What's an "IP Address" these days?

Geoff Huston AM
APNIC

Connected Devices

Connected IoT device count 2019 - 2030

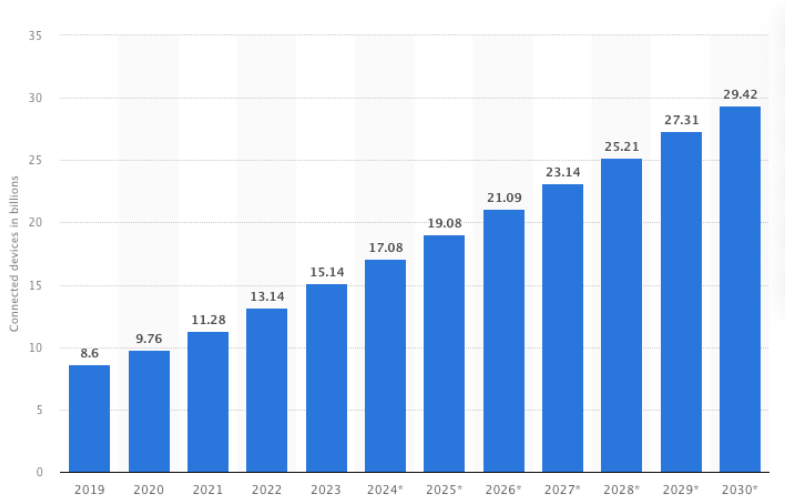

https://www.statista.com/statistics/1183457/iot-connected-devices-worldwide/

Addresses for Connected Devices

Connected IoT device count 2019 - 2030

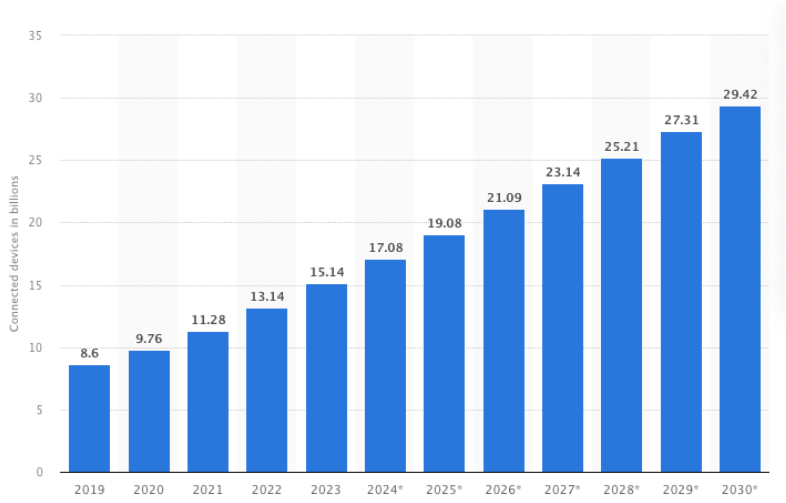

https://www.statista.com/statistics/1183457/iot-connected-devices-worldwide/

Advertised Host Addresses 2019 - 2030

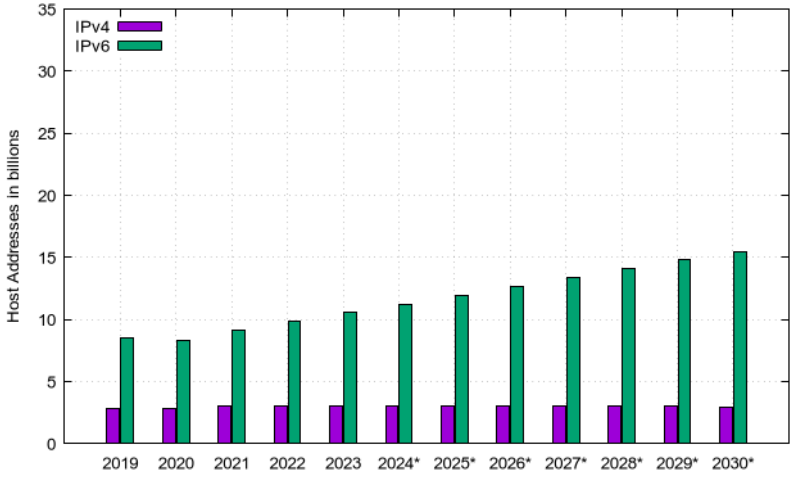

APNIC Data Projections

Addresses for Connected Devices

Connected IoT device count 2019 - 2030

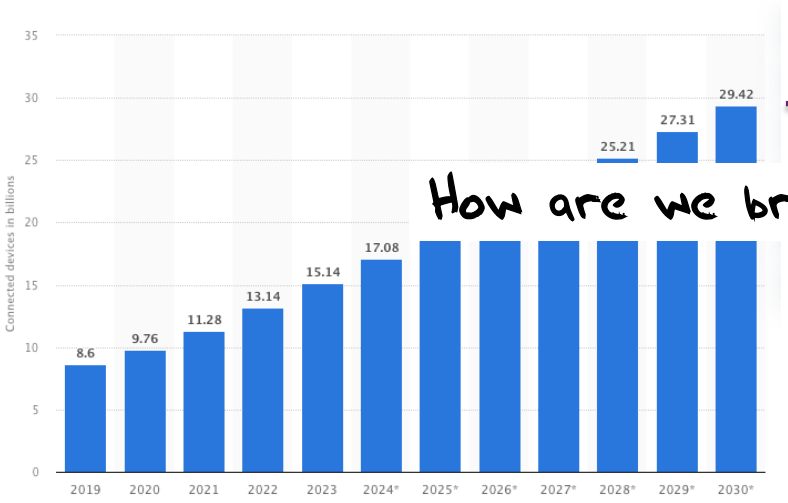

Advertised Host Addresses 2019 - 2030

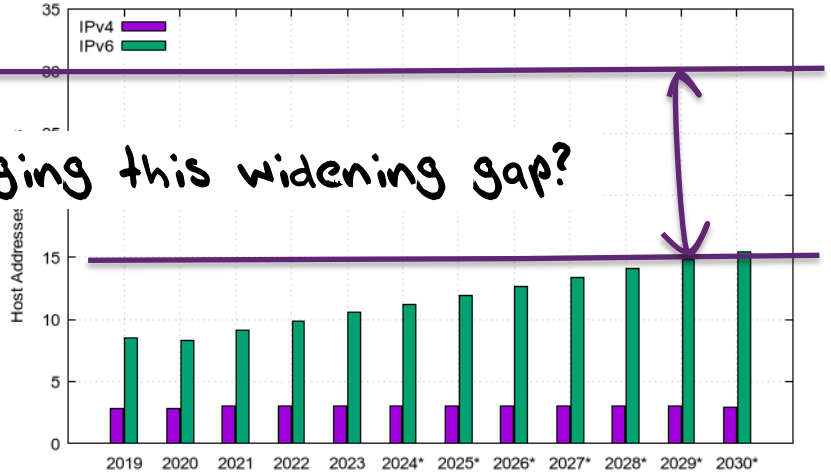

How are we bridging this widening gap?

https://www.statista.com/statistics/1183457/iot-connected-devices-worldwide/

APNIC Data Projections

What did we need from Addresses?

Identity:

- WHO - Every unique end point has a unique address value

Location:

- WHERE - Every unique address is associated with a location within the network

Reachability:

- HOW – Every address informs the network how to direct packets towards it

What did we need from Addresses?

Identity:

– WHO - Every unique en

Location:

ociated with a location within

R

– HOW – Every address informs the network how to direct packets towards it

The IP architecture uses an overloaded semantic framework for addresses

Running Short of Addresses

- Address scarcity has been a feature of public Internet services from the outset
  - “sharing” a dial-up modem bank across a larger pool of subscribers
  - “sharing” a single connection service across a set of connected local computers
- We responded to this scarcity by:
  - using common address pools and assigning addresses from the pool only when they were connected
  - Sharing a single external address across multiple internal devices (NATs)

Client/Server Architecture

- The expansion of the computer industry into “low end” client devices motivated a change in the network service architecture to differentiate between *clients* and *servers*
- Client platforms have no strict need for a persistent identity, and hence no strict need for a persistent address
- This adoption of a client/server architecture has taken the pressure off the addressing system in having to provide a permanent unique IP address to every client

Clients

- Only need a persistent address to use for the lifetime of a connection

Clients

