Building a

Mostly **Open-Source**

Tier 1 - Carrier Voice Network

NETSIP

Aussie Broadband

NETSIP
connect.enable

Over the Wire

Jay Binks

CTO - Collaboration Products

Aussie
Broadband

Why ?

- Over The Wire was providing hosted PBX services
- At the time, many of the wholesale SIP providers (Including big carriers) had less than reliable IP voice platforms
- Why not, NetSIP already had its first handful of customers

What we built

We built what other carriers couldn’t provide us at the time.

A Reliable network

We had grown tired of :

- Outbound calls failing for seemingly no reason
- Inbound numbers, being unreachable
- IP DDoS’s in a carrier network, affecting our private voice services
- Carriers transcoding calls to low bandwidth codecs to save bandwidth in their network

What we built

We built what other carriers couldn’t provide us at the time.

Usable customer portals & API’s

We gave our customers the ability to Self-Serve :

- Customers could order numbers from our inventory in seconds
- Place number ports without signing paper
- Route numbers to the SIP device of their choosing
- Route calls to other offices or to their mobile if their SIP trunk was down

What we built

We built what other carriers couldn’t provide us at the time.

Peace of mind & Protection

Every time we saw a toll fraud incident we added more protection to our system.

- GeoIP restrict SIP traffic to Australia, by default
- Only allow a small number of concurrent international calls by default
- Allowed customers managed ACL’s for every SIP endpoint
- Block known bad UserAgents (*I know… but it works*)
- Suspend customer accounts instead of billing them tens of thousands of dollars

What we built

We built what other carriers couldn't provide us at the time.

Hid the ugliness from our customers

- We built automated systems to test a port at completion, from every mobile network and fixed carrier
- We simplified the porting process for CSP's
  with overarching porting agreements
- We Actively expand our API's based on customer feedback.

NETSIP
connect.enable

We embraced Open Source
from the beginning

Aussie Broadband

Freedom vs Free

We started using Open Source for the cost saving

It definitely was lower cost, at the beginning.

We wanted to avoid paying $50k for a commercial softswitch
before we had supporting revenue

Freedom vs Free

- We have committed patches to open source projects in our stack
- We have built custom modules on top of open source projects
- Having access to the source code enables us to better understand strange behaviors and better manage our platform.

Freedom vs Free

We still use Open Source because

There is no compelling reason to change !

- Great reliability
- No scaling issues to date

  (stress testing indicates no we have plenty of future capacity)
- We were able to solve our problems without vendor support

KAMAILIO

FreeSWITCH

What does it look like ?

php

GO

MariaDB

MEMCACHED

NGINX

squire

sip:wise
an ALE Company
RTPEngine

VOIP MONITOR

debian
GNU/Linux

PROXMOX

ANSIBLE

puppet

Where does all this run?

Everyone else is building service in public cloud, we chose not to.

Our NetSIP voice network consists of :

- 5 Data centers, with customer facing equipment
  (currently : 2xBNE, 2xSYD, 1xMEL)
- 10+ Carrier aggregation racks around the country
  (2 in each major capital city)
- Mostly NextDC sites, as well as Equinix
  (there are a few smaller DC operators where we had specific location requirements. EG ADL)

Avoid single points of failure

While avoiding complexity

- Dual top of rack L3 switches
  (full mesh multiple 10gig uplinks into multiple OTW core PE’s)
- SIP Load Balancers *(2 x with VRRP IP failover)*
- Customer facing SBC’s (3+)
- Virtualisation hosts (2+)
- Warm spare servers in each rack

Avoid single points of failure

While avoiding complexity

- Dual top of rack L3 switches
  (full mesh multiple 10gig uplinks into multiple OTW core PE's)
- SIP Load Balancers *(2 x with VRRP IP failover)*
- Customer facing SBC's (3+)
- Virtualisation hosts (2+)
- Warm spare servers in each rack

NETSIP
connect.enable

Avoid single points of failure

While avoiding complexity

- Dual top of rack L3 switches
  (full mesh multiple 10gig uplinks into multiple OTW core PE's)
- SIP Load Balancers *(2 x with VRRP IP failover)*
- Customer facing SBC's (3+)
- Virtualisation hosts (2+)
- Warm spare servers in each rack

Aussie Broadband

Assume everything fails,
and work backwards

Engineer things around how they fail

We made design decisions based failure modes

- Prioritise reliability *over profitability*
  (A failures should still allow customers to use our service, even if we can’t bill them for the traffic)

- Prioritise simplicity *over perfection*
  (We chose not to build live call failover between SBC’s, even though it would have been easy to do)

NETSIP
connect.enable

OK, Who cares…

How did you become a Tier 1 carrier ?

