

NFSv4 Working Group
Internet-Draft
Intended status: Standards Track
Expires: October 17, 2008

Tom Talpey
NetApp
Brent Callaghan
Apple
April 16, 2008

NFS Direct Data Placement
draft-ietf-nfsv4-nfsdirect-08

Status of this Memo

By submitting this Internet-Draft, each author represents that any applicable patent or other IPR claims of which he or she is aware have been or will be disclosed, and any of which he or she becomes aware will be disclosed, in accordance with [Section 6 of BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

The list of current Internet-Drafts can be accessed at
<http://www.ietf.org/ietf/lid-abstracts.txt>

The list of Internet-Draft Shadow Directories can be accessed at
<http://www.ietf.org/shadow.html>.

Abstract

This draft defines the bindings of the various Network File System (NFS) versions to the Remote Direct Memory Access (RDMA) operations supported by the RPC/RDMA transport protocol. It describes the use of direct data placement by means of server-initiated RDMA operations into client-supplied buffers for implementations of NFS versions 2, 3, 4 and 4.1 over such an RDMA transport.

Table of Contents

1.	Introduction	2
2.	Transfers from NFS Client to NFS Server	2
3.	Transfers from NFS Server to NFS Client	3
4.	NFS Versions 2 and 3 Mapping	4
5.	NFS Version 4 Mapping	5
5.1.	NFS Version 4 Callbacks	7
6.	Port Usage Considerations	8
7.	Security Considerations	8
8.	IANA Considerations	8
9.	Acknowledgments	9
10.	Normative References	9
11.	Informative References	10
12.	Authors' Addresses	10
13.	Intellectual Property and Copyright Statements	11
	Acknowledgment	11

Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [[RFC2119](#)].

[1.](#) Introduction

The Remote Direct Memory Access (RDMA) Transport for Remote Procedure Calls (RPC) [[RPCRDMA](#)] allows an RPC client application to post buffers in a Chunk list for specific arguments and results from an RPC call. The RDMA transport header conveys this list of client buffer addresses to the server where the application can associate them with client data and use RDMA operations to transfer the results directly to and from the posted buffers on the client. The client and server must agree on a consistent mapping of posted buffers to RPC. This document details the mapping for each version of the NFS protocol [[RFC1094](#)] [[RFC1813](#)] [[RFC3530](#)] [[NFSv4.1](#)].

[2.](#) Transfers from NFS Client to NFS Server

The RDMA Read list, in the RDMA transport header, allows an RPC client to marshal RPC call data selectively. Large chunks of data, such as the file data of an NFS WRITE request, MAY be referenced by

an RDMA Read list and be moved efficiently and directly-placed by an RDMA Read operation initiated by the server.

The process of identifying these chunks for the RDMA Read list can be implemented entirely within the RPC layer. It is transparent to the

upper-level protocol, such as NFS. For instance, the file data portion of an NFS WRITE request can be selected as an RDMA "chunk" within the eXternal Data Representation (XDR) marshaling code of RPC based on a size criterion, independently of the NFS protocol layer. The XDR unmarshaling on the receiving system can identify the correspondence between Read chunks and protocol elements via the XDR position value encoded in the Read chunk entry.

RPC RDMA Read chunks are employed by this NFS mapping to convey specific NFS data to the server in a manner which may be directly placed. The following sections describe this mapping for versions of the NFS protocol.

[3.](#) Transfers from NFS Server to NFS Client

The RDMA Write list, in the RDMA transport header, allows the client to post one or more buffers into which the server will RDMA Write designated result chunks directly. If the client sends a null Write list, then results from the RPC call will be returned as either an inline reply, as chunks in an RDMA Read list of server-posted buffers, or in a client-posted reply buffer.

Each posted buffer in a Write list is represented as an array of memory segments. This allows the client some flexibility in submitting discontinuous memory segments into which the server will scatter the result. Each segment is described by a triplet consisting of the segment handle or steering tag (STag), segment length, and memory address or offset.

```
struct xdr_rdma_segment {
    uint32 handle;    /* Registered memory handle */
    uint32 length;    /* Length of the chunk in bytes */
    uint64 offset;    /* Chunk virtual address or offset */
};
```

```

struct xdr_write_chunk {
    struct xdr_rdma_segment target<>;
};

struct xdr_write_list {
    struct xdr_write_chunk entry;
    struct xdr_write_list *next;
};

```

The sum of the segment lengths yields the total size of the buffer, which **MUST** be large enough to accept the result. If the buffer is too small, the server **MUST** return an XDR encode error. The server

MUST return the result data for a posted buffer by progressively filling its segments, perhaps leaving some trailing segments unfilled or partially full if the size of the result is less than the total size of the buffer segments.

