# IPv6 Site Connection to Many Carriers

draft-fbnvv-v6ops-site-multihoming-03

Nick Buraglio buraglio@forwardingplane.net
Klaus Frank klaus.frank@posteo.de
Paolo Nero oselists@gmail.com
Eduard Vasilenko vasilenko.eduard@huawei.com
**Paolo Volpato** paolo.volpato@huawei.com

# Problem Statement and Solutions Considered

- The draft's target is **to present the currently available IPv6 options to support MHMP deployments** and **to discuss their strengths and weaknesses**.
  - Consider options used in the field – No endorsement to any of them
  - Said that, we prefer the native solutions (end-to-end connectivity). Desirable avoiding NAT in multi-homed deployments.
- The reason for that is that enterprises typically require carrier resilience for running their business and forecast is quite high [1]. IPv6 solutions, however, have drawbacks.
- Looked at the following solutions, both native and *stop-gap*:
  1. Static PI address space to the site
  2. Dynamic PA addresses distribution from carriers
  3. *Static ULA with NPTv6 translation*
  4. *Static ULA with NAT66 translation*
  5. Shifting Internet access resilience to a central site
  6. Application proxy.

[1] Report from the IAB Workshop on Routing and Addressing", RFC 4984, https://www.rfceditor.org/info/rfc4984

# Characteristics Considered in the Analysis

- Requirements reflect section 3.1 of RFC 3582.
- All solutions have different technical advantages and disadvantages. Not considered aspects based on geography, market, and organization sizes.

