
RPMsg RTOS Layer User's Guide

NXP Semiconductors

Document Number: RPMSGRTOSLAYERUG

Rev. 1

Jul 2016



Contents

Chapter 1 RPMMsg Component

1.1	Motivation to create RTOS aware API	1
1.2	Implementation	2
1.3	Usage	2
1.4	RPMMsg porting sub-layers	5

Chapter 2 RPMMsg Zero-copy Layer

2.1	Overview	7
2.2	Function Documentation	7
2.2.1	rpmsg_hold_rx_buffer	7
2.2.2	rpmsg_release_rx_buffer	7
2.2.3	rpmsg_alloc_tx_buffer	9
2.2.4	rpmsg_send_offchannel_nocopy	9
2.2.5	rpmsg_sendto_nocopy	10
2.2.6	rpmsg_send_nocopy	11

Chapter 3 RPMMsg RTOS Layer

3.1	Overview	13
3.2	Function Documentation	13
3.2.1	rpmsg_rtos_init	13
3.2.2	rpmsg_rtos_deinit	14
3.2.3	rpmsg_rtos_create_ept	14
3.2.4	rpmsg_rtos_destroy_ept	15
3.2.5	rpmsg_rtos_recv	15

Section number	Title	Page
3.2.6	rpmsg_rtos_recv_nocopy	16
3.2.7	rpmsg_rtos_recv_nocopy_free	16
3.2.8	rpmsg_rtos_alloc_tx_buffer	17
3.2.9	rpmsg_rtos_send	17
3.2.10	rpmsg_rtos_send_nocopy	18

Chapter 4 Revision History

Chapter 1

RPMMsg Component

The Remote Processor Messaging (RPMMsg) is a virtio-based messaging bus that allows Inter Processor Communications (IPC) between independent software contexts running on homogeneous or heterogeneous cores present in an Asymmetric Multi Processing (AMP) system. The RPMMsg component source code has been recently published as a part of the Open Asymmetric Multi Processing (OpenAMP) Framework. The RPMMsg API is compliant with the RPMMsg bus infrastructure present in upstream Linux 3.4.x kernel onward.

This document describes the extension of the RPMMsg API designed and implemented by NXP. It discusses the motivation for these changes in the RPMMsg as well as the advantages of the extension. This document also serves as the API reference, covering all newly added API functions that can be use in an RPMMsg-based application.

The RPMMsg extension is based on the OpenAMP repository code. See <https://github.com/OpenAMP/open-amp.git> / SHA1 ID 44b5f3c0a6458f3cf80. The documentation for the legacy RPMMsg core code can be found in the *docs* folder of this repository.

1.1 Motivation to create RTOS aware API

The original RPMMsg API is based on processing the transmitted data (messages) in the interrupt context via a registered receive callback, which is called when a message is received. The inconvenience of this approach is that either all the processing of received data must be done in the interrupt context, or that the message must be copied in a temporary application buffer for later processing. Both usages of the available API are not compatible with the concept of a Real-Time Operating System (RTOS), since the interrupt always preempts the running task, whatever its current priority, and this interruption can occur at a random date and can take a random amount of time to execute. This can introduce additional jitter in the real-time system timing. Additionally, the practice in application development using an RTOS is to have multiple independent sequential contexts. It is more natural and convenient to have a blocking sequential API, which was not available in the original RPMMsg API. A good example of a blocking API is a socket interface or any POSIX-like interface. Therefore, the natural trend is to provide this a kind of interface to the application programmer. To summarize, the advantages of the RTOS-aware extension of RPMMsg API are the following:

- No data processing in the interrupt context
- Blocking receive API
- Zero-copy send and receive API
- Receive with timeout provided by RTOS
- Compatible with Linux OS upstream

1.2 Implementation

The NXP contribution to the RPMsg consists of two additional layers that are created above the origin base RPMsg layer.

- The RPMsg Extension layer allows users to allocate and release virtio tx buffers, as well as implements the zero-copy send functionality. [The RPMsg Extension layer API](#) is intended to be used in Bare Metal applications.
- The RPMsg RTOS layer addresses RTOS-based application needs discussed above (handling received data outside the interrupt context, blocking receive API implementation, zero-copy mechanisms). See [RPMsg RTOS layer API](#). This RTOS aware RPMsg API layer is split into multiple C modules. The module *rpmsg_rtos.c/h* contains a generic implementation, which does not depend on the used RTOS nor on the used platform. In */porting/<device>/platform.c/h* and *platform_info.c*, there are platform (SoC) dependent functions. In */porting/ew/<rtos name>/rpmsg_porting.c/h*, the RTOS abstraction is implemented using functions from the *platform.h* to make connection with the hardware. However, the *rpmsg_porting.c/h* module itself is hardware-independent.

