Providing a Secure Rdma Protocol

Motivation and Initial Discussion

David Noveck

For nfsv4 WG Meeting at IETF-117

July 28, 2023

Overview

- Motivations
- Likely Protocol Layerings
- Document(s) that might be needed
- Implementation work to consider

Motivation
Up Until Now …

- In the past, RDMA security kind of sucked but:
  - It had a lot of company, since we had to rely on per-principal encryption which is hard to offload
  - Even though its security was worse, in that privacy defeated copy elimination, assumed RDMA was only for within-machine –room use.
- Now, non-RDMA has RPC using TLS
  - And there is no RDMA equivalent
  -

Motivation
Now and Going Forward

- Need to address this gap *somehow*
  - Don’t want to have RDMA users to be forced to accept inferior security
- Will be a need for RDMA outside a within-room environment
  - Requires TLS-equivalent security with Iwarp-equivalent performance to support RDMA over distance, including stretched (metropolitan-area-wide) clusters many miles in diameter

Protocol Layering
Simplest choice – adding TLS to TCP

- TCP would be replaced by TLS-over-TCP
  - DDP would be applied to the unencoded/decoded TCP stream
- Would need a negotiation/upgrade mechanism:
  - Could add a new pseudo-flavor to go directly to “TLSwarp”
  - Or could allow upgrade of RPC-with-TLS to RDMA mode
- Probably not doable
  - Buffering issues as explained in RFC5042

Protocol Layering
Alternatve – Replace TLS/TCP by QUIC

- QUIC replaces the combination of TCP/DDP
  - QUIC frame boundaries eliminate the need for DDP
- Need to investigate use of QUIC Streams
  - Will probably look at treatment of SCTP in RFC5040/5043
  - Might even try to handle each chunk as a separate stream
- Possible advantages:
  - Use of multiple paths
  - Better congestion control?
  - Better recovery from lost packets?

Documents that might be Needed

- Mapping document, to indicate how TLS or QUIC would fit in the framework established by RFC5040
  - Would function like NVMe mapping document with respect to SCSI
- Simple extension of RFC5040 to include QUIC , on the same level as TCP and SCTP
  - Plus a new document on same level as RFC 5043/5044.
  - Will probably need a revision of RFC5042 as well
  - More work but a better result
- Looks like RFC7306c an be left alone.

Implementation Work
What will be Needed and When

- Some Implementation work needed ASAP
  - i.e. as soon as basic layering decisions are made
  - Better not to wait for decisions on document approach, and IETF processes to get through adoption, WGLC, and IESG review, and editing/publication
- Desirable to get interoperabillity testing done early.
  - primarily to make sure we haven't missed anything we need from lower layers (e.g. QUIC)

Implementation Work
What do we Have to Start From

- We have an open-source software iWarp implementation with a BSD license.
- We need to combine it with a software QUIC implementation.
  - Need to find out if one exists.
  - There might be licensing issues for some.