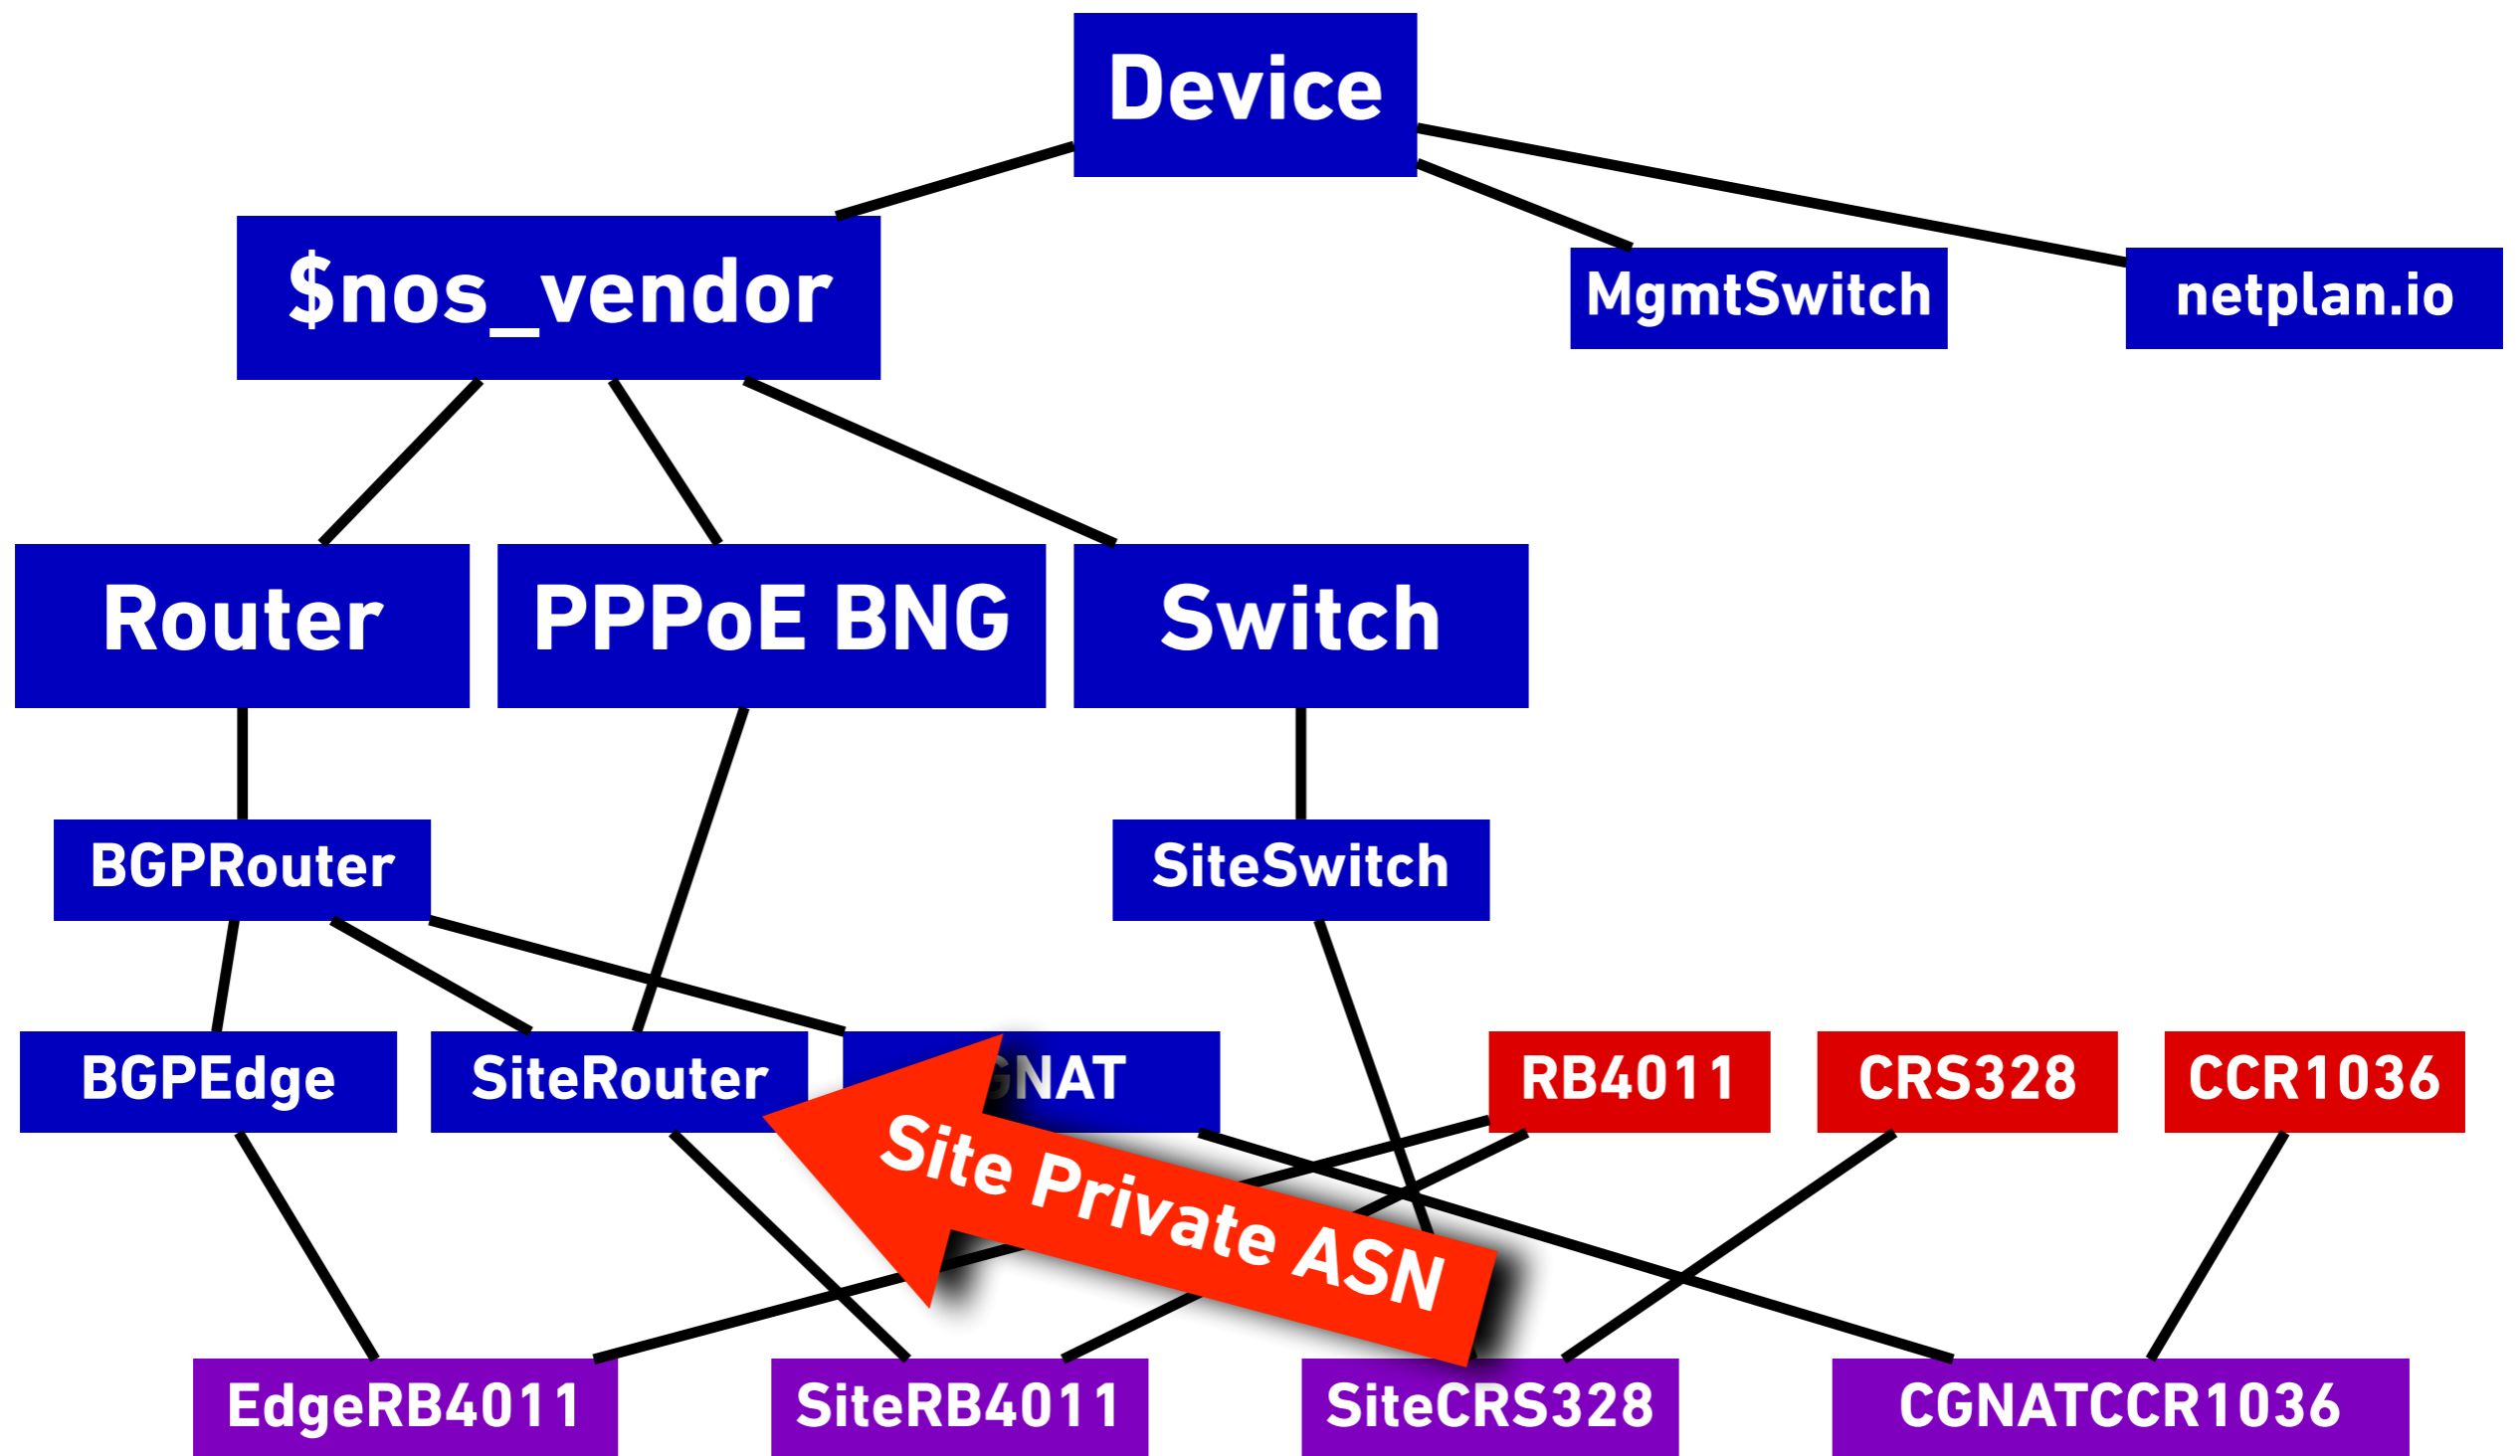
Object Orientated + Mixins



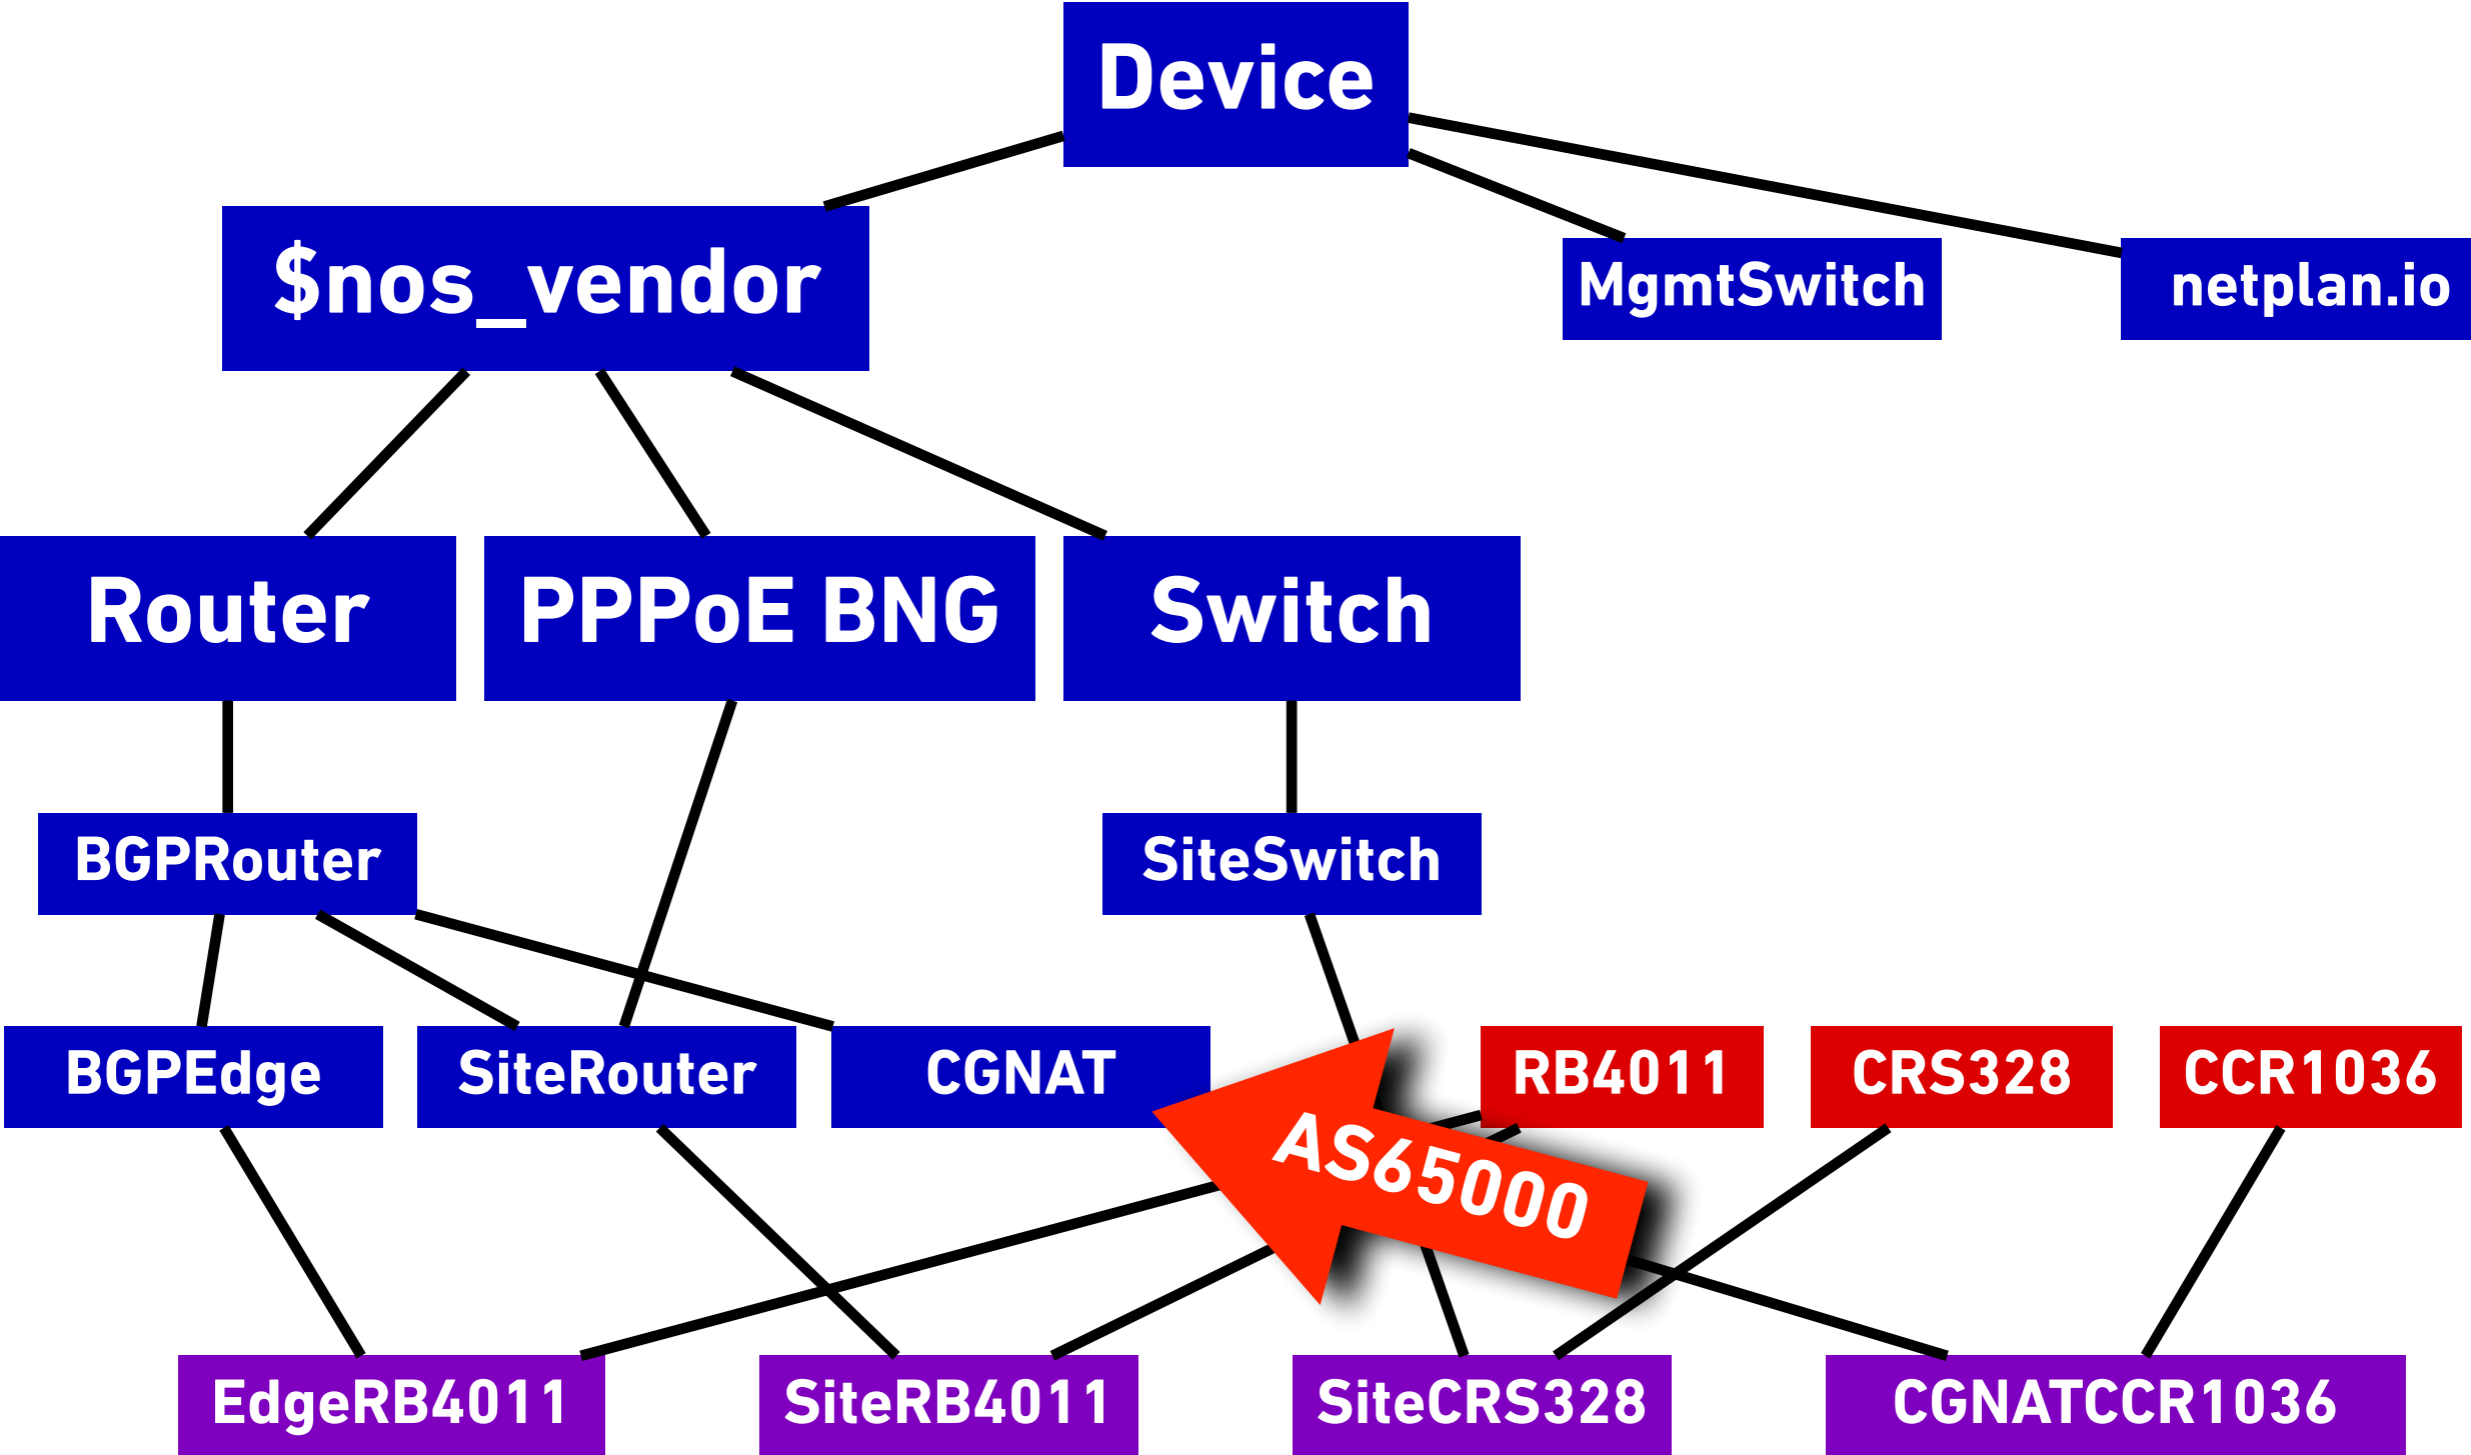
Object Orientated + Mixins



