

# **LISP Map Server Reliable Transport**

## **draft-kouvelas-lisp-map-server-reliable-transport-06**

J. Leong  
D. Lewis  
B. Pitta  
C. Cassar  
I. Kouvelas  
J. Arango

IETF 112 – Online  
Nov, 2021

# Background

- The LISP Map-Server Reliable transport is extensively used in deployments
- Experimentation showed rapid benefits to scale deployments, and it's been key to support operation at scale.
  - Deployment with large number of EIDs
  - Mobility at scale
  - Redistribution of database-mappings to interact with other systems
- In practical terms, since it proposes message reuse, it was implemented as an extension of the registration process

# From Periodic Registration to Reliable Transport

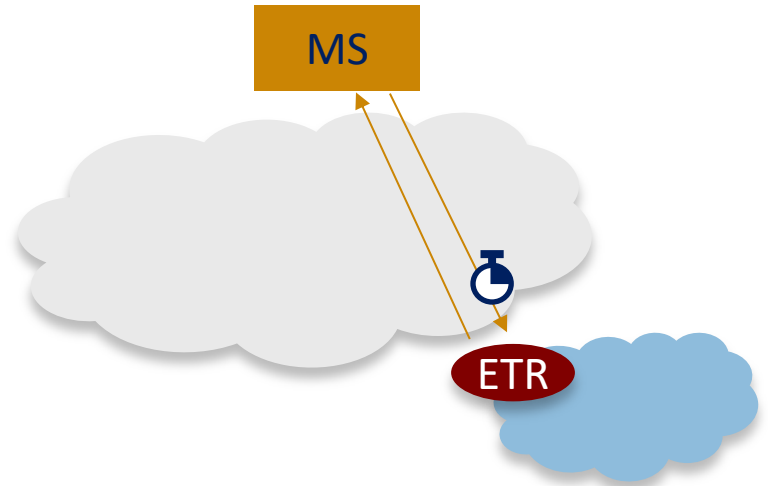
- Periodic UDP communication between xTR and Map-Server to maintain soft state
- Practical concerns in experimentation with some of LISP use cases with large number of EID records (database redistribution & mobility)
- Constant communication load on LISP control plane. Scaling issue with 1000s of records per xTR
- Lack of flow control for communication spanning multiple network hops

# From Periodic Registration to Reliable Transport

- Establish TCP or SCTP based reliable session between the xTR and the Map-Server
- Use session to communicate EID to RLOC registrations and mapping notifications
- Optional alternative to UDP based registration (existing UDP mechanism must be supported)

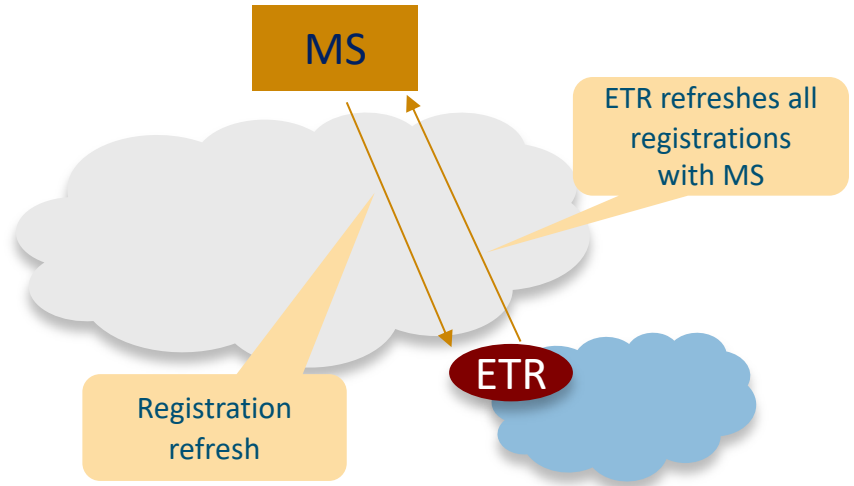
# Reliable Transport operation

- An ETR starts with the periodic UDP registration
- The periodic UDP registration is maintained until an optional reliable session is established




# Reliable Transport operation

- Once the session is established the MS sends a *registration refresh* message to the ETR
- The ETR refreshes all registrations with the MS
- From that instant, registrations are considered active as long as the session is up, and the ETR does not need to re-send them periodically



# Reliable Transport operation

- The MS can reject registrations when it is not ready, or configuration or policy does not allow them.
- The registration-refresh message can be used to notify the ETR about MS changes that may allow new registration attempts
- The registration-refresh message can be scoped (all, IID, specific EID prefix) to trigger specific refreshes.



