

# Unified Communication X (UCX)

API Standard  
Version 1.13





# Contents

<b>1</b>	<b>Preface</b>	<b>1</b>
1.1	Scope of the Document	1
1.2	Audience	1
1.3	Document Status	1
1.4	License	1
<b>2</b>	<b>Introduction</b>	<b>3</b>
2.1	Motivation	3
2.2	UCX	3
<b>3</b>	<b>Design</b>	<b>5</b>
3.1	UCS	5
3.2	UCT	5
3.3	UCP	6
<b>4</b>	<b>Conventions and Notations</b>	<b>7</b>
4.1	Blocking Behavior	7
4.2	Non-blocking Behavior	7
4.3	Fairness	7
4.4	Interaction with Signal Handler Functions	7
<b>5</b>	<b>Deprecated List</b>	<b>9</b>
<b>6</b>	<b>Module Documentation</b>	<b>11</b>
6.1	Unified Communication Protocol (UCP) API	11
6.1.1	Detailed Description	11
6.2	UCP Application Context	12
6.2.1	Detailed Description	13
6.2.2	Data Structure Documentation	13
6.2.2.1	struct ucp_lib_attr	13
6.2.2.2	struct ucp_context_attr	14
6.2.2.3	struct ucp_tag_rcv_info	14
6.2.2.4	struct ucp_request_param_t	14
6.2.2.5	union ucp_request_param_t.cb	15
6.2.2.6	union ucp_request_param_t.rcv_info	16
6.2.3	Macro Definition Documentation	16
6.2.3.1	UCP_ENTITY_NAME_MAX	16
6.2.4	Typedef Documentation	16
6.2.4.1	ucp_lib_attr_t	16
6.2.4.2	ucp_context_attr_t	16
6.2.4.3	ucp_tag_rcv_info_t	16
6.2.4.4	ucp_context_h	16
6.2.4.5	ucp_request_init_callback_t	17
6.2.4.6	ucp_request_cleanup_callback_t	17
6.2.5	Enumeration Type Documentation	17
6.2.5.1	ucp_params_field	17
6.2.5.2	ucp_feature	18

6.2.5.3	ucp_lib_attr_field	18
6.2.5.4	ucp_context_attr_field	18
6.2.6	Function Documentation	19
6.2.6.1	ucp_lib_query()	19
6.2.6.2	ucp_get_version()	19
6.2.6.3	ucp_get_version_string()	19
6.2.6.4	ucp_init()	19
6.2.6.5	ucp_cleanup()	20
6.2.6.6	ucp_context_query()	21
6.2.6.7	ucp_context_print_info()	21
6.3	UCP Worker	22
6.3.1	Detailed Description	24
6.3.2	Data Structure Documentation	25
6.3.2.1	struct ucp_worker_attr	25
6.3.2.2	struct ucp_worker_params	25
6.3.2.3	struct ucp_worker_address_attr	27
6.3.2.4	struct ucp_listener_attr	27
6.3.2.5	struct ucp_conn_request_attr	27
6.3.2.6	struct ucp_listener_params	27
6.3.2.7	struct ucp_am_handler_param	28
6.3.2.8	struct ucp_am_recv_param	29
6.3.2.9	struct ucp_listener_accept_handler	29
6.3.2.10	struct ucp_listener_conn_handler	29
6.3.3	Typedef Documentation	29
6.3.3.1	ucp_worker_attr_t	29
6.3.3.2	ucp_worker_params_t	30
6.3.3.3	ucp_worker_address_attr_t	30
6.3.3.4	ucp_listener_attr_t	30
6.3.3.5	ucp_conn_request_attr_t	30
6.3.3.6	ucp_listener_params_t	30
6.3.3.7	ucp_am_handler_param_t	30
6.3.3.8	ucp_listener_accept_handler_t	30
6.3.3.9	ucp_am_recv_param_t	30
6.3.3.10	ucp_address_t	30
6.3.3.11	ucp_listener_h	31
6.3.3.12	ucp_worker_h	31
6.3.3.13	ucp_listener_accept_callback_t	31
6.3.3.14	ucp_listener_conn_callback_t	31
6.3.3.15	ucp_listener_conn_handler_t	32
6.3.3.16	ucp_wakeup_event_t	32
6.3.4	Enumeration Type Documentation	32
6.3.4.1	ucp_worker_params_field	32
6.3.4.2	ucp_worker_flags_t	33
6.3.4.3	ucp_listener_params_field	33
6.3.4.4	ucp_worker_address_flags_t	33
6.3.4.5	ucp_worker_attr_field	33
6.3.4.6	ucp_worker_address_attr_field	34
6.3.4.7	ucp_listener_attr_field	34
6.3.4.8	ucp_conn_request_attr_field	34
6.3.4.9	ucp_am_cb_flags	35
6.3.4.10	ucp_send_am_flags	35
6.3.4.11	ucp_wakeup_event_types	35
6.3.5	Function Documentation	36
6.3.5.1	ucp_worker_create()	36
6.3.5.2	ucp_worker_destroy()	36
6.3.5.3	ucp_worker_query()	38
6.3.5.4	ucp_worker_print_info()	38
6.3.5.5	ucp_worker_get_address()	38

6.3.5.6	ucp_worker_release_address()	39
6.3.5.7	ucp_worker_address_query()	39
6.3.5.8	ucp_worker_progress()	40
6.3.5.9	ucp_stream_worker_poll()	40
6.3.5.10	ucp_listener_create()	41
6.3.5.11	ucp_listener_destroy()	41
6.3.5.12	ucp_listener_query()	42
6.3.5.13	ucp_conn_request_query()	42
6.3.5.14	ucp_listener_reject()	42
6.3.5.15	ucp_worker_set_am_handler()	43
6.3.5.16	ucp_worker_set_am_recv_handler()	43
6.3.5.17	ucp_worker_fence()	44
6.3.5.18	ucp_worker_flush_nb()	44
6.3.5.19	ucp_worker_flush_nbx()	45
6.3.5.20	ucp_worker_flush()	45
6.4	UCP Memory routines	47
6.4.1	Detailed Description	48
6.4.2	Data Structure Documentation	48
6.4.2.1	struct ucp_mem_map_params	48
6.4.2.2	struct ucp_mem_advise_params	49
6.4.2.3	struct ucp_mem_attr	49
6.4.3	Typedef Documentation	49
6.4.3.1	ucp_mem_map_params_t	49
6.4.3.2	ucp_mem_advise_t	50
6.4.3.3	ucp_mem_advise_params_t	50
6.4.3.4	ucp_rkey_h	50
6.4.3.5	ucp_mem_h	50
6.4.3.6	ucp_mem_attr_t	50
6.4.4	Enumeration Type Documentation	50
6.4.4.1	ucp_mem_map_params_field	50
6.4.4.2	ucp_mem_advise_params_field	51
6.4.4.3	anonymous enum	51
6.4.4.4	anonymous enum	51
6.4.4.5	ucp_mem_advise	51
6.4.4.6	ucp_mem_attr_field	52
6.4.5	Function Documentation	52
6.4.5.1	ucp_mem_map()	52
6.4.5.2	ucp_mem_unmap()	53
6.4.5.3	ucp_mem_query()	54
6.4.5.4	ucp_mem_print_info()	54
6.4.5.5	ucp_mem_advise()	55
6.4.5.6	ucp_rkey_pack()	55
6.4.5.7	ucp_rkey_buffer_release()	56
6.4.5.8	ucp_ep_rkey_unpack()	56
6.4.5.9	ucp_rkey_ptr()	57
6.4.5.10	ucp_rkey_destroy()	57
6.5	UCP Wake-up routines	58
6.5.1	Detailed Description	58
6.5.2	Function Documentation	58
6.5.2.1	ucp_worker_get_efd()	58
6.5.2.2	ucp_worker_wait()	59
6.5.2.3	ucp_worker_wait_mem()	59
6.5.2.4	ucp_worker_arm()	60
6.5.2.5	ucp_worker_signal()	61
6.6	UCP Endpoint	62
6.6.1	Detailed Description	63
6.6.2	Data Structure Documentation	64
6.6.2.1	struct ucp_ep_evaluate_perf_param_t	64

6.6.2.2	struct ucp_ep_evaluate_perf_attr_t	64
6.6.2.3	struct ucp_stream_poll_ep	64
6.6.2.4	struct ucp_ep_attr	64
6.6.2.5	struct ucp_ep_params	65
6.6.3	Typedef Documentation	66
6.6.3.1	ucp_ep_perf_param_field_t	66
6.6.3.2	ucp_ep_perf_attr_field_t	66
6.6.3.3	ucp_stream_poll_ep_t	66
6.6.3.4	ucp_ep_attr_t	66
6.6.3.5	ucp_ep_h	66
6.6.3.6	ucp_conn_request_h	67
6.6.3.7	ucp_am_callback_t	67
6.6.3.8	ucp_am_recv_callback_t	67
6.6.3.9	ucp_ep_params_t	68
6.6.4	Enumeration Type Documentation	68
6.6.4.1	ucp_ep_params_field	68
6.6.4.2	ucp_ep_params_flags_field	69
6.6.4.3	ucp_ep_close_flags_t	69
6.6.4.4	ucp_ep_close_mode	69
6.6.4.5	ucp_ep_perf_param_field	70
6.6.4.6	ucp_ep_perf_attr_field	70
6.6.4.7	ucp_cb_param_flags	70
6.6.4.8	ucp_ep_attr_field	71
6.6.4.9	ucp_err_handling_mode_t	71
6.6.5	Function Documentation	71
6.6.5.1	ucp_ep_create()	71
6.6.5.2	ucp_ep_close_nb()	72
6.6.5.3	ucp_ep_close_nbx()	73
6.6.5.4	ucp_ep_print_info()	73
6.6.5.5	ucp_ep_flush_nb()	73
6.6.5.6	ucp_ep_flush_nbx()	74
6.6.5.7	ucp_ep_evaluate_perf()	74
6.6.5.8	ucp_ep_query()	75
6.6.5.9	ucp_request_release()	75
6.6.5.10	ucp_ep_destroy()	75
6.6.5.11	ucp_disconnect_nb()	76
6.6.5.12	ucp_request_test()	76
6.6.5.13	ucp_ep_flush()	76
6.6.5.14	ucp_ep_modify_nb()	76
6.7	UCP Communication routines	78
6.7.1	Detailed Description	82
6.7.2	Data Structure Documentation	82
6.7.2.1	struct ucp_request_attr_t	82
6.7.2.2	struct ucp_err_handler	82
6.7.3	Typedef Documentation	82
6.7.3.1	ucp_tag_t	82
6.7.3.2	ucp_tag_message_h	82
6.7.3.3	ucp_datatype_t	83
6.7.3.4	ucp_send_callback_t	83
6.7.3.5	ucp_send_nbx_callback_t	83
6.7.3.6	ucp_err_handler_cb_t	83
6.7.3.7	ucp_err_handler_t	84
6.7.3.8	ucp_stream_recv_callback_t	84
6.7.3.9	ucp_stream_recv_nbx_callback_t	84
6.7.3.10	ucp_tag_recv_callback_t	84
6.7.3.11	ucp_tag_recv_nbx_callback_t	85
6.7.3.12	ucp_am_recv_data_nbx_callback_t	85
6.7.4	Enumeration Type Documentation	85

6.7.4.1	ucp_atomic_post_op_t	86
6.7.4.2	ucp_atomic_fetch_op_t	86
6.7.4.3	ucp_atomic_op_t	86
6.7.4.4	ucp_stream_recv_flags_t	87
6.7.4.5	ucp_op_attr_t	87
6.7.4.6	ucp_req_attr_field	88
6.7.4.7	ucp_am_recv_attr_t	88
6.7.4.8	ucp_am_handler_param_field	88
6.7.5	Function Documentation	89
6.7.5.1	ucp_request_query()	89
6.7.5.2	ucp_am_send_nb()	89
6.7.5.3	ucp_am_send_nbx()	90
6.7.5.4	ucp_am_recv_data_nbx()	91
6.7.5.5	ucp_am_data_release()	91
6.7.5.6	ucp_stream_send_nb()	92
6.7.5.7	ucp_stream_send_nbx()	93
6.7.5.8	ucp_tag_send_nb()	93
6.7.5.9	ucp_tag_send_nbr()	94
6.7.5.10	ucp_tag_send_sync_nb()	95
6.7.5.11	ucp_tag_send_nbx()	96
6.7.5.12	ucp_tag_send_sync_nbx()	97
6.7.5.13	ucp_stream_recv_nb()	97
6.7.5.14	ucp_stream_recv_nbx()	98
6.7.5.15	ucp_stream_recv_data_nb()	99
6.7.5.16	ucp_tag_recv_nb()	100
6.7.5.17	ucp_tag_recv_nbr()	100
6.7.5.18	ucp_tag_recv_nbx()	101
6.7.5.19	ucp_tag_probe_nb()	102
6.7.5.20	ucp_tag_msg_recv_nb()	103
6.7.5.21	ucp_tag_msg_recv_nbx()	103
6.7.5.22	ucp_put_nbi()	104
6.7.5.23	ucp_put_nb()	105
6.7.5.24	ucp_put_nbx()	105
6.7.5.25	ucp_get_nbi()	106
6.7.5.26	ucp_get_nb()	107
6.7.5.27	ucp_get_nbx()	108
6.7.5.28	ucp_atomic_post()	109
6.7.5.29	ucp_atomic_fetch_nb()	109
6.7.5.30	ucp_atomic_op_nbx()	110
6.7.5.31	ucp_request_check_status()	111
6.7.5.32	ucp_tag_recv_request_test()	112
6.7.5.33	ucp_stream_recv_request_test()	112
6.7.5.34	ucp_request_cancel()	112
6.7.5.35	ucp_stream_data_release()	113
6.7.5.36	ucp_request_free()	113
6.7.5.37	ucp_request_alloc()	113
6.7.5.38	ucp_request_is_completed()	114
6.7.5.39	ucp_put()	114
6.7.5.40	ucp_get()	115
6.7.5.41	ucp_atomic_add32()	115
6.7.5.42	ucp_atomic_add64()	116
6.7.5.43	ucp_atomic_fadd32()	117
6.7.5.44	ucp_atomic_fadd64()	117
6.7.5.45	ucp_atomic_swap32()	118
6.7.5.46	ucp_atomic_swap64()	119
6.7.5.47	ucp_atomic_cswap32()	120
6.7.5.48	ucp_atomic_cswap64()	120
6.8	UCP Configuration	122

6.8.1	Detailed Description	122
6.8.2	Data Structure Documentation	122
6.8.2.1	struct ucp_params	122
6.8.3	Typedef Documentation	124
6.8.3.1	ucp_params_t	124
6.8.3.2	ucp_config_t	124
6.8.4	Function Documentation	125
6.8.4.1	ucp_config_read()	125
6.8.4.2	ucp_config_release()	125
6.8.4.3	ucp_config_modify()	126
6.8.4.4	ucp_config_print()	126
6.9	UCP Data type routines	127
6.9.1	Detailed Description	128
6.9.2	Data Structure Documentation	128
6.9.2.1	struct ucp_dt_iov	128
6.9.3	Macro Definition Documentation	128
6.9.3.1	ucp_dt_make_contig	128
6.9.3.2	ucp_dt_make_iov	129
6.9.4	Typedef Documentation	129
6.9.4.1	ucp_dt_iov_t	129
6.9.4.2	ucp_generic_dt_ops_t	129
6.9.5	Enumeration Type Documentation	129
6.9.5.1	ucp_dt_type	130
6.9.6	Function Documentation	130
6.9.6.1	ucp_dt_create_generic()	130
6.9.6.2	ucp_dt_destroy()	130
6.9.7	Variable Documentation	131
6.9.7.1	start_pack	131
6.9.7.2	start_unpack	131
6.9.7.3	packed_size	131
6.9.7.4	pack	132
6.9.7.5	unpack	132
6.9.7.6	finish	132
6.10	Unified Communication Transport (UCT) API	135
6.10.1	Detailed Description	135
6.11	UCT Communication Resource	136
6.11.1	Detailed Description	140
6.11.2	Data Structure Documentation	140
6.11.2.1	struct uct_md_resource_desc	140
6.11.2.2	struct uct_component_attr	141
6.11.2.3	struct uct_tl_resource_desc	141
6.11.2.4	struct uct_iface_attr	142
6.11.2.5	struct uct_iface_attr.cap	142
6.11.2.6	struct uct_iface_attr.cap.put	143
6.11.2.7	struct uct_iface_attr.cap.get	143
6.11.2.8	struct uct_iface_attr.cap.am	143
6.11.2.9	struct uct_iface_attr.cap.tag	143
6.11.2.10	struct uct_iface_attr.cap.tag.recv	144
6.11.2.11	struct uct_iface_attr.cap.tag.eager	144
6.11.2.12	struct uct_iface_attr.cap.tag.rndv	144
6.11.2.13	struct uct_iface_attr.cap.atomic32	144
6.11.2.14	struct uct_iface_attr.cap.atomic64	144
6.11.2.15	struct uct_iface_params	144
6.11.2.16	union uct_iface_params.mode	145
6.11.2.17	struct uct_iface_params.mode.device	146
6.11.2.18	struct uct_iface_params.mode.sockaddr	146
6.11.2.19	struct uct_ep_params	146
6.11.2.20	struct uct_completion	149



6.11.2.21	struct uct_pending_req	149
6.11.2.22	struct uct_iov	149
6.11.3	Typedef Documentation	150
6.11.3.1	uct_md_resource_desc_t	150
6.11.3.2	uct_component_attr_t	150
6.11.3.3	uct_tl_resource_desc_t	150
6.11.3.4	uct_component_h	150
6.11.3.5	uct_iface_h	151
6.11.3.6	uct_iface_config_t	151
6.11.3.7	uct_md_config_t	151
6.11.3.8	uct_cm_config_t	151
6.11.3.9	uct_ep_h	151
6.11.3.10	uct_mem_h	151
6.11.3.11	uct_rkey_t	151
6.11.3.12	uct_md_h	151
6.11.3.13	uct_md_ops_t	151
6.11.3.14	uct_rkey_ctx_h	152
6.11.3.15	uct_iface_attr_t	152
6.11.3.16	uct_iface_params_t	152
6.11.3.17	uct_ep_attr_t	152
6.11.3.18	uct_md_attr_t	152
6.11.3.19	uct_completion_t	152
6.11.3.20	uct_pending_req_t	152
6.11.3.21	uct_worker_h	152
6.11.3.22	uct_md_t	152
6.11.3.23	uct_am_trace_type_t	153
6.11.3.24	uct_device_addr_t	153
6.11.3.25	uct_iface_addr_t	153
6.11.3.26	uct_ep_addr_t	153
6.11.3.27	uct_ep_params_t	153
6.11.3.28	uct_ep_connect_params_t	153
6.11.3.29	uct_cm_attr_t	153
6.11.3.30	uct_cm_t	153
6.11.3.31	uct_cm_h	153
6.11.3.32	uct_listener_attr_t	154
6.11.3.33	uct_listener_h	154
6.11.3.34	uct_listener_params_t	154
6.11.3.35	uct_tag_context_t	154
6.11.3.36	uct_tag_t	154
6.11.3.37	uct_worker_cb_id_t	154
6.11.3.38	uct_conn_request_h	154
6.11.3.39	uct_iov_t	154
6.11.3.40	uct_completion_callback_t	155
6.11.3.41	uct_pending_callback_t	155
6.11.3.42	uct_error_handler_t	155
6.11.3.43	uct_pending_purge_callback_t	156
6.11.3.44	uct_pack_callback_t	156
6.11.3.45	uct_unpack_callback_t	156
6.11.3.46	uct_async_event_cb_t	156
6.11.4	Enumeration Type Documentation	157
6.11.4.1	uct_component_attr_field	157
6.11.4.2	anonymous enum	157
6.11.4.3	uct_device_type_t	157
6.11.4.4	uct_iface_event_types	158
6.11.4.5	uct_flush_flags	158
6.11.4.6	uct_progress_types	158
6.11.4.7	uct_cb_flags	159
6.11.4.8	uct_iface_open_mode	159

6.11.4.9	uct_iface_params_field	159
6.11.4.10	uct_ep_params_field	160
6.11.4.11	uct_iface_feature	160
6.11.4.12	anonymous enum	161
6.11.4.13	uct_cb_param_flags	161
6.11.5	Function Documentation	161
6.11.5.1	uct_query_components()	162
6.11.5.2	uct_release_component_list()	162
6.11.5.3	uct_component_query()	162
6.11.5.4	uct_md_open()	163
6.11.5.5	uct_md_close()	163
6.11.5.6	uct_md_query_tl_resources()	164
6.11.5.7	uct_release_tl_resource_list()	164
6.11.5.8	uct_md_iface_config_read()	164
6.11.5.9	uct_config_release()	165
6.11.5.10	uct_iface_open()	165
6.11.5.11	uct_iface_close()	166
6.11.5.12	uct_iface_query()	166
6.11.5.13	uct_iface_get_device_address()	166
6.11.5.14	uct_iface_get_address()	167
6.11.5.15	uct_iface_is_reachable()	167
6.11.5.16	uct_ep_check()	168
6.11.5.17	uct_iface_event_fd_get()	168
6.11.5.18	uct_iface_event_arm()	168
6.11.5.19	uct_iface_mem_alloc()	169
6.11.5.20	uct_iface_mem_free()	169
6.11.5.21	uct_ep_create()	170
6.11.5.22	uct_ep_destroy()	170
6.11.5.23	uct_ep_get_address()	171
6.11.5.24	uct_ep_connect_to_ep()	171
6.11.5.25	uct_iface_flush()	171
6.11.5.26	uct_iface_fence()	172
6.11.5.27	uct_ep_pending_add()	172
6.11.5.28	uct_ep_pending_purge()	173
6.11.5.29	uct_ep_flush()	173
6.11.5.30	uct_ep_fence()	173
6.11.5.31	uct_iface_progress_enable()	174
6.11.5.32	uct_iface_progress_disable()	174
6.11.5.33	uct_iface_progress()	175
6.11.5.34	uct_completion_update_status()	175
6.12	UCT Communication Context	176
6.12.1	Detailed Description	176
6.12.2	Enumeration Type Documentation	176
6.12.2.1	uct_alloc_method_t	176
6.12.3	Function Documentation	177
6.12.3.1	uct_worker_create()	177
6.12.3.2	uct_worker_destroy()	177
6.12.3.3	uct_worker_progress_register_safe()	177
6.12.3.4	uct_worker_progress_unregister_safe()	178
6.12.3.5	uct_config_get()	178
6.12.3.6	uct_config_modify()	179
6.12.3.7	uct_worker_progress()	179
6.13	UCT Memory Domain	181
6.13.1	Detailed Description	183
6.13.2	Data Structure Documentation	183
6.13.2.1	struct uct_md_attr	183
6.13.2.2	struct uct_md_attr.cap	183
6.13.2.3	struct uct_md_mem_attr	183

6.13.2.4	struct uct_allocated_memory	184
6.13.2.5	struct uct_rkey_bundle	184
6.13.2.6	struct uct_mem_alloc_params_t	184
6.13.2.7	struct uct_mem_alloc_params_t.mds	185
6.13.3	Typedef Documentation	185
6.13.3.1	uct_md_mem_attr_field_t	185
6.13.3.2	uct_md_mem_attr_t	185
6.13.3.3	uct_allocated_memory_t	185
6.13.3.4	uct_rkey_bundle_t	186
6.13.4	Enumeration Type Documentation	186
6.13.4.1	uct_sockaddr_accessibility_t	186
6.13.4.2	anonymous enum	186
6.13.4.3	uct_md_mem_flags	186
6.13.4.4	uct_mem_advice_t	187
6.13.4.5	uct_md_mem_attr_field	187
6.13.4.6	uct_mem_alloc_params_field_t	188
6.13.5	Function Documentation	188
6.13.5.1	uct_md_mem_query()	188
6.13.5.2	uct_md_query()	188
6.13.5.3	uct_md_mem_advise()	189
6.13.5.4	uct_md_mem_reg()	189
6.13.5.5	uct_md_mem_dereg()	190
6.13.5.6	uct_md_detect_memory_type()	190
6.13.5.7	uct_mem_alloc()	190
6.13.5.8	uct_mem_free()	192
6.13.5.9	uct_md_config_read()	192
6.13.5.10	uct_md_is_sockaddr_accessible()	193
6.13.5.11	uct_md_mkey_pack()	193
6.13.5.12	uct_rkey_unpack()	193
6.13.5.13	uct_rkey_ptr()	194
6.13.5.14	uct_rkey_release()	194
6.14	UCT Active messages	195
6.14.1	Detailed Description	195
6.14.2	Typedef Documentation	195
6.14.2.1	uct_am_callback_t	196
6.14.2.2	uct_am_tracer_t	196
6.14.3	Enumeration Type Documentation	196
6.14.3.1	uct_msg_flags	196
6.14.3.2	uct_am_trace_type	197
6.14.4	Function Documentation	197
6.14.4.1	uct_iface_set_am_handler()	197
6.14.4.2	uct_iface_set_am_tracer()	198
6.14.4.3	uct_iface_release_desc()	198
6.14.4.4	uct_ep_am_short()	198
6.14.4.5	uct_ep_am_short_iov()	199
6.14.4.6	uct_ep_am_bcopy()	199
6.14.4.7	uct_ep_am_zcopy()	199
6.15	UCT Remote memory access operations	201
6.15.1	Detailed Description	201
6.15.2	Function Documentation	201
6.15.2.1	uct_ep_put_short()	201
6.15.2.2	uct_ep_put_bcopy()	201
6.15.2.3	uct_ep_put_zcopy()	201
6.15.2.4	uct_ep_get_short()	202
6.15.2.5	uct_ep_get_bcopy()	202
6.15.2.6	uct_ep_get_zcopy()	202
6.16	UCT Atomic operations	204
6.16.1	Detailed Description	204

6.16.2	Function Documentation	204
6.16.2.1	uct_ep_atomic_cswap64()	204
6.16.2.2	uct_ep_atomic_cswap32()	204
6.16.2.3	uct_ep_atomic32_post()	204
6.16.2.4	uct_ep_atomic64_post()	205
6.16.2.5	uct_ep_atomic32_fetch()	205
6.16.2.6	uct_ep_atomic64_fetch()	205
6.17	UCT Tag matching operations	206
6.17.1	Detailed Description	206
6.17.2	Typedef Documentation	206
6.17.2.1	uct_tag_unexp_eager_cb_t	207
6.17.2.2	uct_tag_unexp_rndv_cb_t	207
6.17.3	Function Documentation	208
6.17.3.1	uct_ep_tag_eager_short()	208
6.17.3.2	uct_ep_tag_eager_bcopy()	209
6.17.3.3	uct_ep_tag_eager_zcopy()	209
6.17.3.4	uct_ep_tag_rndv_zcopy()	210
6.17.3.5	uct_ep_tag_rndv_cancel()	211
6.17.3.6	uct_ep_tag_rndv_request()	211
6.17.3.7	uct_iface_tag_rcv_zcopy()	212
6.17.3.8	uct_iface_tag_rcv_cancel()	212
6.18	UCT client-server operations	214
6.18.1	Detailed Description	216
6.18.2	Data Structure Documentation	217
6.18.2.1	struct uct_ep_connect_params	217
6.18.2.2	struct uct_cm_attr	217
6.18.2.3	struct uct_listener_attr	217
6.18.2.4	struct uct_listener_params	217
6.18.2.5	struct uct_cm_ep_priv_data_pack_args	218
6.18.2.6	struct uct_cm_ep_resolve_args	218
6.18.2.7	struct uct_cm_remote_data	219
6.18.2.8	struct uct_cm_listener_conn_request_args	219
6.18.2.9	struct uct_cm_ep_client_connect_args	219
6.18.2.10	struct uct_cm_ep_server_conn_notify_args	220
6.18.3	Typedef Documentation	220
6.18.3.1	uct_cm_ep_priv_data_pack_args_t	220
6.18.3.2	uct_cm_ep_resolve_args_t	220
6.18.3.3	uct_cm_remote_data_t	220
6.18.3.4	uct_cm_listener_conn_request_args_t	220
6.18.3.5	uct_cm_ep_client_connect_args_t	221
6.18.3.6	uct_cm_ep_server_conn_notify_args_t	221
6.18.3.7	uct_sockaddr_conn_request_callback_t	221
6.18.3.8	uct_cm_listener_conn_request_callback_t	221
6.18.3.9	uct_cm_ep_server_conn_notify_callback_t	222
6.18.3.10	uct_cm_ep_client_connect_callback_t	222
6.18.3.11	uct_ep_disconnect_cb_t	222
6.18.3.12	uct_cm_ep_priv_data_pack_callback_t	223
6.18.3.13	uct_cm_ep_resolve_callback_t	223
6.18.4	Enumeration Type Documentation	223
6.18.4.1	uct_cm_attr_field	223
6.18.4.2	uct_listener_attr_field	224
6.18.4.3	uct_listener_params_field	224
6.18.4.4	uct_ep_connect_params_field	224
6.18.4.5	uct_cm_ep_priv_data_pack_args_field	224
6.18.4.6	uct_cm_ep_resolve_args_field	225
6.18.4.7	uct_cm_remote_data_field	225
6.18.4.8	uct_cm_listener_conn_request_args_field	225
6.18.4.9	uct_cm_ep_client_connect_args_field	226

6.18.4.10	uct_cm_ep_server_conn_notify_args_field	226
6.18.5	Function Documentation	227
6.18.5.1	uct_iface_accept()	227
6.18.5.2	uct_iface_reject()	227
6.18.5.3	uct_ep_connect()	227
6.18.5.4	uct_ep_disconnect()	228
6.18.5.5	uct_cm_open()	228
6.18.5.6	uct_cm_close()	229
6.18.5.7	uct_cm_query()	229
6.18.5.8	uct_cm_config_read()	229
6.18.5.9	uct_cm_client_ep_conn_notify()	230
6.18.5.10	uct_listener_create()	230
6.18.5.11	uct_listener_destroy()	231
6.18.5.12	uct_listener_reject()	231
6.18.5.13	uct_listener_query()	231
6.19	UCT interface operations and capabilities	232
6.19.1	Detailed Description	232
6.19.2	Macro Definition Documentation	232
6.19.2.1	UCT_IFACE_FLAG_AM_SHORT	233
6.19.2.2	UCT_IFACE_FLAG_AM_BCOPY	233
6.19.2.3	UCT_IFACE_FLAG_AM_ZCOPY	233
6.19.2.4	UCT_IFACE_FLAG_PENDING	233
6.19.2.5	UCT_IFACE_FLAG_PUT_SHORT	233
6.19.2.6	UCT_IFACE_FLAG_PUT_BCOPY	233
6.19.2.7	UCT_IFACE_FLAG_PUT_ZCOPY	233
6.19.2.8	UCT_IFACE_FLAG_GET_SHORT	234
6.19.2.9	UCT_IFACE_FLAG_GET_BCOPY	234
6.19.2.10	UCT_IFACE_FLAG_GET_ZCOPY	234
6.19.2.11	UCT_IFACE_FLAG_ATOMIC_CPU	234
6.19.2.12	UCT_IFACE_FLAG_ATOMIC_DEVICE	234
6.19.2.13	UCT_IFACE_FLAG_ERRHANDLE_SHORT_BUF	234
6.19.2.14	UCT_IFACE_FLAG_ERRHANDLE_BCOPY_BUF	234
6.19.2.15	UCT_IFACE_FLAG_ERRHANDLE_ZCOPY_BUF	234
6.19.2.16	UCT_IFACE_FLAG_ERRHANDLE_AM_ID	234
6.19.2.17	UCT_IFACE_FLAG_ERRHANDLE_REMOTE_MEM	235
6.19.2.18	UCT_IFACE_FLAG_ERRHANDLE_BCOPY_LEN	235
6.19.2.19	UCT_IFACE_FLAG_ERRHANDLE_PEER_FAILURE	235
6.19.2.20	UCT_IFACE_FLAG_EP_CHECK	235
6.19.2.21	UCT_IFACE_FLAG_CONNECT_TO_IFACE	235
6.19.2.22	UCT_IFACE_FLAG_CONNECT_TO_EP	235
6.19.2.23	UCT_IFACE_FLAG_CONNECT_TO_SOCKADDR	235
6.19.2.24	UCT_IFACE_FLAG_AM_DUP	236
6.19.2.25	UCT_IFACE_FLAG_CB_SYNC	236
6.19.2.26	UCT_IFACE_FLAG_CB_ASYNC	236
6.19.2.27	UCT_IFACE_FLAG_EP_KEEPALIVE	236
6.19.2.28	UCT_IFACE_FLAG_TAG_EAGER_SHORT	236
6.19.2.29	UCT_IFACE_FLAG_TAG_EAGER_BCOPY	236
6.19.2.30	UCT_IFACE_FLAG_TAG_EAGER_ZCOPY	236
6.19.2.31	UCT_IFACE_FLAG_TAG_RNDV_ZCOPY	236
6.20	UCT interface for asynchronous event capabilities	237
6.20.1	Detailed Description	237
6.20.2	Macro Definition Documentation	237
6.20.2.1	UCT_IFACE_FLAG_EVENT_SEND_COMP	237
6.20.2.2	UCT_IFACE_FLAG_EVENT_RECV	237
6.20.2.3	UCT_IFACE_FLAG_EVENT_RECV_SIG	237
6.20.2.4	UCT_IFACE_FLAG_EVENT_FD	237
6.20.2.5	UCT_IFACE_FLAG_EVENT_ASYNC_CB	237
6.21	Unified Communication Services (UCS) API	238

6.21.1	Detailed Description	238
6.22	UCS Communication Resource	239
6.22.1	Detailed Description	240
6.22.2	Data Structure Documentation	240
6.22.2.1	struct ucs_sock_addr	240
6.22.3	Typedef Documentation	240
6.22.3.1	ucs_async_event_cb_t	240
6.22.3.2	ucs_sock_addr_t	240
6.22.3.3	ucs_memory_type_t	241
6.22.3.4	ucs_time_t	241
6.22.3.5	ucs_status_ptr_t	241
6.22.4	Enumeration Type Documentation	241
6.22.4.1	ucs_callbackq_flags	241
6.22.4.2	ucs_memory_type	241
6.22.4.3	ucs_status_t	242
6.22.4.4	ucs_thread_mode_t	243
6.22.5	Function Documentation	243
6.22.5.1	ucs_async_set_event_handler()	243
6.22.5.2	ucs_async_add_timer()	244
6.22.5.3	ucs_async_remove_handler()	244
6.22.5.4	ucs_async_modify_handler()	245
6.22.5.5	ucs_async_context_create()	245
6.22.5.6	ucs_async_context_destroy()	245
6.22.5.7	ucs_async_poll()	246
<b>7</b>	<b>Data Structure Documentation</b>	<b>247</b>
7.1	ucp_generic_dt_ops Struct Reference	247
7.1.1	Detailed Description	247
7.2	uct_tag_context Struct Reference	247
7.2.1	Detailed Description	248
7.2.2	Field Documentation	248
7.2.2.1	tag_consumed_cb	248
7.2.2.2	completed_cb	248
7.2.2.3	rndv_cb	249
7.2.2.4	priv	249
<b>8</b>	<b>Example Documentation</b>	<b>251</b>
8.1	ucp_client_server.c	251
8.2	ucp_hello_world.c	261
8.3	uct_hello_world.c	268
<b>Index</b>		<b>277</b>

# Chapter 1

## Preface

### 1.1 Scope of the Document

This document describes the UCX programming interface. The programming interface exposes a high performance communication API, which provides basic building blocks for PGAS, Message Passing Interface (MPI), Big-Data, Analytics, File I/O, and storage library developers.

### 1.2 Audience

This manual is intended for programmers who want to develop parallel programming models like OpenSHMEM, MPI, UPC, Chapel, etc. The manual assumes that the reader is familiar with the following:

- Basic concepts of two-sided, one-sided, atomic, and collective operations
- C programming language

### 1.3 Document Status

This section briefly describes a list of open issues in the UCX specification.

- UCP API - work in progress
- UCT API - work in progress

### 1.4 License

UCX project follows open source development model and the software is licensed under BSD-3 license.





## Chapter 2

# Introduction

### 2.1 Motivation

A communication middleware abstracts the vendor-specific software and hardware interfaces. They bridge the semantic and functionality gap between the programming models and the software and hardware network interfaces by providing data transfer interfaces and implementation, optimized protocols for data transfer between various memories, and managing network resources. There are many communication middleware APIs and libraries to support parallel programming models such as MPI, OpenSHMEM, and task-based models.

Current communication middleware designs typically take two approaches. First, communication middleware such as Intel's PSM (previously Qlogic), Mellanox's MXM, and IBM's PAMI provide high-performance implementations for specific network hardware. Second, communication middleware such as VMI, Cactus, ARMCI, GASNet, and Open MPI are tightly coupled to a specific programming model. Communication middleware designed with either of this design approach requires significant porting effort to move a new network interface or programming model.

To achieve functional and performance portability across architectures and programming models, we introduce Unified Communication X (UCX).

### 2.2 UCX

Unified Communication X (UCX) is a set of network APIs and their implementations for high throughput computing. UCX is a combined effort of national laboratories, industry, and academia to design and implement a high-performing and highly-scalable network stack for next generation applications and systems. UCX design provides the ability to tailor its APIs and network functionality to suit a wide variety of application domains. We envision that these APIs will satisfy the networking needs of many programming models such as the Message Passing Interface (MPI), OpenSHMEM, Partitioned Global Address Space (PGAS) languages, task-based paradigms, and I/O bound applications.

The initial focus is on supporting semantics such as point-to-point communications (one-sided and two-sided), collective communication, and remote atomic operations required for popular parallel programming models. Also, the initial UCX reference implementation is targeted to support current network technologies such as:

- Open Fabrics - InfiniBand (Mellanox, Qlogic, IBM), iWARP, RoCE
- Cray uGNI - GEMINI and ARIES interconnects
- Shared memory (MMAP, Posix, CMA, KNEM, XPMEM, etc.)
- Ethernet (TCP/UDP)

UCX design goals are focused on performance and scalability, while efficiently supporting popular and emerging programming models.

UCX's API and design do not impose architectural constraints on the network hardware nor require any specific capabilities to support the programming model functionality. This is achieved by keeping the API flexible and ability to support the missing functionality efficiently in the software.

Extreme scalability is an important design goal for UCX. To achieve this, UCX follows these design principles:

- Minimal memory consumption : Design avoids data-structures that scale with the number of processing elements (i.e., order N data structures), and share resources among multiple programming models.
- Low-latency Interfaces: Design provides at least two sets of APIs with one set focused on the performance, and the other focused on functionality.
- High bandwidth - With minimal software overhead combined and support for multi-rail and multi-device capabilities, the design provides all the hooks that are necessary for exploiting hardware bandwidth capabilities.
- Asynchronous Progress: API provides non-blocking communication interfaces and design supports asynchronous progress required for communication and computation overlap
- Resilience - the API exposes communication control hooks required for fault tolerant communication library implementation.

UCX design provides native support for hybrid programming models. The design enables resource sharing, optimal memory usage, and progress engine coordination to efficiently implement hybrid programming models. For example, hybrid applications that use both OpenSHMEM and MPI programming models will be able to select between a single-shared UCX network context or a stand alone UCX network context for each one of them. Such flexibility, optimized resource sharing, and reduced memory consumption, improve network and application performance.

# Chapter 3

## Design

The UCX framework consists of the three main components: UC-Services (UCS), UC-Transports (UCT), and UC-Protocols (UCP). Each one of these components exports a public API, and can be used as a stand-alone library.

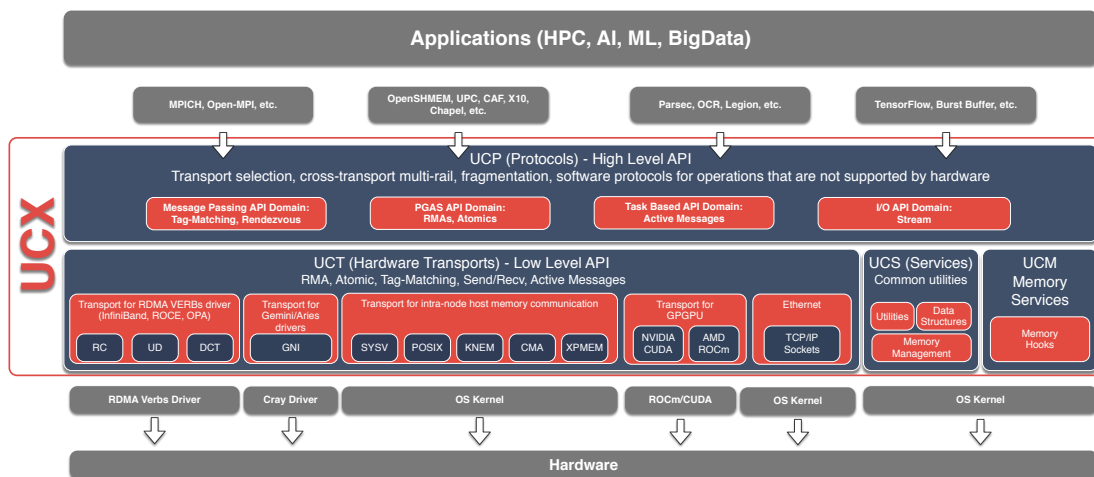


Figure 3.1: UCX Framework Architecture

### 3.1 UCS

UCS is a service layer that provides the necessary functionality for implementing portable and efficient utilities. This layer includes the following services:

- an abstraction for accessing platform specific functionality (atomic operations, thread safety, etc.),
- tools for efficient memory management (memory pools, memory allocators, and memory allocators hooks),
- commonly used data structures (hashes, trees, lists).

### 3.2 UCT

UCT is a transport layer that abstracts the differences across various hardware architectures and provides a low-level API that enables the implementation of communication protocols. The primary goal of the layer is to provide

direct and efficient access to hardware network functionality. For this purpose, UCT relies on vendor provided low-level drivers such as uGNI, Verbs, shared memory, ROCM, CUDA. In addition, the layer provides constructs for communication context management (thread-based and application level), and allocation and management of device-specific memories including those found in accelerators. In terms of communication APIs, UCT defines interfaces for immediate (short), buffered copy-and-send (bcopy), and zero-copy (zcopy) communication operations.

**Short:** This type of operation is optimized for small messages that can be posted and completed in place.

**Bcopy:** This type of operation is optimized for medium size messages that are typically sent through a so-called bouncing-buffer. This auxiliary buffer is typically allocated given network constraints and ready for immediate utilization by the hardware. Since a custom data packing routine could be provided, this method can be used for non-contiguous i/o.

**Zcopy:** This type of operation exposes zero-copy memory-to-memory communication semantics, which means that message is sent directly from user buffer, or received directly to user buffer, without being copied between the network layers.

### 3.3 UCP

UCP implements higher-level protocols that are typically used by message passing (MPI) and PGAS programming models by using lower-level capabilities exposed through the UCT layer. UCP provides the following functionality: ability to select different transports for communication, message fragmentation, multi-rail communication, and initializing and finalizing the library. Currently, the API has the following classes of interfaces: Initialization, Remote Memory Access (RMA) communication, Atomic Memory Operations (AMO), Active Message, Tag-Matching, and Collectives.

**Initialization:** This subset of interfaces defines the communication context setup, queries the network capabilities, and initializes the local communication endpoints. The context represented by the UCX context is an abstraction of the network transport resources. The communication endpoint setup interfaces initialize the UCP endpoint, which is an abstraction of all the necessary resources associated with a particular connection. The communication endpoints are used as input to all communication operations to describe the source and destination of the communication.

**RMA:** This subset of interfaces defines one-sided communication operations such as PUT and GET, required for implementing low overhead, direct memory access communications constructs needed by both distributed and shared memory programming models. UCP includes a separate set of interfaces for communicating non-contiguous data. This functionality was included to support various programming models' communication requirements and leverage the scatter/gather capabilities of modern network hardware.

**AMO:** This subset of interfaces provides support for atomically performing operations on the remote memory, an important class of operations for PGAS programming models, particularly OpenSHMEM.

**Tag Matching:** This interface supports tag-matching for send-receive semantics which is a key communication semantic defined by the MPI specification.

**Stream :** The API provides order and reliable communication semantics. Data is treated as an ordered sequence of bytes pushed through the connection. In contrast of tag-matching interface, the size of each individual send does not necessarily have to match the size of each individual receive, as long as the total number of bytes is the same. This API is designed to match widely used BSD-socket based programming models.

**Active Message:** A subset of functionality where the incoming packet invokes a sender-specified callback in order to be processed by the receiving process. As an example, the two-sided MPI interface can easily be implemented on top of such a concept (TBD: cite openmpi ). However, these interfaces are more general and suited for other programming paradigms where the receiver process does not prepost receives, but expects to react to incoming packets directly. Like RMA and tag-matching interfaces, the active message interface provides separate APIs for different message types and non-contiguous data.

**Collectives:** This subset of interfaces defines group communication and synchronization operations. The collective operations include barrier, all-to-one, all-to-all, and reduction operations. When possible, we will take advantage of hardware acceleration for collectives (e.g., InfiniBand Switch collective acceleration).

## Chapter 4

# Conventions and Notations

This section describes the conventions and notations in the UCX specification.

### 4.1 Blocking Behavior

The blocking UCX routines return only when a UCX operation is complete. After the return, the resources used in the UCX routine are available for reuse.

### 4.2 Non-blocking Behavior

The non-blocking UCX routines return immediately, independent of operation completion. After the return, the resources used for the routines are not necessarily available for reuse.

### 4.3 Fairness

UCX routines do not guarantee fairness. However, the routines enable UCX consumers to write efficient and fair programs.

### 4.4 Interaction with Signal Handler Functions

If UCX routines are invoked from a signal handler function, the behavior of the program is undefined.



## Chapter 5

# Deprecated List

Global **ucp\_atomic\_add32** (ucp\_ep\_h ep, uint32\_t add, uint64\_t remote\_addr, ucp\_rkey\_h rkey)

Replaced by [ucp\\_atomic\\_post](#) with opcode UCP\_ATOMIC\_POST\_OP\_ADD.

Global **ucp\_atomic\_add64** (ucp\_ep\_h ep, uint64\_t add, uint64\_t remote\_addr, ucp\_rkey\_h rkey)

Replaced by [ucp\\_atomic\\_post](#) with opcode UCP\_ATOMIC\_POST\_OP\_ADD.

Global **ucp\_atomic\_cswap32** (ucp\_ep\_h ep, uint32\_t compare, uint32\_t swap, uint64\_t remote\_addr, ucp\_rkey\_h rkey, uint32\_t \*result)

Replaced by [ucp\\_atomic\\_fetch\\_nb](#) with opcode UCP\_ATOMIC\_FETCH\_OP\_CSWAP.

Global **ucp\_atomic\_cswap64** (ucp\_ep\_h ep, uint64\_t compare, uint64\_t swap, uint64\_t remote\_addr, ucp\_rkey\_h rkey, uint64\_t \*result)

Replaced by [ucp\\_atomic\\_fetch\\_nb](#) with opcode UCP\_ATOMIC\_FETCH\_OP\_CSWAP.

Global **ucp\_atomic\_fadd32** (ucp\_ep\_h ep, uint32\_t add, uint64\_t remote\_addr, ucp\_rkey\_h rkey, uint32\_t \*result)

Replaced by [ucp\\_atomic\\_fetch\\_nb](#) with opcode UCP\_ATOMIC\_FETCH\_OP\_FADD.

Global **ucp\_atomic\_fadd64** (ucp\_ep\_h ep, uint64\_t add, uint64\_t remote\_addr, ucp\_rkey\_h rkey, uint64\_t \*result)

Replaced by [ucp\\_atomic\\_fetch\\_nb](#) with opcode UCP\_ATOMIC\_FETCH\_OP\_FADD.

Global **ucp\_atomic\_swap32** (ucp\_ep\_h ep, uint32\_t swap, uint64\_t remote\_addr, ucp\_rkey\_h rkey, uint32\_t \*result)

Replaced by [ucp\\_atomic\\_fetch\\_nb](#) with opcode UCP\_ATOMIC\_FETCH\_OP\_SWAP.

Global **ucp\_atomic\_swap64** (ucp\_ep\_h ep, uint64\_t swap, uint64\_t remote\_addr, ucp\_rkey\_h rkey, uint64\_t \*result)

Replaced by [ucp\\_atomic\\_fetch\\_nb](#) with opcode UCP\_ATOMIC\_FETCH\_OP\_SWAP.

Global **ucp\_disconnect\_nb** (ucp\_ep\_h ep)

Replaced by [ucp\\_ep\\_close\\_nb](#).

Global **ucp\_ep\_destroy** (ucp\_ep\_h ep)

Replaced by [ucp\\_ep\\_close\\_nb](#).

Global **ucp\_ep\_flush** (ucp\_ep\_h ep)

Replaced by [ucp\\_ep\\_flush\\_nb](#).

Global **ucp\_ep\_modify\_nb** (ucp\_ep\_h ep, const ucp\_ep\_params\_t \*params)

Use [ucp\\_listener\\_conn\\_handler\\_t](#) instead of [ucp\\_listener\\_accept\\_handler\\_t](#), if you have other use case please submit an issue on <https://github.com/openucx/ucx> or report to [ucx-group@elist.ornl.gov](mailto:ucx-group@elist.ornl.gov)

Global **ucp\_get** (ucp\_ep\_h ep, void \*buffer, size\_t length, uint64\_t remote\_addr, ucp\_rkey\_h rkey)

Replaced by [ucp\\_get\\_nb](#).

**Global [ucp\\_listener\\_accept\\_handler\\_t](#)**

Replaced by [ucp\\_listener\\_conn\\_handler\\_t](#).

**Global [ucp\\_listener\\_accept\\_handler\\_t](#)**

Replaced by [ucp\\_listener\\_conn\\_handler\\_t](#).

**Global [ucp\\_put](#) ([ucp\\_ep\\_h](#) ep, [const void \\*buffer](#), [size\\_t length](#), [uint64\\_t remote\\_addr](#), [ucp\\_rkey\\_h](#) rkey)**

Replaced by [ucp\\_put\\_nb](#). The following example implements the same functionality using [ucp\\_put\\_nb](#) :

**Global [ucp\\_request\\_is\\_completed](#) ([void \\*request](#))**

Replaced by [ucp\\_request\\_test](#).

**Global [ucp\\_request\\_release](#) ([void \\*request](#))**

Replaced by [ucp\\_request\\_free](#).

**Global [ucp\\_request\\_test](#) ([void \\*request](#), [ucp\\_tag\\_rcv\\_info\\_t \\*info](#))**

Replaced by [ucp\\_tag\\_rcv\\_request\\_test](#) and [ucp\\_request\\_check\\_status](#) depends on use case.

**Global [ucp\\_worker\\_flush](#) ([ucp\\_worker\\_h](#) worker)**

Replaced by [ucp\\_worker\\_flush\\_nb](#). The following example implements the same functionality using [ucp\\_worker\\_flush\\_nb](#) :



## Chapter 6

# Module Documentation

### 6.1 Unified Communication Protocol (UCP) API

This section describes UCP API.

#### Modules

- [UCP Application Context](#)
- [UCP Worker](#)
- [UCP Memory routines](#)
- [UCP Wake-up routines](#)
- [UCP Endpoint](#)
- [UCP Communication routines](#)
- [UCP Configuration](#)
- [UCP Data type routines](#)

#### 6.1.1 Detailed Description

## 6.2 UCP Application Context

### Data Structures

- struct `ucp_lib_attr`  
*Lib attributes. [More...](#)*
- struct `ucp_context_attr`  
*Context attributes. [More...](#)*
- struct `ucp_tag_recv_info`  
*UCP receive information descriptor. [More...](#)*
- struct `ucp_request_param_t`  
*Operation parameters passed to `ucp_tag_send_nbx`, `ucp_tag_send_sync_nbx`, `ucp_tag_recv_nbx`, `ucp_put_nbx`, `ucp_get_nbx`, `ucp_am_send_nbx` and `ucp_am_recv_data_nbx`. [More...](#)*
- union `ucp_request_param_t.cb`
- union `ucp_request_param_t.recv_info`

### Macros

- `#define UCP_ENTITY_NAME_MAX 32`  
*Maximum size of the UCP entity name in structure of entity attributes provided by a query method.*

### Typedefs

- typedef struct `ucp_lib_attr` `ucp_lib_attr_t`  
*Lib attributes.*
- typedef struct `ucp_context_attr` `ucp_context_attr_t`  
*Context attributes.*
- typedef struct `ucp_tag_recv_info` `ucp_tag_recv_info_t`  
*UCP receive information descriptor.*
- typedef struct `ucp_context` \* `ucp_context_h`  
*UCP Application Context.*
- typedef void(\* `ucp_request_init_callback_t`) (void \*request)  
*Request initialization callback.*
- typedef void(\* `ucp_request_cleanup_callback_t`) (void \*request)  
*Request cleanup callback.*

### Enumerations

- enum `ucp_params_field` {  
`UCP_PARAM_FIELD_FEATURES` = UCS\_BIT(0), `UCP_PARAM_FIELD_REQUEST_SIZE` = UCS\_BIT(1),  
`UCP_PARAM_FIELD_REQUEST_INIT` = UCS\_BIT(2), `UCP_PARAM_FIELD_REQUEST_CLEANUP` = UCS\_BIT(3),  
`UCP_PARAM_FIELD_TAG_SENDER_MASK` = UCS\_BIT(4), `UCP_PARAM_FIELD_MT_WORKERS_SHARED`  
= UCS\_BIT(5), `UCP_PARAM_FIELD_ESTIMATED_NUM_EPS` = UCS\_BIT(6), `UCP_PARAM_FIELD_ESTIMATED_NUM_PPNS`  
= UCS\_BIT(7),  
`UCP_PARAM_FIELD_NAME` = UCS\_BIT(8) }  
*UCP context parameters field mask.*
- enum `ucp_feature` {  
`UCP_FEATURE_TAG` = UCS\_BIT(0), `UCP_FEATURE_RMA` = UCS\_BIT(1), `UCP_FEATURE_AMO32` =  
UCS\_BIT(2), `UCP_FEATURE_AMO64` = UCS\_BIT(3),  
`UCP_FEATURE_WAKEUP` = UCS\_BIT(4), `UCP_FEATURE_STREAM` = UCS\_BIT(5), `UCP_FEATURE_AM`  
= UCS\_BIT(6) }

*UCP configuration features.*

- enum `ucp_lib_attr_field` { `UCP_LIB_ATTR_FIELD_MAX_THREAD_LEVEL` = `UCS_BIT(0)` }

*UCP library attributes field mask.*

- enum `ucp_context_attr_field` { `UCP_ATTR_FIELD_REQUEST_SIZE` = `UCS_BIT(0)`, `UCP_ATTR_FIELD_THREAD_MODE` = `UCS_BIT(1)`, `UCP_ATTR_FIELD_MEMORY_TYPES` = `UCS_BIT(2)`, `UCP_ATTR_FIELD_NAME` = `UCS_BIT(3)` }

*UCP context attributes field mask.*

## Functions

- `ucs_status_t ucp_lib_query` (`ucp_lib_attr_t *attr`)  
*Get attributes of the UCP library.*
- void `ucp_get_version` (`unsigned *major_version`, `unsigned *minor_version`, `unsigned *release_number`)  
*Get UCP library version.*
- const char \* `ucp_get_version_string` (`void`)  
*Get UCP library version as a string.*
- static `ucs_status_t ucp_init` (`const ucp_params_t *params`, `const ucp_config_t *config`, `ucp_context_h *context_p`)  
*UCP context initialization.*
- void `ucp_cleanup` (`ucp_context_h context_p`)  
*Release UCP application context.*
- `ucs_status_t ucp_context_query` (`ucp_context_h context_p`, `ucp_context_attr_t *attr`)  
*Get attributes specific to a particular context.*
- void `ucp_context_print_info` (`const ucp_context_h context`, `FILE *stream`)  
*Print context information.*

### 6.2.1 Detailed Description

Application context is a primary concept of UCP design which provides an isolation mechanism, allowing resources associated with the context to separate or share network communication context across multiple instances of applications.

This section provides a detailed description of this concept and routines associated with it.

### 6.2.2 Data Structure Documentation

#### 6.2.2.1 struct ucp\_lib\_attr

The structure defines the attributes that characterize the Library.

##### Data Fields

<code>uint64_t</code>	<code>field_mask</code>	Mask of valid fields in this structure, using bits from <code>ucp_lib_attr_field</code> . Fields not specified in this mask will be ignored. Provides ABI compatibility with respect to adding new fields.
<code>ucs_thread_mode_t</code>	<code>max_thread_level</code>	Maximum level of thread support of the library, which is permanent throughout the lifetime of the library. Accordingly, the user can call <code>ucp_worker_create</code> with appropriate <code>ucp_worker_params_t::thread_mode</code> . For supported thread levels please see <code>ucs_thread_mode_t</code> .

### 6.2.2.2 struct ucp\_context\_attr

The structure defines the attributes that characterize the particular context.

#### Data Fields

uint64_t	field_mask	Mask of valid fields in this structure, using bits from <a href="#">ucp_context_attr_field</a> . Fields not specified in this mask will be ignored. Provides ABI compatibility with respect to adding new fields.
size_t	request_size	Size of UCP non-blocking request. When pre-allocated request is used (e.g. in <a href="#">ucp_tag_rcv_nbr</a> ) it should have enough space to fit UCP request data, which is defined by this value.
ucs_thread_mode_t	thread_mode	Thread safe level of the context. For supported thread levels please see <a href="#">ucs_thread_mode_t</a> .
uint64_t	memory_types	Mask of which memory types are supported, for supported memory types please see <a href="#">ucs_memory_type_t</a> .
char	name[UCP_ENTITY_NAME_MAX]	Tracing and analysis tools can use name to identify this UCX context.

### 6.2.2.3 struct ucp\_tag\_rcv\_info

The UCP receive information descriptor is allocated by application and filled in with the information about the received message by [ucp\\_tag\\_probe\\_nb](#) or [ucp\\_tag\\_rcv\\_request\\_test](#) routines or [ucp\\_tag\\_rcv\\_callback\\_t](#) callback argument.

#### Examples

[ucp\\_client\\_server.c](#), and [ucp\\_hello\\_world.c](#).

#### Data Fields

<a href="#">ucp_tag_t</a>	sender_tag	Sender tag
size_t	length	The size of the received data

### 6.2.2.4 struct ucp\_request\_param\_t

The structure [ucp\\_request\\_param\\_t](#) is used to specify datatype of operation, provide user request in case the external request is used, set completion callback and custom user data passed to this callback.

Example: implementation of function to send contiguous buffer to ep and invoke callback function at operation completion. If the operation completed immediately (status == UCS\_OK) then callback is not called.

```
ucs_status_ptr_t send_data(ucp_ep_h ep, void *buffer, size_t length,
                          ucp_tag_t tag, void *request)
{
    ucp_request_param_t param = {
        .op_attr_mask      = UCP_OP_ATTR_FIELD_CALLBACK |
                             UCP_OP_ATTR_FIELD_REQUEST,
        .request           = request,
        .cb.send           = custom_send_callback_f,
        .user_data         = pointer_to_user_context_passed_to_cb
    };
    ucs_status_ptr_t status;
    status = ucp_tag_send_nbx(ep, buffer, length, tag, &param);
    if (UCS_PTR_IS_ERR(status)) {
```

```

    handle_error(status);
} else if (status == UCS_OK) {
    // operation is completed
}
return status;
}

```

### Examples

[ucp\\_client\\_server.c](#), and [ucp\\_hello\\_world.c](#).

### Data Fields

uint32_t	op_attr_mask	Mask of valid fields in this structure and operation flags, using bits from <a href="#">ucp_op_attr_t</a> . Fields not specified in this mask will be ignored. Provides ABI compatibility with respect to adding new fields.
uint32_t	flags	
void *	request	Request handle allocated by the user. There should be at least UCP request size bytes of available space before the <i>request</i> . The size of the UCP request can be obtained by <a href="#">ucp_context_query</a> function.
union <a href="#">ucp_request_param_t</a>	cb	Callback function that is invoked whenever the send or receive operation is completed.
<a href="#">ucp_datatype_t</a>	datatype	Datatype descriptor for the elements in the buffer. In case the op_attr_mask & UCP_OP_ATTR_FIELD_DATATYPE bit is not set, then use default datatype <a href="#">ucp_dt_make_contig(1)</a>
void *	user_data	Pointer to user data passed to callback function.
void *	reply_buffer	Reply buffer. Can be used for storing operation result, for example by <a href="#">ucp_atomic_op_nbx</a> .
<a href="#">ucs_memory_type_t</a>	memory_type	Memory type of the buffer. see <a href="#">ucs_memory_type_t</a> for possible memory types. An optimization hint to avoid memory type detection for request buffer. If this value is not set (along with its corresponding bit in the op_attr_mask - UCP_OP_ATTR_FIELD_MEMORY_TYPE), then use default <a href="#">UCS_MEMORY_TYPE_UNKNOWN</a> which means the memory type will be detected internally.
union <a href="#">ucp_request_param_t</a>	recv_info	Pointer to the information where received data details are stored in case of an immediate completion of receive operation. The user has to provide a pointer to valid memory/variable which will be updated on function return.
<a href="#">ucp_mem_h</a>	memh	Memory handle for pre-registered buffer. If the handle is provided, protocols that require registered memory can skip the registration step. As a result, the communication request overhead can be reduced and the request can be completed faster. The memory handle should be obtained by calling <a href="#">ucp_mem_map</a> .

#### 6.2.2.5 union [ucp\\_request\\_param\\_t.cb](#)

Callback function that is invoked whenever the send or receive operation is completed.

### Data Fields

<a href="#">ucp_send_nbx_callback_t</a>	send	
<a href="#">ucp_tag_recv_nbx_callback_t</a>	recv	

## Data Fields

<a href="#">ucp_stream_rcv_nbx_callback_t</a>	recv_stream	
<a href="#">ucp_am_rcv_data_nbx_callback_t</a>	recv_am	

**6.2.2.6 union ucp\_request\_param\_t.recv\_info**

Pointer to the information where received data details are stored in case of an immediate completion of receive operation. The user has to provide a pointer to valid memory/variable which will be updated on function return.

## Data Fields

<a href="#">size_t*</a>	length	
<a href="#">ucp_tag_rcv_info_t*</a>	tag_info	

**6.2.3 Macro Definition Documentation****6.2.3.1 UCP\_ENTITY\_NAME\_MAX**

```
#define UCP_ENTITY_NAME_MAX 32
```

**6.2.4 Typedef Documentation****6.2.4.1 ucp\_lib\_attr\_t**

```
typedef struct ucp_lib_attr ucp_lib_attr_t
```

The structure defines the attributes that characterize the Library.

**6.2.4.2 ucp\_context\_attr\_t**

```
typedef struct ucp_context_attr ucp_context_attr_t
```

The structure defines the attributes that characterize the particular context.

**6.2.4.3 ucp\_tag\_rcv\_info\_t**

```
typedef struct ucp_tag_rcv_info ucp_tag_rcv_info_t
```

The UCP receive information descriptor is allocated by application and filled in with the information about the received message by [ucp\\_tag\\_probe\\_nb](#) or [ucp\\_tag\\_rcv\\_request\\_test](#) routines or [ucp\\_tag\\_rcv\\_callback\\_t](#) callback argument.

**6.2.4.4 ucp\_context\_h**

```
typedef struct ucp_context* ucp_context_h
```

UCP application context (or just a context) is an opaque handle that holds a UCP communication instance's global information. It represents a single UCP communication instance. The communication instance could be an OS process (an application) that uses UCP library. This global information includes communication resources, endpoints, memory, temporary file storage, and other communication information directly associated with a specific UCP instance. The context also acts as an isolation mechanism, allowing resources associated with the context to manage multiple concurrent communication instances. For example, users using both MPI and OpenSHMEM sessions simultaneously can isolate their communication by allocating and using separate contexts for each of them. Alternatively, users can share the communication resources (memory, network resource context, etc.) between them by using the same application context. A message sent or a RMA operation performed in one application context cannot be received in any other application context.

#### 6.2.4.5 ucp\_request\_init\_callback\_t

```
typedef void(* ucp_request_init_callback_t) (void *request)
```

This callback routine is responsible for the request initialization.

##### Parameters

in	<i>request</i>	Request handle to initialize.
----	----------------	-------------------------------

#### 6.2.4.6 ucp\_request\_cleanup\_callback\_t

```
typedef void(* ucp_request_cleanup_callback_t) (void *request)
```

This callback routine is responsible for cleanup of the memory associated with the request.

##### Parameters

in	<i>request</i>	Request handle to cleanup.
----	----------------	----------------------------

## 6.2.5 Enumeration Type Documentation

### 6.2.5.1 ucp\_params\_field

```
enum ucp_params_field
```

The enumeration allows specifying which fields in `ucp_params_t` are present. It is used to enable backward compatibility support.

##### Enumerator

UCP_PARAM_FIELD_FEATURES	features
UCP_PARAM_FIELD_REQUEST_SIZE	request_size
UCP_PARAM_FIELD_REQUEST_INIT	request_init
UCP_PARAM_FIELD_REQUEST_CLEANUP	request_cleanup
UCP_PARAM_FIELD_TAG_SENDER_MASK	tag_sender_mask
UCP_PARAM_FIELD_MT_WORKERS_SHARED	mt_workers_shared
UCP_PARAM_FIELD_ESTIMATED_NUM_EPS	estimated_num_eps

## Enumerator

UCP_PARAM_FIELD_ESTIMATED_NUM_PPN	estimated_num_ppn
UCP_PARAM_FIELD_NAME	name

**6.2.5.2 ucp\_feature**

enum `ucp_feature`

The enumeration list describes the features supported by UCP. An application can request the features using [UCP parameters](#) during [UCP initialization](#) process.

## Enumerator

UCP_FEATURE_TAG	Request tag matching support
UCP_FEATURE_RMA	Request remote memory access support
UCP_FEATURE_AMO32	Request 32-bit atomic operations support
UCP_FEATURE_AMO64	Request 64-bit atomic operations support
UCP_FEATURE_WAKEUP	Request interrupt notification support
UCP_FEATURE_STREAM	Request stream support
UCP_FEATURE_AM	Request Active Message support

**6.2.5.3 ucp\_lib\_attr\_field**

enum `ucp_lib_attr_field`

The enumeration allows specifying which fields in `ucp_lib_attr_t` are present. It is used to enable backward compatibility support.

## Enumerator

UCP_LIB_ATTR_FIELD_MAX_THREAD_LEVEL	UCP library maximum supported thread level flag
-------------------------------------	---

**6.2.5.4 ucp\_context\_attr\_field**

enum `ucp_context_attr_field`

The enumeration allows specifying which fields in `ucp_context_attr_t` are present. It is used to enable backward compatibility support.

## Enumerator

UCP_ATTR_FIELD_REQUEST_SIZE	UCP request size
UCP_ATTR_FIELD_THREAD_MODE	UCP context thread flag
UCP_ATTR_FIELD_MEMORY_TYPES	UCP supported memory types
UCP_ATTR_FIELD_NAME	UCP context name



## 6.2.6 Function Documentation

### 6.2.6.1 ucp\_lib\_query()

```
ucs_status_t ucp_lib_query (
    ucp_lib_attr_t * attr )
```

This routine fetches information about the UCP library attributes.

#### Parameters

out	<i>attr</i>	Filled with attributes of the UCP library.
-----	-------------	--

#### Returns

Error code as defined by [ucs\\_status\\_t](#)

### 6.2.6.2 ucp\_get\_version()

```
void ucp_get_version (
    unsigned * major_version,
    unsigned * minor_version,
    unsigned * release_number )
```

This routine returns the UCP library version.

#### Parameters

out	<i>major_version</i>	Filled with library major version.
out	<i>minor_version</i>	Filled with library minor version.
out	<i>release_number</i>	Filled with library release number.

### 6.2.6.3 ucp\_get\_version\_string()

```
const char* ucp_get_version_string (
    void )
```

This routine returns the UCP library version as a string which consists of: "major.minor.release".

### 6.2.6.4 ucp\_init()

```
static ucs_status_t ucp_init (
    const ucp_params_t * params,
    const ucp_config_t * config,
    ucp_context_h * context_p ) [inline], [static]
```

This routine creates and initializes a [UCP application context](#).

**Warning**

This routine must be called before any other UCP function call in the application.

This routine checks API version compatibility, then discovers the available network interfaces, and initializes the network resources required for discovering of the network and memory related devices. This routine is responsible for initialization all information required for a particular application scope, for example, MPI application, OpenSHMEM application, etc.

**Note**

- Higher level protocols can add additional communication isolation, as MPI does with its communicator object. A single communication context may be used to support multiple MPI communicators.
- The context can be used to isolate the communication that corresponds to different protocols. For example, if MPI and OpenSHMEM are using UCP to isolate the MPI communication from the OpenSHMEM communication, users should use different application context for each of the communication libraries.

**Parameters**

in	<i>config</i>	UCP configuration descriptor allocated through <a href="#">ucp_config_read()</a> routine.
in	<i>params</i>	User defined <a href="#">ucp_params_t</a> configurations for the <a href="#">UCP application context</a> .
out	<i>context</i> <i>_p</i>	Initialized <a href="#">UCP application context</a> .

**Returns**

Error code as defined by [ucs\\_status\\_t](#)

**Examples**

[ucp\\_client\\_server.c](#), and [ucp\\_hello\\_world.c](#).

**6.2.6.5 ucp\_cleanup()**

```
void ucp_cleanup (
    ucp\_context\_h context_p )
```

This routine finalizes and releases the resources associated with a [UCP application context](#).

**Warning**

An application cannot call any UCP routine once the UCP application context released.

The cleanup process releases and shuts down all resources associated with the application context. After calling this routine, calling any UCP routine without calling [UCP initialization routine](#) is invalid.

**Parameters**

in	<i>context</i> <i>_p</i>	Handle to <a href="#">UCP application context</a> .
----	-----------------------------	---

**Examples**

[ucp\\_client\\_server.c](#), and [ucp\\_hello\\_world.c](#).

**6.2.6.6 ucp\_context\_query()**

```
ucs_status_t ucp_context_query (
    ucp_context_h context_p,
    ucp_context_attr_t * attr )
```

This routine fetches information about the context.

**Parameters**

in	<i>context_p</i>	Handle to <a href="#">UCP application context</a> .
out	<i>attr</i>	Filled with attributes of <code>context_p</code> context.

**Returns**

Error code as defined by [ucs\\_status\\_t](#)

**6.2.6.7 ucp\_context\_print\_info()**

```
void ucp_context_print_info (
    const ucp_context_h context,
    FILE * stream )
```

This routine prints information about the context configuration: including memory domains, transport resources, and other useful information associated with the context.

**Parameters**

in	<i>context</i>	Print this context object's configuration.
in	<i>stream</i>	Output stream on which to print the information.

