

# Recommendation for Prefix Binding in the Software DS-Lite Context

**draft-vinapamula-softwire-dslite-prefix-binding**

IETF 89-London, March 2014

S. Vinapamula ([sureshk@juniper.net](mailto:sureshk@juniper.net))

M. Boucadair ([mohamed.boucadair@orange.com](mailto:mohamed.boucadair@orange.com))

**Presenter: Y. Lee ([Yiu\\_Lee@Cable.Comcast.com](mailto:Yiu_Lee@Cable.Comcast.com))**

# Problem

- ISPs want to limit the usage of AFTR resources on per-subscriber basis for fair usage of resources
  - Examples of policies: Preserve external IPv4 address assigned in the AFTR, Port Quota, PCP mappings, etc.
  - These policies are used for dimensioning purposes and to minimize the risk of AFTR resource exhaustion
  - Relying on the B4 address is not sufficient nor reliable (multiple softwires can be established, B4 address may change, etc.)
- When the B4 IPv6 address changes, associated mappings created in the AFTR are no more valid
  - Stale mappings hanging around in the system, consume not only system resources, but also reduce the available quota of resources per subscriber
  - Maintaining these stale mappings may result in the creation of a new set of mappings
- When services are hosted behind B4 element, these services have to advertise about their change, whenever there is a change of the B4 address
  - Means to discover the change of B4 address are required
  - Needed to trigger updates to a rendez-vous server

# Introducing Prefix-Mask

- Prefix-Mask is an AFTR system-wide configuration parameter
- Prefix-Mask is an integer that indicates the length of significant bits to be applied on the source IPv6 address (internal side) to identify a subscriber
  - Generic per-subscriber policies are applied based on the Prefix-Mask
  - Does not require to configure every subscriber prefix
- Prefix-Mask must be configurable
  - Default value is 56

# Recommendations (1)

- A policy SHOULD be enforced at the AFTR to limit the amount of active softwires per subscriber
  - The default value is 1
- Resource contexts created at the AFTR level SHOULD be based on the Prefix-Mask, and not based on the full B4 address
  - Administrators SHOULD configure per-subscriber limits of resource usage, instead of per-tunnel limits
  - These resources include: number of flows, maximum authorized mappings including PCP, NAT pool resources, etc.

# Recommendations (2)

- If a new IPv6 address is assigned to B4, the AFTR SHOULD migrate existing state to be bound to the new B4's IP address
  - This ensures the traffic destined to the previous IPv6 address will be redirected to the new IPv6 address
  - The destination address for the encapsulated return traffic SHOULD be the last seen address from the CPE (i.e., matching the same Prefix-Mask)
- *Justifications*
  - *Avoid stale mappings*
  - *To minimize the risk of service disruption*

# Recommendation (3)

- In the event of change of the CPE WAN IPv6 prefix, unsolicited PCP ANNOUNCE messages SHOULD be sent by the B4 element to internal hosts to update their mappings
  - *This is valid for PCP-enabled CPEs*
  - *Justifications:*
    - *Allows internal PCP clients to update their mappings with the new B4 IPv6 address*
    - *Trigger updates to rendez-vous server (e.g., dyndns)*

# Recommendation (4)

- When a new prefix is assigned to the CPE, stale mappings may exist in the AFTR. To avoid such issues, stable IPv6 prefix assignments are **RECOMMENDED**
  - *Justification: Stable prefix assignment allows to avoid consuming both implicit and explicit resources*

# Recommendations (5)

- In case an IPv6 prefix has to be reassigned for any reason, it is **RECOMMENDED** to reassign a prefix only when all the resources in use associated with that prefix are cleared from the AFTR
  - *Justification: Avoid to redirect traffic, destined to the previous prefix owner, to the new one*

# Next Step

- This short document provides a set of recommendations aiming to:
  - Enhance DS-Lite serviceability
  - Ease AFTR resources management
  - Enforce generic per-subscriber policies without requiring explicit configuration of every CPE IPv6 prefix to the AFTR, nor any additional interfaces (e.g., RADIUS)
- This is a missing piece of work
- Request WG adoption