Aussie
Broadband

On to the Tier 1 Carrier stuff

Some time around 2012, I said to my team :

"we are going to build a Tier 1 carrier interconnect".

At the time, our entire company probably only had 20 something staff.

Everyone laughed & thought i was crazy, Including my founding business partners :)

In my mind, this is where I knew the direction I was headed.

Tier 1 - Timeline

2012

Started ingesting routing data.

Tier 1 - Timeline

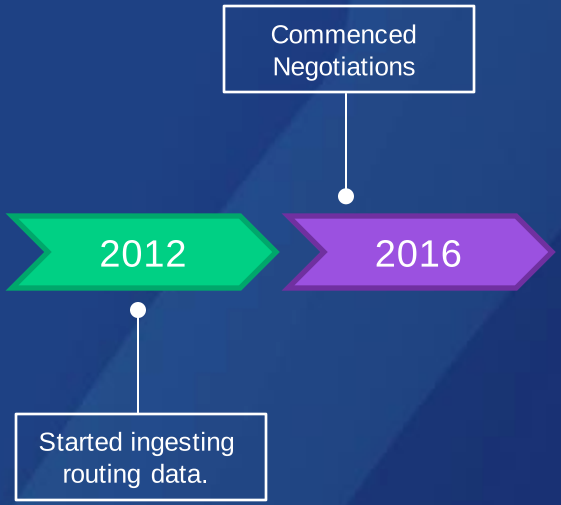

2012

Started ingesting routing data.

2016

Commenced Negotiations

Tier 1 - Timeline

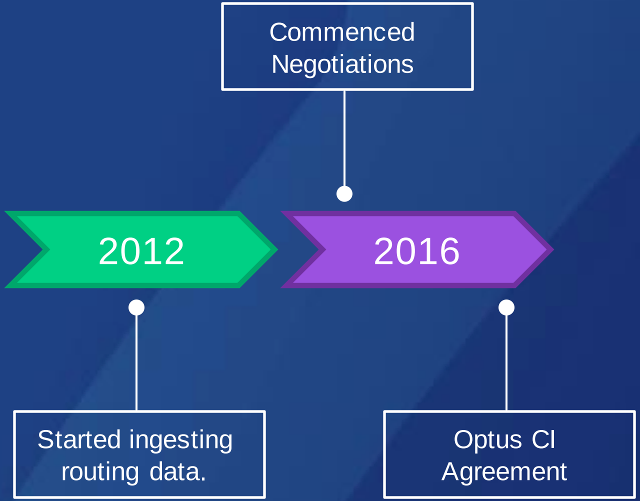

Commenced Negotiations

2012

2016

Started ingesting routing data.

Optus CI Agreement

This is where the fun begins

Every time we asked questions, the answer was something along the lines of ;

“Nobody has built an interconnect network in the last 10 years, we don’t know.”

Tier 1 - Timeline

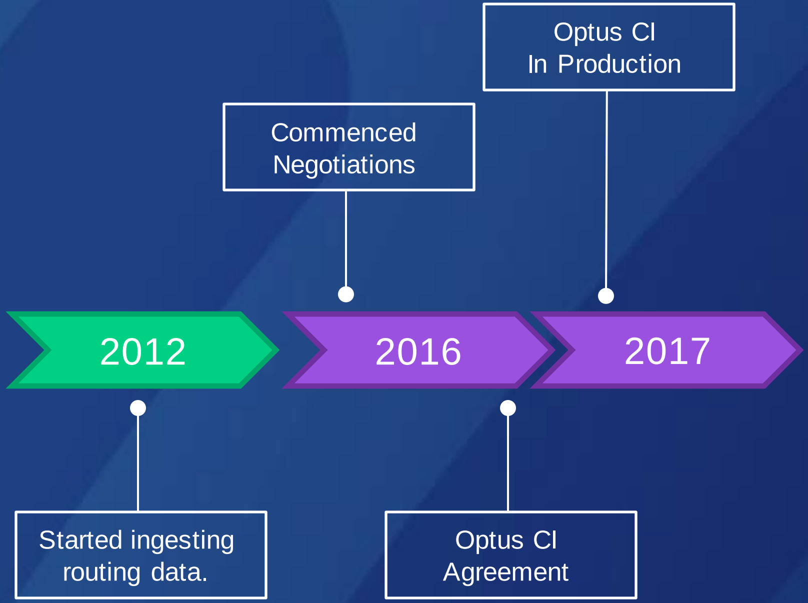

2012

2016

2017

Commenced Negotiations

Optus CI In Production

Started ingesting routing data.