The server returns the RDMA Write list to the client with the segment length fields overwritten to indicate the amount of data RDMA Written to each segment. Results returned by direct placement **MUST NOT** be returned by other methods, e.g., by Read chunk list or inline. If no result data at all is returned for the element, the server places no data in the buffer(s), but does return zeroes in the segment length fields corresponding to the result.

The RDMA Write list allows the client to provide multiple result buffers - each buffer maps to a specific result in the reply. The NFS client and server implementations agree by specifying the mapping of results to buffers for each RPC procedure. The following sections describe this mapping for versions of the NFS protocol.

Through the use of RDMA Write lists in NFS requests, it is not necessary to employ the RDMA Read lists in the NFS replies, as described in the RPC/RDMA protocol. This enables more efficient operation, by avoiding the need for the server to expose buffers for RDMA, and also avoiding "RDMA_DONE" exchanges. Clients **MAY** additionally employ RDMA Reply chunks to receive entire messages, as described in [[RPCRDMA](#)].

[4.](#) NFS Versions 2 and 3 Mapping

A single RDMA Write list entry MAY be posted by the client to receive either the opaque file data from a READ request or the pathname from a READLINK request. The server MUST ignore a Write list for any other NFS procedure, as well as any Write list entries beyond the first in the list.

Similarly, a single RDMA Read list entry MAY be posted by the client to supply the opaque file data for a WRITE request or the pathname for a SYMLINK request. The server MUST ignore any Read list for other NFS procedures, as well as additional Read list entries beyond the first in the list.

Because there are no NFS version 2 or 3 requests that transfer bulk data in both directions, it is not necessary to post requests containing both Write and Read lists. Any unneeded Read or Write lists are ignored by the server.

In the case where the outgoing request or expected incoming reply is

larger than the maximum size supported on the connection, it is possible for the RPC layer to post the entire message or result in a special "RDMA_NOMSG" message type which is transferred entirely by RDMA. This is implemented in RPC, below NFS and therefore has no effect on the message contents.

Non-RDMA (inline) WRITE transfers MAY OPTIONALLY employ the "RDMA_MSGP" padding method described in the RPC/RDMA protocol, if the appropriate value for the server is known to the client. Padding allows the opaque file data to arrive at the server in an aligned fashion, which may improve server performance.

The NFS version 2 and 3 protocols are frequently limited in practice to requests containing less than or equal to 8 kilobytes and 32 kilobytes of data, respectively. In these cases, it is often practical to support basic operation without employing a configuration exchange as discussed in [\[RPCRDMA\]](#). The server MUST post buffers large enough to receive the largest possible incoming message (approximately 12KB for NFS version 2, or 36KB for NFS version 3, would be vastly sufficient), and the client can post buffers large enough to receive replies based on the "rsize" it is using to the server, plus a fixed overhead for the RPC and NFS

headers. Because the server MUST NOT return data in excess of this size, the client can be assured of the adequacy of its posted buffer sizes.

Flow control is handled dynamically by the RPC RDMA protocol, and write padding is OPTIONAL and therefore MAY remain unused.

Alternatively, if the server is administratively configured to values appropriate for all its clients, the same assurance of interoperability within the domain can be made.

The use of a configuration protocol with NFS v2 and v3 is therefore OPTIONAL. Employing a configuration exchange may allow some advantage to server resource management through accurately sizing buffers, enabling the server to know exactly how many RDMA Reads may be in progress at once on the client connection, and enabling client write padding which may be desirable for certain servers when RDMA Read is impractical.

[5.](#) NFS Version 4 Mapping

This specification applies to the first minor version of NFS version 4 (NFSv4.0) and any subsequent minor versions that do not override this mapping.

The Write list MUST be considered only for the COMPOUND procedure. This procedure returns results from a sequence of operations. Only the opaque file data from an NFS READ operation, and the pathname from a READLINK operation MUST utilize entries from the Write list.

If there is no Write list, i.e., the list is null, then any READ or READLINK operations in the COMPOUND MUST return their data inline. The NFSv4.0 client MUST ensure in this case that any result of its READ and READLINK requests will fit within its receive buffers, in order to avoid a resulting RDMA transport error upon transfer. The server is not required to detect this.

The first entry in the Write list MUST be used by the first READ or READLINK in the COMPOUND request. The next Write list entry by the by the next READ or READLINK, and so on. If there are more READ or

READLINK operations than Write list entries, then any remaining operations MUST return their results inline.