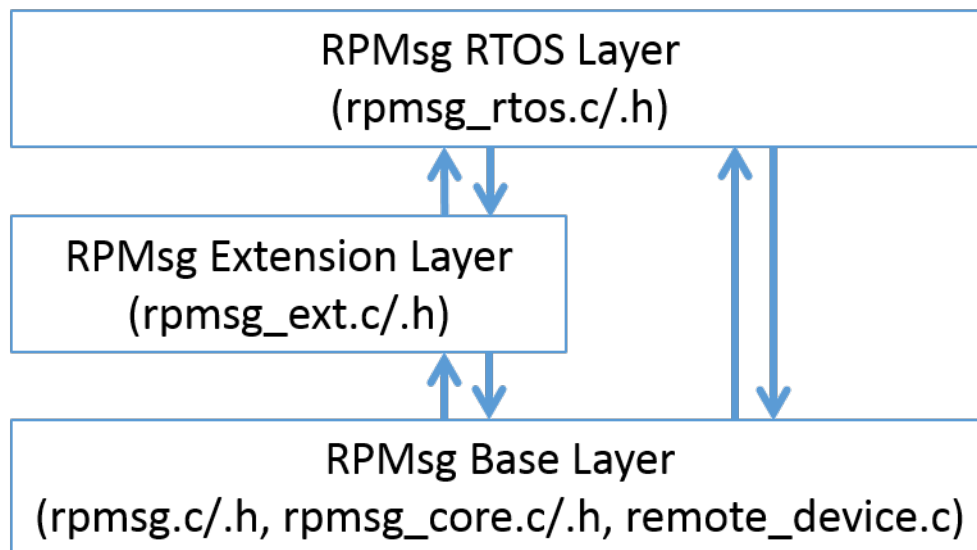


Figure 1.2.1: RPMsg layers

1.3 Usage

To access the RPMsg RTOS layer API, it is necessary to include the *rpmsg_rtos.h* file in the application C module (i.e., *main.c*). After the RTOS startup, the user should call the [rpmsg_rtos_init\(\)](#) function to initialize the RPMsg and to synchronize with the opposite side (other core). After this, application endpoints can be created in any RTOS threads by calling [rpmsg_rtos_create_ept\(\)](#). Consequently, the [rpmsg_rtos_send\(\)](#) function is used to send data from an endpoint to a remote endpoint, whose address is specified in the function call. The [rpmsg_rtos_recv\(\)](#) is then used to receive data on an endpoint or to wait for data to be received with a certain timeout (or the timeout can be set to wait forever).

If the application is low on memory or needs to be more memory efficient and faster, the no-copy mechanism can be used. The RPMsg RTOS layer implements no-copy mechanisms for both sending and

receiving operations. These methods require specifics that have to be considered when used in an application.

no-copy-send mechanism: This mechanism allows sending messages without the cost for copying data from the application buffer to the RPMsg/virtio buffer in the shared memory. The sequence of no-copy sending steps to be performed is as follows:

- Call the `rpmsg_rtos_alloc_tx_buffer()` function to get the virtio buffer and provide the buffer pointer to the application.
- Fill the data to be sent into the pre-allocated virtio buffer. Ensure that the filled data does not exceed the buffer size (provided as the `rpmsg_rtos_alloc_tx_buffer()` *size* output parameter).
- Call the `rpmsg_rtos_send_nocopy()` function to send the message to the destination endpoint. Consider the cache functionality and the virtio buffer alignment. See the `rpmsg_rtos_send_nocopy()` function description below.

no-copy-receive mechanism: This mechanism allows reading messages without the cost for copying data from the virtio buffer in the shared memory to the application buffer. The sequence of no-copy receiving steps to be performed is as follows:

- Call the `rpmsg_rtos_recv_nocopy()` function to get the virtio buffer pointer to the received data.
- Read received data directly from the shared memory.
- Call the `rpmsg_rtos_recv_nocopy_free()` function to release the virtio buffer and to make it available for the next data transfer.

