IPv4 as a service

Deployment models and their consequences

About Retevia

- Founded in 2017 by Lee Howard
- Business focus:
  - IPv4 as a Service
    - MAP-E, MAP-T and Lw4o6
    - NAT64
    - DS-Lite
  - IXP and on-premise deployments
  - IPv6 consulting

Choosing a technology

Different choices for different use-cases

NAT64

- **Advantages:**
  - Well-known technology
  - Adapts to user activity
  - Works with 464xlat
  - Single solution for mobile + residential
  - Simple provisioning
- **Disadvantages:**
  - Stateful
  - Difficult to scale *
  - Single point of failure *
  - Less traceable
  - Requires lot of logging

* Retevia can cluster NAT64 boxes for scaling and redundancy

DS-Lite

- **Advantages:**
  - Well-known technology
  - Adapts to user activity
  - Lot of CPE support
  - Simple provisioning
- **Disadvantages:**
  - Stateful
  - Difficult to scale
  - Single point of failure
  - Less traceable
  - Requires lot of logging

MAP-E

- **Advantages:**
  - Stateless
  - Scales horizontally
  - No single point of failure
  - Reversible algorithm
  - Simple provisioning
  - No logging required
  - Clean packet format
- **Disadvantages:**
  - Not adaptable per user
  - Lack of CPE support

MAP-T

- Advantages:
  - Stateless
  - Scales horizontally
  - No single point of failure
  - Reversible algorithm
  - Simple provisioning
  - No logging required
  - Per-destination exits
- Disadvantages:
  - Not adaptable per user
  - Lack of CPE support
  - Packet mangling and reconstruction

Lightweight 4over6

- **Advantages:**
  - Semi-stateless
  - Adaptable per user
  - Scales horizontally
  - No single point of failure
  - Reversible algorithm
  - Clean packet format
- **Disadvantages:**
  - Semi-stateless
  - Per-user provisioning
  - Logging of per-user provisioning required
  - Lack of CPE support