Optus CI Agreement

NETSIP connect.enable

Aussie Broadband

Inter carrier call routing

There is no equivalent of BGP For routing phone numbers.

- Inter carrier call routing is pretty much entirely a bunch of static routes
- You cant aggregate routes, a single number in a 100 range might be ported
- The routing protocol is basically email
  (Ported numbers are exchanged in a text file)

On flight from Chicago to Brisbane I wrote the basis of our carrier routing module.

- Fast, Efficient, Thread safe In memory data structure for every number in Australia
- Lookups in approx. 1ms
- With caching for frequently called numbers, resulting in 100µs route lookup times

In some cases, it would take longer for LCR lookup response to traverse the network, than it takes to lookup the result.

By the end of August 2017, we had our new routing engine in production

Inter carrier call routing

There is no equivalent of BGP For routing phone numbers.

- Inter carrier call routing is pretty much entirely a bunch of static routes
- You cant aggregate routes, a single number in a 100 range might be ported
- The routing protocol is basically email
  (Ported numbers are exchanged in a text file)

On flight from Chicago to Brisbane I wrote the basis of our carrier routing module.

- Fast, Efficient, Thread safe In memory data structure for every number in Australia
- Lookups in approx. 1ms
- With caching for frequently called numbers, resulting in 100μs route lookup times

In some cases, it would take longer for LCR lookup response to traverse the network, than it takes to lookup the result.

By the end of August 2017, we had our new routing engine in production

Tier 1 - Timeline

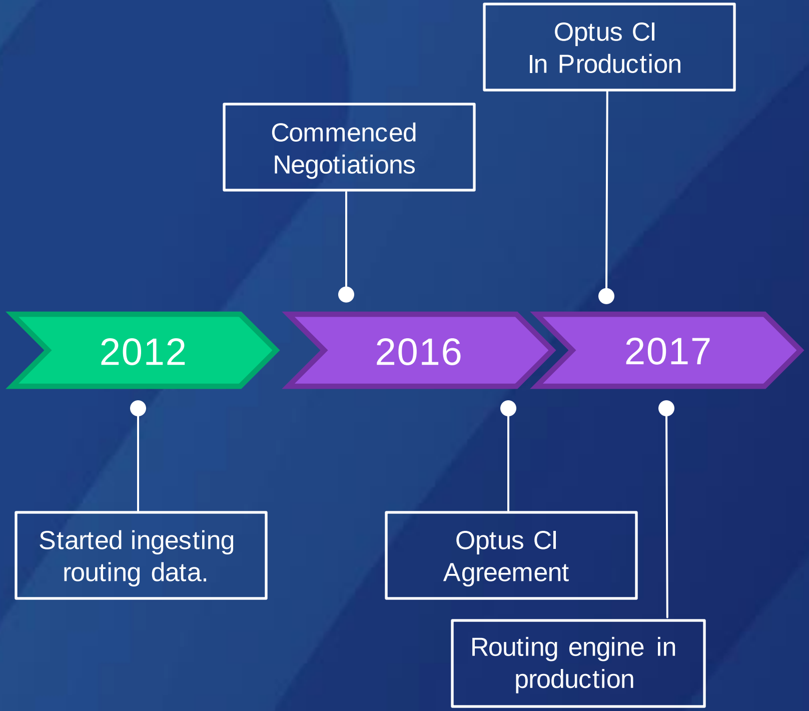

Optus CI
In Production

Commenced
Negotiations

Started ingesting
routing data.