# Reliable EID registration

- Registration message identical to UDP registration
- Use Positive or negative acknowledgement to provide explicit feedback to the xTR
- Ability for MS to request selective refresh of information to deal with configuration changes (registration-refresh)
- Mapping notifications are no longer needed for registration acknowledgement and just convey latest (potentially merged) Map-Server view of the mapping

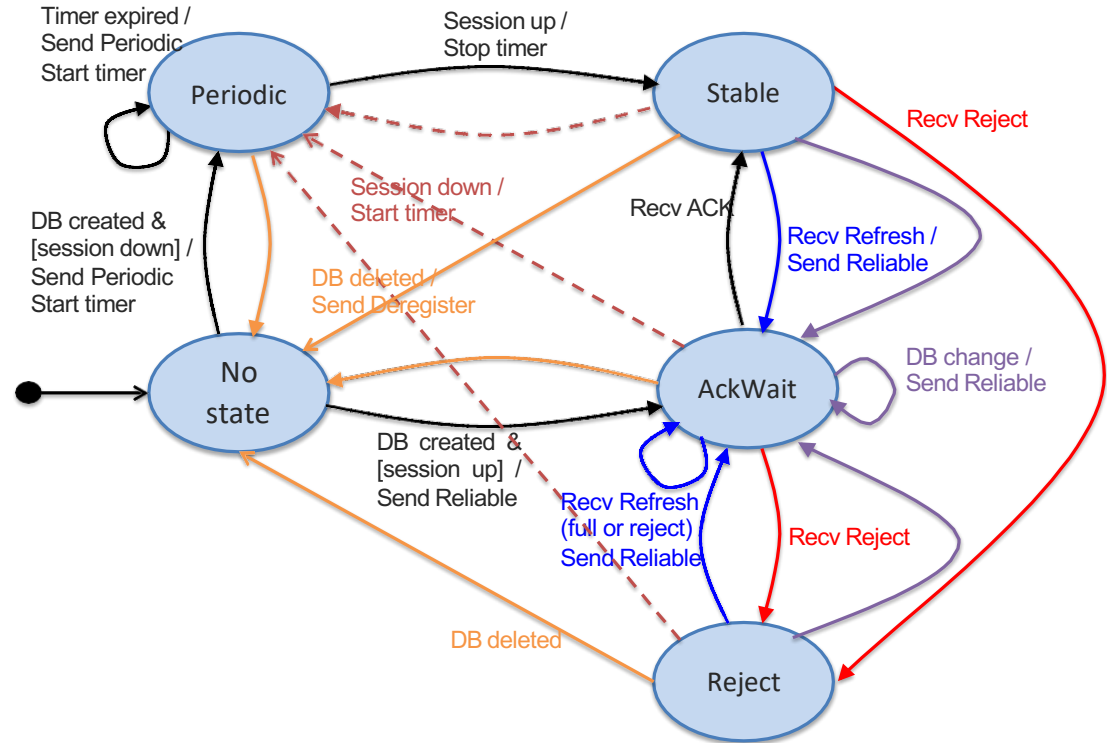


# Map-Server operation

- Received registrations create/update or delete mapping registration state (no timeout)
- Registration state is not discarded when session goes down (falls back to timer-based expiration)
- Registrations are rejected if
  - Authentication fails
  - EID prefix is not configured
  - Mapping locator set not allowed
- Refresh issued to ETR to obtain initial state and to re-request specific prefixes on configuration change. Scope field determine the set of registrations being refreshed.

# ETR Operation

- ETR in periodic UDP registration mode until a reliable transport session is established with the MS
- While in reliable transport mode registrations are only sent in response to refresh requests by the MS



# Implementation notes

- Implementation has been running for a while now with stable code and has been very effective supporting large deployments
- The implementation revealed that only minor extensions are needed to support the EID registration procedure (including message reuse) to run over a reliable transport
- The specification allows running over any reliable transport

## Next Steps

- The Authors would like to request adoption of the document by the WG.

Comments, Questions