

lispers.net NAT Implementation

draft-farinacci-lisp-simple-nat-00

IETF Madrid (Virtual) LISP WG
July 2020

Dino Farinacci

Purpose of Draft

- To describe a simple version of NAT-traversal
- Based on a subset of procedures and message formats from:

draft-ermagan-lisp-nat-traversal

- There has been *lispers.net* deployments since 2014

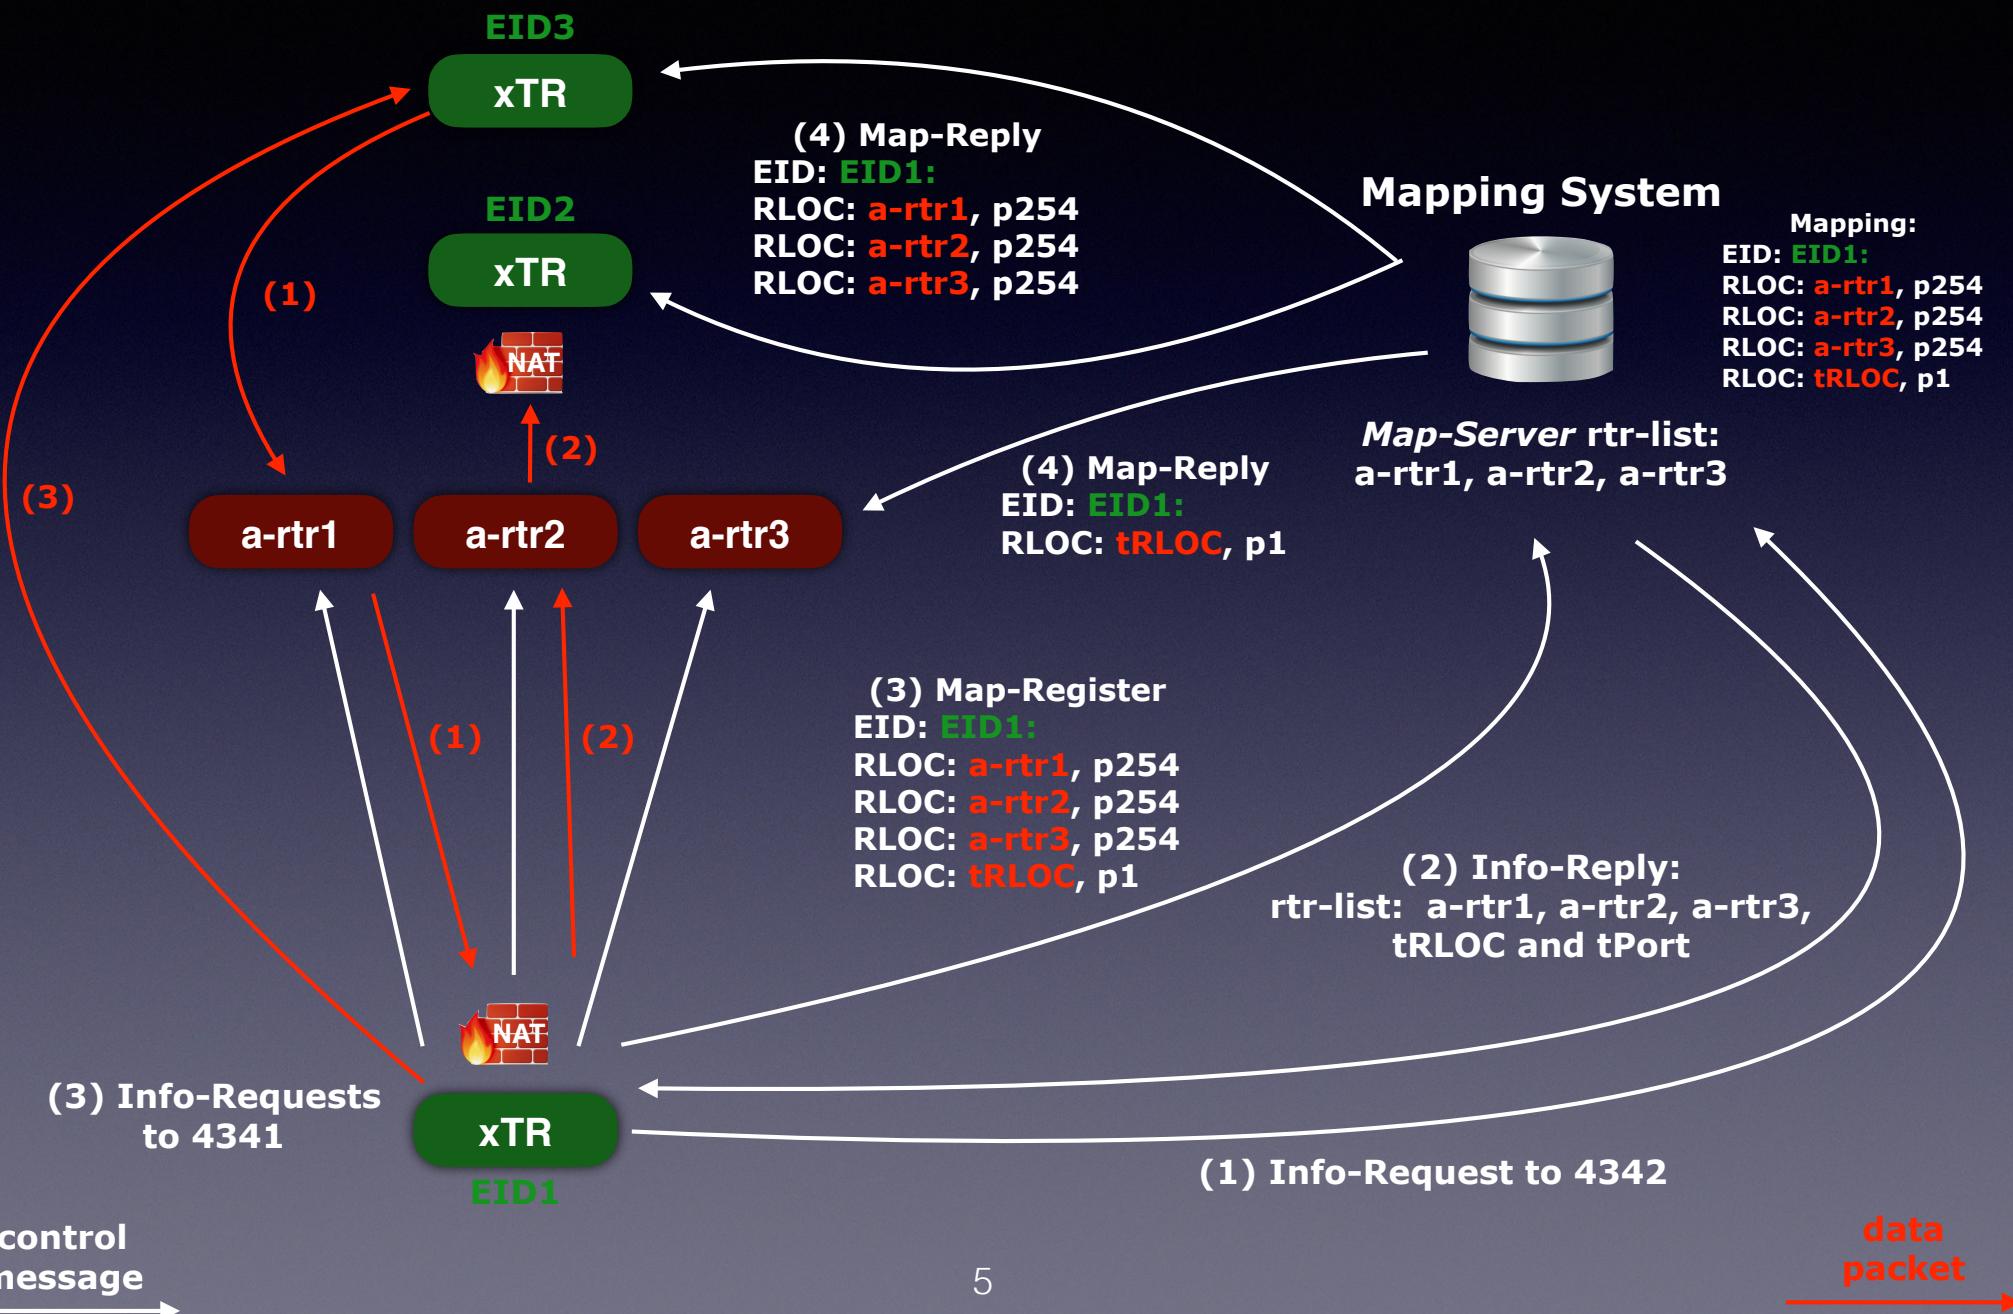
Requirements

- The xTR must function when it resides 1 or more hops away from a NAT device
- The xTR must function when it runs as a LISP-MN, in a VM, or a container (logically or physically)
- The xTR must function when it resides behind multiple NAT devices between itself and an RTR
- The xTR must function when it is multi-homed behind different NAT devices
- RTRs and PxTRs must function behind NATs
- *lig* must function when requests are sent from behind NATs
- IPv4 unicast/multicast and IPv6 unicast/multicast EIDs must be supported
- IPv6 RLOCs are not supported and don't need to be NATed

Protocol Exchange Sequence

- xTR sends Info-Request port **4342** to Map-Server
- Map-Server replies with list of RTRs in Info-Reply
- xTR sends Info-Request port **4341** to each RTR
- NATs install entries so RTRs can send from port **4341** to xTR