Optus CI
Agreement

Routing engine in
production

Tier 1 - Timeline

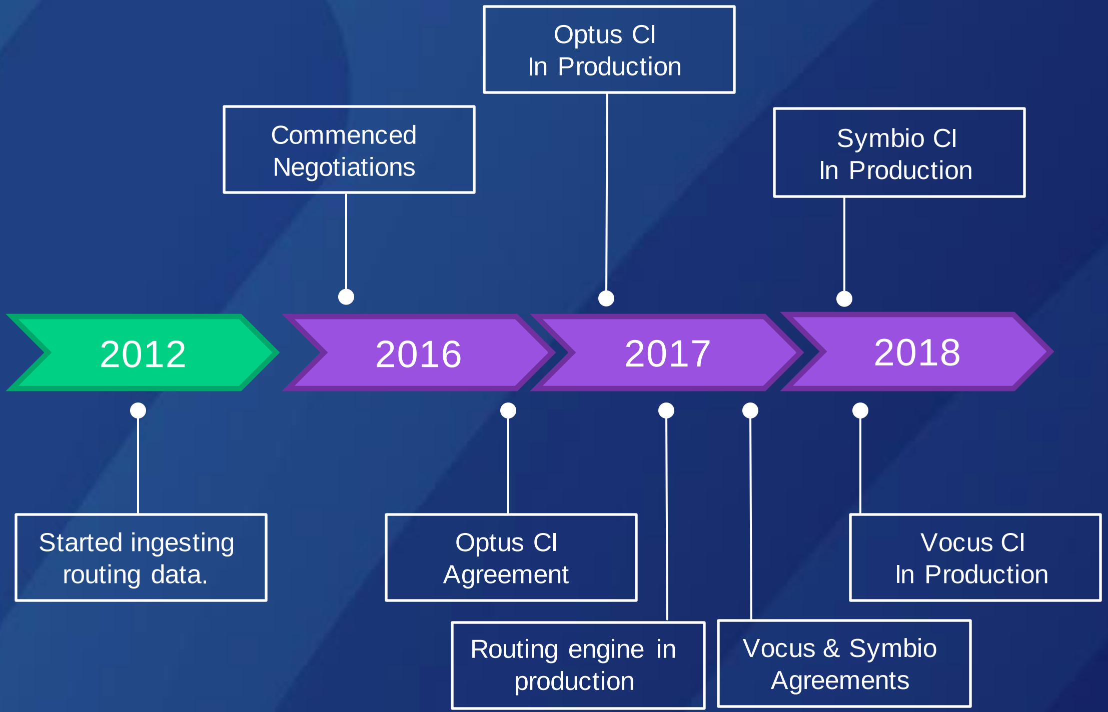

2012

Started ingesting routing data.

2016

Commenced Negotiations

Optus CI Agreement

2017

Optus CI In Production

Routing engine in production

Vocus & Symbio Agreements

2018

Symbio CI In Production

Vocus CI In Production

This is where we hit a speed bump.

Remaining carriers - 2018 Status

- **$RemainingCarrier_AAxx** took their sweet time responding to any requests.
  Struggled to find “internal resources” to fulfill our requests for interconnection
  AKA, their regulatory obligations
- **$RemainingCarrier_Vxxxxxxx** had outsourced everything to an external vendor,
  We waited months before they informed us that they wanted $500k for “billing validation”
- **$RemainingCarrier_Txxxxxx** had “different interpretations” of the regulations for interconnect

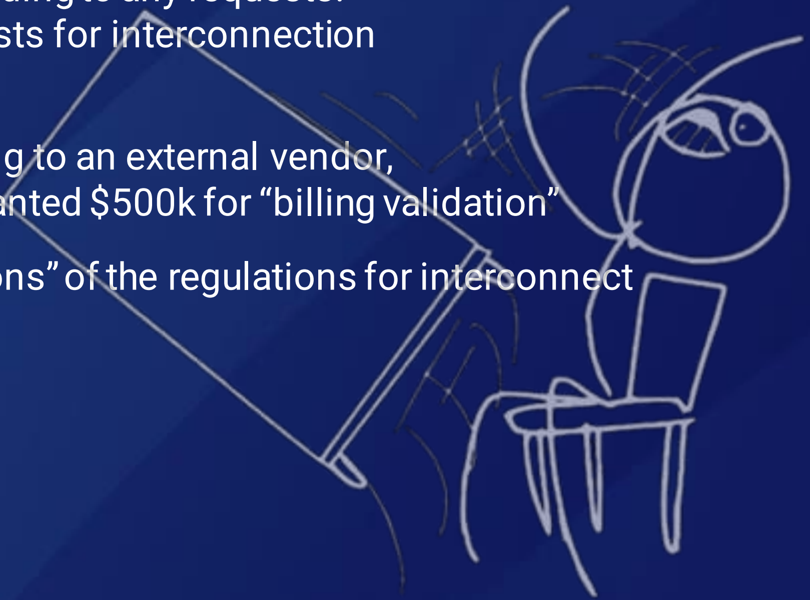

We had numerous meetings with regulators, lawyers, carriers.

(all time I will never get back)

