NAME

rdma_accept - Called to accept a connection request.

SYNOPSIS

#include <rdma/rdma_cma.h>

int rdma_accept (struct rdma_cm_id **id*, struct rdma_conn_param **conn_param*);

ARGUMENTS

id Connection identifier associated with the request.

conn_param Information needed to establish the connection. See CONNECTION PROPERTIES below for details.

DESCRIPTION

Called from the listening side to accept a connection or datagram service lookup request.

RETURN VALUE

Returns 0 on success, or -1 on error. If an error occurs, errno will be set to indicate the failure reason.

NOTES

Unlike the socket accept routine, rdma_accept is not called on a listening rdma_cm_id. Instead, after calling rdma_listen, the user waits for an RDMA_CM_EVENT_CONNECT_REQUEST event to occur. Connection request events give the user a newly created rdma_cm_id, similar to a new socket, but the rdma_cm_id is bound to a specific RDMA device. rdma_accept is called on the new rdma_cm_id.

CONNECTION PROPERTIES

The following properties are used to configure the communication and specified by the conn_param parameter when accepting a connection or datagram communication request. Users should use the rdma_conn_param values reported in the connection request event to determine appropriate values for these fields when accepting. Users may reference the rdma_conn_param structure in the connection event directly, or can reference their own structure. If the rdma_conn_param structure from an event is referenced, the event must not be acked until after this call returns.

If the conn_param parameter is NULL, the values reported in the connection request event are used, adjusted down based on local hardware restrictions.

private_data

References a user-controlled data buffer. The contents of the buffer are copied and transparently

passed to the remote side as part of the communication request. May be NULL if private_data is not required.

private_data_len
Specifies the size of the user-controlled data buffer. Note that the actual amount of data transferred to the remote side is transport dependent and may be larger than that requested.

responder_resources
The maximum number of outstanding RDMA read and atomic operations that the local side will accept from the remote side. Applies only to RDMA_PS_TCP. This value must be less than or equal to the local RDMA device attribute max_qp_rd_atom, but preferably greater than or equal to the responder_resources value reported in the connect request event.

initiator_depth
The maximum number of outstanding RDMA read and atomic operations that the local side will have to the remote side. Applies only to RDMA_PS_TCP. This value must be less than or equal to the local RDMA device attribute max_qp_init_rd_atom and the initiator_depth value reported in the connect request event.

flow_control
Specifies if hardware flow control is available. This value is exchanged with the remote peer and is not used to configure the QP. Applies only to RDMA_PS_TCP.

retry_count
This value is ignored.

rnr_retry_count
The maximum number of times that a send operation from the remote peer should be retried on a connection after receiving a receiver not ready (RNR) error. RNR errors are generated when a send request arrives before a buffer has been posted to receive the incoming data. Applies only to RDMA_PS_TCP.

srq Specifies if the QP associated with the connection is using a shared receive queue. This field is ignored by the library if a QP has been created on the rdma_cm_id. Applies only to RDMA_PS_TCP.

qp_num
Specifies the QP number associated with the connection. This field is ignored by the library if a QP has been created on the rdma_cm_id.

INFINIBAND SPECIFIC

In addition to the connection properties defined above, InfiniBand QPs are configured with minimum RNR NAK timer and local ACK timeout values. The minimum RNR NAK timer value is set to 0, for a delay of 655 ms. The local ACK timeout is calculated based on the packet lifetime and local HCA ACK delay. The packet lifetime is determined by the InfiniBand Subnet Administrator and is part of the route (path record) information obtained by the active side of the connection. The HCA ACK delay is a property of the locally used HCA.

The RNR retry count is a 3-bit value.

The length of the private data provided by the user is limited to 196 bytes for RDMA_PS_TCP, or 136 bytes for RDMA_PS_UDP.

SEE ALSO

rdma_listen(3), rdma_reject(3), rdma_get_cm_event(3)