If a Write list entry is presented, then the corresponding READ or READLINK MUST return its data via an RDMA Write to the buffer indicated by the Write list entry. If the Write list entry has zero RDMA segments, or if the total size of the segments is zero, then the corresponding READ or READLINK operation MUST return its result inline.

The following example shows an RDMA Write list with three posted buffers A, B, and C. The designated operations in the compound request, READ and READLINK, consume the posted buffers by writing their results back to each buffer.

RDMA Write list:

A --> B --> C

Compound request:

PUTFH	LOOKUP	READ	PUTFH	LOOKUP	READLINK	PUTFH	LOOKUP	READ
		v			v			v
		A			B			C

If the client does not want to have the READLINK result returned directly, then it provides a zero length array of segment triplets for buffer B or sets the values in the segment triplet for buffer B to zeros so that the READLINK result MUST be returned inline.

The situation is similar for RDMA Read lists sent by the client and applies to the NFSv4.0 WRITE and SYMLINK procedures as for v3. Additionally, inline segments too large to fit in posted buffers MAY be transferred in special "RDMA_NOMSG" messages.

Non-RDMA (inline) WRITE transfers MAY OPTIONALLY employ the "RDMA_MSGP" padding method described in the RPC/RDMA protocol, if the appropriate value for the server is known to the client. Padding

allows the opaque file data to arrive at the server in an aligned fashion, which may improve server performance. In order to ensure accurate alignment for all data, it is likely that the client will restrict its use of OPTIONAL padding to COMPOUND requests containing only a single WRITE operation.

Unlike NFS versions 2 and 3, the maximum size of an NFS version 4 COMPOUND is not bounded, even when RDMA chunks are in use. While it might appear that a configuration protocol exchange (such as the one described in [[RPCRDMA](#)]) would help, in fact the layering issues involved in building COMPOUNDS by NFS make such a mechanism unworkable.

However, typical NFS version 4 clients rarely issue such problematic requests. In practice, they behave in much more predictable ways, in fact most still support the traditional rsize/wsize mount parameters. Therefore, most NFS version 4 clients function over RPC/RDMA in the same way as NFS versions 2 and 3, operationally.

There are however advantages to allowing both client and server to operate with prearranged size constraints, for example use of the sizes to better manage the server's response cache. An extension to NFS version 4 supporting a more comprehensive exchange of upper layer parameters is part of [[NFSv4.1](#)].

[5.1](#). NFS Version 4 Callbacks

The NFS version 4 protocols support server-initiated callbacks to selected clients, in order to notify them of events such as recalled delegations, etc. These callbacks present no particular issue to being framed over RPC/RDMA, since such callbacks do not carry bulk data such as NFS READ or NFS WRITE. They MAY be transmitted inline via RDMA_MSG, or if the callback message or its reply overflow the negotiated buffer sizes for a callback connection, they MAY be transferred via the RDMA_NOMSG method as described above for other exchanges.

One special case is noteworthy: in NFS version 4.1, the callback channel is optionally negotiated to be on the same connection as one used for client requests. In this case, and because the XID is

present in the RPC/RDMA header, the client MUST ascertain whether the

message is in fact an RPC REPLY, and therefore a reply to a prior request and carrying its XID, before processing it as such. By the same token, the server MUST ascertain whether an incoming message on such a callback-eligible connection is an RPC CALL, before optionally processing the XID.

In the callback case, the XID present in the RPC/RDMA header will potentially have any value which may (or may not) collide with an XID used by the client for a previous or future request. The client and server MUST inspect the RPC component of the message to determine its potential disposition as either an RPC CALL or RPC REPLY, prior to processing this XID, and MUST NOT reject or accept it without also determining the proper context.

6. Port Usage Considerations

NFS use of direct data placement introduces a need for an additional NFS port number assignment for networks which share traditional UDP and TCP port spaces with RDMA services. The iWARP [[RFC5041](#)] [[RFC5040](#)] protocol is such an example (Infiniband is not).

NFS servers for versions 2 and 3 [[RFC1094](#)] [[RFC1813](#)] traditionally listen for clients on UDP and TCP port 2049, and additionally, they register these with the portmapper and/or rpcbind [[RFC1833](#)] service. However, [[RFC3530](#)] requires NFS servers for version 4 to listen on TCP port 2049, and they are not required to register.

An NFS version 2 or version 3 server supporting RPC/RDMA on such a network and registering itself with the RPC portmapper MAY choose an arbitrary port, or MAY use the alternative well-known port number for its RPC/RDMA service. The chosen port MAY be registered with the RPC portmapper under the netid assigned by the requirement in [[RPCRDMA](#)].