## 6.3 UCP Worker

### Data Structures

- struct [ucp\\_worker\\_attr](#)  
*UCP worker attributes. [More...](#)*
- struct [ucp\\_worker\\_params](#)  
*Tuning parameters for the UCP worker. [More...](#)*
- struct [ucp\\_worker\\_address\\_attr](#)  
*UCP worker address attributes. [More...](#)*
- struct [ucp\\_listener\\_attr](#)  
*UCP listener attributes. [More...](#)*
- struct [ucp\\_conn\\_request\\_attr](#)  
*UCP listener's connection request attributes. [More...](#)*
- struct [ucp\\_listener\\_params](#)  
*Parameters for a UCP listener object. [More...](#)*
- struct [ucp\\_am\\_handler\\_param](#)  
*Active Message handler parameters passed to [ucp\\_worker\\_set\\_am\\_rcv\\_handler](#) routine. [More...](#)*
- struct [ucp\\_am\\_rcv\\_param](#)  
*Operation parameters provided in [ucp\\_am\\_rcv\\_callback\\_t](#) callback. [More...](#)*
- struct [ucp\\_listener\\_accept\\_handler](#)
- struct [ucp\\_listener\\_conn\\_handler](#)  
*UCP callback to handle the connection request in a client-server connection establishment flow. [More...](#)*

### Typedefs

- typedef struct [ucp\\_worker\\_attr](#) [ucp\\_worker\\_attr\\_t](#)  
*UCP worker attributes.*
- typedef struct [ucp\\_worker\\_params](#) [ucp\\_worker\\_params\\_t](#)  
*Tuning parameters for the UCP worker.*
- typedef struct [ucp\\_worker\\_address\\_attr](#) [ucp\\_worker\\_address\\_attr\\_t](#)  
*UCP worker address attributes.*
- typedef struct [ucp\\_listener\\_attr](#) [ucp\\_listener\\_attr\\_t](#)  
*UCP listener attributes.*
- typedef struct [ucp\\_conn\\_request\\_attr](#) [ucp\\_conn\\_request\\_attr\\_t](#)  
*UCP listener's connection request attributes.*
- typedef struct [ucp\\_listener\\_params](#) [ucp\\_listener\\_params\\_t](#)  
*Parameters for a UCP listener object.*
- typedef struct [ucp\\_am\\_handler\\_param](#) [ucp\\_am\\_handler\\_param\\_t](#)  
*Active Message handler parameters passed to [ucp\\_worker\\_set\\_am\\_rcv\\_handler](#) routine.*
- typedef struct [ucp\\_listener\\_accept\\_handler](#) [ucp\\_listener\\_accept\\_handler\\_t](#)
- typedef struct [ucp\\_am\\_rcv\\_param](#) [ucp\\_am\\_rcv\\_param\\_t](#)  
*Operation parameters provided in [ucp\\_am\\_rcv\\_callback\\_t](#) callback.*
- typedef struct [ucp\\_address](#) [ucp\\_address\\_t](#)  
*UCP worker address.*
- typedef struct [ucp\\_listener](#) \* [ucp\\_listener\\_h](#)  
*UCP listen handle.*
- typedef struct [ucp\\_worker](#) \* [ucp\\_worker\\_h](#)  
*UCP Worker.*
- typedef void(\* [ucp\\_listener\\_accept\\_callback\\_t](#)) ([ucp\\_ep\\_h](#) ep, void \*arg)  
*A callback for accepting client/server connections on a listener [ucp\\_listener\\_h](#).*

- typedef void(\* [ucp\\_listener\\_conn\\_callback\\_t](#)) ([ucp\\_conn\\_request\\_h](#) conn\_request, void \*arg)  
*A callback for handling of incoming connection request conn\_request from a client.*
- typedef struct [ucp\\_listener\\_conn\\_handler](#) [ucp\\_listener\\_conn\\_handler\\_t](#)  
*UCP callback to handle the connection request in a client-server connection establishment flow.*
- typedef enum [ucp\\_wakeup\\_event\\_types](#) [ucp\\_wakeup\\_event\\_t](#)  
*UCP worker wakeup events mask.*

## Enumerations

- enum [ucp\\_worker\\_params\\_field](#) {  
[UCP\\_WORKER\\_PARAM\\_FIELD\\_THREAD\\_MODE](#) = UCS\_BIT(0), [UCP\\_WORKER\\_PARAM\\_FIELD\\_CPU\\_MASK](#) = UCS\_BIT(1), [UCP\\_WORKER\\_PARAM\\_FIELD\\_EVENTS](#) = UCS\_BIT(2), [UCP\\_WORKER\\_PARAM\\_FIELD\\_USER\\_DATA](#) = UCS\_BIT(3),  
[UCP\\_WORKER\\_PARAM\\_FIELD\\_EVENT\\_FD](#) = UCS\_BIT(4), [UCP\\_WORKER\\_PARAM\\_FIELD\\_FLAGS](#) = UCS\_BIT(5), [UCP\\_WORKER\\_PARAM\\_FIELD\\_NAME](#) = UCS\_BIT(6), [UCP\\_WORKER\\_PARAM\\_FIELD\\_AM\\_ALIGNMENT](#) = UCS\_BIT(7),  
[UCP\\_WORKER\\_PARAM\\_FIELD\\_CLIENT\\_ID](#) = UCS\_BIT(8) }  
*UCP worker parameters field mask.*
- enum [ucp\\_worker\\_flags\\_t](#) { [UCP\\_WORKER\\_FLAG\\_IGNORE\\_REQUEST\\_LEAK](#) = UCS\_BIT(0) }  
*UCP worker flags.*
- enum [ucp\\_listener\\_params\\_field](#) { [UCP\\_LISTENER\\_PARAM\\_FIELD\\_SOCK\\_ADDR](#) = UCS\_BIT(0), [UCP\\_LISTENER\\_PARAM\\_FIELD\\_ACCEPT\\_HANDLER](#) = UCS\_BIT(1), [UCP\\_LISTENER\\_PARAM\\_FIELD\\_CONN\\_HANDLER](#) = UCS\_BIT(2) }  
*UCP listener parameters field mask.*
- enum [ucp\\_worker\\_address\\_flags\\_t](#) { [UCP\\_WORKER\\_ADDRESS\\_FLAG\\_NET\\_ONLY](#) = UCS\_BIT(0) }  
*UCP worker address flags.*
- enum [ucp\\_worker\\_attr\\_field](#) {  
[UCP\\_WORKER\\_ATTR\\_FIELD\\_THREAD\\_MODE](#) = UCS\_BIT(0), [UCP\\_WORKER\\_ATTR\\_FIELD\\_ADDRESS](#) = UCS\_BIT(1), [UCP\\_WORKER\\_ATTR\\_FIELD\\_ADDRESS\\_FLAGS](#) = UCS\_BIT(2), [UCP\\_WORKER\\_ATTR\\_FIELD\\_MAX\\_AM](#) = UCS\_BIT(3),  
[UCP\\_WORKER\\_ATTR\\_FIELD\\_NAME](#) = UCS\_BIT(4), [UCP\\_WORKER\\_ATTR\\_FIELD\\_MAX\\_INFO\\_STRING](#) = UCS\_BIT(5) }  
*UCP worker attributes field mask.*
- enum [ucp\\_worker\\_address\\_attr\\_field](#) { [UCP\\_WORKER\\_ADDRESS\\_ATTR\\_FIELD\\_UID](#) = UCS\_BIT(0) }  
*UCP worker address attributes field mask.*
- enum [ucp\\_listener\\_attr\\_field](#) { [UCP\\_LISTENER\\_ATTR\\_FIELD\\_SOCKADDR](#) = UCS\_BIT(0) }  
*UCP listener attributes field mask.*
- enum [ucp\\_conn\\_request\\_attr\\_field](#) { [UCP\\_CONN\\_REQUEST\\_ATTR\\_FIELD\\_CLIENT\\_ADDR](#) = UCS\_BIT(0), [UCP\\_CONN\\_REQUEST\\_ATTR\\_FIELD\\_CLIENT\\_ID](#) = UCS\_BIT(1) }  
*UCP listener's connection request attributes field mask.*
- enum [ucp\\_am\\_cb\\_flags](#) { [UCP\\_AM\\_FLAG\\_WHOLE\\_MSG](#) = UCS\_BIT(0), [UCP\\_AM\\_FLAG\\_PERSISTENT\\_DATA](#) = UCS\_BIT(1) }  
*Flags for a UCP Active Message callback.*
- enum [ucp\\_send\\_am\\_flags](#) { [UCP\\_AM\\_SEND\\_FLAG\\_REPLY](#) = UCS\_BIT(0), [UCP\\_AM\\_SEND\\_FLAG\\_EAGER](#) = UCS\_BIT(1), [UCP\\_AM\\_SEND\\_FLAG\\_RNDV](#) = UCS\_BIT(2), [UCP\\_AM\\_SEND\\_REPLY](#) = UCP\_AM\_SEND\_FLAG\_REPLY }  
*Flags for sending a UCP Active Message.*
- enum [ucp\\_wakeup\\_event\\_types](#) {  
[UCP\\_WAKEUP\\_RMA](#) = UCS\_BIT(0), [UCP\\_WAKEUP\\_AMO](#) = UCS\_BIT(1), [UCP\\_WAKEUP\\_TAG\\_SEND](#) = UCS\_BIT(2), [UCP\\_WAKEUP\\_TAG\\_RECV](#) = UCS\_BIT(3),  
[UCP\\_WAKEUP\\_TX](#) = UCS\_BIT(10), [UCP\\_WAKEUP\\_RX](#) = UCS\_BIT(11), [UCP\\_WAKEUP\\_EDGE](#) = UCS\_BIT(16) }  
*UCP worker wakeup events mask.*

## Functions

- [ucs\\_status\\_t ucp\\_worker\\_create](#) ([ucp\\_context\\_h](#) context, const [ucp\\_worker\\_params\\_t](#) \*params, [ucp\\_worker\\_h](#) \*worker\_p)  
*Create a worker object.*
- void [ucp\\_worker\\_destroy](#) ([ucp\\_worker\\_h](#) worker)  
*Destroy a worker object.*
- [ucs\\_status\\_t ucp\\_worker\\_query](#) ([ucp\\_worker\\_h](#) worker, [ucp\\_worker\\_attr\\_t](#) \*attr)  
*Get attributes specific to a particular worker.*
- void [ucp\\_worker\\_print\\_info](#) ([ucp\\_worker\\_h](#) worker, FILE \*stream)  
*Print information about the worker.*
- [ucs\\_status\\_t ucp\\_worker\\_get\\_address](#) ([ucp\\_worker\\_h](#) worker, [ucp\\_address\\_t](#) \*\*address\_p, [size\\_t](#) \*address\_length\_p)  
*Get the address of the worker object.*
- void [ucp\\_worker\\_release\\_address](#) ([ucp\\_worker\\_h](#) worker, [ucp\\_address\\_t](#) \*address)  
*Release an address of the worker object.*
- [ucs\\_status\\_t ucp\\_worker\\_address\\_query](#) ([ucp\\_address\\_t](#) \*address, [ucp\\_worker\\_address\\_attr\\_t](#) \*attr)  
*Get attributes of the particular worker address.*
- unsigned [ucp\\_worker\\_progress](#) ([ucp\\_worker\\_h](#) worker)  
*Progress all communications on a specific worker.*
- [ssize\\_t ucp\\_stream\\_worker\\_poll](#) ([ucp\\_worker\\_h](#) worker, [ucp\\_stream\\_poll\\_ep\\_t](#) \*poll\_eps, [size\\_t](#) max\_eps, unsigned flags)  
*Poll for endpoints that are ready to consume streaming data.*
- [ucs\\_status\\_t ucp\\_listener\\_create](#) ([ucp\\_worker\\_h](#) worker, const [ucp\\_listener\\_params\\_t](#) \*params, [ucp\\_listener\\_h](#) \*listener\_p)  
*Create a listener to accept connections on. Connection requests on the listener will arrive at a local address specified by the user.*
- void [ucp\\_listener\\_destroy](#) ([ucp\\_listener\\_h](#) listener)  
*Stop accepting connections on a local address of the worker object.*
- [ucs\\_status\\_t ucp\\_listener\\_query](#) ([ucp\\_listener\\_h](#) listener, [ucp\\_listener\\_attr\\_t](#) \*attr)  
*Get attributes specific to a particular listener.*
- [ucs\\_status\\_t ucp\\_conn\\_request\\_query](#) ([ucp\\_conn\\_request\\_h](#) conn\_request, [ucp\\_conn\\_request\\_attr\\_t](#) \*attr)  
*Get attributes specific to a particular connection request received on the server side.*
- [ucs\\_status\\_t ucp\\_listener\\_reject](#) ([ucp\\_listener\\_h](#) listener, [ucp\\_conn\\_request\\_h](#) conn\_request)  
*Reject an incoming connection request.*
- [ucs\\_status\\_t ucp\\_worker\\_set\\_am\\_handler](#) ([ucp\\_worker\\_h](#) worker, [uint16\\_t](#) id, [ucp\\_am\\_callback\\_t](#) cb, void \*arg, [uint32\\_t](#) flags)  
*Add user defined callback for Active Message.*
- [ucs\\_status\\_t ucp\\_worker\\_set\\_am\\_rcv\\_handler](#) ([ucp\\_worker\\_h](#) worker, const [ucp\\_am\\_handler\\_param\\_t](#) \*param)  
*Add user defined callback for Active Message.*
- [ucs\\_status\\_t ucp\\_worker\\_fence](#) ([ucp\\_worker\\_h](#) worker)  
*Assures ordering between non-blocking operations.*
- [ucs\\_status\\_ptr\\_t ucp\\_worker\\_flush\\_nb](#) ([ucp\\_worker\\_h](#) worker, unsigned flags, [ucp\\_send\\_callback\\_t](#) cb)  
*Flush outstanding AMO and RMA operations on the worker.*
- [ucs\\_status\\_ptr\\_t ucp\\_worker\\_flush\\_nbx](#) ([ucp\\_worker\\_h](#) worker, const [ucp\\_request\\_param\\_t](#) \*param)  
*Flush outstanding AMO and RMA operations on the worker.*
- [ucs\\_status\\_t ucp\\_worker\\_flush](#) ([ucp\\_worker\\_h](#) worker)  
*Flush outstanding AMO and RMA operations on the worker.*

### 6.3.1 Detailed Description

UCP Worker routines

## 6.3.2 Data Structure Documentation

### 6.3.2.1 struct ucp\_worker\_attr

The structure defines the attributes which characterize the particular worker.

#### Data Fields

uint64_t	field_mask	Mask of valid fields in this structure, using bits from <a href="#">ucp_worker_attr_field</a> . Fields not specified in this mask will be ignored. Provides ABI compatibility with respect to adding new fields.
<a href="#">ucs_thread_mode_t</a>	thread_mode	Thread safe level of the worker.
uint32_t	address_flags	Flags indicating requested details of the worker address. If <a href="#">UCP_WORKER_ATTR_FIELD_ADDRESS_FLAGS</a> bit is set in the <code>field_mask</code> , this value should be set as well. Possible flags are specified in <a href="#">ucp_worker_address_flags_t</a> .  <b>Note</b> This is an input attribute.
<a href="#">ucp_address_t</a> *	address	Worker address, which can be passed to remote instances of the UCP library in order to connect to this worker. The memory for the address handle is allocated by <a href="#">ucp_worker_query()</a> routine, and must be released by using <a href="#">ucp_worker_release_address()</a> routine.
size_t	address_length	Size of worker address in bytes.
size_t	max_am_header	Maximum allowed header size for <a href="#">ucp_am_send_nbx</a> routine.
char	name[ <a href="#">UCP_ENTITY_NAME_MAX</a> ]	Tracing and analysis tools can identify the worker using this name.
size_t	max_debug_string	Maximum debug string size that can be filled with <a href="#">ucp_request_query</a> .

### 6.3.2.2 struct ucp\_worker\_params

The structure defines the parameters that are used for the UCP worker tuning during the UCP worker [creation](#).

#### Examples

[ucp\\_client\\_server.c](#), and [ucp\\_hello\\_world.c](#).

#### Data Fields

uint64_t	field_mask	Mask of valid fields in this structure, using bits from <a href="#">ucp_worker_params_field</a> . Fields not specified in this mask will be ignored. Provides ABI compatibility with respect to adding new fields.
----------	------------	--

## Data Fields

<a href="#">ucs_thread_mode_t</a>	thread_mode	The parameter <code>thread_mode</code> suggests the thread safety mode which worker and the associated resources should be created with. This is an optional parameter. The default value is <code>UCS_THREAD_MODE_SINGLE</code> and it is used when the value of the parameter is not set. When this parameter along with its corresponding bit in the <code>field_mask</code> - <code>UCP_WORKER_PARAM_FIELD_THREAD_MODE</code> is set, the <a href="#">ucp_worker_create</a> attempts to create worker with this thread mode. The thread mode with which worker is created can differ from the suggested mode. The actual thread mode of the worker should be obtained using the query interface <a href="#">ucp_worker_query</a> .
<a href="#">ucs_cpu_set_t</a>	cpu_mask	Mask of which CPUs worker resources should preferably be allocated on. This value is optional. If it's not set (along with its corresponding bit in the <code>field_mask</code> - <code>UCP_WORKER_PARAM_FIELD_CPU_MASK</code> ), resources are allocated according to system's default policy.
unsigned	events	Mask of events ( <a href="#">ucp_wakeup_event_t</a> ) which are expected on wakeup. This value is optional. If it's not set (along with its corresponding bit in the <code>field_mask</code> - <code>UCP_WORKER_PARAM_FIELD_EVENTS</code> ), all types of events will trigger on wakeup.
void *	user_data	User data associated with the current worker. This value is optional. If it's not set (along with its corresponding bit in the <code>field_mask</code> - <code>UCP_WORKER_PARAM_FIELD_USER_DATA</code> ), it will default to <code>NULL</code> .
int	event_fd	External event file descriptor. This value is optional. If <a href="#">UCP_WORKER_PARAM_FIELD_EVENT_FD</a> is set in the <code>field_mask</code> , events on the worker will be reported on the provided event file descriptor. In this case, calling <a href="#">ucp_worker_get_efd</a> will result in an error. The provided file descriptor must be capable of aggregating notifications for arbitrary events, for example <code>epoll(7)</code> on Linux systems. <code>user_data</code> will be used as the event user-data on systems which support it. For example, on Linux, it will be placed in <code>epoll_data_t::ptr</code> , when returned from <code>epoll_wait(2)</code> . Otherwise, events will be reported to the event file descriptor returned from <a href="#">ucp_worker_get_efd()</a> .
uint64_t	flags	Worker flags. This value is optional. If <a href="#">UCP_WORKER_PARAM_FIELD_FLAGS</a> is not set in the <code>field_mask</code> , the value of this field will default to 0.
const char *	name	Tracing and analysis tools can identify the worker using this name. To retrieve the worker's name, use <a href="#">ucp_worker_query</a> , as the name you supply may be changed by UCX under some circumstances, e.g. a name conflict. This field is only assigned if you set <a href="#">UCP_WORKER_PARAM_FIELD_NAME</a> in the field mask. If not, then a default unique name will be created for you.
size_t	am_alignment	Minimal address alignment of the active message data pointer as passed in argument <code>data</code> to the active message handler, defined as <a href="#">ucp_am_recv_callback_t</a> .
uint64_t	client_id	Client id that is sent as part of the connection request payload when connecting to a remote socket address. On the remote side, this value can be obtained from <a href="#">ucp_conn_request_h</a> using <a href="#">ucp_conn_request_query</a> .



### 6.3.2.3 struct ucp\_worker\_address\_attr

The structure defines the attributes of the particular worker address.

#### Data Fields

uint64_t	field_mask	Mask of valid fields in this structure, using bits from <a href="#">ucp_worker_address_attr_field</a> . Fields not specified in this mask will be ignored. Provides ABI compatibility with respect to adding new fields.
uint64_t	worker_uid	Unique id of the worker this address belongs to.

### 6.3.2.4 struct ucp\_listener\_attr

The structure defines the attributes which characterize the particular listener.

#### Examples

[ucp\\_client\\_server.c](#).

#### Data Fields

uint64_t	field_mask	Mask of valid fields in this structure, using bits from <a href="#">ucp_listener_attr_field</a> . Fields not specified in this mask will be ignored. Provides ABI compatibility with respect to adding new fields.
struct sockaddr_storage	sockaddr	Sockaddr on which this listener is listening for incoming connection requests.

### 6.3.2.5 struct ucp\_conn\_request\_attr

The structure defines the attributes that characterize the particular connection request received on the server side.

#### Examples

[ucp\\_client\\_server.c](#).

#### Data Fields

uint64_t	field_mask	Mask of valid fields in this structure, using bits from <a href="#">ucp_conn_request_attr_field</a> . Fields not specified in this mask will be ignored. Provides ABI compatibility with respect to adding new fields.
struct sockaddr_storage	client_address	The address of the remote client that sent the connection request to the server.
uint64_t	client_id	Remote client id if remote endpoint's flag <a href="#">UCP_EP_PARAMS_FLAGS_SEND_CLIENT_ID</a> is set.

### 6.3.2.6 struct ucp\_listener\_params

This structure defines parameters for [ucp\\_listener\\_create](#), which is used to listen for incoming client/server connections.

## Examples

[ucp\\_client\\_server.c](#).

## Data Fields

<a href="#">uint64_t</a>	<code>field_mask</code>	Mask of valid fields in this structure, using bits from <a href="#">ucp_listener_params_field</a> . Fields not specified in this mask will be ignored. Provides ABI compatibility with respect to adding new fields.
<a href="#">ucs_sock_addr_t</a>	<code>sockaddr</code>	An address in the form of a <code>sockaddr</code> . This field is mandatory for filling (along with its corresponding bit in the <code>field_mask</code> - <a href="#">UCP_LISTENER_PARAM_FIELD_SOCK_ADDR</a> ). The <a href="#">ucp_listener_create</a> routine will return with an error if <code>sockaddr</code> is not specified.
<a href="#">ucp_listener_accept_handler_t</a>	<code>accept_handler</code>	Handler to endpoint creation in a client-server connection flow. In order for the callback inside this handler to be invoked, the <a href="#">UCP_LISTENER_PARAM_FIELD_ACCEPT_HANDLER</a> needs to be set in the <code>field_mask</code> .
<a href="#">ucp_listener_conn_handler_t</a>	<code>conn_handler</code>	<p>Handler of an incoming connection request in a client-server connection flow. In order for the callback inside this handler to be invoked, the <a href="#">UCP_LISTENER_PARAM_FIELD_CONN_HANDLER</a> needs to be set in the <code>field_mask</code>.</p> <p><b>Note</b></p> <p>User is expected to call <code>ucp_ep_create</code> with set <a href="#">UCP_EP_PARAM_FIELD_CONN_REQUEST</a> flag to <code>ucp_ep_params_t::field_mask</code> and <a href="#">ucp_ep_params_t::conn_request</a> in order to be able to receive communications.</p>

6.3.2.7 struct `ucp_am_handler_param`

## Examples

[ucp\\_client\\_server.c](#).

## Data Fields

<a href="#">uint64_t</a>	<code>field_mask</code>	Mask of valid fields in this structure, using bits from <a href="#">ucp_am_handler_param_field</a> . Fields not specified in this mask will be ignored. Provides ABI compatibility with respect to adding new fields.
<code>unsigned</code>	<code>id</code>	Active Message id.
<a href="#">uint32_t</a>	<code>flags</code>	Handler flags as defined by <a href="#">ucp_am_cb_flags</a> .
<a href="#">ucp_am_recv_callback_t</a>	<code>cb</code>	Active Message callback. To clear the already set callback, this value should be set to <code>NULL</code> .
<code>void *</code>	<code>arg</code>	Active Message argument, which will be passed in to every invocation of <a href="#">ucp_am_recv_callback_t</a> function as the <code>arg</code> argument.

### 6.3.2.8 struct ucp\_am\_recv\_param

Examples

[ucp\\_client\\_server.c](#).

Data Fields

<a href="#">uint64_t</a>	<code>recv_attr</code>	Mask of valid fields in this structure and receive operation flags, using bits from <a href="#">ucp_am_recv_attr_t</a> . Fields not specified in this mask will be ignored. Provides ABI compatibility with respect to adding new fields.
<a href="#">ucp_ep_h</a>	<code>reply_ep</code>	Endpoint, which can be used for the reply to this message.

### 6.3.2.9 struct ucp\_listener\_accept\_handler

**Deprecated** Replaced by [ucp\\_listener\\_conn\\_handler\\_t](#).

Data Fields

<a href="#">ucp_listener_accept_callback_t</a>	<code>cb</code>	Endpoint creation callback
	<code>void *</code>	<code>arg</code> User defined argument for the callback

### 6.3.2.10 struct ucp\_listener\_conn\_handler

This structure is used for handling an incoming connection request on the listener. Setting this type of handler allows creating an endpoint on any other worker and not limited to the worker on which the listener was created.

Note

- Other than communication progress routines, it is allowed to call all other communication routines from the callback in the struct.
- The callback is thread safe with respect to the worker it is invoked on.
- It is the user's responsibility to avoid potential dead lock accessing different worker.

Data Fields

<a href="#">ucp_listener_conn_callback_t</a>	<code>cb</code>	Connection request callback
	<code>void *</code>	<code>arg</code> User defined argument for the callback

## 6.3.3 Typedef Documentation

### 6.3.3.1 ucp\_worker\_attr\_t

```
typedef struct ucp_worker_attr ucp_worker_attr_t
```

The structure defines the attributes which characterize the particular worker.

### 6.3.3.2 `ucp_worker_params_t`

```
typedef struct ucp_worker_params ucp_worker_params_t
```

The structure defines the parameters that are used for the UCP worker tuning during the UCP worker [creation](#).

### 6.3.3.3 `ucp_worker_address_attr_t`

```
typedef struct ucp_worker_address_attr ucp_worker_address_attr_t
```

The structure defines the attributes of the particular worker address.

### 6.3.3.4 `ucp_listener_attr_t`

```
typedef struct ucp_listener_attr ucp_listener_attr_t
```

The structure defines the attributes which characterize the particular listener.

### 6.3.3.5 `ucp_conn_request_attr_t`

```
typedef struct ucp_conn_request_attr ucp_conn_request_attr_t
```

The structure defines the attributes that characterize the particular connection request received on the server side.

### 6.3.3.6 `ucp_listener_params_t`

```
typedef struct ucp_listener_params ucp_listener_params_t
```

This structure defines parameters for [ucp\\_listener\\_create](#), which is used to listen for incoming client/server connections.

### 6.3.3.7 `ucp_am_handler_param_t`

```
typedef struct ucp_am_handler_param ucp_am_handler_param_t
```

### 6.3.3.8 `ucp_listener_accept_handler_t`

```
typedef struct ucp_listener_accept_handler ucp_listener_accept_handler_t
```

**Deprecated** Replaced by [ucp\\_listener\\_conn\\_handler\\_t](#).

### 6.3.3.9 `ucp_am_recv_param_t`

```
typedef struct ucp_am_recv_param ucp_am_recv_param_t
```

### 6.3.3.10 `ucp_address_t`

```
typedef struct ucp_address ucp_address_t
```

The address handle is an opaque object that is used as an identifier for a [worker](#) instance.

### 6.3.3.11 ucp\_listener\_h

```
typedef struct ucp_listener* ucp_listener_h
```

The listener handle is an opaque object that is used for listening on a specific address and accepting connections from clients.

### 6.3.3.12 ucp\_worker\_h

```
typedef struct ucp_worker* ucp_worker_h
```

UCP worker is an opaque object representing the communication context. The worker represents an instance of a local communication resource and the progress engine associated with it. The progress engine is a construct that is responsible for asynchronous and independent progress of communication directives. The progress engine could be implemented in hardware or software. The worker object abstracts an instance of network resources such as a host channel adapter port, network interface, or multiple resources such as multiple network interfaces or communication ports. It could also represent virtual communication resources that are defined across multiple devices. Although the worker can represent multiple network resources, it is associated with a single [UCX application context](#). All communication functions require a context to perform the operation on the dedicated hardware resource(s) and an [endpoint](#) to address the destination.

#### Note

Worker are parallel "threading points" that an upper layer may use to optimize concurrent communications.

### 6.3.3.13 ucp\_listener\_accept\_callback\_t

```
typedef void(* ucp_listener_accept_callback_t) (ucp_ep_h ep, void *arg)
```

This callback routine is invoked on the server side upon creating a connection to a remote client. The user can pass an argument to this callback. The user is responsible for releasing the *ep* handle using the [ucp\\_ep\\_destroy\(\)](#) routine.

#### Parameters

in	<i>ep</i>	Handle to a newly created endpoint which is connected to the remote peer which has initiated the connection.
in	<i>arg</i>	User's argument for the callback.

### 6.3.3.14 ucp\_listener\_conn\_callback\_t

```
typedef void(* ucp_listener_conn_callback_t) (ucp_conn_request_h conn_request, void *arg)
```

This callback routine is invoked on the server side to handle incoming connections from remote clients. The user can pass an argument to this callback. The *conn\_request* handle has to be released, either by [ucp\\_ep\\_create](#) or [ucp\\_listener\\_reject](#) routine.

#### Parameters

in	<i>conn_request</i>	Connection request handle.
in	<i>arg</i>	User's argument for the callback.

### 6.3.3.15 `ucp_listener_conn_handler_t`

```
typedef struct ucp_listener_conn_handler ucp_listener_conn_handler_t
```

This structure is used for handling an incoming connection request on the listener. Setting this type of handler allows creating an endpoint on any other worker and not limited to the worker on which the listener was created.

#### Note

- Other than communication progress routines, it is allowed to call all other communication routines from the callback in the struct.
- The callback is thread safe with respect to the worker it is invoked on.
- It is the user's responsibility to avoid potential dead lock accessing different worker.

### 6.3.3.16 `ucp_wakeup_event_t`

```
typedef enum ucp_wakeup_event_types ucp_wakeup_event_t
```

The enumeration allows specifying which events are expected on wakeup. Empty events are possible for any type of event except for `UCP_WAKEUP_TX` and `UCP_WAKEUP_RX`.

#### Note

Send completions are reported by POLLIN-like events (see poll man page). Since outgoing operations can be initiated at any time, UCP does not generate POLLOUT-like events, although it must be noted that outgoing operations may be queued depending upon resource availability.

## 6.3.4 Enumeration Type Documentation

### 6.3.4.1 `ucp_worker_params_field`

```
enum ucp_worker_params_field
```

The enumeration allows specifying which fields in `ucp_worker_params_t` are present. It is used to enable backward compatibility support.

#### Enumerator

<code>UCP_WORKER_PARAM_FIELD_THREAD_MODE</code>	UCP thread mode
<code>UCP_WORKER_PARAM_FIELD_CPU_MASK</code>	Worker's CPU bitmap
<code>UCP_WORKER_PARAM_FIELD_EVENTS</code>	Worker's events bitmap
<code>UCP_WORKER_PARAM_FIELD_USER_DATA</code>	User data
<code>UCP_WORKER_PARAM_FIELD_EVENT_FD</code>	External event file descriptor
<code>UCP_WORKER_PARAM_FIELD_FLAGS</code>	Worker flags
<code>UCP_WORKER_PARAM_FIELD_NAME</code>	Worker name
<code>UCP_WORKER_PARAM_FIELD_AM_ALIGNMENT</code>	Alignment of active messages on the receiver
<code>UCP_WORKER_PARAM_FIELD_CLIENT_ID</code>	Client id

### 6.3.4.2 ucp\_worker\_flags\_t

enum `ucp_worker_flags_t`

This enumeration allows specifying flags for `ucp_worker_params_t::flags`, which is used as parameter for `ucp_worker_create`.

Enumerator

UCP_WORKER_FLAG_IGNORE_REQUEST_LEAK	Do not print warnings about request leaks
-------------------------------------	---

### 6.3.4.3 ucp\_listener\_params\_field

enum `ucp_listener_params_field`

The enumeration allows specifying which fields in `ucp_listener_params_t` are present. It is used to enable backward compatibility support.

Enumerator

UCP_LISTENER_PARAM_FIELD_SOCK_ADDR	Socket address and length.
UCP_LISTENER_PARAM_FIELD_ACCEPT_HANDLER	User's callback and argument for handling the creation of an endpoint. User's callback and argument for handling the incoming connection request.
UCP_LISTENER_PARAM_FIELD_CONN_HANDLER	

### 6.3.4.4 ucp\_worker\_address\_flags\_t

enum `ucp_worker_address_flags_t`

The enumeration list describes possible UCP worker address flags, indicating what needs to be included to the worker address returned by `ucp_worker_query()` routine.

Enumerator

UCP_WORKER_ADDRESS_FLAG_NET_ONLY	Pack addresses of network devices only. Using such shortened addresses for the remote node peers will reduce the amount of wireup data being exchanged during connection establishment phase.
----------------------------------	---

### 6.3.4.5 ucp\_worker\_attr\_field

enum `ucp_worker_attr_field`

The enumeration allows specifying which fields in `ucp_worker_attr_t` are present. It is used to enable backward compatibility support.

## Enumerator

UCP_WORKER_ATTR_FIELD_THREAD_MODE	UCP thread mode
UCP_WORKER_ATTR_FIELD_ADDRESS	UCP address
UCP_WORKER_ATTR_FIELD_ADDRESS_FLAGS	UCP address flags
UCP_WORKER_ATTR_FIELD_MAX_AM_HEADER	Maximum header size used by UCP AM API
UCP_WORKER_ATTR_FIELD_NAME	UCP worker name
UCP_WORKER_ATTR_FIELD_MAX_INFO_STRING	Maximum size of info string

**6.3.4.6 ucp\_worker\_address\_attr\_field**

enum `ucp_worker_address_attr_field`

The enumeration allows specifying which fields in `ucp_worker_address_attr_t` are present. It is used to enable backward compatibility support.

## Enumerator

UCP_WORKER_ADDRESS_ATTR_FIELD_UID	Unique id of the worker
-----------------------------------	-------------------------

**6.3.4.7 ucp\_listener\_attr\_field**

enum `ucp_listener_attr_field`

The enumeration allows specifying which fields in `ucp_listener_attr_t` are present. It is used to enable backward compatibility support.

## Enumerator

UCP_LISTENER_ATTR_FIELD_SOCKADDR	Sockaddr used for listening
----------------------------------	-----------------------------

**6.3.4.8 ucp\_conn\_request\_attr\_field**

enum `ucp_conn_request_attr_field`

The enumeration allows specifying which fields in `ucp_conn_request_attr_t` are present. It is used to enable backward compatibility support.

## Enumerator

UCP_CONN_REQUEST_ATTR_FIELD_CLIENT_ADDR	Client's address
UCP_CONN_REQUEST_ATTR_FIELD_CLIENT_ID	Remote client id



### 6.3.4.9 ucp\_am\_cb\_flags

enum `ucp_am_cb_flags`

Flags that indicate how to handle UCP Active Messages.

#### Enumerator

<code>UCP_AM_FLAG_WHOLE_MSG</code>	Indicates that the entire message will be handled in one callback.
<code>UCP_AM_FLAG_PERSISTENT_DATA</code>	Guarantees that the specified <code>ucp_am_recv_callback_t</code> callback, will always be called with <code>UCP_AM_RECV_ATTR_FLAG_DATA</code> flag set, so the data will be accessible outside the callback, until <code>ucp_am_data_release</code> is called.

### 6.3.4.10 ucp\_send\_am\_flags

enum `ucp_send_am_flags`

Flags dictate the behavior of `ucp_am_send_nb` and `ucp_am_send_nbx` routines.

#### Enumerator

<code>UCP_AM_SEND_FLAG_REPLY</code>	Force relevant reply endpoint to be passed to the data callback on the receiver.
<code>UCP_AM_SEND_FLAG_EAGER</code>	Force UCP to use only eager protocol for AM sends.
<code>UCP_AM_SEND_FLAG_RNDV</code>	Force UCP to use only rendezvous protocol for AM sends.
<code>UCP_AM_SEND_REPLY</code>	Backward compatibility.

### 6.3.4.11 ucp\_wakeup\_event\_types

enum `ucp_wakeup_event_types`

The enumeration allows specifying which events are expected on wakeup. Empty events are possible for any type of event except for `UCP_WAKEUP_TX` and `UCP_WAKEUP_RX`.

#### Note

Send completions are reported by POLLIN-like events (see poll man page). Since outgoing operations can be initiated at any time, UCP does not generate POLLOUT-like events, although it must be noted that outgoing operations may be queued depending upon resource availability.

#### Enumerator

<code>UCP_WAKEUP_RMA</code>	Remote memory access send completion
<code>UCP_WAKEUP_AMO</code>	Atomic operation send completion
<code>UCP_WAKEUP_TAG_SEND</code>	Tag send completion
<code>UCP_WAKEUP_TAG_RECV</code>	Tag receive completion
<code>UCP_WAKEUP_TX</code>	This event type will generate an event on completion of any outgoing operation (complete or partial, according to the underlying protocol) for any type of transfer (send, atomic, or RMA).

## Enumerator

UCP_WAKEUP_RX	This event type will generate an event on completion of any receive operation (complete or partial, according to the underlying protocol).
UCP_WAKEUP_EDGE	Use edge-triggered wakeup. The event file descriptor will be signaled only for new events, rather than existing ones.

## 6.3.5 Function Documentation

### 6.3.5.1 ucp\_worker\_create()

```
ucs_status_t ucp_worker_create (
    ucp_context_h context,
    const ucp_worker_params_t * params,
    ucp_worker_h * worker_p )
```

This routine allocates and initializes a [worker](#) object. Each worker is associated with one and only one [application](#) context. In the same time, an application context can create multiple [workers](#) in order to enable concurrent access to communication resources. For example, application can allocate a dedicated worker for each application thread, where every worker can be progressed independently of others.

#### Note

The worker object is allocated within context of the calling thread

#### Parameters

in	<i>context</i>	Handle to <a href="#">UCP application context</a> .
in	<i>params</i>	User defined <a href="#">ucp_worker_params_t</a> configurations for the <a href="#">UCP worker</a> .
out	<i>worker↔ _p</i>	A pointer to the worker object allocated by the UCP library

#### Returns

Error code as defined by [ucs\\_status\\_t](#)

#### Examples

[ucp\\_client\\_server.c](#), and [ucp\\_hello\\_world.c](#).

### 6.3.5.2 ucp\_worker\_destroy()

```
void ucp_worker_destroy (
    ucp_worker_h worker )
```

This routine releases the resources associated with a [UCP worker](#).

**Warning**

Once the UCP worker destroy the worker handle cannot be used with any UCP routine.

The destroy process releases and shuts down all resources associated with the [worker](#).

## Parameters

in	<i>worker</i>	Worker object to destroy.
----	---------------	---------------------------

## Examples

[ucp\\_client\\_server.c](#), and [ucp\\_hello\\_world.c](#).

6.3.5.3 `ucp_worker_query()`

```
ucs_status_t ucp_worker_query (
    ucp_worker_h worker,
    ucp_worker_attr_t * attr )
```

This routine fetches information about the worker.

## Parameters

in	<i>worker</i>	Worker object to query.
out	<i>attr</i>	Filled with attributes of worker.

## Returns

Error code as defined by [ucs\\_status\\_t](#)

6.3.5.4 `ucp_worker_print_info()`

```
void ucp_worker_print_info (
    ucp_worker_h worker,
    FILE * stream )
```

This routine prints information about the protocols being used, thresholds, UCT transport methods, and other useful information associated with the worker.

## Parameters

in	<i>worker</i>	Worker object to print information for.
in	<i>stream</i>	Output stream to print the information to.

6.3.5.5 `ucp_worker_get_address()`

```
ucs_status_t ucp_worker_get_address (
    ucp_worker_h worker,
    ucp_address_t ** address_p,
    size_t * address_length_p )
```

This routine returns the address of the worker object. This address can be passed to remote instances of the UCP library in order to connect to this worker. The memory for the address handle is allocated by this function, and must be released by using [ucp\\_worker\\_release\\_address\(\)](#) routine.

## Parameters

in	<i>worker</i>	Worker object whose address to return.
out	<i>address_p</i>	A pointer to the worker address.
out	<i>address_length↔ _p</i>	The size in bytes of the address.

## Returns

Error code as defined by [ucs\\_status\\_t](#)

## Examples

[ucp\\_hello\\_world.c](#).

6.3.5.6 `ucp_worker_release_address()`

```
void ucp_worker_release_address (
    ucp_worker_h worker,
    ucp_address_t * address )
```

This routine release an [address handle](#) associated within the [worker](#) object.

## Warning

Once the address released the address handle cannot be used with any UCP routine.

## Parameters

in	<i>worker</i>	Worker object that is associated with the address object.
in	<i>address</i>	Address to release; the address object has to be allocated using <a href="#">ucp_worker_get_address()</a> routine.

## Examples

[ucp\\_hello\\_world.c](#).

6.3.5.7 `ucp_worker_address_query()`

```
ucs_status_t ucp_worker_address_query (
    ucp_address_t * address,
    ucp_worker_address_attr_t * attr )
```

This routine fetches information about the worker address. The address can be either of local or remote worker.

## Parameters

in	<i>address</i>	Worker address to query.
out	<i>attr</i>	Filled with attributes of the worker address.

**Returns**

Error code as defined by [ucs\\_status\\_t](#).

**6.3.5.8 ucp\_worker\_progress()**

```
unsigned ucp_worker_progress (
    ucp_worker_h worker )
```

This routine explicitly progresses all communication operations on a worker.

**Note**

- Typically, request wait and test routines call [this routine](#) to progress any outstanding operations.
- Transport layers, implementing asynchronous progress using threads, require callbacks and other user code to be thread safe.
- The state of communication can be advanced (progressed) by blocking routines. Nevertheless, the non-blocking routines can not be used for communication progress.

**Parameters**

in	<i>worker</i>	Worker to progress.
----	---------------	---------------------

**Returns**

Non-zero if any communication was progressed, zero otherwise.

**Examples**

[ucp\\_client\\_server.c](#), and [ucp\\_hello\\_world.c](#).

**6.3.5.9 ucp\_stream\_worker\_poll()**

```
ssize_t ucp_stream_worker_poll (
    ucp_worker_h worker,
    ucp_stream_poll_ep_t * poll_eps,
    size_t max_eps,
    unsigned flags )
```

This non-blocking routine returns endpoints on a worker which are ready to consume streaming data. The ready endpoints are placed in *poll\_eps* array, and the function return value indicates how many are there.

**Parameters**

in	<i>worker</i>	Worker to poll.
out	<i>poll_eps</i>	Pointer to array of endpoints, should be allocated by user.
in	<i>max_eps</i>	Maximum number of endpoints that should be filled in <i>poll_eps</i> .
in	<i>flags</i>	Reserved for future use.

**Returns**

Negative value indicates an error according to [ucs\\_status\\_t](#). On success, non-negative value (less or equal *max\_eps*) indicates actual number of endpoints filled in *poll\_eps* array.

**6.3.5.10 ucp\_listener\_create()**

```
ucs_status_t ucp_listener_create (
    ucp_worker_h worker,
    const ucp_listener_params_t * params,
    ucp_listener_h * listener_p )
```

This routine creates a new listener object that is bound to a specific local address. The listener will listen to incoming connection requests. After receiving a request from the remote peer, an endpoint to this peer will be created - either right away or by calling [ucp\\_ep\\_create](#), as specified by the callback type in [ucp\\_listener\\_params\\_t](#). The user's callback will be invoked once the endpoint is created.

**Parameters**

in	<i>worker</i>	Worker object to create the listener on.
in	<i>params</i>	User defined <a href="#">ucp_listener_params_t</a> configurations for the <a href="#">ucp_listener_h</a> .
out	<i>listener↔ _p</i>	A handle to the created listener, can be released by calling <a href="#">ucp_listener_destroy</a>

**Returns**

Error code as defined by [ucs\\_status\\_t](#)

**Note**

[ucp\\_listener\\_params\\_t::conn\\_handler](#) or [ucp\\_listener\\_params\\_t::accept\\_handler](#) must be provided to be able to handle incoming connections.

**Examples**

[ucp\\_client\\_server.c](#).

**6.3.5.11 ucp\_listener\_destroy()**

```
void ucp_listener_destroy (
    ucp_listener_h listener )
```

This routine unbinds the worker from the given handle and stops listening for incoming connection requests on it.

**Parameters**

in	<i>listener</i>	A handle to the listener to stop listening on.
----	-----------------	--

**Examples**

[ucp\\_client\\_server.c](#).

### 6.3.5.12 ucp\_listener\_query()

```
ucs_status_t ucp_listener_query (
    ucp_listener_h listener,
    ucp_listener_attr_t * attr )
```

This routine fetches information about the listener.

#### Parameters

in	<i>listener</i>	listener object to query.
out	<i>attr</i>	Filled with attributes of the listener.

#### Returns

Error code as defined by [ucs\\_status\\_t](#)

#### Examples

[ucp\\_client\\_server.c](#).

### 6.3.5.13 ucp\_conn\_request\_query()

```
ucs_status_t ucp_conn_request_query (
    ucp_conn_request_h conn_request,
    ucp_conn_request_attr_t * attr )
```

This routine fetches information about the connection request.

#### Parameters

in	<i>conn_request</i>	connection request object to query.
out	<i>attr</i>	Filled with attributes of the connection request.

#### Returns

Error code as defined by [ucs\\_status\\_t](#)

#### Examples

[ucp\\_client\\_server.c](#).

### 6.3.5.14 ucp\_listener\_reject()

```
ucs_status_t ucp_listener_reject (
    ucp_listener_h listener,
    ucp_conn_request_h conn_request )
```

Reject the incoming connection request and release associated resources. If the remote initiator endpoint has set an [ucp\\_ep\\_params\\_t::err\\_handler](#), it will be invoked with status [UCS\\_ERR\\_REJECTED](#).



## Parameters

in	<i>listener</i>	Handle to the listener on which the connection request was received.
in	<i>conn_request</i>	Handle to the connection request to reject.

## Returns

Error code as defined by [ucs\\_status\\_t](#)

## Examples

[ucp\\_client\\_server.c](#).

6.3.5.15 `ucp_worker_set_am_handler()`

```
ucs_status_t ucp_worker_set_am_handler (
    ucp_worker_h worker,
    uint16_t id,
    ucp_am_callback_t cb,
    void * arg,
    uint32_t flags )
```

This routine installs a user defined callback to handle incoming Active Messages with a specific id. This callback is called whenever an Active Message that was sent from the remote peer by [ucp\\_am\\_send\\_nb](#) is received on this worker.

## Parameters

in	<i>worker</i>	UCP worker on which to set the Active Message handler.
in	<i>id</i>	Active Message id.
in	<i>cb</i>	Active Message callback. NULL to clear.
in	<i>arg</i>	Active Message argument, which will be passed in to every invocation of the callback as the arg argument.
in	<i>flags</i>	Dictates how an Active Message is handled on the remote endpoint. Currently only UCP_AM_FLAG_WHOLE_MSG is supported, which indicates the callback will not be invoked until all data has arrived.

## Returns

error code if the worker does not support Active Messages or requested callback flags.

6.3.5.16 `ucp_worker_set_am_recv_handler()`

```
ucs_status_t ucp_worker_set_am_recv_handler (
    ucp_worker_h worker,
    const ucp_am_handler_param_t * param )
```

This routine installs a user defined callback to handle incoming Active Messages with a specific id. This callback is called whenever an Active Message that was sent from the remote peer by [ucp\\_am\\_send\\_nbx](#) is received on this worker.

**Warning**

Handlers set by this function are not compatible with [ucp\\_am\\_send\\_nb](#) routine.

**Parameters**

in	<i>worker</i>	UCP worker on which to set the Active Message handler.
in	<i>param</i>	Active Message handler parameters, as defined by <a href="#">ucp_am_handler_param_t</a> .

**Returns**

error code if the worker does not support Active Messages or requested callback flags.

**Examples**

[ucp\\_client\\_server.c](#).

**6.3.5.17 ucp\_worker\_fence()**

```
ucs_status_t ucp_worker_fence (
    ucp_worker_h worker )
```

This routine ensures ordering of non-blocking communication operations on the [UCP worker](#). Communication operations issued on the *worker* prior to this call are guaranteed to be completed before any subsequent communication operations to the same [worker](#) which follow the call to [fence](#).

**Note**

The primary difference between [ucp\\_worker\\_fence\(\)](#) and the [ucp\\_worker\\_flush\\_nb\(\)](#) is the fact the fence routine does not guarantee completion of the operations on the call return but only ensures the order between communication operations. The [flush](#) operation on return guarantees that all operations are completed and corresponding memory regions were updated.

**Parameters**

in	<i>worker</i>	UCP worker.
----	---------------	-------------

**Returns**

Error code as defined by [ucs\\_status\\_t](#)

**6.3.5.18 ucp\_worker\_flush\_nb()**

```
ucs_status_ptr_t ucp_worker_flush_nb (
    ucp_worker_h worker,
    unsigned flags,
    ucp_send_callback_t cb )
```

This routine flushes all outstanding AMO and RMA communications on the [worker](#). All the AMO and RMA operations issued on the *worker* prior to this call are completed both at the origin and at the target when this call returns.

**Note**

For description of the differences between [flush](#) and [fence](#) operations please see [ucp\\_worker\\_fence\(\)](#)

**Parameters**

in	<i>worker</i>	UCP worker.
in	<i>flags</i>	Flags for flush operation. Reserved for future use.
in	<i>cb</i>	Callback which will be called when the flush operation completes.

**Returns**

NULL - The flush operation was completed immediately.

UCS\_PTR\_IS\_ERR(\_ptr) - The flush operation failed.

otherwise - Flush operation was scheduled and can be completed in any point in time. The request handle is returned to the application in order to track progress. The application is responsible for releasing the handle using [ucp\\_request\\_free\(\)](#) routine.

**6.3.5.19 ucp\_worker\_flush\_nbx()**

```
ucs_status_ptr_t ucp_worker_flush_nbx (
    ucp_worker_h worker,
    const ucp_request_param_t * param )
```

This routine flushes all outstanding AMO and RMA communications on the [worker](#). All the AMO and RMA operations issued on the *worker* prior to this call are completed both at the origin and at the target when this call returns.

**Note**

For description of the differences between [flush](#) and [fence](#) operations please see [ucp\\_worker\\_fence\(\)](#)

**Parameters**

in	<i>worker</i>	UCP worker.
in	<i>param</i>	Operation parameters, see <a href="#">ucp_request_param_t</a>

**Returns**

NULL - The flush operation was completed immediately.

UCS\_PTR\_IS\_ERR(\_ptr) - The flush operation failed.

otherwise - Flush operation was scheduled and can be completed in any point in time. The request handle is returned to the application in order to track progress.

**6.3.5.20 ucp\_worker\_flush()**

```
ucs_status_t ucp_worker_flush (
    ucp_worker_h worker )
```

**Deprecated** Replaced by [ucp\\_worker\\_flush\\_nb](#). The following example implements the same functionality using [ucp\\_worker\\_flush\\_nb](#) :

```
ucs_status_t worker_flush(ucp_worker_h worker)
{
    void *request = ucp_worker_flush_nb(worker);
    if (request == NULL) {
        return UCS_OK;
    } else if (UCS_PTR_IS_ERR(request)) {
        return UCS_PTR_STATUS(request);
    } else {
        ucs_status_t status;
        do {
            ucp_worker_progress(worker);
            status = ucp_request_check_status(request);
        } while (status == UCS_INPROGRESS);
        ucp_request_release(request);
        return status;
    }
}
```

This routine flushes all outstanding AMO and RMA communications on the [worker](#). All the AMO and RMA operations issued on the *worker* prior to this call are completed both at the origin and at the target when this call returns.

#### Note

For description of the differences between [flush](#) and [fence](#) operations please see [ucp\\_worker\\_fence\(\)](#)

#### Parameters

in	<i>worker</i>	UCP worker.
----	---------------	-------------

#### Returns

Error code as defined by [ucs\\_status\\_t](#)

## 6.4 UCP Memory routines

### Data Structures

- struct `ucp_mem_map_params`  
*Tuning parameters for the UCP memory mapping. [More...](#)*
- struct `ucp_mem_advise_params`  
*Tuning parameters for the UCP memory advice. [More...](#)*
- struct `ucp_mem_attr`  
*Attributes of the UCP Memory handle, filled by `ucp_mem_query` function. [More...](#)*

### Typedefs

- typedef struct `ucp_mem_map_params` `ucp_mem_map_params_t`  
*Tuning parameters for the UCP memory mapping.*
- typedef enum `ucp_mem_advise` `ucp_mem_advise_t`  
*list of UCP memory use advice.*
- typedef struct `ucp_mem_advise_params` `ucp_mem_advise_params_t`  
*Tuning parameters for the UCP memory advice.*
- typedef struct `ucp_rkey` \* `ucp_rkey_h`  
*UCP Remote memory handle.*
- typedef struct `ucp_mem` \* `ucp_mem_h`  
*UCP Memory handle.*
- typedef struct `ucp_mem_attr` `ucp_mem_attr_t`  
*Attributes of the UCP Memory handle, filled by `ucp_mem_query` function.*

### Enumerations

- enum `ucp_mem_map_params_field` {  
`UCP_MEM_MAP_PARAM_FIELD_ADDRESS` = UCS\_BIT(0), `UCP_MEM_MAP_PARAM_FIELD_LENGTH`  
= UCS\_BIT(1), `UCP_MEM_MAP_PARAM_FIELD_FLAGS` = UCS\_BIT(2), `UCP_MEM_MAP_PARAM_FIELD_PROT`  
= UCS\_BIT(3),  
`UCP_MEM_MAP_PARAM_FIELD_MEMORY_TYPE` = UCS\_BIT(4) }  
*UCP memory mapping parameters field mask.*
- enum `ucp_mem_advise_params_field` { `UCP_MEM_ADVISE_PARAM_FIELD_ADDRESS` = UCS\_BIT(0),  
`UCP_MEM_ADVISE_PARAM_FIELD_LENGTH` = UCS\_BIT(1), `UCP_MEM_ADVISE_PARAM_FIELD_ADVICE`  
= UCS\_BIT(2) }  
*UCP memory advice parameters field mask.*
- enum { `UCP_MEM_MAP_NONBLOCK` = UCS\_BIT(0), `UCP_MEM_MAP_ALLOCATE` = UCS\_BIT(1),  
`UCP_MEM_MAP_FIXED` = UCS\_BIT(2) }  
*UCP memory mapping flags.*
- enum { `UCP_MEM_MAP_PROT_LOCAL_READ` = UCS\_BIT(0), `UCP_MEM_MAP_PROT_LOCAL_WRITE`  
= UCS\_BIT(1), `UCP_MEM_MAP_PROT_REMOTE_READ` = UCS\_BIT(8), `UCP_MEM_MAP_PROT_REMOTE_WRITE`  
= UCS\_BIT(9) }  
*UCP memory mapping protection mode.*
- enum `ucp_mem_advise` { `UCP_MADV_NORMAL` = 0, `UCP_MADV_WILLNEED` }  
*list of UCP memory use advice.*
- enum `ucp_mem_attr_field` { `UCP_MEM_ATTR_FIELD_ADDRESS` = UCS\_BIT(0), `UCP_MEM_ATTR_FIELD_LENGTH`  
= UCS\_BIT(1), `UCP_MEM_ATTR_FIELD_MEM_TYPE` = UCS\_BIT(2) }  
*UCP Memory handle attributes field mask.*

## Functions

- `ucs_status_t ucp_mem_map` (`ucp_context_h` context, const `ucp_mem_map_params_t` \*params, `ucp_mem_h` \*memh\_p)  
*Map or allocate memory for zero-copy operations.*
- `ucs_status_t ucp_mem_unmap` (`ucp_context_h` context, `ucp_mem_h` memh)  
*Unmap memory segment.*
- `ucs_status_t ucp_mem_query` (const `ucp_mem_h` memh, `ucp_mem_attr_t` \*attr)  
*query mapped memory segment*
- void `ucp_mem_print_info` (const char \*mem\_size, `ucp_context_h` context, FILE \*stream)  
*Print memory mapping information.*
- `ucs_status_t ucp_mem_advise` (`ucp_context_h` context, `ucp_mem_h` memh, `ucp_mem_advise_params_t` \*params)  
*give advice about the use of memory*
- `ucs_status_t ucp_rkey_pack` (`ucp_context_h` context, `ucp_mem_h` memh, void \*\*rkey\_buffer\_p, `size_t` \*size\_p)  
*Pack memory region remote access key.*
- void `ucp_rkey_buffer_release` (void \*rkey\_buffer)  
*Release packed remote key buffer.*
- `ucs_status_t ucp_ep_rkey_unpack` (`ucp_ep_h` ep, const void \*rkey\_buffer, `ucp_rkey_h` \*rkey\_p)  
*Create remote access key from packed buffer.*
- `ucs_status_t ucp_rkey_ptr` (`ucp_rkey_h` rkey, `uint64_t` raddr, void \*\*addr\_p)  
*Get a local pointer to remote memory.*
- void `ucp_rkey_destroy` (`ucp_rkey_h` rkey)  
*Destroy the remote key.*

### 6.4.1 Detailed Description

UCP Memory routines

### 6.4.2 Data Structure Documentation

#### 6.4.2.1 struct ucp\_mem\_map\_params

The structure defines the parameters that are used for the UCP memory mapping tuning during the `ucp_mem_map` routine.

##### Data Fields

<code>uint64_t</code>	<code>field_mask</code>	Mask of valid fields in this structure, using bits from <code>ucp_mem_map_params_field</code> . Fields not specified in this mask will be ignored. Provides ABI compatibility with respect to adding new fields.
<code>void *</code>	<code>address</code>	If the address is not NULL, the routine maps (registers) the memory segment pointed to by this address. If the pointer is NULL, the library allocates mapped (registered) memory segment and returns its address in this argument. Therefore, this value is optional. If it's not set (along with its corresponding bit in the <code>field_mask</code> - <code>UCP_MEM_MAP_PARAM_FIELD_ADDRESS</code> ), the <code>ucp_mem_map</code> routine will consider address as set to NULL and will allocate memory.
<code>size_t</code>	<code>length</code>	Length (in bytes) to allocate or map (register). This field is mandatory for filling (along with its corresponding bit in the <code>field_mask</code> - <code>UCP_MEM_MAP_PARAM_FIELD_LENGTH</code> ). The <code>ucp_mem_map</code> routine will return with an error if the length isn't specified.

## Data Fields

unsigned	flags	Allocation flags, e.g. <a href="#">UCP_MEM_MAP_NONBLOCK</a> . This value is optional. If it's not set (along with its corresponding bit in the <code>field_mask</code> - <a href="#">UCP_MEM_MAP_PARAM_FIELD_FLAGS</a> ), the <code>ucp_mem_map</code> routine will consider the flags as set to zero.
unsigned	prot	Memory protection mode, e.g. <a href="#">UCP_MEM_MAP_PROT_LOCAL_READ</a> . This value is optional. If it's not set, the <code>ucp_mem_map</code> routine will consider the flags as set to <code>UCP_MEM_MAP_PROT_LOCAL_READ UCP_MEM_MAP_PROT_LOCAL_WRITE UCP_MEM_MAP_PROT_REMOTE_READ UCP_MEM_MAP_PROT_REMOTE_WRITE</code> .
<a href="#">ucs_memory_type_t</a>	memory_type	

## 6.4.2.2 struct ucp\_mem\_advise\_params

This structure defines the parameters that are used for the UCP memory advice tuning during the `ucp_mem_advise` routine.

## Data Fields

uint64_t	field_mask	Mask of valid fields in this structure, using bits from <a href="#">ucp_mem_advise_params_field</a> . All fields are mandatory. Provides ABI compatibility with respect to adding new fields.
void *	address	Memory base address.
size_t	length	Length (in bytes) to allocate or map (register).
<a href="#">ucp_mem_advise_t</a>	advise	Memory use advice <a href="#">ucp_mem_advise</a>

## 6.4.2.3 struct ucp\_mem\_attr

## Data Fields

uint64_t	field_mask	Mask of valid fields in this structure, using bits from <a href="#">ucp_mem_attr_field</a> . Fields not specified in this mask will be ignored. Provides ABI compatibility with respect to adding new fields.
void *	address	Address of the memory segment.
size_t	length	Size of the memory segment.
<a href="#">ucs_memory_type_t</a>	mem_type	Type of allocated or registered memory

## 6.4.3 Typedef Documentation

## 6.4.3.1 ucp\_mem\_map\_params\_t

```
typedef struct ucp_mem_map_params ucp_mem_map_params_t
```

The structure defines the parameters that are used for the UCP memory mapping tuning during the `ucp_mem_map` routine.

### 6.4.3.2 ucp\_mem\_advise\_t

```
typedef enum ucp_mem_advise ucp_mem_advise_t
```

The enumeration list describes memory advice supported by `ucp_mem_advise()` function.

### 6.4.3.3 ucp\_mem\_advise\_params\_t

```
typedef struct ucp_mem_advise_params ucp_mem_advise_params_t
```

This structure defines the parameters that are used for the UCP memory advice tuning during the `ucp_mem_advise` routine.

### 6.4.3.4 ucp\_rkey\_h

```
typedef struct ucp_rkey* ucp_rkey_h
```

Remote memory handle is an opaque object representing remote memory access information. Typically, the handle includes a memory access key and other network hardware specific information, which are input to remote memory access operations, such as PUT, GET, and ATOMIC. The object is communicated to remote peers to enable an access to the memory region.

### 6.4.3.5 ucp\_mem\_h

```
typedef struct ucp_mem* ucp_mem_h
```

Memory handle is an opaque object representing a memory region allocated through UCP library, which is optimized for remote memory access operations (zero-copy operations). The memory handle is a self-contained object, which includes the information required to access the memory region locally, while `remote key` is used to access it remotely. The memory could be registered to one or multiple network resources that are supported by UCP, such as InfiniBand, Gemini, and others.

### 6.4.3.6 ucp\_mem\_attr\_t

```
typedef struct ucp_mem_attr ucp_mem_attr_t
```

## 6.4.4 Enumeration Type Documentation

### 6.4.4.1 ucp\_mem\_map\_params\_field

```
enum ucp_mem_map_params_field
```

The enumeration allows specifying which fields in `ucp_mem_map_params_t` are present. It is used to enable backward compatibility support.

#### Enumerator

UCP_MEM_MAP_PARAM_FIELD_ADDRESS	Address of the memory that will be used in the <code>ucp_mem_map</code> routine.
UCP_MEM_MAP_PARAM_FIELD_LENGTH	The size of memory that will be allocated or registered in the <code>ucp_mem_map</code> routine.
UCP_MEM_MAP_PARAM_FIELD_FLAGS	Allocation flags.
UCP_MEM_MAP_PARAM_FIELD_PROT	Memory protection mode.
UCP_MEM_MAP_PARAM_FIELD_MEMORY_TYPE	Memory type.



#### 6.4.4.2 ucp\_mem\_advise\_params\_field

enum `ucp_mem_advise_params_field`

The enumeration allows specifying which fields in `ucp_mem_advise_params_t` are present. It is used to enable backward compatibility support.

Enumerator

UCP_MEM_ADVISE_PARAM_FIELD_ADDRESS	Address of the memory
UCP_MEM_ADVISE_PARAM_FIELD_LENGTH	The size of memory
UCP_MEM_ADVISE_PARAM_FIELD_ADVICE	Advice on memory usage

#### 6.4.4.3 anonymous enum

anonymous enum

The enumeration list describes the memory mapping flags supported by `ucp_mem_map()` function.

Enumerator

UCP_MEM_MAP_NONBLOCK	Complete the mapping faster, possibly by not populating the pages in the mapping up-front, and mapping them later when they are accessed by communication routines.
UCP_MEM_MAP_ALLOCATE	Identify requirement for allocation, if passed address is not a null-pointer then it will be used as a hint or direct address for allocation.
UCP_MEM_MAP_FIXED	Don't interpret address as a hint: place the mapping at exactly that address. The address must be a multiple of the page size.

#### 6.4.4.4 anonymous enum

anonymous enum

The enumeration list describes the memory mapping protections supported by the `ucp_mem_map()` function.

Enumerator

UCP_MEM_MAP_PROT_LOCAL_READ	Enable local read access.
UCP_MEM_MAP_PROT_LOCAL_WRITE	Enable local write access.
UCP_MEM_MAP_PROT_REMOTE_READ	Enable remote read access.
UCP_MEM_MAP_PROT_REMOTE_WRITE	Enable remote write access.

#### 6.4.4.5 ucp\_mem\_advice

enum `ucp_mem_advice`

The enumeration list describes memory advice supported by `ucp_mem_advise()` function.

#### Enumerator

UCP_MADV_NORMAL	No special treatment
UCP_MADV_WILLNEED	can be used on the memory mapped with <a href="#">UCP_MEM_MAP_NONBLOCK</a> to speed up memory mapping and to avoid page faults when the memory is accessed for the first time.

#### 6.4.4.6 ucp\_mem\_attr\_field

```
enum ucp_mem_attr_field
```

The enumeration allows specifying which fields in `ucp_mem_attr_t` are present. It is used to enable backward compatibility support.

#### Enumerator

UCP_MEM_ATTR_FIELD_ADDRESS	Virtual address
UCP_MEM_ATTR_FIELD_LENGTH	The size of memory region
UCP_MEM_ATTR_FIELD_MEM_TYPE	Type of allocated or registered memory

### 6.4.5 Function Documentation

#### 6.4.5.1 ucp\_mem\_map()

```
ucs_status_t ucp_mem_map (
    ucp_context_h context,
    const ucp_mem_map_params_t * params,
    ucp_mem_h * memh_p )
```

This routine maps or/and allocates a user-specified memory segment with [UCP application context](#) and the network resources associated with it. If the application specifies NULL as an address for the memory segment, the routine allocates a mapped memory segment and returns its address in the `address_p` argument. The network stack associated with an application context can typically send and receive data from the mapped memory without CPU intervention; some devices and associated network stacks require the memory to be mapped to send and receive data. The [memory handle](#) includes all information required to access the memory locally using UCP routines, while [remote registration handle](#) provides an information that is necessary for remote memory access.

#### Note

Another well know terminology for the "map" operation that is typically used in the context of networking is memory "registration" or "pinning". The UCP library registers the memory the available hardware so it can be assessed directly by the hardware.

Memory mapping assumptions:

- A given memory segment can be mapped by several different communication stacks, if these are compatible.
- The `memh_p` handle returned may be used with any sub-region of the mapped memory.

- If a large segment is registered, and then segmented for subsequent use by a user, then the user is responsible for segmentation and subsequent management.

Table 6.71: Matrix of behavior

parameter/flag	NONBLOCK	ALLOCATE	FIXED	address	result
value	0/1 - the value only affects the register/map phase	0	0	0	error if length > 0
		1	0	0	alloc+register
		0	1	0	error
		0	0	defined	register
		1	1	0	error
		1	0	defined	alloc+register,hint
		0	1	defined	error
		1	1	defined	alloc+register,fixed

#### Note

- **register** means that the memory will be registered in corresponding transports for RMA/AMO operations. This case intends that the memory was allocated by user before.
- **alloc+register** means that the memory will be allocated in the memory provided by the system and registered in corresponding transports for RMA/AMO operations.
- **alloc+register,hint** means that the memory will be allocated with using `ucp_mem_map_params::address` as a hint and registered in corresponding transports for RMA/AMO operations.
- **alloc+register,fixed** means that the memory will be allocated and registered in corresponding transports for RMA/AMO operations.
- **error** is an erroneous combination of the parameters.

#### Parameters

in	<i>context</i>	Application <a href="#">context</a> to map (register) and allocate the memory on.
in	<i>params</i>	User defined <code>ucp_mem_map_params_t</code> configurations for the <a href="#">UCP memory handle</a> .
out	<i>memh</i> <i>_p</i>	UCP <a href="#">handle</a> for the allocated segment.

#### Returns

Error code as defined by [ucs\\_status\\_t](#)

#### 6.4.5.2 ucp\_mem\_unmap()

```
ucs_status_t ucp_mem_unmap (
    ucp_context_h context,
    ucp_mem_h memh )
```

This routine unmaps a user specified memory segment, that was previously mapped using the `ucp_mem_map()` routine. The unmap routine will also release the resources associated with the memory [handle](#). When the function returns, the `ucp_mem_h` and associated [remote key](#) will be invalid and cannot be used with any UCP routine.

**Note**

Another well know terminology for the "unmap" operation that is typically used in the context of networking is memory "de-registration". The UCP library de-registers the memory the available hardware so it can be returned back to the operation system.

**Error cases:**

- Once memory is unmapped a network access to the region may cause a failure.

**Parameters**

in	<i>context</i>	Application <a href="#">context</a> which was used to allocate/map the memory.
in	<i>memh</i>	<a href="#">Handle</a> to memory region.

**Returns**

Error code as defined by [ucs\\_status\\_t](#)

**6.4.5.3 ucp\_mem\_query()**

```
ucs_status_t ucp_mem_query (
    const ucp_mem_h memh,
    ucp_mem_attr_t * attr )
```

This routine returns address and length of memory segment mapped with [ucp\\_mem\\_map\(\)](#) routine.

**Parameters**

in	<i>memh</i>	<a href="#">Handle</a> to memory region.
out	<i>attr</i>	Filled with attributes of the <a href="#">UCP memory handle</a> .

**Returns**

Error code as defined by [ucs\\_status\\_t](#)

**6.4.5.4 ucp\_mem\_print\_info()**

```
void ucp_mem_print_info (
    const char * mem_size,
    ucp_context_h context,
    FILE * stream )
```

This routine maps memory and prints information about the created memory handle: including the mapped memory length, the allocation method, and other useful information associated with the memory handle.

**Parameters**

in	<i>mem_size</i>	Size of the memory to map.
in	<i>context</i>	The context on which the memory is mapped.
in	<i>stream</i>	Output stream on which to print the information.

### 6.4.5.5 ucp\_mem\_advise()

```
ucs_status_t ucp_mem_advise (
    ucp_context_h context,
    ucp_mem_h memh,
    ucp_mem_advise_params_t * params )
```

This routine advises the UCP about how to handle memory range beginning at address and size of length bytes. This call does not influence the semantics of the application, but may influence its performance. The UCP may ignore the advice.

#### Parameters

in	<i>context</i>	Application <a href="#">context</a> which was used to allocate/map the memory.
in	<i>memh</i>	<a href="#">Handle</a> to memory region.
in	<i>params</i>	Memory base address and length. The advice field is used to pass memory use advice as defined in the <a href="#">ucp_mem_advise</a> list The memory range must belong to the <i>memh</i>

#### Returns

Error code as defined by [ucs\\_status\\_t](#)

### 6.4.5.6 ucp\_rkey\_pack()

```
ucs_status_t ucp_rkey_pack (
    ucp_context_h context,
    ucp_mem_h memh,
    void ** rkey_buffer_p,
    size_t * size_p )
```

This routine allocates memory buffer and packs into the buffer a remote access key (RKEY) object. RKEY is an opaque object that provides the information that is necessary for remote memory access. This routine packs the RKEY object in a portable format such that the object can be [unpacked](#) on any platform supported by the UCP library. In order to release the memory buffer allocated by this routine the application is responsible for calling the [ucp\\_rkey\\_buffer\\_release\(\)](#) routine.

#### Note

- RKEYs for InfiniBand and Cray Aries networks typically includes InfiniBand and Aries key.
- In order to enable remote direct memory access to the memory associated with the memory handle the application is responsible for sharing the RKEY with the peers that will initiate the access.

#### Parameters

in	<i>context</i>	Application <a href="#">context</a> which was used to allocate/map the memory.
in	<i>memh</i>	<a href="#">Handle</a> to memory region.
out	<i>rkey_buffer↔ _p</i>	Memory buffer allocated by the library. The buffer contains packed RKEY.
out	<i>size_p</i>	Size (in bytes) of the packed RKEY.

**Returns**

Error code as defined by [ucs\\_status\\_t](#)

**6.4.5.7 ucp\_rkey\_buffer\_release()**

```
void ucp_rkey_buffer_release (
    void * rkey_buffer )
```

This routine releases the buffer that was allocated using [ucp\\_rkey\\_pack\(\)](#).

**Warning**

- Once memory is released an access to the memory may cause a failure.
- If the input memory address was not allocated using [ucp\\_rkey\\_pack\(\)](#) routine the behaviour of this routine is undefined.

**Parameters**

in	<i>rkey_buffer</i>	Buffer to release.
----	--------------------	--------------------

**6.4.5.8 ucp\_ep\_rkey\_unpack()**

```
ucs_status_t ucp_ep_rkey_unpack (
    ucp_ep_h ep,
    const void * rkey_buffer,
    ucp_rkey_h * rkey_p )
```

This routine unpacks the remote key (RKEY) object into the local memory such that it can be accessed and used by UCP routines. The RKEY object has to be packed using the [ucp\\_rkey\\_pack\(\)](#) routine. Application code should not make any changes to the content of the RKEY buffer.

**Note**

The application is responsible for releasing the RKEY object when it is no longer needed, by calling the [ucp\\_rkey\\_destroy\(\)](#) routine.

The remote key object can be used for communications only on the endpoint on which it was unpacked.

**Parameters**

in	<i>ep</i>	Endpoint to access using the remote key.
in	<i>rkey_buffer</i>	Packed rkey.
out	<i>rkey_p</i>	Remote key handle.

**Returns**

Error code as defined by [ucs\\_status\\_t](#)

**6.4.5.9 ucp\_rkey\_ptr()**

```
ucs_status_t ucp_rkey_ptr (
    ucp_rkey_h rkey,
    uint64_t raddr,
    void ** addr_p )
```

This routine returns a local pointer to the remote memory described by the rkey.

**Note**

This routine can return a valid pointer only for the endpoints that are reachable via shared memory.

**Parameters**

in	<i>rkey</i>	A remote key handle.
in	<i>raddr</i>	A remote memory address within the memory area described by the rkey.
out	<i>addr↔ _p</i>	A pointer that can be used for direct access to the remote memory.

**Returns**

Error code as defined by [ucs\\_status\\_t](#) if the remote memory cannot be accessed directly or the remote memory address is not valid.

**6.4.5.10 ucp\_rkey\_destroy()**

```
void ucp_rkey_destroy (
    ucp_rkey_h rkey )
```

This routine destroys the RKEY object and the memory that was allocated using the [ucp\\_ep\\_rkey\\_unpack\(\)](#) routine. This routine also releases any resources that are associated with the RKEY object.

**Warning**

- Once the RKEY object is released an access to the memory will cause an undefined failure.
- If the RKEY object was not created using [ucp\\_ep\\_rkey\\_unpack\(\)](#) routine the behavior of this routine is undefined.
- The RKEY object must be destroyed after all outstanding operations which are using it are flushed, and before the endpoint on which it was unpacked is destroyed.

**Parameters**

in	<i>rkey</i>	Remote key to destroy.
----	-------------	------------------------

## 6.5 UCP Wake-up routines

### Functions

- `ucs_status_t ucp_worker_get_efd` (`ucp_worker_h` worker, `int *fd`)  
*Obtain an event file descriptor for event notification.*
- `ucs_status_t ucp_worker_wait` (`ucp_worker_h` worker)  
*Wait for an event of the worker.*
- `void ucp_worker_wait_mem` (`ucp_worker_h` worker, `void *address`)  
*Wait for memory update on the address.*
- `ucs_status_t ucp_worker_arm` (`ucp_worker_h` worker)  
*Turn on event notification for the next event.*
- `ucs_status_t ucp_worker_signal` (`ucp_worker_h` worker)  
*Cause an event of the worker.*

### 6.5.1 Detailed Description

UCP Wake-up routines

### 6.5.2 Function Documentation

#### 6.5.2.1 `ucp_worker_get_efd()`

```
ucs_status_t ucp_worker_get_efd (
    ucp_worker_h worker,
    int * fd )
```

This routine returns a valid file descriptor for polling functions. The file descriptor will get signaled when an event occurs, as part of the wake-up mechanism. Signaling means a call to `poll()` or `select()` with this file descriptor will return at this point, with this descriptor marked as the reason (or one of the reasons) the function has returned. The user does not need to release the obtained file descriptor.

The wake-up mechanism exists to allow for the user process to register for notifications on events of the underlying interfaces, and wait until such occur. This is an alternative to repeated polling for request completion. The goal is to allow for waiting while consuming minimal resources from the system. This is recommended for cases where traffic is infrequent, and latency can be traded for lower resource consumption while waiting for it.

There are two alternative ways to use the wakeup mechanism: the first is the file descriptor obtained per worker (this function) and the second is the `ucp_worker_wait` function for waiting on the next event internally.

#### Note

UCP [features](#) have to be triggered with `UCP_FEATURE_WAKEUP` to select proper transport

#### Parameters

in	<i>worker</i>	Worker of notified events.
out	<i>fd</i>	File descriptor.



**Returns**

Error code as defined by [ucs\\_status\\_t](#)

**Examples**

[ucp\\_hello\\_world.c](#).

**6.5.2.2 ucp\_worker\_wait()**

```
ucs_status_t ucp_worker_wait (
    ucp_worker_h worker )
```

This routine waits (blocking) until an event has happened, as part of the wake-up mechanism.

This function is guaranteed to return only if new communication events occur on the *worker*. Therefore one must drain all existing events before waiting on the file descriptor. This can be achieved by calling [ucp\\_worker\\_progress](#) repeatedly until it returns 0.

There are two alternative ways to use the wakeup mechanism. The first is by polling on a per-worker file descriptor obtained from [ucp\\_worker\\_get\\_efd](#). The second is by using this function to perform an internal wait for the next event associated with the specified worker.

**Note**

During the blocking call the wake-up mechanism relies on other means of notification and may not progress some of the requests as it would when calling [ucp\\_worker\\_progress](#) (which is not invoked in that duration).

UCP [features](#) have to be triggered with [UCP\\_FEATURE\\_WAKEUP](#) to select proper transport

**Parameters**

in	<i>worker</i>	Worker to wait for events on.
----	---------------	-------------------------------

**Returns**

Error code as defined by [ucs\\_status\\_t](#)

**Examples**

[ucp\\_hello\\_world.c](#).

**6.5.2.3 ucp\_worker\_wait\_mem()**

```
void ucp_worker_wait_mem (
    ucp_worker_h worker,
    void * address )
```

This routine waits for a memory update at the local memory *address*. This is a blocking routine. The routine returns when the memory address is updated ("write") or an event occurs in the system.

This function is guaranteed to return only if new communication events occur on the *worker* or *address* is modified. Therefore one must drain all existing events before waiting on the file descriptor. This can be achieved by calling [ucp\\_worker\\_progress](#) repeatedly until it returns 0.

**Note**

This routine can be used by an application that executes busy-waiting loop checking for a memory update. Instead of continuous busy-waiting on an address the application can use `ucp_worker_wait_mem`, which may suspend execution until the memory is updated. The goal of the routine is to provide an opportunity for energy savings for architectures that support this functionality.

