

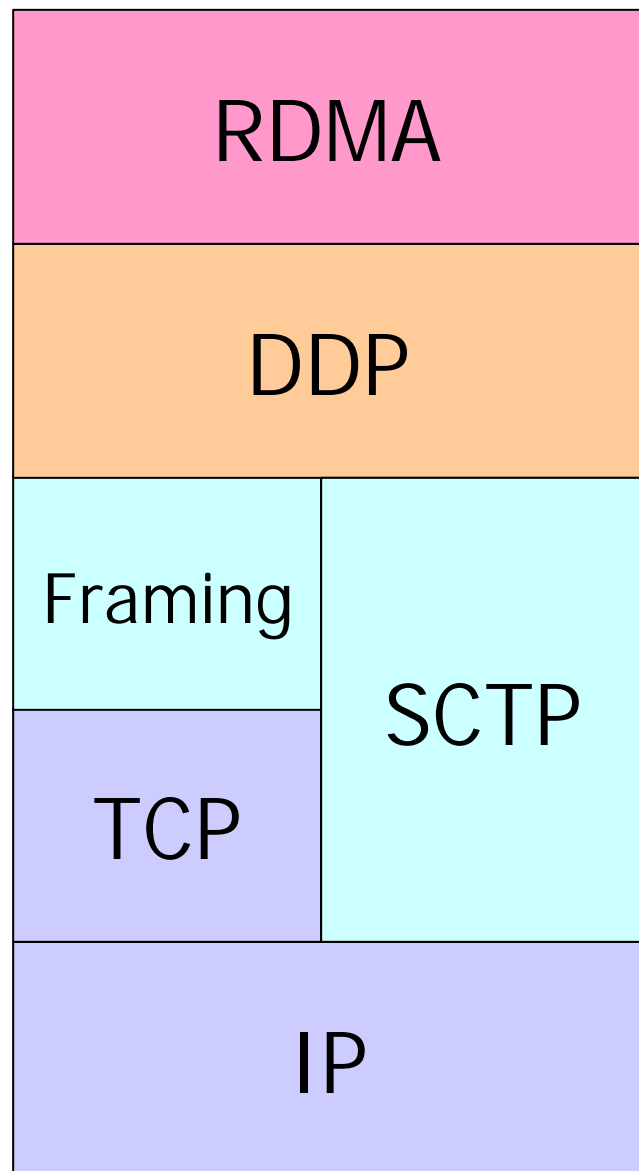
RDMAP and DDP Overview

Renato Recio

Introduction

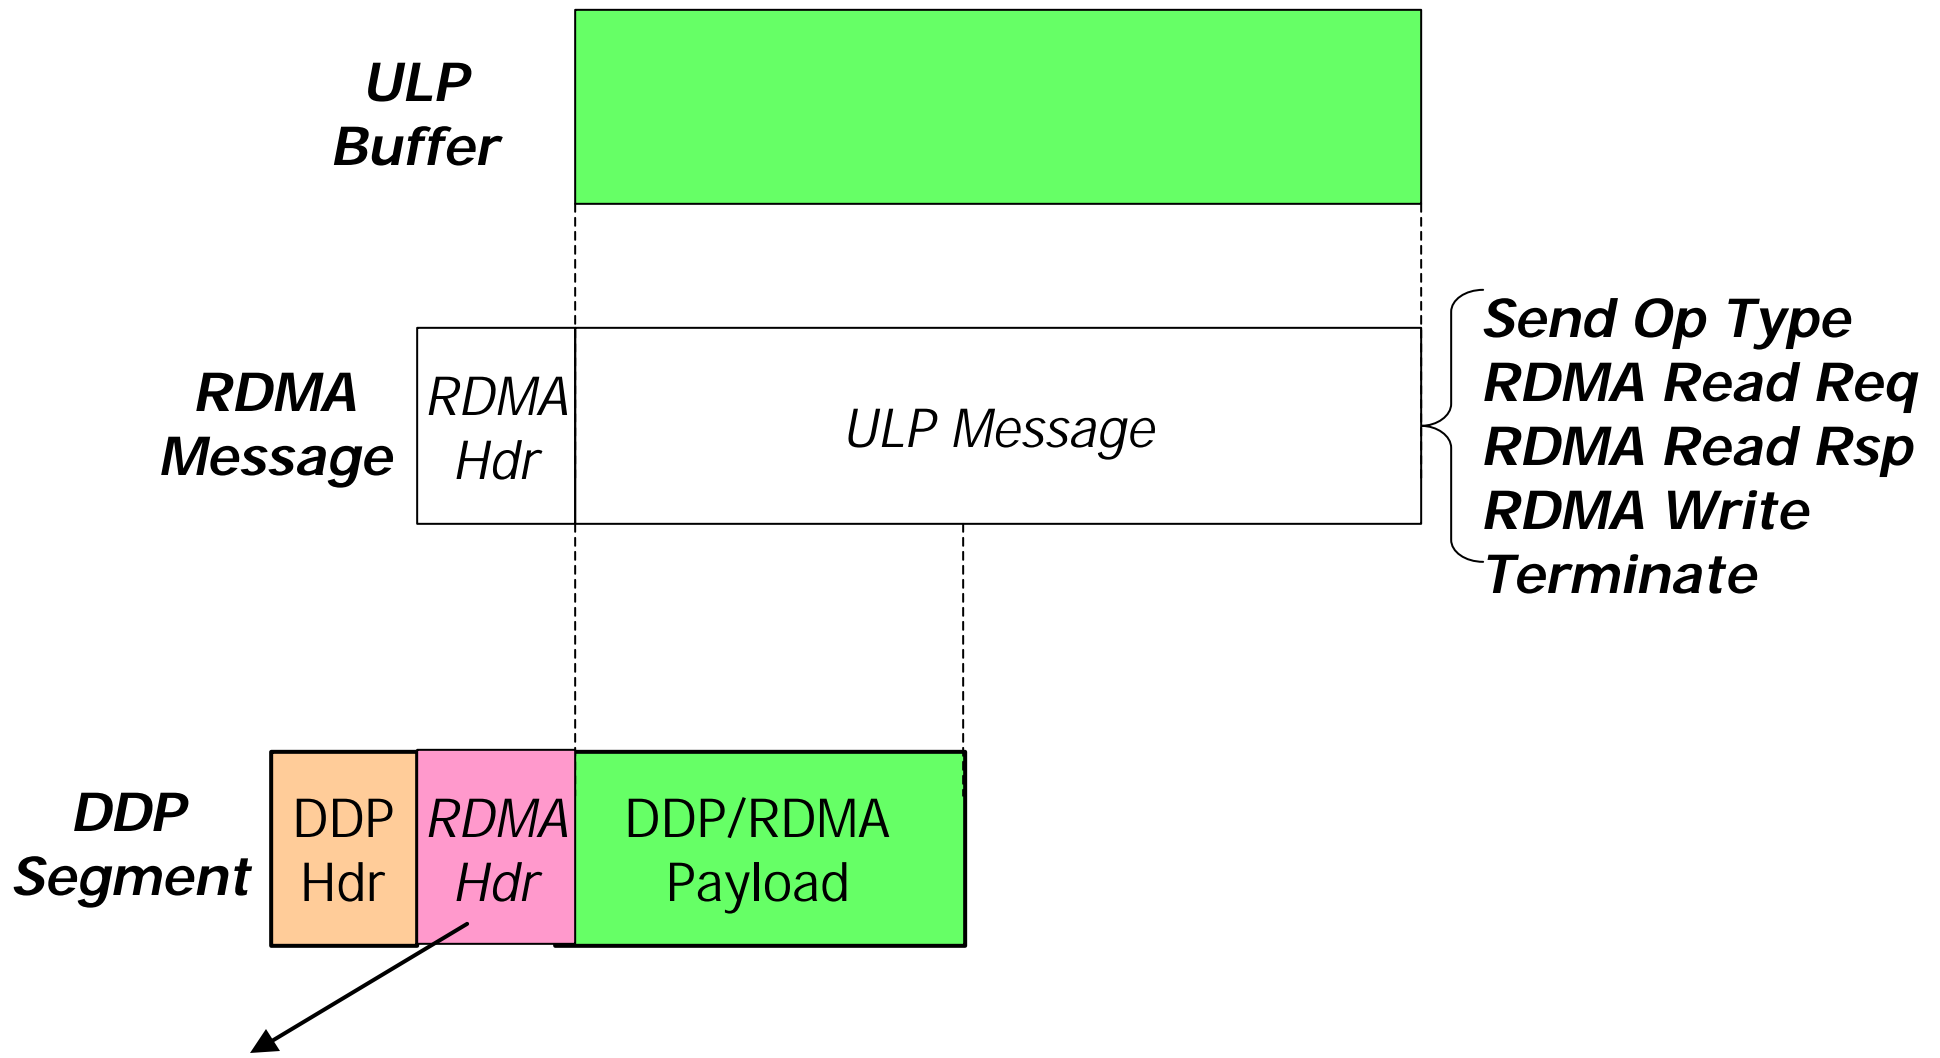
- **Direct Data Placement** – A mechanism whereby **ULP data** contained within DDP Segments **may be Placed directly into its final destination** in memory **without ULP processing**, even when the DDP Segments arrive out of order.
- **DDP** – A wire protocol that supports Direct Data Placement by associating explicit memory buffer placement information with the LLP payload units.
 - ◆ Alone, **supports write only**.
- **Remote Direct Memory Access** – A method of accessing memory on a remote system in which the local system specifies the remote location of the data to be transferred.
 - ◆ Supports, **read and write**.
- **RDMA** – A wire protocol that supports RDMA Operations to transfer ULP data between a Local Peer and the Remote Peer.