- Only need a persistent address to use for the lifetime of a connection
  - QUIC's form of client address agility allows the client to be assigned a different external address each 2 x RTT interval
  - We've shifted session identification away from the old 5-tuple of addresses + ports to an application-level token

Servers

- What's the minimum address requirement for a multi-location service delivery platform?
  - 1 anycast address
- How does the platform differentiate between different hosted services?
  - The DNS
- How do we identify service endpoints?
  - The DNS
- How do we secure service identity?
  - DNS Name certificates!

Architectural Evolution

- The 1980's network architecture was an address-based architecture where every attached endpoint was uniquely addresses by its network attachment using a persistent address
- Names were seen as an alias for addresses as part of the application-level framework intended to improve ease of use
- Addresses were isomorphic to identity

Evolutionary Changes

- We shifted to an asymmetric architecture of clients and servers
- We dispensed with persistent network-wide level identity for clients and used local context addresses instead – clients do not have an address-level identity
- We also dispensed with identity semantics for servers. The prevailing use of anycast service platforms implies that addresses are used for location and forwarding, but not for identity

Today's Name-based Network

- Today's service network is a name-based network
- Names provide identity, names underpin authenticity and security
- Addresses retain the semantics of location and forwarding, but have largely dispensed with the role of endpoint identification

Some Questions

- How well does IPv6 reflect this evolving network architecture that has shifted identity from addresses to names?

Some Questions

- How well does IPv6 reflect this evolving network architecture that has shifted identity from addresses to names?
- Anycast provides a local scope for replicated service platforms. Can anycast also allow local-scoped client-side addresses?

Some Questions

- How well does IPv6 reflect this evolving network architecture that has shifted identity from addresses to names?
- Anycast provides a local scope for replicated service platforms. Can anycast also allow local-scoped client-side addresses?
- If every connected device does **not** need a permanent global address, then how many addresses do we need?

2024 APRICOT APNIC 57

BANGKOK, THAILAND

21 February – 1 March 2024

#apricot2024