Considerations on NAT64 Load-Balancing

draft-zhang-behave-nat64-load-balancing

Dacheng Zhang
Xiaohu Xu
Mohamed Boucadair
Xuewei Wang
Yan Wang
Cameron Byrne
Dong Zhang

Change Log

- This I-D is a merge of http://tools.ietf.org/html/draft-wang-behave-nat64-load-balancer-02 and http://tools.ietf.org/html/draft-xu-behave-stateful-nat-standby-05#section-5

- New text structure

Overview

- **Objectives**
  - Investigate NAT64 load-balancing approaches
  - Analyze the advantages and disadvantages of various prefix selection policies
- **Scope**
  - Both stateless and stateful NAT64 schemes
  - Both Inbound and Outbound load-balancing
  - Failover and redundancy are out of scope

Methodology

- Identify a list of **requirements**
- Identify and describe **various alternatives** to implement load-balancing
  - Stateless NAT64: e.g., Anycast-based, DHCPv6-based and NAT64 farm
  - Stateful NAT64: Anycast-based or rely on various methods for prefix selection such as: source-based, destination-based, Round-Robin or dynamic
- For each load-balancing option, **Pros & Cons** are listed

Requirements Excerpt

- The operations of distributing the load among multiple NAT64 devices should be covered from end-users
  - Load balancing should not lead to (severe) QoS degradation between potential paths
  - The introduction of load balancing function should not change the operations of end-users (i.e., should be transparent)
- An efficient load balancing system should not redirect the traffic to a congested NAT64 device while other NAT64 resources are available

Requirements Excerpt

- The outgoing traffic and the associated incoming traffic should be stuck to the same stateful NAT64 device
- The assignment of the same external IPv4 address should be preserved for all active sessions initiated by an IPv6-only host
- A load balancing solution should be deterministic
  - A load balancing system should execute according the operator's expectation when she configures the system

Prefix64 Selection Policy

- Source-Based Prefix Selection Policy
- Destination-Based Prefix Selection Policy
- Round-Robin Prefix Selection Policy
- Dynamic Prefix Selection Policy

Load Balancer Implementing Options

- DNS64 Servers -- Synthesizing AAAA RR according a Prefix64 Selection Policy and returning it to the requestor
- Prefix64 Assigners -- After getting a Prefix64 from an assigner, a IPv6 only hosts can synthesize appropriate IPv6 address locally
  - DNS64 Servers (We have implemented an experimental system to evaluate the performance)
  - Default Gateways
  - DHCPv6 Servers
  - IPv6 Clients

To Do List

- Add a recommendation section
  - Based on the Pros & Cons analysis
  - Some experiments are undertaken to help assessing the validity of some proposed options
- Example of the recommendation would be (see the appendix for more details)

Stateful	Identical External IPv4 Pools	Dedicated External IPv4 Pools
Identical Prefix64s	Works only when sync is enabled	Works but have oscillation issues
Dedicated Prefix64s	Infeasible due to asymmetric path issues	**Recommended**

Next Steps

- Comments, suggestions and contributions to enrich the document are more than welcome
- Complete the recommendation section
- Handle (hopefully) received comments

Question?

Appendix

Load Balancing Deployment Options (1)

- **Stateful NAT64s**

	Identical External IPv4 Pools	**Dedicated External IPv4 Pools**
Identical Pref64s	Works only when sync is enabled	Works but have oscillation issues
Dedicated Pref64s	Infeasible due to asymmetric path issues	**Recommended**

- It is difficult to determine the paths that traffic between two hosts passes through
- Synchronization of mapping information between two NAT64s is required

Load Balancing Deployment Options (2)

- **Stateful NAT64s**

	Identical External IPv4 Pools	**Dedicated External IPv4 Pools**
Identical Pref64s	Works only when sync is enabled	Works but have oscillation issues
Dedicated Pref64s	Infeasible due to asymmetric path issues	**Recommended**

- When the topology of the IPv6 network is modified, the path that the traffic between two hosts may change. In such a case, the communication between hosts will be broken

Load Balancing Deployment Options (3)

- **Stateful NAT64s**

	Identical External IPv4 Pools	Dedicated External IPv4 Pools
Identical Pref64s	Works only when sync is enabled	Works but have oscillation issues
Dedicated Pref64s	Infeasible due to asymmetric path issues	**Recommended**

- The incoming and outgoing traffic may pass different NAT64s
- Even if the mapping information is provided, the incoming traffic will be associated with a Prefix64 different from what the receiver expects
- It is difficult to avoid the issue of assign the same IPv4 address for different IPv6 hosts

Load Balancing Deployment Options (4)

- **Stateful NAT64s**

	Identical External IPv4 Pools	**Dedicated External IPv4 Pools**
Identical Pref64s	Works only when sync is enabled	Works but have oscillation issues
Dedicated Pref64s	Infeasible due to asymmetric path issues	**Recommended**

- Asymmetric path issue can be avoided

Load Balancing Deployment Options (1)

- **Stateless NAT64s**

	Identical External IPv4 Pools	Dedicated External IPv4 Pools
Identical Pref64s	**Recommended**	works but may fail when uRPF is deployed in the IPv4 network
Dedicated Pref64s	Infeasible due to asymmetric path issues	works but the prefix64 selection policy is required

- The outgoing traffic from a IPv6 hosts to an IPv4 host may pass through different NAT64s
- Therefore, if uRPF is deployed, some packets may be regarded as suspicious and then be discarded

Load Balancing Deployment Options (2)

- **Stateless NAT64s**

	Identical External IPv4 Pools	Dedicated External IPv4 Pools
Identical Pref64s	**Recommended**	works but may fail when uRPF is deployed in the IPv4 network
Dedicated Pref64s	Infeasible due to asymmetric path issues	works but the prefix64 selection policy is required

