IPv6 Site connection to many Carriers

draft-fbnvv-v6ops-site-multihoming

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

The problem: the Internet Redundancy

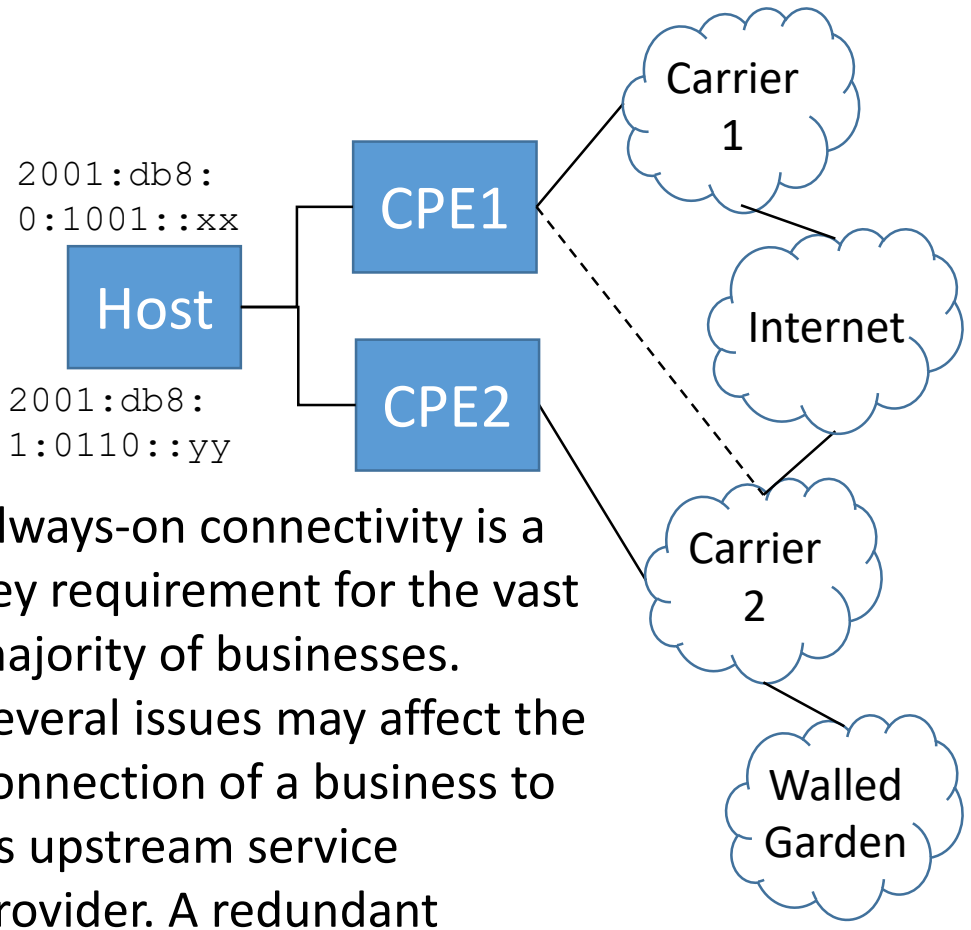

Always-on connectivity is a key requirement for the vast majority of businesses. Several issues may affect the connection of a business to its upstream service provider. A redundant connection to two or more providers is becoming the norm for business.

Solution requirements:

1. Site resiliency to the arbitrary number of carriers.
2. End-to-end connectivity,
3. Possibility for internal communication using any prefixes distributed by local routers, irrespective of the connectivity to carriers.
4. Sub-second convergence on the site after the connectivity is lost to the particular carrier.
5. The potential complex topology of the site with many internal on-site hops (that needs many routers and links)
6. Access resources on the carrier's “walled garden” that is permitted only for the address space distributed by the particular carrier.
   It may need the host to check DNS resolvers from all carriers to discover the restricted resource.
   It may need to choose the source address that would be accepted by the particular carrier.
7. Possibility for traffic steering between different paths (including site internal and the Internet) based on bandwidth, cost, load, latency, packet loss, hop count, etc.