An NFS version 4 server supporting RPC/RDMA on such a network MUST use the alternative well-known port number for its RPC/RDMA service. Clients SHOULD connect to this well-known port without consulting the RPC portmapper (as for NFSv4/TCP).

The port number assigned to an NFS service over an RPC/RDMA transport is available from the IANA port registry [[RFC3232](#)].

7. Security Considerations

The RDMA transport for RPC [[RPCRDMA](#)] supports all RPC [[RFC1831bis](#)] security models, including RPCSEC_GSS [[RFC2203](#)] security and link-level security. The choice of RDMA Read and RDMA Write to return RPC argument and results, respectively, does not affect this, since it

only changes the method of data transfer. Specifically, the requirements of [RPCRDMA] ensure that this choice does not introduce new vulnerabilities.

Because this document defines only the binding of the NFS protocols atop [RPCRDMA], all relevant security considerations are therefore to be described at that layer.

8. IANA Considerations

This document has no IANA considerations.

9. Acknowledgments

The authors would like to thank Dave Noveck and Chet Juszczak for their contributions to this document.

10. Normative References

[RFC2119]

S. Bradner, "Key words for use in RFCs to Indicate Requirement Levels",
Best Current Practice,
[BCP 14](#), [RFC 2119](#), March 1997.

[RFC1094]

"NFS: Network File System Protocol Specification",
(NFS version 2) Informational RFC,
<http://www.ietf.org/rfc/rfc1094.txt>

[RFC1831bis]

R. Thurlow, Ed., "RPC: Remote Procedure Call Protocol
Specification Version 2",
Standards Track RFC

[RFC1813]

B. Callaghan, B. Pawlowski, P. Staubach, "NFS Version 3 Protocol
Specification",
Informational RFC,
<http://www.ietf.org/rfc/rfc1813.txt>

[RFC1833]

R. Srinivasan, "Binding Protocols for ONC RPC Version 2",
Standards Track RFC,
<http://www.ietf.org/rfc/rfc1833.txt>

[RFC3530]

S. Shepler, et al., "NFS version 4 Protocol",

Expires: October 2008

Talpey and Callaghan

[Page 9]

Internet-Draft

NFS Direct Data Placement

April 2008

Standards Track RFC,

<http://www.ietf.org/rfc/rfc3530.txt>

[NFSv4.1]

S. Shepler et al., ed., "NFSv4 Minor Version 1"

Internet Draft Work in Progress,

[draft-ietf-nfsv4-minorversion1](#)

[RFC2203]

M. Eisler, A. Chiu, L. Ling, "RPCSEC_GSS Protocol Specification",
Standards Track RFC,

<http://www.ietf.org/rfc/rfc2203.txt>

11. Informative References

[RFC3232]

Internet Assigned Numbers Authority (IANA),
Port Registry database,

<http://www.ietf.org/rfc/rfc3232.txt>

<http://www.iana.org/assignments/port-numbers>

[RPCRDMA]

T. Talpey, B. Callaghan, "Remote Direct Memory Access Transport
for Remote Procedure Call"

Internet Draft Work in Progress,

[draft-ietf-nfsv4-rpcrdma](#)

[RFC5041]

H. Shah et al., "Direct Data Placement over Reliable Transports",
Standards Track RFC

[RFC5040]

R. Recio et al., "A Remote Direct Memory Access Protocol
Specification",

Standards Track RFC

12. Authors' Addresses

Tom Talpey
Network Appliance, Inc.
1601 Trapelo Road, #16
Waltham, MA 02451 USA

Phone: +1 781 768 5329
EMail: thomas.talpey@netapp.com

Expires: October 2008

Talpey and Callaghan

[Page 10]

Internet-Draft

NFS Direct Data Placement

April 2008

Brent Callaghan
Apple Computer, Inc.
MS: 302-4K
2 Infinite Loop
Cupertino, CA 95014 USA

EMail: brentc@apple.com

13. Intellectual Property and Copyright Statements

Full Copyright Statement

Copyright (C) The IETF Trust (2008).

This document is subject to the rights, licenses and restrictions contained in [BCP 78](#), and except as set forth therein, the authors retain all their rights.

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY, THE IETF TRUST AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Intellectual Property

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights or other rights that might be claimed

to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in [BCP 78](#) and [BCP 79](#).

Copies of IPR disclosures made to the IETF Secretariat and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the IETF on-line IPR repository at <http://www.ietf.org/ipr>.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary

Expires: October 2008

Talpey and Callaghan

[Page 11]

Internet-Draft

NFS Direct Data Placement

April 2008

rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at ietf-ipr@ietf.org.

Acknowledgment

Funding for the RFC Editor function is provided by the IETF Administrative Support Activity (IASA).