**Parameters**

in	<i>worker</i>	Worker to wait for updates on.
in	<i>address</i>	Local memory address

**6.5.2.4 ucp\_worker\_arm()**

```
ucs_status_t ucp_worker_arm (
    ucp_worker_h worker )
```

This routine needs to be called before waiting on each notification on this worker, so will typically be called once the processing of the previous event is over, as part of the wake-up mechanism.

The worker must be armed before waiting on an event (must be re-armed after it has been signaled for re-use) with `ucp_worker_arm`. The events triggering a signal of the file descriptor from `ucp_worker_get_efd` depend on the interfaces used by the worker and defined in the transport layer, and typically represent a request completion or newly available resources. It can also be triggered by calling `ucp_worker_signal`.

The file descriptor is guaranteed to become signaled only if new communication events occur on the *worker*. Therefore one must drain all existing events before waiting on the file descriptor. This can be achieved by calling `ucp_worker_progress` repeatedly until it returns 0.

```
void application_initialization() {
    // should be called once in application init flow and before
    // process_communication() is used
    ...
    status = ucp_worker_get_efd(worker, &fd);
    ...
}

void process_communication() {
    // should be called every time need to wait for some condition such as
    // ucp request completion in sleep mode.
    for (;;) {
        // check for stop condition as long as progress is made
        if (check_for_events()) {
            break;
        } else if (ucp_worker_progress(worker)) {
            continue; // some progress happened but condition not met
        }
        // arm the worker and clean-up fd
        status = ucp_worker_arm(worker);
        if (UCS_OK == status) {
            poll(&fds, nfds, timeout); // wait for events (sleep mode)
        } else if (UCS_ERR_BUSY == status) {
            continue; // could not arm, need to progress more
        } else {
            abort();
        }
    }
}
```

**Note**

UCP features have to be triggered with `UCP_FEATURE_WAKEUP` to select proper transport

**Parameters**

in	<i>worker</i>	Worker of notified events.
----	---------------	----------------------------

**Returns**

[UCS\\_OK](#) The operation completed successfully. File descriptor will be signaled by new events.

[UCS\\_ERR\\_BUSY](#) There are unprocessed events which prevent the file descriptor from being armed. These events should be removed by calling [ucp\\_worker\\_progress\(\)](#). The operation is not completed. File descriptor will not be signaled by new events.

**Other** different error codes in case of issues.

**Examples**

[ucp\\_hello\\_world.c](#).

**6.5.2.5 ucp\_worker\_signal()**

```
ucs_status_t ucp_worker_signal (
    ucp_worker_h worker )
```

This routine signals that the event has happened, as part of the wake-up mechanism. This function causes a blocking call to [ucp\\_worker\\_wait](#) or waiting on a file descriptor from [ucp\\_worker\\_get\\_efd](#) to return, even if no event from the underlying interfaces has taken place.

**Note**

It's safe to use this routine from any thread, even if UCX is compiled without multi-threading support and/or initialized with any value of [ucp\\_params\\_t::mt\\_workers\\_shared](#) and [ucp\\_worker\\_params\\_t::thread\\_mode](#) parameters

**Parameters**

in	<i>worker</i>	Worker to wait for events on.
----	---------------	-------------------------------

**Returns**

Error code as defined by [ucs\\_status\\_t](#)

## 6.6 UCP Endpoint

### Data Structures

- struct `ucp_ep_evaluate_perf_param_t`  
*UCP endpoint performance evaluation request attributes. [More...](#)*
- struct `ucp_ep_evaluate_perf_attr_t`  
*UCP endpoint performance evaluation result attributes. [More...](#)*
- struct `ucp_stream_poll_ep`  
*Output parameter of `ucp_stream_worker_poll` function. [More...](#)*
- struct `ucp_ep_attr`  
*UCP endpoint attributes. [More...](#)*
- struct `ucp_ep_params`  
*Tuning parameters for the UCP endpoint. [More...](#)*

### Typedefs

- typedef enum `ucp_ep_perf_param_field` `ucp_ep_perf_param_field_t`  
*UCP performance fields and flags.*
- typedef enum `ucp_ep_perf_attr_field` `ucp_ep_perf_attr_field_t`  
*UCP performance fields and flags.*
- typedef struct `ucp_stream_poll_ep` `ucp_stream_poll_ep_t`  
*Output parameter of `ucp_stream_worker_poll` function.*
- typedef struct `ucp_ep_attr` `ucp_ep_attr_t`  
*UCP endpoint attributes.*
- typedef struct `ucp_ep` \* `ucp_ep_h`  
*UCP Endpoint.*
- typedef struct `ucp_conn_request` \* `ucp_conn_request_h`  
*UCP connection request.*
- typedef `ucs_status_t`(\* `ucp_am_callback_t`) (void \*arg, void \*data, size\_t length, `ucp_ep_h` reply\_ep, unsigned flags)  
*Callback to process incoming Active Message.*
- typedef `ucs_status_t`(\* `ucp_am_rcv_callback_t`) (void \*arg, const void \*header, size\_t header\_length, void \*data, size\_t length, const `ucp_am_rcv_param_t` \*param)  
*Callback to process incoming Active Message sent by `ucp_am_send_nbx` routine.*
- typedef struct `ucp_ep_params` `ucp_ep_params_t`  
*Tuning parameters for the UCP endpoint.*

### Enumerations

- enum `ucp_ep_params_field` {  
`UCP_EP_PARAM_FIELD_REMOTE_ADDRESS` = UCS\_BIT(0), `UCP_EP_PARAM_FIELD_ERR_HANDLING_MODE`  
= UCS\_BIT(1), `UCP_EP_PARAM_FIELD_ERR_HANDLER` = UCS\_BIT(2), `UCP_EP_PARAM_FIELD_USER_DATA`  
= UCS\_BIT(3),  
`UCP_EP_PARAM_FIELD SOCK_ADDR` = UCS\_BIT(4), `UCP_EP_PARAM_FIELD_FLAGS` = UCS\_BIT(5),  
`UCP_EP_PARAM_FIELD_CONN_REQUEST` = UCS\_BIT(6), `UCP_EP_PARAM_FIELD_NAME` = UCS\_↔  
BIT(7),  
`UCP_EP_PARAM_FIELD_LOCAL SOCK_ADDR` = UCS\_BIT(8) }  
*UCP endpoint parameters field mask.*
- enum `ucp_ep_params_flags_field` { `UCP_EP_PARAMS_FLAGS_CLIENT_SERVER` = UCS\_BIT(0),  
`UCP_EP_PARAMS_FLAGS_NO_LOOPBACK` = UCS\_BIT(1), `UCP_EP_PARAMS_FLAGS_SEND_CLIENT_ID`  
= UCS\_BIT(2) }

*UCP endpoint parameters flags.*

- enum `ucp_ep_close_flags_t` { `UCP_EP_CLOSE_FLAG_FORCE` = `UCS_BIT(0)` }

*Close UCP endpoint modes.*

- enum `ucp_ep_close_mode` { `UCP_EP_CLOSE_MODE_FORCE` = 0, `UCP_EP_CLOSE_MODE_FLUSH` = 1 }

*Close UCP endpoint modes.*

- enum `ucp_ep_perf_param_field` { `UCP_EP_PERF_PARAM_FIELD_MESSAGE_SIZE` = `UCS_BIT(0)` }

*UCP performance fields and flags.*

- enum `ucp_ep_perf_attr_field` { `UCP_EP_PERF_ATTR_FIELD_ESTIMATED_TIME` = `UCS_BIT(0)` }

*UCP performance fields and flags.*

- enum `ucp_cb_param_flags` { `UCP_CB_PARAM_FLAG_DATA` = `UCS_BIT(0)` }

*Descriptor flags for Active Message callback.*

- enum `ucp_ep_attr_field` { `UCP_EP_ATTR_FIELD_NAME` = `UCS_BIT(0)`, `UCP_EP_ATTR_FIELD_LOCAL_SOCKADDR` = `UCS_BIT(1)`, `UCP_EP_ATTR_FIELD_REMOTE_SOCKADDR` = `UCS_BIT(2)` }

*UCP endpoint attributes field mask.*

- enum `ucp_err_handling_mode_t` { `UCP_ERR_HANDLING_MODE_NONE`, `UCP_ERR_HANDLING_MODE_PEER` }

*Error handling mode for the UCP endpoint.*

## Functions

- `ucs_status_t ucp_ep_create` (`ucp_worker_h` worker, const `ucp_ep_params_t` \*params, `ucp_ep_h` \*ep\_p)

*Create and connect an endpoint.*

- `ucs_status_ptr_t ucp_ep_close_nb` (`ucp_ep_h` ep, unsigned mode)

*Non-blocking endpoint closure.*

- `ucs_status_ptr_t ucp_ep_close_nbx` (`ucp_ep_h` ep, const `ucp_request_param_t` \*param)

*Non-blocking endpoint closure.*

- void `ucp_ep_print_info` (`ucp_ep_h` ep, FILE \*stream)

*Print endpoint information.*

- `ucs_status_ptr_t ucp_ep_flush_nb` (`ucp_ep_h` ep, unsigned flags, `ucp_send_callback_t` cb)

*Non-blocking flush of outstanding AMO and RMA operations on the endpoint.*

- `ucs_status_ptr_t ucp_ep_flush_nbx` (`ucp_ep_h` ep, const `ucp_request_param_t` \*param)

*Non-blocking flush of outstanding AMO and RMA operations on the endpoint.*

- `ucs_status_t ucp_ep_evaluate_perf` (`ucp_ep_h` ep, const `ucp_ep_evaluate_perf_param_t` \*param, `ucp_ep_evaluate_perf_attr_t` \*attr)

*Estimate performance characteristics of a specific endpoint.*

- `ucs_status_t ucp_ep_query` (`ucp_ep_h` ep, `ucp_ep_attr_t` \*attr)

*Get attributes of a given endpoint.*

- void `ucp_request_release` (void \*request)

- void `ucp_ep_destroy` (`ucp_ep_h` ep)

- `ucs_status_ptr_t ucp_disconnect_nb` (`ucp_ep_h` ep)

- `ucs_status_t ucp_request_test` (void \*request, `ucp_tag_rcv_info_t` \*info)

- `ucs_status_t ucp_ep_flush` (`ucp_ep_h` ep)

- `ucs_status_ptr_t ucp_ep_modify_nb` (`ucp_ep_h` ep, const `ucp_ep_params_t` \*params)

*Modify endpoint parameters.*

### 6.6.1 Detailed Description

UCP Endpoint routines

## 6.6.2 Data Structure Documentation

### 6.6.2.1 struct ucp\_ep\_evaluate\_perf\_param\_t

The structure defines the attributes which characterize the request for performance estimation of a particular endpoint.

#### Data Fields

uint64_t	field_mask	Mask of valid fields in this structure, using bits from <a href="#">ucp_ep_perf_param_field_t</a> . Fields not specified in this mask will be ignored. Provides ABI compatibility with respect to adding new fields.
size_t	message_size	Message size to use for determining performance. This field must be initialized by the caller.

### 6.6.2.2 struct ucp\_ep\_evaluate\_perf\_attr\_t

The structure defines the attributes which characterize the result of performance estimation of a particular endpoint.

#### Data Fields

uint64_t	field_mask	Mask of valid fields in this structure, using bits from <a href="#">ucp_ep_perf_attr_field_t</a> . Fields not specified in this mask will be ignored. Provides ABI compatibility with respect to adding new fields.
double	estimated_time	Estimated time (in seconds) required to send a message of a given size on this endpoint. This field is set by the <a href="#">ucp_ep_evaluate_perf</a> function.

### 6.6.2.3 struct ucp\_stream\_poll\_ep

The structure defines the endpoint and its user data.

#### Data Fields

<a href="#">ucp_ep_h</a>	ep	Endpoint handle.
void *	user_data	User data associated with an endpoint passed in <a href="#">ucp_ep_params_t::user_data</a> .
unsigned	flags	Reserved for future use.
uint8_t	reserved[16]	Reserved for future use.

### 6.6.2.4 struct ucp\_ep\_attr

The structure defines the attributes that characterize the particular endpoint.

#### Data Fields

uint64_t	field_mask	Mask of valid fields in this structure, using bits from <a href="#">ucp_ep_attr_field_t</a> . Fields not specified in this mask will be ignored. Provides ABI compatibility with respect to adding new fields.
char	name[UCP_ENTITY_NAME_MAX]	Endpoint name. Tracing and analysis tools can identify the endpoint using this name.

## Data Fields

struct sockaddr_storage	local_sockaddr	Local socket address for this endpoint. Valid only for endpoints created by connecting to a socket address. If this field is specified for an endpoint not connected to a socket address, UCS_ERR_NOT_CONNECTED will be returned.
struct sockaddr_storage	remote_sockaddr	Remote socket address this endpoint is connected to. Valid only for endpoints created by connecting to a socket address. If this field is specified for an endpoint not connected to a socket address, UCS_ERR_NOT_CONNECTED will be returned.

## 6.6.2.5 struct ucp\_ep\_params

The structure defines the parameters that are used for the UCP endpoint tuning during the UCP ep [creation](#).

## Examples

[ucp\\_client\\_server.c](#), and [ucp\\_hello\\_world.c](#).

## Data Fields

uint64_t	field_mask	Mask of valid fields in this structure, using bits from <a href="#">ucp_ep_params_field</a> . Fields not specified in this mask will be ignored. Provides ABI compatibility with respect to adding new fields.
const ucp_address_t *	address	Destination address; this field should be set along with its corresponding bit in the field_mask - <a href="#">UCP_EP_PARAM_FIELD_REMOTE_ADDRESS</a> and must be obtained using <a href="#">ucp_worker_get_address</a> .
ucp_err_handling_mode_t	err_mode	Desired error handling mode, optional parameter. Default value is <a href="#">UCP_ERR_HANDLING_MODE_NONE</a> .
ucp_err_handler_t	err_handler	Handler to process transport level failure.
void *	user_data	User data associated with an endpoint. See <a href="#">ucp_stream_poll_ep_t</a> and <a href="#">ucp_err_handler_t</a>
unsigned	flags	Endpoint flags from <a href="#">ucp_ep_params_flags_field</a> . This value is optional. If it's not set (along with its corresponding bit in the field_mask - <a href="#">UCP_EP_PARAM_FIELD_FLAGS</a> ), the <a href="#">ucp_ep_create()</a> routine will consider the flags as set to zero.
ucs_sock_addr_t	sockaddr	Destination address in the form of a sockaddr; this field should be set along with its corresponding bit in the field_mask - <a href="#">UCP_EP_PARAM_FIELD_SOCK_ADDR</a> and must be obtained from the user, it means that this type of the endpoint creation is possible only on client side in client-server connection establishment flow.
ucp_conn_request_h	conn_request	Connection request from client; this field should be set along with its corresponding bit in the field_mask - <a href="#">UCP_EP_PARAM_FIELD_CONN_REQUEST</a> and must be obtained from <a href="#">ucp_listener_conn_callback_t</a> , it means that this type of the endpoint creation is possible only on server side in client-server connection establishment flow.

## Data Fields

<code>const char *</code>	<code>name</code>	Endpoint name. Tracing and analysis tools can identify the endpoint using this name. To retrieve the endpoint's name, use <a href="#">ucp_ep_query</a> , as the name you supply may be changed by UCX under some circumstances, e.g. a name conflict. This field is only assigned if you set <a href="#">UCP_EP_PARAM_FIELD_NAME</a> in the field mask. If not, then a default unique name will be created for you.
<a href="#">ucp_sock_addr_t</a>	<code>local_sockaddr</code>	The sockaddr to bind locally. Specifies the associated network device to bind locally to establish new connections. To retrieve the endpoint's <code>local_sockaddr</code> , use <a href="#">ucp_ep_query</a> . This setting is optional. To enable it, the corresponding - <a href="#">UCP_EP_PARAM_FIELD_LOCAL_SOCK_ADDR</a> bit in the field mask must be set.

## 6.6.3 Typedef Documentation

### 6.6.3.1 `ucp_ep_perf_param_field_t`

```
typedef enum ucp_ep_perf_param_field ucp_ep_perf_param_field_t
```

The enumeration allows specifying which fields in [ucp\\_ep\\_evaluate\\_perf\\_param\\_t](#) are present and operation flags are used. It is used to enable backward compatibility support.

### 6.6.3.2 `ucp_ep_perf_attr_field_t`

```
typedef enum ucp_ep_perf_attr_field ucp_ep_perf_attr_field_t
```

The enumeration allows specifying which fields in [ucp\\_ep\\_evaluate\\_perf\\_attr\\_t](#) are present and operation flags are used. It is used to enable backward compatibility support.

### 6.6.3.3 `ucp_stream_poll_ep_t`

```
typedef struct ucp_stream_poll_ep ucp_stream_poll_ep_t
```

The structure defines the endpoint and its user data.

### 6.6.3.4 `ucp_ep_attr_t`

```
typedef struct ucp_ep_attr ucp_ep_attr_t
```

The structure defines the attributes that characterize the particular endpoint.

### 6.6.3.5 `ucp_ep_h`

```
typedef struct ucp_ep* ucp_ep_h
```

The endpoint handle is an opaque object that is used to address a remote [worker](#). It typically provides a description of source, destination, or both. All UCP communication routines address a destination with the endpoint handle. The endpoint handle is associated with only one [UCP context](#). UCP provides the [endpoint create](#) routine to create the endpoint handle and the [destroy](#) routine to destroy the endpoint handle.



### 6.6.3.6 ucp\_conn\_request\_h

```
typedef struct ucp_conn_request* ucp_conn_request_h
```

A server-side handle to incoming connection request. Can be used to create an endpoint which connects back to the client.

### 6.6.3.7 ucp\_am\_callback\_t

```
typedef ucs_status_t(* ucp_am_callback_t) (void *arg, void *data, size_t length, ucp_ep_h reply←_ep, unsigned flags)
```

When the callback is called, *flags* indicates how *data* should be handled.

#### Parameters

in	<i>arg</i>	User-defined argument.
in	<i>data</i>	Points to the received data. This data may persist after the callback returns and needs to be freed with <a href="#">ucp_am_data_release</a> .
in	<i>length</i>	Length of data.
in	<i>reply_ep</i>	If the Active Message is sent with the UCP_AM_SEND_FLAG_REPLY flag, the sending ep will be passed in. If not, NULL will be passed.
in	<i>flags</i>	If this flag is set to UCP_CB_PARAM_FLAG_DATA, the callback can return UCS_INPROGRESS and data will persist after the callback returns.

#### Returns

UCS\_OK *data* will not persist after the callback returns.

UCS\_INPROGRESS Can only be returned if *flags* is set to UCP\_CB\_PARAM\_FLAG\_DATA. If UCP\_INPR←OGRESS is returned, *data* will persist after the callback has returned. To free the memory, a pointer to the *data* must be passed into [ucp\\_am\\_data\\_release](#).

#### Note

This callback should be set and released by [ucp\\_worker\\_set\\_am\\_handler](#) function.

### 6.6.3.8 ucp\_am\_recv\_callback\_t

```
typedef ucs_status_t(* ucp_am_recv_callback_t) (void *arg, const void *header, size_t header←length, void *data, size_t length, const ucp_am_recv_param_t *param)
```

The callback is always called from the progress context, therefore calling [ucp\\_worker\\_progress\(\)](#) is not allowed. It is recommended to define callbacks with relatively short execution time to avoid blocking of communication progress.

#### Parameters

in	<i>arg</i>	User-defined argument.
in	<i>header</i>	User defined active message header. If <i>header_length</i> is 0, this value is undefined and must not be accessed.
in	<i>header_length</i>	Active message header length in bytes.
in	<i>data</i>	Points to the received data if UCP_AM_RECV_ATTR_FLAG_RNDV flag is not set in <a href="#">ucp_am_recv_param_t::recv_attr</a> . Otherwise it points to the internal UCP descriptor which can further be used for initiating data receive by using <a href="#">ucp_am_recv_data_nbx</a> routine.

## Parameters

in	<i>length</i>	Length of data. If <code>UCP_AM_RECV_ATTR_FLAG_RNDV</code> flag is set in <code>ucp_am_recv_param_t::recv_attr</code> , it indicates the required receive buffer size for initiating rendezvous protocol.
in	<i>param</i>	Data receive parameters.

## Returns

`UCS_OK` *data* will not persist after the callback returns. If `UCP_AM_RECV_ATTR_FLAG_RNDV` flag is set in `param->recv_attr` and `ucp_am_recv_data_nbx` was not called for this data, the data descriptor will be dropped and the corresponding `ucp_am_send_nbx` call will complete with `UCS_OK` status.

`UCS_INPROGRESS` Can only be returned if `param->recv_attr` flags contains `UCP_AM_RECV_ATTR_FLAG_DATA` or `UCP_AM_RECV_ATTR_FLAG_RNDV`. The *data* will persist after the callback has returned. To free the memory, a pointer to the data must be passed into `ucp_am_data_release` or data receive is initiated by `ucp_am_recv_data_nbx`.

otherwise Can only be returned if `param->recv_attr` contains `UCP_AM_RECV_ATTR_FLAG_RNDV`. In this case data descriptor *data* will be dropped and the corresponding `ucp_am_send_nbx` call on the sender side will complete with the status returned from the callback.

## Note

This callback should be set and released by `ucp_worker_set_am_recv_handler` function.

6.6.3.9 `ucp_ep_params_t`

```
typedef struct ucp_ep_params ucp_ep_params_t
```

The structure defines the parameters that are used for the UCP endpoint tuning during the UCP ep [creation](#).

## 6.6.4 Enumeration Type Documentation

6.6.4.1 `ucp_ep_params_field`

```
enum ucp_ep_params_field
```

The enumeration allows specifying which fields in `ucp_ep_params_t` are present. It is used to enable backward compatibility support.

## Enumerator

<code>UCP_EP_PARAM_FIELD_REMOTE_ADDRESS</code>	Address of remote peer
<code>UCP_EP_PARAM_FIELD_ERR_HANDLING_MODE</code>	Error handling mode. <a href="#">ucp_err_handling_mode_t</a>
<code>UCP_EP_PARAM_FIELD_ERR_HANDLER</code>	Handler to process transport level errors
<code>UCP_EP_PARAM_FIELD_USER_DATA</code>	User data pointer
<code>UCP_EP_PARAM_FIELD_SOCKET_ADDR</code>	Socket address field
<code>UCP_EP_PARAM_FIELD_FLAGS</code>	Endpoint flags Connection request field
<code>UCP_EP_PARAM_FIELD_CONN_REQUEST</code>	
<code>UCP_EP_PARAM_FIELD_NAME</code>	Endpoint name
<code>UCP_EP_PARAM_FIELD_LOCAL_SOCKET_ADDR</code>	Local socket Address

### 6.6.4.2 ucp\_ep\_params\_flags\_field

enum [ucp\\_ep\\_params\\_flags\\_field](#)

The enumeration list describes the endpoint's parameters flags supported by [ucp\\_ep\\_create\(\)](#) function.

#### Enumerator

UCP_EP_PARAMS_FLAGS_CLIENT_SERVER	Using a client-server connection establishment mechanism. <a href="#">ucp_sock_addr_t</a> sockaddr field must be provided and contain the address of the remote peer
UCP_EP_PARAMS_FLAGS_NO_LOOPBACK	Avoid connecting the endpoint to itself when connecting the endpoint to the same worker it was created on. Affects protocols which send to a particular remote endpoint, for example stream
UCP_EP_PARAMS_FLAGS_SEND_CLIENT_ID	Send client id when connecting to remote socket address as part of the connection request payload. On the remote side value can be obtained from <a href="#">ucp_conn_request_h</a> using <a href="#">ucp_conn_request_query</a>

### 6.6.4.3 ucp\_ep\_close\_flags\_t

enum [ucp\\_ep\\_close\\_flags\\_t](#)

The enumeration is used to specify the behavior of [ucp\\_ep\\_close\\_nbx](#).

#### Enumerator

UCP_EP_CLOSE_FLAG_FORCE	<p><a href="#">ucp_ep_close_nbx</a> releases the endpoint without any confirmation from the peer. All outstanding requests will be completed with <a href="#">UCS_ERR_CANCELED</a> error.</p> <p><b>Note</b></p> <p>This mode may cause transport level errors on remote side, so it requires set <a href="#">UCP_ERR_HANDLING_MODE_PEER</a> for all endpoints created on both (local and remote) sides to avoid undefined behavior. If this flag is not set then <a href="#">ucp_ep_close_nbx</a> schedules flushes on all outstanding operations.</p> <p>this flag is incompatible with <a href="#">UCP_OP_ATTR_FLAG_NO_IMM_CMPL</a>, since it forces immediate completion.</p>
-------------------------	---

### 6.6.4.4 ucp\_ep\_close\_mode

enum [ucp\\_ep\\_close\\_mode](#)

The enumeration is used to specify the behavior of [ucp\\_ep\\_close\\_nb](#).

## Enumerator

UCP_EP_CLOSE_MODE_FORCE	<p><a href="#">ucp_ep_close_nb</a> releases the endpoint without any confirmation from the peer. All outstanding requests will be completed with <a href="#">UCS_ERR_CANCELED</a> error.</p> <p><b>Note</b></p> <p>This mode may cause transport level errors on remote side, so it requires set <a href="#">UCP_ERR_HANDLING_MODE_PEER</a> for all endpoints created on both (local and remote) sides to avoid undefined behavior.</p>
UCP_EP_CLOSE_MODE_FLUSH	<a href="#">ucp_ep_close_nb</a> schedules flushes on all outstanding operations.

## 6.6.4.5 ucp\_ep\_perf\_param\_field

enum [ucp\\_ep\\_perf\\_param\\_field](#)

The enumeration allows specifying which fields in [ucp\\_ep\\_evaluate\\_perf\\_param\\_t](#) are present and operation flags are used. It is used to enable backward compatibility support.

## Enumerator

UCP_EP_PERF_PARAM_FIELD_MESSAGE_SIZE	Enables <a href="#">ucp_ep_evaluate_perf_param_t::message_size</a>
--------------------------------------	--

## 6.6.4.6 ucp\_ep\_perf\_attr\_field

enum [ucp\\_ep\\_perf\\_attr\\_field](#)

The enumeration allows specifying which fields in [ucp\\_ep\\_evaluate\\_perf\\_attr\\_t](#) are present and operation flags are used. It is used to enable backward compatibility support.

## Enumerator

UCP_EP_PERF_ATTR_FIELD_ESTIMATED_TIME	Enables <a href="#">ucp_ep_evaluate_perf_attr_t::estimated_time</a>
---------------------------------------	---

## 6.6.4.7 ucp\_cb\_param\_flags

enum [ucp\\_cb\\_param\\_flags](#)

In a callback, if flags is set to [UCP\\_CB\\_PARAM\\_FLAG\\_DATA](#) in a callback then data was allocated, so if [UCS\\_INPROGRESS](#) is returned from the callback, the data parameter will persist and the user has to call [ucp\\_am\\_data\\_release](#) when data is no longer needed.

## Enumerator

UCP_CB_PARAM_FLAG_DATA	
------------------------	--

### 6.6.4.8 ucp\_ep\_attr\_field

enum `ucp_ep_attr_field`

The enumeration allows specifying which fields in `ucp_ep_attr_t` are present. It is used to enable backward compatibility support.

Enumerator

UCP_EP_ATTR_FIELD_NAME	UCP endpoint name
UCP_EP_ATTR_FIELD_LOCAL_SOCKADDR	Socketaddr used by the endpoint
UCP_EP_ATTR_FIELD_REMOTE_SOCKADDR	Socketaddr the endpoint is connected to

### 6.6.4.9 ucp\_err\_handling\_mode\_t

enum `ucp_err_handling_mode_t`

Specifies error handling mode for the UCP endpoint.

Enumerator

UCP_ERR_HANDLING_MODE_NONE	<p>No guarantees about error reporting, imposes minimal overhead from a performance perspective.</p> <p><b>Note</b></p> <p>In this mode, any error reporting will not generate calls to <code>ucp_ep_params_t::err_handler</code>.</p>
UCP_ERR_HANDLING_MODE_PEER	<p>Guarantees that send requests are always completed (successfully or error) even in case of remote failure, disables protocols and APIs which may cause a hang or undefined behavior in case of peer failure, may affect performance and memory footprint</p>

## 6.6.5 Function Documentation

### 6.6.5.1 ucp\_ep\_create()

```
ucs_status_t ucp_ep_create (
    ucp_worker_h worker,
    const ucp_ep_params_t * params,
    ucp_ep_h * ep_p )
```

This routine creates and connects an [endpoint](#) on a [local worker](#) for a destination [address](#) that identifies the remote [worker](#). This function is non-blocking, and communications may begin immediately after it returns. If the connection process is not completed, communications may be delayed. The created [endpoint](#) is associated with one and only one [worker](#).

## Parameters

in	<i>worker</i>	Handle to the worker; the endpoint is associated with the worker.
in	<i>params</i>	User defined <a href="#">ucp_ep_params_t</a> configurations for the <a href="#">UCP endpoint</a> .
out	<i>ep_p</i>	A handle to the created endpoint.

## Returns

Error code as defined by [ucs\\_status\\_t](#)

## Note

One of the following fields has to be specified:

- [ucp\\_ep\\_params\\_t::address](#)
- [ucp\\_ep\\_params\\_t::sockaddr](#)
- [ucp\\_ep\\_params\\_t::conn\\_request](#)

By default, [ucp\\_ep\\_create\(\)](#) will connect an endpoint to itself if the endpoint is destined to the same *worker* on which it was created, i.e. *params.address* belongs to *worker*. This behavior can be changed by passing the [UCP\\_EP\\_PARAMS\\_FLAGS\\_NO\\_LOOPBACK](#) flag in *params.flags*. In that case, the endpoint will be connected to the *next* endpoint created in the same way on the same *worker*.

## Examples

[ucp\\_client\\_server.c](#), and [ucp\\_hello\\_world.c](#).

6.6.5.2 [ucp\\_ep\\_close\\_nb\(\)](#)

```
ucs_status_ptr_t ucp_ep_close_nb (
    ucp_ep_h ep,
    unsigned mode )
```

This routine releases the [endpoint](#). The endpoint closure process depends on the selected *mode*.

## Parameters

in	<i>ep</i>	Handle to the endpoint to close.
in	<i>mode</i>	One from <a href="#">ucp_ep_close_mode</a> value.

## Returns

UCS\_OK - The endpoint is closed successfully.

UCS\_PTR\_IS\_ERR(\_ptr) - The closure failed and an error code indicates the transport level status. However, resources are released and the *endpoint* can no longer be used.

otherwise - The closure process is started, and can be completed at any point in time. A request handle is returned to the application in order to track progress of the endpoint closure. The application is responsible for releasing the handle using the [ucp\\_request\\_free](#) routine.

## Note

[ucp\\_ep\\_close\\_nb](#) replaces deprecated [ucp\\_disconnect\\_nb](#) and [ucp\\_ep\\_destroy](#)

### 6.6.5.3 ucp\_ep\_close\_nbx()

```
ucs_status_ptr_t ucp_ep_close_nbx (
    ucp_ep_h ep,
    const ucp_request_param_t * param )
```

#### Parameters

in	<i>ep</i>	Handle to the endpoint to close.
in	<i>param</i>	Operation parameters, see <a href="#">ucp_request_param_t</a> . This operation supports specific flags, which can be passed in <i>param</i> by <a href="#">ucp_request_param_t::flags</a> . The exact set of flags is defined by <a href="#">ucp_ep_close_flags_t</a> .

#### Returns

NULL - The endpoint is closed successfully.

UCS\_PTR\_IS\_ERR(\_ptr) - The closure failed and an error code indicates the transport level status. However, resources are released and the *endpoint* can no longer be used.

otherwise - The closure process is started, and can be completed at any point in time. A request handle is returned to the application in order to track progress of the endpoint closure.

### 6.6.5.4 ucp\_ep\_print\_info()

```
void ucp_ep_print_info (
    ucp_ep_h ep,
    FILE * stream )
```

This routine prints information about the endpoint transport methods, their thresholds, and other useful information associated with the endpoint.

#### Parameters

in	<i>ep</i>	Endpoint object whose configuration to print.
in	<i>stream</i>	Output stream to print the information to.

### 6.6.5.5 ucp\_ep\_flush\_nb()

```
ucs_status_ptr_t ucp_ep_flush_nb (
    ucp_ep_h ep,
    unsigned flags,
    ucp_send_callback_t cb )
```

This routine flushes all outstanding AMO and RMA communications on the [endpoint](#). All the AMO and RMA operations issued on the *ep* prior to this call are completed both at the origin and at the target [endpoint](#) when this call returns.

#### Parameters

in	<i>ep</i>	UCP endpoint.
in	<i>flags</i>	Flags for flush operation. Reserved for future use.
in	<i>cb</i>	Callback which will be called when the flush operation completes.

**Returns**

NULL - The flush operation was completed immediately.

UCS\_PTR\_IS\_ERR(\_ptr) - The flush operation failed.

otherwise - Flush operation was scheduled and can be completed in any point in time. The request handle is returned to the application in order to track progress. The application is responsible for releasing the handle using `ucp_request_free()` routine.

The following example demonstrates how blocking flush can be implemented using non-blocking flush:

```
void empty_function(void *request, ucs_status_t status)
{
}
ucs_status_t blocking_ep_flush(ucp_ep_h ep, ucp_worker_h worker)
{
    void *request;
    request = ucp_ep_flush_nb(ep, 0, empty_function);
    if (request == NULL) {
        return UCS_OK;
    } else if (UCS_PTR_IS_ERR(request)) {
        return UCS_PTR_STATUS(request);
    } else {
        ucs_status_t status;
        do {
            ucp_worker_progress(worker);
            status = ucp_request_check_status(request);
        } while (status == UCS_INPROGRESS);
        ucp_request_free(request);
        return status;
    }
}
```

**6.6.5.6 ucp\_ep\_flush\_nbx()**

```
ucs_status_ptr_t ucp_ep_flush_nbx (
    ucp_ep_h ep,
    const ucp_request_param_t * param )
```

This routine flushes all outstanding AMO and RMA communications on the [endpoint](#). All the AMO and RMA operations issued on the *ep* prior to this call are completed both at the origin and at the target [endpoint](#) when this call returns.

**Parameters**

in	<i>ep</i>	UCP endpoint.
in	<i>param</i>	Operation parameters, see <a href="#">ucp_request_param_t</a> .

**Returns**

NULL - The flush operation was completed immediately.

UCS\_PTR\_IS\_ERR(\_ptr) - The flush operation failed.

otherwise - Flush operation was scheduled and can be completed in any point in time. The request handle is returned to the application in order to track progress.

**Examples**

[ucp\\_hello\\_world.c](#).

**6.6.5.7 ucp\_ep\_evaluate\_perf()**

```
ucs_status_t ucp_ep_evaluate_perf (
    ucp_ep_h ep,
```



```
const ucp_ep_evaluate_perf_param_t * param,
      ucp_ep_evaluate_perf_attr_t * attr )
```

This routine fetches information about the endpoint.

#### Parameters

in	<i>ep</i>	Endpoint to query.
in	<i>param</i>	Filled by the user with request params.
out	<i>attr</i>	Filled with performance estimation of the given operation on the endpoint.

#### Returns

Error code as defined by [ucs\\_status\\_t](#)

#### 6.6.5.8 ucp\_ep\_query()

```
ucs_status_t ucp_ep_query (
    ucp_ep_h ep,
    ucp_ep_attr_t * attr )
```

This routine fetches information about the endpoint.

#### Parameters

in	<i>ep</i>	Endpoint object to query.
out	<i>attr</i>	Filled with attributes of the endpoint.

#### Returns

Error code as defined by [ucs\\_status\\_t](#)

#### 6.6.5.9 ucp\_request\_release()

```
void ucp_request_release (
    void * request )
```

**Deprecated** Replaced by [ucp\\_request\\_free](#).

#### 6.6.5.10 ucp\_ep\_destroy()

```
void ucp_ep_destroy (
    ucp_ep_h ep )
```

**Deprecated** Replaced by [ucp\\_ep\\_close\\_nb](#).

### 6.6.5.11 ucp\_disconnect\_nb()

```
ucs_status_ptr_t ucp_disconnect_nb (
    ucp_ep_h ep )
```

**Deprecated** Replaced by [ucp\\_ep\\_close\\_nb](#).

### 6.6.5.12 ucp\_request\_test()

```
ucs_status_t ucp_request_test (
    void * request,
    ucp_tag_recv_info_t * info )
```

**Deprecated** Replaced by [ucp\\_tag\\_recv\\_request\\_test](#) and [ucp\\_request\\_check\\_status](#) depends on use case.

#### Note

Please use [ucp\\_request\\_check\\_status](#) for cases that only need to check the completion status of an outstanding request. [ucp\\_request\\_check\\_status](#) can be used for any type of request. [ucp\\_tag\\_recv\\_request\\_test](#) should only be used for requests returned by [ucp\\_tag\\_recv\\_nb](#) (or request allocated by user for [ucp\\_tag\\_recv\\_nbr](#)) for which additional information (returned via the *info* pointer) is needed.

### 6.6.5.13 ucp\_ep\_flush()

```
ucs_status_t ucp_ep_flush (
    ucp_ep_h ep )
```

**Deprecated** Replaced by [ucp\\_ep\\_flush\\_nb](#).

### 6.6.5.14 ucp\_ep\_modify\_nb()

```
ucs_status_ptr_t ucp_ep_modify_nb (
    ucp_ep_h ep,
    const ucp_ep_params_t * params )
```

**Deprecated** Use [ucp\\_listener\\_conn\\_handler\\_t](#) instead of [ucp\\_listener\\_accept\\_handler\\_t](#), if you have other use case please submit an issue on <https://github.com/openucx/ucx> or report to [ucx-group@elist.ornl.gov](mailto:ucx-group@elist.ornl.gov)

This routine modifies [endpoint](#) created by [ucp\\_ep\\_create](#) or [ucp\\_listener\\_accept\\_callback\\_t](#). For example, this API can be used to setup custom parameters like [ucp\\_ep\\_params\\_t::user\\_data](#) or [ucp\\_ep\\_params\\_t::err\\_handler](#) to endpoint created by [ucp\\_listener\\_accept\\_callback\\_t](#).

#### Parameters

in	<i>ep</i>	A handle to the endpoint.
in	<i>params</i>	User defined <a href="#">ucp_ep_params_t</a> configurations for the <a href="#">UCP endpoint</a> .

**Returns**

NULL - The endpoint is modified successfully.

UCS\_PTR\_IS\_ERR(\_ptr) - The reconfiguration failed and an error code indicates the status. However, the *endpoint* is not modified and can be used further.

otherwise - The reconfiguration process is started, and can be completed at any point in time. A request handle is returned to the application in order to track progress of the endpoint modification. The application is responsible for releasing the handle using the [ucp\\_request\\_free](#) routine.

**Note**

See the documentation of [ucp\\_ep\\_params\\_t](#) for details, only some of the parameters can be modified.

## 6.7 UCP Communication routines

### Data Structures

- struct [ucp\\_request\\_attr\\_t](#)  
*Attributes of a particular request. [More...](#)*
- struct [ucp\\_err\\_handler](#)  
*UCP endpoint error handling context. [More...](#)*

### Typedefs

- typedef uint64\_t [ucp\\_tag\\_t](#)  
*UCP Tag Identifier.*
- typedef struct ucp\_recv\_desc \* [ucp\\_tag\\_message\\_h](#)  
*UCP Message descriptor.*
- typedef uint64\_t [ucp\\_datatype\\_t](#)  
*UCP Datatype Identifier.*
- typedef void(\* [ucp\\_send\\_callback\\_t](#)) (void \*request, [ucs\\_status\\_t](#) status)  
*Completion callback for non-blocking sends.*
- typedef void(\* [ucp\\_send\\_nbx\\_callback\\_t](#)) (void \*request, [ucs\\_status\\_t](#) status, void \*user\_data)  
*Completion callback for non-blocking sends.*
- typedef void(\* [ucp\\_err\\_handler\\_cb\\_t](#)) (void \*arg, [ucp\\_ep\\_h](#) ep, [ucs\\_status\\_t](#) status)  
*Callback to process peer failure.*
- typedef struct [ucp\\_err\\_handler](#) [ucp\\_err\\_handler\\_t](#)  
*UCP endpoint error handling context.*
- typedef void(\* [ucp\\_stream\\_recv\\_callback\\_t](#)) (void \*request, [ucs\\_status\\_t](#) status, size\_t length)  
*Completion callback for non-blocking stream oriented receives.*
- typedef void(\* [ucp\\_stream\\_recv\\_nbx\\_callback\\_t](#)) (void \*request, [ucs\\_status\\_t](#) status, size\_t length, void \*user\_data)  
*Completion callback for non-blocking stream receives `ucp_stream_recv_nbx` call.*
- typedef void(\* [ucp\\_tag\\_recv\\_callback\\_t](#)) (void \*request, [ucs\\_status\\_t](#) status, [ucp\\_tag\\_recv\\_info\\_t](#) \*info)  
*Completion callback for non-blocking tag receives.*
- typedef void(\* [ucp\\_tag\\_recv\\_nbx\\_callback\\_t](#)) (void \*request, [ucs\\_status\\_t](#) status, const [ucp\\_tag\\_recv\\_info\\_t](#) \*tag\_info, void \*user\_data)  
*Completion callback for non-blocking tag receives `ucp_tag_recv_nbx` call.*
- typedef void(\* [ucp\\_am\\_recv\\_data\\_nbx\\_callback\\_t](#)) (void \*request, [ucs\\_status\\_t](#) status, size\_t length, void \*user\_data)  
*Completion callback for non-blocking Active Message receives.*

### Enumerations

- enum [ucp\\_atomic\\_post\\_op\\_t](#) {  
    [UCP\\_ATOMIC\\_POST\\_OP\\_ADD](#),     [UCP\\_ATOMIC\\_POST\\_OP\\_AND](#),     [UCP\\_ATOMIC\\_POST\\_OP\\_OR](#),  
    [UCP\\_ATOMIC\\_POST\\_OP\\_XOR](#),  
    [UCP\\_ATOMIC\\_POST\\_OP\\_LAST](#) }  
*Atomic operation requested for `ucp_atomic_post`.*
- enum [ucp\\_atomic\\_fetch\\_op\\_t](#) {  
    [UCP\\_ATOMIC\\_FETCH\\_OP\\_FADD](#), [UCP\\_ATOMIC\\_FETCH\\_OP\\_SWAP](#), [UCP\\_ATOMIC\\_FETCH\\_OP\\_CSWAP](#),  
    [UCP\\_ATOMIC\\_FETCH\\_OP\\_FAND](#),  
    [UCP\\_ATOMIC\\_FETCH\\_OP\\_FOR](#), [UCP\\_ATOMIC\\_FETCH\\_OP\\_FXOR](#), [UCP\\_ATOMIC\\_FETCH\\_OP\\_LAST](#)  
    }  
*Atomic operation requested for `ucp_atomic_fetch`.*

*Atomic operation requested for `ucp_atomic_fetch`.*

- enum `ucp_atomic_op_t` {  
`UCP_ATOMIC_OP_ADD`, `UCP_ATOMIC_OP_SWAP`, `UCP_ATOMIC_OP_CSWAP`, `UCP_ATOMIC_OP_AND`,  
`UCP_ATOMIC_OP_OR`, `UCP_ATOMIC_OP_XOR`, `UCP_ATOMIC_OP_LAST` }

*Atomic operation requested for `ucp_atomic_op_nbx`.*

- enum `ucp_stream_recv_flags_t` { `UCP_STREAM_RECV_FLAG_WAITALL` = `UCS_BIT(0)` }

*Flags to define behavior of `ucp_stream_recv_nb` function.*

- enum `ucp_op_attr_t` {  
`UCP_OP_ATTR_FIELD_REQUEST` = `UCS_BIT(0)`, `UCP_OP_ATTR_FIELD_CALLBACK` = `UCS_BIT(1)`,  
`UCP_OP_ATTR_FIELD_USER_DATA` = `UCS_BIT(2)`, `UCP_OP_ATTR_FIELD_DATATYPE` = `UCS_BIT(3)`,  
`UCP_OP_ATTR_FIELD_FLAGS` = `UCS_BIT(4)`, `UCP_OP_ATTR_FIELD_REPLY_BUFFER` = `UCS_BIT(5)`,  
`UCP_OP_ATTR_FIELD_MEMORY_TYPE` = `UCS_BIT(6)`, `UCP_OP_ATTR_FIELD_RECV_INFO` = `UCS_BIT(7)`,  
`UCP_OP_ATTR_FIELD_MEMH` = `UCS_BIT(8)`, `UCP_OP_ATTR_FLAG_NO_IMM_CMPL` = `UCS_BIT(16)`,  
`UCP_OP_ATTR_FLAG_FAST_CMPL` = `UCS_BIT(17)`, `UCP_OP_ATTR_FLAG_FORCE_IMM_CMPL` = `UCS_BIT(18)`,  
`UCP_OP_ATTR_FLAG_MULTI_SEND` = `UCS_BIT(19)` }

*UCP operation fields and flags.*

- enum `ucp_req_attr_field` { `UCP_REQUEST_ATTR_FIELD_INFO_STRING` = `UCS_BIT(0)`, `UCP_REQUEST_ATTR_FIELD_INFO_SIZE` = `UCS_BIT(1)`, `UCP_REQUEST_ATTR_FIELD_STATUS` = `UCS_BIT(2)`, `UCP_REQUEST_ATTR_FIELD_MEM_TYPE` = `UCS_BIT(3)` }

*UCP request query attributes.*

- enum `ucp_am_recv_attr_t` { `UCP_AM_RECV_ATTR_FIELD_REPLY_EP` = `UCS_BIT(0)`, `UCP_AM_RECV_ATTR_FLAG_DATA_SYNC` = `UCS_BIT(16)`, `UCP_AM_RECV_ATTR_FLAG_RNDV` = `UCS_BIT(17)` }

*UCP AM receive data parameter fields and flags.*

- enum `ucp_am_handler_param_field` { `UCP_AM_HANDLER_PARAM_FIELD_ID` = `UCS_BIT(0)`, `UCP_AM_HANDLER_PARAM_FIELD_CB` = `UCS_BIT(1)`, `UCP_AM_HANDLER_PARAM_FIELD_ARG` = `UCS_BIT(2)`, `UCP_AM_HANDLER_PARAM_FIELD_ARG_SIZE` = `UCS_BIT(3)` }

*UCP AM receive data parameters fields and flags.*

## Functions

- `ucs_status_t ucp_request_query` (void \*request, `ucp_request_attr_t` \*attr)

*Get information about `ucp_request`.*

- `ucs_status_ptr_t ucp_am_send_nb` (`ucp_ep_h` ep, `uint16_t` id, const void \*buffer, `size_t` count, `ucp_datatype_t` datatype, `ucp_send_callback_t` cb, unsigned flags)

*Send Active Message.*

- `ucs_status_ptr_t ucp_am_send_nbx` (`ucp_ep_h` ep, unsigned id, const void \*header, `size_t` header\_length, const void \*buffer, `size_t` count, const `ucp_request_param_t` \*param)

*Send Active Message.*

- `ucs_status_ptr_t ucp_am_recv_data_nbx` (`ucp_worker_h` worker, void \*data\_desc, void \*buffer, `size_t` count, const `ucp_request_param_t` \*param)

*Receive Active Message as defined by provided data descriptor.*

- void `ucp_am_data_release` (`ucp_worker_h` worker, void \*data)

*Releases Active Message data.*

- `ucs_status_ptr_t ucp_stream_send_nb` (`ucp_ep_h` ep, const void \*buffer, `size_t` count, `ucp_datatype_t` datatype, `ucp_send_callback_t` cb, unsigned flags)

*Non-blocking stream send operation.*

- `ucs_status_ptr_t ucp_stream_send_nbx` (`ucp_ep_h` ep, const void \*buffer, `size_t` count, const `ucp_request_param_t` \*param)

*Non-blocking stream send operation.*

- `ucs_status_ptr_t ucp_tag_send_nb` (`ucp_ep_h` ep, const void \*buffer, `size_t` count, `ucp_datatype_t` datatype, `ucp_tag_t` tag, `ucp_send_callback_t` cb)

*Non-blocking tagged-send operations.*

- `ucs_status_t ucp_tag_send_nbr` (`ucp_ep_h` ep, const void \*buffer, size\_t count, `ucp_datatype_t` datatype, `ucp_tag_t` tag, void \*req)

*Non-blocking tagged-send operations with user provided request.*

- `ucs_status_ptr_t ucp_tag_send_sync_nb` (`ucp_ep_h` ep, const void \*buffer, size\_t count, `ucp_datatype_t` datatype, `ucp_tag_t` tag, `ucp_send_callback_t` cb)

*Non-blocking synchronous tagged-send operation.*

- `ucs_status_ptr_t ucp_tag_send_nbx` (`ucp_ep_h` ep, const void \*buffer, size\_t count, `ucp_tag_t` tag, const `ucp_request_param_t` \*param)

*Non-blocking tagged-send operation.*

- `ucs_status_ptr_t ucp_tag_send_sync_nbx` (`ucp_ep_h` ep, const void \*buffer, size\_t count, `ucp_tag_t` tag, const `ucp_request_param_t` \*param)

*Non-blocking synchronous tagged-send operation.*

- `ucs_status_ptr_t ucp_stream_rcv_nb` (`ucp_ep_h` ep, void \*buffer, size\_t count, `ucp_datatype_t` datatype, `ucp_stream_rcv_callback_t` cb, size\_t \*length, unsigned flags)

*Non-blocking stream receive operation of structured data into a user-supplied buffer.*

- `ucs_status_ptr_t ucp_stream_rcv_nbx` (`ucp_ep_h` ep, void \*buffer, size\_t count, size\_t \*length, const `ucp_request_param_t` \*param)

*Non-blocking stream receive operation of structured data into a user-supplied buffer.*

- `ucs_status_ptr_t ucp_stream_rcv_data_nb` (`ucp_ep_h` ep, size\_t \*length)

*Non-blocking stream receive operation of unstructured data into a UCP-supplied buffer.*

- `ucs_status_ptr_t ucp_tag_rcv_nb` (`ucp_worker_h` worker, void \*buffer, size\_t count, `ucp_datatype_t` datatype, `ucp_tag_t` tag, `ucp_tag_t` tag\_mask, `ucp_tag_rcv_callback_t` cb)

*Non-blocking tagged-recv operation.*

- `ucs_status_t ucp_tag_rcv_nbr` (`ucp_worker_h` worker, void \*buffer, size\_t count, `ucp_datatype_t` datatype, `ucp_tag_t` tag, `ucp_tag_t` tag\_mask, void \*req)

*Non-blocking tagged-recv operation.*

- `ucs_status_ptr_t ucp_tag_rcv_nbx` (`ucp_worker_h` worker, void \*buffer, size\_t count, `ucp_tag_t` tag, `ucp_tag_t` tag\_mask, const `ucp_request_param_t` \*param)

*Non-blocking tagged-recv operation.*

- `ucp_tag_message_h ucp_tag_probe_nb` (`ucp_worker_h` worker, `ucp_tag_t` tag, `ucp_tag_t` tag\_mask, int remove, `ucp_tag_rcv_info_t` \*info)

*Non-blocking probe and return a message.*

- `ucs_status_ptr_t ucp_tag_msg_rcv_nb` (`ucp_worker_h` worker, void \*buffer, size\_t count, `ucp_datatype_t` datatype, `ucp_tag_message_h` message, `ucp_tag_rcv_callback_t` cb)

*Non-blocking receive operation for a probed message.*

- `ucs_status_ptr_t ucp_tag_msg_rcv_nbx` (`ucp_worker_h` worker, void \*buffer, size\_t count, `ucp_tag_message_h` message, const `ucp_request_param_t` \*param)

*Non-blocking receive operation for a probed message.*

- `ucs_status_t ucp_put_nbi` (`ucp_ep_h` ep, const void \*buffer, size\_t length, uint64\_t remote\_addr, `ucp_rkey_h` rkey)

*Non-blocking implicit remote memory put operation.*

- `ucs_status_ptr_t ucp_put_nb` (`ucp_ep_h` ep, const void \*buffer, size\_t length, uint64\_t remote\_addr, `ucp_rkey_h` rkey, `ucp_send_callback_t` cb)

*Non-blocking remote memory put operation.*

- `ucs_status_ptr_t ucp_put_nbx` (`ucp_ep_h` ep, const void \*buffer, size\_t count, uint64\_t remote\_addr, `ucp_rkey_h` rkey, const `ucp_request_param_t` \*param)

*Non-blocking remote memory put operation.*

- `ucs_status_t ucp_get_nbi` (`ucp_ep_h` ep, void \*buffer, size\_t length, uint64\_t remote\_addr, `ucp_rkey_h` rkey)

*Non-blocking implicit remote memory get operation.*

- `ucs_status_ptr_t ucp_get_nb` (`ucp_ep_h` ep, void \*buffer, size\_t length, uint64\_t remote\_addr, `ucp_rkey_h` rkey, `ucp_send_callback_t` cb)

*Non-blocking remote memory get operation.*

- `ucs_status_ptr_t ucp_get_nbx` (`ucp_ep_h` ep, void \*buffer, size\_t count, uint64\_t remote\_addr, `ucp_rkey_h` rkey, const `ucp_request_param_t` \*param)  
*Non-blocking remote memory get operation.*
- `ucs_status_t ucp_atomic_post` (`ucp_ep_h` ep, `ucp_atomic_post_op_t` opcode, uint64\_t value, size\_t op\_size, uint64\_t remote\_addr, `ucp_rkey_h` rkey)  
*Post an atomic memory operation.*
- `ucs_status_ptr_t ucp_atomic_fetch_nb` (`ucp_ep_h` ep, `ucp_atomic_fetch_op_t` opcode, uint64\_t value, void \*result, size\_t op\_size, uint64\_t remote\_addr, `ucp_rkey_h` rkey, `ucp_send_callback_t` cb)  
*Post an atomic fetch operation.*
- `ucs_status_ptr_t ucp_atomic_op_nbx` (`ucp_ep_h` ep, `ucp_atomic_op_t` opcode, const void \*buffer, size\_t count, uint64\_t remote\_addr, `ucp_rkey_h` rkey, const `ucp_request_param_t` \*param)  
*Post an atomic memory operation.*
- `ucs_status_t ucp_request_check_status` (void \*request)  
*Check the status of non-blocking request.*
- `ucs_status_t ucp_tag_recv_request_test` (void \*request, `ucp_tag_recv_info_t` \*info)  
*Check the status and currently available state of non-blocking request returned from `ucp_tag_recv_nb` routine.*
- `ucs_status_t ucp_stream_recv_request_test` (void \*request, size\_t \*length\_p)  
*Check the status and currently available state of non-blocking request returned from `ucp_stream_recv_nb` routine.*
- void `ucp_request_cancel` (`ucp_worker_h` worker, void \*request)  
*Cancel an outstanding communications request.*
- void `ucp_stream_data_release` (`ucp_ep_h` ep, void \*data)  
*Release UCP data buffer returned by `ucp_stream_recv_data_nb`.*
- void `ucp_request_free` (void \*request)  
*Release a communications request.*
- void \* `ucp_request_alloc` (`ucp_worker_h` worker)  
*Create an empty communications request.*
- int `ucp_request_is_completed` (void \*request)
- `ucs_status_t ucp_put` (`ucp_ep_h` ep, const void \*buffer, size\_t length, uint64\_t remote\_addr, `ucp_rkey_h` rkey)  
*Blocking remote memory put operation.*
- `ucs_status_t ucp_get` (`ucp_ep_h` ep, void \*buffer, size\_t length, uint64\_t remote\_addr, `ucp_rkey_h` rkey)  
*Blocking remote memory get operation.*
- `ucs_status_t ucp_atomic_add32` (`ucp_ep_h` ep, uint32\_t add, uint64\_t remote\_addr, `ucp_rkey_h` rkey)  
*Blocking atomic add operation for 32 bit integers.*
- `ucs_status_t ucp_atomic_add64` (`ucp_ep_h` ep, uint64\_t add, uint64\_t remote\_addr, `ucp_rkey_h` rkey)  
*Blocking atomic add operation for 64 bit integers.*
- `ucs_status_t ucp_atomic_fadd32` (`ucp_ep_h` ep, uint32\_t add, uint64\_t remote\_addr, `ucp_rkey_h` rkey, uint32\_t \*result)  
*Blocking atomic fetch and add operation for 32 bit integers.*
- `ucs_status_t ucp_atomic_fadd64` (`ucp_ep_h` ep, uint64\_t add, uint64\_t remote\_addr, `ucp_rkey_h` rkey, uint64\_t \*result)  
*Blocking atomic fetch and add operation for 64 bit integers.*
- `ucs_status_t ucp_atomic_swap32` (`ucp_ep_h` ep, uint32\_t swap, uint64\_t remote\_addr, `ucp_rkey_h` rkey, uint32\_t \*result)  
*Blocking atomic swap operation for 32 bit values.*
- `ucs_status_t ucp_atomic_swap64` (`ucp_ep_h` ep, uint64\_t swap, uint64\_t remote\_addr, `ucp_rkey_h` rkey, uint64\_t \*result)  
*Blocking atomic swap operation for 64 bit values.*
- `ucs_status_t ucp_atomic_cswap32` (`ucp_ep_h` ep, uint32\_t compare, uint32\_t swap, uint64\_t remote\_addr, `ucp_rkey_h` rkey, uint32\_t \*result)  
*Blocking atomic conditional swap (cswap) operation for 32 bit values.*
- `ucs_status_t ucp_atomic_cswap64` (`ucp_ep_h` ep, uint64\_t compare, uint64\_t swap, uint64\_t remote\_addr, `ucp_rkey_h` rkey, uint64\_t \*result)  
*Blocking atomic conditional swap (cswap) operation for 64 bit values.*

## 6.7.1 Detailed Description

UCP Communication routines

## 6.7.2 Data Structure Documentation

### 6.7.2.1 struct ucp\_request\_attr\_t

Data Fields

uint64_t	field_mask	Mask of valid fields in this structure, using bits from <a href="#">ucp_req_attr_field</a> . Fields not specified in this mask will be ignored. Provides ABI compatibility with respect to adding new fields.
char *	debug_string	Pointer to allocated string of size <a href="#">debug_string_size</a> that will be filled with debug information about transports and protocols that were selected to complete the request.
size_t	debug_string_size	Size of the <a href="#">debug_string</a> . String will be filled up to this size. Maximum possible size debug string can be obtained by querying the worker via <a href="#">ucp_worker_query</a> .
<a href="#">ucs_status_t</a>	status	Status of the request. The same as <a href="#">ucp_request_check_status</a> .
<a href="#">ucs_memory_type_t</a>	mem_type	Detected memory type of the buffer passed to the operation.

### 6.7.2.2 struct ucp\_err\_handler

This structure should be initialized in [ucp\\_ep\\_params\\_t](#) to handle peer failure

Data Fields

<a href="#">ucp_err_handler_cb_t</a>	cb	Error handler callback, if NULL, will not be called.
void *	arg	User defined argument associated with an endpoint, it will be overridden by <a href="#">ucp_ep_params_t::user_data</a> if both are set.

## 6.7.3 Typedef Documentation

### 6.7.3.1 ucp\_tag\_t

```
typedef uint64_t ucp_tag_t
```

UCP tag identifier is a 64bit object used for message identification. UCP tag send and receive operations use the object for an implementation tag matching semantics (derivative of MPI tag matching semantics).

### 6.7.3.2 ucp\_tag\_message\_h

```
typedef struct ucp_recv_desc* ucp_tag_message_h
```

UCP Message descriptor is an opaque handle for a message returned by [ucp\\_tag\\_probe\\_nb](#). This handle can be passed to [ucp\\_tag\\_msg\\_rcv\\_nb](#) in order to receive the message data to a specific buffer.



### 6.7.3.3 ucp\_datatype\_t

```
typedef uint64_t ucp_datatype_t
```

UCP datatype identifier is a 64bit object used for datatype identification. Predefined UCP identifiers are defined by [ucp\\_dt\\_type](#).

### 6.7.3.4 ucp\_send\_callback\_t

```
typedef void(* ucp_send_callback_t) (void *request, ucs_status_t status)
```

This callback routine is invoked whenever the [send operation](#) is completed. It is important to note that the call-back is only invoked in a case when the operation cannot be completed in place.

#### Parameters

in	<i>request</i>	The completed send request.
in	<i>status</i>	Completion status. If the send operation was completed successfully UCS_OK is returned. If send operation was canceled UCS_ERR_CANCELED is returned. Otherwise, an <a href="#">error status</a> is returned.

### 6.7.3.5 ucp\_send\_nbx\_callback\_t

```
typedef void(* ucp_send_nbx_callback_t) (void *request, ucs_status_t status, void *user_data)
```

This callback routine is invoked whenever the [ucp\\_tag\\_send\\_nbx](#), [ucp\\_am\\_send\\_nbx](#), [ucp\\_stream\\_send\\_nbx](#), [ucp\\_put\\_nbx](#), [ucp\\_get\\_nbx](#), [ucp\\_atomic\\_op\\_nbx](#) or any other "send operation" is completed.

#### Parameters

in	<i>request</i>	The completed send request.
in	<i>status</i>	Completion status. If the send operation was completed successfully UCS_OK is returned. If send operation was canceled UCS_ERR_CANCELED is returned. Otherwise, an <a href="#">error status</a> is returned.
in	<i>user_data</i>	User data passed to "user_data" value, see <a href="#">ucp_request_param_t</a>

#### Examples

[ucp\\_client\\_server.c](#).

### 6.7.3.6 ucp\_err\_handler\_cb\_t

```
typedef void(* ucp_err_handler_cb_t) (void *arg, ucp_ep_h ep, ucs_status_t status)
```

This callback routine is invoked when transport level error detected.

#### Parameters

in	<i>arg</i>	User argument to be passed to the callback.
in	<i>ep</i>	Endpoint to handle transport level error. Upon return from the callback, this <i>ep</i> is no longer usable and all subsequent operations on this <i>ep</i> will fail with the error code passed in <i>status</i> .
in	<i>status</i>	<a href="#">error status</a> .

### 6.7.3.7 ucp\_err\_handler\_t

```
typedef struct ucp_err_handler ucp_err_handler_t
```

This structure should be initialized in `ucp_ep_params_t` to handle peer failure

### 6.7.3.8 ucp\_stream\_recv\_callback\_t

```
typedef void(* ucp_stream_recv_callback_t) (void *request, ucs_status_t status, size_t length)
```

This callback routine is invoked whenever the [receive operation](#) is completed and the data is ready in the receive buffer.

#### Parameters

in	<i>request</i>	The completed receive request.
in	<i>status</i>	Completion status. If the send operation was completed successfully UCS_OK is returned. Otherwise, an <a href="#">error status</a> is returned.
in	<i>length</i>	The size of the received data in bytes, always boundary of base datatype size. The value is valid only if the status is UCS_OK.

### 6.7.3.9 ucp\_stream\_recv\_nbx\_callback\_t

```
typedef void(* ucp_stream_recv_nbx_callback_t) (void *request, ucs_status_t status, size_t length, void *user_data)
```

This callback routine is invoked whenever the [receive operation](#) is completed and the data is ready in the receive buffer.

#### Parameters

in	<i>request</i>	The completed receive request.
in	<i>status</i>	Completion status. If the send operation was completed successfully UCS_OK is returned. Otherwise, an <a href="#">error status</a> is returned.
in	<i>length</i>	The size of the received data in bytes, always on the boundary of base datatype size. The value is valid only if the status is UCS_OK.
in	<i>user_data</i>	User data passed to "user_data" value, see <a href="#">ucp_request_param_t</a> .

### 6.7.3.10 ucp\_tag\_recv\_callback\_t

```
typedef void(* ucp_tag_recv_callback_t) (void *request, ucs_status_t status, ucp_tag_recv_info_t *info)
```

This callback routine is invoked whenever the [receive operation](#) is completed and the data is ready in the receive buffer.

#### Parameters

in	<i>request</i>	The completed receive request.
----	----------------	--------------------------------

## Parameters

in	<i>status</i>	Completion status. If the send operation was completed successfully UCS_OK is returned. If send operation was canceled UCS_ERR_CANCELED is returned. If the data can not fit into the receive buffer the <a href="#">UCS_ERR_MESSAGE_TRUNCATED</a> error code is returned. Otherwise, an <a href="#">error status</a> is returned.
in	<i>info</i>	<a href="#">Completion information</a> The <i>info</i> descriptor is Valid only if the status is UCS_OK.

6.7.3.11 `ucp_tag_rcv_nbx_callback_t`

```
typedef void(* ucp_tag_rcv_nbx_callback_t) (void *request, ucs_status_t status, const ucp_tag_rcv_info_t *tag_info, void *user_data)
```

This callback routine is invoked whenever the [receive operation](#) is completed and the data is ready in the receive buffer.

## Parameters

in	<i>request</i>	The completed receive request.
in	<i>status</i>	Completion status. If the receive operation was completed successfully UCS_OK is returned. If send operation was canceled, UCS_ERR_CANCELED is returned. If the data can not fit into the receive buffer the <a href="#">UCS_ERR_MESSAGE_TRUNCATED</a> error code is returned. Otherwise, an <a href="#">error status</a> is returned.
in	<i>tag_info</i>	<a href="#">Completion information</a> The <i>info</i> descriptor is Valid only if the status is UCS_OK.
in	<i>user_data</i>	User data passed to "user_data" value, see <a href="#">ucp_request_param_t</a>

6.7.3.12 `ucp_am_rcv_data_nbx_callback_t`

```
typedef void(* ucp_am_rcv_data_nbx_callback_t) (void *request, ucs_status_t status, size_t length, void *user_data)
```

This callback routine is invoked whenever the [receive operation](#) is completed and the data is ready in the receive buffer.

## Parameters

in	<i>request</i>	The completed receive request.
in	<i>status</i>	Completion status. If the receive operation was completed successfully UCS_OK is returned. Otherwise, an <a href="#">error status</a> is returned.
in	<i>length</i>	The size of the received data in bytes, always boundary of base datatype size. The value is valid only if the status is UCS_OK.
in	<i>user_data</i>	User data passed to "user_data" value, see <a href="#">ucp_request_param_t</a>

## 6.7.4 Enumeration Type Documentation

### 6.7.4.1 ucp\_atomic\_post\_op\_t

enum `ucp_atomic_post_op_t`

This enumeration defines which atomic memory operation should be performed by the `ucp_atomic_post` family of functions. All of these are non-fetching atomics and will not result in a request handle.

#### Enumerator

UCP_ATOMIC_POST_OP_ADD	Atomic add
UCP_ATOMIC_POST_OP_AND	Atomic and
UCP_ATOMIC_POST_OP_OR	Atomic or
UCP_ATOMIC_POST_OP_XOR	Atomic xor
UCP_ATOMIC_POST_OP_LAST	

### 6.7.4.2 ucp\_atomic\_fetch\_op\_t

enum `ucp_atomic_fetch_op_t`

This enumeration defines which atomic memory operation should be performed by the `ucp_atomic_fetch` family of functions. All of these functions will fetch data from the remote node.

#### Enumerator

UCP_ATOMIC_FETCH_OP_FADD	Atomic Fetch and add
UCP_ATOMIC_FETCH_OP_SWAP	Atomic swap
UCP_ATOMIC_FETCH_OP_CSWAP	Atomic conditional swap
UCP_ATOMIC_FETCH_OP_FAND	Atomic Fetch and and
UCP_ATOMIC_FETCH_OP_FOR	Atomic Fetch and or
UCP_ATOMIC_FETCH_OP_FXOR	Atomic Fetch and xor
UCP_ATOMIC_FETCH_OP_LAST	

### 6.7.4.3 ucp\_atomic\_op\_t

enum `ucp_atomic_op_t`

This enumeration defines which atomic memory operation should be performed by the `ucp_atomic_op_nbx` routine.

#### Enumerator

UCP_ATOMIC_OP_ADD	Atomic add
UCP_ATOMIC_OP_SWAP	Atomic swap
UCP_ATOMIC_OP_CSWAP	Atomic conditional swap
UCP_ATOMIC_OP_AND	Atomic and

## Enumerator

UCP_ATOMIC_OP_OR	Atomic or
UCP_ATOMIC_OP_XOR	Atomic xor
UCP_ATOMIC_OP_LAST	

6.7.4.4 `ucp_stream_recv_flags_t`

```
enum ucp_stream_recv_flags_t
```

This enumeration defines behavior of `ucp_stream_recv_nb` function.

## Enumerator

UCP_STREAM_RECV_FLAG_WAITALL	This flag requests that the operation will not be completed until all requested data is received and placed in the user buffer.
------------------------------	---

6.7.4.5 `ucp_op_attr_t`

```
enum ucp_op_attr_t
```

The enumeration allows specifying which fields in `ucp_request_param_t` are present and operation flags are used. It is used to enable backward compatibility support.

## Enumerator

UCP_OP_ATTR_FIELD_REQUEST	request field
UCP_OP_ATTR_FIELD_CALLBACK	cb field
UCP_OP_ATTR_FIELD_USER_DATA	user_data field
UCP_OP_ATTR_FIELD_DATATYPE	datatype field
UCP_OP_ATTR_FIELD_FLAGS	operation-specific flags
UCP_OP_ATTR_FIELD_REPLY_BUFFER	reply_buffer field
UCP_OP_ATTR_FIELD_MEMORY_TYPE	memory type field
UCP_OP_ATTR_FIELD_RECV_INFO	recv_info field
UCP_OP_ATTR_FIELD_MEMH	memory handle field
UCP_OP_ATTR_FLAG_NO_IMM_CMPL	deny immediate completion
UCP_OP_ATTR_FLAG_FAST_CMPL	expedite local completion, even if it delays remote data delivery. Note for implementer: this option can disable zero copy and/or rendezvous protocols which require synchronization with the remote peer before releasing the local send buffer
UCP_OP_ATTR_FLAG_FORCE_IMM_CMPL	force immediate complete operation, fail if the operation cannot be completed immediately
UCP_OP_ATTR_FLAG_MULTI_SEND	optimize for bandwidth of multiple in-flight operations, rather than for the latency of a single operation

#### 6.7.4.6 ucp\_req\_attr\_field

enum `ucp_req_attr_field`

The enumeration allows specifying which fields in `ucp_request_attr_t` are present. It is used to enable backward compatibility support.

Enumerator

UCP_REQUEST_ATTR_FIELD_INFO_STRING	
UCP_REQUEST_ATTR_FIELD_INFO_STRING_SIZE	
UCP_REQUEST_ATTR_FIELD_STATUS	
UCP_REQUEST_ATTR_FIELD_MEM_TYPE	

#### 6.7.4.7 ucp\_am\_recv\_attr\_t

enum `ucp_am_recv_attr_t`

The enumeration allows specifying which fields in `ucp_am_recv_param_t` are present and receive operation flags are used. It is used to enable backward compatibility support.

Enumerator

UCP_AM_RECV_ATTR_FIELD_REPLY_EP	reply_ep field
UCP_AM_RECV_ATTR_FLAG_DATA	Indicates that the data provided in <code>ucp_am_recv_callback_t</code> callback can be held by the user. If UCS_INPROGRESS is returned from the callback, the data parameter will persist and the user has to call <code>ucp_am_data_release</code> when data is no longer needed. This flag is mutually exclusive with <code>UCP_AM_RECV_ATTR_FLAG_RNDV</code> .
UCP_AM_RECV_ATTR_FLAG_RNDV	Indicates that the arriving data was sent using rendezvous protocol. In this case <code>data</code> parameter of the <code>ucp_am_recv_callback_t</code> points to the internal UCP descriptor, which can be used for obtaining the actual data by calling <code>ucp_am_recv_data_nbx</code> routine. This flag is mutually exclusive with <code>UCP_AM_RECV_ATTR_FLAG_DATA</code> .

#### 6.7.4.8 ucp\_am\_handler\_param\_field

enum `ucp_am_handler_param_field`

The enumeration allows specifying which fields in `ucp_am_handler_param_t` are present. It is used to enable backward compatibility support.

Enumerator

UCP_AM_HANDLER_PARAM_FIELD_ID	Indicates that <code>ucp_am_handler_param_t::id</code> field is valid.
UCP_AM_HANDLER_PARAM_FIELD_FLAGS	Indicates that <code>ucp_am_handler_param_t::flags</code> field is valid.
UCP_AM_HANDLER_PARAM_FIELD_CB	Indicates that <code>ucp_am_handler_param_t::cb</code> field is valid.
UCP_AM_HANDLER_PARAM_FIELD_ARG	Indicates that <code>ucp_am_handler_param_t::arg</code> field is valid.

## 6.7.5 Function Documentation

### 6.7.5.1 ucp\_request\_query()

```
ucs_status_t ucp_request_query (
    void * request,
    ucp_request_attr_t * attr )
```

#### Parameters

in	<i>request</i>	Non-blocking request to query.
out	<i>attr</i>	Filled with attributes of the request.

#### Returns

Error code as defined by [ucs\\_status\\_t](#)

### 6.7.5.2 ucp\_am\_send\_nb()

```
ucs_status_ptr_t ucp_am_send_nb (
    ucp_ep_h ep,
    uint16_t id,
    const void * buffer,
    size_t count,
    ucp_datatype_t datatype,
    ucp_send_callback_t cb,
    unsigned flags )
```

This routine sends an Active Message to an ep. It does not support CUDA memory.

#### Parameters

in	<i>ep</i>	UCP endpoint where the Active Message will be run.
in	<i>id</i>	Active Message id. Specifies which registered callback to run.
in	<i>buffer</i>	Pointer to the data to be sent to the target node of the Active Message.
in	<i>count</i>	Number of elements to send.
in	<i>datatype</i>	Datatype descriptor for the elements in the buffer.
in	<i>cb</i>	Callback that is invoked upon completion of the data transfer if it is not completed immediately.
in	<i>flags</i>	Operation flags as defined by <a href="#">ucp_send_am_flags</a> .

#### Returns

NULL Active Message was sent immediately.

UCS\_PTR\_IS\_ERR(\_ptr) Error sending Active Message.

otherwise Pointer to request, and Active Message is known to be completed after cb is run.

### 6.7.5.3 ucp\_am\_send\_nbx()

```
ucs_status_ptr_t ucp_am_send_nbx (
    ucp_ep_h ep,
    unsigned id,
    const void * header,
    size_t header_length,
    const void * buffer,
    size_t count,
    const ucp_request_param_t * param )
```

This routine sends an Active Message to an ep. If the operation completes immediately, then the routine returns NULL and the callback function is ignored, even if specified. Otherwise, if no error is reported and a callback is requested (i.e. the UCP\_OP\_ATTR\_FIELD\_CALLBACK flag is set in the op\_attr\_mask field of *param*), then the UCP library will schedule invocation of the callback routine *param->cb.send* upon completion of the operation.

#### Note

If UCP\_OP\_ATTR\_FLAG\_NO\_IMM\_CMPL flag is set in the op\_attr\_mask field of *param*, then the operation will return a request handle, even if it completes immediately.

This operation supports specific flags, which can be passed in *param* by [ucp\\_request\\_param\\_t::flags](#). The exact set of flags is defined by [ucp\\_send\\_am\\_flags](#).