PI-based solution

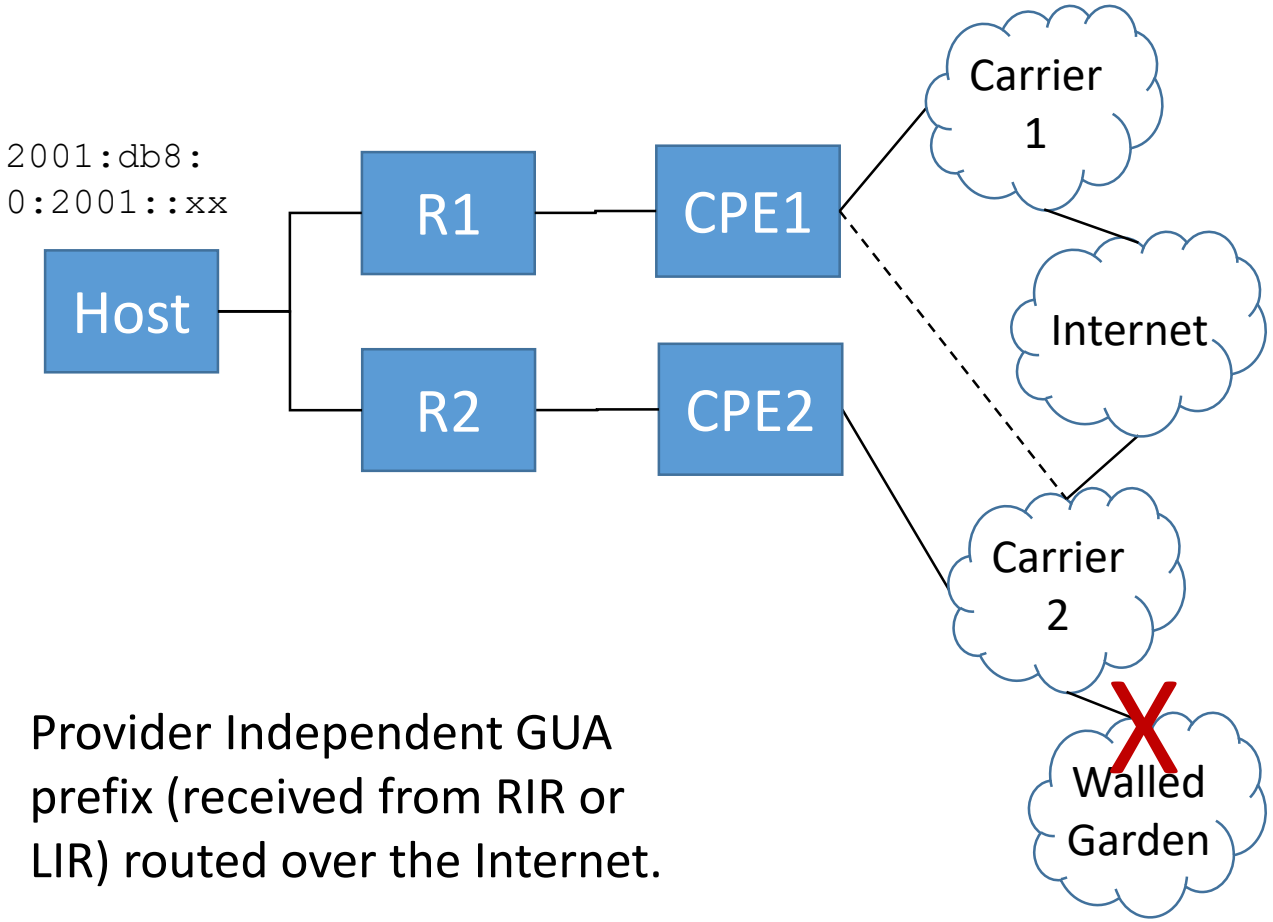

Provider Independent GUA prefix (received from RIR or LIR) routed over the Internet.

Advantages:

- Preserve end-to-end communication,
- Do not request any functionality from the host,
- Easy to implement,
- Support sites with complex topology.

Disadvantages:

- Need to pay and liaise with RIR or LIR (for the PI address space),
- Need carriers to accept PI prefix advertisements, carriers typically charge more for such type of attachment,
- Blow Internet routing table.