2 Years Later

Tier 1 - Timeline

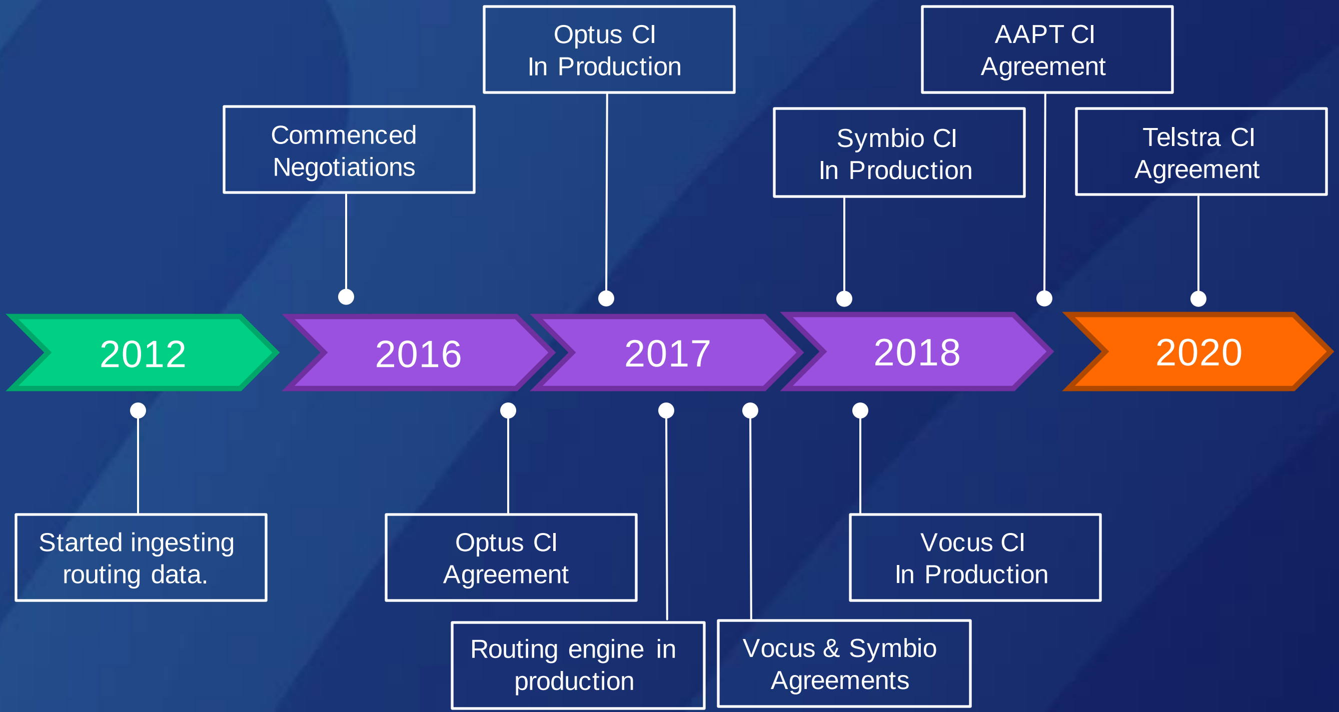

2012

2016

2017

2018

2020

Started ingesting routing data.

Commenced Negotiations

Optus CI Agreement

Optus CI In Production

Routing engine in production

Vocus & Symbio Agreements

Symbio CI In Production

Vocus CI In Production

AAPT CI Agreement

Telstra CI Agreement

Lots and lots of TDM !!!!!!!!

Not just TDM, but COPPER !!!!

- Imagine asking DC Providers for hundreds of copper XConnects !
- Krone punch downs, and contractors mixing TX & RX pairs
- You can't use STM-1 … no sir, you have to take 32+ copper E1's

- After installing hundreds of copper pairs in multiple DC’s
- After 8 sites were complete, and in testing

$RemainingCarrier_T delivered a TDM mux to one of our racks !