#### Parameters

in	<i>ep</i>	UCP endpoint where the Active Message will be run.
in	<i>id</i>	Active Message id. Specifies which registered callback to run.
in	<i>header</i>	User defined Active Message header. NULL value is allowed if no header needed. In this case <i>header_length</i> should be set to 0.
in	<i>header_length</i>	Active message header length in bytes.
in	<i>buffer</i>	Pointer to the data to be sent to the target node of the Active Message.
in	<i>count</i>	Number of elements to send.
in	<i>param</i>	Operation parameters, see <a href="#">ucp_request_param_t</a> .

#### Note

Sending only header without actual data is allowed and is recommended for transferring a latency-critical amount of data.

The maximum allowed header size can be obtained by querying worker attributes by the [ucp\\_worker\\_query](#) routine.

#### Returns

NULL - Active Message was sent immediately.

UCS\_PTR\_IS\_ERR(\_ptr) - Error sending Active Message.

otherwise - Operation was scheduled for send and can be completed at any point in time. The request handle is returned to the application in order to track progress of the message. If user request was not provided in *param->request*, the application is responsible for releasing the handle using [ucp\\_request\\_free](#) routine.

#### Examples

[ucp\\_client\\_server.c](#).



#### 6.7.5.4 ucp\_am\_recv\_data\_nbx()

```
ucs_status_ptr_t ucp_am_recv_data_nbx (
    ucp_worker_h worker,
    void * data_desc,
    void * buffer,
    size_t count,
    const ucp_request_param_t * param )
```

This routine receives a message that is described by the data descriptor *data\_desc*, local address *buffer*, size *count* and *param* parameters on the *worker*. The routine is non-blocking and therefore returns immediately. The receive operation is considered completed when the message is delivered to the *buffer*. If the receive operation cannot be started the routine returns an error.

#### Note

This routine can be performed on any valid data descriptor delivered in [ucp\\_am\\_recv\\_callback\\_t](#). Data descriptor is considered to be valid if:

- It is a rendezvous request (*UCP\_AM\_RECV\_ATTR\_FLAG\_RNDV* is set in [ucp\\_am\\_recv\\_param\\_t::recv\\_attr](#)) or
- It is a persistent data pointer (*UCP\_AM\_RECV\_ATTR\_FLAG\_DATA* is set in [ucp\\_am\\_recv\\_param\\_t::recv\\_attr](#)). In this case receive operation may be needed to unpack data to device memory (for example GPU device) or some specific datatype.

After this call UCP takes ownership of *data\_desc* descriptor, so there is no need to release it even if the operation fails. The routine returns a request handle instead, which can be used for tracking operation progress.

#### Parameters

in	<i>worker</i>	Worker that is used for the receive operation.
in	<i>data_desc</i>	Data descriptor, provided in <a href="#">ucp_am_recv_callback_t</a> routine.
in	<i>buffer</i>	Pointer to the buffer to receive the data.
in	<i>count</i>	Number of elements to receive into <i>buffer</i> .
in	<i>param</i>	Operation parameters, see <a href="#">ucp_request_param_t</a> .

#### Returns

NULL - The receive operation was completed immediately. In this case, if *param->recv\_info.length* is specified in the *param*, the value to which it points is updated with the size of the received message.

UCS\_PTR\_IS\_ERR(\_ptr) - The receive operation failed.

otherwise - Receive operation was scheduled and can be completed at any point in time. The request handle is returned to the application in order to track operation progress. If user request was not provided in *param->request*, the application is responsible for releasing the handle using [ucp\\_request\\_free](#) routine.

#### Examples

[ucp\\_client\\_server.c](#).

#### 6.7.5.5 ucp\_am\_data\_release()

```
void ucp_am_data_release (
    ucp_worker_h worker,
    void * data )
```

This routine releases data that persisted through an Active Message callback because that callback returned UCS\_INPROGRESS.

#### Parameters

in	<i>worker</i>	Worker which received the Active Message.
in	<i>data</i>	Pointer to data that was passed into the Active Message callback as the data parameter.

#### 6.7.5.6 ucp\_stream\_send\_nb()

```
ucs_status_ptr_t ucp_stream_send_nb (
    ucp_ep_h ep,
    const void * buffer,
    size_t count,
    ucp_datatype_t datatype,
    ucp_send_callback_t cb,
    unsigned flags )
```

This routine sends data that is described by the local address *buffer*, size *count*, and *datatype* object to the destination endpoint *ep*. The routine is non-blocking and therefore returns immediately, however the actual send operation may be delayed. The send operation is considered completed when it is safe to reuse the source *buffer*. If the send operation is completed immediately the routine returns UCS\_OK and the callback function *cb* is **not** invoked. If the operation is **not** completed immediately and no error reported, then the UCP library will schedule invocation of the callback *cb* upon completion of the send operation. In other words, the completion of the operation will be signaled either by the return code or by the callback.

#### Note

The user should not modify any part of the *buffer* after this operation is called, until the operation completes.

#### Parameters

in	<i>ep</i>	Destination endpoint handle.
in	<i>buffer</i>	Pointer to the message buffer (payload).
in	<i>count</i>	Number of elements to send.
in	<i>datatype</i>	Datatype descriptor for the elements in the buffer.
in	<i>cb</i>	Callback function that is invoked whenever the send operation is completed. It is important to note that the callback is only invoked in the event that the operation cannot be completed in place.
in	<i>flags</i>	Reserved for future use.

#### Returns

NULL - The send operation was completed immediately.

UCS\_PTR\_IS\_ERR(\_ptr) - The send operation failed.

otherwise - Operation was scheduled for send and can be completed in any point in time. The request handle is returned to the application in order to track progress of the message. The application is responsible for releasing the handle using [ucp\\_request\\_free](#) routine.

### 6.7.5.7 ucp\_stream\_send\_nbx()

```
ucs_status_ptr_t ucp_stream_send_nbx (
    ucp_ep_h ep,
    const void * buffer,
    size_t count,
    const ucp_request_param_t * param )
```

This routine sends data that is described by the local address *buffer*, size *count* object to the destination endpoint *ep*. The routine is non-blocking and therefore returns immediately, however the actual send operation may be delayed. The send operation is considered completed when it is safe to reuse the source *buffer*. If the send operation is completed immediately the routine returns UCS\_OK.

#### Note

The user should not modify any part of the *buffer* after this operation is called, until the operation completes.

#### Parameters

in	<i>ep</i>	Destination endpoint handle.
in	<i>buffer</i>	Pointer to the message buffer (payload).
in	<i>count</i>	Number of elements to send.
in	<i>param</i>	Operation parameters, see <a href="#">ucp_request_param_t</a> .

#### Returns

NULL - The send operation was completed immediately.

UCS\_PTR\_IS\_ERR(\_ptr) - The send operation failed.

otherwise - Operation was scheduled for send and can be completed at any point in time. The request handle is returned to the application in order to track progress of the message.

#### Examples

[ucp\\_client\\_server.c](#).

### 6.7.5.8 ucp\_tag\_send\_nb()

```
ucs_status_ptr_t ucp_tag_send_nb (
    ucp_ep_h ep,
    const void * buffer,
    size_t count,
    ucp_datatype_t datatype,
    ucp_tag_t tag,
    ucp_send_callback_t cb )
```

This routine sends a messages that is described by the local address *buffer*, size *count*, and *datatype* object to the destination endpoint *ep*. Each message is associated with a *tag* value that is used for message matching on the [receiver](#). The routine is non-blocking and therefore returns immediately, however the actual send operation may be delayed. The send operation is considered completed when it is safe to reuse the source *buffer*. If the send operation is completed immediately the routine return UCS\_OK and the call-back function *cb* is **not** invoked. If the operation is **not** completed immediately and no error reported then the UCP library will schedule to invoke the call-back *cb* whenever the send operation will be completed. In other words, the completion of a message can be signaled by the return code or the call-back.

**Note**

The user should not modify any part of the *buffer* after this operation is called, until the operation completes.

**Parameters**

in	<i>ep</i>	Destination endpoint handle.
in	<i>buffer</i>	Pointer to the message buffer (payload).
in	<i>count</i>	Number of elements to send
in	<i>datatype</i>	Datatype descriptor for the elements in the buffer.
in	<i>tag</i>	Message tag.
in	<i>cb</i>	Callback function that is invoked whenever the send operation is completed. It is important to note that the call-back is only invoked in a case when the operation cannot be completed in place.

**Returns**

NULL - The send operation was completed immediately.

UCS\_PTR\_IS\_ERR(\_ptr) - The send operation failed.

otherwise - Operation was scheduled for send and can be completed in any point in time. The request handle is returned to the application in order to track progress of the message. The application is responsible for releasing the handle using [ucp\\_request\\_free\(\)](#) routine.

**6.7.5.9 ucp\_tag\_send\_nbr()**

```
ucs_status_t ucp_tag_send_nbr (
    ucp_ep_h ep,
    const void * buffer,
    size_t count,
    ucp_datatype_t datatype,
    ucp_tag_t tag,
    void * req )
```

This routine provides a convenient and efficient way to implement a blocking send pattern. It also completes requests faster than [ucp\\_tag\\_send\\_nb\(\)](#) because:

- it always uses eager protocol to send data up to the rendezvous threshold.
- its rendezvous threshold is higher than the one used by the [ucp\\_tag\\_send\\_nb\(\)](#). The threshold is controlled by the **UCX\_SEND\_NBR\_RNDV\_THRESH** environment variable.
- its request handling is simpler. There is no callback and no need to allocate and free requests. In fact request can be allocated by caller on the stack.

This routine sends a messages that is described by the local address *buffer*, size *count*, and *datatype* object to the destination endpoint *ep*. Each message is associated with a *tag* value that is used for message matching on the [receiver](#).

The routine is non-blocking and therefore returns immediately, however the actual send operation may be delayed. The send operation is considered completed when it is safe to reuse the source *buffer*. If the send operation is completed immediately the routine returns UCS\_OK.

If the operation is **not** completed immediately and no error reported then the UCP library will fill a user provided *req* and return UCS\_INPROGRESS status. In order to monitor completion of the operation [ucp\\_request\\_check\\_status\(\)](#) should be used.

Following pseudo code implements a blocking send function:

```
MPI_send(...)
{
    char *request;
    ucs_status_t status;
    // allocate request on the stack
    // ucp_context_query() was used to get ucp_request_size
    request = alloca(ucp_request_size);
    // note: make sure that there is enough memory before the
    // request handle
    status = ucp_tag_send_nbr(ep, ..., request + ucp_request_size);
    if (status != UCS_INPROGRESS) {
        return status;
    }
    do {
        ucp_worker_progress(worker);
        status = ucp_request_check_status(request + ucp_request_size);
    } while (status == UCS_INPROGRESS);
    return status;
}
```

#### Note

The user should not modify any part of the *buffer* after this operation is called, until the operation completes.

#### Parameters

in	<i>ep</i>	Destination endpoint handle.
in	<i>buffer</i>	Pointer to the message buffer (payload).
in	<i>count</i>	Number of elements to send
in	<i>datatype</i>	Datatype descriptor for the elements in the buffer.
in	<i>tag</i>	Message tag.
in	<i>req</i>	Request handle allocated by the user. There should be at least UCP request size bytes of available space before the <i>req</i> . The size of UCP request can be obtained by <a href="#">ucp_context_query</a> function.

#### Returns

UCS\_OK - The send operation was completed immediately.

UCS\_INPROGRESS - The send was not completed and is in progress. [ucp\\_request\\_check\\_status\(\)](#) should be used to monitor *req* status.

Error code as defined by [ucs\\_status\\_t](#)

#### 6.7.5.10 ucp\_tag\_send\_sync\_nb()

```
ucs_status_ptr_t ucp_tag_send_sync_nb (
    ucp_ep_h ep,
    const void * buffer,
    size_t count,
    ucp_datatype_t datatype,
    ucp_tag_t tag,
    ucp_send_callback_t cb )
```

Same as [ucp\\_tag\\_send\\_nb](#), except the request completes only after there is a remote tag match on the message (which does not always mean the remote receive has been completed). This function never completes "in-place", and always returns a request handle.

**Note**

The user should not modify any part of the *buffer* after this operation is called, until the operation completes.

Returns `UCS_ERR_UNSUPPORTED` if `UCP_ERR_HANDLING_MODE_PEER` is enabled. This is a temporary implementation-related constraint that will be addressed in future releases.

**Parameters**

in	<i>ep</i>	Destination endpoint handle.
in	<i>buffer</i>	Pointer to the message buffer (payload).
in	<i>count</i>	Number of elements to send
in	<i>datatype</i>	Datatype descriptor for the elements in the buffer.
in	<i>tag</i>	Message tag.
in	<i>cb</i>	Callback function that is invoked whenever the send operation is completed.

**Returns**

`UCS_PTR_IS_ERR(_ptr)` - The send operation failed.

otherwise - Operation was scheduled for send and can be completed in any point in time. The request handle is returned to the application in order to track progress of the message. The application is responsible for releasing the handle using `ucp_request_free()` routine.

**6.7.5.11 ucp\_tag\_send\_nbx()**

```
ucs_status_ptr_t ucp_tag_send_nbx (
    ucp_ep_h ep,
    const void * buffer,
    size_t count,
    ucp_tag_t tag,
    const ucp_request_param_t * param )
```

This routine sends a messages that is described by the local address *buffer*, size *count* object to the destination endpoint *ep*. Each message is associated with a *tag* value that is used for message matching on the `ucp_tag_recv_nb` or `receiver`. The routine is non-blocking and therefore returns immediately, however the actual send operation may be delayed. The send operation is considered completed when it is safe to reuse the source *buffer*. If the send operation is completed immediately the routine returns `UCS_OK` and the call-back function is **not** invoked. If the operation is **not** completed immediately and no error reported then the UCP library will schedule to invoke the call-back whenever the send operation is completed. In other words, the completion of a message can be signaled by the return code or the call-back. Immediate completion signals can be fine-tuned via the `ucp_request_param_t::op_attr_mask` field in the `ucp_request_param_t` structure. The values of this field are a bit-wise OR of the `ucp_op_attr_t` enumeration.

**Note**

The user should not modify any part of the *buffer* after this operation is called, until the operation completes.

**Parameters**

in	<i>ep</i>	Destination endpoint handle.
in	<i>buffer</i>	Pointer to the message buffer (payload).
in	<i>count</i>	Number of elements to send
in	<i>tag</i>	Message tag.
in	<i>param</i>	Operation parameters, see <code>ucp_request_param_t</code>

**Returns**

UCS\_OK - The send operation was completed immediately.

UCS\_PTR\_IS\_ERR(\_ptr) - The send operation failed.

otherwise - Operation was scheduled for send and can be completed in any point in time. The request handle is returned to the application in order to track progress of the message.

**Examples**

[ucp\\_client\\_server.c](#), and [ucp\\_hello\\_world.c](#).

**6.7.5.12 ucp\_tag\_send\_sync\_nbx()**

```
ucs_status_ptr_t ucp_tag_send_sync_nbx (
    ucp_ep_h ep,
    const void * buffer,
    size_t count,
    ucp_tag_t tag,
    const ucp_request_param_t * param )
```

Same as [ucp\\_tag\\_send\\_nbx](#), except the request completes only after there is a remote tag match on the message (which does not always mean the remote receive has been completed). This function never completes "in-place", and always returns a request handle.

**Note**

The user should not modify any part of the *buffer* after this operation is called, until the operation completes.

Returns [UCS\\_ERR\\_UNSUPPORTED](#) if [UCP\\_ERR\\_HANDLING\\_MODE\\_PEER](#) is enabled. This is a temporary implementation-related constraint that will be addressed in future releases.

**Parameters**

in	<i>ep</i>	Destination endpoint handle.
in	<i>buffer</i>	Pointer to the message buffer (payload).
in	<i>count</i>	Number of elements to send
in	<i>tag</i>	Message tag.
in	<i>param</i>	Operation parameters, see <a href="#">ucp_request_param_t</a>

**Returns**

UCS\_OK - The send operation was completed immediately.

UCS\_PTR\_IS\_ERR(\_ptr) - The send operation failed.

otherwise - Operation was scheduled for send and can be completed in any point in time. The request handle is returned to the application in order to track progress of the message.

**6.7.5.13 ucp\_stream\_recv\_nb()**

```
ucs_status_ptr_t ucp_stream_recv_nb (
    ucp_ep_h ep,
    void * buffer,
```

```

size_t count,
ucp_datatype_t datatype,
ucp_stream_recv_callback_t cb,
size_t * length,
unsigned flags )

```

This routine receives data that is described by the local address *buffer*, *size count*, and *datatype* object on the endpoint *ep*. The routine is non-blocking and therefore returns immediately. The receive operation is considered complete when the message is delivered to the buffer. If data is not immediately available, the operation will be scheduled for receive and a request handle will be returned. In order to notify the application about completion of a scheduled receive operation, the UCP library will invoke the call-back *cb* when data is in the receive buffer and ready for application access. If the receive operation cannot be started, the routine returns an error.

#### Parameters

in	<i>ep</i>	UCP endpoint that is used for the receive operation.
in	<i>buffer</i>	Pointer to the buffer to receive the data.
in	<i>count</i>	Number of elements to receive into <i>buffer</i> .
in	<i>datatype</i>	Datatype descriptor for the elements in the buffer.
in	<i>cb</i>	Callback function that is invoked whenever the receive operation is completed and the data is ready in the receive <i>buffer</i> . It is important to note that the call-back is only invoked in a case when the operation cannot be completed immediately.
out	<i>length</i>	Size of the received data in bytes. The value is valid only if return code is UCS_OK.

#### Note

The amount of data received, in bytes, is always an integral multiple of the *datatype* size.

#### Parameters

in	<i>flags</i>	Flags defined in <a href="#">ucp_stream_recv_flags_t</a> .
----	--------------	--

#### Returns

NULL - The receive operation was completed immediately.

UCS\_PTR\_IS\_ERR(*ptr*) - The receive operation failed.

otherwise - Operation was scheduled for receive. A request handle is returned to the application in order to track progress of the operation. The application is responsible for releasing the handle by calling the [ucp\\_request\\_free](#) routine.

#### 6.7.5.14 ucp\_stream\_recv\_nbx()

```

ucs_status_ptr_t ucp_stream_recv_nbx (
    ucp_ep_h ep,
    void * buffer,
    size_t count,
    size_t * length,
    const ucp_request_param_t * param )

```

This routine receives data that is described by the local address *buffer*, *size count* object on the endpoint *ep*. The routine is non-blocking and therefore returns immediately. The receive operation is considered complete when the message is delivered to the buffer. If the receive operation cannot be started, the routine returns an error.



## Parameters

in	<i>ep</i>	UCP endpoint that is used for the receive operation.
in	<i>buffer</i>	Pointer to the buffer that will receive the data.
in	<i>count</i>	Number of elements to receive into <i>buffer</i> .
out	<i>length</i>	Size of the received data in bytes. The value is valid only if return code is NULL.
in	<i>param</i>	Operation parameters, see <a href="#">ucp_request_param_t</a> . This operation supports specific flags, which can be passed in <i>param</i> by <a href="#">ucp_request_param_t::flags</a> . The exact set of flags is defined by <a href="#">ucp_stream_rcv_flags_t</a> .

## Returns

NULL - The receive operation was completed immediately. In this case the value pointed by *length* is updated by the size of received data. Note *param->rcv\_info* is not relevant for this function.

UCS\_PTR\_IS\_ERR(\_ptr) - The receive operation failed.

otherwise - Operation was scheduled for receive. A request handle is returned to the application in order to track progress of the operation.

## Note

The amount of data received, in bytes, is always an integral multiple of the *datatype* size.

## Examples

[ucp\\_client\\_server.c](#).

## 6.7.5.15 ucp\_stream\_rcv\_data\_nb()

```
ucs_status_ptr_t ucp_stream_rcv_data_nb (
    ucp_ep_h ep,
    size_t * length )
```

This routine receives any available data from endpoint *ep*. Unlike [ucp\\_stream\\_rcv\\_nb](#), the returned data is unstructured and is treated as an array of bytes. If data is immediately available, UCS\_STATUS\_PTR(\_ptr) is returned as a pointer to the data, and *length* is set to the size of the returned data buffer. The routine is non-blocking and therefore returns immediately.

## Parameters

in	<i>ep</i>	UCP endpoint that is used for the receive operation.
out	<i>length</i>	Length of received data.

## Returns

NULL - No received data available on the *ep*.

UCS\_PTR\_IS\_ERR(\_ptr) - the receive operation failed and UCS\_PTR\_STATUS(\_ptr) indicates an error.

otherwise - The pointer to the data UCS\_STATUS\_PTR(\_ptr) is returned to the application. After the data is processed, the application is responsible for releasing the data buffer by calling the [ucp\\_stream\\_data\\_release](#) routine.

**Note**

This function returns packed data (equivalent to [ucp\\_dt\\_make\\_contig\(1\)](#)).

This function returns a pointer to a UCP-supplied buffer, whereas [ucp\\_stream\\_recv\\_nb](#) places the data into a user-provided buffer. In some cases, receiving data directly into a UCP-supplied buffer can be more optimal, for example by processing the incoming data in-place and thus avoiding extra memory copy operations.

**6.7.5.16 ucp\_tag\_recv\_nb()**

```
ucs_status_ptr_t ucp_tag_recv_nb (
    ucp_worker_h worker,
    void * buffer,
    size_t count,
    ucp_datatype_t datatype,
    ucp_tag_t tag,
    ucp_tag_t tag_mask,
    ucp_tag_recv_callback_t cb )
```

This routine receives a message that is described by the local address *buffer*, size *count*, and *datatype* object on the *worker*. The tag value of the receive message has to match the *tag* and *tag\_mask* values, where the *tag\_mask* indicates which bits of the tag have to be matched. The routine is non-blocking and therefore returns immediately. The receive operation is considered completed when the message is delivered to the *buffer*. In order to notify the application about completion of the receive operation the UCP library will invoke the call-back *cb* when the received message is in the receive buffer and ready for application access. If the receive operation cannot be stated the routine returns an error.

**Note**

This routine cannot return UCS\_OK. It always returns a request handle or an error.

**Parameters**

in	<i>worker</i>	UCP worker that is used for the receive operation.
in	<i>buffer</i>	Pointer to the buffer to receive the data.
in	<i>count</i>	Number of elements to receive
in	<i>datatype</i>	Datatype descriptor for the elements in the buffer.
in	<i>tag</i>	Message tag to expect.
in	<i>tag_mask</i>	Bit mask that indicates the bits that are used for the matching of the incoming tag against the expected tag.
in	<i>cb</i>	Callback function that is invoked whenever the receive operation is completed and the data is ready in the receive <i>buffer</i> .

**Returns**

UCS\_PTR\_IS\_ERR(\_ptr) - The receive operation failed.

otherwise - Operation was scheduled for receive. The request handle is returned to the application in order to track progress of the operation. The application is responsible for releasing the handle using [ucp\\_request\\_free\(\)](#) routine.

**6.7.5.17 ucp\_tag\_recv\_nbr()**

```
ucs_status_t ucp_tag_recv_nbr (
```

```

ucp_worker_h worker,
void * buffer,
size_t count,
ucp_datatype_t datatype,
ucp_tag_t tag,
ucp_tag_t tag_mask,
void * req )

```

This routine receives a message that is described by the local address *buffer*, size *count*, and *datatype* object on the *worker*. The tag value of the receive message has to match the *tag* and *tag\_mask* values, where the *tag\_mask* indicates which bits of the tag have to be matched. The routine is non-blocking and therefore returns immediately. The receive operation is considered completed when the message is delivered to the *buffer*. In order to monitor completion of the operation [ucp\\_request\\_check\\_status](#) or [ucp\\_tag\\_rcv\\_request\\_test](#) should be used.

#### Parameters

in	<i>worker</i>	UCP worker that is used for the receive operation.
in	<i>buffer</i>	Pointer to the buffer to receive the data.
in	<i>count</i>	Number of elements to receive
in	<i>datatype</i>	Datatype descriptor for the elements in the buffer.
in	<i>tag</i>	Message tag to expect.
in	<i>tag_mask</i>	Bit mask that indicates the bits that are used for the matching of the incoming tag against the expected tag.
in	<i>req</i>	Request handle allocated by the user. There should be at least UCP request size bytes of available space before the <i>req</i> . The size of UCP request can be obtained by <a href="#">ucp_context_query</a> function.

#### Returns

Error code as defined by [ucs\\_status\\_t](#)

#### 6.7.5.18 ucp\_tag\_rcv\_nbx()

```

ucs_status_ptr_t ucp_tag_rcv_nbx (
    ucp_worker_h worker,
    void * buffer,
    size_t count,
    ucp_tag_t tag,
    ucp_tag_t tag_mask,
    const ucp_request_param_t * param )

```

This routine receives a message that is described by the local address *buffer*, size *count*, and *info* object on the *worker*. The tag value of the receive message has to match the *tag* and *tag\_mask* values, where the *tag\_mask* indicates what bits of the tag have to be matched. The routine is a non-blocking and therefore returns immediately. The receive operation is considered completed when the message is delivered to the *buffer*. In order to notify the application about completion of the receive operation the UCP library will invoke the call-back *cb* when the received message is in the receive buffer and ready for application access. If the receive operation cannot be stated the routine returns an error.

#### Parameters

in	<i>worker</i>	UCP worker that is used for the receive operation.
in	<i>buffer</i>	Pointer to the buffer to receive the data.
in	<i>count</i>	Number of elements to receive
in	<i>tag</i>	Message tag to expect.

## Parameters

in	<i>tag_mask</i>	Bit mask that indicates the bits that are used for the matching of the incoming tag against the expected tag.
in	<i>param</i>	Operation parameters, see <a href="#">ucp_request_param_t</a>

## Returns

NULL - The receive operation was completed immediately. In this case, if *param->recv\_info.tag\_info* is specified in the *param*, the value to which it points is updated with the information about the received message.

UCS\_PTR\_IS\_ERR(\_ptr) - The receive operation failed.

otherwise - Operation was scheduled for receive. The request handle is returned to the application in order to track progress of the operation. The application is responsible for releasing the handle using [ucp\\_request\\_free\(\)](#) routine.

## Examples

[ucp\\_client\\_server.c](#).

## 6.7.5.19 ucp\_tag\_probe\_nb()

```
ucp_tag_message_h ucp_tag_probe_nb (
    ucp_worker_h worker,
    ucp_tag_t tag,
    ucp_tag_t tag_mask,
    int remove,
    ucp_tag_recv_info_t * info )
```

This routine probes (checks) if a messages described by the *tag* and *tag\_mask* was received (fully or partially) on the *worker*. The tag value of the received message has to match the *tag* and *tag\_mask* values, where the *tag\_mask* indicates what bits of the tag have to be matched. The function returns immediately and if the message is matched it returns a handle for the message.

## Parameters

in	<i>worker</i>	UCP worker that is used for the probe operation.
in	<i>tag</i>	Message tag to probe for.
in	<i>tag_mask</i>	Bit mask that indicates the bits that are used for the matching of the incoming tag against the expected tag.
in	<i>remove</i>	The flag indicates if the matched message has to be removed from UCP library. If true (1), the message handle is removed from the UCP library and the application is responsible to call <a href="#">ucp_tag_msg_recv_nb()</a> in order to receive the data and release the resources associated with the message handle. If false (0), the return value is merely an indication to whether a matching message is present, and it cannot be used in any other way, and in particular it cannot be passed to <a href="#">ucp_tag_msg_recv_nb()</a> .
out	<i>info</i>	If the matching message is found the descriptor is filled with the details about the message.

## Returns

NULL - No match found.

Message handle (not NULL) - If message is matched the message handle is returned.

**Note**

This function does not advance the communication state of the network. If this routine is used in busy-poll mode, need to make sure `ucp_worker_progress()` is called periodically to extract messages from the transport.

**Examples**

[ucp\\_hello\\_world.c](#).

**6.7.5.20 ucp\_tag\_msg\_rcv\_nb()**

```
ucs_status_ptr_t ucp_tag_msg_rcv_nb (
    ucp_worker_h worker,
    void * buffer,
    size_t count,
    ucp_datatype_t datatype,
    ucp_tag_message_h message,
    ucp_tag_rcv_callback_t cb )
```

This routine receives a message that is described by the local address *buffer*, size *count*, *message* handle, and *datatype* object on the *worker*. The *message* handle can be obtained by calling the `ucp_tag_probe_nb()` routine. The `ucp_tag_msg_rcv_nb()` routine is non-blocking and therefore returns immediately. The receive operation is considered completed when the message is delivered to the *buffer*. In order to notify the application about completion of the receive operation the UCP library will invoke the call-back *cb* when the received message is in the receive buffer and ready for application access. If the receive operation cannot be started the routine returns an error.

**Parameters**

in	<i>worker</i>	UCP worker that is used for the receive operation.
in	<i>buffer</i>	Pointer to the buffer that will receive the data.
in	<i>count</i>	Number of elements to receive
in	<i>datatype</i>	Datatype descriptor for the elements in the buffer.
in	<i>message</i>	Message handle.
in	<i>cb</i>	Callback function that is invoked whenever the receive operation is completed and the data is ready in the receive <i>buffer</i> .

**Returns**

UCS\_PTR\_IS\_ERR(\_ptr) - The receive operation failed.

otherwise - Operation was scheduled for receive. The request handle is returned to the application in order to track progress of the operation. The application is responsible for releasing the handle using `ucp_request_free()` routine.

**Examples**

[ucp\\_hello\\_world.c](#).

**6.7.5.21 ucp\_tag\_msg\_rcv\_nbx()**

```
ucs_status_ptr_t ucp_tag_msg_rcv_nbx (
    ucp_worker_h worker,
    void * buffer,
```

```

size_t count,
ucp_tag_message_h message,
const ucp_request_param_t * param )

```

This routine receives a message that is described by the local address *buffer*, size *count*, and *message* handle on the *worker*. The *message* handle can be obtained by calling the [ucp\\_tag\\_probe\\_nb\(\)](#) routine. The [ucp\\_tag\\_msg\\_rcv\\_nbx\(\)](#) routine is non-blocking and therefore returns immediately. The receive operation is considered completed when the message is delivered to the *buffer*. In order to notify the application about completion of the receive operation the UCP library will invoke the call-back *cb* when the received message is in the receive buffer and ready for application access. If the receive operation cannot be started the routine returns an error.

#### Parameters

in	<i>worker</i>	UCP worker that is used for the receive operation.
in	<i>buffer</i>	Pointer to the buffer that will receive the data.
in	<i>count</i>	Number of elements to receive
in	<i>message</i>	Message handle.
in	<i>param</i>	Operation parameters, see <a href="#">ucp_request_param_t</a>

#### Returns

UCS\_PTR\_IS\_ERR(\_ptr) - The receive operation failed.

otherwise - Operation was scheduled for receive. The request handle is returned to the application in order to track progress of the operation. The application is responsible for releasing the handle using [ucp\\_request\\_free\(\)](#) routine.

#### 6.7.5.22 ucp\_put\_nbi()

```

ucs_status_t ucp_put_nbi (
    ucp_ep_h ep,
    const void * buffer,
    size_t length,
    uint64_t remote_addr,
    ucp_rkey_h rkey )

```

This routine initiates a storage of contiguous block of data that is described by the local address *buffer* in the remote contiguous memory region described by *remote\_addr* address and the [memory](#) handle *rkey*. The routine returns immediately and **does not** guarantee re-usability of the source address *buffer*. If the operation is completed immediately the routine return UCS\_OK, otherwise UCS\_INPROGRESS or an error is returned to user.

#### Note

A user can use [ucp\\_worker\\_flush\\_nb\(\)](#) in order to guarantee re-usability of the source address *buffer*.

#### Parameters

in	<i>ep</i>	Remote endpoint handle.
in	<i>buffer</i>	Pointer to the local source address.
in	<i>length</i>	Length of the data (in bytes) stored under the source address.
in	<i>remote_addr</i>	Pointer to the destination remote memory address to write to.
in	<i>rkey</i>	Remote memory key associated with the remote memory address.

## Returns

Error code as defined by [ucs\\_status\\_t](#)

## 6.7.5.23 ucp\_put\_nb()

```
ucs_status_ptr_t ucp_put_nb (
    ucp_ep_h ep,
    const void * buffer,
    size_t length,
    uint64_t remote_addr,
    ucp_rkey_h rkey,
    ucp_send_callback_t cb )
```

This routine initiates a storage of contiguous block of data that is described by the local address *buffer* in the remote contiguous memory region described by *remote\_addr* address and the [memory handle](#) *rkey*. The routine returns immediately and **does not** guarantee re-usability of the source address *buffer*. If the operation is completed immediately the routine return UCS\_OK, otherwise UCS\_INPROGRESS or an error is returned to user. If the put operation completes immediately, the routine returns UCS\_OK and the call-back routine *cb* is **not** invoked. If the operation is **not** completed immediately and no error is reported, then the UCP library will schedule invocation of the call-back routine *cb* upon completion of the put operation. In other words, the completion of a put operation can be signaled by the return code or execution of the call-back.

## Note

A user can use [ucp\\_worker\\_flush\\_nb\(\)](#) in order to guarantee re-usability of the source address *buffer*.

## Parameters

in	<i>ep</i>	Remote endpoint handle.
in	<i>buffer</i>	Pointer to the local source address.
in	<i>length</i>	Length of the data (in bytes) stored under the source address.
in	<i>remote_addr</i>	Pointer to the destination remote memory address to write to.
in	<i>rkey</i>	Remote memory key associated with the remote memory address.
in	<i>cb</i>	Call-back function that is invoked whenever the put operation is completed and the local buffer can be modified. Does not guarantee remote completion.

## Returns

NULL - The operation was completed immediately.

UCS\_PTR\_IS\_ERR(\_ptr) - The operation failed.

otherwise - Operation was scheduled and can be completed at any point in time. The request handle is returned to the application in order to track progress of the operation. The application is responsible for releasing the handle using [ucp\\_request\\_free\(\)](#) routine.

## 6.7.5.24 ucp\_put\_nbx()

```
ucs_status_ptr_t ucp_put_nbx (
    ucp_ep_h ep,
    const void * buffer,
    size_t count,
```

```
uint64_t remote_addr,
ucp_rkey_h rkey,
const ucp_request_param_t * param )
```

This routine initiates a storage of contiguous block of data that is described by the local address *buffer* in the remote contiguous memory region described by *remote\_addr* address and the [memory handle](#) *rkey*. The routine returns immediately and **does not** guarantee re-usability of the source address *buffer*. If the operation is completed immediately the routine return UCS\_OK, otherwise UCS\_INPROGRESS or an error is returned to user. If the put operation completes immediately, the routine returns UCS\_OK and the call-back routine *param.cb.send* is **not** invoked. If the operation is **not** completed immediately and no error is reported, then the UCP library will schedule invocation of the call-back routine *param.cb.send* upon completion of the put operation. In other words, the completion of a put operation can be signaled by the return code or execution of the call-back. Immediate completion signals can be fine-tuned via the [ucp\\_request\\_param\\_t::op\\_attr\\_mask](#) field in the [ucp\\_request\\_param\\_t](#) structure. The values of this field are a bit-wise OR of the [ucp\\_op\\_attr\\_t](#) enumeration.

#### Note

A user can use [ucp\\_worker\\_flush\\_nb\(\)](#) in order to guarantee re-usability of the source address *buffer*.

#### Parameters

in	<i>ep</i>	Remote endpoint handle.
in	<i>buffer</i>	Pointer to the local source address.
in	<i>count</i>	Number of elements of type <a href="#">ucp_request_param_t::datatype</a> to put. If <a href="#">ucp_request_param_t::datatype</a> is not specified, the type defaults to <a href="#">ucp_dt_make_contig(1)</a> , which corresponds to byte elements.
in	<i>remote_addr</i>	Pointer to the destination remote memory address to write to.
in	<i>rkey</i>	Remote memory key associated with the remote memory address.
in	<i>param</i>	Operation parameters, see <a href="#">ucp_request_param_t</a>

#### Returns

UCS\_OK - The operation was completed immediately.

UCS\_PTR\_IS\_ERR(\_ptr) - The operation failed.

otherwise - Operation was scheduled and can be completed at any point in time. The request handle is returned to the application in order to track progress of the operation. The application is responsible for releasing the handle using [ucp\\_request\\_free\(\)](#) routine.

#### Note

Only the datatype [ucp\\_dt\\_make\\_contig\(1\)](#) is supported for *param->datatype*, see [ucp\\_dt\\_make\\_contig](#).

### 6.7.5.25 ucp\_get\_nbi()

```
ucs_status_t ucp_get_nbi (
    ucp_ep_h ep,
    void * buffer,
    size_t length,
    uint64_t remote_addr,
    ucp_rkey_h rkey )
```

This routine initiate a load of contiguous block of data that is described by the remote memory address *remote\_addr* and the [memory handle](#) *rkey* in the local contiguous memory region described by *buffer* address. The routine returns immediately and **does not** guarantee that remote data is loaded and stored under the local address *buffer*.



**Note**

A user can use `ucp_worker_flush_nb()` in order guarantee that remote data is loaded and stored under the local address *buffer*.

**Parameters**

in	<i>ep</i>	Remote endpoint handle.
in	<i>buffer</i>	Pointer to the local destination address.
in	<i>length</i>	Length of the data (in bytes) stored under the destination address.
in	<i>remote_addr</i>	Pointer to the source remote memory address to read from.
in	<i>rkey</i>	Remote memory key associated with the remote memory address.

**Returns**

Error code as defined by `ucs_status_t`

**6.7.5.26 ucp\_get\_nb()**

```
ucs_status_ptr_t ucp_get_nb (
    ucp_ep_h ep,
    void * buffer,
    size_t length,
    uint64_t remote_addr,
    ucp_rkey_h rkey,
    ucp_send_callback_t cb )
```

This routine initiates a load of a contiguous block of data that is described by the remote memory address *remote\_addr* and the [memory handle](#) *rkey* in the local contiguous memory region described by *buffer* address. The routine returns immediately and **does not** guarantee that remote data is loaded and stored under the local address *buffer*. If the operation is completed immediately the routine return UCS\_OK, otherwise UCS\_INPROGRESS or an error is returned to user. If the get operation completes immediately, the routine returns UCS\_OK and the call-back routine *cb* is **not** invoked. If the operation is **not** completed immediately and no error is reported, then the UCP library will schedule invocation of the call-back routine *cb* upon completion of the get operation. In other words, the completion of a get operation can be signaled by the return code or execution of the call-back.

**Note**

A user can use `ucp_worker_flush_nb()` in order to guarantee re-usability of the source address *buffer*.

**Parameters**

in	<i>ep</i>	Remote endpoint handle.
in	<i>buffer</i>	Pointer to the local destination address.
in	<i>length</i>	Length of the data (in bytes) stored under the destination address.
in	<i>remote_addr</i>	Pointer to the source remote memory address to read from.
in	<i>rkey</i>	Remote memory key associated with the remote memory address.
in	<i>cb</i>	Call-back function that is invoked whenever the get operation is completed and the data is visible to the local process.

**Returns**

NULL - The operation was completed immediately.

UCS\_PTR\_IS\_ERR(\_ptr) - The operation failed.

otherwise - Operation was scheduled and can be completed at any point in time. The request handle is returned to the application in order to track progress of the operation. The application is responsible for releasing the handle using [ucp\\_request\\_free\(\)](#) routine.

**6.7.5.27 ucp\_get\_nbx()**

```
ucs_status_ptr_t ucp_get_nbx (
    ucp_ep_h ep,
    void * buffer,
    size_t count,
    uint64_t remote_addr,
    ucp_rkey_h rkey,
    const ucp_request_param_t * param )
```

This routine initiates a load of a contiguous block of data that is described by the remote memory address *remote\_addr* and the **memory handle** *rkey* in the local contiguous memory region described by *buffer* address. The routine returns immediately and **does not** guarantee that remote data is loaded and stored under the local address *buffer*. If the operation is completed immediately the routine return UCS\_OK, otherwise UCS\_INPROGRESS or an error is returned to user. If the get operation completes immediately, the routine returns UCS\_OK and the call-back routine *param.cb.send* is **not** invoked. If the operation is **not** completed immediately and no error is reported, then the UCP library will schedule invocation of the call-back routine *param.cb.send* upon completion of the get operation. In other words, the completion of a get operation can be signaled by the return code or execution of the call-back.

**Note**

A user can use [ucp\\_worker\\_flush\\_nb\(\)](#) in order to guarantee re-usability of the source address *buffer*.

**Parameters**

in	<i>ep</i>	Remote endpoint handle.
in	<i>buffer</i>	Pointer to the local destination address.
in	<i>count</i>	Number of elements of type <a href="#">ucp_request_param_t::datatype</a> to put. If <a href="#">ucp_request_param_t::datatype</a> is not specified, the type defaults to <a href="#">ucp_dt_make_contig(1)</a> , which corresponds to byte elements.
in	<i>remote_addr</i>	Pointer to the source remote memory address to read from.
in	<i>rkey</i>	Remote memory key associated with the remote memory address.
in	<i>param</i>	Operation parameters, see <a href="#">ucp_request_param_t</a> .

**Returns**

UCS\_OK - The operation was completed immediately.

UCS\_PTR\_IS\_ERR(\_ptr) - The operation failed.

otherwise - Operation was scheduled and can be completed at any point in time. The request handle is returned to the application in order to track progress of the operation. The application is responsible for releasing the handle using [ucp\\_request\\_free\(\)](#) routine.

## Note

Only the datatype `ucp_dt_make_contig(1)` is supported for `param->datatype`, see `ucp_dt_make_contig`.

6.7.5.28 `ucp_atomic_post()`

```
ucs_status_t ucp_atomic_post (
    ucp_ep_h ep,
    ucp_atomic_post_op_t opcode,
    uint64_t value,
    size_t op_size,
    uint64_t remote_addr,
    ucp_rkey_h rkey )
```

This routine posts an atomic memory operation to a remote value. The remote value is described by the combination of the remote memory address `remote_addr` and the [remote memory handle](#) `rkey`. Return from the function does not guarantee completion. A user must call `ucp_ep_flush_nb` or `ucp_worker_flush_nb` to guarantee that the remote value has been updated.

## Parameters

in	<code>ep</code>	UCP endpoint.
in	<code>opcode</code>	One of <code>ucp_atomic_post_op_t</code> .
in	<code>value</code>	Source operand for the atomic operation.
in	<code>op_size</code>	Size of value in bytes
in	<code>remote_addr</code>	Remote address to operate on.
in	<code>rkey</code>	Remote key handle for the remote memory address.

## Returns

Error code as defined by `ucs_status_t`

6.7.5.29 `ucp_atomic_fetch_nb()`

```
ucs_status_ptr_t ucp_atomic_fetch_nb (
    ucp_ep_h ep,
    ucp_atomic_fetch_op_t opcode,
    uint64_t value,
    void * result,
    size_t op_size,
    uint64_t remote_addr,
    ucp_rkey_h rkey,
    ucp_send_callback_t cb )
```

This routine will post an atomic fetch operation to remote memory. The remote value is described by the combination of the remote memory address `remote_addr` and the [remote memory handle](#) `rkey`. The routine is non-blocking and therefore returns immediately. However the actual atomic operation may be delayed. The atomic operation is not considered complete until the values in remote and local memory are completed. If the atomic operation completes immediately, the routine returns `UCS_OK` and the call-back routine `cb` is **not** invoked. If the operation is **not** completed immediately and no error is reported, then the UCP library will schedule invocation of the call-back routine `cb` upon completion of the atomic operation. In other words, the completion of an atomic operation can be signaled by the return code or execution of the call-back.

**Note**

The user should not modify any part of the *result* after this operation is called, until the operation completes.

**Parameters**

in	<i>ep</i>	UCP endpoint.
in	<i>opcode</i>	One of <a href="#">ucp_atomic_fetch_op_t</a> .
in	<i>value</i>	Source operand for atomic operation. In the case of CSWAP this is the conditional for the swap. For SWAP this is the value to be placed in remote memory.
in, out	<i>result</i>	Local memory address to store resulting fetch to. In the case of CSWAP the value in result will be swapped into the <i>remote_addr</i> if the condition is true.
in	<i>op_size</i>	Size of value in bytes and pointer type for result
in	<i>remote_addr</i>	Remote address to operate on.
in	<i>rkey</i>	Remote key handle for the remote memory address.
in	<i>cb</i>	Call-back function that is invoked whenever the send operation is completed. It is important to note that the call-back function is only invoked in a case when the operation cannot be completed in place.

**Returns**

NULL - The operation was completed immediately.

UCS\_PTR\_IS\_ERR(\_ptr) - The operation failed.

otherwise - Operation was scheduled and can be completed at any point in time. The request handle is returned to the application in order to track progress of the operation. The application is responsible for releasing the handle using [ucp\\_request\\_free\(\)](#) routine.

**6.7.5.30 ucp\_atomic\_op\_nbx()**

```
ucs_status_ptr_t ucp_atomic_op_nbx (
    ucp_ep_h ep,
    ucp_atomic_op_t opcode,
    const void * buffer,
    size_t count,
    uint64_t remote_addr,
    ucp_rkey_h rkey,
    const ucp_request_param_t * param )
```

This routine will post an atomic operation to remote memory. The remote value is described by the combination of the remote memory address *remote\_addr* and the [remote memory handle](#) *rkey*. The routine is non-blocking and therefore returns immediately. However, the actual atomic operation may be delayed. In order to enable fetching semantics for atomic operations user has to specify *param.reply\_buffer*. Please see [6.160](#) below for more details.

**Note**

The user should not modify any part of the *buffer* (or also *param->reply\_buffer* for fetch operations), until the operation completes.

Only [ucp\\_dt\\_make\\_config\(4\)](#) and [ucp\\_dt\\_make\\_contig\(8\)](#) are supported in *param->datatype*, see [ucp\\_dt\\_make\\_contig](#). Also, currently atomic operations can handle one element only. Thus, *count* argument must be set to 1.

Table 6.160: Atomic Operations Semantic

Atomic Operation	Pseudo code	X	Y	Z	Result
<a href="#">UCP_ATOMIC_OP_ADD</a>	Result=Y; Y+=X	buffer	remote_addr	-	param.reply_↔ buffer(optional)
<a href="#">UCP_ATOMIC_OP_SWAP</a>	Result=Y; Y=X	buffer	remote_addr	-	param.reply_buffer
<a href="#">UCP_ATOMIC_OP_CAS</a>	Result=Y; if (X==Y) then Y=Z	buffer	remote_addr	param.reply_buffer	param.reply_buffer
<a href="#">UCP_ATOMIC_OP_AND</a>	Result=Y; Y&=X	buffer	remote_addr	-	param.reply_↔ buffer(optional)
<a href="#">UCP_ATOMIC_OP_OR</a>	Result=Y; Y =X	buffer	remote_addr	-	param.reply_↔ buffer(optional)
<a href="#">UCP_ATOMIC_OP_XOR</a>	Result=Y; Y^=X	buffer	remote_addr	-	param.reply_↔ buffer(optional)

**Parameters**

in	<i>ep</i>	UCP endpoint.
in	<i>opcode</i>	One of <a href="#">ucp_atomic_op_t</a> .
in	<i>buffer</i>	Address of operand for the atomic operation. See <a href="#">6.160</a> for exact usage by different atomic operations.
in	<i>count</i>	Number of elements in <i>buffer</i> and <i>result</i> . The size of each element is specified by <a href="#">ucp_request_param_t::datatype</a>
in	<i>remote_addr</i>	Remote address to operate on.
in	<i>rkey</i>	Remote key handle for the remote memory address.
in	<i>param</i>	Operation parameters, see <a href="#">ucp_request_param_t</a> .

**Returns**

NULL - The operation completed immediately.

UCS\_PTR\_IS\_ERR(ptr) - The operation failed.

otherwise - Operation was scheduled and can be completed at some time in the future. The request handle is returned to the application in order to track progress of the operation.

**6.7.5.31 ucp\_request\_check\_status()**

```
ucs_status_t ucp_request_check_status (
    void * request )
```

This routine checks the state of the request and returns its current status. Any value different from UCS\_INPROGRESS means that request is in a completed state.

**Parameters**

in	<i>request</i>	Non-blocking request to check.
----	----------------	--------------------------------

**Returns**

Error code as defined by [ucs\\_status\\_t](#)

**Examples**

[ucp\\_client\\_server.c](#), and [ucp\\_hello\\_world.c](#).

**6.7.5.32 ucp\_tag\_rcv\_request\_test()**

```
ucs_status_t ucp_tag_rcv_request_test (
    void * request,
    ucp_tag_rcv_info_t * info )
```

This routine checks the state and returns current status of the request returned from [ucp\\_tag\\_rcv\\_nb](#) routine or the user allocated request for [ucp\\_tag\\_rcv\\_nbr](#). Any value different from UCS\_INPROGRESS means that the request is in a completed state.

**Parameters**

in	<i>request</i>	Non-blocking request to check.
out	<i>info</i>	It is filled with the details about the message available at the moment of calling.

**Returns**

Error code as defined by [ucs\\_status\\_t](#)

**6.7.5.33 ucp\_stream\_rcv\_request\_test()**

```
ucs_status_t ucp_stream_rcv_request_test (
    void * request,
    size_t * length_p )
```

This routine checks the state and returns current status of the request returned from [ucp\\_stream\\_rcv\\_nb](#) routine. Any value different from UCS\_INPROGRESS means that the request is in a completed state.

**Parameters**

in	<i>request</i>	Non-blocking request to check.
out	<i>length↔ _p</i>	The size of the received data in bytes. This value is only valid if the status is UCS_OK. If valid, it is always an integral multiple of the datatype size associated with the request.

**Returns**

Error code as defined by [ucs\\_status\\_t](#)

**6.7.5.34 ucp\_request\_cancel()**

```
void ucp_request_cancel (
```

```
ucp_worker_h worker,
void * request )
```

**Parameters**

in	<i>worker</i>	UCP worker.
in	<i>request</i>	Non-blocking request to cancel.

This routine tries to cancel an outstanding communication request. After calling this routine, the *request* will be in completed or canceled (but not both) state regardless of the status of the target endpoint associated with the communication request. If the request is completed successfully, the [send](#) or [receive](#) completion callbacks (based on the type of the request) will be called with the *status* argument of the callback set to UCS\_OK, and in a case it is canceled the *status* argument is set to UCS\_ERR\_CANCELED. It is important to note that in order to release the request back to the library the application is responsible for calling [ucp\\_request\\_free\(\)](#).

**6.7.5.35 ucp\_stream\_data\_release()**

```
void ucp_stream_data_release (
    ucp_ep_h ep,
    void * data )
```

**Parameters**

in	<i>ep</i>	Endpoint <i>data</i> received from.
in	<i>data</i>	Data pointer to release, which was returned from <a href="#">ucp_stream_rcv_data_nb</a> .

This routine releases internal UCP data buffer returned by [ucp\\_stream\\_rcv\\_data\\_nb](#) when *data* is processed, the application can't use this buffer after calling this function.

**6.7.5.36 ucp\_request\_free()**

```
void ucp_request_free (
    void * request )
```

**Parameters**

in	<i>request</i>	Non-blocking request to release.
----	----------------	----------------------------------

This routine releases the non-blocking request back to the library, regardless of its current state. Communications operations associated with this request will make progress internally, however no further notifications or callbacks will be invoked for this request.

**Examples**

[ucp\\_client\\_server.c](#), and [ucp\\_hello\\_world.c](#).

**6.7.5.37 ucp\_request\_alloc()**

```
void* ucp_request_alloc (
    ucp_worker_h worker )
```

## Parameters

in	<i>worker</i>	UCP worker.
----	---------------	-------------

## Returns

Error code as defined by [ucs\\_status\\_t](#)

This routine creates request which may be used in functions [ucp\\_tag\\_send\\_nbx](#), [ucp\\_tag\\_rcv\\_nbx](#), etc. The application is responsible for releasing the handle using the [ucp\\_request\\_free](#) routine

6.7.5.38 [ucp\\_request\\_is\\_completed\(\)](#)

```
int ucp_request_is_completed (
    void * request )
```

**Deprecated** Replaced by [ucp\\_request\\_test](#).

6.7.5.39 [ucp\\_put\(\)](#)

```
ucs_status_t ucp_put (
    ucp_ep_h ep,
    const void * buffer,
    size_t length,
    uint64_t remote_addr,
    ucp_rkey_h rkey )
```

**Deprecated** Replaced by [ucp\\_put\\_nb](#). The following example implements the same functionality using [ucp\\_put\\_nb](#) :

```
void empty_callback(void *request, ucs_status_t status)
{
}
ucs_status_t put(ucp_ep_h ep, const void *buffer, size_t length,
                uint64_t remote_addr, ucp_rkey_h rkey)
{
    void *request = ucp_put_nb(ep, buffer, length, remote_addr, rkey,
                              empty_callback),
    if (request == NULL) {
        return UCS_OK;
    } else if (UCS_PTR_IS_ERR(request)) {
        return UCS_PTR_STATUS(request);
    } else {
        ucs_status_t status;
        do {
            ucp_worker_progress(worker);
            status = ucp_request_check_status(request);
        } while (status == UCS_INPROGRESS);
        ucp_request_release(request);
        return status;
    }
}
```

This routine stores contiguous block of data that is described by the local address *buffer* in the remote contiguous memory region described by *remote\_addr* address and the [memory handle](#) *rkey*. The routine returns when it is safe to reuse the source address *buffer*.

## Parameters

in	<i>ep</i>	Remote endpoint handle.
in	<i>buffer</i>	Pointer to the local source address.
in	<i>length</i>	Length of the data (in bytes) stored under the source address.
in	<i>remote_addr</i>	Pointer to the destination remote address to write to.
in	<i>rkey</i>	Remote memory key associated with the remote address.



**Returns**

Error code as defined by [ucs\\_status\\_t](#)

**6.7.5.40 ucp\_get()**

```
ucs_status_t ucp_get (
    ucp_ep_h ep,
    void * buffer,
    size_t length,
    uint64_t remote_addr,
    ucp_rkey_h rkey )
```

**Deprecated** Replaced by [ucp\\_get\\_nb](#).

**See also**

[ucp\\_put](#).

This routine loads contiguous block of data that is described by the remote address *remote\_addr* and the [memory handle](#) *rkey* in the local contiguous memory region described by *buffer* address. The routine returns when remote data is loaded and stored under the local address *buffer*.

**Parameters**

in	<i>ep</i>	Remote endpoint handle.
in	<i>buffer</i>	Pointer to the local source address.
in	<i>length</i>	Length of the data (in bytes) stored under the source address.
in	<i>remote_addr</i>	Pointer to the destination remote address to write to.
in	<i>rkey</i>	Remote memory key associated with the remote address.

**Returns**

Error code as defined by [ucs\\_status\\_t](#)

**6.7.5.41 ucp\_atomic\_add32()**

```
ucs_status_t ucp_atomic_add32 (
    ucp_ep_h ep,
    uint32_t add,
    uint64_t remote_addr,
    ucp_rkey_h rkey )
```

**Deprecated** Replaced by [ucp\\_atomic\\_post](#) with opcode UCP\_ATOMIC\_POST\_OP\_ADD.

**See also**

[ucp\\_put](#).

This routine performs an add operation on a 32 bit integer value atomically. The remote integer value is described by the combination of the remote memory address *remote\_addr* and the [remote memory handle](#) *rkey*. The *add* value is the value that is used for the add operation. When the operation completes the sum of the original remote value and the operand value (*add*) is stored in remote memory. The call to the routine returns immediately, independent of operation completion.

**Note**

The remote address must be aligned to 32 bit.

**Parameters**

in	<i>ep</i>	Remote endpoint handle.
in	<i>add</i>	Value to add.
in	<i>remote_addr</i>	Pointer to the destination remote address of the atomic variable.
in	<i>rkey</i>	Remote memory key associated with the remote address.

**Returns**

Error code as defined by [ucs\\_status\\_t](#)

**6.7.5.42 ucp\_atomic\_add64()**

```
ucs_status_t ucp_atomic_add64 (
    ucp_ep_h ep,
    uint64_t add,
    uint64_t remote_addr,
    ucp_rkey_h rkey )
```

**Deprecated** Replaced by [ucp\\_atomic\\_post](#) with opcode UCP\_ATOMIC\_POST\_OP\_ADD.

**See also**

[ucp\\_put](#).

This routine performs an add operation on a 64 bit integer value atomically. The remote integer value is described by the combination of the remote memory address *remote\_addr* and the [remote memory handle](#) *rkey*. The *add* value is the value that is used for the add operation. When the operation completes the sum of the original remote value and the operand value (*add*) is stored in remote memory. The call to the routine returns immediately, independent of operation completion.

**Note**

The remote address must be aligned to 64 bit.

**Parameters**

in	<i>ep</i>	Remote endpoint handle.
in	<i>add</i>	Value to add.
in	<i>remote_addr</i>	Pointer to the destination remote address of the atomic variable.
in	<i>rkey</i>	Remote memory key associated with the remote address.

**Returns**

Error code as defined by [ucs\\_status\\_t](#)

**6.7.5.43 ucp\_atomic\_fadd32()**

```
ucs_status_t ucp_atomic_fadd32 (
    ucp_ep_h ep,
    uint32_t add,
    uint64_t remote_addr,
    ucp_rkey_h rkey,
    uint32_t * result )
```

**Deprecated** Replaced by [ucp\\_atomic\\_fetch\\_nb](#) with opcode UCP\_ATOMIC\_FETCH\_OP\_FADD.

**See also**

[ucp\\_put](#).

This routine performs an add operation on a 32 bit integer value atomically. The remote integer value is described by the combination of the remote memory address *remote\_addr* and the [remote memory handle](#) *rkey*. The *add* value is the value that is used for the add operation. When the operation completes, the original remote value is stored in the local memory *result*, and the sum of the original remote value and the operand value is stored in remote memory. The call to the routine returns when the operation is completed and the *result* value is updated.

**Note**

The remote address must be aligned to 32 bit.

**Parameters**

in	<i>ep</i>	Remote endpoint handle.
in	<i>add</i>	Value to add.
in	<i>remote_addr</i>	Pointer to the destination remote address of the atomic variable.
in	<i>rkey</i>	Remote memory key associated with the remote address.
out	<i>result</i>	Pointer to the address that is used to store the previous value of the atomic variable described by the <i>remote_addr</i>

**Returns**

Error code as defined by [ucs\\_status\\_t](#)

**6.7.5.44 ucp\_atomic\_fadd64()**

```
ucs_status_t ucp_atomic_fadd64 (
    ucp_ep_h ep,
    uint64_t add,
    uint64_t remote_addr,
    ucp_rkey_h rkey,
    uint64_t * result )
```

**Deprecated** Replaced by [ucp\\_atomic\\_fetch\\_nb](#) with opcode UCP\_ATOMIC\_FETCH\_OP\_FADD.

See also

[ucp\\_put](#).

This routine performs an add operation on a 64 bit integer value atomically. The remote integer value is described by the combination of the remote memory address *remote\_addr* and the [remote memory handle](#) *rkey*. The *add* value is the value that is used for the add operation. When the operation completes, the original remote value is stored in the local memory *result*, and the sum of the original remote value and the operand value is stored in remote memory. The call to the routine returns when the operation is completed and the *result* value is updated.

**Note**

The remote address must be aligned to 64 bit.

**Parameters**

in	<i>ep</i>	Remote endpoint handle.
in	<i>add</i>	Value to add.
in	<i>remote_addr</i>	Pointer to the destination remote address of the atomic variable.
in	<i>rkey</i>	Remote memory key associated with the remote address.
out	<i>result</i>	Pointer to the address that is used to store the previous value of the atomic variable described by the <i>remote_addr</i>

**Returns**

Error code as defined by [ucs\\_status\\_t](#)

#### 6.7.5.45 ucp\_atomic\_swap32()

```
ucs_status_t ucp_atomic_swap32 (
    ucp_ep_h ep,
    uint32_t swap,
    uint64_t remote_addr,
    ucp_rkey_h rkey,
    uint32_t * result )
```

**Deprecated** Replaced by [ucp\\_atomic\\_fetch\\_nb](#) with opcode UCP\_ATOMIC\_FETCH\_OP\_SWAP.

See also

[ucp\\_put](#).

This routine swaps a 32 bit value between local and remote memory. The remote value is described by the combination of the remote memory address *remote\_addr* and the [remote memory handle](#) *rkey*. The *swap* value is the value that is used for the swap operation. When the operation completes, the remote value is stored in the local memory *result*, and the operand value (*swap*) is stored in remote memory. The call to the routine returns when the operation is completed and the *result* value is updated.

**Note**

The remote address must be aligned to 32 bit.

## Parameters

in	<i>ep</i>	Remote endpoint handle.
in	<i>swap</i>	Value to swap.
in	<i>remote_addr</i>	Pointer to the destination remote address of the atomic variable.
in	<i>rkey</i>	Remote memory key associated with the remote address.
out	<i>result</i>	Pointer to the address that is used to store the previous value of the atomic variable described by the <i>remote_addr</i>

## Returns

Error code as defined by [ucs\\_status\\_t](#)

## 6.7.5.46 ucp\_atomic\_swap64()

```
ucs_status_t ucp_atomic_swap64 (
    ucp_ep_h ep,
    uint64_t swap,
    uint64_t remote_addr,
    ucp_rkey_h rkey,
    uint64_t * result )
```

**Deprecated** Replaced by [ucp\\_atomic\\_fetch\\_nb](#) with opcode UCP\_ATOMIC\_FETCH\_OP\_SWAP.

## See also

[ucp\\_put](#).

This routine swaps a 64 bit value between local and remote memory. The remote value is described by the combination of the remote memory address *remote\_addr* and the [remote memory handle](#) *rkey*. The *swap* value is the value that is used for the swap operation. When the operation completes, the remote value is stored in the local memory *result*, and the operand value (*swap*) is stored in remote memory. The call to the routine returns when the operation is completed and the *result* value is updated.

## Note

The remote address must be aligned to 64 bit.

## Parameters

in	<i>ep</i>	Remote endpoint handle.
in	<i>swap</i>	Value to swap.
in	<i>remote_addr</i>	Pointer to the destination remote address of the atomic variable.
in	<i>rkey</i>	Remote memory key associated with the remote address.
out	<i>result</i>	Pointer to the address that is used to store the previous value of the atomic variable described by the <i>remote_addr</i>

**Returns**

Error code as defined by [ucs\\_status\\_t](#)

**6.7.5.47 ucp\_atomic\_cswap32()**

```
ucs_status_t ucp_atomic_cswap32 (
    ucp_ep_h ep,
    uint32_t compare,
    uint32_t swap,
    uint64_t remote_addr,
    ucp_rkey_h rkey,
    uint32_t * result )
```

**Deprecated** Replaced by [ucp\\_atomic\\_fetch\\_nb](#) with opcode UCP\_ATOMIC\_FETCH\_OP\_CSWARE.

**See also**

[ucp\\_put](#).

This routine conditionally swaps a 32 bit value between local and remote memory. The swap occurs only if the condition value (*continue*) is equal to the remote value, otherwise the remote memory is not modified. The remote value is described by the combination of the remote memory address `remote_addr` and the [remote memory handle](#) `rkey`. The `swap` value is the value that is used to update the remote memory if the condition is true. The call to the routine returns when the operation is completed and the `result` value is updated.

**Note**

The remote address must be aligned to 32 bit.

**Parameters**

in	<i>ep</i>	Remote endpoint handle.
in	<i>compare</i>	Value to compare to.
in	<i>swap</i>	Value to swap.
in	<i>remote_addr</i>	Pointer to the destination remote address of the atomic variable.
in	<i>rkey</i>	Remote memory key associated with the remote address.
out	<i>result</i>	Pointer to the address that is used to store the previous value of the atomic variable described by the <i>remote_addr</i>

**Returns**

Error code as defined by [ucs\\_status\\_t](#)

**6.7.5.48 ucp\_atomic\_cswap64()**

```
ucs_status_t ucp_atomic_cswap64 (
    ucp_ep_h ep,
    uint64_t compare,
    uint64_t swap,
```

```
uint64_t remote_addr,
ucp_rkey_h rkey,
uint64_t * result )
```

**Deprecated** Replaced by [ucp\\_atomic\\_fetch\\_nb](#) with opcode UCP\_ATOMIC\_FETCH\_OP\_CSWAP.

#### See also

[ucp\\_put](#).

This routine conditionally swaps a 64 bit value between local and remote memory. The swap occurs only if the condition value (*continue*) is equal to the remote value, otherwise the remote memory is not modified. The remote value is described by the combination of the remote memory address *remote\_addr* and the [remote memory handle](#) *rkey*. The *swap* value is the value that is used to update the remote memory if the condition is true. The call to the routine returns when the operation is completed and the *result* value is updated.

#### Note

The remote address must be aligned to 64 bit.

#### Parameters

in	<i>ep</i>	Remote endpoint handle.
in	<i>compare</i>	Value to compare to.
in	<i>swap</i>	Value to swap.
in	<i>remote_addr</i>	Pointer to the destination remote address of the atomic variable.
in	<i>rkey</i>	Remote memory key associated with the remote address.
out	<i>result</i>	Pointer to the address that is used to store the previous value of the atomic variable described by the <i>remote_addr</i>

#### Returns

Error code as defined by [ucs\\_status\\_t](#)

## 6.8 UCP Configuration

### Data Structures

- struct [ucp\\_params](#)  
*Tuning parameters for UCP library. [More...](#)*

### Typedefs

- typedef struct [ucp\\_params](#) [ucp\\_params\\_t](#)  
*Tuning parameters for UCP library.*
- typedef struct [ucp\\_config](#) [ucp\\_config\\_t](#)  
*UCP configuration descriptor.*

### Functions

- [ucs\\_status\\_t ucp\\_config\\_read](#) (const char \*env\_prefix, const char \*filename, [ucp\\_config\\_t](#) \*\*config\_p)  
*Read UCP configuration descriptor.*
- void [ucp\\_config\\_release](#) ([ucp\\_config\\_t](#) \*config)  
*Release configuration descriptor.*
- [ucs\\_status\\_t ucp\\_config\\_modify](#) ([ucp\\_config\\_t](#) \*config, const char \*name, const char \*value)  
*Modify context configuration.*
- void [ucp\\_config\\_print](#) (const [ucp\\_config\\_t](#) \*config, FILE \*stream, const char \*title, [ucs\\_config\\_print\\_flags\\_t](#) print\_flags)  
*Print configuration information.*

#### 6.8.1 Detailed Description

This section describes routines for configuration of the UCP network layer

#### 6.8.2 Data Structure Documentation

##### 6.8.2.1 struct ucp\_params

The structure defines the parameters that are used for UCP library tuning during UCP library [initialization](#).

##### Note

UCP library implementation uses the [features](#) parameter to optimize the library functionality that minimize memory footprint. For example, if the application does not require send/receive semantics UCP library may avoid allocation of expensive resources associated with send/receive queues.

##### Examples

[ucp\\_client\\_server.c](#), and [ucp\\_hello\\_world.c](#).

##### Data Fields

uint64_t	field_mask	Mask of valid fields in this structure, using bits from <a href="#">ucp_params_field</a> . Fields not specified in this mask will be ignored. Provides ABI compatibility with respect to adding new fields.
----------	------------	---



## Data Fields

uint64_t	features	UCP <a href="#">features</a> that are used for library initialization. It is recommended for applications only to request the features that are required for an optimal functionality. This field must be specified.
size_t	request_size	The size of a reserved space in a non-blocking requests. Typically applications use this space for caching own structures in order to avoid costly memory allocations, pointer dereferences, and cache misses. For example, MPI implementation can use this memory for caching MPI descriptors. This field defaults to 0 if not specified.
<a href="#">ucp_request_init_callback_t</a>	request_init	Pointer to a routine that is used for the request initialization. This function will be called only on the very first time a request memory is initialized, and may not be called again if a request is reused. If a request should be reset before the next reuse, it can be done before calling <a href="#">ucp_request_free</a> . <i>NULL</i> can be used if no such is function required, which is also the default if this field is not specified by <a href="#">field_mask</a> .
<a href="#">ucp_request_cleanup_callback_t</a>	request_cleanup	Pointer to a routine that is responsible for final cleanup of the memory associated with the request. This routine may not be called every time a request is released. For some implementations, the cleanup call may be delayed and only invoked at <a href="#">ucp_worker_destroy</a> . <i>NULL</i> can be used if no such function is required, which is also the default if this field is not specified by <a href="#">field_mask</a> .
uint64_t	tag_sender_mask	Mask which specifies particular bits of the tag which can uniquely identify the sender (UCP endpoint) in tagged operations. This field defaults to 0 if not specified.
int	mt_workers_shared	This flag indicates if this context is shared by multiple workers from different threads. If so, this context needs thread safety support; otherwise, the context does not need to provide thread safety. For example, if the context is used by single worker, and that worker is shared by multiple threads, this context does not need thread safety; if the context is used by worker 1 and worker 2, and worker 1 is used by thread 1 and worker 2 is used by thread 2, then this context needs thread safety. Note that actual thread mode may be different from mode passed to <a href="#">ucp_init</a> . To get actual thread mode use <a href="#">ucp_context_query</a> .

## Data Fields

size_t	estimated_num_eps	An optimization hint of how many endpoints will be created on this context. For example, when used from MPI or SHMEM libraries, this number will specify the number of ranks (or processing elements) in the job. Does not affect semantics, but only transport selection criteria and the resulting performance. The value can be also set by UCX_NUM_EPS environment variable. In such case it will override the number of endpoints set by <i>estimated_num_eps</i>
size_t	estimated_num_ppn	An optimization hint for a single node. For example, when used from MPI or OpenSHMEM libraries, this number will specify the number of Processes Per Node (PPN) in the job. Does not affect semantics, only transport selection criteria and the resulting performance. The value can be also set by the UCX_NUM_PPN environment variable, which will override the number of endpoints set by <i>estimated_num_ppn</i>
const char *	name	Tracing and analysis tools can identify the context using this name. To retrieve the context's name, use <a href="#">ucp_context_query</a> , as the name you supply may be changed by UCX under some circumstances, e.g. a name conflict. This field is only assigned if you set <a href="#">UCP_PARAM_FIELD_NAME</a> in the field mask. If not, then a default unique name will be created for you.

## 6.8.3 Typedef Documentation

### 6.8.3.1 ucp\_params\_t

```
typedef struct ucp_params ucp_params_t
```

The structure defines the parameters that are used for UCP library tuning during UCP library [initialization](#).

#### Note

UCP library implementation uses the [features](#) parameter to optimize the library functionality that minimize memory footprint. For example, if the application does not require send/receive semantics UCP library may avoid allocation of expensive resources associated with send/receive queues.

### 6.8.3.2 ucp\_config\_t

```
typedef struct ucp_config ucp_config_t
```

This descriptor defines the configuration for [UCP application context](#). The configuration is loaded from the runtime environment (using configuration files of environment variables) using [ucp\\_config\\_read](#) routine and can be

printed using `ucp_config_print` routine. In addition, application is responsible to release the descriptor using `ucp_config_release` routine.

## 6.8.4 Function Documentation

### 6.8.4.1 `ucp_config_read()`

```
ucs_status_t ucp_config_read (
    const char * env_prefix,
    const char * filename,
    ucp_config_t ** config_p )
```

The routine fetches the information about UCP library configuration from the run-time environment. Then, the fetched descriptor is used for UCP library [initialization](#). The Application can print out the descriptor using [print](#) routine. In addition the application is responsible for [releasing](#) the descriptor back to the UCP library.

#### Parameters

in	<i>env_prefix</i>	If non-NULL, the routine searches for the environment variables that start with <code>&lt;env_prefix&gt;_UCX_</code> prefix. Otherwise, the routine searches for the environment variables that start with <code>UCX_</code> prefix.
in	<i>filename</i>	If non-NULL, read configuration from the file defined by <i>filename</i> . If the file does not exist, it will be ignored and no error reported to the application.
out	<i>config_p</i>	Pointer to configuration descriptor as defined by <a href="#">ucp_config_t</a> .

#### Returns

Error code as defined by [ucs\\_status\\_t](#)

#### Examples

[ucp\\_hello\\_world.c](#).

### 6.8.4.2 `ucp_config_release()`

```
void ucp_config_release (
    ucp_config_t * config )
```

The routine releases the configuration descriptor that was allocated through [ucp\\_config\\_read\(\)](#) routine.

#### Parameters

out	<i>config</i>	Configuration descriptor as defined by <a href="#">ucp_config_t</a> .
-----	---------------	---

#### Examples

[ucp\\_hello\\_world.c](#).

### 6.8.4.3 ucp\_config\_modify()

```
ucs_status_t ucp_config_modify (
    ucp_config_t * config,
    const char * name,
    const char * value )
```

The routine changes one configuration setting stored in [configuration](#) descriptor.

#### Parameters

in	<i>config</i>	Configuration to modify.
in	<i>name</i>	Configuration variable name.
in	<i>value</i>	Value to set.

#### Returns

Error code.

### 6.8.4.4 ucp\_config\_print()

```
void ucp_config_print (
    const ucp_config_t * config,
    FILE * stream,
    const char * title,
    ucs_config_print_flags_t print_flags )
```

The routine prints the configuration information that is stored in [configuration](#) descriptor.

#### Parameters

in	<i>config</i>	<a href="#">Configuration descriptor</a> to print.
in	<i>stream</i>	Output stream to print the configuration to.
in	<i>title</i>	Configuration title to print.
in	<i>print_flags</i>	Flags that control various printing options.

#### Examples

[ucp\\_hello\\_world.c](#).

## 6.9 UCP Data type routines

### Data Structures

- struct `ucp_dt_iov`  
*Structure for scatter-gather I/O. [More...](#)*
- struct `ucp_generic_dt_ops`  
*UCP generic data type descriptor.*

### Macros

- `#define ucp_dt_make_contig(_elem_size) (((ucp_datatype_t)(_elem_size) << UCP_DATATYPE_SHIFT) | UCP_DATATYPE_CONTIG)`  
*Generate an identifier for contiguous data type.*
- `#define ucp_dt_make_iov() ((ucp_datatype_t)UCP_DATATYPE_IOV)`  
*Generate an identifier for Scatter-gather IOV data type.*

### Typedefs

- typedef struct `ucp_dt_iov` `ucp_dt_iov_t`  
*Structure for scatter-gather I/O.*
- typedef struct `ucp_generic_dt_ops` `ucp_generic_dt_ops_t`  
*UCP generic data type descriptor.*

### Enumerations

- enum `ucp_dt_type` {  
`UCP_DATATYPE_CONTIG = 0, UCP_DATATYPE_STRIDED = 1, UCP_DATATYPE_IOV = 2,`  
`UCP_DATATYPE_GENERIC = 7,`  
`UCP_DATATYPE_SHIFT = 3, UCP_DATATYPE_CLASS_MASK = UCS_MASK(UCP_DATATYPE_SHIFT)`  
 }  
*UCP data type classification.*

### Functions

- `ucs_status_t ucp_dt_create_generic` (const `ucp_generic_dt_ops_t` \*ops, void \*context, `ucp_datatype_t` \*datatype\_p)  
*Create a generic datatype.*
- void `ucp_dt_destroy` (`ucp_datatype_t` datatype)  
*Destroy a datatype and release its resources.*

### Variables

- void \*(\* `ucp_generic_dt_ops::start_pack` )(void \*context, const void \*buffer, size\_t count)  
*Start a packing request.*
- void \*(\* `ucp_generic_dt_ops::start_unpack` )(void \*context, void \*buffer, size\_t count)  
*Start an unpacking request.*
- size\_t(\* `ucp_generic_dt_ops::packed_size` )(void \*state)  
*Get the total size of packed data.*
- size\_t(\* `ucp_generic_dt_ops::pack` )(void \*state, size\_t offset, void \*dest, size\_t max\_length)

*Pack data.*

- `ucs_status_t(* ucp_generic_dt_ops::unpack)(void *state, size_t offset, const void *src, size_t length)`

*Unpack data.*

- `void(* ucp_generic_dt_ops::finish)(void *state)`

*Finish packing/unpacking.*

## 6.9.1 Detailed Description

UCP Data type routines

## 6.9.2 Data Structure Documentation

### 6.9.2.1 struct ucp\_dt\_iov

This structure is used to specify a list of buffers which can be used within a single data transfer function call. This list should remain valid until the data transfer request is completed.