Decision tree

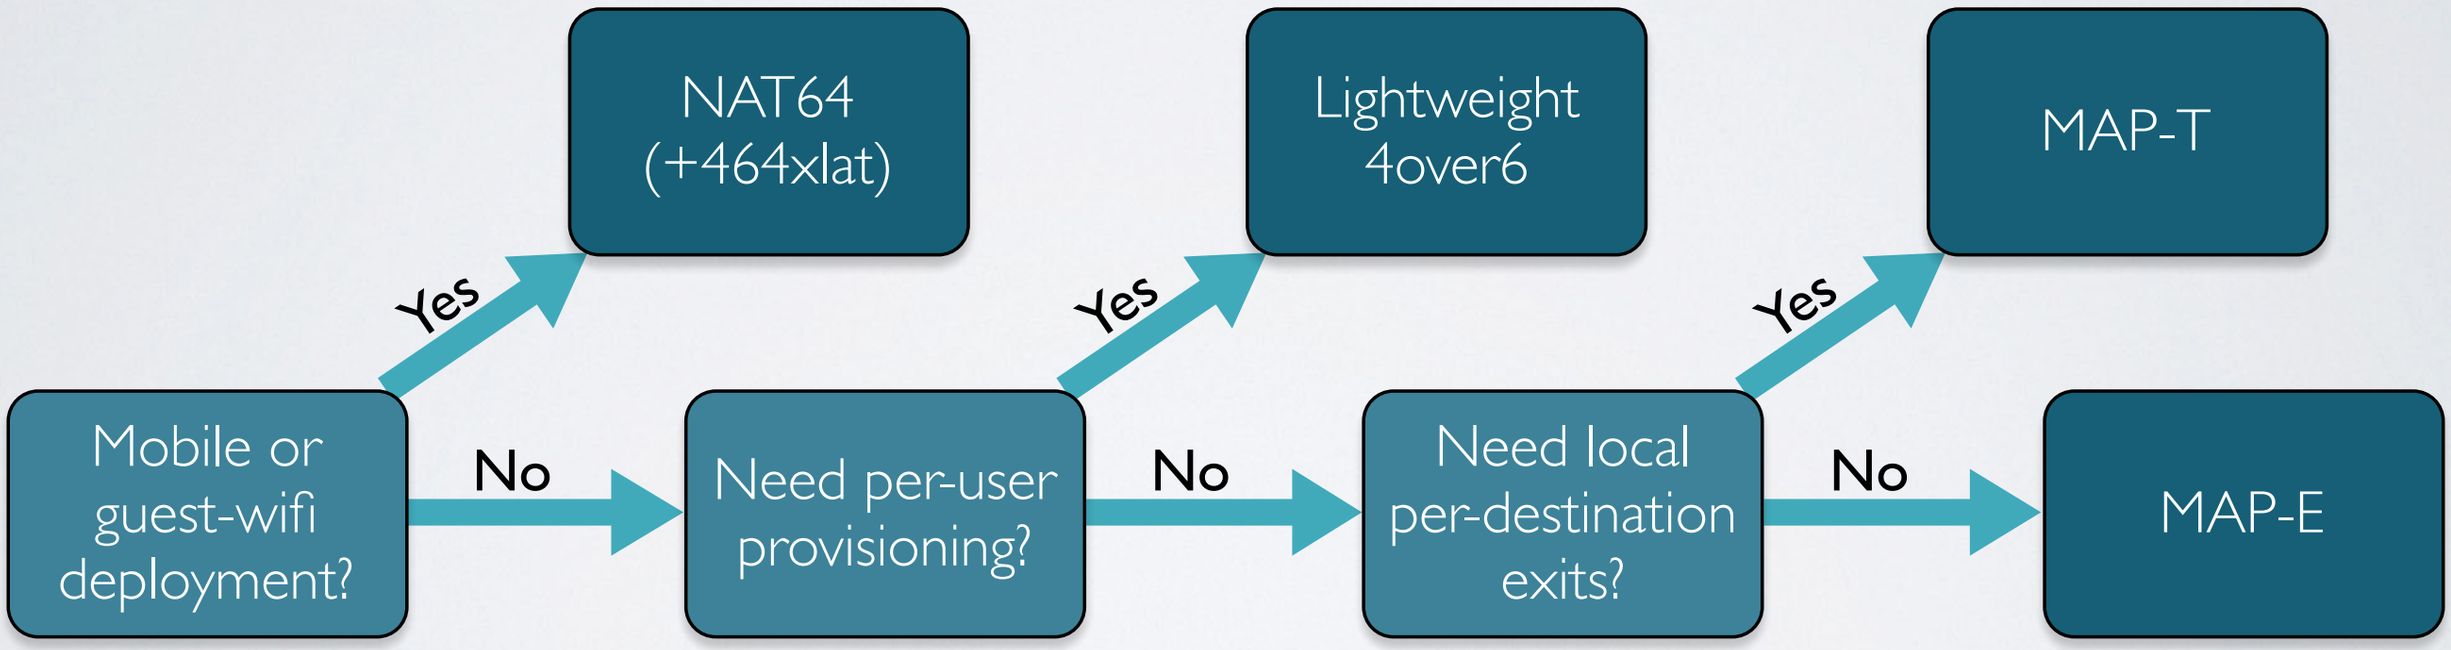

Decision tree - summary

- I would recommend MAP-E or MAP-T
- Lw4o6 provides per-user provisioning
  - But that brings a lot of complexity with it
- **NAT64 is well-supported on mobile devices**
  - But it requires stateful NAT equipment

For data centres

- SIIT-DC
  - Run the whole data centre on IPv6
  - Expose the IPv6 service on IPv4
    - Stateless, scales horizontally, no single point of failure, reversible algorithm, simple provisioning, no logging required, etc.
- **Basically an inverse 1-to-1 MAP-T mapping**

