Prefixes as Probabilities

A Modest Proposal to Radically Extend the Life of IPv4

Todd Underwood
toddunder@gmail.com

The Problem as We Know It

- v4 Prefix space running out quickly
- v6 not operationally equivalent to v4 yet
- Time is running out and mechanisms are needed to extend the useful life of v4
- Possibilities so far:
  - Address market/trading
  - Rationing
  - Flat out

Resource Allocation Now

- RIRs allocate resources on a 0/1 scale (they either allocate or not to a single entity).
  - Odds that someone else has been allocated the same resource are very, very low (mistakes only)
- Disadvantages:
  - Extremely discrete: yes/no allocation
  - Requires external factors (price change, e.g.) to tune allocation/depletion rate
  - No way to deal with varying value (1/8)
- Market (transfer) policy allows limited "selling" of resources, but at high transactional friction.

Resources as Probabilties

A thought experiment:

Allocate unique resources to those who need them at a premium and massively discount multiply allocated resources.

Unique Resource (current model):	$$$$$$$
<5% chance of collision:	$$$$$
~20% chance of collision:	$$$
~50% chance of collision:	$cheap!

- Continuous space, market-based pricing.
- Need for uniqueness not universal

Advantages

- Massively extend the lifetime of IPv4 space by permitting multiplicative reuse of remaining prefixes
- Derivatives market possible: trades on prefix-probabilities futures will provide direct predictions on address exhaustion and value
- Trivial ability to use imperfect-quality space (1/8, 240/4, ex-spammer space and others)
- Space expansion effect could be as high as 100x or 10000x but modeling needed.

Modeling the Space Advantage

Cr - count of reuse; the total number of times a given resource has been allocated
Oop - Operationally Overhead on a Prefix
    not just collisions, but any operational impact
Ap - Association domain for a prefix
    characterize the domain in which a prefix operates on the internet -- access type, geography provider, and use pattern

p(oop) is obviously a function of Cr & Ap

Future work: develop a useful model of Ap domains and characterize possible Cr within each.

Collision Avoidance and Detection

Promising techniques for collision avoidance:

Time Limited Address (TLA) - An address intended to be used in particular times only

Geography Limited Address Domain (GLAD) - The larger the committed geographical uniqueness the higher the price and guarantee

Market Address Domain (MAD) - Industry segment (NAICS or other system) or develop a more industry-appropriate categorization system.

Provider Isolation and Selection System - Specifically allow other providers to reuse a resource

Each needs to be modeled for cost/benefit trade-offs

Engineering Challenges

- Prefixes increasingly multiply originated
  - Possible to explore time, geography and provider as collision-reduction spaces
  - Many conflicting uses will be minimally disruptive (NAT, different providers, little use overlap)
  - Market mechanisms needed to re-price actual conflicting prefix use (future work)
- Actual detection of collision mitigation will likely require ARIN et al to flow-instrument several internet core networks and do high-rate sampling.
- Difficult work for high payoff

Questions?

Offers to help with this work

toddunder@gmail.com