Usage

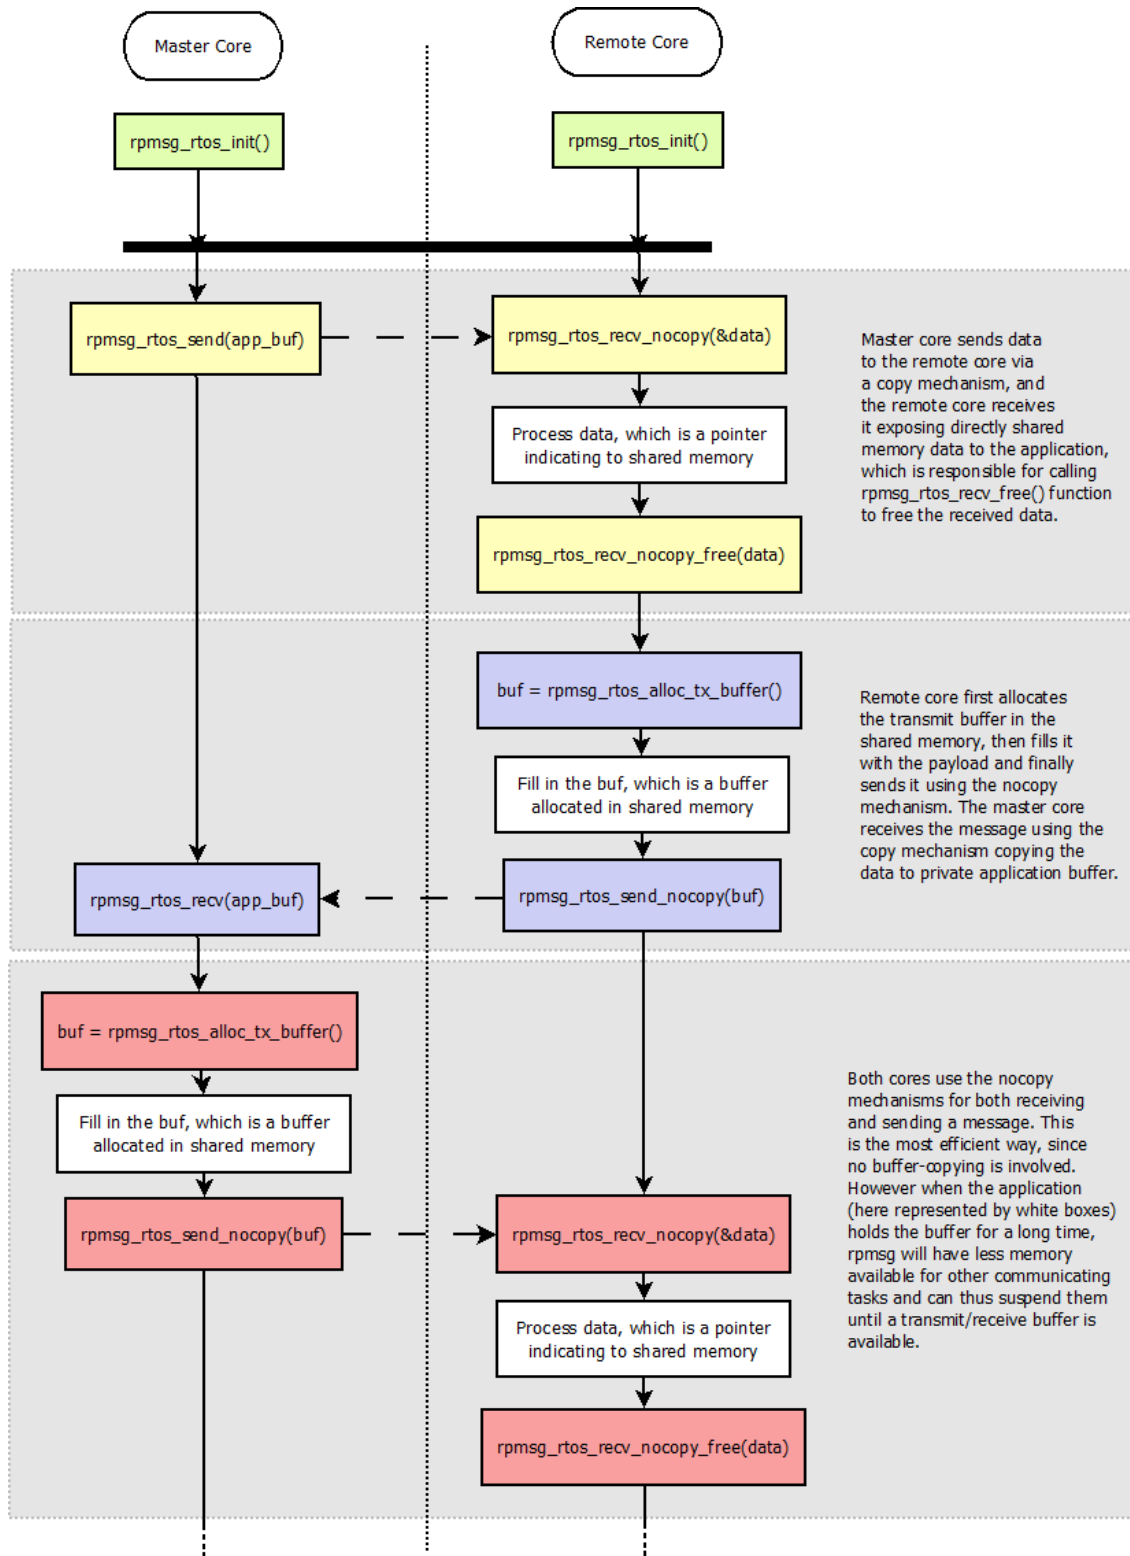


Figure 1.3.1: RPMsg send/receive mechanisms

When deinitializing the RPMMsg communication, the master side calls the `rpmsg_rtos_deinit()` function that deinitializes all on the master side, and also triggers the Name Service (NS) destroy callback on the remote side, which destroys the default channel and the default endpoint. From that time onwards, any call of send or receive API on the remote side returns an error. It is up to the user application to gracefully stop the RPMMsg, i.e., to destroy all application-created endpoints (`rpmsg_rtos_destroy_ept()`) first, then destroy the RPMMsg component (`rpmsg_rtos_deinit()`).

