

Storage Maintenance (storm) Working Group
Internet Draft
Intended status: Standards Track
Expires: January 2013

Hemal Shah
Broadcom Corporation
Felix Marti
Wael Nouredine
Asgeir Eiriksson
Chelsio Communications, Inc.
Robert Sharp
Intel Corporation
July 9, 2012

RDMA Protocol Extensions
draft-ietf-storm-rdmap-ext-03.txt

Status of this Memo

This Internet-Draft is submitted to IETF in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current>.

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."

This Internet-Draft will expire on January 9, 2013.

Copyright Notice

Copyright (c) 2012 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in

Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Abstract

This document specifies extensions to the IETF Remote Direct Memory Access Protocol (RDMA [RFC5040]). RDMA provides read and write services directly to applications and enables data to be transferred directly into Upper Layer Protocol (ULP) Buffers without intermediate data copies. The extensions specified in this document provide the following capabilities and/or improvements: Atomic Operations and Immediate Data.

Table of Contents

1. Introduction.....	3
2. Requirements Language.....	3
3. Glossary.....	3
4. Header Format Extensions.....	5
4.1. RDMA Control and Invalidate STag Fields.....	5
4.2. RDMA Message Definitions.....	6
5. Atomic Operations.....	7
5.1. Atomic Operation Details.....	8
5.1.1. FetchAdd.....	8
5.1.2. Swap.....	9
5.1.3. CmpSwap.....	10
5.2. Atomic Operations.....	11
5.2.1. Atomic Operation Request Message.....	12
5.2.2. Atomic Operation Response Message.....	15
5.3. Atomicity Guarantees.....	17
5.4. Atomic Operations Ordering and Completion Rules.....	17
6. Immediate Data.....	18
6.1. RDMA Interactions with ULP for Immediate Data.....	18
6.2. Immediate Data Header Format.....	18
6.3. Immediate Data or Immediate Data with SE Message.....	19
6.4. Ordering and Completions.....	20
7. Ordering and Completions Table.....	20
8. Error Processing.....	23
8.1. Errors Detected at the Local Peer.....	23
8.2. Errors Detected at the Remote Peer.....	24

9. Security Considerations.....	24
10. IANA Considerations.....	24
10.1. Atomic Operation Codes.....	25
11. References.....	26
11.1. Normative References.....	26
11.2. Informative References.....	26
12. Acknowledgments.....	26
Appendix A. DDP Segment Formats for RDMA Messages.....	27
A.1. DDP Segment for Atomic Operation Request.....	27
A.2. DDP Segment for Atomic Response.....	29
A.3. DDP Segment for Immediate Data and Immediate Data with SE29	

1. Introduction

The RDMA Protocol [RFC5040] provides capabilities for zero copy and kernel bypass data communications. This document specifies the following extensions to the RDMA Protocol (RDMA):

- o Atomic operations on remote memory locations. Support for atomic operation enhances the usability of RDMA in distributed shared memory environments.
- o Immediate Data messages allow the ULP at the sender to provide a small amount of data following an RDMA Message.

Other RDMA transport protocols define the functionality added by these extensions leading to differences in RDMA applications and/or Upper Layer Protocols. Removing these differences in the transport protocols simplifies these applications and ULPs and that is the main motivation for the extensions specified in this document.

2. 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 RFC-2119 [RFC2119].

3. Glossary

This document is an extension of [RFC5040] and key words are defined in the glossary of the referenced document.

Atomic Operation - is an operation that results in an execution of a 64-bit operation at a specific address on a remote node. The

consumer can use Atomic Operations to read, modify and write at the destination address while at the same time guarantee that no other read or write operation will occur across any other RDMAP Streams on an RNIC at the Data Sink.

Atomic Operation Request - An RDMA Message used by the Data Source to perform an Atomic Operation at the Data Sink.

Atomic Operation Response - An RDMA Message used by the Data Sink to describe the completion of an Atomic Operation at the Data Sink.

CmpSwap - is an Atomic Operation that is used to compare and swap a value at a specific address on a remote node.

