

RPC-over-RDMA

Version Two

Chuck Lever
<chuck.lever@oracle.com>

RPC-over-RDMA v2 Goals

- Original goals were narrow:
 - Integrate support for transport properties
 - Enable extensibility via an RFC 8178-like process
- Later, we added:
 - Enable support for generic remote invalidation
 - Richer error reporting

Potential Extensions

- Making use of Send/Receive for in-place data transfer (proposed in `draft-dnoveck-nfsv4-rpcrdma-rtrext`)
- Reducing/eliminating the use of Reply chunks (proposed in `draft-cel-nfsv4-rpcrdma-reliable-reply`)
- Machine authentication similar to TLS (not yet proposed)

RPC-over-RDMA v2 Goals

- One more goal was discussed at IETF 102: Instead of accounting for RPC Call/Reply round-trips, account for RPC messages. This would provide support for:
 - Retransmits and Call-only operation
 - Messages not associated with an RPC, like transport control plane messages
 - Other possibilities arising in future extensions

Prototyping

- Linux NFS client and server support transport properties as a la `draft-ietf-nfsv4-rpcrdma-cm-pvt-data`
- This provides the ability to detect support for basic remote invalidation, also now implemented in Linux NFS client and server
- There is interest in these areas in other implementations
- Extensibility and rich error reporting have not yet been prototyped

Document Status

- Document reached original goals by revision -07
- No recent changes due to:
 - Focus on higher priority matters such as RPC on TLS
 - `draft-ietf-nfsv4-rpcrdma-cm-pvt-data` relieves pressing issues with RPC-over-RDMA v1

Discussion

- Do we need deeper prototype implementations before proceeding?
- Is the document ready for promotion to WG status?
- Will the charter milestone be met, or should it be adjusted?
- Is the initial transport property list complete?
- Should an extension be included in the document?