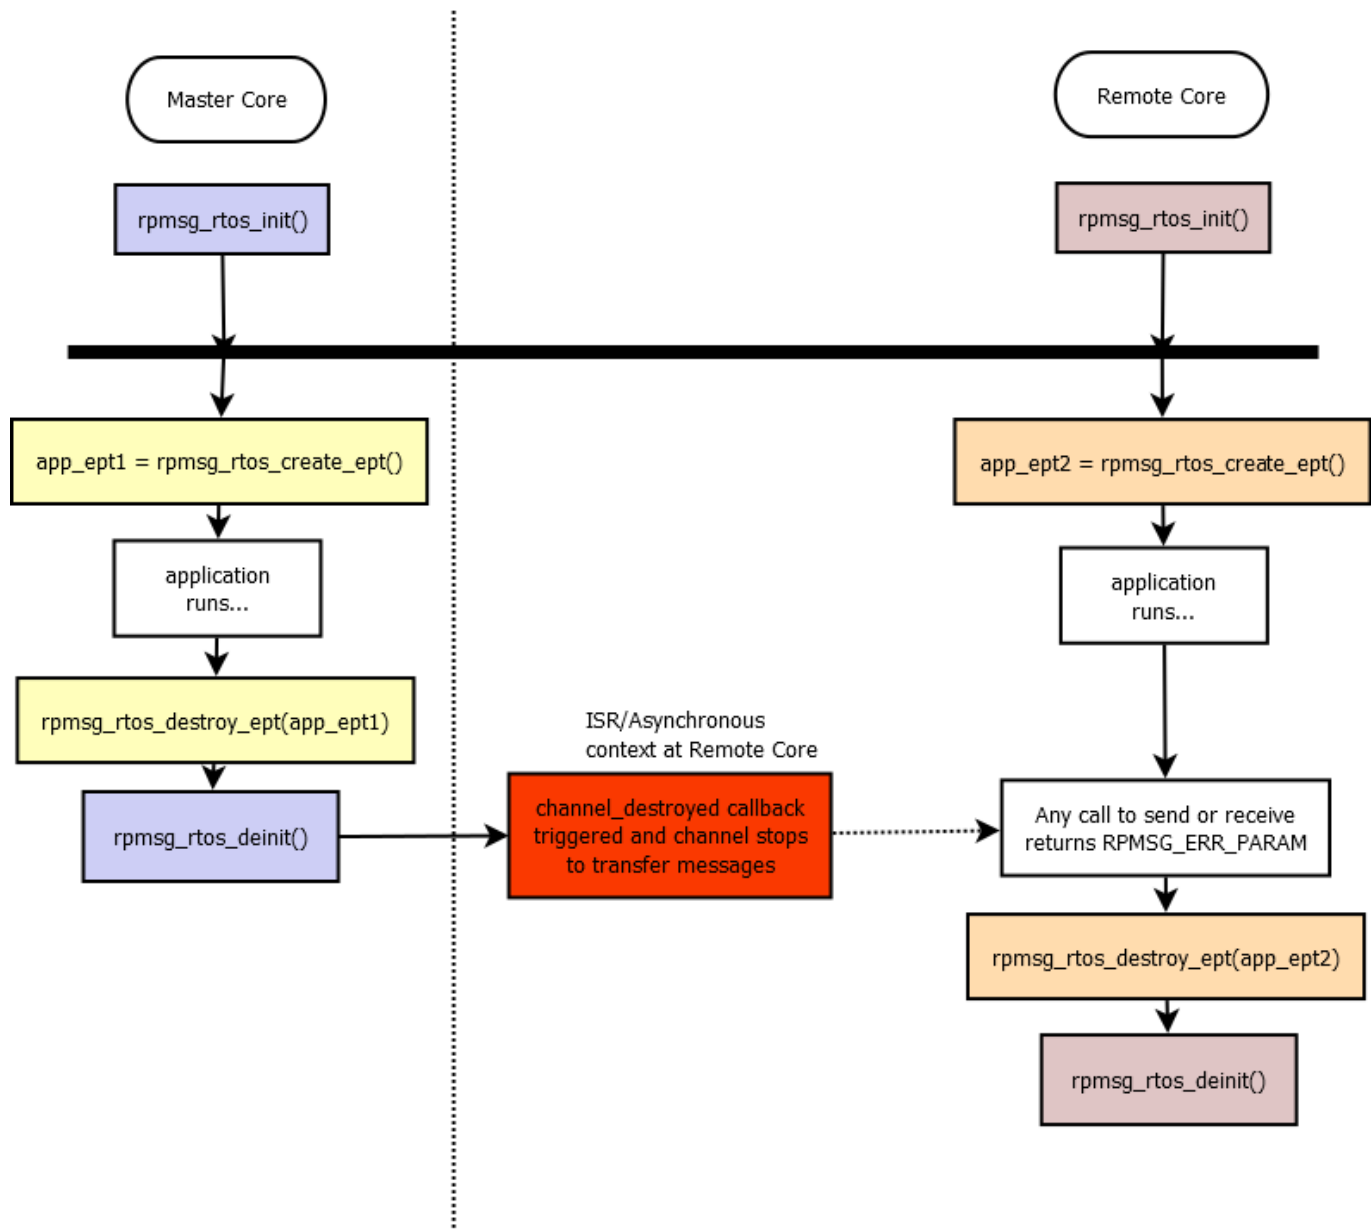


Figure 1.3.2: RPMMsg init and deinit process in RTOS environment

1.4 RPMMsg porting sub-layers

The RPMMsg porting layers have been also modified and consolidated in order to

RPMMsg porting sub-layers

- Strictly separate platform-related (multicore device) and environment-related (Bare Metal, RTOS) layers.
- Update the environment layer API by functions requested by the RTOS layer. The following *env* functions have been introduced:
 - *int env_create_queue(void** queue, int length, int element_size)*
 - *void env_delete_queue(void* queue)*
 - *int env_put_queue(void* queue, void* msg, int timeout_ms)*
 - *int env_get_queue(void* queue, void* msg, int timeout_ms)*

Currently, the environment layer is implemented for Bare Metal and FreeRTOS. To support other RTOSes, it is necessary to create (clone) the *rpmsg_porting.c/h* sub-layer using the desired RTOS API, put this code into the */porting/env/<rtos name>* folder, and to include this path into the list of the project include paths.