FetchAdd - is an Atomic Operation that is used to atomically increment a value at a specific address on a remote node.

Immediate Data - a small fixed size portion of data sent from the Data Source to a Data Sink

Immediate Data Message - An RDMA Message used by the Data Source to send Immediate Data to the Data Sink

Immediate Data with Solicited Event (SE) Message - An RDMA Message used by the Data Source to send Immediate Data with Solicited Event to the Data Sink

Requester - the sender of an RDMA Atomic Operation request.

Responder - the receiver of an RDMA Atomic Operation request.

Swap - is an Atomic Operation that is used to swap a value at a specific address on a remote node.

4. Header Format Extensions

The control information of RDMA Messages is included in DDP protocol [RFC5041] defined header fields, with the following new formats:

- Four new RDMA Messages carry additional RDMAP headers. The Immediate Data operation and Immediate Data with Solicited Event operation include 8 bytes of data following the RDMAP header. Atomic Operations include Atomic Request or Atomic Response headers following the RDMAP header.

4.1. RDMAP Control and Invalidate STag Fields

The RDMA Messages defined by this specification use all 8 bits of the RDMAP Control Field. The first octet reserved for ULP use in the DDP Protocol MUST be used by the RDMAP to carry the RDMAP Control Field. The ordering of the bits in the first octet MUST be as shown in Figure 1.

Figure 1 depicts the format of the DDP Control and RDMAP Control fields, in the style and convention of [RFC5040]:




Figure 1 DDP Control and RDMAP Control Fields

Figure 2 defines the values of RDMA Opcode field that MUST be used for the RDMA Messages defined in this specification.

Figure 2 also defines when the STag, Tagged Offset, and Queue Number fields MUST be provided for the RDMA Messages defined in this specification.

All RDMA Messages defined in this specification MUST have:

The RDMA Version (RV) field: 01b.

Opcode field: See Figure 2.

Invalidate STag: MUST be set to zero by the sender, ignored by the receiver.

RDMA Opcode	Message Type	Tagged Flag	STag and TO	Queue Number	Invalidate STag	Message Length Communicated between DDP and RDMAP
1000b	Immediate Data	0	N/A	0	N/A	Yes
1001b	Immediate Data with SE	0	N/A	0	N/A	Yes
1010b	Atomic Request	0	N/A	1	N/A	Yes
1011b	Atomic Response	0	N/A	3	N/A	Yes

Figure 2 Additional RDMA Usage of DDP Fields

Note: N/A means Not Applicable.

All other DDP and RDMAP control fields MUST be set as described in [RFC5040].

4.2. RDMA Message Definitions

The following figure defines which RDMA Headers MUST be used on each new RDMA Message and which new RDMA Messages are allowed to carry ULP payload:

RDMA Message OpCode	Message Type	RDMA Header Used	ULP Message allowed in the RDMA Message
1000b	Immediate Data	Immediate Data Header	No
1001b	Immediate Data with SE	Immediate Data Header	No
1010b	Atomic Request	Atomic Request Header	No
1011b	Atomic Response	Atomic Response Header	No

Figure 3 RDMA Message Definitions

5. Atomic Operations

The RDMA Protocol Specification in [RFC5040] does not include support for Atomic Operations which are an important building block for implementing distributed shared memory.

This document extends the RDMA Protocol specification with a set of basic Atomic Operations, and specifies their resource and ordering rules.

Atomic operations as specified in this document execute a 64-bit operation at a specified destination address on a remote node. The operations atomically read, modify and write back the contents of the destination address and guarantee that Atomic Operations on this address by other RDMAP Streams on the same RNIC do not occur between the read and the write. Atomic Operations as specified in this document MAY be implemented. The discovery of whether the Atomic Operations are implemented or not is outside the scope of this specification and it should be handled by the ULPs or applications. The advertisement of Tagged Buffer information for Atomic Operations is outside the scope of this specification and it must be handled by the ULPs.

Implementation note: It is recommended that the applications do not use the buffer addresses used for Atomic Operations for other RDMA operations.

