

“RDMA Commit” Draft status

Tom Talpey

IETF 116 Yokohama

March 31, 2023

draft-talpey-rdma-commit

- Specifies extensions for Flush, Verify and Atomic Write to RDMA protocol RFC5040
- Enhanced placement semantics for RDMA to persistent media
 - E.g. Persistent memory, and other nonvolatile RDMA-addressable devices
 - Also provides remote content verification and a transaction operation
- First published February 2016
 - Updated March 2020, January 2023

<https://datatracker.ietf.org/doc/draft-talpey-rdma-commit/>

Consider Working Group adoption

- In 2020, rough WG consensus was positive for adoption in NFSv4
 - Discussion of appropriate WG:
 - STORM/RDDP previously closed
 - TSVWG was (perhaps) too broad, and had few RDMA experts
 - NFSv4 was a likely consumer, and RDMA expertise present
 - Comments received from NFSv4 discussion
 - No formal action was taken
- In late 2022/early 2023, draft was refreshed
 - Incorporating previous comments, other updates
 - Ready for consideration

If adopted, there are things to consider

- ~Half the document is informative
 - Is this content still useful? Remove it? Split into two documents?
- Retitle?
 - “Commit” is largely obsolete term
 - Now “flush” commonly used, but is only part of the picture
 - “RDMA Placement Extensions” proposed
- Prototypes have been created, but...
 - They are not publicly shared/sharable
 - Not appropriate to include
- Infiniband Trade Association (IB/RoCE) version 1.6 has since standardized the same extensions
 - IBTA spec is closed and members-only
 - Protocol and Verb (application) semantics are identical to this draft
 - Experimental Linux prototype exists, extends software RoCE “rxe”
 - Are these relevant to IETF work, or to this document?
- Software iWARP (Linux “siw”) implementation
 - ...is certainly possible
 - But hasn’t been written yet!

Discussion