NETSIP
connect.enable

Tier 1 - Timeline

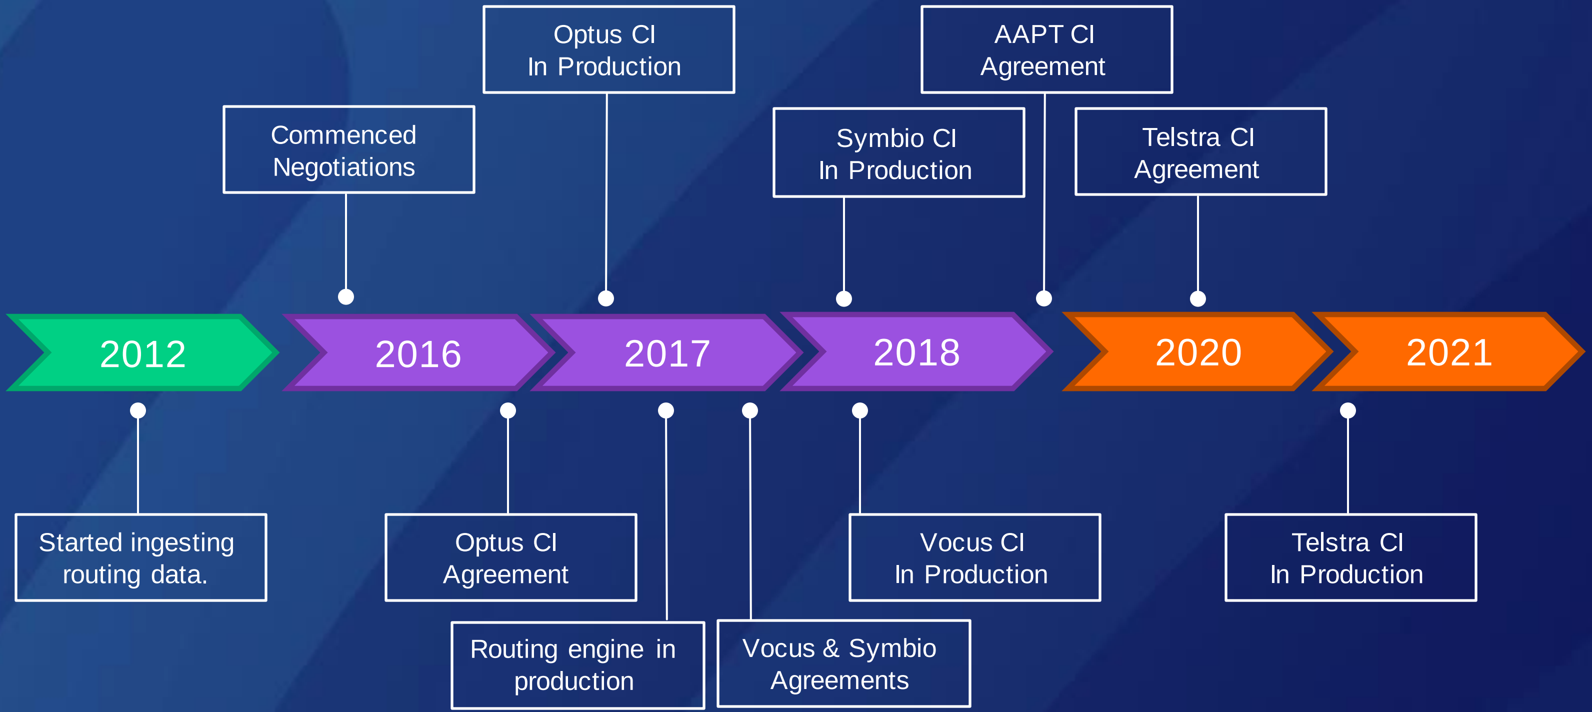

2012

Started ingesting routing data.