- The reason is identical to stateful NAT64s

Load Balancing Deployment Options (3)

- **Stateless NAT64s**

	Identical External IPv4 Pools	**Dedicated External IPv4 Pools**
Identical Pref64s	**Recommended**	works but may fail when uRPF is deployed in the IPv4 network
Dedicated Pref64s	Infeasible due to asymmetric path issues	works but the prefix64 selection policy is required

- Source based policy will be required

Load Balancing Deployment Options (4)

- **Stateless NAT64s**

	Identical External IPv4 Pools	Dedicated External IPv4 Pools
Identical Pref64s	**Recommended**	works but may fail when uRPF is deployed in the IPv4 network
Dedicated Pref64s	Infeasible due to asymmetric path issues	works but the prefix64 selection policy is required

Experimental Evaluation (1)

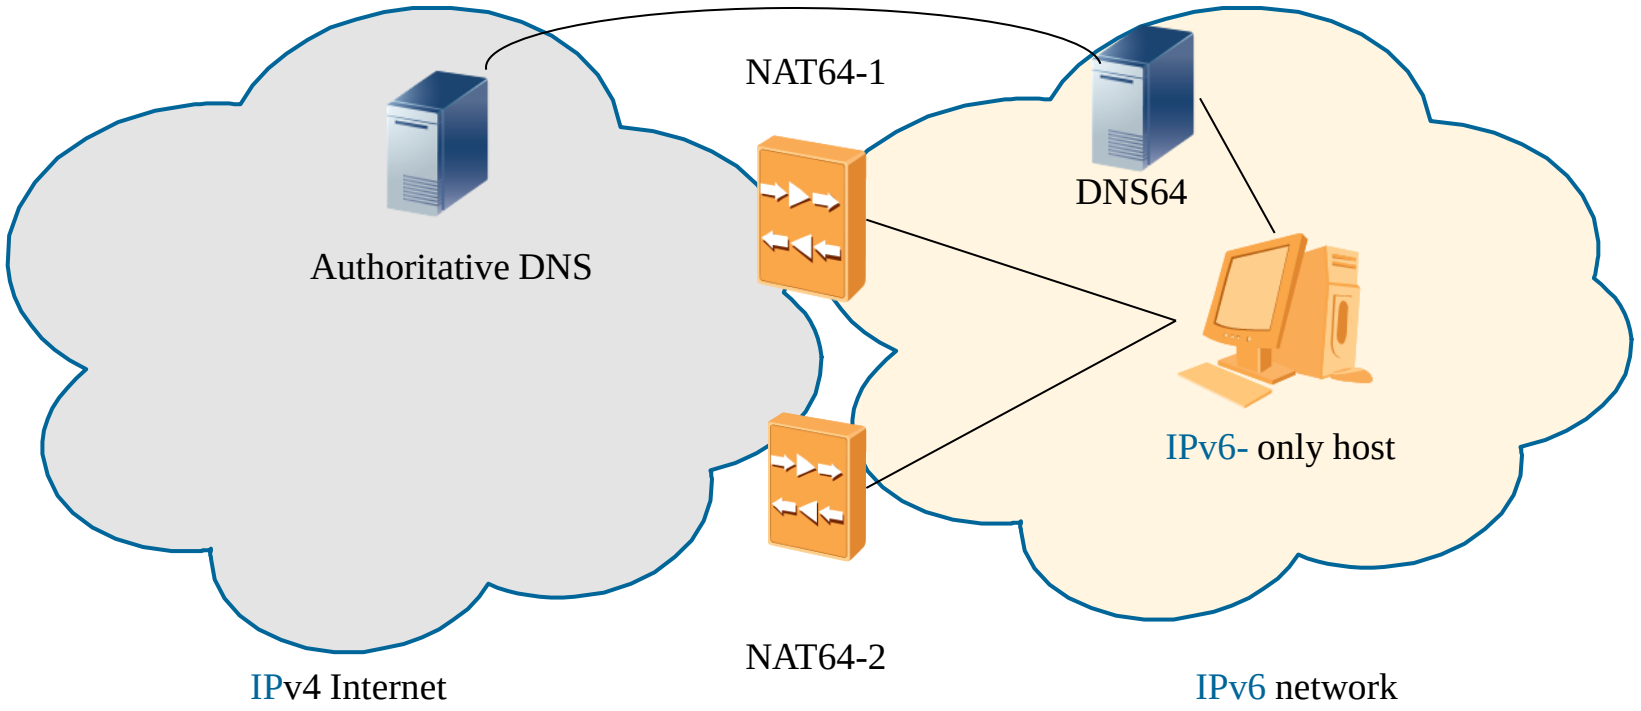

Experimental Evaluation (2)

Application		Source Address Based Policy	Destination Address Based Policy	Round Robin
Instant Messaging	Gmail chat on the web (http)	ok	ok	ok
	MSN Web (http)	ok	ok	ok
Game	Web-based (e.g. armorgames)	ok	ok	ok
Website	Google	ok	ok	ok
	Amazon	ok	ok	ok
	Baidu	ok	ok	ok
	Sina	ok	ok	ok
	Hi5	ok	ok	ok
	BBC	ok	ok	ok

Security Considerations

- All the traffic between an IPv6 host and an IPv4 host should be intercepted and processed by a same NAT64 device so as to benefit intrusion detection/prevision systems to monitor the operations of users
- The source-based prefix selection policy can fulfill this requirement, while the destination-based prefix selection policy and the Round-Robin prefix selection policy cannot fulfill this requirement