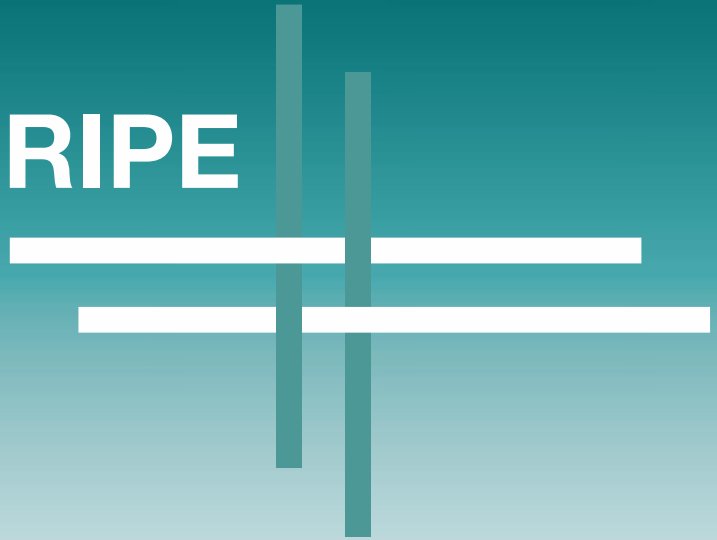

IPv6 Connectivity to Many Carriers

IETF v6ops draft: “IPv6 Site Connection to Many Carriers”
draft-fbnvv-v6ops-site-multihoming-01

Authors: Nick Buraglio (Energy Sciences Network), Klaus Frank (Individual), Paolo Nero (Individual), Eduard Vasilenko (Huawei), Paolo Volpato (Huawei)

Problem Scope and Its Importance

- Carrier resilience is a typical business requirement
  - An enterprise site may be homed to multiple carriers
- IPv4 deployments have solved multi-homing through private internal site addressing in combination with separate NAT engines
- With IPv6, support for true end-to-end connectivity on the Internet is desirable avoiding NAT in multi-homed deployments
  - Native IPv6 solutions for carrier resilience, however, have drawbacks
  - The draft's target is to present the currently-available options and discuss their strengths and weaknesses
- IAB foresees as many as 10M multi-homed sites by 2050
- The views come from technical talks, are not the authors' preference.