Atomic Operations use the same remote addressing mechanism as RDMA Reads and Writes. The buffer address specified in the request is in the address space of the Remote Peer that the Atomic Operation is targeted at.

5.1. Atomic Operation Details

The following sub-sections describe the Atomic Operations in more details.

5.1.1. FetchAdd

The FetchAdd Atomic Operation requests the Responder to read a 64-bit Original Remote Data value at a 64-bit aligned buffer address in the Responder's memory, to perform FetchAdd operation on multiple fields of selectable length specified by 64-bit "Add Mask", and write the result back to the same virtual address. The Atomic addition is performed independently on each one of these fields. A bit set in the Add Mask field specifies the field boundary. FetchAdd Atomic Operations MUST target buffer addresses that are 64-bit aligned. FetchAdd Atomic Operations that target buffer addresses that are not 64-bit aligned MUST be surfaced as errors and the Responder's memory MUST NOT be modified in such cases. Additionally an error MUST be surfaced and a terminate message MUST be generated. The setting of "Add Mask" field to 0x0000000000000000 results in Atomic Add of 64-bit Original Remote Data Value and 64-bit "Add Data".

The pseudo code below describes masked FetchAdd Atomic Operation.

```
bit_location = 1
carry = 0
Remote Data Value = 0
for bit = 0 to 63
{
    if (bit != 0 ) bit_location = bit_location << 1
```

```
    val1 = (!((Original Remote Data Value & bit_location))
    val2 = (!((Add Data & bit_location))
    sum = carry + val1 + val2
    carry = (!((sum & 2))
    sum = sum & 1
    if (sum)
        Remote Data Value |= bit_location
    carry = ((carry) && (!(Add Mask & bit_location)))
}
```

The FetchAdd operation is performed in the endian format of the target memory. The "Original Remote Data" is converted from the endian format of the target memory for return and returned to the Requester. The fields are in big-endian format on the wire.

The Requester specifies:

- o Remote STag
- o Remote Tagged Offset
- o Add Data
- o Add Mask

The Responder returns:

- o Original Remote Data

5.1.2. Swap

The Swap Atomic Operation requires the Responder to read a 64-bit value at a 64-bit aligned buffer address in the Responder's memory, then to write the "Swap Data" fields into the same buffer address. The "Original Remote Data" is converted from the endian format of the target memory for return and returned to the Requester. The fields are in big-endian format on the wire.

The Requester specifies:

- o Remote STag
- o Remote Tagged Offset
- o Swap Data

The Responder returns:

- o Original Remote Data

After the successful completion of Swap operation, the Responder's memory at the specified buffer address MUST contain the "Swap Data" field in the header. Swap Atomic Operations MUST target buffer addresses that are 64-bit aligned. Swap Atomic Operations that target buffer addresses that are not 64-bit aligned MUST be surfaced as errors and the Responder's memory MUST NOT be modified in such cases. Additionally an error MUST be surfaced and a terminate message MUST be generated.

5.1.3. CmpSwap

The CmpSwap Atomic Operation requires the Responder to read a 64-bit value at a 64-bit aligned buffer address in the Responder's memory, to perform an AND logical operation using the 64 bit "Compare Mask" field in the Atomic Operation Request header, then to compare it with the result of a logical AND operation of the "Compare Mask" and the "Compare Data" fields in the header, and, if the two values are equal, to swap masked bits in the same buffer address with the masked Swap Data. If the two masked compare values are not equal, the contents of the Responder's memory are not changed. In either case, the original value read from the buffer address is converted from the endian format of the target memory for return and returned to the Requester. The fields are in big-endian format on the wire.

The Requester specifies:

- o Remote STag
- o Remote Tagged Offset
- o Swap Data
- o Swap Mask

- o Compare Data
- o Compare Mask

The Responder returns:

- o Original Remote Data Value

The following pseudo code describes the masked CmpSwap operation result.

```
if (!((Compare Data ^ Original Remote Data value) & Compare Mask)
then
    Remote Data Value =
        (Original Remote Data Value & ~(Swap Mask))
        | (Swap Data & Swap Mask)
else
    Remote Data Value = Original Remote Data Value
```

After the operation, the remote data buffer MUST contain the "Original Remote Data Value" (if comparison did not match) or the masked "Swap Data" (if the comparison did match). CmpSwap Atomic Operations MUST target buffer addresses that are 64-bit aligned. CmpSwap Atomic Operations that target buffer addresses that are not 64-bit aligned MUST be surfaced as errors and the remote data buffer MUST NOT be modified in such cases. Additionally an error MUST be surfaced and a terminate message MUST be generated.