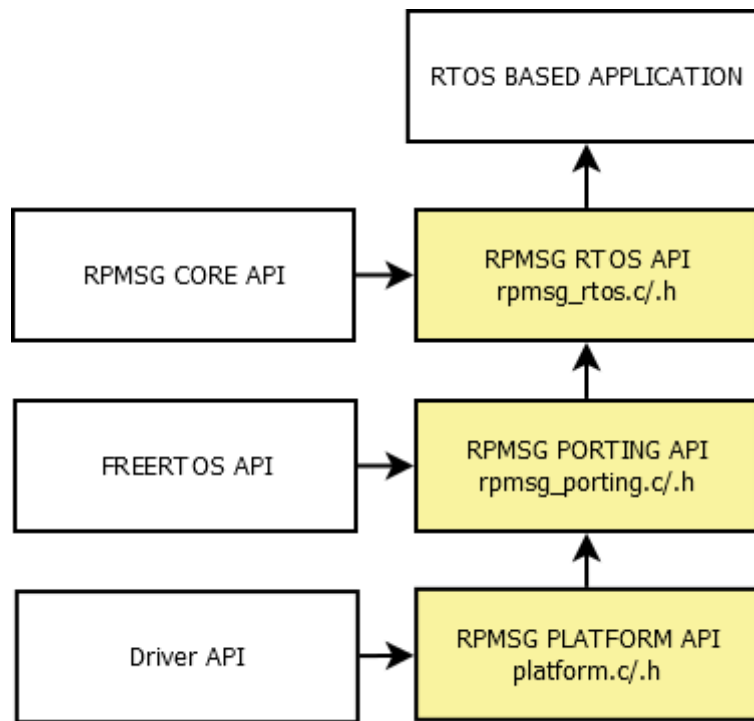


Figure 1.4.1: Rewrite *rpmsg_porting.c/h* sub-layer

Chapter 2 RPMMsg Zero-copy Layer

2.1 Overview

This part describes the RPMMsg extension layer that allows:

- Allocation/release of the virtio tx buffer.
- Zero-copy send functionality.

Functions

- void [rpmmsg_hold_rx_buffer](#) (struct rpmmsg_channel *rpdev, void *rxbuf)
- void [rpmmsg_release_rx_buffer](#) (struct rpmmsg_channel *rpdev, void *rxbuf)
- void * [rpmmsg_alloc_tx_buffer](#) (struct rpmmsg_channel *rpdev, unsigned long *size, int wait)
- int [rpmmsg_send_offchannel_nocopy](#) (struct rpmmsg_channel *rpdev, unsigned long src, unsigned long dst, void *txbuf, int len)
- static int [rpmmsg_sendto_nocopy](#) (struct rpmmsg_channel *rpdev, void *txbuf, int len, unsigned long dst)
- static int [rpmmsg_send_nocopy](#) (struct rpmmsg_channel *rpdev, void *txbuf, int len)

2.2 Function Documentation

2.2.1 void rpmmsg_hold_rx_buffer (struct rpmmsg_channel * rpdev, void * rxbuf)

Holds the rx buffer for usage outside the receive callback.

Calling this function prevents the RPMMsg receive buffer from being released back to the pool of shmem buffers. This API can only be called at rx callback context ([rpmmsg_rx_cb_t](#)). With this API, the application doesn't need to copy the message in rx callback. Instead, the rx buffer base address is saved in application context and further processed in application process. After the message is processed, the application can release the rx buffer for future reuse in vring by calling the [rpmmsg_release_rx_buffer\(\)](#) function.

Parameters

in	<i>rpdev</i>	The RPMMsg channel
in	<i>rxbuf</i>	RX buffer with message payload

See Also

[rpmmsg_release_rx_buffer](#)

2.2.2 void rpmmsg_release_rx_buffer (struct rpmmsg_channel * rpdev, void * rxbuf)

Releases the rx buffer for future reuse in vring.

Function Documentation

This API can be called at process context when the message in rx buffer is processed.

Parameters

<i>rpdev</i>	- the RPMsg channel
<i>rxbuf</i>	- rx buffer with message payload

See Also

[rpmsg_hold_rx_buffer](#)

2.2.3 void* rpmsg_alloc_tx_buffer (struct rpmsg_channel * *rpdev*, unsigned long * *size*, int *wait*)

Allocates the tx buffer for message payload.

This API can only be called at process context to get the tx buffer in vring. By this way, the application can directly put its message into the vring tx buffer without copy from an application buffer. It is the application responsibility to correctly fill the allocated tx buffer by data and passing correct parameters to the [rpmsg_send_nocopy\(\)](#) or [rpmsg_sendto_nocopy\(\)](#) function to perform data no-copy-send mechanism.

Parameters

in	<i>rpdev</i>	Pointer to RPMsg channel
in	<i>size</i>	Pointer to store tx buffer size
in	<i>wait</i>	Boolean, wait or not for buffer to become available

Returns

The tx buffer address on success and NULL on failure

See Also

[rpmsg_send_offchannel_nocopy](#)
[rpmsg_sendto_nocopy](#)
[rpmsg_send_nocopy](#)

2.2.4 int rpmsg_send_offchannel_nocopy (struct rpmsg_channel * *rpdev*, unsigned long *src*, unsigned long *dst*, void * *txbuf*, int *len*)

Sends a message in tx buffer allocated by [rpmsg_alloc_tx_buffer\(\)](#) using explicit src/dst addresses.