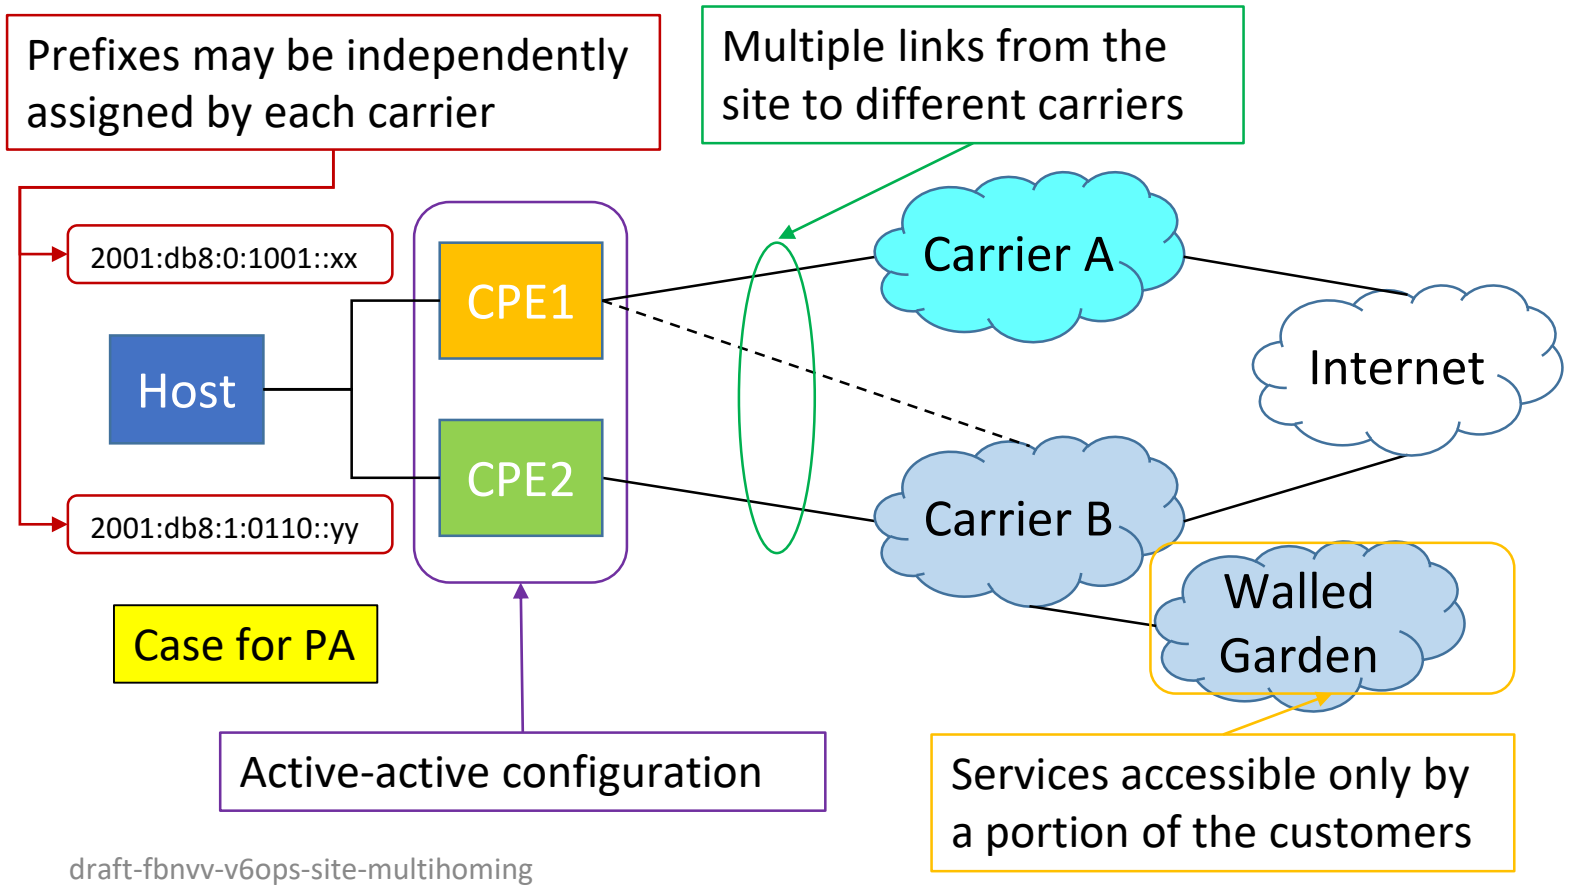


| |
|---|
| Site resiliency to an arbitrary number of carriers, with an arbitrary number of routers on the link |
| End-to-end connectivity wherever possible |
| Possibility for internal communication using any prefixes distributed by local routers, irrespective of the status of the connectivity to the carriers |
| The speed of convergence for the prefix deprecation on the site, after connectivity is lost, should be comparable to the speed of routing convergence on the site |
| Support for sites with complex topologies, including multiple internal on-site hops |
| Access to carrier's "subscriber-only services" allowed using the address space distributed by the particular carrier |
| Possibility for traffic steering between different paths based on bandwidth, cost, load, latency, etc. |
| … |

# Update from version -02

- Reviewed section 3 (problem history for the host-driven solution) to consider the advancement of draft-ietf-v6ops-dhcp-pd-per-device.
- Clarified that the cases Static ULA+NPTv6 (section 5.3) and Static ULA+NAT66 (section 5.4) document operational deployment (and are not promoted by the authors).
- Resulting order of preference, based on technical motivations, shows that “PI” is preferred over “PA” that is preferred over “ULA+NPT” that is preferred over “ULA+NAT”.
- Added Appendix A to show what is used in the field.
- Editorial changes.

# Survey on Existing IPv6 Multi-homing Deployments I

https://docs.google.com/forms/d/e/1FAIpQLSdvj4VtixaoXpMpfXhUJawXdQ60MzBKKp6aZ3i9FkKTvynqSg/viewform?usp=sf_link