PA-based solution

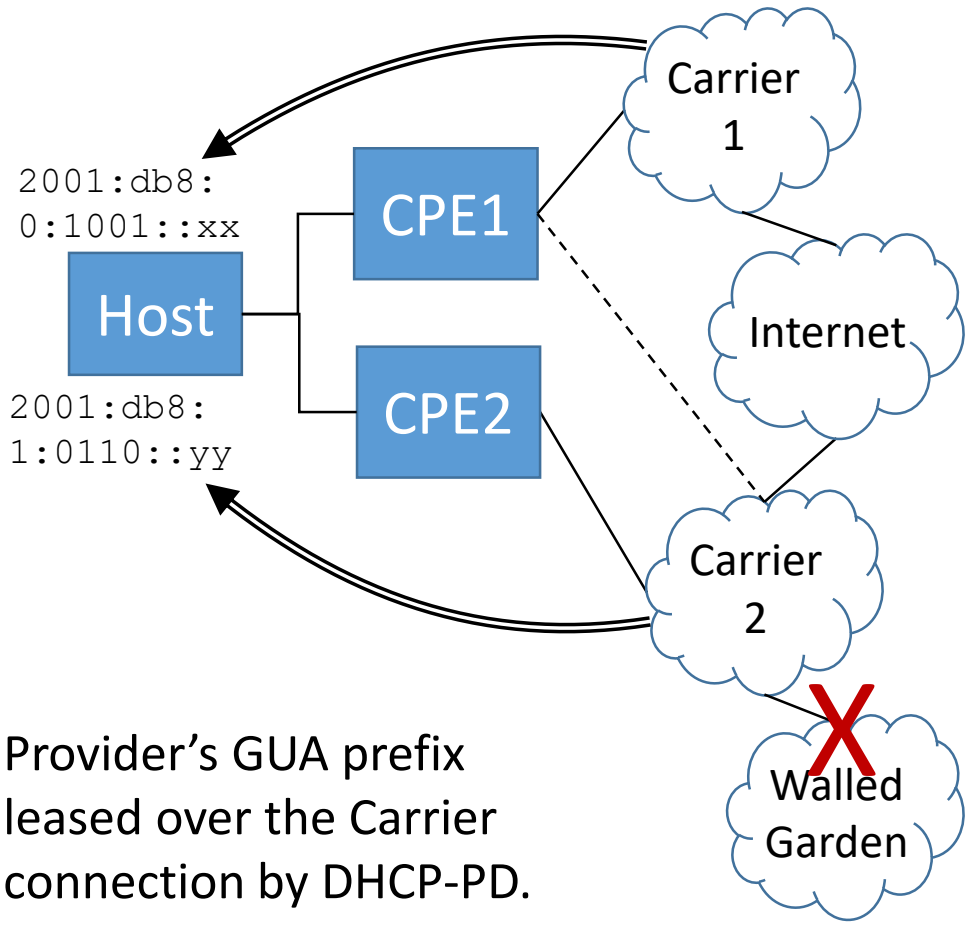

Provider's GUA prefix leased over the Carrier connection by DHCP-PD.

Advantages:

- No need for a registered address space,
- Preserve end-to-end communication.

Disadvantages:

- Very complex,
- Not all issues are resolved yet, only the simplest scenario possible (simple topology, unpredictable traffic distribution),
- Carrier may frequently change the prefix that may disrupt communication (flush renumbering),
- Sites with complex topologies are not supported yet,
- Traffic steering by any policies (including the capability to access a "walled garden") is not supported yet.