Characteristics Considered in the Analysis

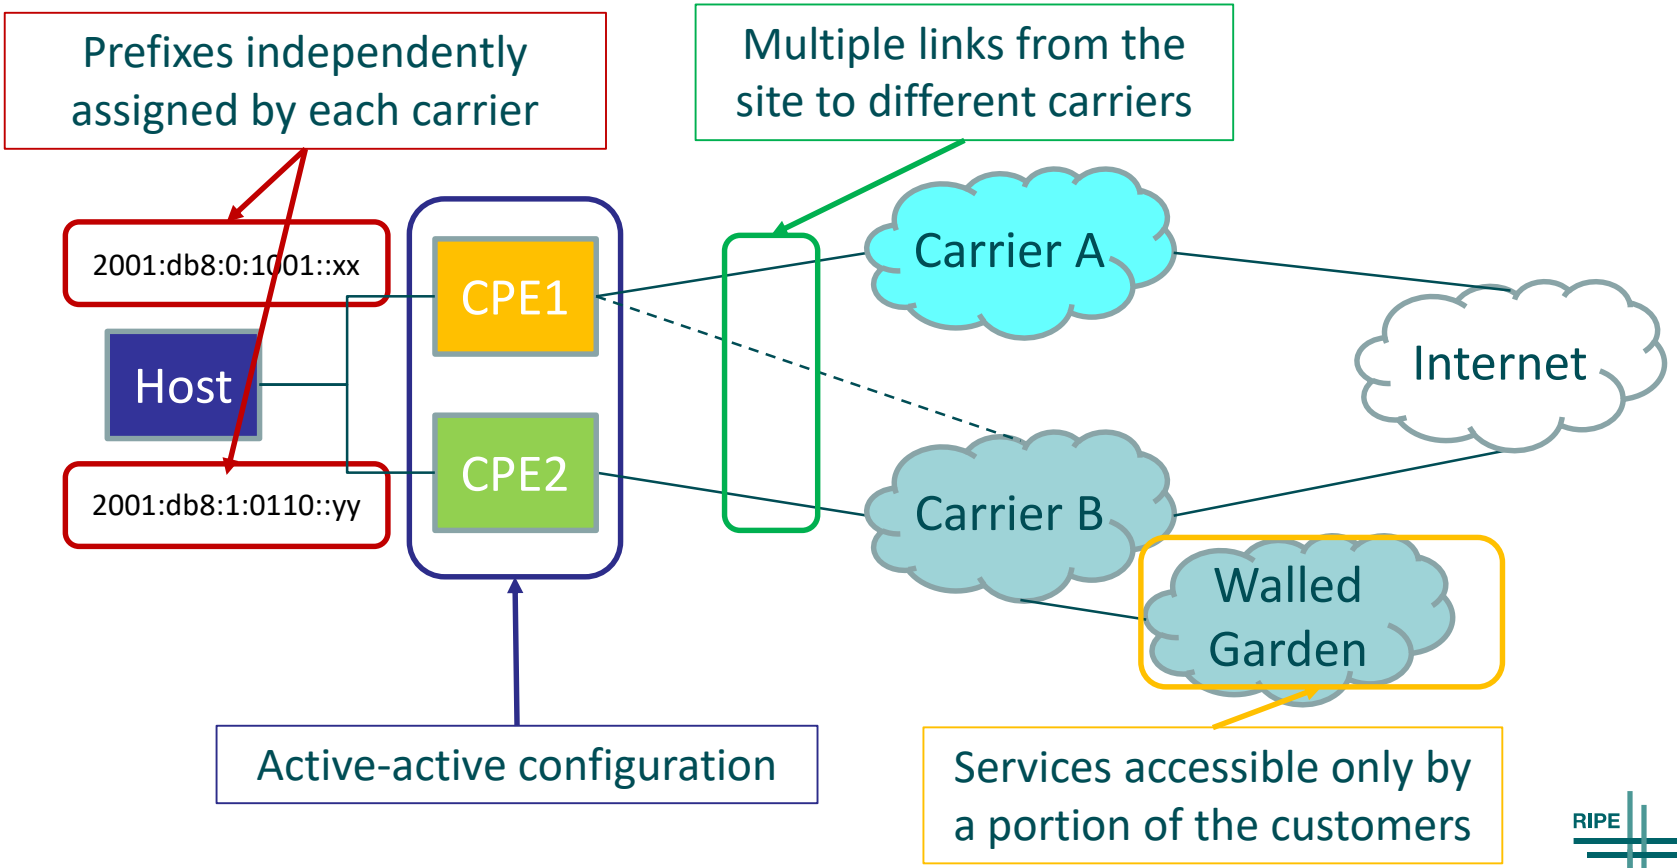

The Role of the Host in Multi-Homing

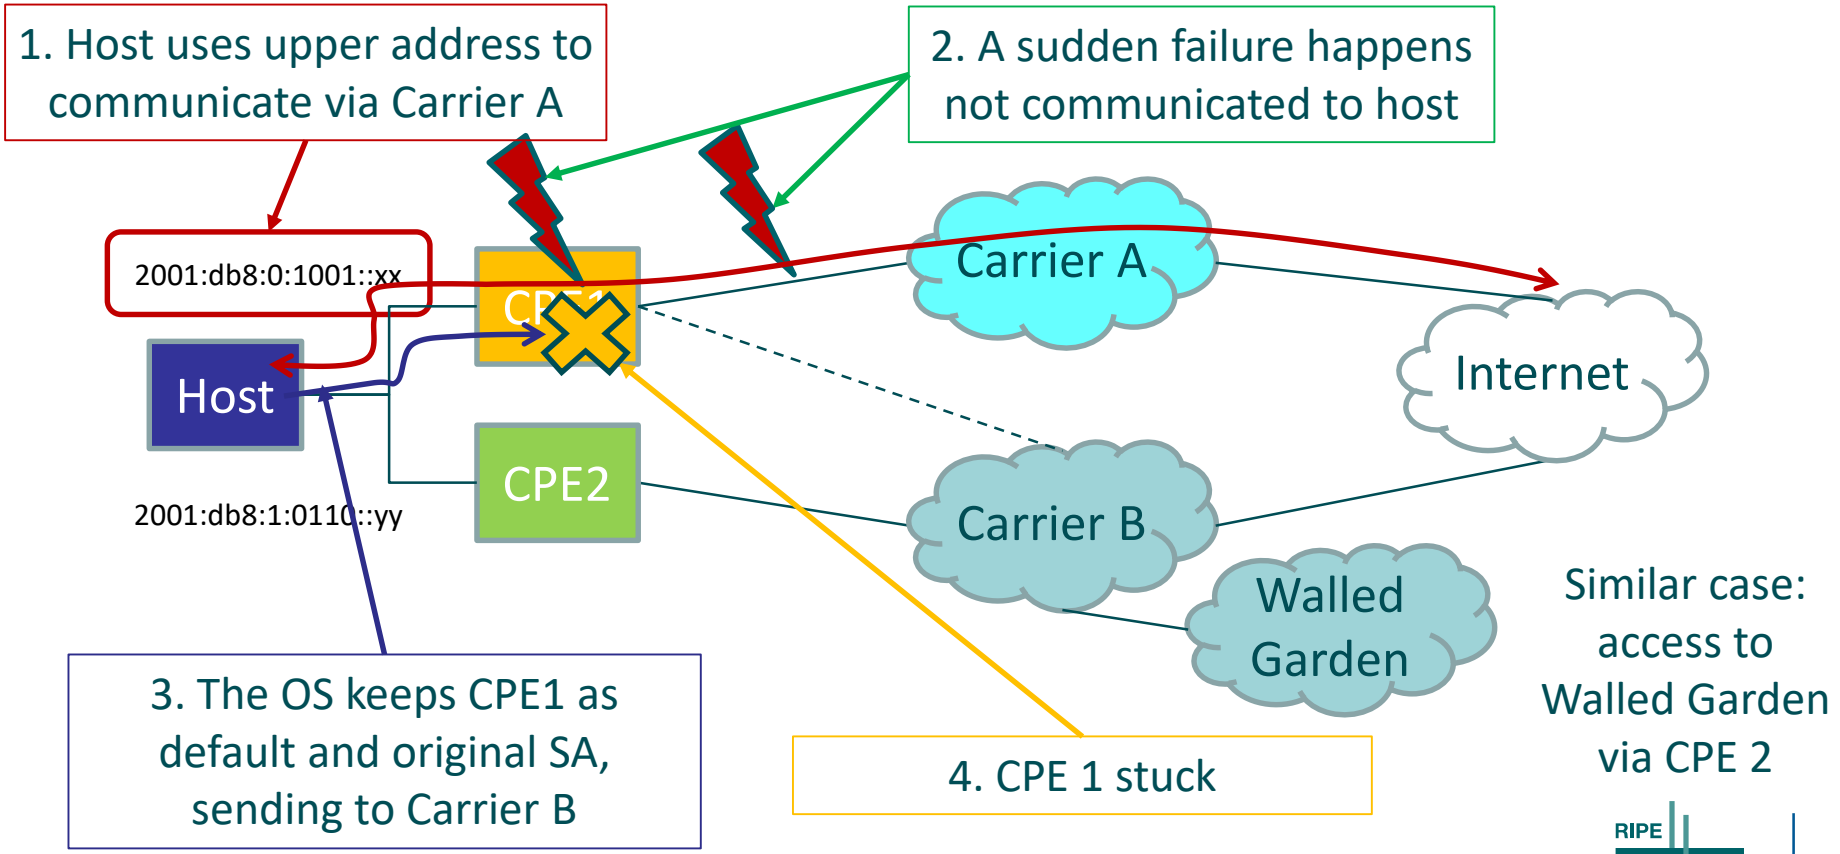

Solutions Considered

Method	Description
Static PI address space to the site	Provider Independent (PI) addresses are allocated to the site, while routing announcements are propagated by carriers on behalf of the client
Dynamic PA addresses distribution from carriers	An IPv6 host gets different Provider Aggregatable (PA) addresses for its interfaces, possibly from different carriers. It is the host that properly chooses the combination of a source address and the relevant next hop to communicate with the destination
Static ULA with NPTv6	Unique Local Addresses (ULA) assigned to the site, then NPTv6 translation is adopted to communicate with the external destination
Static ULA with NAT66	As the previous one, but NAT66 translation is combined with ULA
Move access resilience to a hub site	A branch site is granted redundant connectivity to a central hub location where the aspects related to resilient Internet connectivity are handled
Application proxy	Combines the need for policy/authentication/traffic filtering with Internet access for clients

How to Compare the Mechanisms

- Requirements reflect section 3.1 of RFC 3582
- All solutions have different advantages and disadvantages 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** that distributed such prefixes	**The speed of convergence for the prefix deprecation on the site**, after connectivity is lost to any particular carrier, should be comparable to the speed of routing convergence on the site
Support for sites with complex topologies, including multiple internal on-site hops requiring many routers and links	**Access to carrier's "subscriber-only services"** allowed using the address space distributed by the particular carrier. A given host may need to choose the correct source address accepted by the particular carrier
Possibility for **traffic steering** between different paths (including both internal to the site, and the Internet) based on bandwidth, cost, load, latency, etc.	……..