5.2. Atomic Operations

The Atomic Operation Request and Response are RDMA Messages. An Atomic Operation makes use of the DDP Untagged Buffer Model. Atomic Operations use the same Queue Number as RDMA Read Requests (QN=1). Reusing the same Queue Number allows the Atomic Operations to reuse the same infrastructure (e.g. ORD/IRD flow control) as defined for RDMA Read Requests.

The RDMA Message OpCode for an Atomic Request Message is 1010b. The RDMA Message OpCode for an Atomic Response Message is 1011b.

5.2.1. Atomic Operation Request Message

The Atomic Operation Request Message carries an Atomic Operation Header that describes the buffer address in the Responder's memory. The Atomic Operation Request header immediately follows the DDP header. The RDMAP layer passes to the DDP layer a RDMAP Control Field. The following figure depicts the Atomic Operation Request Header that MUST be used for all Atomic Operation Request Messages:




Figure 4 Atomic Operation Request Header

Reserved (Not Used): 28 bits

This field MUST be set to zero on transmit, ignored on receive.

Atomic Operation Code (AOpCode): 4 bits.

See Figure 5.

Request Identifier: 32 bits.

The Request Identifier specifies a number that is used to identify Atomic Operation Request Message. The use of this field is implementation dependent and outside the scope of this specification.

Remote STag: 32 bits.

The Remote STag identifies the Remote Peer's Tagged Buffer targeted by the Atomic Operation. The Remote STag is associated with the RDMAP Stream through a mechanism that is outside the scope of the RDMAP specification.

Remote Tagged Offset: 64 bits.

The Remote Tagged Offset specifies the starting offset, in octets, from the base of the Remote Peer's Tagged Buffer targeted by the Atomic Operation. The Remote Tagged Offset MAY start at an arbitrary offset.

Add or Swap Data: 64 bits.

The Add or Swap Data field specifies the 64-bit "Add Data" value in an Atomic FetchAdd Operation or the 64-bit "Swap Data" value in an Atomic Swap or CmpSwap Operation.

Add or Swap Mask: 64 bits

This field is used in masked Atomic Operations (FetchAdd and CmpSwap) to perform a bitwise logical AND operation as specified in the definition of these operations. For non-masked Atomic Operations (Swap), this field MUST be set to ffffffffffffffffh on transmit and ignored by the receiver.

Compare Data: 64 bits.

The Compare Data field specifies the 64-bit "Compare Data" value in an Atomic CmpSwap Operation. For Atomic FetchAdd and Atomic Swap operation, the Compare Data field MUST be set to zero on transmit and ignored by the receiver.

Compare Mask: 64 bits

This field is used in masked Atomic Operation CmpSwap to perform a bitwise logical AND operation as specified in the definition of these operations. For Atomic Operations FetchAdd and Swap, this field MUST be set to ffffffffffffffffh on transmit and ignored by the receiver.

Atomic Operation Code	Atomic Operation	Add or Swap Data	Add or Swap Mask	Compare Data	Compare Mask
0000b	FetchAdd	Add Data	Add Mask	N/A	N/A
0001b	Swap	Swap Data	N/A	N/A	N/A
0010b	CmpSwap	Swap Data	Swap Mask	Valid	Valid
0011b to 1111b	Reserved	Not Specified			

Figure 5 Atomic Operation Message Definitions

The Atomic Operation Request Message has the following semantics:

1. An Atomic Operation Request Message MUST reference an Untagged Buffer. That is, the Local Peer's RDMAP layer MUST request that the DDP mark the Message as Untagged.
2. One Atomic Operation Request Message MUST consume one Untagged Buffer.