Layering Overview



- RDMA – Converts RDMA Write, RDMA Read, and Sends into a DDP Message(s).
- DDP – Segments outbound DDP Messages into 1 or more DDP Segments; reassembles 1 or more DDP Segments into a DDP Message.
- DDP places a set of requirements on the transport, which include:
 - ◆ Reliable delivery
 - ◆ Preservation of DDP Segment and Message boundaries
 - ◆ Strong digest
 - ◆ May provide data out of order, but must specify the sender specified order.
 - ◆ Length of incoming DDP Segments

RDMA/DDP Header Format




Note: The RDMA header is not needed for all RDMA Messages.

DDP Overview

■ DDP supports two data transfer models:


◆ Tagged Buffer data transfer model

- ★ Data Sink advertises an identifier (STag) for the ULP buffer.
- ★ Data Source specifies STag and Tagged Offset (TO) to transfer data to a portion of the Tagged Buffer.
- ★ Allows multiple DDP Messages targeted to a Tagged Buffer with a single buffer advertisement.



DDP Overview (Continued)

- ◆ Untagged Buffer data transfer model
 - ★ Enables data transfer without requiring buffer advertisement.
 - ★ Receiver can queue up a series of ULP buffers to specify the order in which the buffers will be consumed.
 - ★ Each Untagged DDP Message from the Data Source consumes an Untagged Buffer at the Data Sink
 - ★ Requires associating a receive ULP buffer for each DDP Message
 - ▲ If an Untagged DDP Message arrives without an associated Untagged Buffer, the DDP Message is dropped and DDP Stream is terminated .