IPv4-as-a-service or on-premise

Responsibilities and scalability

IPv4-as-a-Service

- **No need for IPv4 at all**
  - Transit and peering managed by Retevia
  - DDOS protection is provided as part of the service
- **Managed infrastructure**
  - Translation, DHCPv6 and DNS are provided
  - Retevia provides admin panels
- **Great for enterprises and smaller ISPs**

Example deployment

Router
Router
Switch
Switch
Management
Logging
Translator
Translator
Translator
Translator
Translator
Translator
Translator
Etc…

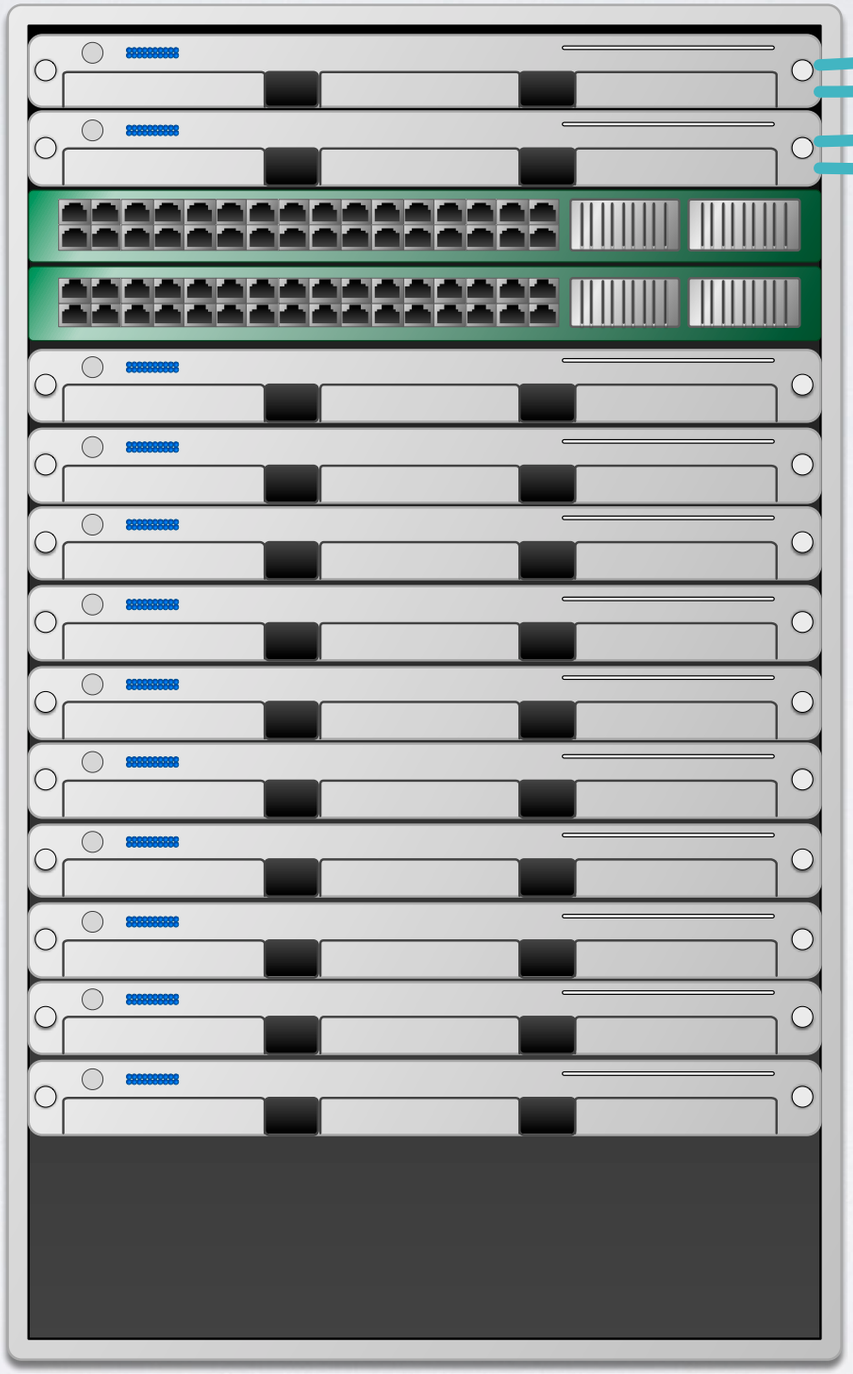

On-premise deployment

- **Manage your own IPv4**
  - Keep using your existing transit, peering and DDOS protection
- **Part of your own infrastructure**
  - Integration with your existing DHCPv6 and DNS
  - Retevia provides admin panels
- **Great for larger enterprises and ISPs**