- 3. The Responder's RDMAP layer MUST process an Atomic Operation Request Message. A valid Atomic Operation Request Message MUST NOT be delivered to the Responder's ULP (i.e., it is processed by the RDMAP layer).
- 4. At the Responder, when an invalid Atomic Operation Request Message is delivered to the Remote Peer's RDMAP layer, an error is surfaced.
- 5. An Atomic Operation Request Message MUST reference the RDMA Read Request Queue. That is, the Requester's RDMAP layer MUST request that the DDP layer set the Queue Number field to one.
- 6. The Requester MUST pass to the DDP layer Atomic Operation Request Messages in the order they were submitted by the ULP.
- 7. The Responder MUST process the Atomic Operation Request Messages in the order they were sent.
- 8. If the Responder receives a valid Atomic Operation Request Message, it MUST respond with a valid Atomic Operation Response Message.

5.2.2. Atomic Operation Response Message

The Atomic Operation Response Message carries an Atomic Operation Response Header that contains the "Original Request Identifier" and "Original Remote Data Value". The Atomic Operation Response Header immediately follows the DDP header. The RDMAP layer passes to the DDP layer a RDMAP Control Field. The following figure depicts the Atomic Operation Response header that MUST be used for all Atomic Operation Response Messages:




Figure 6 Atomic Operation Response Header

Original Request Identifier: 32 bits.

The Original Request Identifier MUST be set to the value specified in the Request Identifier field that was originally provided in the corresponding Atomic Operation Request Message.

Original Remote Data Value: 64 bits.

The Original Remote Value specifies the original 64-bit value stored at the buffer address targeted by the Atomic Operation.

The Atomic Operation Response Message has the following semantics:

1. The Atomic Operation Response Message for the associated Atomic Operation Request Message travels in the opposite direction.
2. An Atomic Operation Response Message MUST consume an Untagged Buffer. That is, the Responder RDMAP layer MUST request that the DDP mark the Message as Untagged.
3. An Atomic Operation Response Message MUST reference the Queue Number 3. That is, the Responder's RDMAP layer MUST request that the DDP layer set the Queue Number field to 3.
4. The Responder MUST ensure that a sufficient number of Untagged Buffers are available on the RDMA Read Request Queue (Queue with DDP Queue Number 1) to support the maximum number of Atomic Operation Requests negotiated by the ULP.
5. The RDMAP layer MUST Deliver the Atomic Operation Response Message to the ULP.
6. At the Requester, when an invalid Atomic Operation Response Message is delivered to the Remote Peer's RDMAP layer, an error is surfaced.
7. The Responder RDMAP layer MUST pass Atomic Operation Response Messages to the DDP layer, in the order that the Atomic Operation Request Messages were received by the RDMAP layer, at the Responder.

5.3. Atomicity Guarantees

Atomicity of the Read-Modify-Write (RMW) on the Responder's node by the Atomic Operation MUST be assured in the presence of concurrent atomic accesses by other RDMAP Streams on the same RNIC.

5.4. Atomic Operations Ordering and Completion Rules

In addition to the ordering and completion rules described in [RFC5040], the following rules apply to implementations of the Atomic operations.

1. For an Atomic operation, the contents of the Tagged Buffer at the Responder MAY be indeterminate until the Atomic Operation Response Message has been Delivered at the Requester.
2. Atomic Operation Request Messages MUST NOT start processing at the Responder until they have been Delivered to RDMAP by DDP.
3. Atomic Operation Response Messages MAY be generated at the Responder after subsequent RDMA Write Messages or Send Messages have been Placed or Delivered.
4. Atomic Operation Response Message processing at the Responder MUST be started only after the Atomic Operation Request Message has been Delivered by the DDP layer (thus, all previous RDMA Messages have been properly submitted for ordered Placement).
5. Send Messages MAY be Completed at the Responder before prior incoming Atomic Operation Request Messages have completed their response processing.
6. An Atomic Operation MUST NOT be Completed at the Requester until the DDP layer Delivers the associated incoming Atomic Operation Response Message.
7. If more than one outstanding Atomic Request Messages are supported by both peers, the Atomic Operation Request Messages MUST be processed in the order they were delivered by the DDP layer on the Responder. Atomic Operation Response Messages MUST be submitted to the DDP layer on the Responder in the order the Atomic Operation Request Messages were Delivered by DDP.