ULA+NPT solution

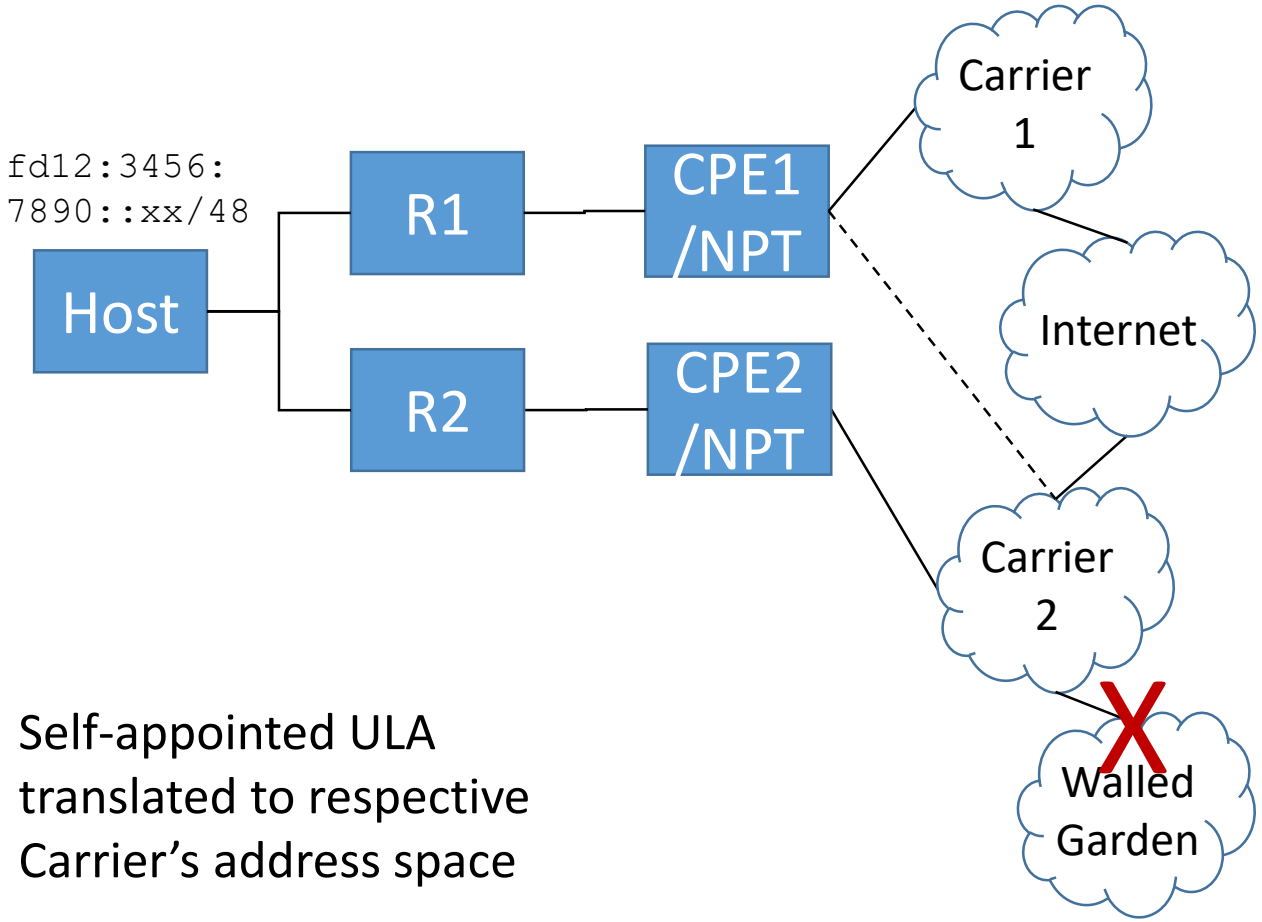

Self-appointed ULA translated to respective Carrier's address space

- Advantages:
- No need for official address space, the ULA prefix is a random self-generated,
- Easy to implement, the same practice as for the current IPv4 carrier resiliency,
- Possibility for traffic distribution policy between different carriers, but without the possibility to access a filtered resource ("walled garden").

Disadvantages:

- Challenge to automate ULA prioritization on hosts above IPv4,
- NPT may break some applications with address referrals at the application level, some additional solutions may be needed (STUN, ALG),
- Session initiation from the outside is practical only for cases when the carrier prefix is stable or DNS is dynamically updated,
- Currently limited to one subnet per site in the mobile environment.