#### Note

If *length* is zero, the memory pointed to by *buffer* will not be accessed. Otherwise, *buffer* must point to valid memory.

#### Examples

[ucp\\_client\\_server.c](#).

#### Data Fields

<code>void *</code>	<code>buffer</code>	Pointer to a data buffer
<code>size_t</code>	<code>length</code>	Length of the <i>buffer</i> in bytes

## 6.9.3 Macro Definition Documentation

### 6.9.3.1 ucp\_dt\_make\_contig

```
#define ucp_dt_make_contig(
    _elem_size ) (((ucp_datatype_t)(_elem_size) << UCP_DATATYPE_SHIFT) | UCP_DATATYPE_CONTIG)
```

This macro creates an identifier for contiguous datatype that is defined by the size of the basic element.

#### Parameters

<code>in</code>	<code>_elem_size</code>	Size of the basic element of the type.
-----------------	-------------------------	--

#### Returns

Data-type identifier.

**Note**

In case of partial receive, the buffer will be filled with integral count of elements.

**Examples**

[ucp\\_client\\_server.c](#), and [ucp\\_hello\\_world.c](#).

**6.9.3.2 ucp\_dt\_make\_iov**

```
#define ucp_dt_make_iov( ) ((ucp_datatype_t)UCP_DATATYPE_IOV)
```

This macro creates an identifier for datatype of scatter-gather list with multiple pointers

**Returns**

Data-type identifier.

**Note**

In the event of partial receive, [ucp\\_dt\\_iov\\_t::buffer](#) can be filled with any number of bytes according to its [ucp\\_dt\\_iov\\_t::length](#).

**6.9.4 Typedef Documentation****6.9.4.1 ucp\_dt\_iov\_t**

```
typedef struct ucp_dt_iov ucp_dt_iov_t
```

This structure is used to specify a list of buffers which can be used within a single data transfer function call. This list should remain valid until the data transfer request is completed.

**Note**

If *length* is zero, the memory pointed to by *buffer* will not be accessed. Otherwise, *buffer* must point to valid memory.

**6.9.4.2 ucp\_generic\_dt\_ops\_t**

```
typedef struct ucp_generic_dt_ops ucp_generic_dt_ops_t
```

This structure provides a generic datatype descriptor that is used for definition of application defined datatypes.

Typically, the descriptor is used for an integration with datatype engines implemented within MPI and SHMEM implementations.

**Note**

In case of partial receive, any amount of received data is acceptable which matches buffer size.

**6.9.5 Enumeration Type Documentation**

### 6.9.5.1 ucp\_dt\_type

enum [ucp\\_dt\\_type](#)

The enumeration list describes the datatypes supported by UCP.

#### Enumerator

UCP_DATATYPE_CONTIG	Contiguous datatype
UCP_DATATYPE_STRIDED	Strided datatype
UCP_DATATYPE_IOV	Scatter-gather list with multiple pointers
UCP_DATATYPE_GENERIC	Generic datatype with user-defined pack/unpack routines
UCP_DATATYPE_SHIFT	Number of bits defining the datatype classification
UCP_DATATYPE_CLASS_MASK	Data-type class mask

## 6.9.6 Function Documentation

### 6.9.6.1 ucp\_dt\_create\_generic()

```
ucs_status_t ucp_dt_create_generic (
    const ucp_generic_dt_ops_t * ops,
    void * context,
    ucp_datatype_t * datatype_p )
```

This routine create a generic datatype object. The generic datatype is described by the *ops* object which provides a table of routines defining the operations for generic datatype manipulation. Typically, generic datatypes are used for integration with datatype engines provided with MPI implementations (MPICH, Open MPI, etc). The application is responsible for releasing the *datatype\_p* object using [ucp\\_dt\\_destroy\(\)](#) routine.

#### Parameters

in	<i>ops</i>	Generic datatype function table as defined by <a href="#">ucp_generic_dt_ops_t</a> .
in	<i>context</i>	Application defined context passed to this routine. The context is passed as a parameter to the routines in the <i>ops</i> table.
out	<i>datatype</i> <sub>↔</sub> <i>_p</i>	A pointer to datatype object.

#### Returns

Error code as defined by [ucs\\_status\\_t](#)

### 6.9.6.2 ucp\_dt\_destroy()

```
void ucp_dt_destroy (
    ucp_datatype_t datatype )
```

This routine destroys the *datatype* object and releases any resources that are associated with the object. The *datatype* object must be allocated using [ucp\\_dt\\_create\\_generic\(\)](#) routine.



**Warning**

- Once the *datatype* object is released an access to this object may cause an undefined failure.

**Parameters**

in	<i>datatype</i>	Datatype object to destroy.
----	-----------------	-----------------------------

**6.9.7 Variable Documentation****6.9.7.1 start\_pack**

```
void>(* ucp_generic_dt_ops::start_pack) (void *context, const void *buffer, size_t count)
```

The pointer refers to application defined start-to-pack routine. It will be called from the [ucp\\_tag\\_send\\_nb](#) routine.

**Parameters**

in	<i>context</i>	User-defined context.
in	<i>buffer</i>	Buffer to pack.
in	<i>count</i>	Number of elements to pack into the buffer.

**Returns**

A custom state that is passed to the following [pack\(\)](#) routine.

**6.9.7.2 start\_unpack**

```
void>(* ucp_generic_dt_ops::start_unpack) (void *context, void *buffer, size_t count)
```

The pointer refers to application defined start-to-unpack routine. It will be called from the [ucp\\_tag\\_recv\\_nb](#) routine.

**Parameters**

in	<i>context</i>	User-defined context.
in	<i>buffer</i>	Buffer to unpack to.
in	<i>count</i>	Number of elements to unpack in the buffer.

**Returns**

A custom state that is passed later to the following [unpack\(\)](#) routine.

**6.9.7.3 packed\_size**

```
size_t(* ucp_generic_dt_ops::packed_size) (void *state)
```

The pointer refers to user defined routine that returns the size of data in a packed format.

**Parameters**

in	<i>state</i>	State as returned by <a href="#">start_pack()</a> routine.
----	--------------	--

**Returns**

The size of the data in a packed form.

**6.9.7.4 pack**

```
size_t(* ucp_generic_dt_ops::pack) (void *state, size_t offset, void *dest, size_t max_length)
```

The pointer refers to application defined pack routine.

**Parameters**

in	<i>state</i>	State as returned by <a href="#">start_pack()</a> routine.
in	<i>offset</i>	Virtual offset in the output stream.
in	<i>dest</i>	Destination buffer to pack the data.
in	<i>max_length</i>	Maximum length to pack.

**Returns**

The size of the data that was written to the destination buffer. Must be less than or equal to *max\_length*.

**6.9.7.5 unpack**

```
ucs_status_t(* ucp_generic_dt_ops::unpack) (void *state, size_t offset, const void *src, size_t length)
```

The pointer refers to application defined unpack routine.

**Parameters**

in	<i>state</i>	State as returned by <a href="#">start_unpack()</a> routine.
in	<i>offset</i>	Virtual offset in the input stream.
in	<i>src</i>	Source to unpack the data from.
in	<i>length</i>	Length to unpack.

**Returns**

UCS\_OK or an error if unpacking failed.

**6.9.7.6 finish**

```
void(* ucp_generic_dt_ops::finish) (void *state)
```

The pointer refers to application defined finish routine.

## Parameters

in	<i>state</i>	State as returned by <a href="#">start_pack()</a> and <a href="#">start_unpack()</a> routines.
----	--------------	--

## 6.10 Unified Communication Transport (UCT) API

This section describes UCT API.

### Modules

- [UCT Communication Resource](#)
- [UCT Communication Context](#)
- [UCT Memory Domain](#)
- [UCT Active messages](#)
- [UCT Remote memory access operations](#)
- [UCT Atomic operations](#)
- [UCT Tag matching operations](#)
- [UCT client-server operations](#)

### 6.10.1 Detailed Description

## 6.11 UCT Communication Resource

### Modules

- [UCT interface operations and capabilities](#)  
*List of capabilities supported by UCX API.*
- [UCT interface for asynchronous event capabilities](#)  
*List of capabilities supported by UCT iface event API.*

### Data Structures

- struct [uct\\_md\\_resource\\_desc](#)  
*Memory domain resource descriptor. [More...](#)*
- struct [uct\\_component\\_attr](#)  
*UCT component attributes. [More...](#)*
- struct [uct\\_tl\\_resource\\_desc](#)  
*Communication resource descriptor. [More...](#)*
- struct [uct\\_iface\\_attr](#)  
*Interface attributes: capabilities and limitations. [More...](#)*
- struct [uct\\_iface\\_attr.cap](#)
- struct [uct\\_iface\\_attr.cap.put](#)
- struct [uct\\_iface\\_attr.cap.get](#)
- struct [uct\\_iface\\_attr.cap.am](#)
- struct [uct\\_iface\\_attr.cap.tag](#)
- struct [uct\\_iface\\_attr.cap.tag.recv](#)
- struct [uct\\_iface\\_attr.cap.tag.eager](#)
- struct [uct\\_iface\\_attr.cap.tag.rndv](#)
- struct [uct\\_iface\\_attr.cap.atomic32](#)
- struct [uct\\_iface\\_attr.cap.atomic64](#)
- struct [uct\\_iface\\_params](#)  
*Parameters used for interface creation. [More...](#)*
- union [uct\\_iface\\_params.mode](#)
- struct [uct\\_iface\\_params.mode.device](#)
- struct [uct\\_iface\\_params.mode.sockaddr](#)
- struct [uct\\_ep\\_params](#)  
*Parameters for creating a UCT endpoint by [uct\\_ep\\_create](#). [More...](#)*
- struct [uct\\_completion](#)  
*Completion handle. [More...](#)*
- struct [uct\\_pending\\_req](#)  
*Pending request. [More...](#)*
- struct [uct\\_iov](#)  
*Structure for scatter-gather I/O. [More...](#)*

### Typedefs

- typedef struct [uct\\_md\\_resource\\_desc](#) [uct\\_md\\_resource\\_desc\\_t](#)  
*Memory domain resource descriptor.*
- typedef struct [uct\\_component\\_attr](#) [uct\\_component\\_attr\\_t](#)  
*UCT component attributes.*
- typedef struct [uct\\_tl\\_resource\\_desc](#) [uct\\_tl\\_resource\\_desc\\_t](#)  
*Communication resource descriptor.*

- typedef struct uct\_component \* uct\_component\_h
- typedef struct uct\_iface \* uct\_iface\_h
- typedef struct uct\_iface\_config uct\_iface\_config\_t
- typedef struct uct\_md\_config uct\_md\_config\_t
- typedef struct uct\_cm\_config uct\_cm\_config\_t
- typedef struct uct\_ep \* uct\_ep\_h
- typedef void \* uct\_mem\_h
- typedef uintptr\_t uct\_rkey\_t
- typedef struct uct\_md \* uct\_md\_h
- Memory domain handler.*
- typedef struct uct\_md\_ops uct\_md\_ops\_t
- typedef void \* uct\_rkey\_ctx\_h
- typedef struct uct\_iface\_attr uct\_iface\_attr\_t
- typedef struct uct\_iface\_params uct\_iface\_params\_t
- typedef struct uct\_ep\_attr uct\_ep\_attr\_t
- typedef struct uct\_md\_attr uct\_md\_attr\_t
- typedef struct uct\_completion uct\_completion\_t
- typedef struct uct\_pending\_req uct\_pending\_req\_t
- typedef struct uct\_worker \* uct\_worker\_h
- typedef struct uct\_md uct\_md\_t
- typedef enum uct\_am\_trace\_type uct\_am\_trace\_type\_t
- typedef struct uct\_device\_addr uct\_device\_addr\_t
- typedef struct uct\_iface\_addr uct\_iface\_addr\_t
- typedef struct uct\_ep\_addr uct\_ep\_addr\_t
- typedef struct uct\_ep\_params uct\_ep\_params\_t
- typedef struct uct\_ep\_connect\_params uct\_ep\_connect\_params\_t
- typedef struct uct\_cm\_attr uct\_cm\_attr\_t
- typedef struct uct\_cm uct\_cm\_t
- typedef uct\_cm\_t \* uct\_cm\_h
- typedef struct uct\_listener\_attr uct\_listener\_attr\_t
- typedef struct uct\_listener \* uct\_listener\_h
- typedef struct uct\_listener\_params uct\_listener\_params\_t
- typedef struct uct\_tag\_context uct\_tag\_context\_t
- typedef uint64\_t uct\_tag\_t
- typedef int uct\_worker\_cb\_id\_t
- typedef void \* uct\_conn\_request\_h
- typedef struct uct\_iov uct\_iov\_t
- Structure for scatter-gather I/O.*
- typedef void(\* uct\_completion\_callback\_t) (uct\_completion\_t \*self)
- Callback to process send completion.*
- typedef ucs\_status\_t(\* uct\_pending\_callback\_t) (uct\_pending\_req\_t \*self)
- Callback to process pending requests.*
- typedef ucs\_status\_t(\* uct\_error\_handler\_t) (void \*arg, uct\_ep\_h ep, ucs\_status\_t status)
- Callback to process peer failure.*
- typedef void(\* uct\_pending\_purge\_callback\_t) (uct\_pending\_req\_t \*self, void \*arg)
- Callback to purge pending requests.*
- typedef size\_t(\* uct\_pack\_callback\_t) (void \*dest, void \*arg)
- Callback for producing data.*
- typedef void(\* uct\_unpack\_callback\_t) (void \*arg, const void \*data, size\_t length)
- Callback for consuming data.*
- typedef void(\* uct\_async\_event\_cb\_t) (void \*arg, unsigned flags)
- Callback to process asynchronous events.*

## Enumerations

- enum `uct_component_attr_field` { `UCT_COMPONENT_ATTR_FIELD_NAME` = UCS\_BIT(0), `UCT_COMPONENT_ATTR_FIELD_MD_RESOURCES` = UCS\_BIT(1), `UCT_COMPONENT_ATTR_FIELD_FLUSH_FLAGS` = UCS\_BIT(2), `UCT_COMPONENT_ATTR_FIELD_FLUSH_THRESHOLD` = UCS\_BIT(3) }  
*UCT component attributes field mask.*
- enum { `UCT_COMPONENT_FLAG_CM` = UCS\_BIT(0) }  
*Capability flags of `uct_component_h`.*
- enum `uct_device_type_t` { `UCT_DEVICE_TYPE_NET`, `UCT_DEVICE_TYPE_SHM`, `UCT_DEVICE_TYPE_ACC`, `UCT_DEVICE_TYPE_SELF`, `UCT_DEVICE_TYPE_LAST` }  
*List of UCX device types.*
- enum `uct_iface_event_types` { `UCT_EVENT_SEND_COMP` = UCS\_BIT(0), `UCT_EVENT_RECV` = UCS\_BIT(1), `UCT_EVENT_RECV_SIG` = UCS\_BIT(2) }  
*Asynchronous event types.*
- enum `uct_flush_flags` { `UCT_FLUSH_FLAG_LOCAL` = 0, `UCT_FLUSH_FLAG_CANCEL` = UCS\_BIT(0) }  
*Flush modifiers.*
- enum `uct_progress_types` { `UCT_PROGRESS_SEND` = UCS\_BIT(0), `UCT_PROGRESS_RECV` = UCS\_BIT(1), `UCT_PROGRESS_THREAD_SAFE` = UCS\_BIT(7) }  
*UCT progress types.*
- enum `uct_cb_flags` { `UCT_CB_FLAG_RESERVED` = UCS\_BIT(1), `UCT_CB_FLAG_ASYNC` = UCS\_BIT(2) }  
*Callback flags.*
- enum `uct_iface_open_mode` { `UCT_IFACE_OPEN_MODE_DEVICE` = UCS\_BIT(0), `UCT_IFACE_OPEN_MODE_SOCKADDR` = UCS\_BIT(1), `UCT_IFACE_OPEN_MODE_SOCKADDR_CLIENT` = UCS\_BIT(2) }  
*Mode in which to open the interface.*
- enum `uct_iface_params_field` {  
`UCT_IFACE_PARAM_FIELD_CPU_MASK` = UCS\_BIT(0), `UCT_IFACE_PARAM_FIELD_OPEN_MODE` = UCS\_BIT(1), `UCT_IFACE_PARAM_FIELD_DEVICE` = UCS\_BIT(2), `UCT_IFACE_PARAM_FIELD_SOCKADDR` = UCS\_BIT(3),  
`UCT_IFACE_PARAM_FIELD_STATS_ROOT` = UCS\_BIT(4), `UCT_IFACE_PARAM_FIELD_RX_HEADROOM` = UCS\_BIT(5), `UCT_IFACE_PARAM_FIELD_ERR_HANDLER_ARG` = UCS\_BIT(6), `UCT_IFACE_PARAM_FIELD_ERR_HANDLER_FLAGS` = UCS\_BIT(7),  
`UCT_IFACE_PARAM_FIELD_ERR_HANDLER_FLAGS` = UCS\_BIT(8), `UCT_IFACE_PARAM_FIELD_HW_TM_EAGER_ARG` = UCS\_BIT(9), `UCT_IFACE_PARAM_FIELD_HW_TM_EAGER_CB` = UCS\_BIT(10), `UCT_IFACE_PARAM_FIELD_HW_TM_FLAGS` = UCS\_BIT(11),  
`UCT_IFACE_PARAM_FIELD_HW_TM_RNDV_CB` = UCS\_BIT(12), `UCT_IFACE_PARAM_FIELD_ASYNC_EVENT_ARG` = UCS\_BIT(13), `UCT_IFACE_PARAM_FIELD_ASYNC_EVENT_CB` = UCS\_BIT(14), `UCT_IFACE_PARAM_FIELD_KEEPA_LIVE` = UCS\_BIT(15),  
`UCT_IFACE_PARAM_FIELD_AM_ALIGNMENT` = UCS\_BIT(16), `UCT_IFACE_PARAM_FIELD_AM_ALIGN_OFFSET` = UCS\_BIT(17), `UCT_IFACE_PARAM_FIELD_FEATURES` = UCS\_BIT(18) }  
*UCT interface created by `uct_iface_open` parameters field mask.*
- enum `uct_ep_params_field` {  
`UCT_EP_PARAM_FIELD_IFACE` = UCS\_BIT(0), `UCT_EP_PARAM_FIELD_USER_DATA` = UCS\_BIT(1), `UCT_EP_PARAM_FIELD_DEV_ADDR` = UCS\_BIT(2), `UCT_EP_PARAM_FIELD_IFACE_ADDR` = UCS\_BIT(3),  
`UCT_EP_PARAM_FIELD_SOCKADDR` = UCS\_BIT(4), `UCT_EP_PARAM_FIELD_SOCKADDR_CB_FLAGS` = UCS\_BIT(5), `UCT_EP_PARAM_FIELD_SOCKADDR_PACK_CB` = UCS\_BIT(6), `UCT_EP_PARAM_FIELD_CM` = UCS\_BIT(7),  
`UCT_EP_PARAM_FIELD_CONN_REQUEST` = UCS\_BIT(8), `UCT_EP_PARAM_FIELD_SOCKADDR_CONNECT_CB_CLIENT` = UCS\_BIT(9), `UCT_EP_PARAM_FIELD_SOCKADDR_NOTIFY_CB_SERVER` = UCS\_BIT(10), `UCT_EP_PARAM_FIELD_SOCKADDR_NOTIFY_CB_CLIENT` = UCS\_BIT(11),  
`UCT_EP_PARAM_FIELD_PATH_INDEX` = UCS\_BIT(12), `UCT_EP_PARAM_FIELD_CM_RESOLVE_CB` = UCS\_BIT(13), `UCT_EP_PARAM_FIELD_PRIV_DATA` = UCS\_BIT(14), `UCT_EP_PARAM_FIELD_PRIV_DATA_LENGTH` = UCS\_BIT(15),  
`UCT_EP_PARAM_FIELD_LOCAL_SOCKADDR` = UCS\_BIT(16) }  
*UCT endpoint created by `uct_ep_create` parameters field mask.*



- enum `uct_iface_feature` {  
`UCT_IFACE_FEATURE_AM` = UCS\_BIT(0), `UCT_IFACE_FEATURE_PUT` = UCS\_BIT(1), `UCT_IFACE_FEATURE_GET`  
= UCS\_BIT(2), `UCT_IFACE_FEATURE_AMO32` = UCS\_BIT(3),  
`UCT_IFACE_FEATURE_AMO64` = UCS\_BIT(4), `UCT_IFACE_FEATURE_TAG` = UCS\_BIT(5), `UCT_IFACE_FEATURE_LAST`  
= UCS\_BIT(6) }  
*UCT interface configuration features.*
- enum { `UCT_TAG_RECV_CB_INLINE_DATA` = UCS\_BIT(0) }  
*flags of `uct_tag_context`.*
- enum `uct_cb_param_flags` { `UCT_CB_PARAM_FLAG_DESC` = UCS\_BIT(0), `UCT_CB_PARAM_FLAG_FIRST`  
= UCS\_BIT(1), `UCT_CB_PARAM_FLAG_MORE` = UCS\_BIT(2) }  
*Flags for active message and tag-matching offload callbacks (callback's parameters).*

## Functions

- `ucs_status_t uct_query_components` (`uct_component_h **components_p`, unsigned `*num_components_p`)  
*Query for list of components.*
- void `uct_release_component_list` (`uct_component_h *components`)  
*Release the list of components returned from `uct_query_components`.*
- `ucs_status_t uct_component_query` (`uct_component_h` component, `uct_component_attr_t *component_attr`)  
*Get component attributes.*
- `ucs_status_t uct_md_open` (`uct_component_h` component, const char `*md_name`, const `uct_md_config_t *config`, `uct_md_h *md_p`)  
*Open a memory domain.*
- void `uct_md_close` (`uct_md_h` md)  
*Close a memory domain.*
- `ucs_status_t uct_md_query_tl_resources` (`uct_md_h` md, `uct_tl_resource_desc_t **resources_p`, unsigned `*num_resources_p`)  
*Query for transport resources.*
- void `uct_release_tl_resource_list` (`uct_tl_resource_desc_t *resources`)  
*Release the list of resources returned from `uct_md_query_tl_resources`.*
- `ucs_status_t uct_md_iface_config_read` (`uct_md_h` md, const char `*tl_name`, const char `*env_prefix`, const char `*filename`, `uct_iface_config_t **config_p`)  
*Read transport-specific interface configuration.*
- void `uct_config_release` (void `*config`)  
*Release configuration memory returned from `uct_md_iface_config_read()`, `uct_md_config_read()`, or from `uct_cm_config_read()`.*
- `ucs_status_t uct_iface_open` (`uct_md_h` md, `uct_worker_h` worker, const `uct_iface_params_t *params`, const `uct_iface_config_t *config`, `uct_iface_h *iface_p`)  
*Open a communication interface.*
- void `uct_iface_close` (`uct_iface_h` iface)  
*Close and destroy an interface.*
- `ucs_status_t uct_iface_query` (`uct_iface_h` iface, `uct_iface_attr_t *iface_attr`)  
*Get interface attributes.*
- `ucs_status_t uct_iface_get_device_address` (`uct_iface_h` iface, `uct_device_addr_t *addr`)  
*Get address of the device the interface is using.*
- `ucs_status_t uct_iface_get_address` (`uct_iface_h` iface, `uct_iface_addr_t *addr`)  
*Get interface address.*
- int `uct_iface_is_reachable` (const `uct_iface_h` iface, const `uct_device_addr_t *dev_addr`, const `uct_iface_addr_t *iface_addr`)  
*Check if remote iface address is reachable.*
- `ucs_status_t uct_ep_check` (const `uct_ep_h` ep, unsigned flags, `uct_completion_t *comp`)  
*check if the destination endpoint is alive in respect to UCT library*

- [ucs\\_status\\_t uct\\_iface\\_event\\_fd\\_get](#) ([uct\\_iface\\_h](#) iface, int \*fd\_p)  
*Obtain a notification file descriptor for polling.*
- [ucs\\_status\\_t uct\\_iface\\_event\\_arm](#) ([uct\\_iface\\_h](#) iface, unsigned events)  
*Turn on event notification for the next event.*
- [ucs\\_status\\_t uct\\_iface\\_mem\\_alloc](#) ([uct\\_iface\\_h](#) iface, [size\\_t](#) length, unsigned flags, const char \*name, [uct\\_allocated\\_memory\\_t](#) \*mem)  
*Allocate memory which can be used for zero-copy communications.*
- void [uct\\_iface\\_mem\\_free](#) (const [uct\\_allocated\\_memory\\_t](#) \*mem)  
*Release memory allocated with [uct\\_iface\\_mem\\_alloc\(\)](#).*
- [ucs\\_status\\_t uct\\_ep\\_create](#) (const [uct\\_ep\\_params\\_t](#) \*params, [uct\\_ep\\_h](#) \*ep\_p)  
*Create new endpoint.*
- void [uct\\_ep\\_destroy](#) ([uct\\_ep\\_h](#) ep)  
*Destroy an endpoint.*
- [ucs\\_status\\_t uct\\_ep\\_get\\_address](#) ([uct\\_ep\\_h](#) ep, [uct\\_ep\\_addr\\_t](#) \*addr)  
*Get endpoint address.*
- [ucs\\_status\\_t uct\\_ep\\_connect\\_to\\_ep](#) ([uct\\_ep\\_h](#) ep, const [uct\\_device\\_addr\\_t](#) \*dev\_addr, const [uct\\_ep\\_addr\\_t](#) \*ep\_addr)  
*Connect endpoint to a remote endpoint.*
- [ucs\\_status\\_t uct\\_iface\\_flush](#) ([uct\\_iface\\_h](#) iface, unsigned flags, [uct\\_completion\\_t](#) \*comp)  
*Flush outstanding communication operations on an interface.*
- [ucs\\_status\\_t uct\\_iface\\_fence](#) ([uct\\_iface\\_h](#) iface, unsigned flags)  
*Ensures ordering of outstanding communications on the interface. Operations issued on the interface prior to this call are guaranteed to be completed before any subsequent communication operations to the same interface which follow the call to fence.*
- [ucs\\_status\\_t uct\\_ep\\_pending\\_add](#) ([uct\\_ep\\_h](#) ep, [uct\\_pending\\_req\\_t](#) \*req, unsigned flags)  
*Add a pending request to an endpoint.*
- void [uct\\_ep\\_pending\\_purge](#) ([uct\\_ep\\_h](#) ep, [uct\\_pending\\_purge\\_callback\\_t](#) cb, void \*arg)  
*Remove all pending requests from an endpoint.*
- [ucs\\_status\\_t uct\\_ep\\_flush](#) ([uct\\_ep\\_h](#) ep, unsigned flags, [uct\\_completion\\_t](#) \*comp)  
*Flush outstanding communication operations on an endpoint.*
- [ucs\\_status\\_t uct\\_ep\\_fence](#) ([uct\\_ep\\_h](#) ep, unsigned flags)  
*Ensures ordering of outstanding communications on the endpoint. Operations issued on the endpoint prior to this call are guaranteed to be completed before any subsequent communication operations to the same endpoint which follow the call to fence.*
- void [uct\\_iface\\_progress\\_enable](#) ([uct\\_iface\\_h](#) iface, unsigned flags)  
*Enable synchronous progress for the interface.*
- void [uct\\_iface\\_progress\\_disable](#) ([uct\\_iface\\_h](#) iface, unsigned flags)  
*Disable synchronous progress for the interface.*
- unsigned [uct\\_iface\\_progress](#) ([uct\\_iface\\_h](#) iface)  
*Perform a progress on an interface.*
- static UCS\_F\_ALWAYS\_INLINE void [uct\\_completion\\_update\\_status](#) ([uct\\_completion\\_t](#) \*comp, [ucs\\_status\\_t](#) status)  
*Update status of UCT completion handle.*

### 6.11.1 Detailed Description

This section describes a concept of the Communication Resource and routines associated with the concept.

### 6.11.2 Data Structure Documentation

#### 6.11.2.1 struct uct\_md\_resource\_desc

This structure describes a memory domain resource.

## Data Fields

char	md_name[UCT_MD_NAME_MAX]	Memory domain name
------	--------------------------	--------------------

## 6.11.2.2 struct uct\_component\_attr

This structure defines the attributes for UCT component. It is used for [uct\\_component\\_query](#)

## Examples

[uct\\_hello\\_world.c](#).

## Data Fields

uint64_t	field_mask	Mask of valid fields in this structure, using bits from <a href="#">uct_component_attr_field</a> . Fields not specified in this mask will be ignored. Provides ABI compatibility with respect to adding new fields.
char	name[UCT_COMPONENT_NAME_MAX]	Component name
unsigned	md_resource_count	Number of memory-domain resources
<a href="#">uct_md_resource_desc_t</a> *	md_resources	Array of memory domain resources. When used, it should be initialized prior to calling <a href="#">uct_component_query</a> with a pointer to an array, which is large enough to hold all memory domain resource entries. After the call, this array will be filled with information about existing memory domain resources. In order to allocate this array, you can call <a href="#">uct_component_query</a> twice: The first time would only obtain the amount of entries required, by specifying <a href="#">UCT_COMPONENT_ATTR_FIELD_MD_RESOURCE_COUNT</a> in <code>field_mask</code> . Then the array could be allocated with the returned number of entries, and passed to a second call to <a href="#">uct_component_query</a> , this time setting <code>field_mask</code> to <a href="#">UCT_COMPONENT_ATTR_FIELD_MD_RESOURCES</a> .
uint64_t	flags	Flags as defined by <a href="#">UCT_COMPONENT_FLAG_xx</a> .

## 6.11.2.3 struct uct\_tl\_resource\_desc

Resource descriptor is an object representing the network resource. Resource descriptor could represent a stand-alone communication resource such as an HCA port, network interface, or multiple resources such as multiple network interfaces or communication ports. It could also represent virtual communication resources that are defined over a single physical network interface.

## Examples

[uct\\_hello\\_world.c](#).

## Data Fields

char	tl_name[UCT_TL_NAME_MAX]	Transport name
char	dev_name[UCT_DEVICE_NAME_MAX]	Hardware device name
<a href="#">uct_device_type_t</a>	dev_type	The device represented by this resource (e.g. UCT_DEVICE_TYPE_NET for a network interface)
ucs_sys_device_t	sys_device	The identifier associated with the device bus_id as captured in ucs_sys_bus_id_t struct

## 6.11.2.4 struct uct\_iface\_attr

## Examples

[uct\\_hello\\_world.c](#).

## Data Fields

struct <a href="#">uct_iface_attr</a>	cap	Interface capabilities
size_t	device_addr_len	Size of device address
size_t	iface_addr_len	Size of interface address
size_t	ep_addr_len	Size of endpoint address
size_t	max_conn_priv	Max size of the iface's private data. used for connection establishment with sockaddr
struct sockaddr_storage	listen_sockaddr	Sockaddr on which this iface is listening.
double	overhead	Message overhead, seconds
<a href="#">uct_ppn_bandwidth_t</a>	bandwidth	Bandwidth model
ucs_linear_func_t	latency	Latency as function of number of active endpoints
uint8_t	priority	Priority of device
size_t	max_num_eps	Maximum number of endpoints
unsigned	dev_num_paths	How many network paths can be utilized on the device used by this interface for optimal performance. Endpoints that connect to the same remote address but use different paths can potentially achieve higher total bandwidth compared to using only a single endpoint.

## 6.11.2.5 struct uct\_iface\_attr.cap

## Data Fields

<a href="#">cap</a>	put	Attributes for PUT operations
<a href="#">cap</a>	get	Attributes for GET operations
<a href="#">cap</a>	am	Attributes for AM operations
<a href="#">cap</a>	tag	Attributes for TAG operations
<a href="#">cap</a>	atomic32	
<a href="#">cap</a>	atomic64	Attributes for atomic operations

## Data Fields

uint64_t	flags	Flags from <a href="#">UCT interface operations and capabilities</a>
uint64_t	event_flags	Flags from <a href="#">UCT interface for asynchronous event capabilities</a>

## 6.11.2.6 struct uct\_iface\_attr.cap.put

## Data Fields

size_t	max_short	Maximal size for put_short
size_t	max_bcopy	Maximal size for put_bcopy
size_t	min_zcopy	Minimal size for put_zcopy (total of <a href="#">uct_iov_t::length</a> of the <i>iov</i> parameter)
size_t	max_zcopy	Maximal size for put_zcopy (total of <a href="#">uct_iov_t::length</a> of the <i>iov</i> parameter)
size_t	opt_zcopy_align	Optimal alignment for zero-copy buffer address
size_t	align_mtu	MTU used for alignment
size_t	max_iov	Maximal <i>iovcnt</i> parameter in <a href="#">uct_ep_put_zcopy</a>

## 6.11.2.7 struct uct\_iface\_attr.cap.get

## Data Fields

size_t	max_short	Maximal size for get_short
size_t	max_bcopy	Maximal size for get_bcopy
size_t	min_zcopy	Minimal size for get_zcopy (total of <a href="#">uct_iov_t::length</a> of the <i>iov</i> parameter)
size_t	max_zcopy	Maximal size for get_zcopy (total of <a href="#">uct_iov_t::length</a> of the <i>iov</i> parameter)
size_t	opt_zcopy_align	Optimal alignment for zero-copy buffer address
size_t	align_mtu	MTU used for alignment
size_t	max_iov	Maximal <i>iovcnt</i> parameter in <a href="#">uct_ep_get_zcopy</a>

## 6.11.2.8 struct uct\_iface\_attr.cap.am

## Data Fields

size_t	max_short	Total maximum size (incl. the header)
size_t	max_bcopy	Total maximum size (incl. the header)
size_t	min_zcopy	Minimal size for am_zcopy (incl. the header and total of <a href="#">uct_iov_t::length</a> of the <i>iov</i> parameter)
size_t	max_zcopy	Total max. size (incl. the header and total of <a href="#">uct_iov_t::length</a> of the <i>iov</i> parameter)
size_t	opt_zcopy_align	Optimal alignment for zero-copy buffer address
size_t	align_mtu	MTU used for alignment
size_t	max_hdr	Max. header size for zcopy
size_t	max_iov	Maximal <i>iovcnt</i> parameter in <a href="#">uct_ep_am_zcopy</a>

## 6.11.2.9 struct uct\_iface\_attr.cap.tag

## Data Fields

<a href="#">tag</a>	recv	
<a href="#">tag</a>	eager	Attributes related to eager protocol

## Data Fields

<a href="#">tag</a>	<code>rndv</code>	Attributes related to rendezvous protocol
---------------------	-------------------	---

**6.11.2.10 struct `uct_iface_attr.cap.tag.recv`**

## Data Fields

<code>size_t</code>	<code>min_recv</code>	Minimal allowed length of posted receive buffer
<code>size_t</code>	<code>max_zcopy</code>	Maximal allowed data length in <a href="#">uct_iface_tag_recv_zcopy</a>
<code>size_t</code>	<code>max_iov</code>	Maximal <code>iovcnt</code> parameter in <a href="#">uct_iface_tag_recv_zcopy</a>
<code>size_t</code>	<code>max_outstanding</code>	Maximal number of simultaneous receive operations

**6.11.2.11 struct `uct_iface_attr.cap.tag.eager`**

## Data Fields

<code>size_t</code>	<code>max_short</code>	Maximal allowed data length in <a href="#">uct_ep_tag_eager_short</a>
<code>size_t</code>	<code>max_bcopy</code>	Maximal allowed data length in <a href="#">uct_ep_tag_eager_bcopy</a>
<code>size_t</code>	<code>max_zcopy</code>	Maximal allowed data length in <a href="#">uct_ep_tag_eager_zcopy</a>
<code>size_t</code>	<code>max_iov</code>	Maximal <code>iovcnt</code> parameter in <a href="#">uct_ep_tag_eager_zcopy</a>

**6.11.2.12 struct `uct_iface_attr.cap.tag.rndv`**

## Data Fields

<code>size_t</code>	<code>max_zcopy</code>	Maximal allowed data length in <a href="#">uct_ep_tag_rndv_zcopy</a>
<code>size_t</code>	<code>max_hdr</code>	Maximal allowed header length in <a href="#">uct_ep_tag_rndv_zcopy</a> and <a href="#">uct_ep_tag_rndv_request</a>
<code>size_t</code>	<code>max_iov</code>	Maximal <code>iovcnt</code> parameter in <a href="#">uct_ep_tag_rndv_zcopy</a>

**6.11.2.13 struct `uct_iface_attr.cap.atomic32`**

## Data Fields

<code>uint64_t</code>	<code>op_flags</code>	Attributes for atomic-post operations
<code>uint64_t</code>	<code>fop_flags</code>	Attributes for atomic-fetch operations

**6.11.2.14 struct `uct_iface_attr.cap.atomic64`**

## Data Fields

<code>uint64_t</code>	<code>op_flags</code>	Attributes for atomic-post operations
<code>uint64_t</code>	<code>fop_flags</code>	Attributes for atomic-fetch operations

**6.11.2.15 struct `uct_iface_params`**

This structure should be allocated by the user and should be passed to [uct\\_iface\\_open](#). User has to initialize all fields of this structure.

## Examples

[uct\\_hello\\_world.c](#).

## Data Fields

uint64_t	field_mask	Mask of valid fields in this structure, using bits from <a href="#">uct_iface_params_field</a> . Fields not specified in this mask will be ignored.
ucs_cpu_set_t	cpu_mask	Mask of CPUs to use for resources
uint64_t	open_mode	Interface open mode bitmap. <a href="#">uct_iface_open_mode</a>
union <a href="#">uct_iface_params</a>	mode	Mode-specific parameters
ucs_stats_node_t *	stats_root	Root in the statistics tree. Can be NULL. If non NULL, it will be a root of <i>uct_iface</i> object in the statistics tree.
size_t	rx_headroom	How much bytes to reserve before the receive segment.
void *	err_handler_arg	Custom argument of <i>err_handler</i> .
<a href="#">uct_error_handler_t</a>	err_handler	The callback to handle transport level error.
uint32_t	err_handler_flags	Callback flags to indicate where the <i>err_handler</i> callback can be invoked from. <a href="#">uct_cb_flags</a>
void *	eager_arg	These callbacks are only relevant for HW Tag Matching
<a href="#">uct_tag_unexp_eager_cb_t</a>	eager_cb	Callback for tag matching unexpected eager messages
void *	rndv_arg	
<a href="#">uct_tag_unexp_rndv_cb_t</a>	rndv_cb	Callback for tag matching unexpected rndv messages
void *	async_event_arg	
<a href="#">uct_async_event_cb_t</a>	async_event_cb	Callback for asynchronous event handling. The callback will be invoked from UCT transport when there are new events to be read by user if the iface has <a href="#">UCT_IFACE_FLAG_EVENT_ASYNC_CB</a> capability
<a href="#">ucs_time_t</a>	keepalive_interval	
size_t	am_alignment	Desired alignment for Active Messages on the receiver. Note that only data received in the UCT descriptor can be aligned (i.e. <a href="#">UCT_CB_PARAM_FLAG_DESC</a> flag is provided in the Active Message handler callback). The provided value must be power of 2. The default value is 1.
size_t	am_align_offset	Offset in the Active Message receive buffer, which should be aligned to the <i>am_alignment</i> boundary. Note this parameter has no effect without setting <i>am_alignment</i> parameter. The provided value must be less than the given <i>am_alignment</i> value. The default value is 0. <pre> +++ pointer to data in uct_am_callback_t     + alignment boundary     v v +-----+   align       offset     +-----+ </pre>
uint64_t	features	UCT <a href="#">features</a> that are used for interface initialization.

6.11.2.16 union [uct\\_iface\\_params.mode](#)

Mode-specific parameters

## Data Fields

<a href="#">mode</a>	device	The fields in this structure ( <i>tl_name</i> and <i>dev_name</i> ) need to be set only when the <a href="#">UCT_IFACE_OPEN_MODE_DEVICE</a> bit is set in <a href="#">uct_iface_params_t::open_mode</a> . This will make <a href="#">uct_iface_open</a> open the interface on the specified device.
----------------------	--------	---

## Data Fields

<a href="#">mode</a>	sockaddr	These callbacks and address are only relevant for client-server connection establishment with sockaddr and are needed on the server side. The callbacks and address need to be set when the <a href="#">UCT_IFACE_OPEN_MODE_SOCKADDR_SERVER</a> bit is set in <a href="#">uct_iface_params_t::open_mode</a> . This will make <a href="#">uct_iface_open</a> open the interface on the specified address as a server.
----------------------	----------	--

**6.11.2.17 struct uct\_iface\_params.mode.device**

The fields in this structure ([tl\\_name](#) and [dev\\_name](#)) need to be set only when the [UCT\\_IFACE\\_OPEN\\_MODE\\_DEVICE](#) bit is set in [uct\\_iface\\_params\\_t::open\\_mode](#). This will make [uct\\_iface\\_open](#) open the interface on the specified device.

## Data Fields

const char *	<a href="#">tl_name</a>	Transport name
const char *	<a href="#">dev_name</a>	Device Name

**6.11.2.18 struct uct\_iface\_params.mode.sockaddr**

These callbacks and address are only relevant for client-server connection establishment with sockaddr and are needed on the server side. The callbacks and address need to be set when the [UCT\\_IFACE\\_OPEN\\_MODE\\_SOCKADDR\\_SERVER](#) bit is set in [uct\\_iface\\_params\\_t::open\\_mode](#). This will make [uct\\_iface\\_open](#) open the interface on the specified address as a server.

## Data Fields

<a href="#">ucs_sock_addr_t</a>	<a href="#">listen_sockaddr</a>	
void *	<a href="#">conn_request_arg</a>	Argument for connection request callback
<a href="#">uct_sockaddr_conn_request_callback_t</a>	<a href="#">conn_request_cb</a>	Callback for an incoming connection request on the server
uint32_t	<a href="#">cb_flags</a>	Callback flags to indicate where the callback can be invoked from. <a href="#">uct_cb_flags</a>

**6.11.2.19 struct uct\_ep\_params**

## Examples

[uct\\_hello\\_world.c](#).

## Data Fields

uint64_t	<a href="#">field_mask</a>	Mask of valid fields in this structure, using bits from <a href="#">uct_ep_params_field</a> . Fields not specified by this mask will be ignored.
<a href="#">uct_iface_h</a>	<a href="#">iface</a>	Interface to create the endpoint on. Either <i>iface</i> or <i>cm</i> field must be initialized but not both.
void *	<a href="#">user_data</a>	User data associated with the endpoint.



## Data Fields

<code>const uct_device_addr_t *</code>	<code>dev_addr</code>	The device address to connect to on the remote peer. This must be defined together with <code>uct_ep_params_t::iface_addr</code> to create an endpoint connected to a remote interface.
<code>const uct_iface_addr_t *</code>	<code>iface_addr</code>	This specifies the remote address to use when creating an endpoint that is connected to a remote interface.  <b>Note</b>  This requires <code>UCT_IFACE_FLAG_CONNECT_TO_IFACE</code> capability.
<code>const ucs_sock_addr_t *</code>	<code>sockaddr</code>	The sockaddr to connect to on the remote peer. If set, <code>uct_ep_create</code> will create an endpoint for a connection to the remote peer, specified by its socket address.  <b>Note</b>  The interface in this routine requires the <code>UCT_IFACE_FLAG_CONNECT_TO_SOCKADDR</code> capability.
<code>uint32_t</code>	<code>sockaddr_cb_flags</code>	<code>uct_cb_flags</code> to indicate <code>uct_ep_params_t::sockaddr_pack_cb</code> , <code>uct_ep_params_t::sockaddr_cb_client</code> , <code>uct_ep_params_t::sockaddr_cb_server</code> , <code>uct_ep_params_t::disconnect_cb</code> and <code>uct_ep_params_t::cm_resolve_cb</code> behavior. If none from these are not set, this field will be ignored.
<code>uct_cm_ep_priv_data_pack_callback_t</code>	<code>sockaddr_pack_cb</code>	Callback that will be used for filling the user's private data to be delivered to the remote peer by the callback on the server or client side. This field is only valid if <code>uct_ep_params_t::sockaddr</code> is set.  <b>Note</b>  It is never guaranteed that the callback will be called. If, for example, the endpoint goes into error state before issuing the connection request, the callback will not be invoked.  Can not be set together with <code>uct_ep_params_t::private_data</code> or <code>uct_ep_params_t::cm_resolve_cb</code> .
<code>uct_cm_h</code>	<code>cm</code>	The connection manager object as created by <code>uct_cm_open</code> . Either <code>cm</code> or <code>iface</code> field must be initialized but not both.

## Data Fields

<a href="#">uct_conn_request_h</a>	conn_request	<p>Connection request that was passed to <a href="#">uct_cm_listener_conn_request_args_t::conn_request</a>.</p> <p><b>Note</b></p> <p>After a call to <a href="#">uct_ep_create</a>, <i>params.conn_request</i> is consumed and should not be used anymore, even if the call returns with an error.</p>
<a href="#">uct_cm_ep_client_connect_callback_t</a>	sockaddr_cb_client	<p>Callback that will be invoked when the endpoint on the client side is being connected to the server by a connection manager <a href="#">uct_cm_h</a>.</p>
<a href="#">uct_cm_ep_server_conn_notify_callback_t</a>	sockaddr_cb_server	<p>Callback that will be invoked when the endpoint on the server side is being connected to a client by a connection manager <a href="#">uct_cm_h</a>.</p>
<a href="#">uct_ep_disconnect_cb_t</a>	disconnect_cb	<p>Callback that will be invoked when the endpoint is disconnected.</p>
unsigned	path_index	<p>Index of the path which the endpoint should use, must be in the range 0..(<a href="#">uct_iface_attr_t::dev_num_paths</a> - 1).</p>
<a href="#">uct_cm_ep_resolve_callback_t</a>	cm_resolve_cb	<p>This callback is invoked when the remote server address provided in field <a href="#">uct_ep_params_t::sockaddr</a> is resolved to the local device to be used for connection establishment.</p> <p><b>Note</b></p> <p>This field is mutually exclusive with <a href="#">uct_ep_params::sockaddr_pack_cb</a>.</p>
const void *	private_data	<p>Private data to be passed from server to client. Can be used only along with <a href="#">uct_ep_params::conn_request</a>.</p> <p><b>Note</b></p> <p>This field is mutually exclusive with <a href="#">uct_ep_params::sockaddr_pack_cb</a>.</p>
size_t	private_data_length	<p>Length of <a href="#">uct_ep_params::private_data</a>, the maximal allowed value is indicated by the <a href="#">uct_cm_attr::max_conn_priv</a>.</p>
const <a href="#">ucs_sock_addr_t</a> *	local_sockaddr	<p>The sockaddr to bind locally. If set, <a href="#">uct_ep_create</a> will create an endpoint binding to this local sockaddr.</p> <p><b>Note</b></p> <p>The interface in this routine requires the <a href="#">UCT_IFACE_FLAG_CONNECT_TO_SOCKADDR</a> capability.</p>

### 6.11.2.20 struct uct\_completion

This structure should be allocated by the user and can be passed to communication primitives. The user must initialize all fields of the structure. If the operation returns UCS\_INPROGRESS, this structure will be in use by the transport until the operation completes. When the operation completes, "count" field is decremented by 1, and whenever it reaches 0 - the callback is called.

Notes:

- The same structure can be passed multiple times to communication functions without the need to wait for completion.
- If the number of operations is smaller than the initial value of the counter, the callback will not be called at all, so it may be left undefined.
- status field is required to track the first time the error occurred, and report it via a callback when count reaches 0.

Examples

[uct\\_hello\\_world.c](#).

Data Fields

<a href="#">uct_completion_callback_t</a>	func	User callback function
int	count	Completion counter
<a href="#">ucs_status_t</a>	status	Completion status, this field must be initialized with UCS_OK before first operation is started.

### 6.11.2.21 struct uct\_pending\_req

This structure should be passed to [uct\\_ep\\_pending\\_add\(\)](#) and is used to signal new available resources back to user.

Data Fields

<a href="#">uct_pending_callback_t</a>	func	User callback function
char	priv[UCT_PENDING_REQ_PRIV_LEN]	Used internally by UCT

### 6.11.2.22 struct uct\_iov

Specifies a list of buffers which can be used within a single data transfer function call.

```
buffer
|
|-----+-----+-----+-----+-----+
| payload | empty | payload | empty | payload |
|-----+-----+-----+-----+-----+
|<-length-->|      |<-length-->|      |<-length-->|
|<---- stride ---->|<---- stride ---->|
```

**Note**

The sum of lengths in all iov list must be less or equal to `max_zcopy` of the respective communication operation.

If `length` or `count` are zero, the memory pointed to by `buffer` will not be accessed. Otherwise, `buffer` must point to valid memory.

If `count` is one, every iov entry specifies a single contiguous data block

If `count` > 1, each iov entry specifies a strided block of `count` elements and distance of `stride` byte between consecutive elements

**Examples**

[uct\\_hello\\_world.c](#).

**Data Fields**

void *	buffer	Data buffer
size_t	length	Length of the payload in bytes
<a href="#">uct_mem_h</a>	memh	Local memory key descriptor for the data
size_t	stride	Stride between beginnings of payload elements in the buffer in bytes
unsigned	count	Number of payload elements in the buffer

**6.11.3 Typedef Documentation****6.11.3.1 uct\_md\_resource\_desc\_t**

```
typedef struct uct_md_resource_desc uct_md_resource_desc_t
```

This structure describes a memory domain resource.

**6.11.3.2 uct\_component\_attr\_t**

```
typedef struct uct_component_attr uct_component_attr_t
```

This structure defines the attributes for UCT component. It is used for [uct\\_component\\_query](#)

**6.11.3.3 uct\_tl\_resource\_desc\_t**

```
typedef struct uct_tl_resource_desc uct_tl_resource_desc_t
```

Resource descriptor is an object representing the network resource. Resource descriptor could represent a stand-alone communication resource such as an HCA port, network interface, or multiple resources such as multiple network interfaces or communication ports. It could also represent virtual communication resources that are defined over a single physical network interface.

**6.11.3.4 uct\_component\_h**

```
typedef struct uct_component* uct_component_h
```

**6.11.3.5 uct\_iface\_h**

```
typedef struct uct_iface* uct_iface_h
```

**6.11.3.6 uct\_iface\_config\_t**

```
typedef struct uct_iface_config uct_iface_config_t
```

**6.11.3.7 uct\_md\_config\_t**

```
typedef struct uct_md_config uct_md_config_t
```

**6.11.3.8 uct\_cm\_config\_t**

```
typedef struct uct_cm_config uct_cm_config_t
```

**6.11.3.9 uct\_ep\_h**

```
typedef struct uct_ep* uct_ep_h
```

**6.11.3.10 uct\_mem\_h**

```
typedef void* uct_mem_h
```

**6.11.3.11 uct\_rkey\_t**

```
typedef uintptr_t uct_rkey_t
```

**6.11.3.12 uct\_md\_h**

```
typedef struct uct_md* uct_md_h
```

**6.11.3.13 uct\_md\_ops\_t**

```
typedef struct uct_md_ops uct_md_ops_t
```

**6.11.3.14 uct\_rkey\_ctx\_h**

```
typedef void* uct_rkey_ctx_h
```

**6.11.3.15 uct\_iface\_attr\_t**

```
typedef struct uct_iface_attr uct_iface_attr_t
```

**6.11.3.16 uct\_iface\_params\_t**

```
typedef struct uct_iface_params uct_iface_params_t
```

**6.11.3.17 uct\_ep\_attr\_t**

```
typedef struct uct_ep_attr uct_ep_attr_t
```

**6.11.3.18 uct\_md\_attr\_t**

```
typedef struct uct_md_attr uct_md_attr_t
```

**6.11.3.19 uct\_completion\_t**

```
typedef struct uct_completion uct_completion_t
```

**6.11.3.20 uct\_pending\_req\_t**

```
typedef struct uct_pending_req uct_pending_req_t
```

**6.11.3.21 uct\_worker\_h**

```
typedef struct uct_worker* uct_worker_h
```

**6.11.3.22 uct\_md\_t**

```
typedef struct uct_md uct_md_t
```

**6.11.3.23 uct\_am\_trace\_type\_t**

```
typedef enum uct_am_trace_type uct_am_trace_type_t
```

**6.11.3.24 uct\_device\_addr\_t**

```
typedef struct uct_device_addr uct_device_addr_t
```

**6.11.3.25 uct\_iface\_addr\_t**

```
typedef struct uct_iface_addr uct_iface_addr_t
```

**6.11.3.26 uct\_ep\_addr\_t**

```
typedef struct uct_ep_addr uct_ep_addr_t
```

**6.11.3.27 uct\_ep\_params\_t**

```
typedef struct uct_ep_params uct_ep_params_t
```

**6.11.3.28 uct\_ep\_connect\_params\_t**

```
typedef struct uct_ep_connect_params uct_ep_connect_params_t
```

**6.11.3.29 uct\_cm\_attr\_t**

```
typedef struct uct_cm_attr uct_cm_attr_t
```

**6.11.3.30 uct\_cm\_t**

```
typedef struct uct_cm uct_cm_t
```

**6.11.3.31 uct\_cm\_h**

```
typedef uct_cm_t* uct_cm_h
```

**6.11.3.32 uct\_listener\_attr\_t**

```
typedef struct uct_listener_attr uct_listener_attr_t
```

**6.11.3.33 uct\_listener\_h**

```
typedef struct uct_listener* uct_listener_h
```

**6.11.3.34 uct\_listener\_params\_t**

```
typedef struct uct_listener_params uct_listener_params_t
```

**6.11.3.35 uct\_tag\_context\_t**

```
typedef struct uct_tag_context uct_tag_context_t
```

**6.11.3.36 uct\_tag\_t**

```
typedef uint64_t uct_tag_t
```

**6.11.3.37 uct\_worker\_cb\_id\_t**

```
typedef int uct_worker_cb_id_t
```

**6.11.3.38 uct\_conn\_request\_h**

```
typedef void* uct_conn_request_h
```

**6.11.3.39 uct\_iov\_t**

```
typedef struct uct_iov uct_iov_t
```

Specifies a list of buffers which can be used within a single data transfer function call.

```
buffer
|
+-----+-----+-----+-----+-----+
| payload | empty | payload | empty | payload |
+-----+-----+-----+-----+-----+
|<-length-->|      |<-length-->|      |<-length-->|
|<---- stride ---->|<---- stride ---->|
```



**Note**

The sum of lengths in all iov list must be less or equal to `max_zcopy` of the respective communication operation.

If `length` or `count` are zero, the memory pointed to by `buffer` will not be accessed. Otherwise, `buffer` must point to valid memory.

If `count` is one, every iov entry specifies a single contiguous data block

If `count` > 1, each iov entry specifies a strided block of `count` elements and distance of `stride` byte between consecutive elements

**6.11.3.40 uct\_completion\_callback\_t**

```
typedef void(* uct_completion_callback_t) (uct_completion_t *self)
```

**Parameters**

in	<i>self</i>	Pointer to relevant completion structure, which was initially passed to the operation.
----	-------------	--

**6.11.3.41 uct\_pending\_callback\_t**

```
typedef ucs_status_t(* uct_pending_callback_t) (uct_pending_req_t *self)
```

**Parameters**

in	<i>self</i>	Pointer to relevant pending structure, which was initially passed to the operation.
----	-------------	---

**Returns**

[UCS\\_OK](#) - This pending request has completed and should be removed. [UCS\\_INPROGRESS](#) - Some progress was made, but not completed. Keep this request and keep processing the queue. Otherwise - Could not make any progress. Keep this pending request on the queue, and stop processing the queue.

**6.11.3.42 uct\_error\_handler\_t**

```
typedef ucs_status_t(* uct_error_handler_t) (void *arg, uct_ep_h ep, ucs_status_t status)
```

**Note**

User should purge a pending queue and do not post any TX operations and cancel all possible outstanding operations prior closing a UCT endpoint.

**Parameters**

in	<i>arg</i>	User argument to be passed to the callback.
in	<i>ep</i>	Endpoint which has failed. Upon return from the callback, this <i>ep</i> is no longer usable and all subsequent operations on this <i>ep</i> will fail with the error code passed in <i>status</i> .
in	<i>status</i>	Status indicating error.

**Returns**

[UCT\\_OK](#) - The error was handled successfully. Otherwise - The error was not handled and is returned back to the transport.

**6.11.3.43 uct\_pending\_purge\_callback\_t**

```
typedef void(* uct_pending_purge_callback_t) (uct_pending_req_t *self, void *arg)
```

**Parameters**

in	<i>self</i>	Pointer to relevant pending structure, which was initially passed to the operation.
in	<i>arg</i>	User argument to be passed to the callback.

**6.11.3.44 uct\_pack\_callback\_t**

```
typedef size_t(* uct_pack_callback_t) (void *dest, void *arg)
```

**Parameters**

in	<i>dest</i>	Memory buffer to pack the data to.
in	<i>arg</i>	Custom user-argument.

**Returns**

Size of the data was actually produced.

**6.11.3.45 uct\_unpack\_callback\_t**

```
typedef void(* uct_unpack_callback_t) (void *arg, const void *data, size_t length)
```

**Parameters**

in	<i>arg</i>	Custom user-argument.
in	<i>data</i>	Memory buffer to unpack the data from.
in	<i>length</i>	How much data to consume (size of "data")

**Note**

The arguments for this callback are in the same order as libc's memcpy().

**6.11.3.46 uct\_async\_event\_cb\_t**

```
typedef void(* uct_async_event_cb_t) (void *arg, unsigned flags)
```

## Parameters

in	<i>arg</i>	User argument to be passed to the callback.
in	<i>flags</i>	Flags to be passed to the callback (reserved for future use).

## 6.11.4 Enumeration Type Documentation

### 6.11.4.1 `uct_component_attr_field`

enum `uct_component_attr_field`

The enumeration allows specifying which fields in `uct_component_attr_t` are present. It is used for backward compatibility support.

## Enumerator

UCT_COMPONENT_ATTR_FIELD_NAME	Component name
UCT_COMPONENT_ATTR_FIELD_MD_RESOURCE_COUNT	MD resource count
UCT_COMPONENT_ATTR_FIELD_MD_RESOURCES	MD resources array
UCT_COMPONENT_ATTR_FIELD_FLAGS	Capability flags

### 6.11.4.2 `anonymous enum`

anonymous enum

The enumeration defines bit mask of `uct_component_h` capabilities in `uct_component_attr_t::flags` which is set by `uct_component_query`.

## Enumerator

UCT_COMPONENT_FLAG_CM	If set, the component supports <code>uct_cm_h</code> functionality. See <code>uct_cm_open</code> for details.
-----------------------	---

### 6.11.4.3 `uct_device_type_t`

enum `uct_device_type_t`

## Enumerator

UCT_DEVICE_TYPE_NET	Network devices
UCT_DEVICE_TYPE_SHM	Shared memory devices
UCT_DEVICE_TYPE_ACC	Acceleration devices
UCT_DEVICE_TYPE_SELF	Loop-back device
UCT_DEVICE_TYPE_LAST	

#### 6.11.4.4 `uct_iface_event_types`

enum `uct_iface_event_types`

##### Note

The `UCT_EVENT_RECV` and `UCT_EVENT_RECV_SIG` event types are used to indicate receive-side completions for both tag matching and active messages. If the interface supports signaled receives (`UCT_IFACE_FLAG_EVENT_RECV_SIG`), then for the messages sent with `UCT_SEND_FLAG_SIGNALED` flag, `UCT_EVENT_RECV_SIG` should be triggered on the receiver. Otherwise, `UCT_EVENT_RECV` should be triggered.

##### Enumerator

<code>UCT_EVENT_SEND_COMP</code>	Send completion event
<code>UCT_EVENT_RECV</code>	Tag or active message received
<code>UCT_EVENT_RECV_SIG</code>	Signaled tag or active message received

#### 6.11.4.5 `uct_flush_flags`

enum `uct_flush_flags`

##### Enumerator

<code>UCT_FLUSH_FLAG_LOCAL</code>	Guarantees that the data transfer is completed but the target buffer may not be updated yet.
<code>UCT_FLUSH_FLAG_CANCEL</code>	The library will make a best effort attempt to cancel all uncompleted operations. However, there is a chance that some operations will not be canceled in which case the user will need to handle their completions through the relevant callbacks. After <code>uct_ep_flush</code> with this flag is completed, the endpoint will be set to error state, and it becomes unusable for send operations and should be destroyed.

#### 6.11.4.6 `uct_progress_types`

enum `uct_progress_types`

##### Enumerator

<code>UCT_PROGRESS_SEND</code>	Progress send operations
<code>UCT_PROGRESS_RECV</code>	Progress receive operations
<code>UCT_PROGRESS_THREAD_SAFE</code>	Enable/disable progress while another thread may be calling <code>ucp_worker_progress()</code> .

### 6.11.4.7 uct\_cb\_flags

enum `uct_cb_flags`

List of flags for a callback.

#### Enumerator

<code>UCT_CB_FLAG_RESERVED</code>	Reserved for future use.
<code>UCT_CB_FLAG_ASYNC</code>	Callback is allowed to be called from any thread in the process, and therefore should be thread-safe. For example, it may be called from a transport async progress thread. To guarantee async invocation, the interface must have the <code>UCT_IFACE_FLAG_CB_ASYNC</code> flag set. If async callback is requested on an interface which only supports sync callback (i.e., only the <code>UCT_IFACE_FLAG_CB_SYNC</code> flag is set), the callback will be invoked only from the context that called <code>uct_iface_progress</code> .

### 6.11.4.8 uct\_iface\_open\_mode

enum `uct_iface_open_mode`

#### Enumerator

<code>UCT_IFACE_OPEN_MODE_DEVICE</code>	Interface is opened on a specific device
<code>UCT_IFACE_OPEN_MODE_SOCKADDR_SERVER</code>	Interface is opened on a specific address on the server side. This mode will be deprecated in the near future for a better API.
<code>UCT_IFACE_OPEN_MODE_SOCKADDR_CLIENT</code>	Interface is opened on a specific address on the client side. This mode will be deprecated in the near future for a better API.

### 6.11.4.9 uct\_iface\_params\_field

enum `uct_iface_params_field`

The enumeration allows specifying which fields in `uct_iface_params_t` are present, for backward compatibility support.

#### Enumerator

<code>UCT_IFACE_PARAM_FIELD_CPU_MASK</code>	Enables <code>uct_iface_params_t::cpu_mask</code>
<code>UCT_IFACE_PARAM_FIELD_OPEN_MODE</code>	Enables <code>uct_iface_params_t::open_mode</code>
<code>UCT_IFACE_PARAM_FIELD_DEVICE</code>	Enables <code>uct_iface_params_t::mode::device</code>
<code>UCT_IFACE_PARAM_FIELD_SOCKADDR</code>	Enables <code>uct_iface_params_t::mode::sockaddr</code>
<code>UCT_IFACE_PARAM_FIELD_STATS_ROOT</code>	Enables <code>uct_iface_params_t::stats_root</code>
<code>UCT_IFACE_PARAM_FIELD_RX_HEADROOM</code>	Enables <code>uct_iface_params_t::rx_headroom</code>
<code>UCT_IFACE_PARAM_FIELD_ERR_HANDLER_ARG</code>	Enables <code>uct_iface_params_t::err_handler_arg</code>
<code>UCT_IFACE_PARAM_FIELD_ERR_HANDLER</code>	Enables <code>uct_iface_params_t::err_handler</code>
<code>UCT_IFACE_PARAM_FIELD_ERR_HANDLER_FLAGS</code>	Enables <code>uct_iface_params_t::err_handler_flags</code>
<code>UCT_IFACE_PARAM_FIELD_HW_TM_EAGER_ARG</code>	Enables <code>uct_iface_params_t::eager_arg</code>

## Enumerator

UCT_IFACE_PARAM_FIELD_HW_TM_EAGER_CB	Enables <a href="#">uct_iface_params_t::eager_cb</a>
UCT_IFACE_PARAM_FIELD_HW_TM_RNDV_ARG	Enables <a href="#">uct_iface_params_t::rndv_arg</a>
UCT_IFACE_PARAM_FIELD_HW_TM_RNDV_CB	Enables <a href="#">uct_iface_params_t::rndv_cb</a>
UCT_IFACE_PARAM_FIELD_ASYNC_EVENT_ARG	Enables <a href="#">uct_iface_params_t::async_event_arg</a>
UCT_IFACE_PARAM_FIELD_ASYNC_EVENT_CB	Enables <a href="#">uct_iface_params_t::async_event_cb</a>
UCT_IFACE_PARAM_FIELD_KEEPLIVE_INTERVAL	Enables <a href="#">uct_iface_params_t::keepalive_interval</a>
UCT_IFACE_PARAM_FIELD_AM_ALIGNMENT	Enables <a href="#">uct_iface_params_t::am_alignment</a>
UCT_IFACE_PARAM_FIELD_AM_ALIGN_OFFSET	Enables <a href="#">uct_iface_params_t::am_align_offset</a>
UCT_IFACE_PARAM_FIELD_FEATURES	Enables <a href="#">uct_iface_params_t::features</a>

6.11.4.10 `uct_ep_params_field`

enum [uct\\_ep\\_params\\_field](#)

The enumeration allows specifying which fields in [uct\\_ep\\_params\\_t](#) are present, for backward compatibility support.

## Enumerator

UCT_EP_PARAM_FIELD_IFACE	Enables <a href="#">uct_ep_params::iface</a>
UCT_EP_PARAM_FIELD_USER_DATA	Enables <a href="#">uct_ep_params::user_data</a>
UCT_EP_PARAM_FIELD_DEV_ADDR	Enables <a href="#">uct_ep_params::dev_addr</a>
UCT_EP_PARAM_FIELD_IFACE_ADDR	Enables <a href="#">uct_ep_params::iface_addr</a>
UCT_EP_PARAM_FIELD_SOCKADDR	Enables <a href="#">uct_ep_params::sockaddr</a>
UCT_EP_PARAM_FIELD_SOCKADDR_CB_FLAGS	Enables <a href="#">uct_ep_params::sockaddr_cb_flags</a>
UCT_EP_PARAM_FIELD_SOCKADDR_PACK_CB	Enables <a href="#">uct_ep_params::sockaddr_pack_cb</a>
UCT_EP_PARAM_FIELD_CM	Enables <a href="#">uct_ep_params::cm</a>
UCT_EP_PARAM_FIELD_CONN_REQUEST	Enables <a href="#">uct_ep_params::conn_request</a>
UCT_EP_PARAM_FIELD_SOCKADDR_CONNECT_CB_CLIENT	Enables <a href="#">uct_ep_params::sockaddr_cb_client</a>
UCT_EP_PARAM_FIELD_SOCKADDR_NOTIFY_CB_SERVER	Enables <a href="#">uct_ep_params::sockaddr_cb_server</a>
UCT_EP_PARAM_FIELD_SOCKADDR_DISCONNECT_CB	Enables <a href="#">uct_ep_params::disconnect_cb</a>
UCT_EP_PARAM_FIELD_PATH_INDEX	Enables <a href="#">uct_ep_params::path_index</a>
UCT_EP_PARAM_FIELD_CM_RESOLVE_CB	Enables <a href="#">uct_ep_params::cm_resolve_cb</a>
UCT_EP_PARAM_FIELD_PRIV_DATA	Enables <a href="#">uct_ep_params::private_data</a>
UCT_EP_PARAM_FIELD_PRIV_DATA_LENGTH	Enables <a href="#">uct_ep_params::private_data_length</a>
UCT_EP_PARAM_FIELD_LOCAL_SOCKADDR	Enables <a href="#">uct_ep_params::local_sockaddr</a>

6.11.4.11 `uct_iface_feature`

enum [uct\\_iface\\_feature](#)

The enumeration list describes the features supported by UCT. An application can request the features using [UCT parameters](#) during [UCT iface initialization](#) process.

## Enumerator

UCT_IFACE_FEATURE_AM	Request Active Message support
UCT_IFACE_FEATURE_PUT	Request PUT support
UCT_IFACE_FEATURE_GET	Request GET support
UCT_IFACE_FEATURE_AMO32	Request 32-bit atomic operations support
UCT_IFACE_FEATURE_AMO64	Request 64-bit atomic operations support
UCT_IFACE_FEATURE_TAG	Request tag matching offload support
UCT_IFACE_FEATURE_LAST	Used to determine the number of features

## 6.11.4.12 anonymous enum

anonymous enum

## Enumerator

UCT_TAG_RECV_CB_INLINE_DATA	
-----------------------------	--

## 6.11.4.13 uct\_cb\_param\_flags

enum `uct_cb_param_flags`

If UCT\_CB\_PARAM\_FLAG\_DESC flag is enabled, then data is part of a descriptor which includes the user-defined `rx_headroom`, and the callback may return UCS\_INPROGRESS and hold on to that descriptor. Otherwise, the data can't be used outside the callback. If needed, the data must be copied-out.

```

descriptor      data
|               |
+-----+-----+
| rx_headroom | payload                |
+-----+-----+

```

UCT\_CB\_PARAM\_FLAG\_FIRST and UCT\_CB\_PARAM\_FLAG\_MORE flags are relevant for `uct_tag_unexp_eager_cb_t` callback only. The former value indicates that the data is the first fragment of the message. The latter value means that more fragments of the message yet to be delivered.

## Enumerator

UCT_CB_PARAM_FLAG_DESC	
UCT_CB_PARAM_FLAG_FIRST	
UCT_CB_PARAM_FLAG_MORE	

## 6.11.5 Function Documentation

### 6.11.5.1 `uct_query_components()`

```
ucs_status_t uct_query_components (
    uct_component_h ** components_p,
    unsigned * num_components_p )
```

Obtain the list of transport components available on the current system.

#### Parameters

out	<i>components_p</i>	Filled with a pointer to an array of component handles.
out	<i>num_components_p</i>	Filled with the number of elements in the array.

#### Returns

UCS\_OK if successful, or UCS\_ERR\_NO\_MEMORY if failed to allocate the array of component handles.

#### Examples

[uct\\_hello\\_world.c](#).

### 6.11.5.2 `uct_release_component_list()`

```
void uct_release_component_list (
    uct_component_h * components )
```

This routine releases the memory associated with the list of components allocated by [uct\\_query\\_components](#).

#### Parameters

in	<i>components</i>	Array of component handles to release.
----	-------------------	--

#### Examples

[uct\\_hello\\_world.c](#).

### 6.11.5.3 `uct_component_query()`

```
ucs_status_t uct_component_query (
    uct_component_h component,
    uct_component_attr_t * component_attr )
```

Query various attributes of a component.

#### Parameters

in	<i>component</i>	Component handle to query attributes for. The handle can be obtained from <a href="#">uct_query_components</a> .
in, out	<i>component_attr</i>	Filled with component attributes.



**Returns**

UCS\_OK if successful, or nonzero error code in case of failure.

**Examples**

[uct\\_hello\\_world.c](#).

**6.11.5.4 uct\_md\_open()**

```
ucs_status_t uct_md_open (
    uct_component_h component,
    const char * md_name,
    const uct_md_config_t * config,
    uct_md_h * md_p )
```

Open a specific memory domain. All communications and memory operations are performed in the context of a specific memory domain. Therefore it must be created before communication resources.

**Parameters**

in	<i>component</i>	Component on which to open the memory domain, as returned from <a href="#">uct_query_components</a> .
in	<i>md_name</i>	Memory domain name, as returned from <a href="#">uct_component_query</a> .
in	<i>config</i>	MD configuration options. Should be obtained from <a href="#">uct_md_config_read()</a> function, or point to MD-specific structure which extends <code>uct_md_config_t</code> .
out	<i>md_p</i>	Filled with a handle to the memory domain.

**Returns**

Error code.

**Examples**

[uct\\_hello\\_world.c](#).

**6.11.5.5 uct\_md\_close()**

```
void uct_md_close (
    uct_md_h md )
```

**Parameters**

in	<i>md</i>	Memory domain to close.
----	-----------	-------------------------

**Examples**

[uct\\_hello\\_world.c](#).

### 6.11.5.6 uct\_md\_query\_tl\_resources()

```
ucs_status_t uct_md_query_tl_resources (
    uct_md_h md,
    uct_tl_resource_desc_t ** resources_p,
    unsigned * num_resources_p )
```

This routine queries the [memory domain](#) for communication resources that are available for it.

#### Parameters

in	<i>md</i>	Handle to memory domain.
out	<i>resources_p</i>	Filled with a pointer to an array of resource descriptors.
out	<i>num_resources_p</i>	Filled with the number of resources in the array.

#### Returns

Error code.

#### Examples

[uct\\_hello\\_world.c](#).

### 6.11.5.7 uct\_release\_tl\_resource\_list()

```
void uct_release_tl_resource_list (
    uct_tl_resource_desc_t * resources )
```

This routine releases the memory associated with the list of resources allocated by [uct\\_md\\_query\\_tl\\_resources](#).

#### Parameters

in	<i>resources</i>	Array of resource descriptors to release.
----	------------------	---

#### Examples

[uct\\_hello\\_world.c](#).

### 6.11.5.8 uct\_md\_iface\_config\_read()

```
ucs_status_t uct_md_iface_config_read (
    uct_md_h md,
    const char * tl_name,
    const char * env_prefix,
    const char * filename,
    uct_iface_config_t ** config_p )
```

#### Parameters

in	<i>md</i>	Memory domain on which the transport's interface was registered.
----	-----------	--

## Parameters

in	<i>tl_name</i>	Transport name. If <i>md</i> supports <code>UCT_MD_FLAG_SOCKADDR</code> , the transport name is allowed to be NULL. In this case, the configuration returned from this routine should be passed to <code>uct_iface_open</code> with <code>UCT_IFACE_OPEN_MODE_SOCKADDR_SERVER</code> or <code>UCT_IFACE_OPEN_MODE_SOCKADDR_CLIENT</code> set in <code>uct_iface_params_t::open_mode</code> . In addition, if <i>tl_name</i> is not NULL, the configuration returned from this routine should be passed to <code>uct_iface_open</code> with <code>UCT_IFACE_OPEN_MODE_DEVICE</code> set in <code>uct_iface_params_t::open_mode</code> .
in	<i>env_prefix</i>	If non-NULL, search for environment variables starting with this <code>UCT_&lt;prefix&gt;_</code> . Otherwise, search for environment variables starting with just <code>UCT_</code> .
in	<i>filename</i>	If non-NULL, read configuration from this file. If the file does not exist, it will be ignored.
out	<i>config_p</i>	Filled with a pointer to configuration.