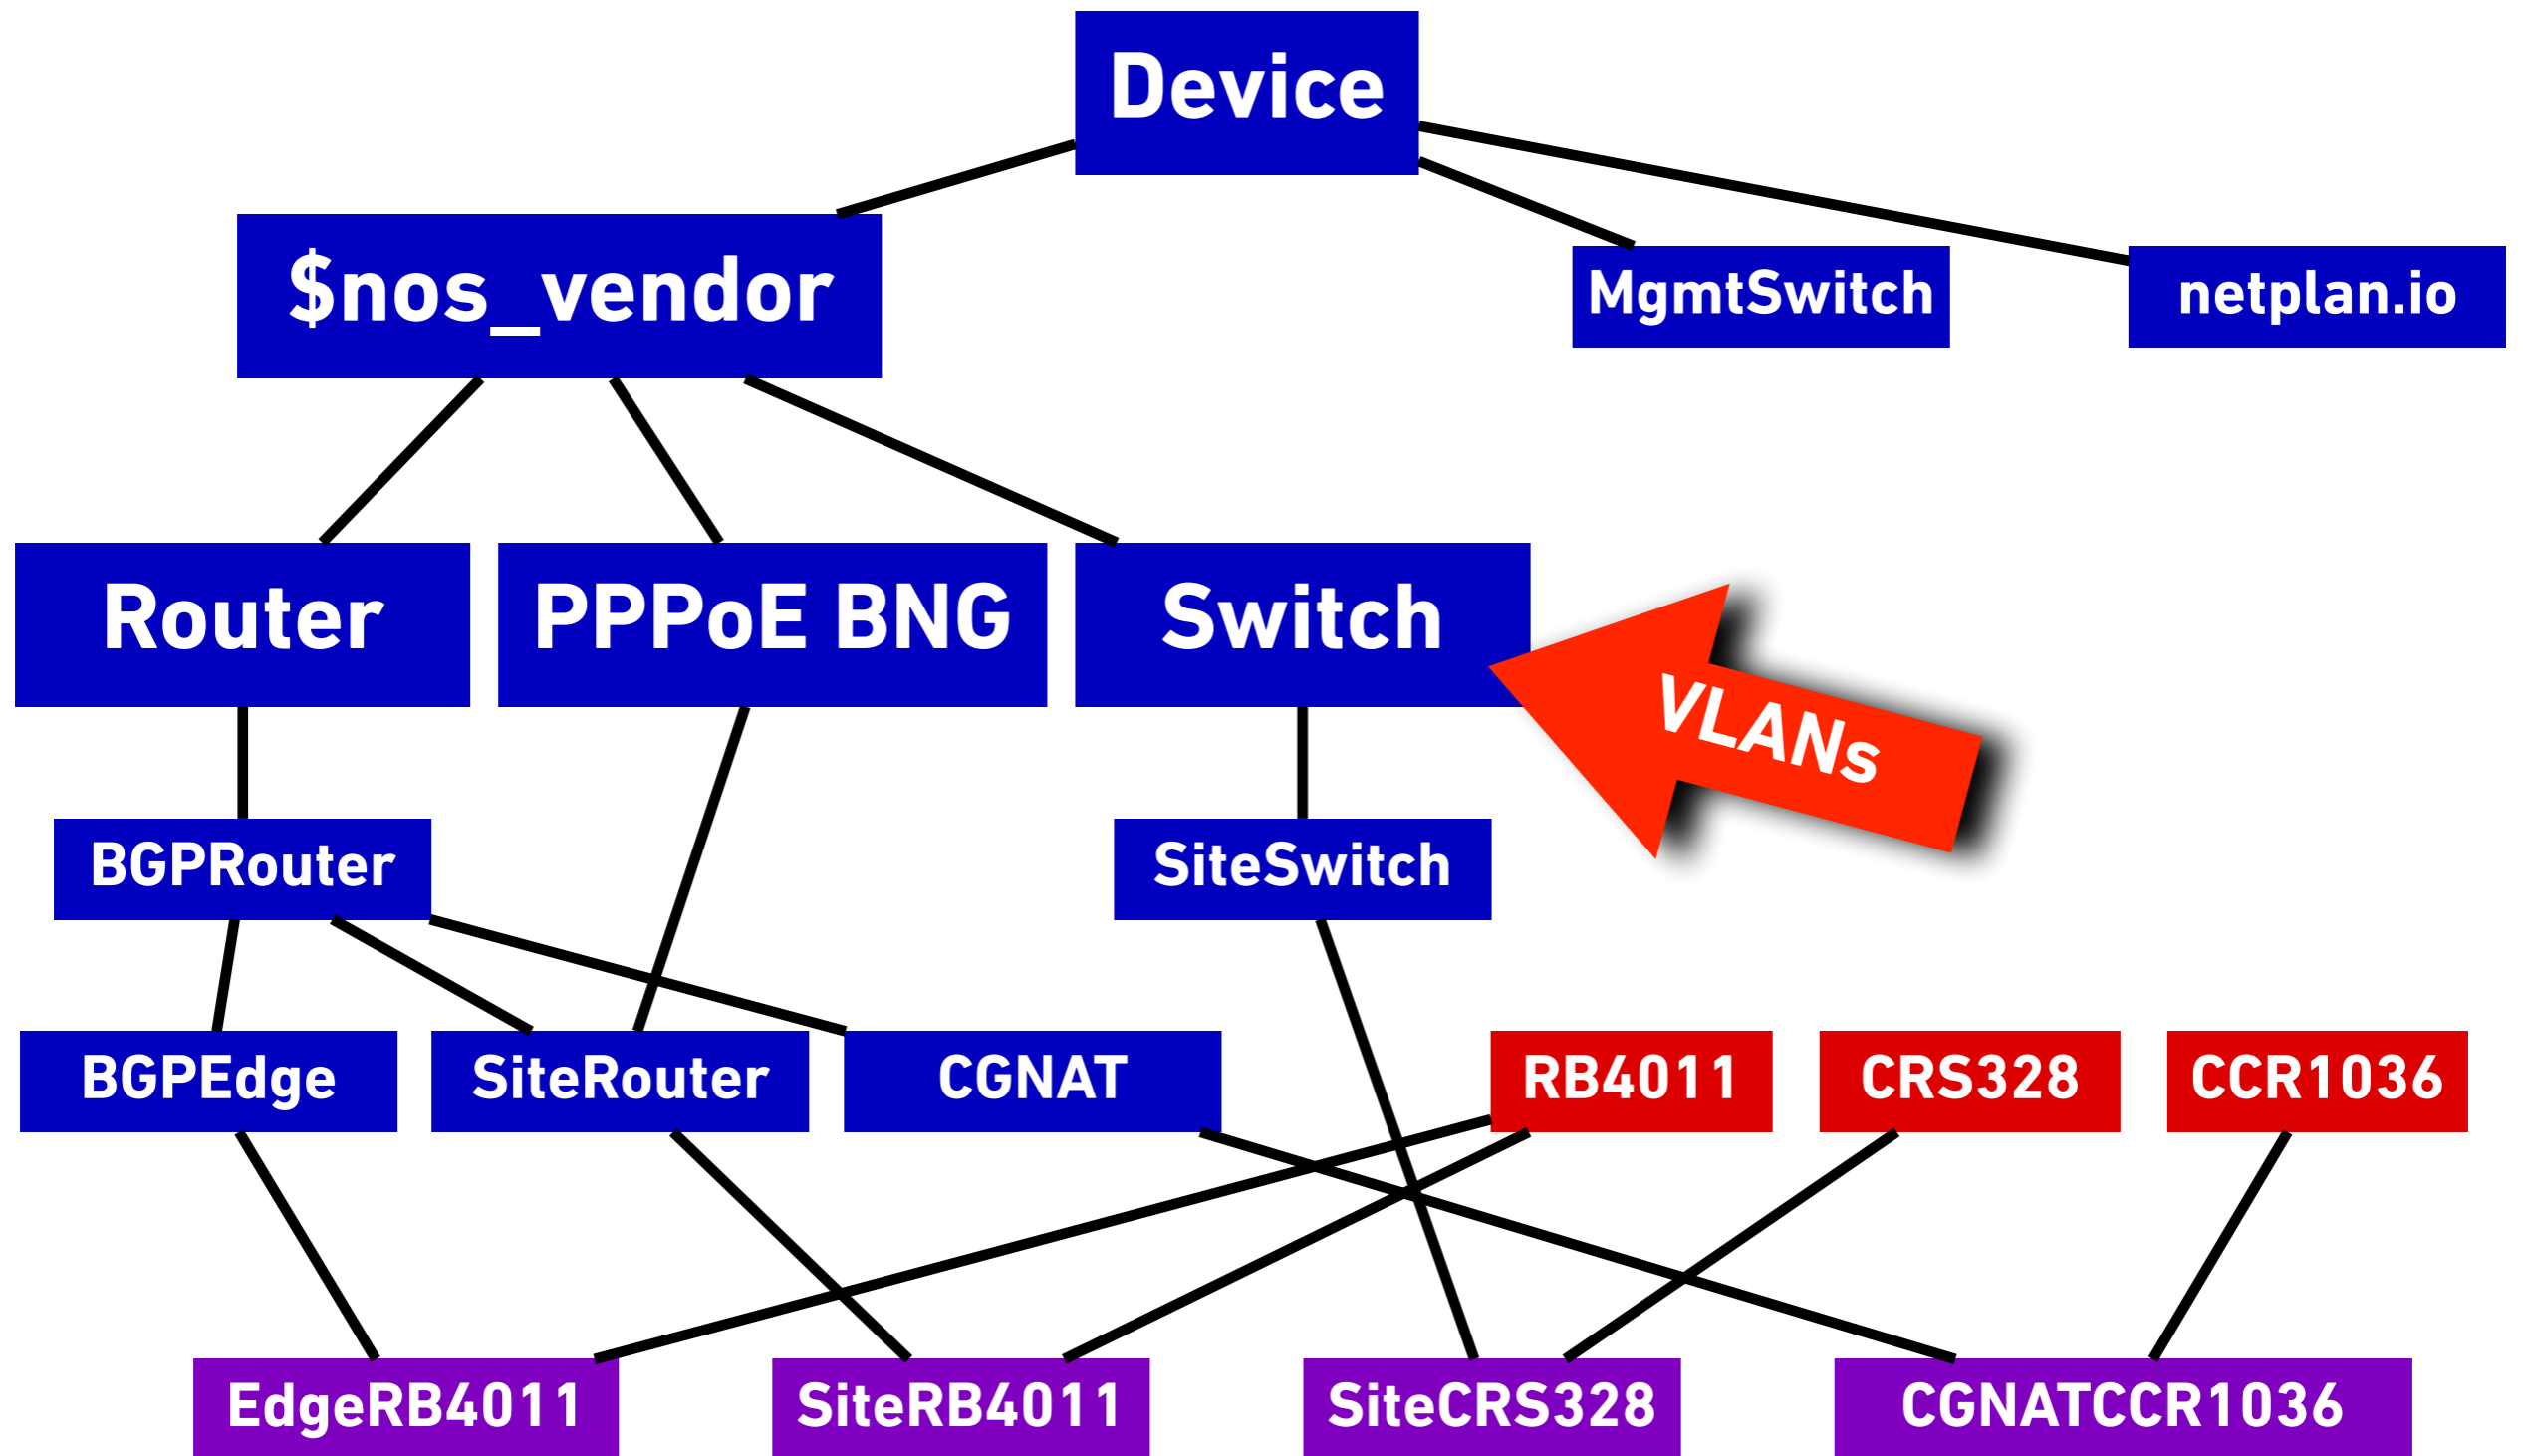
Object Orientated + Mixins



Object Orientated + Mixins



Object Orientated + Mixins



Semi-Automatic Approach

- ✘ Standardise hardware and network software
 - ✘ One device has one function
 - ✘ Eliminate special cases
- ✘ Build configs for all devices with software
 - ✘ Template output for everything
- ✘ /* TODO: how to amend all deployed configs? */



"iterated discovery process"

\$nos_vendor

Router

BNG

Switch

IPsec VPN

BGPRouter

SiteSwitch

ClusterCC

SiteCP

OOBRouter

CGNAT

BGPEdge

SiteRouter

MiniRouter

CCR1036

CRS328

hAP ac

RB4011

hEX PoE

hEX S

OOB/AP

CGN1036

Edge4011

Site4011

Mini

CPVPNC

Site328

CP

Templating In Action

❌ `python ./build.py`



❌ `ls -R configs`

configs/cc:

`cc.ayreofcara.sron.net.cloudnet.scot.conf cc.kirkbrae.west.net.cloudnet.scot.conf`

configs/cgn:

`cgn._future.net.cloudnet.scot.conf cgn.keelylang.main.net.cloudnet.scot.conf cgn.stove.main.net.cloudnet.scot.conf`
`cgn.ayreofcara.sron.net.cloudnet.scot.conf cgn.kirkbrae.west.net.cloudnet.scot.conf`

configs/cp:

`cp.ambulancestation.pwest.net.cloudnet.scot.conf cp.kirkbrae.west.net.cloudnet.scot.conf`
`cp.brekkaskail.net.cloudnet.scot.conf cp.northwalls.hoy.net.cloudnet.scot.conf`
`cp.clestrain.pwest.net.cloudnet.scot.conf cp.southwalls.hoy.net.cloudnet.scot.conf`
`cp.flottaoil.flot.net.cloudnet.scot.conf cp.woo.west.net.cloudnet.scot.conf`

configs/edge:

`edge._future.net.cloudnet.scot.conf edge.keelylang.main.net.cloudnet.scot.conf edge.stove.main.net.cloudnet.scot.conf`
`edge.ayreofcara.sron.net.cloudnet.scot.conf edge.kirkbrae.west.net.cloudnet.scot.conf`

L2 to L3 Migration

- ✘ Architected / HLD in May 2021 (pen and paper!)
- ✘ Low-Level Design in June 2021 (Netbox)
- ✘ Template-automated in July 2021 (Python)

- ✘ Full BOM, costed to the patch cable
- ✘ Deadline is October 2021
 - ✘ Feeling pretty pleased with ourselves!

DELAYS

Quadruple-Whammy

- ✘ Pandemic
 - ✘ Almost everything is harder, takes longer
- ✘ Brexit
 - ✘ Customs charges, import duty, returns to supplier
- ✘ Supply Chain
 - ✘ Not having e.g. PSU capacitors affects all vendors
- ✘ Orkney
 - ✘ Turns out that Orkney is Quite Far

“Space [...] is big. Really big. You just won't believe how vastly hugely mind-bogglingly big it is.”

– Douglas Adams, The Hitchhiker's Guide to the Galaxy

SCOTLAND

**is big. Really big. You just
won't believe how vastly hugely
mind-bogglingly big it is."**

– Douglas Adams, The Hitchhiker's Guide to the Galaxy



0 50 100 mi

FAROE ISLANDS

SWEDEN

NORWAY

ÅLAND ISLANDS

Orkney

Oslo

Stockholm

Baltic Sea

North Sea

DENMARK

Copenhagen

UNITED KINGDOM

Dublin

Hamburg

Berlin

NETHERLANDS



0 50 100 mi

FAROE ISLANDS

SWEDEN

NORWAY

ÅLAND ISLANDS

Oslo

Stockholm

Baltic Sea

Quite Far

North Sea

DENMARK

Copenhagen

LAND

Dublin

Hamburg

Berlin

NETHERLANDS


Almost as Northern as Stockholm




STANDARDISATION

No Pets Allowed!



- ✘ Same equipment at every site
 - ✘ Easier for sparing
 - ✘ Easier for automation/configuration
- ✘ Same cabling at every site
 - ✘ Production-line process for pre-build
- ✘ But not every site is identical?
 - ✘ LLD is a component-based template
 - ✘ What moving parts do we need at this site?



The Most Complicated POP Site




site router




site switch


site switch




IPv4 CGNAT




BGP edge router







5G-Control-Plane VPNC




5G-User-Plane (UPF)



5G-RAN Node (eNB/gNB)






point-to-point link PSUs



OOB/management

Variations on the Theme



Pre-Build, Take Everything

- ✘ Some islands' ferry services are twice per week
 - ✘ Winter timetable has fewer sailings
 - ✘ Bad weather/tide cancellations; also covid
 - ✘ Can send parcels, people, vehicles
 - ✘ Don't be late for your ferry home
- ✘ Islands also served by LoganAir
 - ✘ Popular with tourists (thanks, [Tom Scott](#))
 - ✘ Limited capacity for people and parcels


Unbox and Configure



Build and Box



Less Time Cabling



More Time Hoisting!



`/* TODO */`

Story Isn't Over...

- ✘ New network architecture is mostly deployed
 - ✘ Still some remaining sites being migrated over
 - ✘ A couple of "gotchas" were identified during build
- ✘ How to apply future "patches" to configs?
 - ✘ Vendor's NOS has API; store metadata in config
 - ✘ Query, amend, update version comment/MOTD

Results

- ❌ Customer performance significantly improved
 - ❌ e.g. one customer went from 6 to 150Mbit/sec
 - ❌ (only <1Mbit/sec available from Openreach infra)
- ❌ Network manageability should be simplified
 - ❌ Customer public/CGN IP allocations from RADIUS
- ❌ Loop-free via L3 routing rather than STP
 - ❌ Improved resilience for rain/tide fade, failures

The Team



ASK ME HOW TO SNEAK A MOOSE INTO A DCMS REPORT

E: [marek @ faelix . net](mailto:marek@faelix.net)

T: [@faelix](https://twitter.com/faelix)

W: <https://faelix.net/>

<https://faelix.link/uknof48>

