# 4V6 – aka stateless 4Via6

http://tools.ietf.org/html/draft-dec-stateless-4v6-00

W. Dec (wdec@cisco.com)

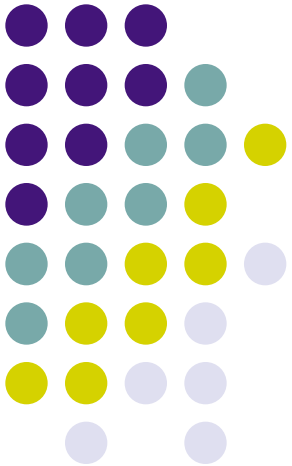

# Introduction

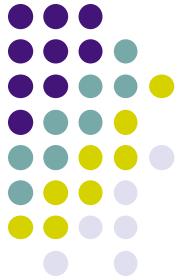

- A number of A+P based techniques for supporting IPv4 using an IPv6 infrastructure have been proposed.
  - 4rd
  - dIVI
- These stateless variants share a lot of characteristics, aka stateless 4V6
- Progress on adoption of the techniques has been delayed at least in part due to issues claimed to apply to *any* A+P proposal.
- This draft & presentation highlights:
  - The general characteristics of a stateless 4V6
  - The applicability of the issues to 4V6

# Stateless 4V6

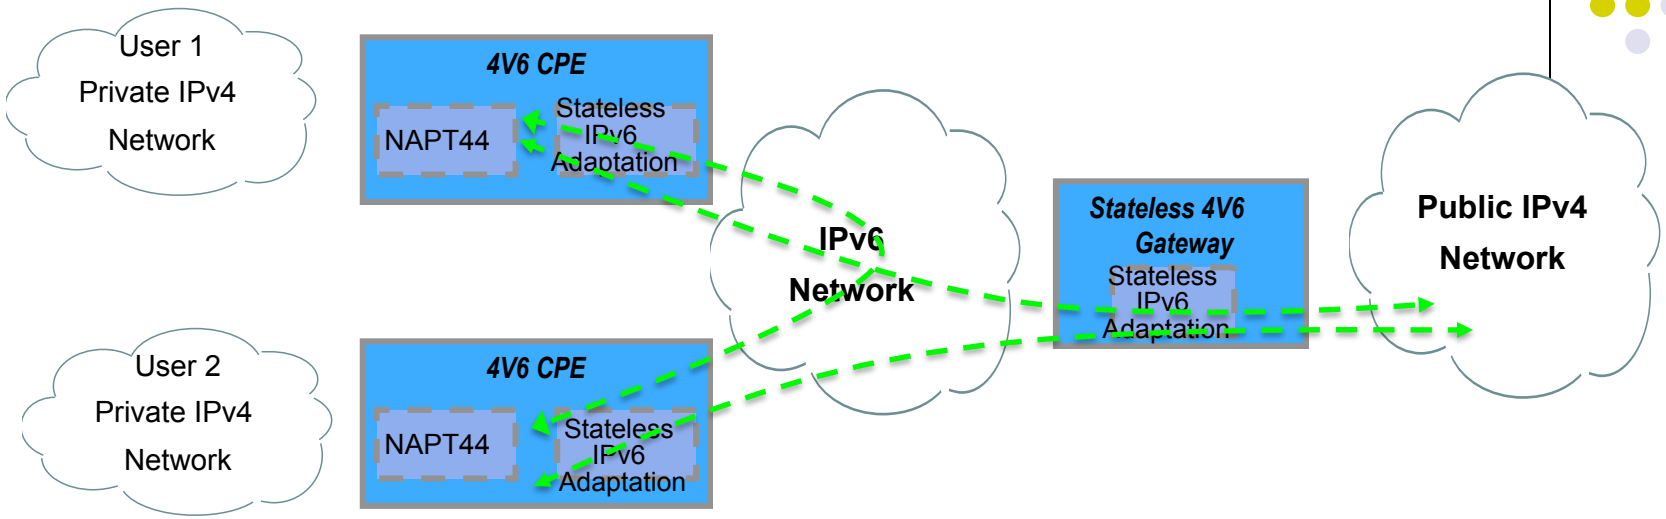


- Dedicated 4V6 CPE: IPv4 interface to the user, IPv6 interface to the SP network. IPv6 adaptation function can be stateless tunnel or stateless NAT
- NAPT44 function derives IPv4 address + port range(s) from CPE's IPv6 address. NAPT44 operate in port range restricted mode
- The CPE has one IPv6 addressed interface for the stateless 4via6 application and IPv6 adaptation function.
- Stateless 4V6 Gateway with matching IPv6 adaptation function. No translation state or logs
- No per user configuration on CPEs or Gateway.
- IPv6 address assignment log contains all info regarding NAT. No dynamic NAT flow/xlate logging
- No DNS64 or IPv6 ALGs are used. Any existing IPv4 ALGs on the CPE work "as are". DNS resolver proxy.
- No changes to end user IPv4 hosts/stacks
- User-User IPv4 traffic can flow directly between CPEs over IPv6 (bypassing gateway)
- End user's native IPv6 traffic/network set-up using regular DHCPv6 PD means (not shown – Native IPv6 works as ships in the night)

# Claimed Issue

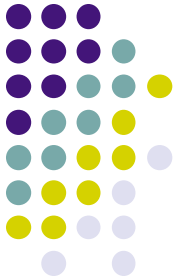

- Unicast addresses, implementation on hosts and ambiguity
  - IPv4 address no longer belongs to a single “host”
  - Host IP stack changes required to support shared addresses
  - Ambiguity with multi interface hosts
- Applicability to 4V6
  - The 4V6 solution does not address end hosts. Shared address is confined to the NAPT44 function
  - One NAPT44 function per CPE with one IPv4 address and unique port ranges supplied via IPv6 address
- Conclusion
  - Issue is not applicable