This function sends txbuf of length len to the remote dst address, and uses src as the source address. The message will be sent to the remote processor which the rpdev channel belongs to. The application has to take the responsibility for:

Function Documentation

1. tx buffer allocation ([rpmmsg_alloc_tx_buffer\(\)](#))
2. filling the data to be sent into the pre-allocated tx buffer
3. not exceeding the buffer size when filling the data
4. data cache coherency

After the [rpmmsg_send_offchannel_nocopy\(\)](#) function is issued the tx buffer is no more owned by the sending task and must not be touched anymore unless the [rpmmsg_send_offchannel_nocopy\(\)](#) function fails and returns an error. In that case the application should try to re-issue the [rpmmsg_send_offchannel_nocopy\(\)](#) again and if it is still not possible to send the message and the application wants to give it up from whatever reasons the [rpmmsg_release_rx_buffer](#) function could be called, passing the pointer to the tx buffer to be released as a parameter.

Parameters

in	<i>rpdev</i>	The RPMsg channel
in	<i>src</i>	Source address
in	<i>dst</i>	Destination address
in	<i>txbuf</i>	TX buffer with message filled
in	<i>len</i>	Length of payload

Returns

0 on success and an appropriate error value on failure

See Also

[rpmmsg_alloc_tx_buffer](#)
[rpmmsg_sendto_nocopy](#)
[rpmmsg_send_nocopy](#)

2.2.5 static int rpmmsg_sendto_nocopy (struct rpmmsg_channel * *rpdev*, void * *txbuf*, int *len*, unsigned long *dst*) [static]

Sends a message in tx buffer allocated by [rpmmsg_alloc_tx_buffer\(\)](#) across to the remote processor, specify *dst*.

This function sends *txbuf* of length *len* to the remote *dst* address. The message will be sent to the remote processor which the *rpdev* channel belongs to, using *rpdev*'s source address. The application has to take the responsibility for:

1. tx buffer allocation ([rpmmsg_alloc_tx_buffer\(\)](#))
2. filling the data to be sent into the pre-allocated tx buffer
3. not exceeding the buffer size when filling the data
4. data cache coherency

After the [rpmsg_sendto_nocopy\(\)](#) function is issued the tx buffer is no more owned by the sending task and must not be touched anymore unless the [rpmsg_sendto_nocopy\(\)](#) function fails and returns an error. In that case the application should try to re-issue the [rpmsg_sendto_nocopy\(\)](#) again and if it is still not possible to send the message and the application wants to give it up from whatever reasons the [rpmsg_release_rx_buffer](#) function could be called, passing the pointer to the tx buffer to be released as a parameter.

Parameters

in	<i>rpdev</i>	The RPMsg channel
in	<i>txbuf</i>	TX buffer with message filled
in	<i>len</i>	Length of payload
in	<i>dst</i>	Destination address

Returns

0 on success and an appropriate error value on failure

See Also

[rpmsg_alloc_tx_buffer](#)
[rpmsg_send_offchannel_nocopy](#)
[rpmsg_send_nocopy](#)

2.2.6 static int rpmsg_send_nocopy (struct rpmsg_channel * *rpdev*, void * *txbuf*, int *len*) [static]

Sends a message in tx buffer allocated by [rpmsg_alloc_tx_buffer\(\)](#) across to the remote processor.

This function sends txbuf of length len on the rpdev channel. The message will be sent to the remote processor which the rpdev channel belongs to, using rpdev’s source and destination addresses. The application has to take the responsibility for:

1. tx buffer allocation ([rpmsg_alloc_tx_buffer\(\)](#))
2. filling the data to be sent into the pre-allocated tx buffer
3. not exceeding the buffer size when filling the data
4. data cache coherency

After the [rpmsg_send_nocopy\(\)](#) function is issued the tx buffer is no more owned by the sending task and must not be touched anymore unless the [rpmsg_send_nocopy\(\)](#) function fails and returns an error. In that case the application should try to re-issue the [rpmsg_send_nocopy\(\)](#) again and if it is still not possible to send the message and the application wants to give it up from whatever reasons the [rpmsg_release_rx_buffer](#) function could be called, passing the pointer to the tx buffer to be released as a parameter.

Function Documentation

Parameters

in	<i>rpdev</i>	The RPMsg channel
in	<i>txbuf</i>	TX buffer with message filled
in	<i>len</i>	Length of payload

Returns

0 on success and an appropriate error value on failure

See Also

[rpmsg_alloc_tx_buffer](#)
[rpmsg_send_offchannel_nocopy](#)
[rpmsg_sendto_nocopy](#)

Chapter 3

RPMMsg RTOS Layer

3.1 Overview

This part describes the RPMMsg RTOS adaptation layer that allows:

- Handling of received messages outside the interrupt context.
- Implementation of blocking API for the RPMMsg receive side.
- Provides zero-copy receive functionality.
- Provides zero-copy send functionality.