Example deployment

Translator

Translator

Etc…

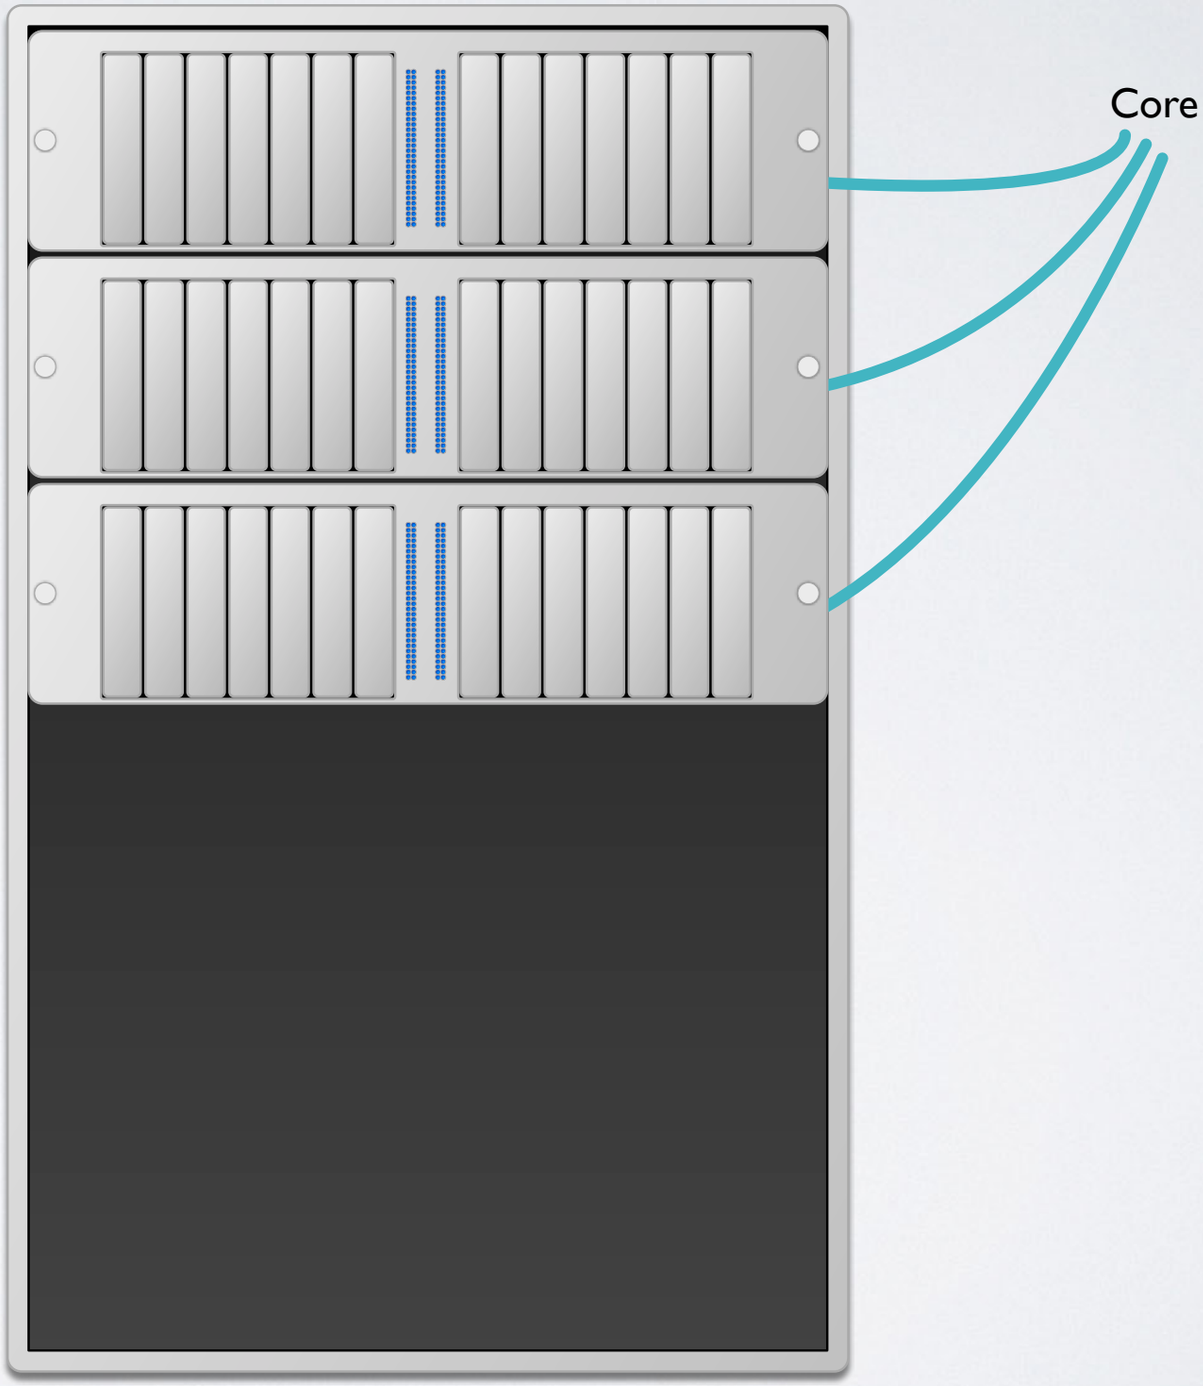

SCALABILITY

Does it go up to 11?

NAT64 and DS-Lite

- The stateful protocols
- Hard to scale
  - Distribute customer groups over separate instances
  - Little scalability within one instance
  - A customer's traffic must pass through the box with that customer's session state
- **Redundancy**
  - NAT64 redundancy still possible up to ±8 Gbps

MAP-E, MAP-T and Lw4o6

- The stateless protocols
- Scales really really well
  - Tested up to 185 Gbps at 83 Mpps
  - Expect upcoming hardware to scale to ±400 Gbps
  - CPU handles this just fine, PCIe busses are bottleneck
- **Redundancy**
  - No per-session state, add redundant boxes as needed
  - Lw4o6 has per-user state, so need to sync the config

LEGAL IMPLICATIONS

What can/must/may you track?

Data retention

Investigatory Powers Act 2016

In this Part "relevant communications data" means communications data which may be used to identify, or assist in identifying, any of the following—

(a) the sender or recipient of a communication (whether or not a person),

(b) the time or duration of a communication,

(c) the type, method or pattern, or fact, of communication,

(d) …

and this expression therefore includes, in particular, internet connection records.

Data retention

Investigatory Powers Act 2016

A retention notice must not require an operator who controls or provides a telecommunication system ("the system operator") to retain data which—

…

(c) is not needed by the system operator for the functioning of the system in relation to that communication, and

(d) is not retained or used by the system operator for any other lawful purpose,

and which it is reasonably practicable to separate from other data which is subject to the notice.