6. Immediate Data

The Immediate Data operation is used in conjunction with an RDMA Operation to improve ULP processing efficiency by allowing 8 bytes of immediate data to be delivered with the completion of the previous operation after the previous operation has been delivered at the Remote Peer.

6.1. RDMAP Interactions with ULP for Immediate Data

For Immediate Data operations, the following are the interactions between the RDMAP Layer and the ULP:

- At the Data Source:
 - The ULP passes to the RDMAP Layer the following:
 - Eight bytes of ULP Immediate Data
 - When the Immediate Data operation Completes, an indication of the Completion results.
- At the Data Sink:
 - If the Immediate Data operation is Completed successfully, the RDMAP Layer passes the following information to the ULP Layer:
 - Eight bytes of Immediate Data
 - An Event, if the Data Sink is configured to generate an Event and the RDMA Message Opcode indicates Message Type Immediate Data with Solicited Event.
 - If the Immediate Data operation is Completed in error, the Data Sink RDMAP Layer will pass up the corresponding error information to the Data Sink ULP and send a Terminate Message to the Data Source RDMAP Layer. The Data Source RDMAP Layer will then pass up the Terminate Message to the ULP.

6.2. Immediate Data Header Format

The Immediate Data and Immediate Data with SE Messages carry immediate data as shown in Figure 7. The RDMAP layer passes to the

DDP layer an RDMAP Control Field and 8 bytes of Immediate Data. The first 8 bytes of the data following the DDP header contains the Immediate Data. See section A.3. for the DDP segment format of an Immediate Data or Immediate Data with SE Message.




Figure 7 Immediate Data or Immediate Data with SE Message Header

Immediate Data: 64 bits.

Eight bytes of data transferred from the Requester to an untagged buffer at the Responder.

6.3. Immediate Data or Immediate Data with SE Message

The Immediate Data or Immediate Data with SE Message uses the DDP Untagged Buffer Model to transfer Immediate data from the Data Source to the Data Sink.

- An Immediate Data or Immediate Data with SE Message MUST reference an Untagged Buffer. That is, the Local Peer's RDMAP Layer MUST request that the DDP layer mark the Message as Untagged.
- One Immediate Data or Immediate Data with SE Message MUST consume one Untagged Buffer.
- At the Remote Peer, the Immediate Data or Immediate Data with SE Message MUST be Delivered to the Remote Peer's ULP in the order they were sent.
- For an Immediate Data or Immediate Data with SE Message, the Local Peer's RDMAP Layer MUST request that the DDP layer set the Queue Number field to zero.

- For an Immediate Data or Immediate Data with SE Message, the Local Peer's RDMAP Layer MUST request that the DDP layer transmit 8 bytes of data.
- The Local Peer MUST issue Immediate Data and Immediate Data with SE Messages in the order they were submitted by the ULP.
- The Remote Peer MUST check that Immediate Data and Immediate Data with SE Messages include exactly 8 bytes of data from the DDP layer.

6.4. Ordering and Completions

Ordering and completion rules for Immediate Data are the same as those for a Send operation as described in section 5.5 of RFC 5040.

7. Ordering and Completions Table

The following table summarizes the ordering relationships for Atomic and Immediate Data operations from the standpoint of Local Peer issuing the Operations. Note that in the table that follows, Send includes Send, Send with Invalidate, Send with Solicited Event, and Send with Solicited Event and Invalidate. Also note that in the table below, Immediate Data includes Immediate Data and Immediate Data with Solicited Event.