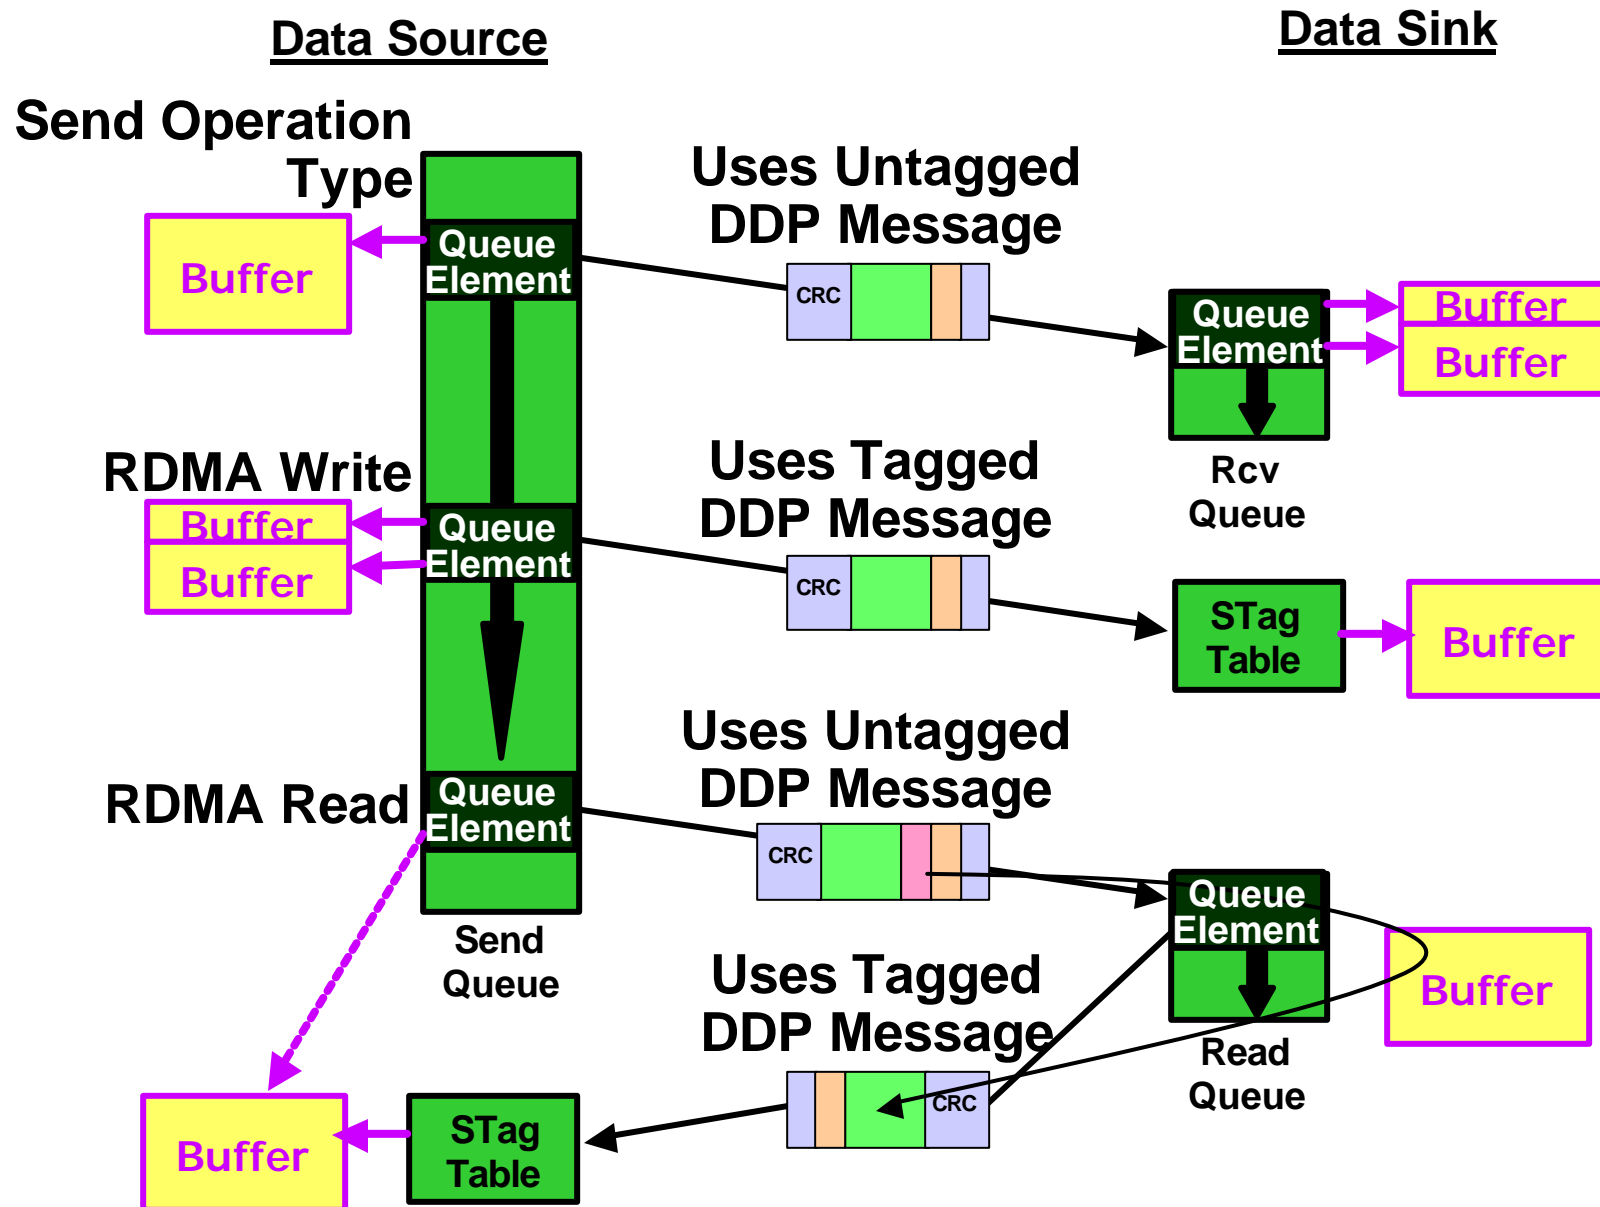
RDMAP Function

- RDMAP enables receive side, data copy removal through the use of the following operations:
 - ◆ Send Operation Type – Transfers data from a local buffer to a remote buffer that has not been explicitly advertised.
 - ◆ RDMA Write – Transfers data from a local buffer to a remote buffer that has been explicitly advertised.
 - ◆ RDMA Read – Retrieves data from a remote buffer that has been explicitly advertised and places it into a local buffer.

For each of the above, the data transfer length can be up to $2^{32} - 1$ octets.

- RDMAP also provides a Terminate operation that transfers information associated with a local error.

RDMA Data Transfer Mechanisms



Documents

■RDMA:

- ◆draft-recio-iwarp-rdma-01.txt

■DDP

- ◆draft-shah-iwarp-ddp-01.txt