***Q1. How many CPEs are connected to ISPs on the average site?***
17 responses

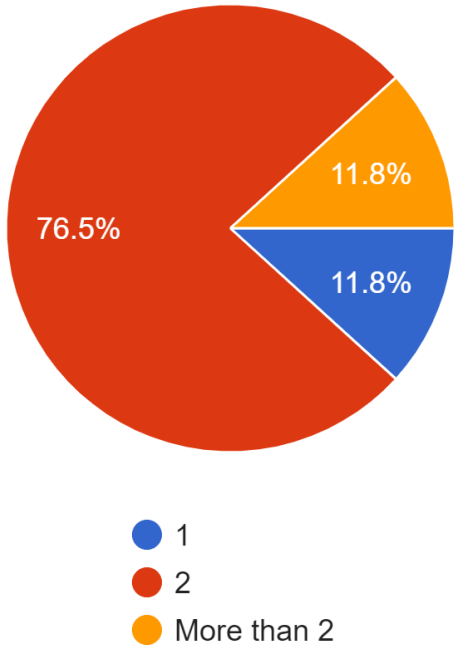


***Q2. How many uplinks are configured per CPEs/CEs in your average site?***
17 responses

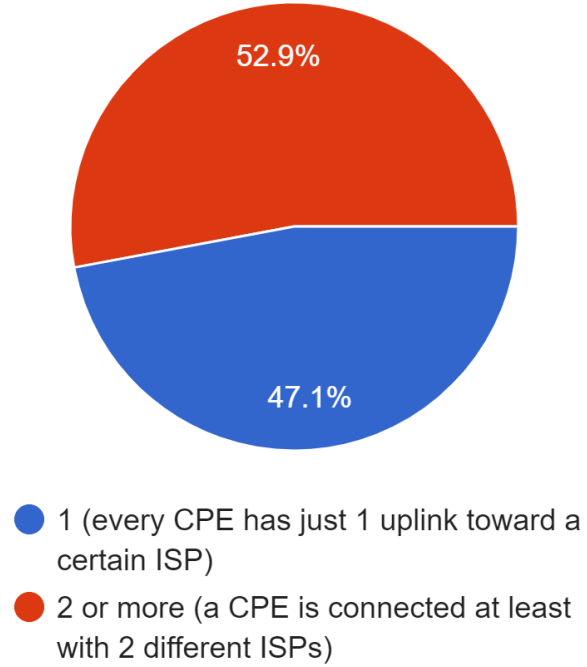


***Q3. Which configuration do you support?***
16 responses

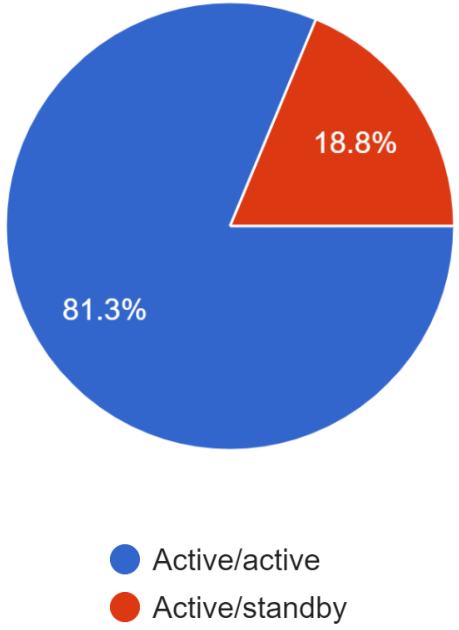

# Survey on Existing IPv6 Multi-homing Deployments II

***Q4. Do your CPEs/CEs implement VRRP or any other dynamic exchange of redundancy control information?***

17 responses

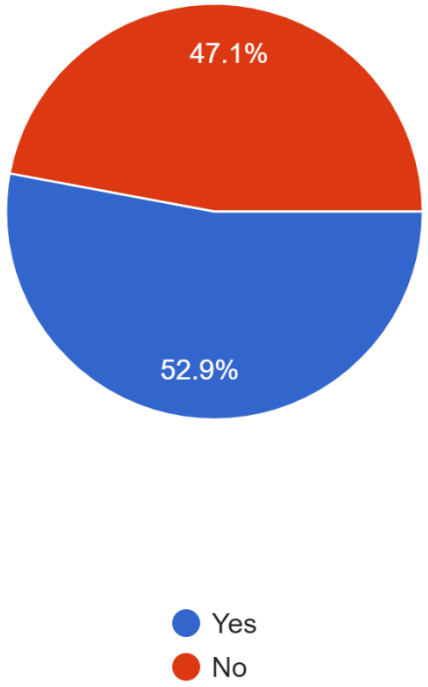


***Q5. If the answer to the previous is "No", which method do you employ for supporting IPv6 multi-homing?***

12 responses

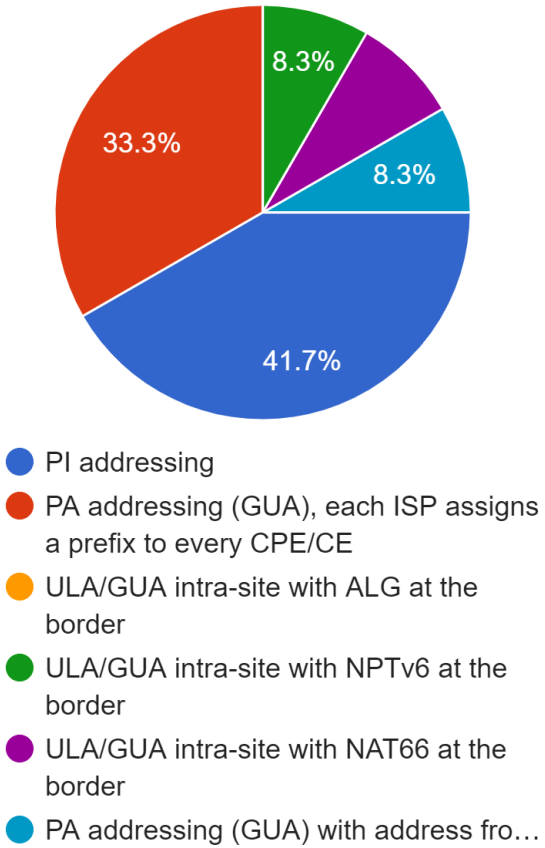

# Conclusion and next steps

- We are open to proposals on:
  - The requirements to evaluate the solutions considered
  - The solutions themselves
  - The survey.

- We’d like to ask the WG if there is interest on progressing the draft:
  - The draft received quite a good amount of comments in the past months
  - This is a clear problem space that is not well understood by the community. Is there any reason to not adopt it?

Thank you

- Backup

# Solutions Considerations

| | Requirement | PI | PA | ULA+NPT | ULA+NAT |
|---|---|---|---|---|---|
| 1 | Carriers Resiliency | + | + | + | + |
| 2 | End-to-End Connectivity | + | + | +/-*1 | - |