Functions

- int `rpmsg_rtos_init` (int `dev_id`, struct `remote_device` `**rdev`, int `role`, struct `rpmsg_channel` `**def_chnl`)
- void `rpmsg_rtos_deinit` (struct `remote_device` `*rdev`)
- struct `rpmsg_endpoint` * `rpmsg_rtos_create_ept` (struct `rpmsg_channel` `*rp_chnl`, unsigned long `addr`)
- void `rpmsg_rtos_destroy_ept` (struct `rpmsg_endpoint` `*rp_ept`)
- int `rpmsg_rtos_recv` (struct `rpmsg_endpoint` `*ept`, void `*data`, int `*len`, int `maxlen`, unsigned long `*src`, int `timeout_ms`)
- int `rpmsg_rtos_recv_nocopy` (struct `rpmsg_endpoint` `*ept`, void `**data`, int `*len`, unsigned long `*src`, int `timeout_ms`)
- int `rpmsg_rtos_recv_nocopy_free` (struct `rpmsg_endpoint` `*ept`, void `*data`)
- void * `rpmsg_rtos_alloc_tx_buffer` (struct `rpmsg_endpoint` `*ept`, unsigned long `*size`)
- int `rpmsg_rtos_send` (struct `rpmsg_endpoint` `*ept`, void `*data`, int `len`, unsigned long `dst`)
- int `rpmsg_rtos_send_nocopy` (struct `rpmsg_endpoint` `*ept`, void `*txbuf`, int `len`, unsigned long `dst`)

3.2 Function Documentation

3.2.1 int `rpmsg_rtos_init` (int `dev_id`, struct `remote_device` `** rdev`, int `role`, struct `rpmsg_channel` `** def_chnl`)

This function allocates and initializes the RPMMsg driver resources for given device ID (cpu id).

The successful return from this function leaves fully enabled IPC link. RTOS aware version.

Parameters

in	<code>dev_id</code>	Remote device for which driver is to be initialized
out	<code>rdev</code>	Pointer to newly created remote device

Function Documentation

in	<i>role</i>	Role of the other device, Master or Remote
out	<i>def_chnl</i>	Pointer to RPMsg channel

Returns

Status of function execution

See Also

[rpmsg_rtos_deinit](#)

3.2.2 void rpmsg_rtos_deinit (struct remote_device * *rdev*)

This function frees RPMsg driver resources for given remote device.

RTOS aware version.

Parameters

in	<i>rdev</i>	Pointer to device to de-init
----	-------------	------------------------------

See Also

[rpmsg_rtos_init](#)

3.2.3 struct rpmsg_endpoint* rpmsg_rtos_create_ept (struct rpmsg_channel * *rp_chnl*, unsigned long *addr*)

This function creates RPMsg endpoint for the RPMsg channel.

RTOS aware version.

Parameters

in	<i>rp_chnl</i>	Pointer to RPMsg channel
in	<i>addr</i>	Endpoint src address

Returns

Pointer to endpoint control block

See Also

[rpmsg_rtos_destroy_ept](#)

3.2.4 void rpmsg_rtos_destroy_ept (struct rpmsg_endpoint * *rp_ept*)

This function deletes RPMMsg endpoint and performs cleanup.

RTOS aware version.

Parameters

in	<i>rp_ept</i>	Pointer to endpoint to destroy
----	---------------	--------------------------------

See Also

[rpmsg_rtos_create_ept](#)

3.2.5 int rpmsg_rtos_recv (struct rpmsg_endpoint * *ept*, void * *data*, int * *len*, int *maxlen*, unsigned long * *src*, int *timeout_ms*)

RTOS receive function - blocking version of the received function that can be called from an RTOS task.

The data is copied from the receive buffer into the user supplied buffer.

This is the "receive with copy" version of the RPMMsg receive function. This version is simple to use but it requires copying data from shared memory into the user space buffer. The user has no obligation or burden to manage the shared memory buffers.

Parameters

in	<i>ept</i>	Pointer to the RPMMsg endpoint on which data is received
in	<i>data</i>	Pointer to the user buffer the received data are copied to
out	<i>len</i>	Pointer to an int variable that will contain the number of bytes actually copied into the buffer
in	<i>maxlen</i>	Maximum number of bytes to copy (received buffer size)
out	<i>src</i>	Pointer to address of the endpoint from which data is received
in	<i>timeout_ms</i>	Timeout, in milliseconds, to wait for a message. A value of 0 means don't wait (non-blocking call). A value of 0xffffffff means wait forever (blocking call).

Returns

Status of function execution

See Also

[rpmsg_rtos_recv_nocopy](#)

Function Documentation

3.2.6 `int rpmsg_rtos_rcv_nocopy (struct rpmsg_endpoint * ept, void ** data, int * len, unsigned long * src, int timeout_ms)`

RTOS receive function - blocking version of the received function that can be called from an RTOS task.

The data is NOT copied into the user-app. buffer.

This is the "zero-copy receive" version of the RPMMsg receive function. No data is copied. Only the pointer to the data is returned. This version is fast, but it requires the user to manage buffer allocation. Specifically, the user must decide when a buffer is no longer in use and make the appropriate API call to free it, see [rpmsg_rtos_rcv_nocopy_free\(\)](#).

Parameters

in	<i>ept</i>	Pointer to the RPMMsg endpoint on which data is received
out	<i>data</i>	Pointer to the RPMMsg buffer of the shared memory where the received data is stored
out	<i>len</i>	Pointer to an int variable that that will contain the number of valid bytes in the RPMMsg buffer
out	<i>src</i>	Pointer to address of the endpoint from which data is received
in	<i>timeout_ms</i>	Timeout, in milliseconds, to wait for a message. A value of 0 means don't wait (non-blocking call). A value of 0xffffffff means wait forever (blocking call).