## Returns

Error code.

## Examples

[uct\\_hello\\_world.c](#).

6.11.5.9 `uct_config_release()`

```
void uct_config_release (
    void * config )
```

## Parameters

in	<i>config</i>	Configuration to release.
----	---------------	---------------------------

## Examples

[uct\\_hello\\_world.c](#).

6.11.5.10 `uct_iface_open()`

```
ucs_status_t uct_iface_open (
    uct_md_h md,
    uct_worker_h worker,
    const uct_iface_params_t * params,
    const uct_iface_config_t * config,
    uct_iface_h * iface_p )
```

## Parameters

in	<i>md</i>	Memory domain to create the interface on.
in	<i>worker</i>	Handle to worker which will be used to progress communications on this interface.
in	<i>params</i>	User defined <code>uct_iface_params_t</code> parameters.
in	<i>config</i>	Interface configuration options. Should be obtained from <code>uct_md_iface_config_read()</code> function, or point to transport-specific structure which extends <code>uct_iface_config_t</code> .

## Parameters

out	<i>iface</i> ↔ _p	Filled with a handle to opened communication interface.
-----	----------------------	---

## Returns

Error code.

## Examples

[uct\\_hello\\_world.c](#).

**6.11.5.11 uct\_iface\_close()**

```
void uct_iface_close (
    uct_iface_h iface )
```

## Parameters

in	<i>iface</i>	Interface to close.
----	--------------	---------------------

## Examples

[uct\\_hello\\_world.c](#).

**6.11.5.12 uct\_iface\_query()**

```
ucs_status_t uct_iface_query (
    uct_iface_h iface,
    uct_iface_attr_t * iface_attr )
```

## Parameters

in	<i>iface</i>	Interface to query.
out	<i>iface_attr</i>	Filled with interface attributes.

## Examples

[uct\\_hello\\_world.c](#).

**6.11.5.13 uct\_iface\_get\_device\_address()**

```
ucs_status_t uct_iface_get_device_address (
    uct_iface_h iface,
    uct_device_addr_t * addr )
```

Get underlying device address of the interface. All interfaces using the same device would return the same address.

## Parameters

in	<i>iface</i>	Interface to query.
out	<i>addr</i>	Filled with device address. The size of the buffer provided must be at least <a href="#">uct_iface_attr_t::device_addr_len</a> .

## Examples

[uct\\_hello\\_world.c](#).

**6.11.5.14 uct\_iface\_get\_address()**

```
ucs_status_t uct_iface_get_address (
    uct_iface_h iface,
    uct_iface_addr_t * addr )
```

requires [UCT\\_IFACE\\_FLAG\\_CONNECT\\_TO\\_IFACE](#).

## Parameters

in	<i>iface</i>	Interface to query.
out	<i>addr</i>	Filled with interface address. The size of the buffer provided must be at least <a href="#">uct_iface_attr_t::iface_addr_len</a> .

## Examples

[uct\\_hello\\_world.c](#).

**6.11.5.15 uct\_iface\_is\_reachable()**

```
int uct_iface_is_reachable (
    const uct_iface_h iface,
    const uct_device_addr_t * dev_addr,
    const uct_iface_addr_t * iface_addr )
```

This function checks if a remote address can be reached from a local interface. If the function returns true, it does not necessarily mean a connection and/or data transfer would succeed, since the reachability check is a local operation it does not detect issues such as network mis-configuration or lack of connectivity.

## Parameters

in	<i>iface</i>	Interface to check reachability from.
in	<i>dev_addr</i>	Device address to check reachability to. It is NULL if <code>iface_attr.dev_addr_len == 0</code> , and must be non-NULL otherwise.
in	<i>iface_addr</i>	Interface address to check reachability to. It is NULL if <code>iface_attr.iface_addr_len == 0</code> , and must be non-NULL otherwise.

**Returns**

Nonzero if reachable, 0 if not.

**Examples**

[uct\\_hello\\_world.c](#).

**6.11.5.16 uct\_ep\_check()**

```
ucs_status_t uct_ep_check (
    const uct_ep_h ep,
    unsigned flags,
    uct_completion_t * comp )
```

This function checks if the destination endpoint is alive with respect to the UCT library. If the status of *ep* is known, either **UCS\_OK** or an error is returned immediately. Otherwise, **UCS\_INPROGRESS** is returned, indicating that synchronization on the status is needed. In this case, the status will be propagated by *comp* callback.

**Parameters**

in	<i>ep</i>	Endpoint to check
in	<i>flags</i>	Flags that define level of check (currently unsupported - set to 0).
in	<i>comp</i>	Handler to process status of <i>ep</i>

**Returns**

Error code.

**6.11.5.17 uct\_iface\_event\_fd\_get()**

```
ucs_status_t uct_iface_event_fd_get (
    uct_iface_h iface,
    int * fd_p )
```

Only interfaces that support at least one of the UCT\_IFACE\_FLAG\_EVENT\* flags will implement this function.

**Parameters**

in	<i>iface</i>	Interface to get the notification descriptor.
out	<i>fd_p</i>	Location to write the notification file descriptor.

**Returns**

Error code.

**6.11.5.18 uct\_iface\_event\_arm()**

```
ucs_status_t uct_iface_event_arm (
```

```
uct_iface_h iface,
unsigned events )
```

This routine needs to be called before waiting on each notification on this interface, so will typically be called once the processing of the previous event is over.

#### Parameters

in	<i>iface</i>	Interface to arm.
in	<i>events</i>	Events to wakeup on. See <a href="#">uct_iface_event_types</a>

#### Returns

[UCS\\_OK](#) The operation completed successfully. File descriptor will be signaled by new events.

[UCS\\_ERR\\_BUSY](#) There are unprocessed events which prevent the file descriptor from being armed. The operation is not completed. File descriptor will not be signaled by new events.

[Other](#) different error codes in case of issues.

#### 6.11.5.19 uct\_iface\_mem\_alloc()

```
ucs_status_t uct_iface_mem_alloc (
    uct_iface_h iface,
    size_t length,
    unsigned flags,
    const char * name,
    uct_allocated_memory_t * mem )
```

Allocate a region of memory which can be used for zero-copy data transfer or remote access on a particular transport interface.

#### Parameters

in	<i>iface</i>	Interface to allocate memory on.
in	<i>length</i>	Size of memory region to allocate.
in	<i>flags</i>	Memory allocation flags, see <a href="#">uct_md_mem_flags</a> .
in	<i>name</i>	Allocation name, for debug purposes.
out	<i>mem</i>	Descriptor of allocated memory.

#### Returns

[UCS\\_OK](#) if allocation was successful, error code otherwise.

#### 6.11.5.20 uct\_iface\_mem\_free()

```
void uct_iface_mem_free (
    const uct_allocated_memory_t * mem )
```

#### Parameters

in	<i>mem</i>	Descriptor of memory to release.
----	------------	----------------------------------

### 6.11.5.21 uct\_ep\_create()

```
ucs_status_t uct_ep_create (
    const uct_ep_params_t * params,
    uct_ep_h * ep_p )
```

Create a UCT endpoint in one of the available modes:

1. Unconnected endpoint: If no any address is present in `uct_ep_params`, this creates an unconnected endpoint. To establish a connection to a remote endpoint, `uct_ep_connect_to_ep` will need to be called. Use of this mode requires `uct_ep_params_t::iface` has the `UCT_IFACE_FLAG_CONNECT_TO_EP` capability flag. It may be obtained by `uct_iface_query`.
2. Connect to a remote interface: If `uct_ep_params_t::dev_addr` and `uct_ep_params_t::iface_addr` are set, this will establish an endpoint that is connected to a remote interface. This requires that `uct_ep_params_t::iface` has the `UCT_IFACE_FLAG_CONNECT_TO_IFACE` capability flag. It may be obtained by `uct_iface_query`.
3. Connect to a remote socket address: If `uct_ep_params_t::sockaddr` is set, this will create an endpoint that is connected to a remote socket. This requires that either `uct_ep_params::cm`, or `uct_ep_params::iface` will be set. In the latter case, the interface has to support `UCT_IFACE_FLAG_CONNECT_TO_SOCKETADDR` flag, which can be checked by calling `uct_iface_query`.

#### Parameters

in	<i>params</i>	User defined <code>uct_ep_params_t</code> configuration for the <i>ep_p</i> .
out	<i>ep_p</i>	Filled with handle to the new endpoint.

#### Returns

UCS\_OK The endpoint is created successfully. This does not guarantee that the endpoint has been connected to the destination defined in *params*; in case of failure, the error will be reported to the interface error handler callback provided to `uct_iface_open` via `uct_iface_params_t::err_handler`.

Error code as defined by `ucs_status_t`

#### Examples

`uct_hello_world.c`.

### 6.11.5.22 uct\_ep\_destroy()

```
void uct_ep_destroy (
    uct_ep_h ep )
```

#### Parameters

in	<i>ep</i>	Endpoint to destroy.
----	-----------	----------------------

#### Examples

`uct_hello_world.c`.



**6.11.5.23 uct\_ep\_get\_address()**

```
ucs_status_t uct_ep_get_address (
    uct_ep_h ep,
    uct_ep_addr_t * addr )
```

**Parameters**

in	<i>ep</i>	Endpoint to query.
out	<i>addr</i>	Filled with endpoint address. The size of the buffer provided must be at least <a href="#">uct_iface_attr_t::ep_addr_len</a> .

**Examples**

[uct\\_hello\\_world.c](#).

**6.11.5.24 uct\_ep\_connect\_to\_ep()**

```
ucs_status_t uct_ep_connect_to_ep (
    uct_ep_h ep,
    const uct_device_addr_t * dev_addr,
    const uct_ep_addr_t * ep_addr )
```

requires [UCT\\_IFACE\\_FLAG\\_CONNECT\\_TO\\_EP](#) capability.

**Parameters**

in	<i>ep</i>	Endpoint to connect.
in	<i>dev_addr</i>	Remote device address.
in	<i>ep_addr</i>	Remote endpoint address.

**Examples**

[uct\\_hello\\_world.c](#).

**6.11.5.25 uct\_iface\_flush()**

```
ucs_status_t uct_iface_flush (
    uct_iface_h iface,
    unsigned flags,
    uct_completion_t * comp )
```

Flushes all outstanding communications issued on the interface prior to this call. The operations are completed at the origin or at the target as well. The exact completion semantic depends on *flags* parameter.

**Note**

Currently only one completion type is supported. It guarantees that the data transfer is completed but the target buffer may not be updated yet.

## Parameters

in	<i>iface</i>	Interface to flush communications from.
in	<i>flags</i>	Flags that control completion semantic (currently only <a href="#">UCT_FLUSH_FLAG_LOCAL</a> is supported).
in, out	<i>comp</i>	Completion handle as defined by <a href="#">uct_completion_t</a> . Can be NULL, which means that the call will return the current state of the interface and no completion will be generated in case of outstanding communications. If it is not NULL completion counter is decremented by 1 when the call completes. Completion callback is called when the counter reaches 0.

## Returns

UCS\_OK - No outstanding communications left. UCS\_INPROGRESS - Some communication operations are still in progress. If non-NULL 'comp' is provided, it will be updated upon completion of these operations.

6.11.5.26 `uct_iface_fence()`

```
ucs_status_t uct_iface_fence (
    uct_iface_h iface,
    unsigned flags )
```

## Parameters

in	<i>iface</i>	Interface to issue communications from.
in	<i>flags</i>	Flags that control ordering semantic (currently unsupported - set to 0).

## Returns

UCS\_OK - Ordering is inserted.

6.11.5.27 `uct_ep_pending_add()`

```
ucs_status_t uct_ep_pending_add (
    uct_ep_h ep,
    uct_pending_req_t * req,
    unsigned flags )
```

Add a pending request to the endpoint pending queue. The request will be dispatched when the endpoint could potentially have additional send resources.

## Parameters

in	<i>ep</i>	Endpoint to add the pending request to.
in	<i>req</i>	Pending request, which would be dispatched when more resources become available. The user is expected to initialize the "func" field. After being passed to the function, the request is owned by UCT, until the callback is called and returns UCS_OK.
in	<i>flags</i>	Flags that control pending request processing (see <a href="#">uct_cb_flags</a> )

**Returns**

UCS\_OK - request added to pending queue UCS\_ERR\_BUSY - request was not added to pending queue, because send resources are available now. The user is advised to retry.

**6.11.5.28 uct\_ep\_pending\_purge()**

```
void uct_ep_pending_purge (
    uct_ep_h ep,
    uct_pending_purge_callback_t cb,
    void * arg )
```

Remove pending requests from the given endpoint and pass them to the provided callback function. The callback return value is ignored.

**Parameters**

in	<i>ep</i>	Endpoint to remove pending requests from.
in	<i>cb</i>	Callback to pass the removed requests to.
in	<i>arg</i>	Argument to pass to the <i>cb</i> callback.

**6.11.5.29 uct\_ep\_flush()**

```
ucs_status_t uct_ep_flush (
    uct_ep_h ep,
    unsigned flags,
    uct_completion_t * comp )
```

Flushes all outstanding communications issued on the endpoint prior to this call. The operations are completed at the origin or at the target as well. The exact completion semantic depends on *flags* parameter.

**Parameters**

in	<i>ep</i>	Endpoint to flush communications from.
in	<i>flags</i>	Flags <a href="#">uct_flush_flags</a> that control completion semantic.
in, out	<i>comp</i>	Completion handle as defined by <a href="#">uct_completion_t</a> . Can be NULL, which means that the call will return the current state of the endpoint and no completion will be generated in case of outstanding communications. If it is not NULL completion counter is decremented by 1 when the call completes. Completion callback is called when the counter reaches 0.

**Returns**

UCS\_OK - No outstanding communications left. UCS\_ERR\_NO\_RESOURCE - Flush operation could not be initiated. A subsequent call to [uct\\_ep\\_pending\\_add](#) would add a pending operation, which provides an opportunity to retry the flush. UCS\_INPROGRESS - Some communication operations are still in progress. If non-NULL 'comp' is provided, it will be updated upon completion of these operations.

**6.11.5.30 uct\_ep\_fence()**

```
ucs_status_t uct_ep_fence (
```

```
uct_ep_h ep,
unsigned flags )
```

**Parameters**

in	<i>ep</i>	Endpoint to issue communications from.
in	<i>flags</i>	Flags that control ordering semantic (currently unsupported - set to 0).

**Returns**

UCS\_OK - Ordering is inserted.

**6.11.5.31 uct\_iface\_progress\_enable()**

```
void uct_iface_progress_enable (
    uct_iface_h iface,
    unsigned flags )
```

Notify the transport that it should actively progress communications during [uct\\_worker\\_progress\(\)](#).

When the interface is created, its progress is initially disabled.

**Parameters**

in	<i>iface</i>	The interface to enable progress.
in	<i>flags</i>	The type of progress to enable as defined by <a href="#">uct_progress_types</a>

**Note**

This function is not thread safe with respect to [ucp\\_worker\\_progress\(\)](#), unless the flag [UCT\\_PROGRESS\\_THREAD\\_SAFE](#) is specified.

**Examples**

[uct\\_hello\\_world.c](#).

**6.11.5.32 uct\_iface\_progress\_disable()**

```
void uct_iface_progress_disable (
    uct_iface_h iface,
    unsigned flags )
```

Notify the transport that it should not progress its communications during [uct\\_worker\\_progress\(\)](#). Thus the latency of other transports may be improved.

By default, progress is disabled when the interface is created.

**Parameters**

in	<i>iface</i>	The interface to disable progress.
in	<i>flags</i>	The type of progress to disable as defined by <a href="#">uct_progress_types</a> .

**Note**

This function is not thread safe with respect to `uct_worker_progress()`, unless the flag `UCT_PROGRESS_THREAD_SAFE` is specified.

**6.11.5.33 `uct_iface_progress()`**

```
unsigned uct_iface_progress (
    uct_iface_h iface )
```

**6.11.5.34 `uct_completion_update_status()`**

```
static UCS_F_ALWAYS_INLINE void uct_completion_update_status (
    uct_completion_t * comp,
    ucs_status_t status ) [static]
```

**Parameters**

<i>comp</i>	[in] Completion handle to update.
<i>status</i>	[in] Status to update <i>comp</i> handle.

## 6.12 UCT Communication Context

### Enumerations

- enum `uct_alloc_method_t` {  
`UCT_ALLOC_METHOD_THP`, `UCT_ALLOC_METHOD_MD`, `UCT_ALLOC_METHOD_HEAP`, `UCT_ALLOC_METHOD_MMAP`,  
`UCT_ALLOC_METHOD_HUGE`, `UCT_ALLOC_METHOD_LAST`, `UCT_ALLOC_METHOD_DEFAULT = UCT_ALLOC_METHOD_LAST` }  
*Memory allocation methods.*

### Functions

- `ucs_status_t uct_worker_create` (`ucs_async_context_t *async`, `uct_thread_mode_t thread_mode`, `uct_worker_h *worker_p`)  
*Create a worker object.*
- void `uct_worker_destroy` (`uct_worker_h worker`)  
*Destroy a worker object.*
- void `uct_worker_progress_register_safe` (`uct_worker_h worker`, `ucs_callback_t func`, `void *arg`, `unsigned flags`, `uct_worker_cb_id_t *id_p`)  
*Add a slow path callback function to a worker progress.*
- void `uct_worker_progress_unregister_safe` (`uct_worker_h worker`, `uct_worker_cb_id_t *id_p`)  
*Remove a slow path callback function from worker's progress.*
- `ucs_status_t uct_config_get` (`void *config`, `const char *name`, `char *value`, `size_t max`)  
*Get value by name from interface configuration (`uct_iface_config_t`), memory domain configuration (`uct_md_config_t`) or connection manager configuration (`uct_cm_config_t`).*
- `ucs_status_t uct_config_modify` (`void *config`, `const char *name`, `const char *value`)  
*Modify interface configuration (`uct_iface_config_t`), memory domain configuration (`uct_md_config_t`) or connection manager configuration (`uct_cm_config_t`).*
- unsigned `uct_worker_progress` (`uct_worker_h worker`)  
*Explicit progress for UCT worker.*

### 6.12.1 Detailed Description

UCT context abstracts all the resources required for network communication. It is designed to enable either share or isolate resources for multiple programming models used by an application.

This section provides a detailed description of this concept and routines associated with it.

### 6.12.2 Enumeration Type Documentation

#### 6.12.2.1 `uct_alloc_method_t`

```
enum uct_alloc_method_t
```

#### Enumerator

<code>UCT_ALLOC_METHOD_THP</code>	Allocate from OS using libc allocator with Transparent Huge Pages enabled
<code>UCT_ALLOC_METHOD_MD</code>	Allocate using memory domain
<code>UCT_ALLOC_METHOD_HEAP</code>	Allocate from heap using libc allocator
<code>UCT_ALLOC_METHOD_MMAP</code>	Allocate from OS using <code>mmap()</code> syscall

## Enumerator

UCT_ALLOC_METHOD_HUGE	Allocate huge pages
UCT_ALLOC_METHOD_LAST	
UCT_ALLOC_METHOD_DEFAULT	Use default method

### 6.12.3 Function Documentation

#### 6.12.3.1 uct\_worker\_create()

```
ucs_status_t uct_worker_create (
    ucs_async_context_t * async,
    ucs_thread_mode_t thread_mode,
    uct_worker_h * worker_p )
```

The worker represents a progress engine. Multiple progress engines can be created in an application, for example to be used by multiple threads. Transports can allocate separate communication resources for every worker, so that every worker can be progressed independently of others.

## Parameters

in	<i>async</i>	Context for async event handlers. Must not be NULL.
in	<i>thread_mode</i>	Thread access mode to the worker and all interfaces and endpoints associated with it.
out	<i>worker_p</i>	Filled with a pointer to the worker object.

## Examples

[uct\\_hello\\_world.c](#).

#### 6.12.3.2 uct\_worker\_destroy()

```
void uct_worker_destroy (
    uct_worker_h worker )
```

## Parameters

in	<i>worker</i>	Worker object to destroy.
----	---------------	---------------------------

## Examples

[uct\\_hello\\_world.c](#).

#### 6.12.3.3 uct\_worker\_progress\_register\_safe()

```
void uct_worker_progress_register_safe (
    uct_worker_h worker,
```

```

ucs_callback_t func,
void * arg,
unsigned flags,
uct_worker_cb_id_t * id_p )

```

If `*id_p` is equal to `UCS_CALLBACKQ_ID_NULL`, this function will add a callback which will be invoked every time progress is made on the worker. `*id_p` will be updated with an id which refers to this callback and can be used in [uct\\_worker\\_progress\\_unregister\\_safe](#) to remove it from the progress path.

#### Parameters

in	<i>worker</i>	Handle to the worker whose progress should invoke the callback.
in	<i>func</i>	Pointer to the callback function.
in	<i>arg</i>	Argument for the callback function.
in	<i>flags</i>	Callback flags, see <a href="#">ucs_callbackq_flags</a> .
in, out	<i>id_p</i>	Points to a location to store a callback identifier. If <code>*id_p</code> is equal to <code>UCS_CALLBACKQ_ID_NULL</code> , a callback will be added and <code>*id_p</code> will be replaced with a callback identifier which can be subsequently used to remove the callback. Otherwise, no callback will be added and <code>*id_p</code> will be left unchanged.

#### Note

This function is thread safe.

#### 6.12.3.4 uct\_worker\_progress\_unregister\_safe()

```

void uct_worker_progress_unregister_safe (
    uct_worker_h worker,
    uct_worker_cb_id_t * id_p )

```

If `*id_p` is not equal to `UCS_CALLBACKQ_ID_NULL`, remove a callback which was previously added by [uct\\_worker\\_progress\\_register\\_safe](#). `*id_p` will be reset to `UCS_CALLBACKQ_ID_NULL`.

#### Parameters

in	<i>worker</i>	Handle to the worker whose progress should invoke the callback.
in, out	<i>id_p</i>	Points to a callback identifier which indicates the callback to remove. If <code>*id_p</code> is not equal to <code>UCS_CALLBACKQ_ID_NULL</code> , the callback will be removed and <code>*id_p</code> will be reset to <code>UCS_CALLBACKQ_ID_NULL</code> . If <code>*id_p</code> is equal to <code>UCS_CALLBACKQ_ID_NULL</code> , no operation will be performed and <code>*id_p</code> will be left unchanged.

#### Note

This function is thread safe.

#### 6.12.3.5 uct\_config\_get()

```

ucs_status_t uct_config_get (
    void * config,
    const char * name,

```



```
char * value,
size_t max )
```

**Parameters**

in	<i>config</i>	Configuration to get from.
in	<i>name</i>	Configuration variable name.
out	<i>value</i>	Pointer to get value. Should be allocated/freed by caller.
in	<i>max</i>	Available memory space at <i>value</i> pointer.

**Returns**

UCS\_OK if found, otherwise UCS\_ERR\_INVALID\_PARAM or UCS\_ERR\_NO\_ELEM if error.

**6.12.3.6 uct\_config\_modify()**

```
ucs_status_t uct_config_modify (
    void * config,
    const char * name,
    const char * value )
```

**Parameters**

in	<i>config</i>	Configuration to modify.
in	<i>name</i>	Configuration variable name.
in	<i>value</i>	Value to set.

**Returns**

Error code.

**6.12.3.7 uct\_worker\_progress()**

```
unsigned uct_worker_progress (
    uct_worker_h worker )
```

This routine explicitly progresses any outstanding communication operations and active message requests.

**Note**

- In the current implementation, users **MUST** call this routine to receive the active message requests.

**Parameters**

in	<i>worker</i>	Handle to worker.
----	---------------	-------------------

**Returns**

Nonzero if any communication was progressed, zero otherwise.

**Examples**

[uct\\_hello\\_world.c](#).

## 6.13 UCT Memory Domain

### Data Structures

- struct `uct_md_attr`  
*Memory domain attributes. [More...](#)*
- struct `uct_md_attr.cap`
- struct `uct_md_mem_attr`  
*Memory domain attributes. [More...](#)*
- struct `uct_allocated_memory`  
*Describes a memory allocated by UCT. [More...](#)*
- struct `uct_rkey_bundle`  
*Remote key with its type. [More...](#)*
- struct `uct_mem_alloc_params_t`  
*Parameters for allocating memory using `uct_mem_alloc`. [More...](#)*
- struct `uct_mem_alloc_params_t.mds`

### Typedefs

- typedef enum `uct_md_mem_attr_field` `uct_md_mem_attr_field_t`  
*UCT MD memory attributes field mask.*
- typedef struct `uct_md_mem_attr` `uct_md_mem_attr_t`  
*Memory domain attributes.*
- typedef struct `uct_allocated_memory` `uct_allocated_memory_t`  
*Describes a memory allocated by UCT.*
- typedef struct `uct_rkey_bundle` `uct_rkey_bundle_t`  
*Remote key with its type.*

### Enumerations

- enum `uct_sockaddr_accessibility_t` { `UCT_SOCKADDR_ACC_LOCAL`, `UCT_SOCKADDR_ACC_REMOTE` }  
*Socket address accessibility type.*
- enum {  
`UCT_MD_FLAG_ALLOC` = `UCS_BIT(0)`, `UCT_MD_FLAG_REG` = `UCS_BIT(1)`, `UCT_MD_FLAG_NEED_MEMH` = `UCS_BIT(2)`, `UCT_MD_FLAG_NEED_RKEY` = `UCS_BIT(3)`,  
`UCT_MD_FLAG_ADVISE` = `UCS_BIT(4)`, `UCT_MD_FLAG_FIXED` = `UCS_BIT(5)`, `UCT_MD_FLAG_RKEY_PTR` = `UCS_BIT(6)`, `UCT_MD_FLAG_SOCKADDR` = `UCS_BIT(7)`,  
`UCT_MD_FLAG_INVALIDATE` = `UCS_BIT(8)` }  
*Memory domain capability flags.*
- enum `uct_md_mem_flags` {  
`UCT_MD_MEM_FLAG_NONBLOCK` = `UCS_BIT(0)`, `UCT_MD_MEM_FLAG_FIXED` = `UCS_BIT(1)`,  
`UCT_MD_MEM_FLAG_LOCK` = `UCS_BIT(2)`, `UCT_MD_MEM_FLAG_HIDE_ERRORS` = `UCS_BIT(3)`,  
`UCT_MD_MEM_ACCESS_REMOTE_PUT` = `UCS_BIT(5)`, `UCT_MD_MEM_ACCESS_REMOTE_GET` = `UCS_BIT(6)`, `UCT_MD_MEM_ACCESS_REMOTE_ATOMIC` = `UCS_BIT(7)`, `UCT_MD_MEM_ACCESS_LOCAL_READ` = `UCS_BIT(8)`,  
`UCT_MD_MEM_ACCESS_LOCAL_WRITE` = `UCS_BIT(9)`, `UCT_MD_MEM_ACCESS_ALL`, `UCT_MD_MEM_ACCESS_RMA` }  
*Memory allocation/registration flags.*
- enum `uct_mem_advice_t` { `UCT_MADV_NORMAL` = 0, `UCT_MADV_WILLNEED` }  
*list of UCT memory use advice*

- enum `uct_md_mem_attr_field` { `UCT_MD_MEM_ATTR_FIELD_MEM_TYPE` = UCS\_BIT(0), `UCT_MD_MEM_ATTR_FIELD_SY` = UCS\_BIT(1), `UCT_MD_MEM_ATTR_FIELD_BASE_ADDRESS` = UCS\_BIT(2), `UCT_MD_MEM_ATTR_FIELD_ALLOC_LEN` = UCS\_BIT(3) }  
*UCT MD memory attributes field mask.*
- enum `uct_mem_alloc_params_field_t` { `UCT_MEM_ALLOC_PARAM_FIELD_FLAGS` = UCS\_BIT(0), `UCT_MEM_ALLOC_PARAM_FIELD_ADDRESS` = UCS\_BIT(1), `UCT_MEM_ALLOC_PARAM_FIELD_MEM_TYPE` = UCS\_BIT(2), `UCT_MEM_ALLOC_PARAM_FIELD_MDS` = UCS\_BIT(3), `UCT_MEM_ALLOC_PARAM_FIELD_NAME` = UCS\_BIT(4) }  
*UCT allocation parameters specification field mask.*

## Functions

- `ucs_status_t uct_md_mem_query` (`uct_md_h` md, `const void *address`, `size_t length`, `uct_md_mem_attr_t *mem_attr`)  
*Query attributes of a given pointer.*
- `ucs_status_t uct_md_query` (`uct_md_h` md, `uct_md_attr_t *md_attr`)  
*Query for memory domain attributes.*
- `ucs_status_t uct_md_mem_advise` (`uct_md_h` md, `uct_mem_h` memh, `void *addr`, `size_t length`, `uct_mem_advise_t advice`)  
*Give advice about the use of memory.*
- `ucs_status_t uct_md_mem_reg` (`uct_md_h` md, `void *address`, `size_t length`, `unsigned flags`, `uct_mem_h *memh_p`)  
*Register memory for zero-copy sends and remote access.*
- `ucs_status_t uct_md_mem_dereg` (`uct_md_h` md, `uct_mem_h` memh)  
*Undo the operation of `uct_md_mem_reg()`.*
- `ucs_status_t uct_md_detect_memory_type` (`uct_md_h` md, `const void *addr`, `size_t length`, `ucs_memory_type_t *mem_type_p`)  
*Detect memory type.*
- `ucs_status_t uct_mem_alloc` (`size_t length`, `const uct_alloc_method_t *methods`, `unsigned num_methods`, `const uct_mem_alloc_params_t *params`, `uct_allocated_memory_t *mem`)  
*Allocate memory for zero-copy communications and remote access.*
- `ucs_status_t uct_mem_free` (`const uct_allocated_memory_t *mem`)  
*Release allocated memory.*
- `ucs_status_t uct_md_config_read` (`uct_component_h` component, `const char *env_prefix`, `const char *filename`, `uct_md_config_t **config_p`)  
*Read the configuration for a memory domain.*
- `int uct_md_is_sockaddr_accessible` (`uct_md_h` md, `const ucs_sock_addr_t *sockaddr`, `uct_sockaddr_accessibility_t mode`)  
*Check if remote sock address is accessible from the memory domain.*
- `ucs_status_t uct_md_rkey_pack` (`uct_md_h` md, `uct_mem_h` memh, `void *rkey_buffer`)  
*Pack a remote key.*
- `ucs_status_t uct_rkey_unpack` (`uct_component_h` component, `const void *rkey_buffer`, `uct_rkey_bundle_t *rkey_ob`)  
*Unpack a remote key.*
- `ucs_status_t uct_rkey_ptr` (`uct_component_h` component, `uct_rkey_bundle_t *rkey_ob`, `uint64_t remote_addr`, `void **addr_p`)  
*Get a local pointer to remote memory.*
- `ucs_status_t uct_rkey_release` (`uct_component_h` component, `const uct_rkey_bundle_t *rkey_ob`)  
*Release a remote key.*

### 6.13.1 Detailed Description

The Memory Domain abstracts resources required for network communication, which typically includes memory, transport mechanisms, compute and network resources. It is an isolation mechanism that can be employed by the applications for isolating resources between multiple programming models. The attributes of the Memory Domain are defined by the structure `uct_md_attr()`. The communication and memory operations are defined in the context of Memory Domain.

### 6.13.2 Data Structure Documentation

#### 6.13.2.1 struct uct\_md\_attr

This structure defines the attributes of a Memory Domain which includes maximum memory that can be allocated, credentials required for accessing the memory, CPU mask indicating the proximity of CPUs, and bitmaps indicating the types of memory (CPU/CUDA/ROCM) that can be detected, allocated and accessed.

#### Examples

[uct\\_hello\\_world.c](#).

#### Data Fields

struct <code>uct_md_attr</code>	<code>cap</code>	
<code>ucs_linear_func_t</code>	<code>reg_cost</code>	Memory registration cost estimation (time,seconds) as a linear function of the buffer size.
<code>char</code>	<code>component_name[UCT_COMPONENT_NAME_MAX]</code>	Component name
<code>size_t</code>	<code>rkey_packed_size</code>	Size of buffer needed for packed rkey
<code>ucs_cpu_set_t</code>	<code>local_cpus</code>	Mask of CPUs near the resource

#### 6.13.2.2 struct uct\_md\_attr.cap

#### Data Fields

<code>uint64_t</code>	<code>max_alloc</code>	Maximal allocation size
<code>size_t</code>	<code>max_reg</code>	Maximal registration size
<code>uint64_t</code>	<code>flags</code>	<code>UCT_MD_FLAG_xx</code>
<code>uint64_t</code>	<code>reg_mem_types</code>	Bitmap of memory types that Memory Domain can be registered with
<code>uint64_t</code>	<code>detect_mem_types</code>	Bitmap of memory types that Memory Domain can detect if address belongs to it
<code>uint64_t</code>	<code>alloc_mem_types</code>	Bitmap of memory types that Memory Domain can allocate memory on
<code>uint64_t</code>	<code>access_mem_types</code>	Memory types that Memory Domain can access

#### 6.13.2.3 struct uct\_md\_mem\_attr

This structure defines the attributes of a memory pointer which may include the memory type of the pointer, and the system device that backs the pointer depending on the bit fields populated in `field_mask`.

#### Data Fields

<code>uint64_t</code>	<code>field_mask</code>	Mask of valid fields in this structure, using bits from <a href="#">uct_md_mem_attr_field_t</a> .
-----------------------	-------------------------	---

## Data Fields

<a href="#">ucs_memory_type_t</a>	mem_type	The type of memory. E.g. CPU/GPU memory or some other valid type. If the md does not support sys_dev query, then UCS_MEMORY_TYPE_UNKNOWN is returned.
<a href="#">ucs_sys_device_t</a>	sys_dev	Index of the system device on which the buffer resides. eg: NUMA/GPU If the md does not support sys_dev query, then UCS_SYS_DEVICE_ID_UNKNOWN is returned.
void *	base_address	Base address of the allocation to which the provided buffer belongs to. If the md not support base address query, then the pointer passed to uct_md_mem_query is returned as is.
size_t	alloc_length	Length of the whole allocation to which the provided buffer belongs to. If the md not support querying allocation length, then the length passed to uct_md_mem_query is returned as is.

**6.13.2.4 struct uct\_allocated\_memory**

This structure describes the memory block which includes the address, size, and Memory Domain used for allocation. This structure is passed to interface and the memory is allocated by memory allocation functions [uct\\_mem\\_alloc](#).

## Data Fields

void *	address	Address of allocated memory
size_t	length	Real size of allocated memory
<a href="#">uct_alloc_method_t</a>	method	Method used to allocate the memory
<a href="#">ucs_memory_type_t</a>	mem_type	type of allocated memory
<a href="#">uct_md_h</a>	md	if method==MD: MD used to allocate the memory
<a href="#">uct_mem_h</a>	memh	if method==MD: MD memory handle

**6.13.2.5 struct uct\_rkey\_bundle**

This structure describes the credentials (typically key) and information required to access the remote memory by the communication interfaces.

## Data Fields

<a href="#">uct_rkey_t</a>	rkey	Remote key descriptor, passed to RMA functions
void *	handle	Handle, used internally for releasing the key
void *	type	Remote key type

**6.13.2.6 struct uct\_mem\_alloc\_params\_t**

## Data Fields

uint64_t	field_mask	Mask of valid fields in this structure, using bits from <a href="#">uct_mem_alloc_params_field_t</a> . Fields not specified in this mask will be ignored.
----------	------------	---

## Data Fields

unsigned	flags	Memory allocation flags, see <a href="#">uct_md_mem_flags</a> If UCT_MEM_ALLOC_PARAM_FIELD_FLAGS is not specified in <code>field_mask</code> , then (UCT_MD_MEM_ACCESS_LOCAL_READ   UCT_MD_MEM_ACCESS_LOCAL_WRITE) is used by default.
void *	address	If <code>address</code> is NULL, the underlying allocation routine will choose the address at which to create the mapping. If <code>address</code> is non-NULL and UCT_MD_MEM_FLAG_FIXED is not set, the address will be interpreted as a hint as to where to establish the mapping. If <code>address</code> is non-NULL and UCT_MD_MEM_FLAG_FIXED is set, then the specified address is interpreted as a requirement. In this case, if the mapping to the exact address cannot be made, the allocation request fails.
<a href="#">ucs_memory_type_t</a>	mem_type	Type of memory to be allocated.
struct <a href="#">uct_mem_alloc_params_t</a>	mds	
const char *	name	Name of the allocated region, used to track memory usage for debugging and profiling. If UCT_MEM_ALLOC_PARAM_FIELD_NAME is not specified in <code>field_mask</code> , then "anonymous-uct_mem_alloc" is used by default.

6.13.2.7 struct [uct\\_mem\\_alloc\\_params\\_t.mds](#)

## Data Fields

const <a href="#">uct_md_h</a> *	mds	Array of memory domains to attempt to allocate the memory with, for MD allocation method.
unsigned	count	Length of 'mds' array. May be empty, in such case 'mds' may be NULL, and MD allocation method will be skipped.

## 6.13.3 Typedef Documentation

6.13.3.1 [uct\\_md\\_mem\\_attr\\_field\\_t](#)

```
typedef enum uct\_md\_mem\_attr\_field uct\_md\_mem\_attr\_field\_t
```

The enumeration allows specifying which fields in [uct\\_md\\_mem\\_attr\\_t](#) are present.

6.13.3.2 [uct\\_md\\_mem\\_attr\\_t](#)

```
typedef struct uct\_md\_mem\_attr uct\_md\_mem\_attr\_t
```

This structure defines the attributes of a memory pointer which may include the memory type of the pointer, and the system device that backs the pointer depending on the bit fields populated in `field_mask`.

6.13.3.3 [uct\\_allocated\\_memory\\_t](#)

```
typedef struct uct\_allocated\_memory uct\_allocated\_memory\_t
```

This structure describes the memory block which includes the address, size, and Memory Domain used for allocation. This structure is passed to interface and the memory is allocated by memory allocation functions [uct\\_mem\\_alloc](#).

#### 6.13.3.4 uct\_rkey\_bundle\_t

```
typedef struct uct_rkey_bundle uct_rkey_bundle_t
```

This structure describes the credentials (typically key) and information required to access the remote memory by the communication interfaces.

### 6.13.4 Enumeration Type Documentation

#### 6.13.4.1 uct\_sockaddr\_accessibility\_t

```
enum uct_sockaddr_accessibility_t
```

Enumerator

UCT_SOCKADDR_ACC_LOCAL	Check if local address exists. Address should belong to a local network interface
UCT_SOCKADDR_ACC_REMOTE	Check if remote address can be reached. Address is routable from one of the local network interfaces

#### 6.13.4.2 anonymous enum

```
anonymous enum
```

Enumerator

UCT_MD_FLAG_ALLOC	MD supports memory allocation
UCT_MD_FLAG_REG	MD supports memory registration
UCT_MD_FLAG_NEED_MEMH	The transport needs a valid local memory handle for zero-copy operations
UCT_MD_FLAG_NEED_RKEY	The transport needs a valid remote memory key for remote memory operations
UCT_MD_FLAG_ADVISE	MD supports memory advice
UCT_MD_FLAG_FIXED	MD supports memory allocation with fixed address
UCT_MD_FLAG_RKEY_PTR	MD supports direct access to remote memory via a pointer that is returned by <a href="#">uct_rkey_ptr</a>
UCT_MD_FLAG_SOCKADDR	MD support for client-server connection establishment via sockaddr
UCT_MD_FLAG_INVALIDATE	MD supports memory invalidation

#### 6.13.4.3 uct\_md\_mem\_flags

```
enum uct_md_mem_flags
```



## Enumerator

UCT_MD_MEM_FLAG_NONBLOCK	Hint to perform non-blocking allocation/registration: page mapping may be deferred until it is accessed by the CPU or a transport.
UCT_MD_MEM_FLAG_FIXED	Place the mapping at exactly defined address
UCT_MD_MEM_FLAG_LOCK	Registered memory should be locked. May incur extra cost for registration, but memory access is usually faster.
UCT_MD_MEM_FLAG_HIDE_ERRORS	Hide errors on memory registration. In some cases registration failure is not an error (e. g. for merged memory regions).
UCT_MD_MEM_ACCESS_REMOTE_PUT	enable remote put access
UCT_MD_MEM_ACCESS_REMOTE_GET	enable remote get access
UCT_MD_MEM_ACCESS_REMOTE_ATOMIC	enable remote atomic access
UCT_MD_MEM_ACCESS_LOCAL_READ	enable local read access
UCT_MD_MEM_ACCESS_LOCAL_WRITE	enable local write access
UCT_MD_MEM_ACCESS_ALL	enable local and remote access for all operations
UCT_MD_MEM_ACCESS_RMA	enable local and remote access for put and get operations

## 6.13.4.4 uct\_mem\_advice\_t

enum [uct\\_mem\\_advice\\_t](#)

## Enumerator

UCT_MADV_NORMAL	No special treatment
UCT_MADV_WILLNEED	can be used on the memory mapped with <a href="#">UCT_MD_MEM_FLAG_NONBLOCK</a> to speed up memory mapping and to avoid page faults when the memory is accessed for the first time.

## 6.13.4.5 uct\_md\_mem\_attr\_field

enum [uct\\_md\\_mem\\_attr\\_field](#)

The enumeration allows specifying which fields in [uct\\_md\\_mem\\_attr\\_t](#) are present.

## Enumerator

UCT_MD_MEM_ATTR_FIELD_MEM_TYPE	Indicate if memory type is populated. E.g. CPU/GPU
UCT_MD_MEM_ATTR_FIELD_SYS_DEV	Indicate if details of system device backing the pointer are populated. E.g. NUMA/GPU
UCT_MD_MEM_ATTR_FIELD_BASE_ADDRESS	Request base address of the allocation to which the buffer belongs.
UCT_MD_MEM_ATTR_FIELD_ALLOC_LENGTH	Request the whole length of the allocation to which the buffer belongs.

### 6.13.4.6 uct\_mem\_alloc\_params\_field\_t

enum `uct_mem_alloc_params_field_t`

The enumeration allows specifying which fields in `uct_mem_alloc_params_t` are present.

#### Enumerator

UCT_MEM_ALLOC_PARAM_FIELD_FLAGS	Enables <code>uct_mem_alloc_params_t::flags</code>
UCT_MEM_ALLOC_PARAM_FIELD_ADDRESS	Enables <code>uct_mem_alloc_params_t::address</code>
UCT_MEM_ALLOC_PARAM_FIELD_MEM_TYPE	Enables <code>uct_mem_alloc_params_t::mem_type</code>
UCT_MEM_ALLOC_PARAM_FIELD_MDS	Enables <code>uct_mem_alloc_params_t::mds</code>
UCT_MEM_ALLOC_PARAM_FIELD_NAME	Enables <code>uct_mem_alloc_params_t::name</code>

## 6.13.5 Function Documentation

### 6.13.5.1 uct\_md\_mem\_query()

```
ucs_status_t uct_md_mem_query (
    uct_md_h md,
    const void * address,
    size_t length,
    uct_md_mem_attr_t * mem_attr )
```

Return attributes such as memory type, base address, allocation length, and system device for the given pointer of specific length.

#### Parameters

in	<i>md</i>	Memory domain to run the query on. This function returns an error if the md does not recognize the pointer.
in	<i>address</i>	The address of the pointer. Must be non-NULL else UCS_ERR_INVALID_PARAM error is returned.
in	<i>length</i>	Length of the memory region to examine. Must be nonzero else UCS_ERR_INVALID_PARAM error is returned.
in, out	<i>mem_attr</i>	If successful, filled with ptr attributes.

#### Returns

UCS\_OK if at least one attribute is successfully queried otherwise an error code as defined by `ucs_status_t` is returned.

### 6.13.5.2 uct\_md\_query()

```
ucs_status_t uct_md_query (
    uct_md_h md,
    uct_md_attr_t * md_attr )
```

## Parameters

in	<i>md</i>	Memory domain to query.
out	<i>md_attr</i>	Filled with memory domain attributes.

## Examples

[uct\\_hello\\_world.c](#).

6.13.5.3 `uct_md_mem_advise()`

```
ucs_status_t uct_md_mem_advise (
    uct_md_h md,
    uct_mem_h memh,
    void * addr,
    size_t length,
    uct_mem_advice_t advice )
```

This routine advises the UCT about how to handle memory range beginning at address and size of length bytes. This call does not influence the semantics of the application, but may influence its performance. The advice may be ignored.

## Parameters

in	<i>md</i>	Memory domain memory was allocated or registered on.
in	<i>memh</i>	Memory handle, as returned from <a href="#">uct_mem_alloc</a>
in	<i>addr</i>	Memory base address. Memory range must belong to the <i>memh</i>
in	<i>length</i>	Length of memory to advise. Must be >0.
in	<i>advice</i>	Memory use advice as defined in the <a href="#">uct_mem_advice_t</a> list

6.13.5.4 `uct_md_mem_reg()`

```
ucs_status_t uct_md_mem_reg (
    uct_md_h md,
    void * address,
    size_t length,
    unsigned flags,
    uct_mem_h * memh_p )
```

Register memory on the memory domain. In order to use this function, MD must support [UCT\\_MD\\_FLAG\\_REG](#) flag.

## Parameters

in	<i>md</i>	Memory domain to register memory on.
out	<i>address</i>	Memory to register.
in	<i>length</i>	Size of memory to register. Must be >0.
in	<i>flags</i>	Memory allocation flags, see <a href="#">uct_md_mem_flags</a> .
out	<i>memh</i> ↔ <i>_p</i>	Filled with handle for allocated region.

## Examples

[uct\\_hello\\_world.c](#).

## 6.13.5.5 uct\_md\_mem\_dereg()

```
ucs_status_t uct_md_mem_dereg (
    uct_md_h md,
    uct_mem_h memh )
```

## Parameters

in	<i>md</i>	Memory domain which was used to register the memory.
in	<i>memh</i>	Local access key to memory region.

## Examples

[uct\\_hello\\_world.c](#).

## 6.13.5.6 uct\_md\_detect\_memory\_type()

```
ucs_status_t uct_md_detect_memory_type (
    uct_md_h md,
    const void * addr,
    size_t length,
    ucs_memory_type_t * mem_type_p )
```

## Parameters

in	<i>md</i>	Memory domain to detect memory type
in	<i>addr</i>	Memory address to detect.
in	<i>length</i>	Size of memory
out	<i>mem_type</i> <sub>↔</sub> <i>_p</i>	Filled with memory type of the address range if function succeeds

## Returns

UCS\_OK If memory type is successfully detected UCS\_ERR\_INVALID\_ADDR If failed to detect memory type

## 6.13.5.7 uct\_mem\_alloc()

```
ucs_status_t uct_mem_alloc (
    size_t length,
    const uct_alloc_method_t * methods,
    unsigned num_methods,
    const uct_mem_alloc_params_t * params,
    uct_allocated_memory_t * mem )
```

Allocate potentially registered memory.

## Parameters

in	<i>length</i>	The minimal size to allocate. The actual size may be larger, for example because of alignment restrictions. Must be >0.
in	<i>methods</i>	Array of memory allocation methods to attempt. Each of the provided allocation methods will be tried in array order, to perform the allocation, until one succeeds. Whenever the MD method is encountered, each of the provided MDs will be tried in array order, to allocate the memory, until one succeeds, or they are exhausted. In this case the next allocation method from the initial list will be attempted.
in	<i>num_methods</i>	Length of 'methods' array.
in	<i>params</i>	Memory allocation characteristics, see <a href="#">uct_mem_alloc_params_t</a> .
out	<i>mem</i>	In case of success, filled with information about the allocated memory. <a href="#">uct_allocated_memory_t</a>

**6.13.5.8 uct\_mem\_free()**

```
ucs_status_t uct_mem_free (
    const uct_allocated_memory_t * mem )
```

Release the memory allocated by [uct\\_mem\\_alloc](#).

## Parameters

in	<i>mem</i>	Description of allocated memory, as returned from <a href="#">uct_mem_alloc</a> .
----	------------	---

**6.13.5.9 uct\_md\_config\_read()**

```
ucs_status_t uct_md_config_read (
    uct_component_h component,
    const char * env_prefix,
    const char * filename,
    uct_md_config_t ** config_p )
```

## Parameters

in	<i>component</i>	Read the configuration of this component.
in	<i>env_prefix</i>	If non-NULL, search for environment variables starting with this UCT_<prefix>_. Otherwise, search for environment variables starting with just UCT_.
in	<i>filename</i>	If non-NULL, read configuration from this file. If the file does not exist, it will be ignored.
out	<i>config_p</i>	Filled with a pointer to the configuration.

## Returns

Error code.

## Examples

[uct\\_hello\\_world.c](#).

**6.13.5.10 uct\_md\_is\_sockaddr\_accessible()**

```
int uct_md_is_sockaddr_accessible (
    uct_md_h md,
    const ucs_sock_addr_t * sockaddr,
    uct_sockaddr_accessibility_t mode )
```

This function checks if a remote sock address can be accessed from a local memory domain. Accessibility can be checked in local or remote mode.

**Parameters**

in	<i>md</i>	Memory domain to check accessibility from. This memory domain must support the <a href="#">UCT_MD_FLAG_SOCKADDR</a> flag.
in	<i>sockaddr</i>	Socket address to check accessibility to.
in	<i>mode</i>	Mode for checking accessibility, as defined in <a href="#">uct_sockaddr_accessibility_t</a> . Indicates if accessibility is tested on the server side - for binding to the given sockaddr, or on the client side - for connecting to the given remote peer's sockaddr.

**Returns**

Nonzero if accessible, 0 if inaccessible.

**6.13.5.11 uct\_md\_mkey\_pack()**

```
ucs_status_t uct_md_mkey_pack (
    uct_md_h md,
    uct_mem_h memh,
    void * rkey_buffer )
```

**Parameters**

in	<i>md</i>	Handle to memory domain.
in	<i>memh</i>	Local key, whose remote key should be packed.
out	<i>rkey_buffer</i>	Filled with packed remote key.

**Returns**

Error code.

**6.13.5.12 uct\_rkey\_unpack()**

```
ucs_status_t uct_rkey_unpack (
    uct_component_h component,
    const void * rkey_buffer,
    uct_rkey_bundle_t * rkey_ob )
```

**Parameters**

in	<i>component</i>	Component on which to unpack the remote key.
in	<i>rkey_buffer</i>	Packed remote key buffer.
out	<i>rkey_ob</i>	Filled with the unpacked remote key and its type.

**Note**

The remote key must be unpacked with the same component that was used to pack it. For example, if a remote device address on the remote memory domain which was used to pack the key is reachable by a transport on a local component, then that component is eligible to unpack the key. If the remote key buffer cannot be unpacked with the given component, UCS\_ERR\_INVALID\_PARAM will be returned.

**Returns**

Error code.

**6.13.5.13 uct\_rkey\_ptr()**

```
ucs_status_t uct_rkey_ptr (
    uct_component_h component,
    uct_rkey_bundle_t * rkey_ob,
    uint64_t remote_addr,
    void ** addr_p )
```

This routine returns a local pointer to the remote memory described by the rkey bundle. The MD must support [UCT\\_MD\\_FLAG\\_RKEY\\_PTR](#) flag.

**Parameters**

in	<i>component</i>	Component on which to obtain the pointer to the remote key.
in	<i>rkey_ob</i>	A remote key bundle as returned by the <a href="#">uct_rkey_unpack</a> function.
in	<i>remote_addr</i>	A remote address within the memory area described by the rkey_ob.
out	<i>addr_p</i>	A pointer that can be used for direct access to the remote memory.

**Note**

The component used to obtain a local pointer to the remote memory must be the same component that was used to pack the remote key. See notes section for [uct\\_rkey\\_unpack](#).

**Returns**

Error code if the remote memory cannot be accessed directly or the remote address is not valid.

**6.13.5.14 uct\_rkey\_release()**

```
ucs_status_t uct_rkey_release (
    uct_component_h component,
    const uct_rkey_bundle_t * rkey_ob )
```

**Parameters**

in	<i>component</i>	Component which was used to unpack the remote key.
in	<i>rkey_ob</i>	Remote key to release.



## 6.14 UCT Active messages

### Typedefs

- typedef `ucs_status_t(* uct_am_callback_t)` (void \*arg, void \*data, size\_t length, unsigned flags)  
*Callback to process incoming active message.*
- typedef void(\* `uct_am_tracer_t`) (void \*arg, `uct_am_trace_type_t` type, uint8\_t id, const void \*data, size\_t length, char \*buffer, size\_t max)  
*Callback to trace active messages.*

### Enumerations

- enum `uct_msg_flags` { `UCT_SEND_FLAG_SIGNALED` = UCS\_BIT(0), `UCT_SEND_FLAG_PEER_CHECK` = UCS\_BIT(1) }  
*Flags for active message send operation.*
- enum `uct_am_trace_type` { `UCT_AM_TRACE_TYPE_SEND`, `UCT_AM_TRACE_TYPE_RECV`, `UCT_AM_TRACE_TYPE_SEND_DROP`, `UCT_AM_TRACE_TYPE_RECV_DROP`, `UCT_AM_TRACE_TYPE_LAST` }  
*Trace types for active message tracer.*

### Functions

- `ucs_status_t uct_iface_set_am_handler` (`uct_iface_h` iface, uint8\_t id, `uct_am_callback_t` cb, void \*arg, uint32\_t flags)  
*Set active message handler for the interface.*
- `ucs_status_t uct_iface_set_am_tracer` (`uct_iface_h` iface, `uct_am_tracer_t` tracer, void \*arg)  
*Set active message tracer for the interface.*
- void `uct_iface_release_desc` (void \*desc)  
*Release AM descriptor.*
- `ucs_status_t uct_ep_am_short` (`uct_ep_h` ep, uint8\_t id, uint64\_t header, const void \*payload, unsigned length)
- `ucs_status_t uct_ep_am_short_iov` (`uct_ep_h` ep, uint8\_t id, const `uct_iov_t` \*iov, size\_t iovcnt)  
*Short io-vector send operation.*
- ssize\_t `uct_ep_am_bcopy` (`uct_ep_h` ep, uint8\_t id, `uct_pack_callback_t` pack\_cb, void \*arg, unsigned flags)
- `ucs_status_t uct_ep_am_zcopy` (`uct_ep_h` ep, uint8\_t id, const void \*header, unsigned header\_length, const `uct_iov_t` \*iov, size\_t iovcnt, unsigned flags, `uct_completion_t` \*comp)  
*Send active message while avoiding local memory copy.*

#### 6.14.1 Detailed Description

Defines active message functions.

#### 6.14.2 Typedef Documentation

### 6.14.2.1 uct\_am\_callback\_t

```
typedef ucs_status_t(* uct_am_callback_t) (void *arg, void *data, size_t length, unsigned flags)
```

When the callback is called, *flags* indicates how *data* should be handled. If *flags* contain `UCT_CB_PARAM_FLAG_DESC` value, it means *data* is part of a descriptor which must be released later by `uct_iface_release_desc` by the user if the callback returns `UCS_INPROGRESS`.

#### Parameters

in	<i>arg</i>	User-defined argument.
in	<i>data</i>	Points to the received data. This may be a part of a descriptor which may be released later.
in	<i>length</i>	Length of data.
in	<i>flags</i>	Mask with <code>uct_cb_param_flags</code>

#### Note

This callback could be set and released by `uct_iface_set_am_handler` function.

#### Return values

<code>UCS_OK</code>	- descriptor was consumed, and can be released by the caller.
<code>UCS_INPROGRESS</code>	- descriptor is owned by the callee, and would be released later. Supported only if <i>flags</i> contain <code>UCT_CB_PARAM_FLAG_DESC</code> value. Otherwise, this is an error.

### 6.14.2.2 uct\_am\_tracer\_t

```
typedef void(* uct_am_tracer_t) (void *arg, uct_am_trace_type_t type, uint8_t id, const void *data, size_t length, char *buffer, size_t max)
```

Writes a string which represents active message contents into 'buffer'.

#### Parameters

in	<i>arg</i>	User-defined argument.
in	<i>type</i>	Message type.
in	<i>id</i>	Active message id.
in	<i>data</i>	Points to the received data.
in	<i>length</i>	Length of data.
out	<i>buffer</i>	Filled with a debug information string.
in	<i>max</i>	Maximal length of the string.

## 6.14.3 Enumeration Type Documentation

### 6.14.3.1 uct\_msg\_flags

```
enum uct_msg_flags
```

## Enumerator

UCT_SEND_FLAG_SINGALED	Trigger <a href="#">UCT_EVENT_RECV_SIG</a> event on remote side. Make best effort attempt to avoid triggering <a href="#">UCT_EVENT_RECV</a> event. Ignored if not supported by interface.
UCT_SEND_FLAG_PEER_CHECK	Forces checking connectivity to a peer. If the connection is not alive, an error callback will be invoked. If the flag is not set, there is no guarantee that a connectivity error could be detected.

6.14.3.2 `uct_am_trace_type`

```
enum uct_am_trace_type
```

## Enumerator

UCT_AM_TRACE_TYPE_SEND	
UCT_AM_TRACE_TYPE_RECV	
UCT_AM_TRACE_TYPE_SEND_DROP	
UCT_AM_TRACE_TYPE_RECV_DROP	
UCT_AM_TRACE_TYPE_LAST	

## 6.14.4 Function Documentation

6.14.4.1 `uct_iface_set_am_handler()`

```
ucs_status_t uct_iface_set_am_handler (
    uct_iface_h iface,
    uint8_t id,
    uct_am_callback_t cb,
    void * arg,
    uint32_t flags )
```

Only one handler can be set of each active message ID, and setting a handler replaces the previous value. If `cb == NULL`, the current handler is removed.

## Parameters

in	<i>iface</i>	Interface to set the active message handler for.
in	<i>id</i>	Active message id. Must be 0..UCT_AM_ID_MAX-1.
in	<i>cb</i>	Active message callback. NULL to clear.
in	<i>arg</i>	Active message argument.
in	<i>flags</i>	Required <a href="#">callback flags</a>

**Returns**

error code if the interface does not support active messages or requested callback flags

**Examples**

[uct\\_hello\\_world.c](#).

**6.14.4.2 uct\_iface\_set\_am\_tracer()**

```
ucs_status_t uct_iface_set_am_tracer (
    uct_iface_h iface,
    uct_am_tracer_t tracer,
    void * arg )
```

Sets a function which dumps active message debug information to a buffer, which is printed every time an active message is sent or received, when data tracing is on. Without the tracer, only transport-level information is printed.

**Parameters**

in	<i>iface</i>	Interface to set the active message tracer for.
in	<i>tracer</i>	Active message tracer. NULL to clear.
in	<i>arg</i>	Tracer custom argument.

**6.14.4.3 uct\_iface\_release\_desc()**

```
void uct_iface_release_desc (
    void * desc )
```

Release active message descriptor *desc*, which was passed to [the active message callback](#), and owned by the callee.

**Parameters**

in	<i>desc</i>	Descriptor to release.
----	-------------	------------------------

**Examples**

[uct\\_hello\\_world.c](#).

**6.14.4.4 uct\_ep\_am\_short()**

```
ucs_status_t uct_ep_am_short (
    uct_ep_h ep,
    uint8_t id,
    uint64_t header,
    const void * payload,
    unsigned length )
```

## Examples

[uct\\_hello\\_world.c](#).

## 6.14.4.5 uct\_ep\_am\_short\_iov()

```
ucs_status_t uct_ep_am_short_iov (
    uct_ep_h ep,
    uint8_t id,
    const uct_iov_t * iov,
    size_t iovcnt )
```

This routine sends a message using [short](#) protocol. The input data in *iov* array of [uct\\_iov\\_t](#) structures is sent to remote side to contiguous buffer keeping the order of the data in the array.

## Parameters

in	<i>ep</i>	Destination endpoint handle.
in	<i>id</i>	Active message id. Must be in range 0..UCT_AM_ID_MAX-1.
in	<i>iov</i>	Points to an array of <a href="#">uct_iov_t</a> structures. The <i>iov</i> pointer must be a valid address of an array of <a href="#">uct_iov_t</a> structures. A particular structure pointer must be a valid address. A NULL terminated array is not required. <i>stride</i> and <i>count</i> fields in <a href="#">uct_iov_t</a> structure are ignored in current implementation. The total size of the data buffers in the array is limited by <a href="#">uct_iface_attr::cap::am::max_short</a> .
in	<i>iovcnt</i>	Size of the <i>iov</i> data <a href="#">uct_iov_t</a> structures array. If <i>iovcnt</i> is zero, the data is considered empty. <i>iovcnt</i> is limited by <a href="#">uct_iface_attr::cap::am::max_iov</a> .

## Returns

UCS\_OK Operation completed successfully.

UCS\_ERR\_NO\_RESOURCE Could not start the operation due to lack of send resources.

otherwise Error code.

## 6.14.4.6 uct\_ep\_am\_bcopy()

```
ssize_t uct_ep_am_bcopy (
    uct_ep_h ep,
    uint8_t id,
    uct_pack_callback_t pack_cb,
    void * arg,
    unsigned flags )
```

## Examples

[uct\\_hello\\_world.c](#).

## 6.14.4.7 uct\_ep\_am\_zcopy()

```
ucs_status_t uct_ep_am_zcopy (
    uct_ep_h ep,
```

```

uint8_t id,
const void * header,
unsigned header_length,
const uct_iov_t * iov,
size_t iovcnt,
unsigned flags,
uct_completion_t * comp )

```

The input data in *iov* array of `uct_iov_t` structures sent to remote side ("gather output"). Buffers in *iov* are processed in array order. This means that the function complete `iov[0]` before proceeding to `iov[1]`, and so on.

#### Parameters

in	<i>ep</i>	Destination endpoint handle.
in	<i>id</i>	Active message id. Must be in range 0..UCT_AM_ID_MAX-1.
in	<i>header</i>	Active message header.
in	<i>header_length</i>	Active message header length in bytes.
in	<i>iov</i>	Points to an array of <code>uct_iov_t</code> structures. The <i>iov</i> pointer must be a valid address of an array of <code>uct_iov_t</code> structures. A particular structure pointer must be a valid address. A NULL terminated array is not required.
in	<i>iovcnt</i>	Size of the <i>iov</i> data <code>uct_iov_t</code> structures array. If <i>iovcnt</i> is zero, the data is considered empty. <i>iovcnt</i> is limited by <code>uct_iface_attr::cap::am::max_iov</code> .
in	<i>flags</i>	Active message flags, see <code>uct_msg_flags</code> .
in	<i>comp</i>	Completion handle as defined by <code>uct_completion_t</code> .

#### Returns

UCS\_OK Operation completed successfully.

UCS\_INPROGRESS Some communication operations are still in progress. If non-NULL *comp* is provided, it will be updated upon completion of these operations.

UCS\_ERR\_NO\_RESOURCE Could not start the operation due to lack of send resources.

#### Note

If the operation returns `UCS_INPROGRESS`, the memory buffers pointed to by *iov* array must not be modified until the operation is completed by *comp*. *header* can be released or changed.

#### Examples

[uct\\_hello\\_world.c](#).

## 6.15 UCT Remote memory access operations

### Functions

- `ucs_status_t uct_ep_put_short` (`uct_ep_h` ep, `const void *buffer`, `unsigned length`, `uint64_t remote_addr`, `uct_rkey_t rkey`)
- `ssize_t uct_ep_put_bcopy` (`uct_ep_h` ep, `uct_pack_callback_t` pack\_cb, `void *arg`, `uint64_t remote_addr`, `uct_rkey_t rkey`)
- `ucs_status_t uct_ep_put_zcopy` (`uct_ep_h` ep, `const uct_iov_t *iov`, `size_t iovcnt`, `uint64_t remote_addr`, `uct_rkey_t rkey`, `uct_completion_t *comp`)

*Write data to remote memory while avoiding local memory copy.*

- `ucs_status_t uct_ep_get_short` (`uct_ep_h` ep, `void *buffer`, `unsigned length`, `uint64_t remote_addr`, `uct_rkey_t rkey`)
- `ucs_status_t uct_ep_get_bcopy` (`uct_ep_h` ep, `uct_unpack_callback_t` unpack\_cb, `void *arg`, `size_t length`, `uint64_t remote_addr`, `uct_rkey_t rkey`, `uct_completion_t *comp`)
- `ucs_status_t uct_ep_get_zcopy` (`uct_ep_h` ep, `const uct_iov_t *iov`, `size_t iovcnt`, `uint64_t remote_addr`, `uct_rkey_t rkey`, `uct_completion_t *comp`)

*Read data from remote memory while avoiding local memory copy.*

### 6.15.1 Detailed Description

Defines remote memory access operations.

### 6.15.2 Function Documentation

#### 6.15.2.1 `uct_ep_put_short()`

```
ucs_status_t uct_ep_put_short (
    uct_ep_h ep,
    const void * buffer,
    unsigned length,
    uint64_t remote_addr,
    uct_rkey_t rkey )
```

#### 6.15.2.2 `uct_ep_put_bcopy()`

```
ssize_t uct_ep_put_bcopy (
    uct_ep_h ep,
    uct_pack_callback_t pack_cb,
    void * arg,
    uint64_t remote_addr,
    uct_rkey_t rkey )
```

#### 6.15.2.3 `uct_ep_put_zcopy()`

```
ucs_status_t uct_ep_put_zcopy (
    uct_ep_h ep,
    const uct_iov_t * iov,
```

```

size_t iovcnt,
uint64_t remote_addr,
uct_rkey_t rkey,
uct_completion_t * comp )

```

The input data in *iov* array of `uct_iov_t` structures sent to remote address ("gather output"). Buffers in *iov* are processed in array order. This means that the function complete `iov[0]` before proceeding to `iov[1]`, and so on.

#### Parameters

in	<i>ep</i>	Destination endpoint handle.
in	<i>iov</i>	Points to an array of <code>uct_iov_t</code> structures. The <i>iov</i> pointer must be a valid address of an array of <code>uct_iov_t</code> structures. A particular structure pointer must be a valid address. A NULL terminated array is not required.
in	<i>iovcnt</i>	Size of the <i>iov</i> data <code>uct_iov_t</code> structures array. If <i>iovcnt</i> is zero, the data is considered empty. <i>iovcnt</i> is limited by <code>uct_iface_attr::cap::put::max_iov</code> .
in	<i>remote_addr</i>	Remote address to place the <i>iov</i> data.
in	<i>rkey</i>	Remote key descriptor provided by <code>uct_rkey_unpack</code>
in	<i>comp</i>	Completion handle as defined by <code>uct_completion_t</code> .

#### Returns

UCS\_INPROGRESS Some communication operations are still in progress. If non-NULL *comp* is provided, it will be updated upon completion of these operations.

#### 6.15.2.4 uct\_ep\_get\_short()

```

ucs_status_t uct_ep_get_short (
    uct_ep_h ep,
    void * buffer,
    unsigned length,
    uint64_t remote_addr,
    uct_rkey_t rkey )

```

#### 6.15.2.5 uct\_ep\_get\_bcopy()

```

ucs_status_t uct_ep_get_bcopy (
    uct_ep_h ep,
    uct_unpack_callback_t unpack_cb,
    void * arg,
    size_t length,
    uint64_t remote_addr,
    uct_rkey_t rkey,
    uct_completion_t * comp )

```

#### 6.15.2.6 uct\_ep\_get\_zcopy()

```

ucs_status_t uct_ep_get_zcopy (
    uct_ep_h ep,
    const uct_iov_t * iov,

```



```

size_t iovcnt,
uint64_t remote_addr,
uct_rkey_t rkey,
uct_completion_t * comp )

```

The output data in *iov* array of [uct\\_iov\\_t](#) structures received from remote address ("scatter input"). Buffers in *iov* are processed in array order. This means that the function complete *iov*[0] before proceeding to *iov*[1], and so on.

#### Parameters

in	<i>ep</i>	Destination endpoint handle.
in	<i>iov</i>	Points to an array of <a href="#">uct_iov_t</a> structures. The <i>iov</i> pointer must be a valid address of an array of <a href="#">uct_iov_t</a> structures. A particular structure pointer must be a valid address. A NULL terminated array is not required.
in	<i>iovcnt</i>	Size of the <i>iov</i> data <a href="#">uct_iov_t</a> structures array. If <i>iovcnt</i> is zero, the data is considered empty. <i>iovcnt</i> is limited by <a href="#">uct_iface_attr::cap::get::max_iov</a> .
in	<i>remote_addr</i>	Remote address of the data placed to the <i>iov</i> .
in	<i>rkey</i>	Remote key descriptor provided by <a href="#">uct_rkey_unpack</a>
in	<i>comp</i>	Completion handle as defined by <a href="#">uct_completion_t</a> .

#### Returns

UCS\_INPROGRESS Some communication operations are still in progress. If non-NULL *comp* is provided, it will be updated upon completion of these operations.

## 6.16 UCT Atomic operations

### Functions

- [ucs\\_status\\_t uct\\_ep\\_atomic\\_cswap64](#) ([uct\\_ep\\_h](#) ep, [uint64\\_t](#) compare, [uint64\\_t](#) swap, [uint64\\_t](#) remote\_addr, [uct\\_rkey\\_t](#) rkey, [uint64\\_t](#) \*result, [uct\\_completion\\_t](#) \*comp)
- [ucs\\_status\\_t uct\\_ep\\_atomic\\_cswap32](#) ([uct\\_ep\\_h](#) ep, [uint32\\_t](#) compare, [uint32\\_t](#) swap, [uint64\\_t](#) remote\_addr, [uct\\_rkey\\_t](#) rkey, [uint32\\_t](#) \*result, [uct\\_completion\\_t](#) \*comp)
- [ucs\\_status\\_t uct\\_ep\\_atomic32\\_post](#) ([uct\\_ep\\_h](#) ep, [uct\\_atomic\\_op\\_t](#) opcode, [uint32\\_t](#) value, [uint64\\_t](#) remote\_addr, [uct\\_rkey\\_t](#) rkey)
- [ucs\\_status\\_t uct\\_ep\\_atomic64\\_post](#) ([uct\\_ep\\_h](#) ep, [uct\\_atomic\\_op\\_t](#) opcode, [uint64\\_t](#) value, [uint64\\_t](#) remote\_addr, [uct\\_rkey\\_t](#) rkey)
- [ucs\\_status\\_t uct\\_ep\\_atomic32\\_fetch](#) ([uct\\_ep\\_h](#) ep, [uct\\_atomic\\_op\\_t](#) opcode, [uint32\\_t](#) value, [uint32\\_t](#) \*result, [uint64\\_t](#) remote\_addr, [uct\\_rkey\\_t](#) rkey, [uct\\_completion\\_t](#) \*comp)
- [ucs\\_status\\_t uct\\_ep\\_atomic64\\_fetch](#) ([uct\\_ep\\_h](#) ep, [uct\\_atomic\\_op\\_t](#) opcode, [uint64\\_t](#) value, [uint64\\_t](#) \*result, [uint64\\_t](#) remote\_addr, [uct\\_rkey\\_t](#) rkey, [uct\\_completion\\_t](#) \*comp)

### 6.16.1 Detailed Description

Defines atomic operations.

### 6.16.2 Function Documentation

#### 6.16.2.1 [uct\\_ep\\_atomic\\_cswap64\(\)](#)

```
ucs_status_t uct_ep_atomic_cswap64 (
    uct_ep_h ep,
    uint64_t compare,
    uint64_t swap,
    uint64_t remote_addr,
    uct_rkey_t rkey,
    uint64_t * result,
    uct_completion_t * comp )
```

#### 6.16.2.2 [uct\\_ep\\_atomic\\_cswap32\(\)](#)

```
ucs_status_t uct_ep_atomic_cswap32 (
    uct_ep_h ep,
    uint32_t compare,
    uint32_t swap,
    uint64_t remote_addr,
    uct_rkey_t rkey,
    uint32_t * result,
    uct_completion_t * comp )
```

#### 6.16.2.3 [uct\\_ep\\_atomic32\\_post\(\)](#)

```
ucs_status_t uct_ep_atomic32_post (
    uct_ep_h ep,
```

```
    uct_atomic_op_t opcode,  
    uint32_t value,  
    uint64_t remote_addr,  
    uct_rkey_t rkey )
```

#### 6.16.2.4 uct\_ep\_atomic64\_post()

```
ucs_status_t uct_ep_atomic64_post (   
    uct_ep_h ep,  
    uct_atomic_op_t opcode,  
    uint64_t value,  
    uint64_t remote_addr,  
    uct_rkey_t rkey )
```

#### 6.16.2.5 uct\_ep\_atomic32\_fetch()

```
ucs_status_t uct_ep_atomic32_fetch (   
    uct_ep_h ep,  
    uct_atomic_op_t opcode,  
    uint32_t value,  
    uint32_t * result,  
    uint64_t remote_addr,  
    uct_rkey_t rkey,  
    uct_completion_t * comp )
```

#### 6.16.2.6 uct\_ep\_atomic64\_fetch()

```
ucs_status_t uct_ep_atomic64_fetch (   
    uct_ep_h ep,  
    uct_atomic_op_t opcode,  
    uint64_t value,  
    uint64_t * result,  
    uint64_t remote_addr,  
    uct_rkey_t rkey,  
    uct_completion_t * comp )
```

## 6.17 UCT Tag matching operations

### Data Structures

- struct `uct_tag_context`

*Posted tag context.*

### Typedefs

- typedef `ucs_status_t(*uct_tag_unexp_eager_cb_t)` (void \*arg, void \*data, size\_t length, unsigned flags, `uct_tag_t` stag, uint64\_t imm, void \*\*context)

*Callback to process unexpected eager tagged message.*

- typedef `ucs_status_t(*uct_tag_unexp_rndv_cb_t)` (void \*arg, unsigned flags, uint64\_t stag, const void \*header, unsigned header\_length, uint64\_t remote\_addr, size\_t length, const void \*rkey\_buf)

*Callback to process unexpected rendezvous tagged message.*

### Functions

- `ucs_status_t uct_ep_tag_eager_short` (`uct_ep_h` ep, `uct_tag_t` tag, const void \*data, size\_t length)  
*Short eager tagged-send operation.*
- `ssize_t uct_ep_tag_eager_bcopy` (`uct_ep_h` ep, `uct_tag_t` tag, uint64\_t imm, `uct_pack_callback_t` pack\_cb, void \*arg, unsigned flags)  
*Bcopy eager tagged-send operation.*
- `ucs_status_t uct_ep_tag_eager_zcopy` (`uct_ep_h` ep, `uct_tag_t` tag, uint64\_t imm, const `uct_iov_t` \*iov, size\_t iovcnt, unsigned flags, `uct_completion_t` \*comp)  
*Zcopy eager tagged-send operation.*
- `ucs_status_ptr_t uct_ep_tag_rndv_zcopy` (`uct_ep_h` ep, `uct_tag_t` tag, const void \*header, unsigned header\_length, const `uct_iov_t` \*iov, size\_t iovcnt, unsigned flags, `uct_completion_t` \*comp)  
*Rendezvous tagged-send operation.*
- `ucs_status_t uct_ep_tag_rndv_cancel` (`uct_ep_h` ep, void \*op)  
*Cancel outstanding rendezvous operation.*
- `ucs_status_t uct_ep_tag_rndv_request` (`uct_ep_h` ep, `uct_tag_t` tag, const void \*header, unsigned header\_length, unsigned flags)  
*Send software rendezvous request.*
- `ucs_status_t uct_iface_tag_rcv_zcopy` (`uct_iface_h` iface, `uct_tag_t` tag, `uct_tag_t` tag\_mask, const `uct_iov_t` \*iov, size\_t iovcnt, `uct_tag_context_t` \*ctx)  
*Post a tag to a transport interface.*
- `ucs_status_t uct_iface_tag_rcv_cancel` (`uct_iface_h` iface, `uct_tag_context_t` \*ctx, int force)  
*Cancel a posted tag.*

#### 6.17.1 Detailed Description

Defines tag matching operations.