Commenced Negotiations

2016

Optus CI Agreement

Optus CI In Production

2017

Routing engine in production

Vocus & Symbio Agreements

Symbio CI In Production

2018

Vocus CI In Production

AAPT CI Agreement

Telstra CI Agreement

2020

2021

Telstra CI In Production

Aussie Broadband

NETSIP
connect.enable

Tier 1 - Timeline

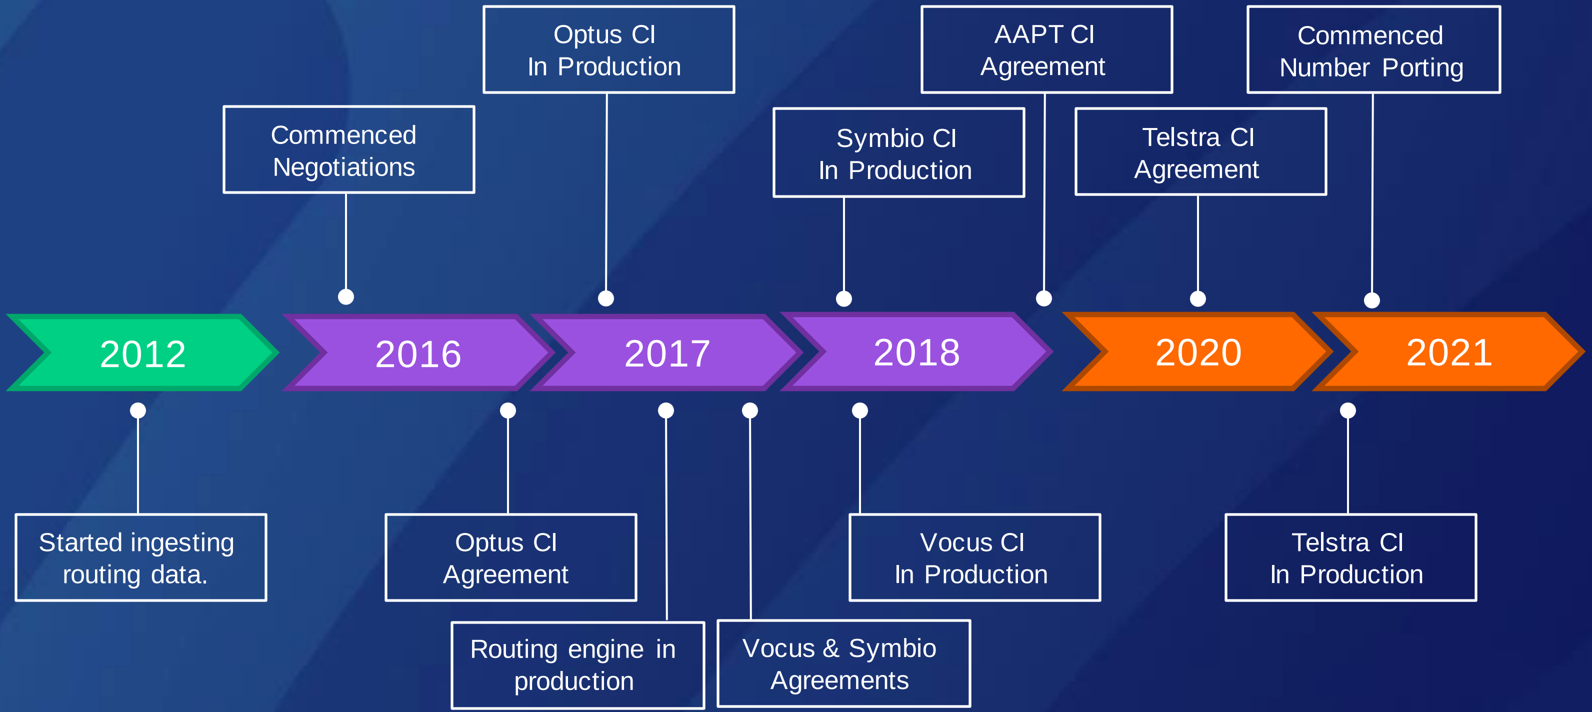

Commenced Negotiations

Optus CI In Production

Symbio CI In Production

AAPT CI Agreement

Telstra CI Agreement

Commenced Number Porting

2012

2016

2017

2018

2020

2021

Started ingesting routing data.