Pros and Cons (Summary)

Method	Pros	Cons
Static PI	Preserves E2E communication, no special host functionality, seamless link failover without transport session re-establishment	Hard to implement for smaller entities, cost of PI space/operations, impact on Internet routing table
Dynamic PA	No need to own PI space, preserves E2E communication, common configuration	Not all issues resolved yet, prefixes may not get deprecated when the CPE fails, complex topologies not well supported yet
ULA+NPTv6	Easy to implement, similar to current IPv4 carrier resiliency techniques	Breaks applications with address referrals, loses the E2E connectivity advantage
ULA+NAT66	Easy to implement, equivalent in practice to current IPv4 carrier resiliency	Breaks applications with address referrals, session initiation blocked from the outside, stateful processing
Hub site	General simplification of network config, no need for special support on hosts	Expected latency increase, more capacity renting, hub becomes single point of failure
Proxy	No NAT, comm terminated and re-established at higher layer using different source address	No E2E, application proxy is an additional point of failure, requires explicit config

Analysis (not the Definitive Guideline)

	Requirement	PI	PA	ULA+ NPTv6	ULA+ NAT66
1	Carriers Resiliency	+	+	+	+
2	End-to-End Connectivity	+	+	+/-*1	-
3	Internal Connectivity	+	+	+/-*2	+/-*2
4	Convergence Speed	+	+/-*3	+	+
5	Complex Topology Support	+	+/-*7	+/-*4	+
6	Subscriber-only Services	-	-	+/-*5	+/-*5
7	Traffic Steering on Router	+/-*6	+/-*7	+	+

*1. Permits initiating connectivity in any directions; address references in application need special treatment

*2. Complexity in promoting ULA above IPv4 in policy table of hosts

*3. DHCP-PD not adopted yet but needed for prefix deprecation propagation

*4. May depend on prefix length

*5. Needs “Routing Information Options” of Route Preferences, not widely supported

*6. Complexity in organizing the steering of incoming traffic

*7. Complex configuration

Poll on IPv6 Multi-homing

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

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

15 responses

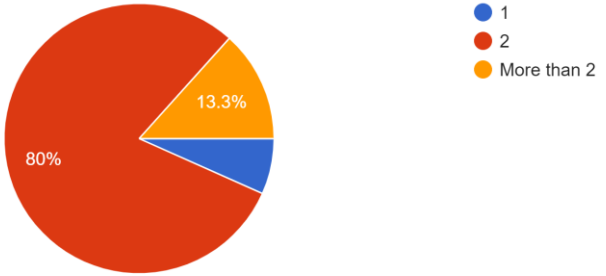

Q2. How many uplinks are configured per CPEs/CEs in your average site?

15 responses

Q3. Which configuration do you support?

14 responses

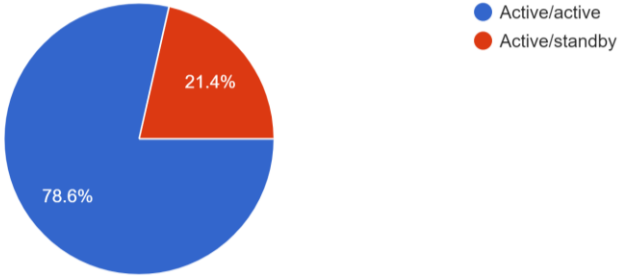

Q5. If the answer to the previous is "No", which method do you employ for supporting IPv6 multi-homing?

10 responses

Key Takeaways

- On a theoretical perspective, PI addressing is preferred over PA, in turn they are both preferred over ULA+NxT
  - If PI widely adopted, consequences may arise
- This does not consider other local factors
  - Many other non-technical requirements could be added to the table that may change the decision logic, including cost
  - More requirements to be considered?
- We would like to have your feedback from your operational experience
  - Please contribute to the survey
  - Feel free to comment on the relevant mailing lists or let's have coffee together.

Questions?

References

1. D. Meyer, L. Zhang, K. Fall, "Report from the IAB Workshop on Routing and Addressing", RFC 4984, https://www.rfceditor.org/info/rfc4984
2. J. Abley, B. Black, V. Gill, "Goals for IPv6 Site-Multihoming Architectures", RFC 3582, https://www.rfc-editor.org/info/rfc3582
3. D. Thaler, R. Draves, A. Matsumoto, T. Chown, "Default Address Selection for Internet Protocol Version 6 (IPv6)", RFC 6724, https://www.rfc-editor.org/info/rfc6724
4. M. Stenberg, S. Barth, P. Pfister, "Home Networking Control Protocol", RFC 7788, https://www.rfc-editor.org/info/rfc7788
5. T. Chown, J. Loughney, T. Winters, "IPv6 Node Requirements", RFC 8504, https://www.rfc-editor.org/info/rfc8504
6. F. Baker, C. Bowers, J. Linkova, "Enterprise Multihoming Using Provider-Assigned IPv6 Addresses without Network Prefix Translation: Requirements and Solutions", RFC 8678, https://www.rfc-editor.org/info/rfc8678
7. F. Baker, B. Carpenter, "First-Hop Router Selection by Hosts in a Multi-Prefix Network", RFC 8028, https://www.rfc-editor.org/info/rfc8028