First Operation	Second Operation	Placement Guarantee at Remote Peer	Placement Guarantee at Local Peer	Ordering Guarantee at Remote Peer
Immediate Data	Send	No Placement Guarantee between Send Payload and Immediate Data	Not Applicable	Completed in Order
Immediate Data	RDMA Write	No Placement Guarantee between RDMA Write Payload and	Not Applicable	Not Applicable

		Immediate Data		
Immediate Data	RDMA Read	No Placement Guarantee between Immediate Data and RDMA Read Request	RDMA Read Response will not be Placed until Immediate Data is Placed at Remote Peer	RDMA Read Response Message will not be generated until Immediate Data has been Completed
Immediate Data	Atomic	No Placement Guarantee between Immediate Data and Atomic Request	Atomic Response will not be Placed until Immediate Data is Placed at Remote Peer	Atomic Response Message will not be generated until Immediate Data has been Completed
Immediate Data or Send	Immediate Data	No Placement Guarantee	Not Applicable	Completed in Order
RDMA Write	Immediate Data	No Placement Guarantee	Not Applicable	Immediate Data is Completed after RDMA Write is Placed and Delivered
RDMA Read	Immediate Data	No Placement Guarantee between Immediate Data and RDMA Read Request	Immediate Data MAY be Placed before RDMA Read Response is generated	Not Applicable
Atomic	Immediate Data	No Placement Guarantee between Immediate	Immediate Data MAY be Placed before	Not Applicable

		Data and Atomic Request	Atomic Response is generated	
Atomic	Send	No Placement Guarantee between Send Payload and Atomic Request	Send Payload MAY be Placed before Atomic Response is generated	Not Applicable
Atomic	RDMA Write	No Placement Guarantee between RDMA Write Payload and Atomic Request	RDMA Write Payload MAY be Placed before Atomic Response is generated	Not Applicable
Atomic	RDMA Read	No Placement Guarantee between Atomic Request and RDMA Read Request	No Placement Guarantee between Atomic Response and RDMA Read Response	RDMA Read Response Message will not be generated until Atomic Response Message has been generated
Atomic	Atomic	No Placement Guarantee between two Atomic Requests	No Placement Guarantee between two Atomic Responses	Second Atomic Response Message will not be generated until first Atomic Response has been generated
Send	Atomic	No Placement Guarantee between Send Payload and Atomic Request	Atomic Response will not be Placed at the Local	Atomic Response Message will not be generated until Send has been Completed

		Request	Peer Until Send Payload is Placed at the Remote Peer	
RDMA Write	Atomic	No Placement Guarantee between RDMA Write Payload and Atomic Request	Atomic Response will not be Placed at the Local Peer Until Send Payload is Placed at the Remote Peer	Not Applicable
RDMA Read	Atomic	No Placement Guarantee between Atomic Request and RDMA Read Request	No Placement Guarantee between Atomic Response and RDMA Read Response	Atomic Response Message will not be generated until RDMA Read Response has been generated

8. Error Processing

In addition to error processing described in section 7 of [RFC5040], the following rules apply for the new RDMA Messages defined in this specification.

8.1. Errors Detected at the Local Peer

The Local Peer MUST send a Terminate Message for each of the following cases:

1. For errors detected while creating an Atomic Request, Atomic Response, Immediate Data, or Immediate Data with SE Message, or other reasons not directly associated with an incoming Message, the Terminate Message and Error code are sent instead of the Message. In this case, the Error Type and Error Code fields are included in the Terminate Message, but the Terminated DDP Header and Terminated RDMA Header fields are set to zero.

2. For errors detected on an incoming Atomic Request, Atomic Response, Immediate Data, or Immediate Data with Solicited Event (after the Message has been Delivered by DDP), the Terminate Message is sent at the earliest possible opportunity, preferably in the next outgoing RDMA Message. In this case, the Error Type, Error Code, and Terminated DDP Header fields are included in the Terminate Message, but the Terminated RDMA Header field is set to zero.

8.2. Errors Detected at the Remote Peer

On incoming Atomic Requests, Atomic Responses, Immediate Data, and Immediate Data with Solicited Event, the following MUST be validated:

- The DDP layer MUST validate all DDP Segment fields.
- The RDMA OpCode MUST be valid.
- The RDMA Version MUST be valid.