#### 6.17.2 Typedef Documentation

### 6.17.2.1 uct\_tag\_unexp\_eager\_cb\_t

```
typedef ucs_status_t(* uct_tag_unexp_eager_cb_t) (void *arg, void *data, size_t length, unsigned
flags, uct_tag_t stag, uint64_t imm, void **context)
```

This callback is invoked when tagged message sent by eager protocol has arrived and no corresponding tag has been posted.

#### Note

The callback is always invoked from the context (thread, process) that called [uct\\_iface\\_progress\(\)](#).

It is allowed to call other communication routines from the callback.

#### Parameters

in	<i>arg</i>	User-defined argument
in	<i>data</i>	Points to the received unexpected data.
in	<i>length</i>	Length of data.
in	<i>flags</i>	Mask with <a href="#">uct_cb_param_flags</a> flags. If it contains <a href="#">UCT_CB_PARAM_FLAG_DESC</a> value, this means <i>data</i> is part of a descriptor which must be released later using <a href="#">uct_iface_release_desc</a> by the user if the callback returns <a href="#">UCS_INPROGRESS</a> .
in	<i>stag</i>	Tag from sender.
in	<i>imm</i>	Immediate data from sender.
in, out	<i>context</i>	Storage for a per-message user-defined context. In this context, the message is defined by the sender side as a single call to <a href="#">uct_ep_tag_eager_short/bcopy/zcopy</a> . On the transport level the message can be fragmented and delivered to the target over multiple fragments. The fragments will preserve the original order of the message. Each fragment will result in invocation of the above callback. The user can use <a href="#">UCT_CB_PARAM_FLAG_FIRST</a> to identify the first fragment, allocate the context object and use the context as a token that is set by the user and passed to subsequent callbacks of the same message. The user is responsible for allocation and release of the context.

#### Note

No need to allocate the context in the case of a single fragment message (i.e. *flags* contains [UCT\\_CB\\_PARAM\\_FLAG\\_FIRST](#), but does not contain [UCT\\_CB\\_PARAM\\_FLAG\\_MORE](#)).

#### Return values

<a href="#">UCS_OK</a>	- data descriptor was consumed, and can be released by the caller.
<a href="#">UCS_INPROGRESS</a>	- data descriptor is owned by the callee, and will be released later.

### 6.17.2.2 uct\_tag\_unexp\_rndv\_cb\_t

```
typedef ucs_status_t(* uct_tag_unexp_rndv_cb_t) (void *arg, unsigned flags, uint64_t stag,
const void *header, unsigned header_length, uint64_t remote_addr, size_t length, const void
*rkey_buf)
```

This callback is invoked when rendezvous send notification has arrived and no corresponding tag has been posted.

**Note**

The callback is always invoked from the context (thread, process) that called `uct_iface_progress()`.  
It is allowed to call other communication routines from the callback.

**Parameters**

in	<i>arg</i>	User-defined argument
in	<i>flags</i>	Mask with <code>uct_cb_param_flags</code>
in	<i>stag</i>	Tag from sender.
in	<i>header</i>	User defined header.
in	<i>header_length</i>	User defined header length in bytes.
in	<i>remote_addr</i>	Sender's buffer virtual address.
in	<i>length</i>	Sender's buffer length.
in	<i>rkey_buf</i>	Sender's buffer packed remote key. It can be passed to <code>uct_rkey_unpack()</code> to create <code>uct_rkey_t</code> .

**Warning**

If the user became the owner of the *desc* (by returning `UCS_INPROGRESS`) the descriptor must be released later by `uct_iface_release_desc` by the user.

**Return values**

<code>UCS_OK</code>	- descriptor was consumed, and can be released by the caller.
<code>UCS_INPROGRESS</code>	- descriptor is owned by the callee, and would be released later.

**6.17.3 Function Documentation****6.17.3.1 `uct_ep_tag_eager_short()`**

```
ucs_status_t uct_ep_tag_eager_short (
    uct_ep_h ep,
    uct_tag_t tag,
    const void * data,
    size_t length )
```

This routine sends a message using `short` eager protocol. Eager protocol means that the whole data is sent to the peer immediately without any preceding notification. The data is provided as buffer and its length, and must not be larger than the corresponding `max_short` value in `uct_iface_attr`. The immediate value delivered to the receiver is implicitly equal to 0. If it's required to pass nonzero imm value, `uct_ep_tag_eager_bcopy` should be used.

**Parameters**

in	<i>ep</i>	Destination endpoint handle.
in	<i>tag</i>	Tag to use for the eager message.
in	<i>data</i>	Data to send.
in	<i>length</i>	Data length.

**Returns**

UCS\_OK - operation completed successfully.

UCS\_ERR\_NO\_RESOURCE - could not start the operation due to lack of send resources.

**6.17.3.2 uct\_ep\_tag\_eager\_bcopy()**

```
ssize_t uct_ep_tag_eager_bcopy (
    uct_ep_h ep,
    uct_tag_t tag,
    uint64_t imm,
    uct_pack_callback_t pack_cb,
    void * arg,
    unsigned flags )
```

This routine sends a message using **bcopy** eager protocol. Eager protocol means that the whole data is sent to the peer immediately without any preceding notification. Custom data callback is used to copy the data to the network buffers.

**Note**

The resulted data length must not be larger than the corresponding *max\_bcopy* value in [uct\\_iface\\_attr](#).

**Parameters**

in	<i>ep</i>	Destination endpoint handle.
in	<i>tag</i>	Tag to use for the eager message.
in	<i>imm</i>	Immediate value which will be available to the receiver.
in	<i>pack_cb</i>	User callback to pack the data.
in	<i>arg</i>	Custom argument to <i>pack_cb</i> .
in	<i>flags</i>	Tag message flags, see <a href="#">uct_msg_flags</a> .

**Returns**

>=0 - The size of the data packed by *pack\_cb*.

otherwise - Error code.

**6.17.3.3 uct\_ep\_tag\_eager\_zcopy()**

```
ucs_status_t uct_ep_tag_eager_zcopy (
    uct_ep_h ep,
    uct_tag_t tag,
    uint64_t imm,
    const uct_iov_t * iov,
    size_t iovcnt,
    unsigned flags,
    uct_completion_t * comp )
```

This routine sends a message using **zcopy** eager protocol. Eager protocol means that the whole data is sent to the peer immediately without any preceding notification. The input data (which has to be previously registered) in *iov* array of [uct\\_iov\\_t](#) structures sent to remote side ("gather output"). Buffers in *iov* are processed in array order, so the function complete *iov*[0] before proceeding to *iov*[1], and so on.

**Note**

The resulted data length must not be larger than the corresponding *max\_zcopy* value in `uct_iface_attr`.

**Parameters**

in	<i>ep</i>	Destination endpoint handle.
in	<i>tag</i>	Tag to use for the eager message.
in	<i>imm</i>	Immediate value which will be available to the receiver.
in	<i>iov</i>	Points to an array of <code>uct_iov_t</code> structures. A particular structure pointer must be a valid address. A NULL terminated array is not required.
in	<i>iovcnt</i>	Size of the <i>iov</i> array. If <i>iovcnt</i> is zero, the data is considered empty. Note that <i>iovcnt</i> is limited by the corresponding <i>max_iov</i> value in <code>uct_iface_attr</code> .
in	<i>flags</i>	Tag message flags, see <code>uct_msg_flags</code> .
in	<i>comp</i>	Completion callback which will be called when the data is reliably received by the peer, and the buffer can be reused or invalidated.

**Returns**

UCS\_OK - operation completed successfully.

UCS\_ERR\_NO\_RESOURCE - could not start the operation due to lack of send resources.

UCS\_INPROGRESS - operation started, and *comp* will be used to notify when it's completed.

**6.17.3.4 uct\_ep\_tag\_rndv\_zcopy()**

```
ucs_status_ptr_t uct_ep_tag_rndv_zcopy (
    uct_ep_h ep,
    uct_tag_t tag,
    const void * header,
    unsigned header_length,
    const uct_iov_t * iov,
    size_t iovcnt,
    unsigned flags,
    uct_completion_t * comp )
```

This routine sends a message using rendezvous protocol. Rendezvous protocol means that only a small notification is sent at first, and the data itself is transferred later (when there is a match) to avoid extra memory copy.

**Note**

The header will be available to the receiver in case of unexpected rendezvous operation only, i.e. the peer has not posted tag for this message yet (by means of `uct_iface_tag_rcv_zcopy`), when it is arrived.

**Parameters**

in	<i>ep</i>	Destination endpoint handle.
in	<i>tag</i>	Tag to use for the eager message.
in	<i>header</i>	User defined header.
in	<i>header_length</i>	User defined header length in bytes. Note that it is limited by the corresponding <i>max_hdr</i> value in <code>uct_iface_attr</code> .
in	<i>iov</i>	Points to an array of <code>uct_iov_t</code> structures. A particular structure pointer must be valid address. A NULL terminated array is not required.



## Parameters

in	<i>iovcnt</i>	Size of the <i>iov</i> array. If <i>iovcnt</i> is zero, the data is considered empty. Note that <i>iovcnt</i> is limited by the corresponding <i>max_iov</i> value in <a href="#">uct_iface_attr</a> .
in	<i>flags</i>	Tag message flags, see <a href="#">uct_msg_flags</a> .
in	<i>comp</i>	Completion callback which will be called when the data is reliably received by the peer, and the buffer can be reused or invalidated.

## Returns

$\geq 0$  - The operation is in progress and the return value is a handle which can be used to cancel the outstanding rendezvous operation.

otherwise - Error code.

6.17.3.5 `uct_ep_tag_rndv_cancel()`

```
ucs_status_t uct_ep_tag_rndv_cancel (
    uct_ep_h ep,
    void * op )
```

This routine signals the underlying transport disregard the outstanding operation without calling completion callback provided in [uct\\_ep\\_tag\\_rndv\\_zcopy](#).

## Note

The operation handle should be valid at the time the routine is invoked. I.e. it should be a handle of the real operation which is not completed yet.

## Parameters

in	<i>ep</i>	Destination endpoint handle.
in	<i>op</i>	Rendezvous operation handle, as returned from <a href="#">uct_ep_tag_rndv_zcopy</a> .

## Returns

UCS\_OK - The operation has been canceled.

6.17.3.6 `uct_ep_tag_rndv_request()`

```
ucs_status_t uct_ep_tag_rndv_request (
    uct_ep_h ep,
    uct_tag_t tag,
    const void * header,
    unsigned header_length,
    unsigned flags )
```

This routine sends a rendezvous request only, which indicates that the data transfer should be completed in software.

## Parameters

in	<i>ep</i>	Destination endpoint handle.
in	<i>tag</i>	Tag to use for matching.
in	<i>header</i>	User defined header
in	<i>header_length</i>	User defined header length in bytes. Note that it is limited by the corresponding <i>max_hdr</i> value in <a href="#">uct_iface_attr</a> .
in	<i>flags</i>	Tag message flags, see <a href="#">uct_msg_flags</a> .

## Returns

UCS\_OK - operation completed successfully.

UCS\_ERR\_NO\_RESOURCE - could not start the operation due to lack of send resources.

## 6.17.3.7 uct\_iface\_tag\_rcv\_zcopy()

```
ucs_status_t uct_iface_tag_rcv_zcopy (
    uct_iface_h iface,
    uct_tag_t tag,
    uct_tag_t tag_mask,
    const uct_iov_t * iov,
    size_t iovcnt,
    uct_tag_context_t * ctx )
```

This routine posts a tag to be matched on a transport interface. When a message with the corresponding tag arrives it is stored in the user buffer (described by *iov* and *iovcnt*) directly. The operation completion is reported using callbacks on the *ctx* structure.

## Parameters

in	<i>iface</i>	Interface to post the tag on.
in	<i>tag</i>	Tag to expect.
in	<i>tag_mask</i>	Mask which specifies what bits of the tag to compare.
in	<i>iov</i>	Points to an array of <a href="#">uct_iov_t</a> structures. The <i>iov</i> pointer must be a valid address of an array of <a href="#">uct_iov_t</a> structures. A particular structure pointer must be a valid address. A NULL terminated array is not required.
in	<i>iovcnt</i>	Size of the <i>iov</i> data <a href="#">uct_iov_t</a> structures array. If <i>iovcnt</i> is zero, the data is considered empty. <i>iovcnt</i> is limited by <a href="#">uct_iface_attr::cap::tag::max_iov</a> .
in, out	<i>ctx</i>	Context associated with this particular tag, "priv" field in this structure is used to track the state internally.

## Returns

UCS\_OK - The tag is posted to the transport.

UCS\_ERR\_NO\_RESOURCE - Could not start the operation due to lack of resources.

UCS\_ERR\_EXCEEDS\_LIMIT - No more room for tags in the transport.

## 6.17.3.8 uct\_iface\_tag\_rcv\_cancel()

```
ucs_status_t uct_iface_tag_rcv_cancel (
```

```
uct_iface_h iface,  
uct_tag_context_t * ctx,  
int force )
```

This routine cancels a tag, which was previously posted by `uct_iface_tag_rcv_zcopy`. The tag would be either matched or canceled, in a bounded time, regardless of the peer actions. The original completion callback of the tag would be called with the status if `force` is not set.

#### Parameters

in	<i>iface</i>	Interface to cancel the tag on.
in	<i>ctx</i>	Tag context which was used for posting the tag. If <code>force</code> is 0, <code>ctx-&gt;completed_cb</code> will be called with either <code>UCS_OK</code> which means the tag was matched and data received despite the cancel request, or <code>UCS_ERR_CANCELED</code> which means the tag was successfully canceled before it was matched.
in	<i>force</i>	Whether to report completions to <code>ctx-&gt;completed_cb</code> . If nonzero, the cancel is assumed to be successful, and the callback is not called.

#### Returns

`UCS_OK` - The tag is canceled in the transport.

## 6.18 UCT client-server operations

### Data Structures

- struct [uct\\_ep\\_connect\\_params](#)  
*Parameters for connecting a UCT endpoint by [uct\\_ep\\_connect](#). [More...](#)*
- struct [uct\\_cm\\_attr](#)  
*Connection manager attributes, capabilities and limitations. [More...](#)*
- struct [uct\\_listener\\_attr](#)  
*UCT listener attributes, capabilities and limitations. [More...](#)*
- struct [uct\\_listener\\_params](#)  
*Parameters for creating a listener object [uct\\_listener\\_h](#) by [uct\\_listener\\_create](#). [More...](#)*
- struct [uct\\_cm\\_ep\\_priv\\_data\\_pack\\_args](#)  
*Arguments to the client-server private data pack callback. [More...](#)*
- struct [uct\\_cm\\_ep\\_resolve\\_args](#)  
*Arguments to the client-server resolved callback. [More...](#)*
- struct [uct\\_cm\\_remote\\_data](#)  
*Data received from the remote peer. [More...](#)*
- struct [uct\\_cm\\_listener\\_conn\\_request\\_args](#)  
*Arguments to the listener's connection request callback. [More...](#)*
- struct [uct\\_cm\\_ep\\_client\\_connect\\_args](#)  
*Arguments to the client's connect callback. [More...](#)*
- struct [uct\\_cm\\_ep\\_server\\_conn\\_notify\\_args](#)  
*Arguments to the server's notify callback. [More...](#)*

### Typedefs

- typedef struct [uct\\_cm\\_ep\\_priv\\_data\\_pack\\_args](#) [uct\\_cm\\_ep\\_priv\\_data\\_pack\\_args\\_t](#)  
*Arguments to the client-server private data pack callback.*
- typedef struct [uct\\_cm\\_ep\\_resolve\\_args](#) [uct\\_cm\\_ep\\_resolve\\_args\\_t](#)  
*Arguments to the client-server resolved callback.*
- typedef struct [uct\\_cm\\_remote\\_data](#) [uct\\_cm\\_remote\\_data\\_t](#)  
*Data received from the remote peer.*
- typedef struct [uct\\_cm\\_listener\\_conn\\_request\\_args](#) [uct\\_cm\\_listener\\_conn\\_request\\_args\\_t](#)  
*Arguments to the listener's connection request callback.*
- typedef struct [uct\\_cm\\_ep\\_client\\_connect\\_args](#) [uct\\_cm\\_ep\\_client\\_connect\\_args\\_t](#)  
*Arguments to the client's connect callback.*
- typedef struct [uct\\_cm\\_ep\\_server\\_conn\\_notify\\_args](#) [uct\\_cm\\_ep\\_server\\_conn\\_notify\\_args\\_t](#)  
*Arguments to the server's notify callback.*
- typedef void(\* [uct\\_sockaddr\\_conn\\_request\\_callback\\_t](#)) ([uct\\_iface\\_h](#) iface, void \*arg, [uct\\_conn\\_request\\_h](#) conn\_request, const void \*conn\_priv\_data, size\_t length)  
*Callback to process an incoming connection request on the server side.*
- typedef void(\* [uct\\_cm\\_listener\\_conn\\_request\\_callback\\_t](#)) ([uct\\_listener\\_h](#) listener, void \*arg, const [uct\\_cm\\_listener\\_conn\\_request\\_args\\_t](#) \*conn\_req\_args)  
*Callback to process an incoming connection request on the server side listener in a connection manager.*
- typedef void(\* [uct\\_cm\\_ep\\_server\\_conn\\_notify\\_callback\\_t](#)) ([uct\\_ep\\_h](#) ep, void \*arg, const [uct\\_cm\\_ep\\_server\\_conn\\_notify\\_args\\_t](#) \*connect\_args)  
*Callback to process an incoming connection establishment acknowledgment on the server side listener, from the client, which indicates that the client side is connected. The callback also notifies the server side of a local error on a not-yet-connected endpoint.*
- typedef void(\* [uct\\_cm\\_ep\\_client\\_connect\\_callback\\_t](#)) ([uct\\_ep\\_h](#) ep, void \*arg, const [uct\\_cm\\_ep\\_client\\_connect\\_args\\_t](#) \*connect\_args)

Callback to process an incoming connection response on the client side from the server or handle a local error on a not-yet-connected endpoint.

- typedef void(\* [uct\\_ep\\_disconnect\\_cb\\_t](#)) (uct\_ep\_h ep, void \*arg)  
Callback to handle the disconnection of the remote peer.
- typedef ssize\_t(\* [uct\\_cm\\_ep\\_priv\\_data\\_pack\\_callback\\_t](#)) (void \*arg, const [uct\\_cm\\_ep\\_priv\\_data\\_pack\\_args\\_t](#) \*pack\_args, void \*priv\_data)  
Callback to fill the user's private data in a client-server flow.
- typedef [ucs\\_status\\_t](#)(\* [uct\\_cm\\_ep\\_resolve\\_callback\\_t](#)) (void \*user\_data, const [uct\\_cm\\_ep\\_resolve\\_args\\_t](#) \*resolve\_args)  
Callback to notify that the client side endpoint is bound to a local device.

## Enumerations

- enum [uct\\_cm\\_attr\\_field](#) { [UCT\\_CM\\_ATTR\\_FIELD\\_MAX\\_CONN\\_PRIV](#) = [UCS\\_BIT\(0\)](#) }  
UCT connection manager attributes field mask.
- enum [uct\\_listener\\_attr\\_field](#) { [UCT\\_LISTENER\\_ATTR\\_FIELD\\_SOCKADDR](#) = [UCS\\_BIT\(0\)](#) }  
UCT listener attributes field mask.
- enum [uct\\_listener\\_params\\_field](#) { [UCT\\_LISTENER\\_PARAM\\_FIELD\\_BACKLOG](#) = [UCS\\_BIT\(0\)](#), [UCT\\_LISTENER\\_PARAM\\_FIELD\\_DEV\\_ADDR](#) = [UCS\\_BIT\(1\)](#), [UCT\\_LISTENER\\_PARAM\\_FIELD\\_USER\\_DATA](#) = [UCS\\_BIT\(2\)](#) }  
UCT listener created by [uct\\_listener\\_create](#) parameters field mask.
- enum [uct\\_ep\\_connect\\_params\\_field](#) { [UCT\\_EP\\_CONNECT\\_PARAM\\_FIELD\\_PRIVATE\\_DATA](#) = [UCS\\_BIT\(0\)](#), [UCT\\_EP\\_CONNECT\\_PARAM\\_FIELD\\_PRIVATE\\_DATA\\_LENGTH](#) = [UCS\\_BIT\(1\)](#) }  
UCT endpoint connected by [uct\\_ep\\_connect](#) parameters field mask.
- enum [uct\\_cm\\_ep\\_priv\\_data\\_pack\\_args\\_field](#) { [UCT\\_CM\\_EP\\_PRIV\\_DATA\\_PACK\\_ARGS\\_FIELD\\_DEVICE\\_NAME](#) = [UCS\\_BIT\(0\)](#) }  
Client-Server private data pack callback arguments field mask.
- enum [uct\\_cm\\_ep\\_resolve\\_args\\_field](#) { [UCT\\_CM\\_EP\\_RESOLVE\\_ARGS\\_FIELD\\_DEV\\_NAME](#) = [UCS\\_BIT\(0\)](#), [UCT\\_CM\\_EP\\_RESOLVE\\_ARGS\\_FIELD\\_STATUS](#) = [UCS\\_BIT\(1\)](#) }  
Client-Server resolve callback arguments field mask.
- enum [uct\\_cm\\_remote\\_data\\_field](#) { [UCT\\_CM\\_REMOTE\\_DATA\\_FIELD\\_DEV\\_ADDR](#) = [UCS\\_BIT\(0\)](#), [UCT\\_CM\\_REMOTE\\_DATA\\_FIELD\\_DEV\\_ADDR\\_LENGTH](#) = [UCS\\_BIT\(1\)](#), [UCT\\_CM\\_REMOTE\\_DATA\\_FIELD\\_CONN\\_PRIV](#) = [UCS\\_BIT\(2\)](#), [UCT\\_CM\\_REMOTE\\_DATA\\_FIELD\\_CONN\\_PRIV\\_DATA\\_LENGTH](#) = [UCS\\_BIT\(3\)](#) }  
Remote data attributes field mask.
- enum [uct\\_cm\\_listener\\_conn\\_request\\_args\\_field](#) { [UCT\\_CM\\_LISTENER\\_CONN\\_REQUEST\\_ARGS\\_FIELD\\_DEV\\_NAME](#) = [UCS\\_BIT\(0\)](#), [UCT\\_CM\\_LISTENER\\_CONN\\_REQUEST\\_ARGS\\_FIELD\\_CONN\\_REQUEST](#) = [UCS\\_BIT\(1\)](#), [UCT\\_CM\\_LISTENER\\_CONN\\_REQUEST\\_ARGS\\_FIELD\\_REMOTE\\_DATA](#) = [UCS\\_BIT\(2\)](#), [UCT\\_CM\\_LISTENER\\_CONN\\_REQUEST\\_ARGS\\_FIELD\\_CLIENT\\_ADDR](#) = [UCS\\_BIT\(3\)](#) }  
Listener's connection request callback arguments field mask.
- enum [uct\\_cm\\_ep\\_client\\_connect\\_args\\_field](#) { [UCT\\_CM\\_EP\\_CLIENT\\_CONNECT\\_ARGS\\_FIELD\\_REMOTE\\_DATA](#) = [UCS\\_BIT\(0\)](#), [UCT\\_CM\\_EP\\_CLIENT\\_CONNECT\\_ARGS\\_FIELD\\_STATUS](#) = [UCS\\_BIT\(1\)](#) }  
Field mask flags for client-side connection established callback.
- enum [uct\\_cm\\_ep\\_server\\_conn\\_notify\\_args\\_field](#) { [UCT\\_CM\\_EP\\_SERVER\\_CONN\\_NOTIFY\\_ARGS\\_FIELD\\_STATUS](#) = [UCS\\_BIT\(0\)](#) }  
Field mask flags for server-side connection established notification callback.

## Functions

- [ucs\\_status\\_t](#) [uct\\_iface\\_accept](#) ([uct\\_iface\\_h](#) iface, [uct\\_conn\\_request\\_h](#) conn\_request)  
Accept connection request.
- [ucs\\_status\\_t](#) [uct\\_iface\\_reject](#) ([uct\\_iface\\_h](#) iface, [uct\\_conn\\_request\\_h](#) conn\_request)  
Reject connection request. Will invoke an error handler [uct\\_error\\_handler\\_t](#) on the remote transport interface, if set.
- [ucs\\_status\\_t](#) [uct\\_ep\\_connect](#) ([uct\\_ep\\_h](#) ep, const [uct\\_ep\\_connect\\_params\\_t](#) \*params)

Connect a client side endpoint after it is bound to a local network device, i.e. `uct_ep_params_t::cm_resolve_cb` was invoked.

- `ucs_status_t uct_ep_disconnect` (`uct_ep_h` ep, unsigned flags)
 

Initiate a disconnection of an endpoint connected to a `sockaddr` by a connection manager `uct_cm_h`.
- `ucs_status_t uct_cm_open` (`uct_component_h` component, `uct_worker_h` worker, const `uct_cm_config_t *config`, `uct_cm_h *cm_p`)
 

Open a connection manager.
- void `uct_cm_close` (`uct_cm_h` cm)
 

Close a connection manager.
- `ucs_status_t uct_cm_query` (`uct_cm_h` cm, `uct_cm_attr_t *cm_attr`)
 

Get connection manager attributes.
- `ucs_status_t uct_cm_config_read` (`uct_component_h` component, const char `*env_prefix`, const char `*filename`, `uct_cm_config_t **config_p`)
 

Read the configuration for a connection manager.
- `ucs_status_t uct_cm_client_ep_conn_notify` (`uct_ep_h` ep)
 

Notify the server about client-side connection establishment.
- `ucs_status_t uct_listener_create` (`uct_cm_h` cm, const struct `sockaddr *saddr`, `socklen_t socklen`, const `uct_listener_params_t *params`, `uct_listener_h *listener_p`)
 

Create a new transport listener object.
- void `uct_listener_destroy` (`uct_listener_h` listener)
 

Destroy a transport listener.
- `ucs_status_t uct_listener_reject` (`uct_listener_h` listener, `uct_conn_request_h` conn\_request)
 

Reject a connection request.
- `ucs_status_t uct_listener_query` (`uct_listener_h` listener, `uct_listener_attr_t *listener_attr`)
 

Get attributes specific to a particular listener.

### 6.18.1 Detailed Description

Defines client-server operations. The client-server API allows the connection establishment between an active side - a client, and its peer - the passive side - a server. The connection can be established through a UCT transport that supports listening and connecting via IP address and port (listening can also be on `INADDR_ANY`).

The following is a general overview of the operations on the server side:

Connecting: `uct_cm_open` Open a connection manager. `uct_listener_create` Create a listener on the CM and start listening on a given IP,port / `INADDR_ANY`. `uct_cm_listener_conn_request_callback_t` This callback is invoked by the UCT transport to handle an incoming connection request from a client. Accept or reject the client's connection request. `uct_ep_create` Connect to the client by creating an endpoint if the request is accepted. The server creates a new endpoint for every connection request that it accepts. `uct_cm_ep_server_conn_notify_callback_t` This callback is invoked by the UCT transport to handle the connection notification from the client.

#### Note

The private data which the server should send to the client can be either provided directly to `uct_ep_create`, or filled by `uct_cm_ep_priv_data_pack_callback_t` provided to `uct_ep_create`.

In order to reject a connection request, can either call `uct_listener_reject` or return failure status as defined by `ucs_status_t` from `uct_cm_ep_priv_data_pack_callback_t`.

Disconnecting: `uct_ep_disconnect` Disconnect the server's endpoint from the client. Can be called when initiating a disconnect or when receiving a disconnect notification from the remote side. `uct_ep_disconnect_cb_t` This callback is invoked by the UCT transport when the client side calls `uct_ep_disconnect` as well. `uct_ep_destroy` Destroy the endpoint connected to the remote peer. If this function is called before the endpoint was disconnected, the `uct_ep_disconnect_cb_t` will not be invoked.

Destroying the server's resources: `uct_listener_destroy` Destroy the listener object. `uct_cm_close` Close the connection manager.

The following is a general overview of the operations on the client side:

Connecting: [uct\\_cm\\_open](#) Open a connection manager. [uct\\_ep\\_create](#) Create an endpoint for establishing a connection to the server. [uct\\_cm\\_ep\\_resolve\\_callback\\_t](#) This callback is invoked on the client side of the connection manager, after the remote server address was resolved to the local device to be used for connection establishment. [uct\\_ep\\_connect](#) This function should be called on the client side, in order to send private data and resume connection establishment, following an address-resolved notification via [uct\\_cm\\_ep\\_resolve\\_callback\\_t](#). [uct\\_cm\\_ep\\_client\\_connect\\_callback\\_t](#) This callback is invoked by the UCT transport to handle a connection response from the server. After invoking this callback, the UCT transport will finalize the client's connection to the server. [uct\\_cm\\_client\\_ep\\_conn\\_notify](#) After the client's connection establishment is completed, the client should call this function in which it sends a notification message to the server stating that it (the client) is connected. The notification message that is sent depends on the transport's implementation.

Disconnecting: [uct\\_ep\\_disconnect](#) Disconnect the client's endpoint from the server. Can be called when initiating a disconnect or when receiving a disconnect notification from the remote side. [uct\\_ep\\_disconnect\\_cb\\_t](#) This callback is invoked by the UCT transport when the server side calls [uct\\_ep\\_disconnect](#) as well. [uct\\_ep\\_destroy](#) Destroy the endpoint connected to the remote peer.

Destroying the client's resources: [uct\\_cm\\_close](#) Close the connection manager.

## 6.18.2 Data Structure Documentation

### 6.18.2.1 struct uct\_ep\_connect\_params

#### Data Fields

uint64_t	field_mask	Mask of valid fields in this structure, using bits from <a href="#">uct_ep_connect_params_field</a> . Fields not specified by this mask will be ignored.
const void *	private_data	User's private data to be passed from client to server.
size_t	private_data_length	Length of <a href="#">uct_ep_connect_params::private_data</a> , the maximal allowed value is indicated by the <a href="#">uct_cm_attr::max_conn_priv</a> .

### 6.18.2.2 struct uct\_cm\_attr

#### Data Fields

uint64_t	field_mask	Mask of valid fields in this structure, using bits from <a href="#">uct_cm_attr_field</a> . Fields not specified by this mask will be ignored.
size_t	max_conn_priv	Max size of the connection manager's private data used for connection establishment with sockaddr.

### 6.18.2.3 struct uct\_listener\_attr

#### Data Fields

uint64_t	field_mask	Mask of valid fields in this structure, using bits from <a href="#">uct_listener_attr_field</a> . Fields not specified by this mask will be ignored.
struct sockaddr_storage	sockaddr	Socket address on which this listener is listening.

### 6.18.2.4 struct uct\_listener\_params

## Data Fields

uint64_t	field_mask	Mask of valid fields in this structure, using bits from <a href="#">uct_listener_params_field</a> . Fields not specified by this mask will be ignored.
int	backlog	Backlog of incoming connection requests. If specified, must be a positive value. If not specified, each CM component will use its maximal allowed value, based on the system's setting.
<a href="#">uct_cm_listener_conn_request_callback_t</a>	conn_request_cb	Callback function for handling incoming connection requests.
void *	user_data	User data associated with the listener.

**6.18.2.5 struct uct\_cm\_ep\_priv\_data\_pack\_args**

Used with the client-server API on a connection manager.

## Data Fields

uint64_t	field_mask	Mask of valid fields in this structure, using bits from <a href="#">uct_cm_ep_priv_data_pack_args_field</a> . Fields not specified by this mask should not be accessed by the callback.
char	dev_name[UCT_DEVICE_NAME_MAX]	Device name. This routine may fill the user's private data according to the given device name. The device name that is passed to this routine, corresponds to <a href="#">uct_tl_resource_desc_t::dev_name</a> as returned from <a href="#">uct_md_query_tl_resources</a> .

**6.18.2.6 struct uct\_cm\_ep\_resolve\_args**

Used with the client-server API on a connection manager.

## Data Fields

uint64_t	field_mask	Mask of valid fields in this structure, using bits from <a href="#">uct_cm_ep_resolve_args_field</a> . Fields not specified by this mask should not be accessed by the callback.
char	dev_name[UCT_DEVICE_NAME_MAX]	Device name indicates the device that the endpoint was bound to during address and route resolution. The device name that is passed to this callback corresponds to <a href="#">uct_tl_resource_desc_t::dev_name</a> as returned from <a href="#">uct_md_query_tl_resources</a> .
<a href="#">ucs_status_t</a>	status	Indicates address resolution status: UCS_OK - address of the remote server was resolved successfully. UCS_ERR_UNREACHABLE - the remote server is unreachable. Otherwise - indicates an internal connection establishment error on the local (client) side.



### 6.18.2.7 struct uct\_cm\_remote\_data

The remote peer's device address, the data received from it and their lengths. Used with the client-server API on a connection manager.

#### Data Fields

uint64_t	field_mask	Mask of valid fields in this structure, using bits from <a href="#">uct_cm_remote_data_field</a> . Fields not specified by this mask will be ignored.
const uct_device_addr_t *	dev_addr	Device address of the remote peer.
size_t	dev_addr_length	Length of the remote device address.
const void *	conn_priv_data	Pointer to the received data. This is the private data that was passed to <a href="#">uct_ep_params_t::sockaddr_pack_cb</a> .
size_t	conn_priv_data_length	Length of the received data from the peer.

### 6.18.2.8 struct uct\_cm\_listener\_conn\_request\_args

The local device name, connection request handle and the data the client sent. Used with the client-server API on a connection manager.

#### Data Fields

uint64_t	field_mask	Mask of valid fields in this structure, using bits from <a href="#">uct_cm_listener_conn_request_args_field</a> . Fields not specified by this mask should not be accessed by the callback.
char	dev_name[UCT_DEVICE_NAME_MAX]	Local device name which handles the incoming connection request.
<a href="#">uct_conn_request_h</a>	conn_request	Connection request handle. Can be passed to this callback from the transport and will be used by it to accept or reject the connection request from the client.
const <a href="#">uct_cm_remote_data_t</a> *	remote_data	Remote data from the client.
<a href="#">ucs_sock_addr_t</a>	client_address	Client's address.

### 6.18.2.9 struct uct\_cm\_ep\_client\_connect\_args

Used with the client-server API on a connection manager.

#### Data Fields

uint64_t	field_mask	Mask of valid fields in this structure, using bits from <a href="#">uct_cm_ep_client_connect_args_field</a> . Fields not specified by this mask should not be accessed by the callback.
const <a href="#">uct_cm_remote_data_t</a> *	remote_data	Remote data from the server.

## Data Fields

<a href="#">ucs_status_t</a>	status	Indicates the connection establishment response from the remote server: UCS_OK - the remote server accepted the connection request. UCS_ERR_REJECTED - the remote server rejected the connection request. UCS_ERR_CONNECTION_RESET - the server's connection was reset during the connection establishment to the client. Otherwise - indicates an internal connection establishment error on the local (client) side.
------------------------------	--------	---

**6.18.2.10 struct uct\_cm\_ep\_server\_conn\_notify\_args**

Used with the client-server API on a connection manager.

## Data Fields

uint64_t	field_mask	Mask of valid fields in this structure, using bits from <a href="#">uct_cm_ep_server_conn_notify_args_field</a> . Fields not specified by this mask should not be accessed by the callback.
<a href="#">ucs_status_t</a>	status	Indicates the client's <a href="#">ucs_status_t</a> status: UCS_OK - the client completed its connection establishment and called <a href="#">uct_cm_client_ep_conn_notify</a> UCS_ERR_CONNECTION_RESET - the client's connection was reset during the connection establishment to the server. Otherwise - indicates an internal connection establishment error on the local (server) side.

**6.18.3 Typedef Documentation****6.18.3.1 uct\_cm\_ep\_priv\_data\_pack\_args\_t**

```
typedef struct uct_cm_ep_priv_data_pack_args uct_cm_ep_priv_data_pack_args_t
```

Used with the client-server API on a connection manager.

**6.18.3.2 uct\_cm\_ep\_resolve\_args\_t**

```
typedef struct uct_cm_ep_resolve_args uct_cm_ep_resolve_args_t
```

Used with the client-server API on a connection manager.

**6.18.3.3 uct\_cm\_remote\_data\_t**

```
typedef struct uct_cm_remote_data uct_cm_remote_data_t
```

The remote peer's device address, the data received from it and their lengths. Used with the client-server API on a connection manager.

**6.18.3.4 uct\_cm\_listener\_conn\_request\_args\_t**

```
typedef struct uct_cm_listener_conn_request_args uct_cm_listener_conn_request_args_t
```

The local device name, connection request handle and the data the client sent. Used with the client-server API on a connection manager.

### 6.18.3.5 `uct_cm_ep_client_connect_args_t`

```
typedef struct uct_cm_ep_client_connect_args uct_cm_ep_client_connect_args_t
```

Used with the client-server API on a connection manager.

### 6.18.3.6 `uct_cm_ep_server_conn_notify_args_t`

```
typedef struct uct_cm_ep_server_conn_notify_args uct_cm_ep_server_conn_notify_args_t
```

Used with the client-server API on a connection manager.

### 6.18.3.7 `uct_sockaddr_conn_request_callback_t`

```
typedef void(* uct_sockaddr_conn_request_callback_t) (uct_iface_h iface, void *arg, uct_conn_request_h
conn_request, const void *conn_priv_data, size_t length)
```

This callback routine will be invoked on the server side upon receiving an incoming connection request. It should be set by the server side while initializing an interface. Incoming data is placed inside the `conn_priv_data` buffer. This callback has to be thread safe. Other than communication progress routines, it is allowed to call other UCT communication routines from this callback.

#### Parameters

in	<i>iface</i>	Transport interface.
in	<i>arg</i>	User defined argument for this callback.
in	<i>conn_request</i>	Transport level connection request. The user should accept or reject the request by calling <code>uct_iface_accept</code> or <code>uct_iface_reject</code> routines respectively. <code>conn_request</code> should not be used outside the scope of this callback.
in	<i>conn_priv_data</i>	Points to the received data. This is the private data that was passed to the <code>uct_ep_params_t::sockaddr_pack_cb</code> on the client side.
in	<i>length</i>	Length of the received data.

### 6.18.3.8 `uct_cm_listener_conn_request_callback_t`

```
typedef void(* uct_cm_listener_conn_request_callback_t) (uct_listener_h listener, void *arg,
const uct_cm_listener_conn_request_args_t *conn_req_args)
```

This callback routine will be invoked on the server side upon receiving an incoming connection request. It should be set by the server side while initializing a listener in a connection manager. This callback has to be thread safe. Other than communication progress routines, it is allowed to call other UCT communication routines from this callback.

#### Parameters

in	<i>listener</i>	Transport listener.
in	<i>arg</i>	User argument for this callback as defined in <code>uct_listener_params_t::user_data</code>
in	<i>conn_req_args</i>	Listener's arguments to handle the connection request from the client.

### 6.18.3.9 uct\_cm\_ep\_server\_conn\_notify\_callback\_t

```
typedef void(* uct_cm_ep_server_conn_notify_callback_t) (uct_ep_h ep, void *arg, const uct_cm_ep_server_conn_notify_args_t *connect_args)
```

This callback routine will be invoked on the server side upon receiving an incoming connection establishment acknowledgment from the client, which is sent from it once the client is connected to the server. Used to connect the server side to the client or handle an error from it - depending on the status field. This callback will also be invoked in the event of an internal local error with a failed [uct\\_cm\\_ep\\_server\\_conn\\_notify\\_args::status](#) if the endpoint was not connected yet. This callback has to be thread safe. Other than communication progress routines, it is permissible to call other UCT communication routines from this callback.

#### Parameters

in	<i>ep</i>	Transport endpoint.
in	<i>arg</i>	User argument for this callback as defined in <a href="#">uct_ep_params_t::user_data</a>
in	<i>connect_args</i>	Server's connect callback arguments.

### 6.18.3.10 uct\_cm\_ep\_client\_connect\_callback\_t

```
typedef void(* uct_cm_ep_client_connect_callback_t) (uct_ep_h ep, void *arg, const uct_cm_ep_client_connect_callback_args_t *connect_args)
```

This callback routine will be invoked on the client side upon receiving an incoming connection response from the server. Used to connect the client side to the server or handle an error from it - depending on the status field. This callback will also be invoked in the event of an internal local error with a failed [uct\\_cm\\_ep\\_client\\_connect\\_args::status](#) if the endpoint was not connected yet. This callback has to be thread safe. Other than communication progress routines, it is permissible to call other UCT communication routines from this callback.

#### Parameters

in	<i>ep</i>	Transport endpoint.
in	<i>arg</i>	User argument for this callback as defined in <a href="#">uct_ep_params_t::user_data</a> .
in	<i>connect_args</i>	Client's connect callback arguments

### 6.18.3.11 uct\_ep\_disconnect\_cb\_t

```
typedef void(* uct_ep_disconnect_cb_t) (uct_ep_h ep, void *arg)
```

This callback routine will be invoked on the client and server sides upon a disconnect of the remote peer. It will disconnect the given endpoint from the remote peer. This callback won't be invoked if the endpoint was not connected to the remote peer yet. This callback has to be thread safe. Other than communication progress routines, it is permissible to call other UCT communication routines from this callback.

#### Parameters

in	<i>ep</i>	Transport endpoint to disconnect.
in	<i>arg</i>	User argument for this callback as defined in <a href="#">uct_ep_params_t::user_data</a> .

### 6.18.3.12 `uct_cm_ep_priv_data_pack_callback_t`

```
typedef ssize_t(* uct_cm_ep_priv_data_pack_callback_t) (void *arg, const uct_cm_ep_priv_data_pack_args_t
*pack_args, void *priv_data)
```

This callback routine will be invoked on the client side, before sending the transport's connection request to the server, or on the server side before sending a connection response to the client. This callback routine can be set when creating an endpoint. The user's private data should be placed inside the `priv_data` buffer to be sent to the remote side. The maximal allowed length of the private data is indicated by the field `max_conn_priv` inside `uct_iface_attr` or inside `uct_cm_attr` when using a connection manager. Communication progress routines should not be called from this callback. It is allowed to call other UCT communication routines from this callback.

#### Parameters

in	<i>arg</i>	User defined argument for this callback.
in	<i>pack_args</i>	Handle for the the private data packing.
out	<i>priv_data</i>	User's private data to be passed to the remote side.

#### Returns

Negative value indicates an error according to `ucs_status_t`. On success, a non-negative value indicates actual number of bytes written to the `priv_data` buffer.

### 6.18.3.13 `uct_cm_ep_resolve_callback_t`

```
typedef ucs_status_t(* uct_cm_ep_resolve_callback_t) (void *user_data, const uct_cm_ep_resolve_args_t
*resolve_args)
```

This callback routine will be invoked, when the client side endpoint is bound to a local device. The callback routine can be set when creating an endpoint. Communication progress routines should not be called from this callback. It is allowed to call other UCT communication routines from this callback.

#### Parameters

in	<i>user_data</i>	User argument as defined in <code>uct_ep_params_t::user_data</code> .
in	<i>resolve_args</i>	Handle for the extra arguments provided by the transport.

#### Returns

UCS\_OK on success or error as defined in `ucs_status_t`.

## 6.18.4 Enumeration Type Documentation

### 6.18.4.1 `uct_cm_attr_field`

```
enum uct_cm_attr_field
```

The enumeration allows specifying which fields in `uct_cm_attr_t` are present, for backward compatibility support.

#### Enumerator

UCT_CM_ATTR_FIELD_MAX_CONN_PRIV	Enables <code>uct_cm_attr::max_conn_priv</code>
---------------------------------	---

### 6.18.4.2 uct\_listener\_attr\_field

enum `uct_listener_attr_field`

The enumeration allows specifying which fields in `uct_listener_attr_t` are present, for backward compatibility support.

#### Enumerator

UCT_LISTENER_ATTR_FIELD_SOCKADDR	Enables <code>uct_listener_attr::sockaddr</code>
----------------------------------	--

### 6.18.4.3 uct\_listener\_params\_field

enum `uct_listener_params_field`

The enumeration allows specifying which fields in `uct_listener_params_t` are present, for backward compatibility support.

#### Enumerator

UCT_LISTENER_PARAM_FIELD_BACKLOG	Enables <code>uct_listener_params::backlog</code>
UCT_LISTENER_PARAM_FIELD_CONN_REQUEST_CB	Enables <code>uct_listener_params::conn_request_cb</code>
UCT_LISTENER_PARAM_FIELD_USER_DATA	Enables <code>uct_listener_params::user_data</code>

### 6.18.4.4 uct\_ep\_connect\_params\_field

enum `uct_ep_connect_params_field`

The enumeration allows specifying which fields in `uct_ep_connect_params_t` are present, for backward compatibility support.

#### Enumerator

UCT_EP_CONNECT_PARAM_FIELD_PRIVATE_↔ DATA	Enables <code>uct_ep_connect_params::private_data</code>
UCT_EP_CONNECT_PARAM_FIELD_PRIVATE_↔ DATA_LENGTH	Enables <code>uct_ep_connect_params::private_data_length</code>

### 6.18.4.5 uct\_cm\_ep\_priv\_data\_pack\_args\_field

enum `uct_cm_ep_priv_data_pack_args_field`

The enumeration allows specifying which fields in `uct_cm_ep_priv_data_pack_args` are present, for backward compatibility support.

## Enumerator

UCT_CM_EP_PRIV_DATA_PACK_ARGS_FIELD↔ _DEVICE_NAME	Enables <code>uct_cm_ep_priv_data_pack_args::dev_name</code> Indicates that <code>dev_name</code> field in <code>uct_cm_ep_priv_data_pack_args_t</code> is valid.
--	---

6.18.4.6 `uct_cm_ep_resolve_args_field`

```
enum uct_cm_ep_resolve_args_field
```

The enumeration allows specifying which fields in `uct_cm_ep_resolve_args` are present, for backward compatibility support.

## Enumerator

UCT_CM_EP_RESOLVE_ARGS_FIELD_DEV_N↔ AME	Indicates that <code>uct_cm_ep_resolve_args::dev_name</code> is valid.
UCT_CM_EP_RESOLVE_ARGS_FIELD_STATUS	Indicates that <code>uct_cm_ep_resolve_args::status</code> is valid.

6.18.4.7 `uct_cm_remote_data_field`

```
enum uct_cm_remote_data_field
```

The enumeration allows specifying which fields in `uct_cm_remote_data` are present, for backward compatibility support.

## Enumerator

UCT_CM_REMOTE_DATA_FIELD_DEV_ADDR	Enables <code>uct_cm_remote_data::dev_addr</code>
UCT_CM_REMOTE_DATA_FIELD_DEV_ADDR↔ LENGTH	Enables <code>uct_cm_remote_data::dev_addr_length</code>
UCT_CM_REMOTE_DATA_FIELD_CONN_PRIV↔ DATA	Enables <code>uct_cm_remote_data::conn_priv_data</code>
UCT_CM_REMOTE_DATA_FIELD_CONN_PRIV↔ DATA_LENGTH	Enables <code>uct_cm_remote_data::conn_priv_data_length</code>

6.18.4.8 `uct_cm_listener_conn_request_args_field`

```
enum uct_cm_listener_conn_request_args_field
```

The enumeration allows specifying which fields in `uct_cm_listener_conn_request_args` are present, for backward compatibility support.

## Enumerator

## Enumerator

UCT_CM_LISTENER_CONN_REQUEST_ARGS_↔ FIELD_DEV_NAME	Enables <a href="#">uct_cm_listener_conn_request_args::dev_name</a> Indicates that dev_name field in uct_cm_listener_conn_request_args_t is valid.
UCT_CM_LISTENER_CONN_REQUEST_ARGS_↔ FIELD_CONN_REQUEST	Enables <a href="#">uct_cm_listener_conn_request_args::conn_request</a> Indicates that conn_request field in uct_cm_listener_conn_request_args_t is valid.
UCT_CM_LISTENER_CONN_REQUEST_ARGS_↔ FIELD_REMOTE_DATA	Enables <a href="#">uct_cm_listener_conn_request_args::remote_data</a> Indicates that remote_data field in uct_cm_listener_conn_request_args_t is valid.
UCT_CM_LISTENER_CONN_REQUEST_ARGS_↔ FIELD_CLIENT_ADDR	Enables <a href="#">uct_cm_listener_conn_request_args::client_address</a> Indicates that client_address field in uct_cm_listener_conn_request_args_t is valid.

**6.18.4.9 uct\_cm\_ep\_client\_connect\_args\_field**

```
enum uct_cm_ep_client_connect_args_field
```

The enumeration allows specifying which fields in [uct\\_cm\\_ep\\_client\\_connect\\_args](#) are present, for backward compatibility support.

## Enumerator

UCT_CM_EP_CLIENT_CONNECT_ARGS_FIELD↔ _REMOTE_DATA	Enables <a href="#">uct_cm_ep_client_connect_args::remote_data</a>
UCT_CM_EP_CLIENT_CONNECT_ARGS_FIELD↔ _STATUS	Enables <a href="#">uct_cm_ep_client_connect_args::status</a>

**6.18.4.10 uct\_cm\_ep\_server\_conn\_notify\_args\_field**

```
enum uct_cm_ep_server_conn_notify_args_field
```

The enumeration allows specifying which fields in [uct\\_cm\\_ep\\_server\\_conn\\_notify\\_args](#) are present, for backward compatibility support.

## Enumerator

UCT_CM_EP_SERVER_CONN_NOTIFY_ARGS_↔ FIELD_STATUS	Enables <a href="#">uct_cm_ep_server_conn_notify_args::status</a> Indicates that status field in uct_cm_ep_server_conn_notify_args_t is valid.
---	---



## 6.18.5 Function Documentation

### 6.18.5.1 `uct_iface_accept()`

```
ucs_status_t uct_iface_accept (
    uct_iface_h iface,
    uct_conn_request_h conn_request )
```

#### Parameters

in	<i>iface</i>	Transport interface which generated connection request <i>conn_request</i> .
in	<i>conn_request</i>	Connection establishment request passed as parameter of <a href="#">uct_sockaddr_conn_request_callback_t</a> .

#### Returns

Error code as defined by [ucs\\_status\\_t](#)

### 6.18.5.2 `uct_iface_reject()`

```
ucs_status_t uct_iface_reject (
    uct_iface_h iface,
    uct_conn_request_h conn_request )
```

#### Parameters

in	<i>iface</i>	Interface which generated connection establishment request <i>conn_request</i> .
in	<i>conn_request</i>	Connection establishment request passed as parameter of <a href="#">uct_sockaddr_conn_request_callback_t</a> .

#### Returns

Error code as defined by [ucs\\_status\\_t](#)

### 6.18.5.3 `uct_ep_connect()`

```
ucs_status_t uct_ep_connect (
    uct_ep_h ep,
    const uct_ep_connect_params_t * params )
```

This non-blocking routine establishes connection of the client side endpoint and sends private data to the peer.

#### Parameters

in	<i>ep</i>	Endpoint to connect.
in	<i>params</i>	Parameters as defined in <a href="#">uct_ep_connect_params_t</a> .

**Returns**

UCS\_OK Operation has been initiated successfully. Other error codes as defined by [ucs\\_status\\_t](#).

**6.18.5.4 uct\_ep\_disconnect()**

```
ucs_status_t uct_ep_disconnect (
    uct_ep_h ep,
    unsigned flags )
```

This non-blocking routine will send a disconnect notification on the endpoint, so that [uct\\_ep\\_disconnect\\_cb\\_t](#) will be called on the remote peer. The remote side should also call this routine when handling the initiator's disconnect. After a call to this function, the given endpoint may not be used for communications anymore. The [uct\\_ep\\_flush](#) / [uct\\_iface\\_flush](#) routines will guarantee that the disconnect notification is delivered to the remote peer. [uct\\_ep\\_destroy](#) should be called on this endpoint after invoking this routine and [uct\\_ep\\_params::disconnect\\_cb](#) was called.

**Parameters**

in	<i>ep</i>	Endpoint to disconnect.
in	<i>flags</i>	Reserved for future use.

**Returns**

UCS\_OK Operation has completed successfully. UCS\_ERR\_BUSY The *ep* is not connected yet (either [uct\\_cm\\_ep\\_client\\_connect\\_callback\\_t](#) or [uct\\_cm\\_ep\\_server\\_conn\\_notify\\_callback\\_t](#) was not invoked). UCS\_INPROGRESS The disconnect request has been initiated, but the remote peer has not yet responded to this request, and consequently the registered callback [uct\\_ep\\_disconnect\\_cb\\_t](#) has not been invoked to handle the request. UCS\_ERR\_NOT\_CONNECTED The *ep* is disconnected locally and remotely. Other error codes as defined by [ucs\\_status\\_t](#).

**6.18.5.5 uct\_cm\_open()**

```
ucs_status_t uct_cm_open (
    uct_component_h component,
    uct_worker_h worker,
    const uct_cm_config_t * config,
    uct_cm_h * cm_p )
```

Open a connection manager. All client server connection establishment operations are performed in the context of a specific connection manager.

**Note**

This is an alternative API for [uct\\_iface\\_open\\_mode::UCT\\_IFACE\\_OPEN\\_MODE\\_SOCKETADDR\\_SERVER](#) and [uct\\_iface\\_open\\_mode::UCT\\_IFACE\\_OPEN\\_MODE\\_SOCKETADDR\\_CLIENT](#).

**Parameters**

in	<i>component</i>	Component on which to open the connection manager, as returned from <a href="#">uct_query_components</a> .
in	<i>worker</i>	Worker on which to open the connection manager.

## Parameters

in	<i>config</i>	CM configuration options. Either obtained from <code>uct_cm_config_read()</code> function, or pointer to CM-specific structure that extends <code>uct_cm_config_t</code> .
out	<i>cm_p</i>	Filled with a handle to the connection manager.

## Returns

Error code.

6.18.5.6 `uct_cm_close()`

```
void uct_cm_close (
    uct_cm_h cm )
```

## Parameters

in	<i>cm</i>	Connection manager to close.
----	-----------	------------------------------

6.18.5.7 `uct_cm_query()`

```
ucs_status_t uct_cm_query (
    uct_cm_h cm,
    uct_cm_attr_t * cm_attr )
```

This routine queries the `cm` for its attributes `uct_cm_attr_t`.

## Parameters

in	<i>cm</i>	Connection manager to query.
out	<i>cm_attr</i>	Filled with connection manager attributes.

6.18.5.8 `uct_cm_config_read()`

```
ucs_status_t uct_cm_config_read (
    uct_component_h component,
    const char * env_prefix,
    const char * filename,
    uct_cm_config_t ** config_p )
```

## Parameters

in	<i>component</i>	Read the configuration of the connection manager on this component.
in	<i>env_prefix</i>	If non-NULL, search for environment variables starting with this UCT_<prefix>_. Otherwise, search for environment variables starting with just UCT_.
in	<i>filename</i>	If non-NULL, read configuration from this file. If the file does not exist, or exists but cannot be opened or read, it will be ignored.

## Parameters

out	<i>config_p</i>	Filled with a pointer to the configuration.
-----	-----------------	---

## Returns

Error code.

**6.18.5.9 uct\_cm\_client\_ep\_conn\_notify()**

```
ucs_status_t uct_cm_client_ep_conn_notify (
    uct_ep_h ep )
```

This routine should be called on the client side after the client completed establishing its connection to the server. The routine will send a notification message to the server indicating that the client is connected.

## Parameters

in	<i>ep</i>	The connected endpoint on the client side.
----	-----------	--

## Returns

Error code.

**6.18.5.10 uct\_listener\_create()**

```
ucs_status_t uct_listener_create (
    uct_cm_h cm,
    const struct sockaddr * saddr,
    socklen_t socklen,
    const uct_listener_params_t * params,
    uct_listener_h * listener_p )
```

This routine creates a new listener on the given CM which will start listening on a given sockaddr.

## Parameters

in	<i>cm</i>	Connection manager on which to open the listener. This cm should not be closed as long as there are open listeners on it.
in	<i>saddr</i>	The socket address to listen on.
in	<i>socklen</i>	The saddr length.
in	<i>params</i>	User defined <a href="#">uct_listener_params_t</a> configurations for the <i>listener_p</i> .
out	<i>listener_p</i>	Filled with handle to the new listener.

## Returns

Error code.

**6.18.5.11 uct\_listener\_destroy()**

```
void uct_listener_destroy (
    uct_listener_h listener )
```

**Parameters**

in	<i>listener</i>	Listener to destroy.
----	-----------------	----------------------

**6.18.5.12 uct\_listener\_reject()**

```
ucs_status_t uct_listener_reject (
    uct_listener_h listener,
    uct_conn_request_h conn_request )
```

This routine can be invoked on the server side. It rejects a connection request from the client.

**Parameters**

in	<i>listener</i>	Listener which will reject the connection request.
in	<i>conn_request</i>	Connection establishment request passed as parameter of <a href="#">uct_cm_listener_conn_request_callback_t</a> in <a href="#">uct_cm_listener_conn_request_args_t::conn_request</a> .

**Returns**

Error code as defined by [ucs\\_status\\_t](#)

**6.18.5.13 uct\_listener\_query()**

```
ucs_status_t uct_listener_query (
    uct_listener_h listener,
    uct_listener_attr_t * listener_attr )
```

This routine queries the [listener](#) for its attributes [uct\\_listener\\_attr\\_t](#).

**Parameters**

in	<i>listener</i>	Listener object to query.
out	<i>listener_attr</i>	Filled with attributes of the listener.

**Returns**

Error code as defined by [ucs\\_status\\_t](#)

## 6.19 UCT interface operations and capabilities

List of capabilities supported by UCX API.

### Macros

- #define UCT\_IFACE\_FLAG\_AM\_SHORT UCS\_BIT(0)
- #define UCT\_IFACE\_FLAG\_AM\_BCOPY UCS\_BIT(1)
- #define UCT\_IFACE\_FLAG\_AM\_ZCOPY UCS\_BIT(2)
- #define UCT\_IFACE\_FLAG\_PENDING UCS\_BIT(3)
- #define UCT\_IFACE\_FLAG\_PUT\_SHORT UCS\_BIT(4)
- #define UCT\_IFACE\_FLAG\_PUT\_BCOPY UCS\_BIT(5)
- #define UCT\_IFACE\_FLAG\_PUT\_ZCOPY UCS\_BIT(6)
- #define UCT\_IFACE\_FLAG\_GET\_SHORT UCS\_BIT(8)
- #define UCT\_IFACE\_FLAG\_GET\_BCOPY UCS\_BIT(9)
- #define UCT\_IFACE\_FLAG\_GET\_ZCOPY UCS\_BIT(10)
- #define UCT\_IFACE\_FLAG\_ATOMIC\_CPU UCS\_BIT(30)
- #define UCT\_IFACE\_FLAG\_ATOMIC\_DEVICE UCS\_BIT(31)
- #define UCT\_IFACE\_FLAG\_ERRHANDLE\_SHORT\_BUF UCS\_BIT(32)
- #define UCT\_IFACE\_FLAG\_ERRHANDLE\_BCOPY\_BUF UCS\_BIT(33)
- #define UCT\_IFACE\_FLAG\_ERRHANDLE\_ZCOPY\_BUF UCS\_BIT(34)
- #define UCT\_IFACE\_FLAG\_ERRHANDLE\_AM\_ID UCS\_BIT(35)
- #define UCT\_IFACE\_FLAG\_ERRHANDLE\_REMOTE\_MEM UCS\_BIT(36)
- #define UCT\_IFACE\_FLAG\_ERRHANDLE\_BCOPY\_LEN UCS\_BIT(37)
- #define UCT\_IFACE\_FLAG\_ERRHANDLE\_PEER\_FAILURE UCS\_BIT(38)
- #define UCT\_IFACE\_FLAG\_EP\_CHECK UCS\_BIT(39)
- #define UCT\_IFACE\_FLAG\_CONNECT\_TO\_IFACE UCS\_BIT(40)
- #define UCT\_IFACE\_FLAG\_CONNECT\_TO\_EP UCS\_BIT(41)
- #define UCT\_IFACE\_FLAG\_CONNECT\_TO\_SOCKADDR UCS\_BIT(42)
- #define UCT\_IFACE\_FLAG\_AM\_DUP UCS\_BIT(43)
- #define UCT\_IFACE\_FLAG\_CB\_SYNC UCS\_BIT(44)
- #define UCT\_IFACE\_FLAG\_CB\_ASYNC UCS\_BIT(45)
- #define UCT\_IFACE\_FLAG\_EP\_KEEPALIVE UCS\_BIT(46)
- #define UCT\_IFACE\_FLAG\_TAG\_EAGER\_SHORT UCS\_BIT(50)
- #define UCT\_IFACE\_FLAG\_TAG\_EAGER\_BCOPY UCS\_BIT(51)
- #define UCT\_IFACE\_FLAG\_TAG\_EAGER\_ZCOPY UCS\_BIT(52)
- #define UCT\_IFACE\_FLAG\_TAG\_RNDV\_ZCOPY UCS\_BIT(53)

### 6.19.1 Detailed Description

The definition list presents a full list of operations and capabilities exposed by UCX API.

### 6.19.2 Macro Definition Documentation

### 6.19.2.1 UCT\_IFACE\_FLAG\_AM\_SHORT

```
#define UCT_IFACE_FLAG_AM_SHORT UCS_BIT(0)
```

Short active message

#### Examples

[uct\\_hello\\_world.c](#).

### 6.19.2.2 UCT\_IFACE\_FLAG\_AM\_BCOPY

```
#define UCT_IFACE_FLAG_AM_BCOPY UCS_BIT(1)
```

Buffered active message

#### Examples

[uct\\_hello\\_world.c](#).

### 6.19.2.3 UCT\_IFACE\_FLAG\_AM\_ZCOPY

```
#define UCT_IFACE_FLAG_AM_ZCOPY UCS_BIT(2)
```

Zero-copy active message

#### Examples

[uct\\_hello\\_world.c](#).

### 6.19.2.4 UCT\_IFACE\_FLAG\_PENDING

```
#define UCT_IFACE_FLAG_PENDING UCS_BIT(3)
```

Pending operations

### 6.19.2.5 UCT\_IFACE\_FLAG\_PUT\_SHORT

```
#define UCT_IFACE_FLAG_PUT_SHORT UCS_BIT(4)
```

Short put

### 6.19.2.6 UCT\_IFACE\_FLAG\_PUT\_BCOPY

```
#define UCT_IFACE_FLAG_PUT_BCOPY UCS_BIT(5)
```

Buffered put

### 6.19.2.7 UCT\_IFACE\_FLAG\_PUT\_ZCOPY

```
#define UCT_IFACE_FLAG_PUT_ZCOPY UCS_BIT(6)
```

Zero-copy put

### 6.19.2.8 UCT\_IFACE\_FLAG\_GET\_SHORT

```
#define UCT_IFACE_FLAG_GET_SHORT UCS_BIT(8)
```

Short get

### 6.19.2.9 UCT\_IFACE\_FLAG\_GET\_BCOPY

```
#define UCT_IFACE_FLAG_GET_BCOPY UCS_BIT(9)
```

Buffered get

### 6.19.2.10 UCT\_IFACE\_FLAG\_GET\_ZCOPY

```
#define UCT_IFACE_FLAG_GET_ZCOPY UCS_BIT(10)
```

Zero-copy get

### 6.19.2.11 UCT\_IFACE\_FLAG\_ATOMIC\_CPU

```
#define UCT_IFACE_FLAG_ATOMIC_CPU UCS_BIT(30)
```

Atomic communications are consistent with respect to CPU operations.

### 6.19.2.12 UCT\_IFACE\_FLAG\_ATOMIC\_DEVICE

```
#define UCT_IFACE_FLAG_ATOMIC_DEVICE UCS_BIT(31)
```

Atomic communications are consistent only with respect to other atomics on the same device.

### 6.19.2.13 UCT\_IFACE\_FLAG\_ERRHANDLE\_SHORT\_BUF

```
#define UCT_IFACE_FLAG_ERRHANDLE_SHORT_BUF UCS_BIT(32)
```

Invalid buffer for short operation

### 6.19.2.14 UCT\_IFACE\_FLAG\_ERRHANDLE\_BCOPY\_BUF

```
#define UCT_IFACE_FLAG_ERRHANDLE_BCOPY_BUF UCS_BIT(33)
```

Invalid buffer for buffered operation

### 6.19.2.15 UCT\_IFACE\_FLAG\_ERRHANDLE\_ZCOPY\_BUF

```
#define UCT_IFACE_FLAG_ERRHANDLE_ZCOPY_BUF UCS_BIT(34)
```

Invalid buffer for zero copy operation

### 6.19.2.16 UCT\_IFACE\_FLAG\_ERRHANDLE\_AM\_ID

```
#define UCT_IFACE_FLAG_ERRHANDLE_AM_ID UCS_BIT(35)
```

Invalid AM id on remote



**6.19.2.17 UCT\_IFACE\_FLAG\_ERRHANDLE\_REMOTE\_MEM**

```
#define UCT_IFACE_FLAG_ERRHANDLE_REMOTE_MEM UCS_BIT(36)
```

Remote memory access

**6.19.2.18 UCT\_IFACE\_FLAG\_ERRHANDLE\_BCOPY\_LEN**

```
#define UCT_IFACE_FLAG_ERRHANDLE_BCOPY_LEN UCS_BIT(37)
```

Invalid length for buffered operation

**6.19.2.19 UCT\_IFACE\_FLAG\_ERRHANDLE\_PEER\_FAILURE**

```
#define UCT_IFACE_FLAG_ERRHANDLE_PEER_FAILURE UCS_BIT(38)
```

Remote peer failures/outage

**6.19.2.20 UCT\_IFACE\_FLAG\_EP\_CHECK**

```
#define UCT_IFACE_FLAG_EP_CHECK UCS_BIT(39)
```

Endpoint check

**6.19.2.21 UCT\_IFACE\_FLAG\_CONNECT\_TO\_IFACE**

```
#define UCT_IFACE_FLAG_CONNECT_TO_IFACE UCS_BIT(40)
```

Supports connecting to interface

**Examples**

[uct\\_hello\\_world.c](#).

**6.19.2.22 UCT\_IFACE\_FLAG\_CONNECT\_TO\_EP**

```
#define UCT_IFACE_FLAG_CONNECT_TO_EP UCS_BIT(41)
```

Supports connecting to specific endpoint

**Examples**

[uct\\_hello\\_world.c](#).

**6.19.2.23 UCT\_IFACE\_FLAG\_CONNECT\_TO\_SOCKADDR**

```
#define UCT_IFACE_FLAG_CONNECT_TO_SOCKADDR UCS_BIT(42)
```

Supports connecting to sockaddr

#### 6.19.2.24 UCT\_IFACE\_FLAG\_AM\_DUP

```
#define UCT_IFACE_FLAG_AM_DUP UCS_BIT(43)
```

Active messages may be received with duplicates This happens if the transport does not keep enough information to detect retransmissions

#### 6.19.2.25 UCT\_IFACE\_FLAG\_CB\_SYNC

```
#define UCT_IFACE_FLAG_CB_SYNC UCS_BIT(44)
```

Interface supports setting a callback which is invoked only from the calling context of [uct\\_worker\\_progress\(\)](#)

#### 6.19.2.26 UCT\_IFACE\_FLAG\_CB\_ASYNC

```
#define UCT_IFACE_FLAG_CB_ASYNC UCS_BIT(45)
```

Interface supports setting a callback which will be invoked within a reasonable amount of time if [uct\\_worker\\_progress\(\)](#) is not being called. The callback can be invoked from any progress context and it may also be invoked when [uct\\_worker\\_progress\(\)](#) is called.

#### 6.19.2.27 UCT\_IFACE\_FLAG\_EP\_KEEPALIVE

```
#define UCT_IFACE_FLAG_EP_KEEPALIVE UCS_BIT(46)
```

Transport endpoint has built-in keepalive feature, which guarantees the error callback on the transport interface will be called if the communication channel with remote peer is broken, even if there are no outstanding send operations

#### 6.19.2.28 UCT\_IFACE\_FLAG\_TAG\_EAGER\_SHORT

```
#define UCT_IFACE_FLAG_TAG_EAGER_SHORT UCS_BIT(50)
```

Hardware tag matching short eager support

#### 6.19.2.29 UCT\_IFACE\_FLAG\_TAG\_EAGER\_BCOPY

```
#define UCT_IFACE_FLAG_TAG_EAGER_BCOPY UCS_BIT(51)
```

Hardware tag matching bcopy eager support

#### 6.19.2.30 UCT\_IFACE\_FLAG\_TAG\_EAGER\_ZCOPY

```
#define UCT_IFACE_FLAG_TAG_EAGER_ZCOPY UCS_BIT(52)
```

Hardware tag matching zcopy eager support

#### 6.19.2.31 UCT\_IFACE\_FLAG\_TAG\_RNDV\_ZCOPY

```
#define UCT_IFACE_FLAG_TAG_RNDV_ZCOPY UCS_BIT(53)
```

Hardware tag matching rendezvous zcopy support

## 6.20 UCT interface for asynchronous event capabilities

List of capabilities supported by UCT iface event API.

### Macros

- #define `UCT_IFACE_FLAG_EVENT_SEND_COMP` UCS\_BIT(0)
- #define `UCT_IFACE_FLAG_EVENT_RECV` UCS\_BIT(1)
- #define `UCT_IFACE_FLAG_EVENT_RECV_SIG` UCS\_BIT(2)
- #define `UCT_IFACE_FLAG_EVENT_FD` UCS\_BIT(3)
- #define `UCT_IFACE_FLAG_EVENT_ASYNC_CB` UCS\_BIT(4)

### 6.20.1 Detailed Description

The definition list presents a full list of operations and capabilities supported by UCT iface event.

### 6.20.2 Macro Definition Documentation

#### 6.20.2.1 UCT\_IFACE\_FLAG\_EVENT\_SEND\_COMP

```
#define UCT_IFACE_FLAG_EVENT_SEND_COMP UCS_BIT(0)
```

Event notification of send completion is supported

#### 6.20.2.2 UCT\_IFACE\_FLAG\_EVENT\_RECV

```
#define UCT_IFACE_FLAG_EVENT_RECV UCS_BIT(1)
```

Event notification of tag and active message receive is supported

#### 6.20.2.3 UCT\_IFACE\_FLAG\_EVENT\_RECV\_SIG

```
#define UCT_IFACE_FLAG_EVENT_RECV_SIG UCS_BIT(2)
```

Event notification of signaled tag and active message is supported

#### 6.20.2.4 UCT\_IFACE\_FLAG\_EVENT\_FD

```
#define UCT_IFACE_FLAG_EVENT_FD UCS_BIT(3)
```

Event notification through File Descriptor is supported

#### 6.20.2.5 UCT\_IFACE\_FLAG\_EVENT\_ASYNC\_CB

```
#define UCT_IFACE_FLAG_EVENT_ASYNC_CB UCS_BIT(4)
```

Event notification through asynchronous callback invocation is supported

## 6.21 Unified Communication Services (UCS) API

This section describes UCS API.