- Remote ITRs encapsulate to RTRs
- RTRs encapsulate to translated address and port
- xTRs Map-Register list of RTR RLOCs and its translated RLOC address
- xTRs have default map-cache entries to RTRs
- xTRs can short-cut map-cache entries to “non-NAT xTR” RLOCs

Mapping System Interface



Design Observations

9. Design Observations

The following benefits and observations can be attributed to this design:

- o An ITR behind a NAT virtually runs no control-plane and a very simple data-plane. All it does is RLOC-probe the RTRs in the common RLOC-set for each default map-cache entry. And maintains a very small map-cache of 4 entries per instance-ID it supports.
- o An xTR behind a NAT can tell if another xTR is behind the same set of NAT devices and use Private RLOCs to reach each other on a short-cut path. It does this by comparing the Global RLOC returned from RTRs in Info-Reply messages.
- o An xTR behind a NAT is free to send a Map-Request to the mapping system for any EID to test to see if there is a direct path to the LISP site versus potentially using a sub-optimal path through an RTR. This happens when xTRs exist that are not behind NAT devices where their RLOCs are global RLOCs.
- o By sending Info-Requests to Map-Servers, an xTR behind a NAT can tell if they are reachable and if those Map-Servers also act as Map-Resolvers, the xTR behind the NAT can avoid sending Map-Requests to unreachable Map-Resolvers.
- o Enhanced data-plane security can be used via LISP-Crypto mechanisms detailed in [RFC8061] using this NAT-Traversal design so both unicast and multicast traffic are encrypted.
- o This design allows for the minimum amount of NAT device state since only RTRs are encapsulating to ETRs behind NAT devices. Therefore, the number of ITRs sending packets to EIDs behind NAT devices is aggregated out for scale. Scale is also achieved when xTRs behind NATs roam and RLOC-set changes need to update only RTR map-caches.
- o The protocol procedures in this document can be used when a LISP site has multiple xTRs each connected via separate NAT devices to the public network. Each xTR registers their Global RLOCs and RTRs with merge semantics to the mapping system so remote ITRs can load-split traffic across a merged RTR set as well as RTRs across each xTR behind different NAT devices.

Draft Destiny

Ask if WG okay with Informational RFC?

Questions/Reactions/Tomatoes?