Returns

Status of function execution

See Also

[rpmsg_rtos_rcv_nocopy_free](#)
[rpmsg_rtos_rcv](#)

3.2.7 `int rpmsg_rtos_rcv_nocopy_free (struct rpmsg_endpoint * ept, void * data)`

This function frees a buffer previously returned by [rpmsg_rtos_rcv_nocopy\(\)](#).

Once the zero-copy mechanism of receiving data is used, this function has to be called to free a buffer and to make it available for the next data transfer.

Parameters

in	<i>ept</i>	Pointer to the RPMsg endpoint that has consumed received data
in	<i>data</i>	Pointer to the RPMsg buffer of the shared memory that has to be freed

Returns

Status of function execution

See Also

[rpmsg_rtos_rcv_nocopy](#)

3.2.8 void* rpmsg_rtos_alloc_tx_buffer (struct rpmsg_endpoint * *ept*, unsigned long * *size*)

Allocates the tx buffer for message payload.

This API can only be called at process context to get the tx buffer in vring. By this way, the application can directly put its message into the vring tx buffer without copy from an application buffer. It is the application responsibility to correctly fill the allocated tx buffer by data and passing correct parameters to the [rpmsg_rtos_send_nocopy\(\)](#) function to perform data no-copy-send mechanism.

Parameters

in	<i>ept</i>	Pointer to the RPMsg endpoint that requests tx buffer allocation
out	<i>size</i>	Pointer to store tx buffer size

Returns

The tx buffer address on success and NULL on failure

See Also

[rpmsg_rtos_send_nocopy](#)

3.2.9 int rpmsg_rtos_send (struct rpmsg_endpoint * *ept*, void * *data*, int *len*, unsigned long *dst*)

Sends a message across to the remote processor.

This function sends data of length *len* to the remote *dst* address. In case there are no TX buffers available, the function will block until one becomes available, or a timeout of 15 seconds elapses. When the latter happens, -ERESTARTSYS is returned.

Function Documentation

Parameters

in	<i>ept</i>	Pointer to the RPMsg endpoint
in	<i>data</i>	Pointer to the application buffer containing data to be sent
in	<i>len</i>	Size of the data, in bytes, to transmit
in	<i>dst</i>	Destination address of the message

Returns

0 on success and an appropriate error value on failure

See Also

[rpmsg_rtos_send_nocopy](#)

3.2.10 int rpmsg_rtos_send_nocopy (struct rpmsg_endpoint * *ept*, void * *txbuf*, int *len*, unsigned long *dst*)

Sends a message in tx buffer allocated by [rpmsg_rtos_alloc_tx_buffer\(\)](#) to the remote processor.

This function sends txbuf of length len to the remote dst address. The application has to take the responsibility for:

1. tx buffer allocation ([rpmsg_rtos_alloc_tx_buffer\(\)](#))
2. filling the data to be sent into the pre-allocated tx buffer
3. not exceeding the buffer size when filling the data
4. data cache coherency

After the [rpmsg_rtos_send_nocopy\(\)](#) function is issued the tx buffer is no more owned by the sending task and must not be touched any more unless the [rpmsg_rtos_send_nocopy\(\)](#) function fails and returns an error. In that case the application should try to re-issue the [rpmsg_rtos_send_nocopy\(\)](#) again and if it is still not possible to send the message and the application wants to give it up from whatever reasons the [rpmsg_rtos_rcv_nocopy_free](#) function could be called, passing the pointer to the tx buffer to be released as a parameter.

Parameters

in	<i>ept</i>	Pointer to the RPMsg endpoint
in	<i>txbuf</i>	Tx buffer with message filled
in	<i>len</i>	Size of the data, in bytes, to transmit
in	<i>dst</i>	Destination address of the message

Returns

0 on success and an appropriate error value on failure

See Also

[rpmmsg_rtos_alloc_tx_buffer](#)
[rpmmsg_rtos_send](#)

Chapter 4

Revision History

This table summarizes revisions to this document.

Revision number	Date	Substantive changes
0	09/2015	Initial release
0.1	11/2015	Update for Extension layer
1	07/2016	Updates for release

How to Reach Us:

Home Page:
nxp.com

Web Support:
nxp.com/support

Information in this document is provided solely to enable system and software implementers to use Freescale products. There are no express or implied copyright licenses granted hereunder to design or fabricate any integrated circuits based on the information in this document.

Freescale reserves the right to make changes without further notice to any products herein. Freescale makes no warranty, representation, or guarantee regarding the suitability of its products for any particular purpose, nor does Freescale assume any liability arising out of the application or use of any product or circuit, and specifically disclaims any and all liability, including without limitation consequential or incidental damages. "Typical" parameters that may be provided in Freescale data sheets and/or specifications can and do vary in different applications, and actual performance may vary over time. All operating parameters, including "typicals," must be validated for each customer application by customer's technical experts. Freescale does not convey any license under its patent rights nor the rights of others. Freescale sells products pursuant to standard terms and conditions of sale, which can be found at the following address: nxp.com/SalesTermsandConditions.

Freescale, the Freescale logo, and Kinetis are trademarks of Freescale Semiconductor, Inc., Reg. U.S. Pat. & Tm. Off. All other product or service names are the property of their respective owners. ARM, ARM Powered Logo, and Cortex are registered trademarks of ARM limited (or its subsidiaries) in the EU and/or elsewhere. All rights reserved.

© 2016 Freescale Semiconductor, Inc.