On incoming Atomic requests the following additional validation MUST be performed:

- The RDMAP layer MUST validate that the Remote Peer's Tagged Buffer address references a 64-bit aligned ULP buffer address. In the case of an error, the RDMAP layer MUST generate a Terminate Message indicating RDMA Layer Remote Operation Error with Error Code Name "Catastrophic Error, Localized to RDMAP Stream" as described in Section 4.8 of [RFC5040]. Implementation Note: A ULP implementation can avoid this error by having the target ULP buffer of an atomic operation 64-bit aligned.

9. Security Considerations

This document specifies extensions to the RDMA Protocol specification in [RFC5040], and as such the Security Considerations discussed in Section 8 of [RFC5040] apply.

10. IANA Considerations

IANA is requested to add the following entries to the "RDMA Message Operation Codes" registry of "RDDP Registries":

0x8, Immediate Data, [RFCXXXX]

0x9, Immediate Data with SE, [RFCXXXX]

0xA, Atomic Request, [RFCXXXX]

0xB, Atomic Response, [RFCXXXX]

In addition, the following registry is requested to be added to "RDDP Registries". The following section specifies the registry, its initial contents and the administration policy in more detail.

10.1. Atomic Operation Codes

Name of the registry: "Atomic Operation Codes"

Namespace details: Atomic Operation Codes are 4-bit values [RFCXXXX].

Information that must be provided to assign a new value: An IESG-approved standards-track specification defining the semantics and interoperability requirements of the proposed new value and the fields to be recorded in the registry.

Assignment policy: If the requested value is not already assigned, it may be assigned to the requester.

Fields to record in the registry: Atomic Operation Code, Atomic Operation, RFC Reference.

Initial registry contents:

0x0, FetchAdd, [RFCXXXX]

0x1, Swap, [RFCXXXX]

0x2, CmpSwap, [RFCXXXX]

All other values are Unassigned and available to IANA for assignment.

Allocation Policy: Standards Action ([RFC5226])

RFC Editor: Please replace XXXX in all instances of [RFCXXXX] above with the RFC number of this document and remove this note.

11. References

11.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC5040] Recio, R. et al., "A Remote Direct Memory Access Protocol Specification", RFC 5040, October 2007.
- [RFC5041] Shah, H. et al., "Direct Data Placement over Reliable Transports", RFC 5041, October 2007.
- [RFC5226] T. Narten and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", RFC 5226, BCP 26, May 2008.

11.2. Informative References

12. Acknowledgments

The authors would like to acknowledge the following contributors who provided valuable comments and suggestions.

- o David Black
- o Arkady Kanevsky
- o Bernard Metzler
- o Jim Pinkerton
- o Tom Talpey
- o Steve Wise


This document was prepared using 2-Word-v2.0.template.dot.

Appendix A. DDP Segment Formats for RDMA Messages

This appendix is for information only and is NOT part of the standard. It simply depicts the DDP Segment format for the various RDMA Messages.

A.1. DDP Segment for Atomic Operation Request

The following figure depicts an Atomic Operation Request, DDP Segment:



Authors' Addresses

Hemal Shah
Broadcom Corporation
5300 California Avenue
Irvine, CA 92617
Phone: 1-949-926-6941
Email: hemal@broadcom.com

Felix Marti
Chelsio Communications, Inc.
370 San Aleso Ave.
Sunnyvale, CA 94085
Phone: 1-408-962-3600
Email: felix@chelsio.com

Asgeir Eiriksson
Chelsio Communications, Inc.
370 San Aleso Ave.
Sunnyvale, CA 94085
Phone: 1-408-962-3600
Email: asgeir@chelsio.com

Wael Nouredine
Chelsio Communications, Inc.
370 San Aleso Ave.
Sunnyvale, CA 94085
Phone: 1-408-962-3600
Email: wael@chelsio.com

Robert Sharp
Intel Corporation
1501 South Mopac, Suite 400, Mailstop: AN1-WTR1
Austin, TX 78746
Phone: 1-512-493-3242
Email: robert.o.sharp@intel.com