Optus CI Agreement

Routing engine in production

Vocus & Symbio Agreements

Vocus CI In Production

Telstra CI In Production

Aussie Broadband

NETSIP
connect.enable

Tier 1 - Timeline

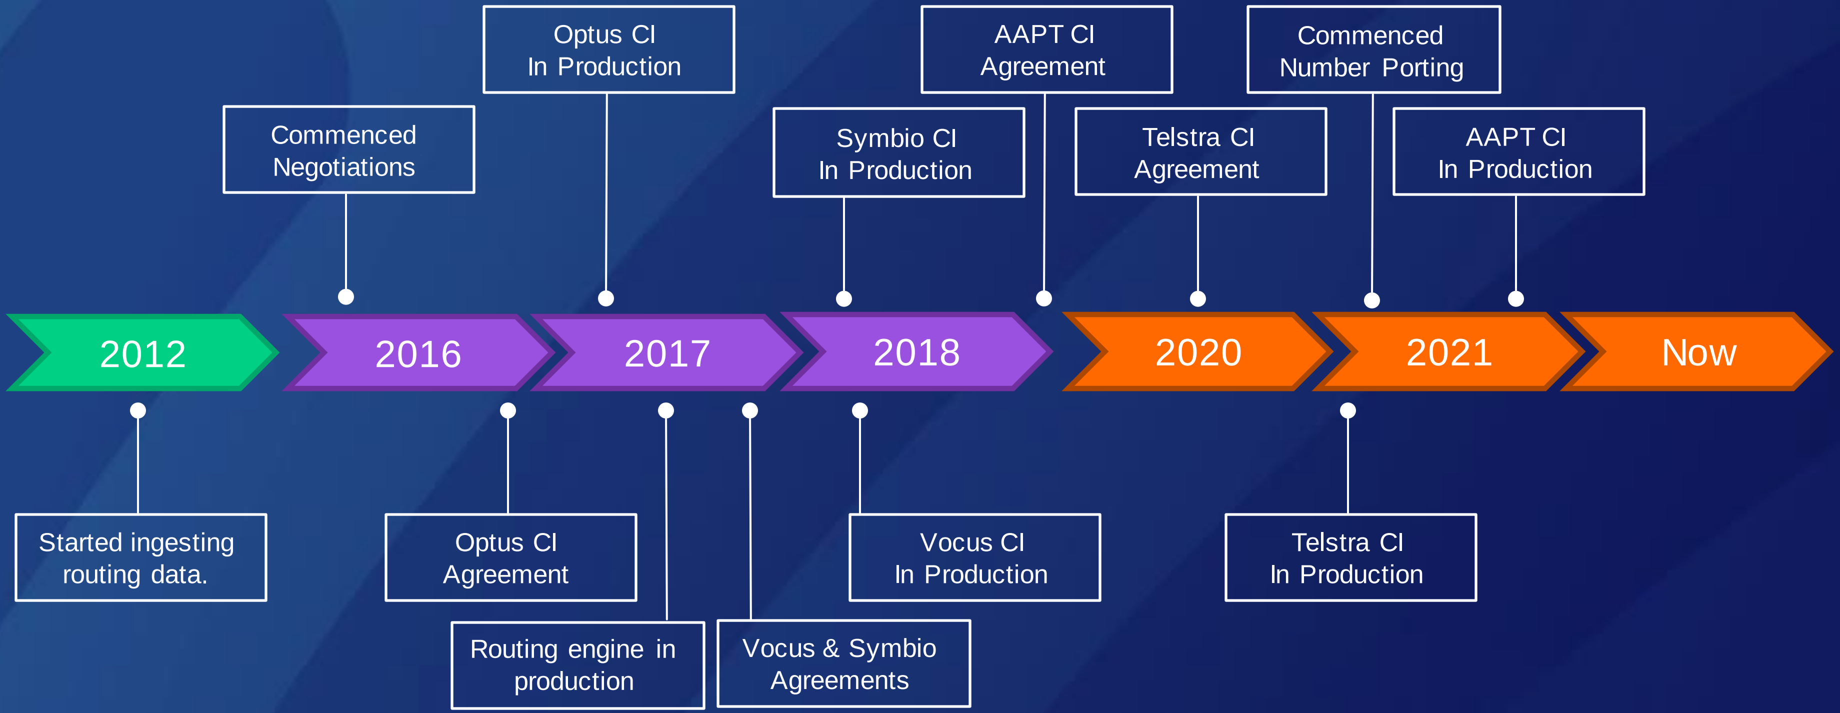

2012

Started ingesting routing data.

Commenced Negotiations

2016

Optus CI Agreement

Optus CI In Production

2017

Routing engine in production

Vocus & Symbio Agreements

Symbio CI In Production

Vocus CI In Production

2018

AAPT CI Agreement

2020

Telstra CI Agreement

Commenced Number Porting

Telstra CI In Production

2021

AAPT CI In Production

Now

Aussie Broadband

Learnings & WTF Moments

(or ranting, maybe)

WTF - #1

- $MostCarriers Why don’t we just give you good CTS Rates, and you forget this Tier 1 stuff

- One carrier had to get $SomeGuy back from retirement, nobody else knows how to pull it together

I cant count how many time I heard :

“Nobody has built an interconnect network in the last 10 years, we don’t know.”

WTF - #2 : Call Routing

Different numbers are routed & ported differently !

I explained routing for geographic numbers before :

- *Static routes*
- *Email based routing*
  (ported numbers are exchanged in text files, kindof)

13XXXX, 1300, 1800 :

Pay a government appointed monopoly (not TLS) money for the routing information

Mobile numbers :

Similar to Geographic, but not … discussion for another day.

WTF - #3 - Pseudo Regulation

- Some things are regulated, other major parts are not
  (such as fixed network costs)

- Some carriers don't believe the regulation applies to them
  (Still too open to interpretation)

- Some carriers simply don't care about their obligation, and know your only recourse is to go to court.

- The regulator wont really enforce anything

WTF - #3 - Pseudo Regulation

- Some call types are regulated

  Eg : Fixed & Mobile calls

- 13 / 1300 / 1800 are not regulated

- Commercials around number porting are unregulated

WTF - #4 - Number Portability

- The ”System” is so broken !!!

- Cat A ports are meant to be automated, but most carriers manually validate every port out request, and get upset when you send them more than 50 per day !!

- Per City Batching Requirements for Cat C

- Timezones (Comms alliance spec = Syd timezone),
  but some carriers do timezone of the number.
  Leading to unexpected out of hours bookings for Perth & ADL.

WTF - #4 - Number Portability Commercials

- **Cat C – PNV process is ridiculous**
  Manual, Expensive ($75 - $350 per batch) and time consuming
  Takes 10 business days, and often has mistakes

- Industry standard rejection charges are punitive !!
  $10-20 PER NUMBER rejection fees for a CNA rejection is common

- Out of hours port fees are hundreds of dollars,
  and faults go unresolved until the next business day
  (What are we paying for ?)

- Out of hours ports shouldn't even be needed, the porting system should guarantee call delivery during the porting process.

NETSIP
connect.enable

Questions ?

Aussie
Broadband

Building a

Mostly **Open-Source**

Tier 1 - Carrier Voice Network

NETSIP

Aussie Broadband