ULA+NAT solution

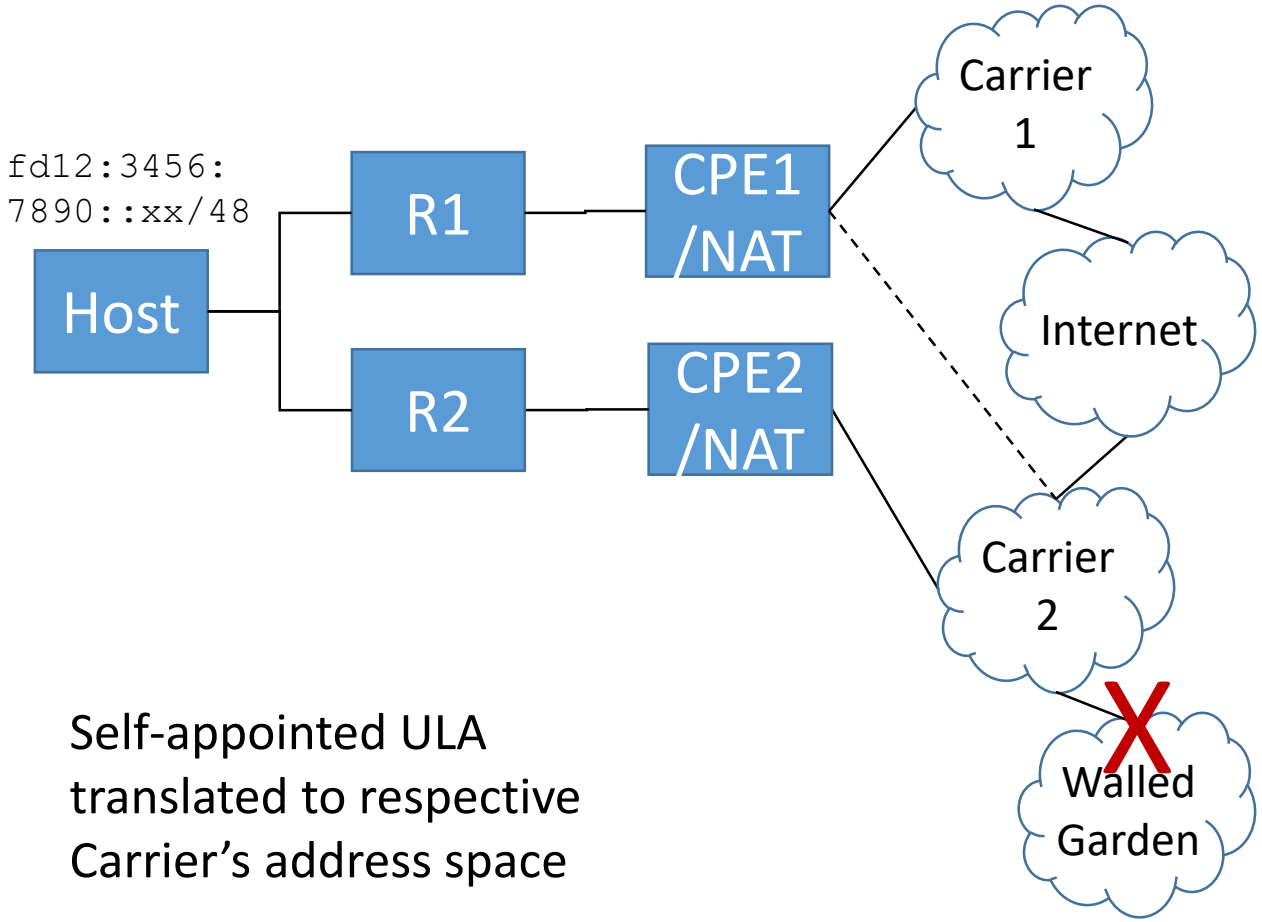

Self-appointed ULA translated to respective Carrier's address space

Advantages:

- No need for official address space, the ULA prefix is a random self-generated,
- Easy to implement, the same practice as for the current IPv4 carrier resiliency,
- NAT may be a normative requirement by itself (that is very disputable but claimed in many discussions anyway),
- Support for sites with complex topologies,
- Possibility for traffic distribution policy between different carriers, but without the possibility to access a filtered resource ("walled garden").

Disadvantages:

- Challenge to automate ULA prioritization on hosts above IPv4,
- NAT may break some applications with address referrals at the application level, some additional solutions may be needed (STUN, ALG),
- Session initiation from the outside is blocked in practice (needs complex configuration),
- NAT needs logs for compliance and troubleshooting,
- Principally bigger cost because of stateful processing.

Solutions Considerations

	Requirement	PI	PA	ULA+NPT	ULA+NAT
1	Carriers Resiliency	+	+	+	+
2	End-to-End Connectivity	+	+	+/-*1	-
3	Internal Connectivity	+	+	+	+
4	Convergence speed	+	+/-*2	+	+
5	Complex Topology support	+	-	+/-*3	+
6	Walled Garden Access	-	-	+/-*4	+/-*4
7	Traffic Steering on Router	+	-	+	+
7	Traffic Steering on Host OS	-	-	-	-
7	Traffic Steering on Application	-	-	-	-

Reasons for partial support:

*1. It permits initiating connectivity in any direction but address references in the application layer would need special treatment like ALG or STUN

*2. HNCP or DHCP-PD has not been adopted by the market but it is needed for prefix deprecation propagation over a complex site

*3. It is not possible (on real products) to get bigger than /64 external prefix in the mobile environment

*4. It needs a routing announcement as “Routing Information Options” of Route Preferences which is not widely supported

- On the pure technical perspectives
  “PI” is preferred over “PA” that is preferred over “ULA+NPT” that is preferred over “ULA+NAT”,
- If IPv6 E2E connectivity is a value then only “PI” or “PA” solutions are acceptable,
- Many other non-technical requirements could be added to the table that may change the decision logic (for example, NAT may be perceived as security or regulatory requirement).

Conclusion

- IPv6 still has a lot of operational issues to work out.
- The issue of redundancy for business is a key challenge.
- Choosing among the available solutions is mandatory for almost any business migrating to IPv6. It is a very complex process now because information is scattered over dozens of RFCs.

- In the short term:
  - Look for more co-authors

- In the medium term:
  - Submit the draft in v6ops to discuss