| 3 | Internal Connectivity | + | + | +/-*2 | +/-*2 |
| 4 | Convergence Speed | + | +/-*3 | + | + |
| 5 | Complex Topology Support*4 | + | -*5 | +/- | + |
| 6 | Subscriber-only Services | - | -*6 | +/- | +/- |
| 7 | Traffic Steering on Router | +/-*7 | -*8 | + | + |
| 7 | Traffic Steering on Host OS*9 | - | - | - | - |
| 7 | Traffic Steering on Application | - | -*10 | - | - |

**Footnotes**

1. Permits initiating connectivity in any direction, poor app support as NAT/NPT on IPv6 is not expected, lacking UPnPv6, and STUN not used/implemented or both ends behind NAT/NPT without external STUN server.
2. Difficulty promoting ULA space above IPv4 in [SASA] policy table of clients.
3. Missing market adoption for HNCP and SEND to instantly deprecate prefix on link failure.
4. Uncommon to get more than one /64 via cellular
5. PE+CPE need DHCP−PD support; CPE needs capability to assign one prefix per upstream to client subnets (uncommon; workaround using one router per upstream for each client subnet) + No src-routing by Router Advertisements alone
6. Needs either unsupported or by default disabled RFC4191 “Route Information Option”
7. High complexity for inbound traffic steering
8. Only through Prefix “choice” restrictions
9. At most managed clients only, e.g. via policy
10. Not practical, e.g. by explicitly binding src IP (+missing multi−homing support in most apps)