# Claimed Issue

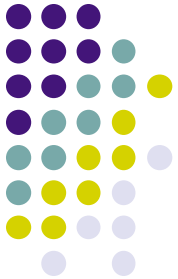

- Non TCP/UDP Protocols
  - Such protocols stop working across NAPT44
  - Such protocols stop working on a shared link with shared addresses
- Applicability to 4V6
  - The 4V6 solution does not address end hosts.
  - Non TCP/UDP protocols stop working across *any* type of NAPT44, without dedicated markup
- Conclusion
  - Issue is not applicable

# Claimed Issue

- Provisioning and OSS
  - Provisioning and OSS systems need to evolve to handle A+P within DHCPv6, Databases, etc
- Applicability to 4V6
  - 4V6 depends on an operational IPv6 network incl provisioning OSS.
  - Provisioning of 4V6 CPEs needs to indeed be addressed, but that goes for *any* CPE such as 6rd, Ds-Lite, SIP client, etc
  - Operators have been doing this for years. DHCP is an SP mainstay also for IPv6, eg DHCPv6 PD
  - Useful to note that the issue was NOT raised against other types of solutions which themselves impose onerous Provisioning and OSS changes. Eg DS-Lite requires the addition of (all of which 4V6 doesn't):
    - NAT logging
    - Per user CPE provisioning
    - New monitoring techniques
- Conclusion
  - It's a deployment trade off., not a technical show stopper
  - 4V6 actually simplifies the evolution of systems needed in comparison to some other techniques eg AFTR.
  - 4V6 provisioning needs are on par with 6rd

# Claimed Issue

- Training and education
  - Developers and support staff need to be trained
- Applicability to 4V6
  - Granted, developers and support staff will need to be trained in IPv6
  - Support staff is *already today* trained in troubleshooting CPE based NAT
  - Many developers unaware of IPv4 address crunch and already current NAPT port restrictions. This is a bigger problem.
- Conclusion
  - 4V6 falls in line with current SP operational practices
  - IPv4 developers should be NAPT port conscious. Applies to all forms of NAPT usage; AFTR, 4V6.

# Claimed Issue

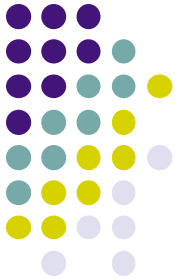

- Security
  - Restricted port range greatly weakens IP TCP/UDP protocol security, eg Random attacks
- Applicability to 4V6
  - Random TCP attack challenge: 2^32 * 2^(port range bits). This is computationally rather taxing with likely port ranges, eg 9-10 bits or more.
  - UDP: Application protocol dependent. DNS: 2^16 * 2^(port range). This is computationally not taxing even with the full port range (16 with bits).
    - Eg. On some of today's common GPUs 2e9 per second keys can be generated; 2^32 keys in ~4 seconds.
  - For UDP/DNS, the CPE is performing resolution over v6, hence port constraint does not apply to DNS attack.
  - Alternative solutions do not guarantee full port range is used
  - Extensions have been proposed which allow further port range randomization.
- Conclusion
  - For practical purposes, 4V6 does not appear to substantially degrade security. Mitigation techniques can be adopted.

# Claimed Issue

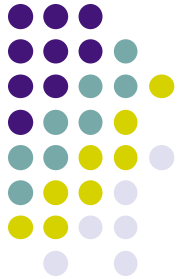

- Port Statistical Multiplexing and Monetization
  - Stateless 4V6 does not allow statistical multiplexing of ports.
  - Port limits will drive operators to set prices for ports
- Applicability to 4V6
  - 4V6 represents a number of tradeoffs compared to centralized NAT:
    - Design simplicity
    - stat muxing vs no-nat-logging
  - For many operators, even a 64x increase of remaining IPv4 addresses is sufficient
  - Monetization of ports is equally easily achievable using other techniques (eg AFTR)
- Conclusion
  - 4V6 represents a trade-off; simplicity vs port use efficiency.
  - Monetization of port space is not a technical matter. The technology to do so is available (CGN), even without 4V6

# Claimed Issue

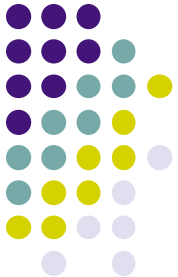

- Re-addressing
  - Changes to port ranges require changes of IPv6 addresses
  - “IPv6 re-addressing is hard” problem
- Applicability to 4V6
  - Granted, changes of port ranges require CPE re-addressing. However, many of today’s operators deal with readdressing on a regular basis (with IPv4)
  - The problem does not quite fall under the “IPv6 re-addressing is hard” class
    - the change of address is confined to the CPE and the 4V6 app. The user’s home is typically not re-addressed.
  - 4V6 (Re-)addressing can be achieved in multiple ways; DHCPv6, TR69, other.
- Conclusion
  - Re-addressing is something that a 4V6 system can do in multiple ways and the scope of impact is limited to the 4V6 CPE.
  - IPv6 re-addressing is something that operators will be exposed to irrespective of 4V6.

# Summary

- Majority of issues attributed to past A+P are not applicable to the characterised 4V6 system
- The remaining few represent classic tradeoffs in areas of:
  - Operations
  - Design & Implementation
  - Scalability
- These are tradeoffs for adopters to determine based on standard solutions
  - Approaches of stateless tunnel and stateless NAT IPv6 adaptation functions represents another tradeoff
- IETF technical community should allow progress of 4V6 solution variants.
  - Where is the home for 4V6 solutions?