Now combine the two…

- **Before CGN/NAT64/DS-Lite**
  - ISP communication systems didn't track source, destination, time, type or method beyond "customer X gets IP address Y"
- **With CGN/NAT64/DS-Lite**
  - The translators need to track session information to perform their normal function
  - Which allows the retention notice to require the ISP to retain it

How to translate?

- Different ways to do NAT
  - Based on classic 5-tuple + private side
    - Protocol
    - Source address & port
    - Destination address & port
  - Based on client-side 3-tuple + private side
    - Protocol
    - Source address & port

How to translate?

- Different ways to do NAT
  - Based on 5-tuple + private side
    - More efficient use of port numbers
    - May reuse source address & port for different destinations
    - Precludes use of original STUN protocol
- Based on client-side 3-tuple + private side
  - Reserve one or more source ports for a user
  - Use those ports for reaching any destination

A+P based protocols

- **Client-side 3-tuple built into protocol**
  - Reserves a set of source ports for a user, based on the MAP algorithm (or Lw4o6 provisioning)
  - The NAT process in the user's CPE uses those ports for reaching any destination

NAT tables, classic 5-tuple

Protocol	Client A+P	Public A+P	Dest A+P
TCP	**10.0.0.1**:5678	**185.54.92.1:7890**	**31.13.91.36:80**
TCP	**10.0.0.2**:5413	**185.54.92.1:7890**	**31.13.91.36:443**
TCP	**10.0.0.3**:8145	**185.54.92.1:7890**	**31.13.91.37:80**
TCP	**10.0.0.4**:51821	**185.54.92.1:7890**	**31.13.91.37:443**
TCP	**10.0.0.5**:8312	**185.54.92.1:7890**	**31.13.91.38:80**
TCP	**10.0.0.6**:13578	**185.54.92.1:7890**	**31.13.91.38:443**
TCP	**10.0.0.7**:62038	**185.54.92.1:7890**	**31.13.91.39:80**
TCP	**10.0.0.8**:12345	**185.54.92.1:7890**	**31.13.91.39:443**
TCP	**10.0.0.9**:9141	**185.54.92.1:7890**	**31.13.91.40:80**
TCP	**10.0.0.10**:5421	**185.54.92.1:7890**	**31.13.91.40:443**

Bold = Log

NAT tables, client 3-tuple

Protocol	Client A+P	Public A+P	Dest A+P
TCP	**10.0.0.1**:5678	**185.54.92.1:7890**	Don't care
TCP	**10.0.0.2**:5413	**185.54.92.1:7891**	Don't care
TCP	**10.0.0.3**:8145	**185.54.92.1:7892**	Don't care
TCP	**10.0.0.4**:51821	**185.54.92.1:7893**	Don't care
TCP	**10.0.0.5**:8312	**185.54.92.1:7894**	Don't care
TCP	**10.0.0.6**:13578	**185.54.92.1:7895**	Don't care
TCP	**10.0.0.7**:62038	**185.54.92.1:7896**	Don't care
TCP	**10.0.0.8**:12345	**185.54.92.1:7897**	Don't care
TCP	**10.0.0.9**:9141	**185.54.92.1:7898**	Don't care
TCP	**10.0.0.10**:5421	**185.54.92.1:7899**	Don't care

Bold = Log

Impact on tracking

- **Full 5-tuple**
  - Every single (TCP, UDP etc) session must be logged
    - Reveals exactly who communicated with which service
    - Also reveals when and how long
- Massive amount of logging data

Impact on tracking

- **Client-side 3-tuple**
  - Only the ports available to each user are known
    - NAT64 and DS-Lite might reveal when a user is active
    - Doesn't reveal sensitive user behavior
- Low (or no) need for logging

Summary

- **5-tuple based logging**
  - Expensive
  - Reveals user's detailed on-line activities
  - Allows for potential re-use of port numbers
- **3-tuple based logging**
  - Cheap
  - Protects user's privacy
  - A little less efficient with port numbers

Q&A

Discussion!