

LISP Reliable Transport

draft-kouvelas-lisp-reliable-transport-00

Chris Cassar
Isidor Kouvelas
Darrel Lewis

IETF 89, London, UK

Current Periodic Registration

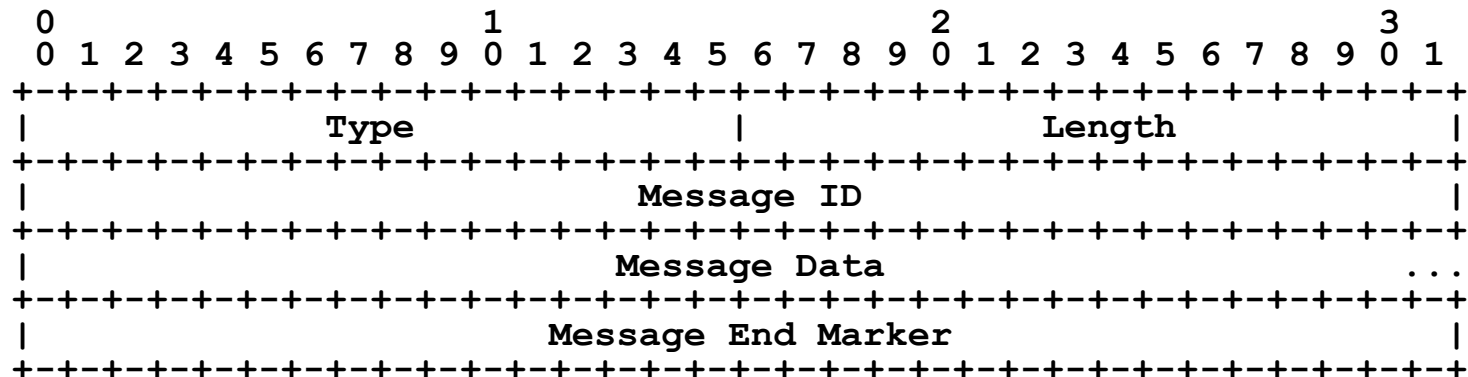
- Periodic UDP communication between xTR and Map-Server to maintain soft state
- Practical concerns in experimentation with new LISP use cases that have large number of EID records (database redistribution & mobility)
- Constant communication load on LISP control plane. Scaling issue with 1000s of records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

Reliable Transport Alternative

- 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)
- Single session between LISP peers can support additional future applications
- TLV message based exchange to support protocol extensibility

TLV Message Format

- 16 bit type and length
- Message-ID used by the receiver to identify the message in replies or notification messages
- End marker allows receiver to validate correct message parsing



Reliable EID Registration

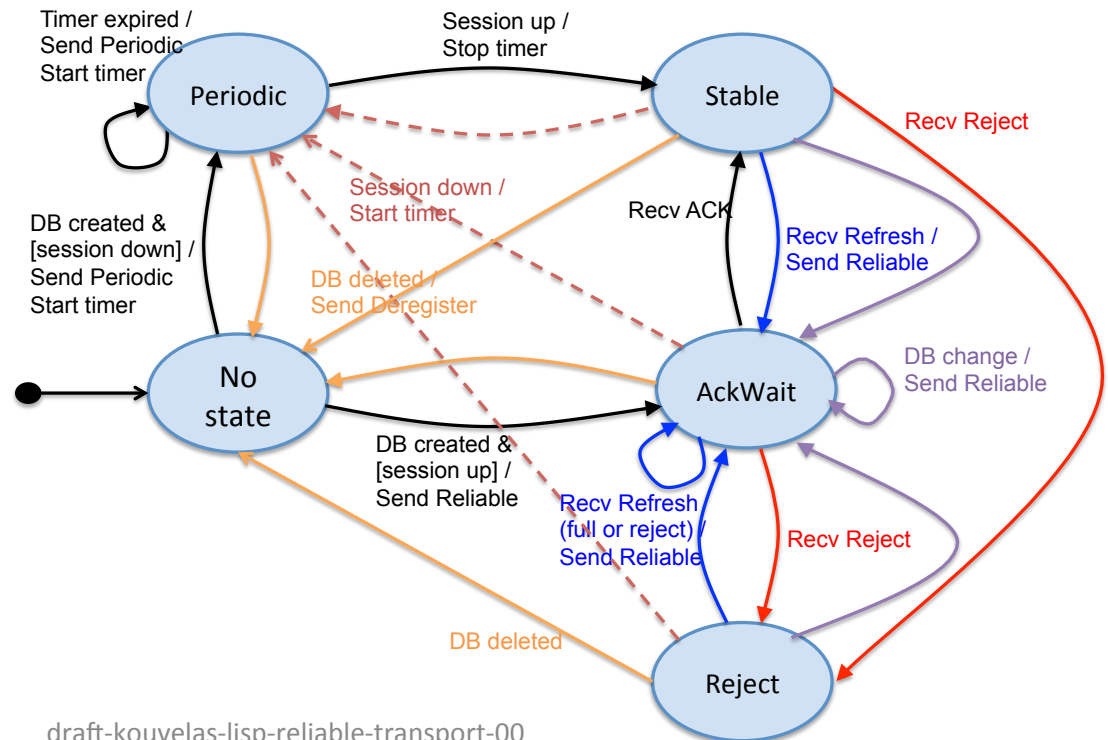
- Registration message identical to UDP registration
- Positive or negative acknowledgement to provide explicit feedback to the xTR
- Ability for MS to request selective refresh of information to deal with configuration changes
- Mapping notifications are no longer needed for registration acknowledgement 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
- Mapping Notification message sent when the mapping for a prefix for which notifications are requested changes

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



Conclusion

- Use reliable transport session between xTR and MS for reliable flow-controlled communication
- Eliminate periodic messaging to improve scaling
- Messages based on extensible TLV format
- Minor extension to EID registration mechanism to work over reliable transport (alternative to UDP registrations)
- Session re-usable for other LISP applications