### Modules

- [UCS Communication Resource](#)

### 6.21.1 Detailed Description

## 6.22 UCS Communication Resource

### Data Structures

- struct [ucs\\_sock\\_addr](#)

### Typedefs

- typedef void(\* [ucs\\_async\\_event\\_cb\\_t](#)) (int id, ucs\_event\_set\_types\_t events, void \*arg)
- typedef struct [ucs\\_sock\\_addr](#) [ucs\\_sock\\_addr\\_t](#)
- typedef enum [ucs\\_memory\\_type](#) [ucs\\_memory\\_type\\_t](#)  
*Memory types.*
- typedef unsigned long [ucs\\_time\\_t](#)
- typedef void \* [ucs\\_status\\_ptr\\_t](#)  
*Status pointer.*

### Enumerations

- enum [ucs\\_callbackq\\_flags](#) { [UCS\\_CALLBACKQ\\_FLAG\\_FAST](#) = [UCS\\_BIT\(0\)](#), [UCS\\_CALLBACKQ\\_FLAG\\_ONESHOT](#) = [UCS\\_BIT\(1\)](#) }
- enum [ucs\\_memory\\_type](#) { [UCS\\_MEMORY\\_TYPE\\_HOST](#), [UCS\\_MEMORY\\_TYPE\\_CUDA](#), [UCS\\_MEMORY\\_TYPE\\_CUDA\\_MANAGED](#), [UCS\\_MEMORY\\_TYPE\\_ROCM](#), [UCS\\_MEMORY\\_TYPE\\_ROCM\\_MANAGED](#), [UCS\\_MEMORY\\_TYPE\\_LAST](#), [UCS\\_MEMORY\\_TYPE\\_UNKNOWN](#) = [UCS\\_MEMORY\\_TYPE\\_LAST](#) }  
*Memory types.*
- enum [ucs\\_status\\_t](#) { [UCS\\_OK](#) = 0, [UCS\\_INPROGRESS](#) = 1, [UCS\\_ERR\\_NO\\_MESSAGE](#) = -1, [UCS\\_ERR\\_NO\\_RESOURCE](#) = -2, [UCS\\_ERR\\_IO\\_ERROR](#) = -3, [UCS\\_ERR\\_NO\\_MEMORY](#) = -4, [UCS\\_ERR\\_INVALID\\_PARAM](#) = -5, [UCS\\_ERR\\_UNREACHABLE](#) = -6, [UCS\\_ERR\\_INVALID\\_ADDR](#) = -7, [UCS\\_ERR\\_NOT\\_IMPLEMENTED](#) = -8, [UCS\\_ERR\\_MESSAGE\\_TRUNCATED](#) = -9, [UCS\\_ERR\\_NO\\_PROGRESS](#) = -10, [UCS\\_ERR\\_BUFFER\\_TOO\\_SMALL](#) = -11, [UCS\\_ERR\\_NO\\_ELEM](#) = -12, [UCS\\_ERR\\_SOME\\_CONNECTS\\_FAILED](#) = -13, [UCS\\_ERR\\_NO\\_DEVICE](#) = -14, [UCS\\_ERR\\_BUSY](#) = -15, [UCS\\_ERR\\_CANCELED](#) = -16, [UCS\\_ERR\\_SHMEM\\_SEGMENT](#) = -17, [UCS\\_ERR\\_ALREADY\\_EXISTS](#) = -18, [UCS\\_ERR\\_OUT\\_OF\\_RANGE](#) = -19, [UCS\\_ERR\\_TIMED\\_OUT](#) = -20, [UCS\\_ERR\\_EXCEEDS\\_LIMIT](#) = -21, [UCS\\_ERR\\_UNSUPPORTED](#) = -22, [UCS\\_ERR\\_REJECTED](#) = -23, [UCS\\_ERR\\_NOT\\_CONNECTED](#) = -24, [UCS\\_ERR\\_CONNECTION\\_RESET](#) = -25, [UCS\\_ERR\\_FIRST\\_LINK\\_FAILURE](#) = -40, [UCS\\_ERR\\_LAST\\_LINK\\_FAILURE](#) = -59, [UCS\\_ERR\\_FIRST\\_ENDPOINT\\_FAILURE](#) = -60, [UCS\\_ERR\\_ENDPOINT\\_TIMEOUT](#) = -80, [UCS\\_ERR\\_LAST\\_ENDPOINT\\_FAILURE](#) = -89, [UCS\\_ERR\\_LAST](#) = -100 }  
*Status codes.*
- enum [ucs\\_thread\\_mode\\_t](#) { [UCS\\_THREAD\\_MODE\\_SINGLE](#), [UCS\\_THREAD\\_MODE\\_SERIALIZED](#), [UCS\\_THREAD\\_MODE\\_MULTI](#), [UCS\\_THREAD\\_MODE\\_LAST](#) }  
*Thread sharing mode.*

### Functions

- [ucs\\_status\\_t ucs\\_async\\_set\\_event\\_handler](#) (ucs\_async\_mode\_t mode, int event\_fd, ucs\_event\_set\_types\_t events, [ucs\\_async\\_event\\_cb\\_t](#) cb, void \*arg, ucs\_async\_context\_t \*async)

- `ucs_status_t ucs_async_add_timer` (`ucs_async_mode_t mode`, `ucs_time_t interval`, `ucs_async_event_cb_t cb`, `void *arg`, `ucs_async_context_t *async`, `int *timer_id_p`)
- `ucs_status_t ucs_async_remove_handler` (`int id`, `int sync`)
- `ucs_status_t ucs_async_modify_handler` (`int fd`, `ucs_event_set_types_t events`)
- `ucs_status_t ucs_async_context_create` (`ucs_async_mode_t mode`, `ucs_async_context_t **async_p`)  
*Create an asynchronous execution context.*
- `void ucs_async_context_destroy` (`ucs_async_context_t *async`)  
*Destroy the asynchronous execution context.*
- `void ucs_async_poll` (`ucs_async_context_t *async`)

### 6.22.1 Detailed Description

This section describes a concept of the Communication Resource and routines associated with the concept.

### 6.22.2 Data Structure Documentation

#### 6.22.2.1 struct ucs\_sock\_addr

BSD socket address specification.

Data Fields

<code>const struct sockaddr *</code>	<code>addr</code>	Pointer to socket address
<code>socklen_t</code>	<code>addrlen</code>	Address length

### 6.22.3 Typedef Documentation

#### 6.22.3.1 ucs\_async\_event\_cb\_t

```
typedef void(* ucs_async_event_cb_t) (int id, ucs_event_set_types_t events, void *arg)
```

Async event callback.

Parameters

<i>id</i>	Event id (timer or file descriptor).
<i>events</i>	The events that triggered the callback.
<i>arg</i>	User-defined argument.

#### 6.22.3.2 ucs\_sock\_addr\_t

```
typedef struct ucs_sock_addr ucs_sock_addr_t
```

BSD socket address specification.

**6.22.3.3 ucs\_memory\_type\_t**

```
typedef enum ucs_memory_type ucs_memory_type_t
```

List of supported memory types.

**6.22.3.4 ucs\_time\_t**

```
typedef unsigned long ucs_time_t
```

UCS time units. These are not necessarily aligned with metric time units. MUST compare short time values with UCS\_SHORT\_TIME\_CMP to handle wrap-around.

**6.22.3.5 ucs\_status\_ptr\_t**

```
typedef void* ucs_status_ptr_t
```

A pointer can represent one of these values:

- NULL / UCS\_OK
- Error code pointer (UCS\_ERR\_xx)
- Valid pointer

**6.22.4 Enumeration Type Documentation****6.22.4.1 ucs\_callbackq\_flags**

```
enum ucs_callbackq_flags
```

Callback flags

Enumerator

UCS_CALLBACKQ_FLAG_FAST	Fast-path (best effort)
UCS_CALLBACKQ_FLAG_ONESHOT	Call the callback only once (cannot be used with FAST)

**6.22.4.2 ucs\_memory\_type**

```
enum ucs_memory_type
```

List of supported memory types.

Enumerator

UCS_MEMORY_TYPE_HOST	Default system memory
UCS_MEMORY_TYPE_CUDA	NVIDIA CUDA memory
UCS_MEMORY_TYPE_CUDA_MANAGED	NVIDIA CUDA managed (or unified) memory
UCS_MEMORY_TYPE_ROCM	AMD ROCM memory
UCS_MEMORY_TYPE_ROCM_MANAGED	AMD ROCM managed system memory
UCS_MEMORY_TYPE_LAST	

## Enumerator

UCS_MEMORY_TYPE_UNKNOWN	
-------------------------	--

## 6.22.4.3 ucs\_status\_t

```
enum ucs_status_t
```

## Note

In order to evaluate the necessary steps to recover from a certain error, all error codes which can be returned by the external API are grouped by the largest entity permanently effected by the error. Each group ranges between its UCS\_ERR\_FIRST\_<name> and UCS\_ERR\_LAST\_<name> enum values. For example, if a link fails it may be sufficient to destroy (and possibly replace) it, in contrast to an endpoint-level error.

## Enumerator

UCS_OK	
UCS_INPROGRESS	
UCS_ERR_NO_MESSAGE	
UCS_ERR_NO_RESOURCE	
UCS_ERR_IO_ERROR	
UCS_ERR_NO_MEMORY	
UCS_ERR_INVALID_PARAM	
UCS_ERR_UNREACHABLE	
UCS_ERR_INVALID_ADDR	
UCS_ERR_NOT_IMPLEMENTED	
UCS_ERR_MESSAGE_TRUNCATED	
UCS_ERR_NO_PROGRESS	
UCS_ERR_BUFFER_TOO_SMALL	
UCS_ERR_NO_ELEM	
UCS_ERR_SOME_CONNECTS_FAILED	
UCS_ERR_NO_DEVICE	
UCS_ERR_BUSY	
UCS_ERR_CANCELED	
UCS_ERR_SHMEM_SEGMENT	
UCS_ERR_ALREADY_EXISTS	
UCS_ERR_OUT_OF_RANGE	
UCS_ERR_TIMED_OUT	
UCS_ERR_EXCEEDS_LIMIT	
UCS_ERR_UNSUPPORTED	
UCS_ERR_REJECTED	
UCS_ERR_NOT_CONNECTED	
UCS_ERR_CONNECTION_RESET	
UCS_ERR_FIRST_LINK_FAILURE	
UCS_ERR_LAST_LINK_FAILURE	
UCS_ERR_FIRST_ENDPOINT_FAILURE	
UCS_ERR_ENDPOINT_TIMEOUT	
UCS_ERR_LAST_ENDPOINT_FAILURE	
UCS_ERR_LAST	



## Examples

[uct\\_hello\\_world.c](#).

6.22.4.4 `ucs_thread_mode_t`

enum `ucs_thread_mode_t`

Specifies thread sharing mode of an object.

## Enumerator

<code>UCS_THREAD_MODE_SINGLE</code>	Only the master thread can access (i.e. the thread that initialized the context; multiple threads may exist and never access)
<code>UCS_THREAD_MODE_SERIALIZED</code>	Multiple threads can access, but only one at a time
<code>UCS_THREAD_MODE_MULTI</code>	Multiple threads can access concurrently
<code>UCS_THREAD_MODE_LAST</code>	

## 6.22.5 Function Documentation

6.22.5.1 `ucs_async_set_event_handler()`

```
ucs_status_t ucs_async_set_event_handler (
    ucs_async_mode_t mode,
    int event_fd,
    ucs_event_set_types_t events,
    ucs_async_event_cb_t cb,
    void * arg,
    ucs_async_context_t * async )
```

Register a file descriptor for monitoring (call handler upon events). Every fd can have only one handler.

## Parameters

<i>mode</i>	Thread or signal.
<i>event↔ _fd</i>	File descriptor to set handler for.
<i>events</i>	Events to wait on (UCS_EVENT_SET_EVxxx bits).
<i>cb</i>	Callback function to execute.
<i>arg</i>	Argument to callback.
<i>async</i>	Async context to which events are delivered. If NULL, safety is up to the user.

**Returns**

Error code as defined by [ucs\\_status\\_t](#).

**6.22.5.2 ucs\_async\_add\_timer()**

```
ucs_status_t ucs_async_add_timer (
    ucs_async_mode_t mode,
    ucs_time_t interval,
    ucs_async_event_cb_t cb,
    void * arg,
    ucs_async_context_t * async,
    int * timer_id_p )
```

Add timer handler.

**Parameters**

<i>mode</i>	Thread or signal.
<i>interval</i>	Timer interval.
<i>cb</i>	Callback function to execute.
<i>arg</i>	Argument to callback.
<i>async</i>	Async context to which events are delivered. If NULL, safety is up to the user.
<i>timer_id</i> ↔ <i>_p</i>	Filled with timer id.

**Returns**

Error code as defined by [ucs\\_status\\_t](#).

**6.22.5.3 ucs\_async\_remove\_handler()**

```
ucs_status_t ucs_async_remove_handler (
    int id,
    int sync )
```

Remove an event handler (Timer or event file).

**Parameters**

<i>id</i>	Timer/FD to remove.
<i>sync</i>	If nonzero, wait until the handler for this event is not running anymore. If called from the context of the callback, the handler will be removed immediately after the current callback returns.

**Returns**

Error code as defined by [ucs\\_status\\_t](#).

**6.22.5.4 ucs\_async\_modify\_handler()**

```
ucs_status_t ucs_async_modify_handler (
    int fd,
    ucs_event_set_types_t events )
```

Modify events mask for an existing event handler (event file).

**Parameters**

<i>fd</i>	File descriptor modify events for.
<i>events</i>	New set of events to wait on (UCS_EVENT_SET_EVxxx bits).

**Returns**

Error code as defined by [ucs\\_status\\_t](#).

**6.22.5.5 ucs\_async\_context\_create()**

```
ucs_status_t ucs_async_context_create (
    ucs_async_mode_t mode,
    ucs_async_context_t ** async_p )
```

Allocate and initialize an asynchronous execution context. This can be used to ensure safe event delivery.

**Parameters**

<i>mode</i>	Indicates whether to use signals or polling threads for waiting.
<i>async_p</i>	Event context pointer to initialize.

**Returns**

Error code as defined by [ucs\\_status\\_t](#).

**Examples**

[uct\\_hello\\_world.c](#).

**6.22.5.6 ucs\_async\_context\_destroy()**

```
void ucs_async_context_destroy (
    ucs_async_context_t * async )
```

Clean up the async context, and release system resources if possible. The context memory released.

**Parameters**

<i>async</i>	Asynchronous context to clean up.
--------------	-----------------------------------

**Examples**

[uct\\_hello\\_world.c](#).

**6.22.5.7 ucs\_async\_poll()**

```
void ucs_async_poll (
    ucs_async_context_t * async )
```

Poll on async context.

**Parameters**

<i>async</i>	Async context to poll on. NULL polls on all.
--------------	--

## Chapter 7

# Data Structure Documentation

### 7.1 ucp\_generic\_dt\_ops Struct Reference

UCP generic data type descriptor.

#### Data Fields

- void [\\*\(\*start\\_pack\*\)](#)(void \*context, const void \*buffer, size\_t count)  
*Start a packing request.*
- void [\\*\(\*start\\_unpack\*\)](#)(void \*context, void \*buffer, size\_t count)  
*Start an unpacking request.*
- size\_t [\(\*packed\\_size\*\)](#)(void \*state)  
*Get the total size of packed data.*
- size\_t [\(\*pack\*\)](#)(void \*state, size\_t offset, void \*dest, size\_t max\_length)  
*Pack data.*
- [ucs\\_status\\_t\(\*unpack\*\)](#)(void \*state, size\_t offset, const void \*src, size\_t length)  
*Unpack data.*
- void [\(\*finish\*\)](#)(void \*state)  
*Finish packing/unpacking.*

#### 7.1.1 Detailed Description

This structure provides a generic datatype descriptor that is used for definition of application defined datatypes.

Typically, the descriptor is used for an integration with datatype engines implemented within MPI and SHMEM implementations.

#### Note

In case of partial receive, any amount of received data is acceptable which matches buffer size.

The documentation for this struct was generated from the following file:

- ucp.h

### 7.2 uct\_tag\_context Struct Reference

Posted tag context.

## Data Fields

- void(\* [tag\\_consumed\\_cb](#))([uct\\_tag\\_context\\_t](#) \*self)
- void(\* [completed\\_cb](#))([uct\\_tag\\_context\\_t](#) \*self, [uct\\_tag\\_t](#) stag, [uint64\\_t](#) imm, [size\\_t](#) length, void \*inline\_data, [ucs\\_status\\_t](#) status)
- void(\* [rndv\\_cb](#))([uct\\_tag\\_context\\_t](#) \*self, [uct\\_tag\\_t](#) stag, const void \*header, unsigned header\_length, [ucs\\_status\\_t](#) status, unsigned flags)
- char [priv](#) [UCT\_TAG\_PRIV\_LEN]

### 7.2.1 Detailed Description

Tag context is an object which tracks a tag posted to the transport. It contains callbacks for matching events on this tag.

### 7.2.2 Field Documentation

#### 7.2.2.1 tag\_consumed\_cb

```
void(* uct_tag_context::tag_consumed_cb) (uct\_tag\_context\_t *self)
```

Tag is consumed by the transport and should not be matched in software.

##### Parameters

in	<i>self</i>	Pointer to relevant context structure, which was initially passed to <a href="#">uct_iface_tag_rcv_zcopy</a> .
----	-------------	--

#### 7.2.2.2 completed\_cb

```
void(* uct_tag_context::completed_cb) (uct\_tag\_context\_t *self, uct\_tag\_t stag, uint64\_t imm, size\_t length, void *inline_data, ucs\_status\_t status)
```

Tag processing is completed by the transport.

##### Parameters

in	<i>self</i>	Pointer to relevant context structure, which was initially passed to <a href="#">uct_iface_tag_rcv_zcopy</a> .
in	<i>stag</i>	Tag from sender.
in	<i>imm</i>	Immediate data from sender. For rendezvous, it's always 0.
in	<i>length</i>	Completed length.
in	<i>inline_data</i>	If non-null, points to a temporary buffer which contains the received data. In this case the received data was not placed directly in the receive buffer. This callback routine is responsible for copy-out the inline data, otherwise it is released.
in	<i>status</i>	Completion status: (a) UCS_OK - Success, data placed in provided buffer. (b) UCS_ERR_TRUNCATED - Sender's length exceed posted buffer, no data is copied. (c) UCS_ERR_CANCELED - Canceled by user.

### 7.2.2.3 rndv\_cb

```
void(* uct_tag_context::rndv_cb) (uct_tag_context_t *self, uct_tag_t stag, const void *header,
unsigned header_length, ucs_status_t status, unsigned flags)
```

Tag was matched by a rendezvous request, which should be completed by the protocol layer.

#### Parameters

in	<i>self</i>	Pointer to relevant context structure, which was initially passed to <a href="#">uct_iface_tag_rcv_zcopy</a> .
in	<i>stag</i>	Tag from sender.
in	<i>header</i>	User defined header.
in	<i>header_length</i>	User defined header length in bytes.
in	<i>status</i>	Completion status.
in	<i>flags</i>	Flags defined by UCT_TAG_RECV_CB_XX.

### 7.2.2.4 priv

```
char uct_tag_context::priv[UCT_TAG_PRIV_LEN]
```

A placeholder for the private data used by the transport

The documentation for this struct was generated from the following file:

- uct.h





## Chapter 8

# Example Documentation

### 8.1 ucp\_client\_server.c

UCP client / server example using different APIs (tag, stream, am) utility.

```
/*
 * UCP client - server example utility
 * -----
 *
 * Server side:
 *   ./ucp_client_server
 *
 * Client side:
 *   ./ucp_client_server -a <server-ip>
 *
 * Notes:
 *
 * - The server will listen to incoming connection requests on INADDR_ANY.
 * - The client needs to pass the IP address of the server side to connect to
 *   as an argument to the test.
 * - Currently, the passed IP needs to be an IPoIB or a RoCE address.
 * - The port which the server side would listen on can be modified with the
 *   '-p' option and should be used on both sides. The default port to use is
 *   13337.
 */
#include "hello_world_util.h"
#include "ucp_util.h"
#include <ucp/api/ucp.h>
#include <string.h> /* memset */
#include <arpa/inet.h> /* inet_addr */
#include <unistd.h> /* getopt */
#include <stdlib.h> /* atoi */
#define DEFAULT_PORT 13337
#define IP_STRING_LEN 50
#define PORT_STRING_LEN 8
#define TAG 0xCAFE
#define COMM_TYPE_DEFAULT "STREAM"
#define PRINT_INTERVAL 2000
#define DEFAULT_NUM_ITERATIONS 1
#define TEST_AM_ID 0
static long test_string_length = 16;
static long iov_cnt = 1;
static uint16_t server_port = DEFAULT_PORT;
static sa_family_t ai_family = AF_INET;
static int num_iterations = DEFAULT_NUM_ITERATIONS;
typedef enum {
    CLIENT_SERVER_SEND_RECV_STREAM = UCS_BIT(0),
    CLIENT_SERVER_SEND_RECV_TAG = UCS_BIT(1),
    CLIENT_SERVER_SEND_RECV_AM = UCS_BIT(2),
    CLIENT_SERVER_SEND_RECV_DEFAULT = CLIENT_SERVER_SEND_RECV_STREAM
} send_recv_type_t;
typedef struct ucx_server_ctx {
    volatile ucp_conn_request_h conn_request;
    ucp_listener_h listener;
} ucx_server_ctx_t;
typedef struct test_req {
    int complete;
} test_req_t;
static struct {
    volatile int complete;
    int is_rndv;

```

```

    void          *desc;
    void          *recv_buf;
} am_data_desc = {0, 0, NULL, NULL};
static void usage(void);
void buffer_free(ucp_dt_iov_t *iovs)
{
    size_t idx;
    for (idx = 0; idx < iov_cnt; idx++) {
        mem_type_free(iovs[idx].buffer);
    }
}
int buffer_malloc(ucp_dt_iov_t *iovs)
{
    size_t idx;
    for (idx = 0; idx < iov_cnt; idx++) {
        iovs[idx].length = test_string_length;
        iovs[idx].buffer = mem_type_malloc(iovs[idx].length);
        if (iovs[idx].buffer == NULL) {
            buffer_free(iovs);
            return -1;
        }
    }
    return 0;
}
int fill_buffer(ucp_dt_iov_t *iovs)
{
    int ret = 0;
    size_t idx;
    for (idx = 0; idx < iov_cnt; idx++) {
        ret = generate_test_string(iovs[idx].buffer, iovs[idx].length);
        if (ret != 0) {
            break;
        }
    }
    CHKERR_ACTION(ret != 0, "generate test string", return -1);
    return 0;
}
static void common_cb(void *user_data, const char *type_str)
{
    test_req_t *ctx;
    if (user_data == NULL) {
        fprintf(stderr, "user_data passed to %s mustn't be NULL\n", type_str);
        return;
    }
    ctx = user_data;
    ctx->complete = 1;
}
static void tag_recv_cb(void *request, ucs_status_t status,
                       const ucp_tag_recv_info_t *info, void *user_data)
{
    common_cb(user_data, "tag_recv_cb");
}
static void stream_recv_cb(void *request, ucs_status_t status, size_t length,
                          void *user_data)
{
    common_cb(user_data, "stream_recv_cb");
}
static void am_recv_cb(void *request, ucs_status_t status, size_t length,
                      void *user_data)
{
    common_cb(user_data, "am_recv_cb");
}
static void send_cb(void *request, ucs_status_t status, void *user_data)
{
    common_cb(user_data, "send_cb");
}
static void err_cb(void *arg, ucp_ep_h ep, ucs_status_t status)
{
    printf("error handling callback was invoked with status %d (%s)\n",
          status, ucs_status_string(status));
}
void set_sock_addr(const char *address_str, struct sockaddr_storage *saddr)
{
    struct sockaddr_in *sa_in;
    struct sockaddr_in6 *sa_in6;
    /* The server will listen on INADDR_ANY */
    memset(saddr, 0, sizeof(*saddr));
    switch (ai_family) {
    case AF_INET:
        sa_in = (struct sockaddr_in*)saddr;
        if (address_str != NULL) {
            inet_pton(AF_INET, address_str, &sa_in->sin_addr);
        } else {
            sa_in->sin_addr.s_addr = INADDR_ANY;
        }
        sa_in->sin_family = AF_INET;
        sa_in->sin_port = htons(server_port);
    }
}

```

```

        break;
    case AF_INET6:
        sa_in6 = (struct sockaddr_in6*)saddr;
        if (address_str != NULL) {
            inet_pton(AF_INET6, address_str, &sa_in6->sin6_addr);
        } else {
            sa_in6->sin6_addr = in6addr_any;
        }
        sa_in6->sin6_family = AF_INET6;
        sa_in6->sin6_port = htons(server_port);
        break;
    default:
        fprintf(stderr, "Invalid address family");
        break;
}
}
static ucs_status_t start_client(ucp_worker_h ucp_worker,
                                const char *address_str, ucp_ep_h *client_ep)
{
    ucp_ep_params_t ep_params;
    struct sockaddr_storage connect_addr;
    ucs_status_t status;
    set_sock_addr(address_str, &connect_addr);
    /*
     * Endpoint field mask bits:
     * UCP_EP_PARAM_FIELD_FLAGS           - Use the value of the 'flags' field.
     * UCP_EP_PARAM_FIELD_SOCK_ADDR       - Use a remote sockaddr to connect
     *                                     to the remote peer.
     * UCP_EP_PARAM_FIELD_ERR_HANDLING_MODE - Error handling mode - this flag
     *                                     is temporarily required since the
     *                                     endpoint will be closed with
     *                                     UCP_EP_CLOSE_MODE_FORCE which
     *                                     requires this mode.
     *                                     Once UCP_EP_CLOSE_MODE_FORCE is
     *                                     removed, the error handling mode
     *                                     will be removed.
     */
    ep_params.field_mask = UCP_EP_PARAM_FIELD_FLAGS |
                          UCP_EP_PARAM_FIELD_SOCK_ADDR |
                          UCP_EP_PARAM_FIELD_ERR_HANDLER |
                          UCP_EP_PARAM_FIELD_ERR_HANDLING_MODE;
    ep_params.err_mode = UCP_ERR_HANDLING_MODE_PEER;
    ep_params.err_handler.cb = err_cb;
    ep_params.err_handler.arg = NULL;
    ep_params.flags = UCP_EP_PARAMS_FLAGS_CLIENT_SERVER;
    ep_params.sockaddr.addr = (struct sockaddr*)&connect_addr;
    ep_params.sockaddr.addrlen = sizeof(connect_addr);
    status = ucp_ep_create(ucp_worker, &ep_params, client_ep);
    if (status != UCS_OK) {
        fprintf(stderr, "failed to connect to %s (%s)\n", address_str,
                ucs_status_string(status));
    }
    return status;
}
static void print_iov(const ucp_dt_iov_t *iovs)
{
    char *msg = alloca(test_string_length);
    size_t idx;
    for (idx = 0; idx < iov_cnt; idx++) {
        /* In case of Non-System memory */
        mem_type_memcpy(msg, iovs[idx].buffer, test_string_length);
        printf("%s.\n", msg);
    }
}
static
void print_result(int is_server, const ucp_dt_iov_t *iovs, int current_iter)
{
    if (is_server) {
        printf("Server: iteration #%d\n", (current_iter + 1));
        printf("UCX data message was received\n");
        printf("\n\n----- UCP TEST SUCCESS ----- \n\n");
    } else {
        printf("Client: iteration #%d\n", (current_iter + 1));
        printf("\n\n----- \n\n");
    }
    print_iov(iovs);
    printf("\n\n----- \n\n");
}
static ucs_status_t request_wait(ucp_worker_h ucp_worker, void *request,
                                test_req_t *ctx)
{
    ucs_status_t status;
    /* if operation was completed immediately */
    if (request == NULL) {
        return UCS_OK;
    }
    if (UCS_PTR_IS_ERR(request)) {

```

```

        return UCS_PTR_STATUS(request);
    }
    while (ctx->complete == 0) {
        ucp_worker_progress(ucp_worker);
    }
    status = ucp_request_check_status(request);
    ucp_request_free(request);
    return status;
}
static int request_finalize(ucp_worker_h ucp_worker, test_req_t *request,
                           test_req_t *ctx, int is_server, ucp_dt_iov_t *iovs,
                           int current_iter)
{
    int ret = 0;
    ucs_status_t status;
    status = request_wait(ucp_worker, request, ctx);
    if (status != UCS_OK) {
        fprintf(stderr, "unable to %s UCX message (%s)\n",
                is_server ? "receive": "send", ucs_status_string(status));
        ret = -1;
        goto release_iovs;
    }
    /* Print the output of the first, last and every PRINT_INTERVAL iteration */
    if ((current_iter == 0) || (current_iter == (num_iterations - 1)) ||
        !((current_iter + 1) % (PRINT_INTERVAL))) {
        print_result(is_server, iovs, current_iter);
    }
release_iovs:
    buffer_free(iovs);
    return ret;
}
static int
fill_request_param(ucp_dt_iov_t *iovs, int is_client,
                  void **msg, size_t *msg_length,
                  test_req_t *ctx, ucp_request_param_t *param)
{
    CHKERR_ACTION(buffer_malloc(iovs) != 0, "allocate memory", return -1);
    if (is_client && (fill_buffer(iovs) != 0)) {
        buffer_free(iovs);
        return -1;
    }
    *msg = (iovs_cnt == 1) ? iovs[0].buffer : iovs;
    *msg_length = (iovs_cnt == 1) ? iovs[0].length : iovs_cnt;
    ctx->complete = 0;
    param->op_attr_mask = UCP_OP_ATTR_FIELD_CALLBACK |
                        UCP_OP_ATTR_FIELD_DATATYPE |
                        UCP_OP_ATTR_FIELD_USER_DATA;
    param->datatype = (iovs_cnt == 1) ? ucp_dt_make_contig(1) :
                    UCP_DATATYPE_IOV;
    param->user_data = ctx;
    return 0;
}
static int send_rcv_stream(ucp_worker_h ucp_worker, ucp_ep_h ep, int is_server,
                           int current_iter)
{
    ucp_dt_iov_t *iovs = alloca(iovs_cnt * sizeof(ucp_dt_iov_t));
    ucp_request_param_t param;
    test_req_t *request;
    size_t msg_length;
    void *msg;
    test_req_t ctx;
    memset(iovs, 0, iovs_cnt * sizeof(*iovs));
    if (fill_request_param(iovs, !is_server, &msg, &msg_length,
                          &ctx, &param) != 0) {
        return -1;
    }
    if (!is_server) {
        /* Client sends a message to the server using the stream API */
        param.cb.send = send_cb;
        request = ucp_stream_send_nbx(ep, msg, msg_length, &param);
    } else {
        /* Server receives a message from the client using the stream API */
        param.op_attr_mask |= UCP_OP_ATTR_FIELD_FLAGS;
        param.flags = UCP_STREAM_RECV_FLAG_WAITALL;
        param.cb.recv_stream = stream_rcv_cb;
        request = ucp_stream_recv_nbx(ep, msg, msg_length,
                                     &msg_length, &param);
    }
    return request_finalize(ucp_worker, request, &ctx, is_server, iovs,
                           current_iter);
}
static int send_rcv_tag(ucp_worker_h ucp_worker, ucp_ep_h ep, int is_server,
                       int current_iter)
{
    ucp_dt_iov_t *iovs = alloca(iovs_cnt * sizeof(ucp_dt_iov_t));
    ucp_request_param_t param;
    void *request;

```

```

size_t msg_length;
void *msg;
test_req_t ctx;
memset(iov, 0, iov_cnt * sizeof(*iov));
if (fill_request_param(iov, !is_server, &msg, &msg_length,
                      &ctx, &param) != 0) {
    return -1;
}
if (!is_server) {
    /* Client sends a message to the server using the Tag-Matching API */
    param.cb.send = send_cb;
    request      = ucp_tag_send_nbx(ep, msg, msg_length, TAG, &param);
} else {
    /* Server receives a message from the client using the Tag-Matching API */
    param.cb.recv = tag_recv_cb;
    request      = ucp_tag_recv_nbx(ucp_worker, msg, msg_length, TAG, 0,
                                   &param);
}
return request_finalize(ucp_worker, request, &ctx, is_server, iov,
                       current_iter);
}
ucp_status_t ucp_am_data_cb(void *arg, const void *header, size_t header_length,
                           void *data, size_t length,
                           const ucp_am_recv_param_t *param)
{
    ucp_dt_iov_t *iovs;
    size_t idx;
    size_t offset;
    if (length != iov_cnt * test_string_length) {
        fprintf(stderr, "received wrong data length %ld (expected %ld)",
                length, iov_cnt * test_string_length);
        return UCS_OK;
    }
    if (header_length != 0) {
        fprintf(stderr, "received unexpected header, length %ld", header_length);
    }
    am_data_desc.complete = 1;
    if (param->recv_attr & UCP_AM_RECV_ATTR_FLAG_RNDV) {
        /* Rendezvous request arrived, data contains an internal UCX descriptor,
         * which has to be passed to ucp_am_recv_data_nbx function to confirm
         * data transfer.
         */
        am_data_desc.is_rndv = 1;
        am_data_desc.desc    = data;
        return UCS_INPROGRESS;
    }
    /* Message delivered with eager protocol, data should be available
     * immediately
     */
    am_data_desc.is_rndv = 0;
    iov = am_data_desc.recv_buf;
    offset = 0;
    for (idx = 0; idx < iov_cnt; idx++) {
        mem_type_memcpy(iovs[idx].buffer, UCS_PTR_BYTE_OFFSET(data, offset),
                        iov[idx].length);
        offset += iov[idx].length;
    }
    return UCS_OK;
}
static int send_recv_am(ucp_worker_h ucp_worker, ucp_ep_h ep, int is_server,
                       int current_iter)
{
    ucp_dt_iov_t *iovs = alloca(iov_cnt * sizeof(ucp_dt_iov_t));
    test_req_t *request;
    ucp_request_param_t params;
    size_t msg_length;
    void *msg;
    test_req_t ctx;
    memset(iov, 0, iov_cnt * sizeof(*iovs));
    if (fill_request_param(iov, !is_server, &msg, &msg_length,
                          &ctx, &params) != 0) {
        return -1;
    }
    if (is_server) {
        am_data_desc.recv_buf = iov;
        /* waiting for AM callback has called */
        while (!am_data_desc.complete) {
            ucp_worker_progress(ucp_worker);
        }
        am_data_desc.complete = 0;
        if (am_data_desc.is_rndv) {
            /* Rendezvous request has arrived, need to invoke receive operation
             * to confirm data transfer from the sender to the "recv_message"
             * buffer.
             */
            params.op_attr_mask |= UCP_OP_ATTR_FLAG_NO_IMM_CMPL;
            params.cb.recv_am    = am_recv_cb;
            request              = ucp_am_recv_data_nbx(ucp_worker,

```

```

am_data_desc.desc,
msg, msg_length,
&params);
} else {
    /* Data has arrived eagerly and is ready for use, no need to
    * initiate receive operation. */
    request = NULL;
}
} else {
    /* Client sends a message to the server using the AM API */
    params.cb.send = (ucp_send_nbx_callback_t)send_cb,
    request = ucp_am_send_nbx(ep, TEST_AM_ID, NULL, 0ul, msg,
    msg_length, &params);
}
return request_finalize(ucp_worker, request, &ctx, is_server, iov,
    current_iter);
}
static void usage()
{
    fprintf(stderr, "Usage: ucp_client_server [parameters]\n");
    fprintf(stderr, "UCP client-server example utility\n");
    fprintf(stderr, "\nParameters are:\n");
    fprintf(stderr, " -a Set IP address of the server "
    "(required for client and should not be specified "
    "for the server)\n");
    fprintf(stderr, " -l Set IP address where server listens "
    "(If not specified, server uses INADDR_ANY; "
    "Irrelevant at client)\n");
    fprintf(stderr, " -p Port number to listen/connect to (default = %d). "
    "0 on the server side means select a random port and print it\n",
    DEFAULT_PORT);
    fprintf(stderr, " -c Communication type for the client and server. "
    "Valid values are:\n"
    "'stream' : Stream API\n"
    "'tag' : Tag API\n"
    "'am' : AM API\n"
    "If not specified, %s API will be used.\n", COMM_TYPE_DEFAULT);
    fprintf(stderr, " -i Number of iterations to run. Client and server must "
    "have the same value. (default = %d).\n",
    num_iterations);
    fprintf(stderr, " -v Number of buffers in a single data "
    "transfer function call. (default = %ld).\n",
    iov_cnt);
    print_common_help();
    fprintf(stderr, "\n");
}
static int parse_cmd(int argc, char *const argv[], char **server_addr,
    char **listen_addr, send_recv_type_t *send_recv_type)
{
    int c = 0;
    int port;
    while ((c = getopt(argc, argv, "a:l:p:c:6i:s:v:m:h")) != -1) {
        switch (c) {
            case 'a':
                *server_addr = optarg;
                break;
            case 'c':
                if (!strcasecmp(optarg, "stream")) {
                    *send_recv_type = CLIENT_SERVER_SEND_RECV_STREAM;
                } else if (!strcasecmp(optarg, "tag")) {
                    *send_recv_type = CLIENT_SERVER_SEND_RECV_TAG;
                } else if (!strcasecmp(optarg, "am")) {
                    *send_recv_type = CLIENT_SERVER_SEND_RECV_AM;
                } else {
                    fprintf(stderr, "Wrong communication type %s. "
                    "Using %s as default\n", optarg, COMM_TYPE_DEFAULT);
                    *send_recv_type = CLIENT_SERVER_SEND_RECV_DEFAULT;
                }
                break;
            case 'l':
                *listen_addr = optarg;
                break;
            case 'p':
                port = atoi(optarg);
                if ((port < 0) || (port > UINT16_MAX)) {
                    fprintf(stderr, "Wrong server port number %d\n", port);
                    return -1;
                }
                server_port = port;
                break;
            case '6':
                ai_family = AF_INET6;
                break;
            case 'i':
                num_iterations = atoi(optarg);
                break;
            case 's':

```

```

        test_string_length = atol(optarg);
        if (test_string_length < 0) {
            fprintf(stderr, "Wrong string size %ld\n", test_string_length);
            return UCS_ERR_UNSUPPORTED;
        }
        break;
    case 'v':
        iov_cnt = atol(optarg);
        if (iov_cnt <= 0) {
            fprintf(stderr, "Wrong iov count %ld\n", iov_cnt);
            return UCS_ERR_UNSUPPORTED;
        }
        break;
    case 'm':
        test_mem_type = parse_mem_type(optarg);
        if (test_mem_type == UCS_MEMORY_TYPE_LAST) {
            return UCS_ERR_UNSUPPORTED;
        }
        break;
    case 'h':
    default:
        usage();
        return -1;
    }
}
return 0;
}
static char* sockaddr_get_ip_str(const struct sockaddr_storage *sock_addr,
                                char *ip_str, size_t max_size)
{
    struct sockaddr_in  addr_in;
    struct sockaddr_in6 addr_in6;
    switch (sock_addr->ss_family) {
    case AF_INET:
        memcpy(&addr_in, sock_addr, sizeof(struct sockaddr_in));
        inet_ntop(AF_INET, &addr_in.sin_addr, ip_str, max_size);
        return ip_str;
    case AF_INET6:
        memcpy(&addr_in6, sock_addr, sizeof(struct sockaddr_in6));
        inet_ntop(AF_INET6, &addr_in6.sin6_addr, ip_str, max_size);
        return ip_str;
    default:
        return "Invalid address family";
    }
}
static char* sockaddr_get_port_str(const struct sockaddr_storage *sock_addr,
                                   char *port_str, size_t max_size)
{
    struct sockaddr_in  addr_in;
    struct sockaddr_in6 addr_in6;
    switch (sock_addr->ss_family) {
    case AF_INET:
        memcpy(&addr_in, sock_addr, sizeof(struct sockaddr_in));
        snprintf(port_str, max_size, "%d", ntohs(addr_in.sin_port));
        return port_str;
    case AF_INET6:
        memcpy(&addr_in6, sock_addr, sizeof(struct sockaddr_in6));
        snprintf(port_str, max_size, "%d", ntohs(addr_in6.sin6_port));
        return port_str;
    default:
        return "Invalid address family";
    }
}
static int client_server_communication(ucp_worker_h worker, ucp_ep_h ep,
                                       send_recv_type_t send_recv_type,
                                       int is_server, int current_iter)
{
    int ret;
    switch (send_recv_type) {
    case CLIENT_SERVER_SEND_RECV_STREAM:
        /* Client-Server communication via Stream API */
        ret = send_recv_stream(worker, ep, is_server, current_iter);
        break;
    case CLIENT_SERVER_SEND_RECV_TAG:
        /* Client-Server communication via Tag-Matching API */
        ret = send_recv_tag(worker, ep, is_server, current_iter);
        break;
    case CLIENT_SERVER_SEND_RECV_AM:
        /* Client-Server communication via AM API. */
        ret = send_recv_am(worker, ep, is_server, current_iter);
        break;
    default:
        fprintf(stderr, "unknown send-recv type %d\n", send_recv_type);
        return -1;
    }
    return ret;
}

```

```

static int init_worker(ucp_context_h ucp_context, ucp_worker_h *ucp_worker)
{
    ucp_worker_params_t worker_params;
    ucs_status_t status;
    int ret = 0;
    memset(&worker_params, 0, sizeof(worker_params));
    worker_params.field_mask = UCP_WORKER_PARAM_FIELD_THREAD_MODE;
    worker_params.thread_mode = UCS_THREAD_MODE_SINGLE;
    status = ucp_worker_create(ucp_context, &worker_params, ucp_worker);
    if (status != UCS_OK) {
        fprintf(stderr, "failed to ucp_worker_create (%s)\n", ucs_status_string(status));
        ret = -1;
    }
    return ret;
}

static void server_conn_handle_cb(ucp_conn_request_h conn_request, void *arg)
{
    ucx_server_ctx_t *context = arg;
    ucp_conn_request_attr_t attr;
    char ip_str[IP_STRING_LEN];
    char port_str[PORT_STRING_LEN];
    ucs_status_t status;
    attr.field_mask = UCP_CONN_REQUEST_ATTR_FIELD_CLIENT_ADDR;
    status = ucp_conn_request_query(conn_request, &attr);
    if (status == UCS_OK) {
        printf("Server received a connection request from client at address %s:%s\n",
              sockaddr_get_ip_str(&attr.client_address, ip_str, sizeof(ip_str)),
              sockaddr_get_port_str(&attr.client_address, port_str, sizeof(port_str)));
    } else if (status != UCS_ERR_UNSUPPORTED) {
        fprintf(stderr, "failed to query the connection request (%s)\n",
              ucs_status_string(status));
    }
    if (context->conn_request == NULL) {
        context->conn_request = conn_request;
    } else {
        /* The server is already handling a connection request from a client,
         * reject this new one */
        printf("Rejecting a connection request. "
              "Only one client at a time is supported.\n");
        status = ucp_listener_reject(context->listener, conn_request);
        if (status != UCS_OK) {
            fprintf(stderr, "server failed to reject a connection request: (%s)\n",
                  ucs_status_string(status));
        }
    }
}

static ucs_status_t server_create_ep(ucp_worker_h data_worker,
                                     ucp_conn_request_h conn_request,
                                     ucp_ep_h *server_ep)
{
    ucp_ep_params_t ep_params;
    ucs_status_t status;
    /* Server creates an ep to the client on the data worker.
     * This is not the worker the listener was created on.
     * The client side should have initiated the connection, leading
     * to this ep's creation */
    ep_params.field_mask = UCP_EP_PARAM_FIELD_ERR_HANDLER |
                          UCP_EP_PARAM_FIELD_CONN_REQUEST;
    ep_params.conn_request = conn_request;
    ep_params.err_handler.cb = err_cb;
    ep_params.err_handler.arg = NULL;
    status = ucp_ep_create(data_worker, &ep_params, server_ep);
    if (status != UCS_OK) {
        fprintf(stderr, "failed to create an endpoint on the server: (%s)\n",
              ucs_status_string(status));
    }
    return status;
}

static ucs_status_t
start_server(ucp_worker_h ucp_worker, ucx_server_ctx_t *context,
            ucp_listener_h *listener_p, const char *address_str)
{
    struct sockaddr_storage listen_addr;
    ucp_listener_params_t params;
    ucp_listener_attr_t attr;
    ucs_status_t status;
    char ip_str[IP_STRING_LEN];
    char port_str[PORT_STRING_LEN];
    set_sock_addr(address_str, &listen_addr);
    params.field_mask = UCP_LISTENER_PARAM_FIELD_SOCK_ADDR |
                      UCP_LISTENER_PARAM_FIELD_CONN_HANDLER;
    params.sockaddr.addr = (const struct sockaddr*)&listen_addr;
    params.sockaddr.addrlen = sizeof(listen_addr);
    params.conn_handler.cb = server_conn_handle_cb;
    params.conn_handler.arg = context;
    /* Create a listener on the server side to listen on the given address.*/
    status = ucp_listener_create(ucp_worker, &params, listener_p);
}

```



```

if (status != UCS_OK) {
    fprintf(stderr, "failed to listen (%s)\n", ucs_status_string(status));
    goto out;
}
/* Query the created listener to get the port it is listening on. */
attr.field_mask = UCP_LISTENER_ATTR_FIELD_SOCKADDR;
status = ucp_listener_query(*listener_p, &attr);
if (status != UCS_OK) {
    fprintf(stderr, "failed to query the listener (%s)\n",
            ucs_status_string(status));
    ucp_listener_destroy(*listener_p);
    goto out;
}
fprintf(stderr, "server is listening on IP %s port %s\n",
        sockaddr_get_ip_str(&attr.sockaddr, ip_str, IP_STRING_LEN),
        sockaddr_get_port_str(&attr.sockaddr, port_str, PORT_STRING_LEN));
printf("Waiting for connection...\n");
out:
return status;
}
static int client_server_do_work(ucp_worker_h ucp_worker, ucp_ep_h ep,
                                send_recv_type_t send_recv_type, int is_server)
{
    int i, ret = 0;
    for (i = 0; i < num_iterations; i++) {
        ret = client_server_communication(ucp_worker, ep, send_recv_type,
                                         is_server, i);
        if (ret != 0) {
            fprintf(stderr, "%s failed on iteration #%d\n",
                    (is_server ? "server": "client"), i + 1);
            goto out;
        }
    }
out:
return ret;
}
static int run_server(ucp_context_h ucp_context, ucp_worker_h ucp_worker,
                     char *listen_addr, send_recv_type_t send_recv_type)
{
    ucx_server_ctx_t context;
    ucp_worker_h ucp_data_worker;
    ucp_am_handler_param_t param;
    ucp_ep_h server_ep;
    ucs_status_t status;
    int ret;
    /* Create a data worker (to be used for data exchange between the server
     * and the client after the connection between them was established) */
    ret = init_worker(ucp_context, &ucp_data_worker);
    if (ret != 0) {
        goto err;
    }
    if (send_recv_type == CLIENT_SERVER_SEND_RECV_AM) {
        /* Initialize Active Message data handler */
        param.field_mask = UCP_AM_HANDLER_PARAM_FIELD_ID |
                          UCP_AM_HANDLER_PARAM_FIELD_CB |
                          UCP_AM_HANDLER_PARAM_FIELD_ARG;
        param.id = TEST_AM_ID;
        param.cb = ucp_am_data_cb;
        param.arg = ucp_data_worker; /* not used in our callback */
        status = ucp_worker_set_am_recv_handler(ucp_data_worker,
                                               &param);
        if (status != UCS_OK) {
            ret = -1;
            goto err_worker;
        }
    }
    /* Initialize the server's context. */
    context.conn_request = NULL;
    /* Create a listener on the worker created at first. The 'connection
     * worker' - used for connection establishment between client and server.
     * This listener will stay open for listening to incoming connection
     * requests from the client */
    status = start_server(ucp_worker, &context, &context.listener, listen_addr);
    if (status != UCS_OK) {
        ret = -1;
        goto err_worker;
    }
    /* Server is always up listening */
    while (1) {
        /* Wait for the server to receive a connection request from the client.
         * If there are multiple clients for which the server's connection request
         * callback is invoked, i.e. several clients are trying to connect in
         * parallel, the server will handle only the first one and reject the rest */
        while (context.conn_request == NULL) {
            ucp_worker_progress(ucp_worker);
        }
        /* Server creates an ep to the client on the data worker.

```

```

    * This is not the worker the listener was created on.
    * The client side should have initiated the connection, leading
    * to this ep's creation */
status = server_create_ep(ucp_data_worker, context.conn_request,
                        &server_ep);
if (status != UCS_OK) {
    ret = -1;
    goto err_listener;
}
/* The server waits for all the iterations to complete before moving on
 * to the next client */
ret = client_server_do_work(ucp_data_worker, server_ep, send_recv_type,
                            1);
if (ret != 0) {
    goto err_ep;
}
/* Close the endpoint to the client */
ep_close(ucp_data_worker, server_ep, UCP_EP_CLOSE_MODE_FORCE);
/* Reinitialize the server's context to be used for the next client */
context.conn_request = NULL;
printf("Waiting for connection...\n");
}
err_ep:
    ep_close(ucp_data_worker, server_ep, UCP_EP_CLOSE_MODE_FORCE);
err_listener:
    ucp_listener_destroy(context.listener);
err_worker:
    ucp_worker_destroy(ucp_data_worker);
err:
    return ret;
}
static int run_client(ucp_worker_h ucp_worker, char *server_addr,
                    send_recv_type_t send_recv_type)
{
    ucp_ep_h    client_ep;
    ucs_status_t status;
    int         ret;
    status = start_client(ucp_worker, server_addr, &client_ep);
    if (status != UCS_OK) {
        fprintf(stderr, "failed to start client (%s)\n", ucs_status_string(status));
        ret = -1;
        goto out;
    }
    ret = client_server_do_work(ucp_worker, client_ep, send_recv_type, 0);
    /* Close the endpoint to the server */
    ep_close(ucp_worker, client_ep, UCP_EP_CLOSE_MODE_FORCE);
out:
    return ret;
}
static int init_context(ucp_context_h *ucp_context, ucp_worker_h *ucp_worker,
                    send_recv_type_t send_recv_type)
{
    /* UCP objects */
    ucp_params_t ucp_params;
    ucs_status_t status;
    int ret = 0;
    memset(&ucp_params, 0, sizeof(ucp_params));
    /* UCP initialization */
    ucp_params.field_mask = UCP_PARAM_FIELD_FEATURES;
    if (send_recv_type == CLIENT_SERVER_SEND_RECV_STREAM) {
        ucp_params.features = UCP_FEATURE_STREAM;
    } else if (send_recv_type == CLIENT_SERVER_SEND_RECV_TAG) {
        ucp_params.features = UCP_FEATURE_TAG;
    } else {
        ucp_params.features = UCP_FEATURE_AM;
    }
    status = ucp_init(&ucp_params, NULL, ucp_context);
    if (status != UCS_OK) {
        fprintf(stderr, "failed to ucp_init (%s)\n", ucs_status_string(status));
        ret = -1;
        goto err;
    }
    ret = init_worker(*ucp_context, ucp_worker);
    if (ret != 0) {
        goto err_cleanup;
    }
    return ret;
err_cleanup:
    ucp_cleanup(*ucp_context);
err:
    return ret;
}
int main(int argc, char **argv)
{
    send_recv_type_t send_recv_type = CLIENT_SERVER_SEND_RECV_DEFAULT;
    char *server_addr = NULL;
    char *listen_addr = NULL;

```

```

int ret;
/* UCP objects */
ucp_context_h ucp_context;
ucp_worker_h ucp_worker;
ret = parse_cmd(argc, argv, &server_addr, &listen_addr, &send_recv_type);
if (ret != 0) {
    goto err;
}
/* Initialize the UCX required objects */
ret = init_context(&ucp_context, &ucp_worker, send_recv_type);
if (ret != 0) {
    goto err;
}
/* Client-Server initialization */
if (server_addr == NULL) {
    /* Server side */
    ret = run_server(ucp_context, ucp_worker, listen_addr, send_recv_type);
} else {
    /* Client side */
    ret = run_client(ucp_worker, server_addr, send_recv_type);
}
ucp_worker_destroy(ucp_worker);
ucp_cleanup(ucp_context);
err:
return ret;
}

```

## 8.2 ucp\_hello\_world.c

UCP hello world client / server example utility.

```

#ifndef HAVE_CONFIG_H
# define HAVE_CONFIG_H /* Force using config.h, so test would fail if header
                        actually tries to use it */
#endif
/*
 * UCP hello world client / server example utility
 * -----
 *
 * Server side:
 *
 *   ./ucp_hello_world
 *
 * Client side:
 *
 *   ./ucp_hello_world -n <server host name>
 *
 * Notes:
 *
 *   - Client acquires Server UCX address via TCP socket
 *
 * Author:
 *
 *   Ilya Nelkenbaum <ilya@nelkenbaum.com>
 *   Sergey Shalnov <sergeysh@mellanox.com> 7-June-2016
 */
#include "hello_world_util.h"
#include "ucp_util.h"
#include <ucp/api/ucp.h>
#include <sys/socket.h>
#include <sys/types.h>
#include <sys/epoll.h>
#include <netinet/in.h>
#include <assert.h>
#include <netdb.h>
#include <stdio.h>
#include <stdlib.h>
#include <string.h>
#include <unistd.h> /* getopt */
#include <pthread.h> /* pthread_self */
#include <errno.h> /* errno */
#include <time.h>
#include <signal.h> /* raise */
struct msg {
    uint64_t data_len;
};
struct ucx_context {
    int completed;
};
enum ucp_test_mode_t {
    TEST_MODE_PROBE,
    TEST_MODE_WAIT,
    TEST_MODE_EVENTFD
}

```

```

} ucp_test_mode = TEST_MODE_PROBE;
typedef enum {
    FAILURE_MODE_NONE,
    FAILURE_MODE_SEND,      /* fail send operation on server */
    FAILURE_MODE_RECV,      /* fail receive operation on client */
    FAILURE_MODE_KEEPLIVE   /* fail without communication on client */
} failure_mode_t;
static struct err_handling {
    ucp_err_handling_mode_t ucp_err_mode;
    failure_mode_t         failure_mode;
} err_handling_opt;
static ucs_status_t ep_status = UCS_OK;
static uint16_t server_port = 13337;
static sa_family_t ai_family = AF_INET;
static long test_string_length = 16;
static const ucp_tag_t tag = 0x1337a880u;
static const ucp_tag_t tag_mask = UINT64_MAX;
static const char *addr_msg_str = "UCX address message";
static const char *data_msg_str = "UCX data message";
static int print_config = 0;
static ucp_address_t *local_addr;
static ucp_address_t *peer_addr;
static size_t local_addr_len;
static size_t peer_addr_len;
static ucs_status_t parse_cmd(int argc, char * const argv[], char **server_name);
static void set_msg_data_len(struct msg *msg, uint64_t data_len)
{
    mem_type_memcpy(&msg->data_len, &data_len, sizeof(data_len));
}
static void request_init(void *request)
{
    struct ucx_context *context = (struct ucx_context *)request;
    context->completed = 0;
}
static void send_handler(void *request, ucs_status_t status, void *ctx)
{
    struct ucx_context *context = (struct ucx_context *)request;
    const char *str = (const char *)ctx;
    context->completed = 1;
    printf("[0x%x] send handler called for \"%s\" with status %d (%s)\n",
           (unsigned int)pthread_self(), str, status,
           ucs_status_string(status));
}
static void failure_handler(void *arg, ucp_ep_h ep, ucs_status_t status)
{
    ucs_status_t *arg_status = (ucs_status_t *)arg;
    printf("[0x%x] failure handler called with status %d (%s)\n",
           (unsigned int)pthread_self(), status, ucs_status_string(status));
    *arg_status = status;
}
static void recv_handler(void *request, ucs_status_t status,
                        ucp_tag_recv_info_t *info)
{
    struct ucx_context *context = (struct ucx_context *)request;
    context->completed = 1;
    printf("[0x%x] receive handler called with status %d (%s), length %lu\n",
           (unsigned int)pthread_self(), status, ucs_status_string(status),
           info->length);
}
static ucs_status_t ucx_wait(ucp_worker_h ucp_worker, struct ucx_context *request,
                            const char *op_str, const char *data_str)
{
    ucs_status_t status;
    if (UCS_PTR_IS_ERR(request)) {
        status = UCS_PTR_STATUS(request);
    } else if (UCS_PTR_IS_PTR(request)) {
        while (!request->completed) {
            ucp_worker_progress(ucp_worker);
        }
        request->completed = 0;
        status = ucp_request_check_status(request);
        ucp_request_free(request);
    } else {
        status = UCS_OK;
    }
    if (status != UCS_OK) {
        fprintf(stderr, "unable to %s %s (%s)\n", op_str, data_str,
                ucs_status_string(status));
    } else {
        printf("finish to %s %s\n", op_str, data_str);
    }
    return status;
}
static ucs_status_t test_poll_wait(ucp_worker_h ucp_worker)
{
    int err = 0;
    ucs_status_t ret = UCS_ERR_NO_MESSAGE;

```

```

int epoll_fd_local = 0;
int epoll_fd      = 0;
ucs_status_t status;
struct epoll_event ev;
ev.data.u64      = 0;
status = ucp_worker_get_efd(ucp_worker, &epoll_fd);
CHKERR_JUMP(UCS_OK != status, "ucp_worker_get_efd", err);
/* It is recommended to copy original fd */
epoll_fd_local = epoll_create(1);
ev.data.fd = epoll_fd;
ev.events = EPOLLIN;
err = epoll_ctl(epoll_fd_local, EPOLL_CTL_ADD, epoll_fd, &ev);
CHKERR_JUMP(err < 0, "add original socket to the new epoll\n", err_fd);
/* Need to prepare ucp_worker before epoll_wait */
status = ucp_worker_arm(ucp_worker);
if (status == UCS_ERR_BUSY) { /* some events are arrived already */
    ret = UCS_OK;
    goto err_fd;
}
CHKERR_JUMP(status != UCS_OK, "ucp_worker_arm\n", err_fd);
do {
    err = epoll_wait(epoll_fd_local, &ev, 1, -1);
} while ((err == -1) && (errno == EINTR));
ret = UCS_OK;
err_fd:
close(epoll_fd_local);
err:
return ret;
}
static void ep_close_err_mode(ucp_worker_h ucp_worker, ucp_ep_h ucp_ep)
{
    uint64_t ep_close_flags;
    if (err_handling_opt.ucp_err_mode == UCP_ERR_HANDLING_MODE_PEER) {
        ep_close_flags = UCP_EP_CLOSE_FLAG_FORCE;
    } else {
        ep_close_flags = 0;
    }
    ep_close(ucp_worker, ucp_ep, ep_close_flags);
}
static int run_ucx_client(ucp_worker_h ucp_worker)
{
    struct msg *msg = NULL;
    size_t msg_len = 0;
    int ret = -1;
    ucp_request_param_t send_param;
    ucp_tag_recv_info_t info_tag;
    ucp_tag_message_h msg_tag;
    ucs_status_t status;
    ucp_ep_h server_ep;
    ucp_ep_params_t ep_params;
    struct ucx_context *request;
    char *str;
    /* Send client UCX address to server */
    ep_params.field_mask = UCP_EP_PARAM_FIELD_REMOTE_ADDRESS |
                          UCP_EP_PARAM_FIELD_ERR_HANDLING_MODE |
                          UCP_EP_PARAM_FIELD_ERR_HANDLER |
                          UCP_EP_PARAM_FIELD_USER_DATA;
    ep_params.address = peer_addr;
    ep_params.err_mode = err_handling_opt.ucp_err_mode;
    ep_params.err_handler.cb = failure_handler;
    ep_params.err_handler.arg = NULL;
    ep_params.user_data = &ep_status;
    status = ucp_ep_create(ucp_worker, &ep_params, &server_ep);
    CHKERR_JUMP(status != UCS_OK, "ucp_ep_create\n", err);
    msg_len = sizeof(*msg) + local_addr_len;
    msg = malloc(msg_len);
    CHKERR_JUMP(msg == NULL, "allocate memory\n", err_ep);
    memset(msg, 0, msg_len);
    msg->data_len = local_addr_len;
    memcpy(msg + 1, local_addr, local_addr_len);
    send_param.op_attr_mask = UCP_OP_ATTR_FIELD_CALLBACK |
                             UCP_OP_ATTR_FIELD_USER_DATA;
    send_param.cb.send = send_handler;
    send_param.user_data = (void*)addr_msg_str;
    request = ucp_tag_send_nbx(server_ep, msg, msg_len, tag,
                              &send_param);
    status = ucx_wait(ucp_worker, request, "send",
                     addr_msg_str);
    if (status != UCS_OK) {
        free(msg);
        goto err_ep;
    }
    free(msg);
    if (err_handling_opt.failure_mode == FAILURE_MODE_RECV) {
        fprintf(stderr, "Emulating failure before receive operation on client side\n");
        raise(SIGKILL);
    }
}

```

```

/* Receive test string from server */
for (;;) {
    CHKERR_JUMP(ep_status != UCS_OK, "receive data: EP disconnected\n", err_ep);
    /* Probing incoming events in non-block mode */
    msg_tag = ucp_tag_probe_nb(ucp_worker, tag, tag_mask, 1, &info_tag);
    if (msg_tag != NULL) {
        /* Message arrived */
        break;
    } else if (ucp_worker_progress(ucp_worker)) {
        /* Some events were polled; try again without going to sleep */
        continue;
    }
    /* If we got here, ucp_worker_progress() returned 0, so we can sleep.
     * Following blocked methods used to polling internal file descriptor
     * to make CPU idle and don't spin loop
     */
    if (ucp_test_mode == TEST_MODE_WAIT) {
        /* Polling incoming events*/
        status = ucp_worker_wait(ucp_worker);
        CHKERR_JUMP(status != UCS_OK, "ucp_worker_wait\n", err_ep);
    } else if (ucp_test_mode == TEST_MODE_EVENTFD) {
        status = test_poll_wait(ucp_worker);
        CHKERR_JUMP(status != UCS_OK, "test_poll_wait\n", err_ep);
    }
}
if (err_handling_opt.failure_mode == FAILURE_MODE_KEEPAALIVE) {
    fprintf(stderr, "Emulating unexpected failure after receive completion "
                "on client side, server should detect error by "
                "keepalive mechanism\n");
    raise(SIGKILL);
}
msg = mem_type_malloc(info_tag.length);
CHKERR_JUMP(msg == NULL, "allocate memory\n", err_ep);
request = ucp_tag_msg_rcv_nb(ucp_worker, msg, info_tag.length,
                            ucp_dt_make_contig(1), msg_tag,
                            rcv_handler);
status = ucx_wait(ucp_worker, request, "receive", data_msg_str);
if (status != UCS_OK) {
    mem_type_free(msg);
    goto err_ep;
}
str = calloc(1, test_string_length);
if (str == NULL) {
    fprintf(stderr, "Memory allocation failed\n");
    ret = -1;
    goto err_msg;
}
mem_type_memcpy(str, msg + 1, test_string_length);
printf("\n\n---- UCP TEST SUCCESS ----\n\n");
printf("%s", str);
printf("\n\n-----\n\n");
free(str);
ret = 0;
err_msg:
    mem_type_free(msg);
err_ep:
    ep_close_err_mode(ucp_worker, server_ep);
err:
    return ret;
}
static void flush_callback(void *request, ucs_status_t status, void *user_data)
{
}
static ucs_status_t flush_ep(ucp_worker_h worker, ucp_ep_h ep)
{
    ucp_request_param_t param;
    void *request;
    param.op_attr_mask = UCP_OP_ATTR_FIELD_CALLBACK;
    param.cb.send      = flush_callback;
    request            = ucp_ep_flush_nbx(ep, &param);
    if (request == NULL) {
        return UCS_OK;
    } else if (UCS_PTR_IS_ERR(request)) {
        return UCS_PTR_STATUS(request);
    } else {
        ucs_status_t status;
        do {
            ucp_worker_progress(worker);
            status = ucp_request_check_status(request);
        } while (status == UCS_INPROGRESS);
        ucp_request_free(request);
        return status;
    }
}
static int run_ucx_server(ucp_worker_h ucp_worker)
{
    struct msg *msg          = NULL;

```

```

struct ucx_context *request = NULL;
size_t msg_len = 0;
ucp_request_param_t send_param;
ucp_tag_recv_info_t info_tag;
ucp_tag_message_h msg_tag;
ucs_status_t status;
ucp_ep_h client_ep;
ucp_ep_params_t ep_params;
int ret;
/* Receive client UCX address */
do {
    /* Progressing before probe to update the state */
    ucp_worker_progress(ucp_worker);
    /* Probing incoming events in non-block mode */
    msg_tag = ucp_tag_probe_nb(ucp_worker, tag, tag_mask, 1, &info_tag);
} while (msg_tag == NULL);
msg = malloc(info_tag.length);
CHKERR_ACTION(msg == NULL, "allocate memory\n", ret = -1; goto err);
request = ucp_tag_msg_recv_nb(ucp_worker, msg, info_tag.length,
                             ucp_dt_make_contig(1), msg_tag, recv_handler);
status = ucx_wait(ucp_worker, request, "receive", addr_msg_str);
if (status != UCS_OK) {
    free(msg);
    ret = -1;
    goto err;
}
if (err_handling_opt.failure_mode == FAILURE_MODE_SEND) {
    fprintf(stderr, "Emulating unexpected failure on server side, client "
                  "should detect error by keepalive mechanism\n");
    free(msg);
    raise(SIGKILL);
    exit(1);
}
peer_addr_len = msg->data_len;
peer_addr = malloc(peer_addr_len);
if (peer_addr == NULL) {
    fprintf(stderr, "unable to allocate memory for peer address\n");
    free(msg);
    ret = -1;
    goto err;
}
memcpy(peer_addr, msg + 1, peer_addr_len);
free(msg);
/* Send test string to client */
ep_params.field_mask = UCP_EP_PARAM_FIELD_REMOTE_ADDRESS |
                      UCP_EP_PARAM_FIELD_ERR_HANDLING_MODE |
                      UCP_EP_PARAM_FIELD_ERR_HANDLER |
                      UCP_EP_PARAM_FIELD_USER_DATA;
ep_params.address = peer_addr;
ep_params.err_mode = err_handling_opt.ucp_err_mode;
ep_params.err_handler.cb = failure_handler;
ep_params.err_handler.arg = NULL;
ep_params.user_data = &ep_status;
status = ucp_ep_create(ucp_worker, &ep_params, &client_ep);
/* If peer failure testing was requested, it could be possible that UCP EP
 * couldn't be created; in this case set 'ret = 0' to report success */
ret = (err_handling_opt.failure_mode != FAILURE_MODE_NONE) ? 0 : -1;
CHKERR_ACTION(status != UCS_OK, "ucp_ep_create\n", goto err);
msg_len = sizeof(*msg) + test_string_length;
msg = mem_type_malloc(msg_len);
CHKERR_ACTION(msg == NULL, "allocate memory\n", ret = -1; goto err_ep);
mem_type_memset(msg, 0, msg_len);
set_msg_data_len(msg, msg_len - sizeof(*msg));
ret = generate_test_string((char *) (msg + 1), test_string_length);
CHKERR_JUMP(ret < 0, "generate test string", err_free_mem_type_msg);
if (err_handling_opt.failure_mode == FAILURE_MODE_RECV) {
    /* Sleep for small amount of time to ensure that client was killed
     * and peer failure handling is covered */
    sleep(5);
}
ucp_worker_progress(ucp_worker);
send_param.op_attr_mask = UCP_OP_ATTR_FIELD_CALLBACK |
                         UCP_OP_ATTR_FIELD_USER_DATA |
                         UCP_OP_ATTR_FIELD_MEMORY_TYPE;
send_param.cb.send = send_handler;
send_param.user_data = (void*)data_msg_str;
send_param.memory_type = test_mem_type;
request = ucp_tag_send_nbx(client_ep, msg, msg_len, tag,
                          &send_param);
status = ucx_wait(ucp_worker, request, "send",
                  data_msg_str);
if (status != UCS_OK) {
    if (err_handling_opt.failure_mode != FAILURE_MODE_NONE) {
        ret = -1;
    } else {
        /* If peer failure testing was requested, set 'ret = 0' to report
         * success from the application */
    }
}

```

```

        ret = 0;
        /* Make sure that failure_handler was called */
        while (ep_status == UCS_OK) {
            ucp_worker_progress(ucp_worker);
        }
    }
    goto err_free_mem_type_msg;
}
if (err_handling_opt.failure_mode == FAILURE_MODE_KEEPAALIVE) {
    fprintf(stderr, "Waiting for client is terminated\n");
    while (ep_status == UCS_OK) {
        ucp_worker_progress(ucp_worker);
    }
}
status = flush_ep(ucp_worker, client_ep);
printf("flush_ep completed with status %d (%s)\n",
       status, ucs_status_string(status));
ret = 0;
err_free_mem_type_msg:
    mem_type_free(msg);
err_ep:
    ep_close_err_mode(ucp_worker, client_ep);
err:
    return ret;
}
static int run_test(const char *client_target_name, ucp_worker_h ucp_worker)
{
    if (client_target_name != NULL) {
        return run_ucx_client(ucp_worker);
    } else {
        return run_ucx_server(ucp_worker);
    }
}
static void progress_worker(void *arg)
{
    ucp_worker_progress((ucp_worker_h)arg);
}
int main(int argc, char **argv)
{
    /* UCP temporary vars */
    ucp_params_t ucp_params;
    ucp_worker_params_t worker_params;
    ucp_config_t *config;
    ucs_status_t status;
    /* UCP handler objects */
    ucp_context_h ucp_context;
    ucp_worker_h ucp_worker;
    /* OOB connection vars */
    uint64_t addr_len = 0;
    char *client_target_name = NULL;
    int oob_sock = -1;
    int ret = -1;
    memset(&ucp_params, 0, sizeof(ucp_params));
    memset(&worker_params, 0, sizeof(worker_params));
    /* Parse the command line */
    status = parse_cmd(argc, argv, &client_target_name);
    CHKERR_JUMP(status != UCS_OK, "parse_cmd\n", err);
    /* UCP initialization */
    status = ucp_config_read(NULL, NULL, &config);
    CHKERR_JUMP(status != UCS_OK, "ucp_config_read\n", err);
    ucp_params.field_mask = UCP_PARAM_FIELD_FEATURES |
                          UCP_PARAM_FIELD_REQUEST_SIZE |
                          UCP_PARAM_FIELD_REQUEST_INIT;
    ucp_params.features = UCP_FEATURE_TAG;
    if (ucp_test_mode == TEST_MODE_WAIT || ucp_test_mode == TEST_MODE_EVENTFD) {
        ucp_params.features |= UCP_FEATURE_WAKEUP;
    }
    ucp_params.request_size = sizeof(struct ucx_context);
    ucp_params.request_init = request_init;
    status = ucp_init(&ucp_params, config, &ucp_context);
    if (print_config) {
        ucp_config_print(config, stdout, NULL, UCS_CONFIG_PRINT_CONFIG);
    }
    ucp_config_release(config);
    CHKERR_JUMP(status != UCS_OK, "ucp_init\n", err);
    worker_params.field_mask = UCP_WORKER_PARAM_FIELD_THREAD_MODE;
    worker_params.thread_mode = UCS_THREAD_MODE_SINGLE;
    status = ucp_worker_create(ucp_context, &worker_params, &ucp_worker);
    CHKERR_JUMP(status != UCS_OK, "ucp_worker_create\n", err_cleanup);
    status = ucp_worker_get_address(ucp_worker, &local_addr, &local_addr_len);
    CHKERR_JUMP(status != UCS_OK, "ucp_worker_get_address\n", err_worker);
    printf("[0x%x] local address length: %lu\n",
           (unsigned int)pthread_self(), local_addr_len);
    /* OOB connection establishment */
    if (client_target_name) {
        peer_addr_len = local_addr_len;
        oob_sock = connect_common(client_target_name, server_port, ai_family);
    }
}

```



```

    CHKERR_JUMP(oob_sock < 0, "client_connect\n", err_addr);
    ret = recv(oob_sock, &addr_len, sizeof(addr_len), MSG_WAITALL);
    CHKERR_JUMP_RETVAL(ret != (int)sizeof(addr_len),
        "receive address length\n", err_addr, ret);
    peer_addr_len = addr_len;
    peer_addr = malloc(peer_addr_len);
    CHKERR_JUMP(!peer_addr, "allocate memory\n", err_addr);
    ret = recv(oob_sock, peer_addr, peer_addr_len, MSG_WAITALL);
    CHKERR_JUMP_RETVAL(ret != (int)peer_addr_len,
        "receive address\n", err_peer_addr, ret);
} else {
    oob_sock = connect_common(NULL, server_port, ai_family);
    CHKERR_JUMP(oob_sock < 0, "server_connect\n", err_peer_addr);
    addr_len = local_addr_len;
    ret = send(oob_sock, &addr_len, sizeof(addr_len), 0);
    CHKERR_JUMP_RETVAL(ret != (int)sizeof(addr_len),
        "send address length\n", err_peer_addr, ret);
    ret = send(oob_sock, local_addr, local_addr_len, 0);
    CHKERR_JUMP_RETVAL(ret != (int)local_addr_len, "send address\n",
        err_peer_addr, ret);
}
ret = run_test(client_target_name, ucp_worker);
if (!ret && (err_handling_opt.failure_mode == FAILURE_MODE_NONE)) {
    /* Make sure remote is disconnected before destroying local worker */
    ret = barrier(oob_sock, progress_worker, ucp_worker);
}
close(oob_sock);
err_peer_addr:
    free(peer_addr);
err_addr:
    ucp_worker_release_address(ucp_worker, local_addr);
err_worker:
    ucp_worker_destroy(ucp_worker);
err_cleanup:
    ucp_cleanup(ucp_context);
err:
    return ret;
}
static void print_usage()
{
    fprintf(stderr, "Usage: ucp_hello_world [parameters]\n");
    fprintf(stderr, "UCP hello world client/server example utility\n");
    fprintf(stderr, "\nParameters are:\n");
    fprintf(stderr, "  -w      Select test mode \"wait\" to test "
        "ucp_worker_wait function\n");
    fprintf(stderr, "  -f      Select test mode \"event fd\" to test "
        "ucp_worker_get_efd function with later poll\n");
    fprintf(stderr, "  -b      Select test mode \"busy polling\" to test "
        "ucp_tag_probe_nb and ucp_worker_progress (default)\n");
    fprintf(stderr, "  -n <name> Set node name or IP address "
        "of the server (required for client and should be ignored "
        "for server)\n");
    fprintf(stderr, "  -e <type> Emulate unexpected failure and handle an "
        "error with enabled UCP_ERR_HANDLING_MODE_PEER\n");
    fprintf(stderr, "      send      - send failure on server side "
        "before send initiated\n");
    fprintf(stderr, "      recv      - receive failure on client side "
        "before receive completed\n");
    fprintf(stderr, "      keepalive - keepalive failure on client side "
        "after communication completed\n");
    fprintf(stderr, "  -c      Print UCP configuration\n");
    print_common_help();
    fprintf(stderr, "\n");
}
ucs_status_t parse_cmd(int argc, char * const argv[], char **server_name)
{
    int c = 0, idx = 0;
    err_handling_opt.ucp_err_mode = UCP_ERR_HANDLING_MODE_NONE;
    err_handling_opt.failure_mode = FAILURE_MODE_NONE;
    while ((c = getopt(argc, argv, "wfben:ps:m:ch")) != -1) {
        switch (c) {
            case 'w':
                ucp_test_mode = TEST_MODE_WAIT;
                break;
            case 'f':
                ucp_test_mode = TEST_MODE_EVENTFD;
                break;
            case 'b':
                ucp_test_mode = TEST_MODE_PROBE;
                break;
            case 'e':
                err_handling_opt.ucp_err_mode = UCP_ERR_HANDLING_MODE_PEER;
                if (!strcmp(optarg, "recv")) {
                    err_handling_opt.failure_mode = FAILURE_MODE_RECV;
                } else if (!strcmp(optarg, "send")) {
                    err_handling_opt.failure_mode = FAILURE_MODE_SEND;
                } else if (!strcmp(optarg, "keepalive")) {

```

```

        err_handling_opt.failure_mode = FAILURE_MODE_KEEPLIVE;
    } else {
        print_usage();
        return UCS_ERR_UNSUPPORTED;
    }
    break;
case 'n':
    *server_name = optarg;
    break;
case '6':
    ai_family = AF_INET6;
    break;
case 'p':
    server_port = atoi(optarg);
    if (server_port <= 0) {
        fprintf(stderr, "Wrong server port number %d\n", server_port);
        return UCS_ERR_UNSUPPORTED;
    }
    break;
case 's':
    test_string_length = atol(optarg);
    if (test_string_length < 0) {
        fprintf(stderr, "Wrong string size %ld\n", test_string_length);
        return UCS_ERR_UNSUPPORTED;
    }
    break;
case 'm':
    test_mem_type = parse_mem_type(optarg);
    if (test_mem_type == UCS_MEMORY_TYPE_LAST) {
        return UCS_ERR_UNSUPPORTED;
    }
    break;
case 'c':
    print_config = 1;
    break;
case 'h':
    print_usage();
    return UCS_ERR_UNSUPPORTED;
}
}
fprintf(stderr, "INFO: UCP_HELLO_WORLD mode = %d server = %s port = %d, pid = %d\n",
        ucp_test_mode, *server_name, server_port, getpid());
for (idx = optind; idx < argc; idx++) {
    fprintf(stderr, "WARNING: Non-option argument %s\n", argv[idx]);
}
return UCS_OK;
}

```

### 8.3 uct\_hello\_world.c

UCT hello world client / server example utility.

```

#include "hello_world_util.h"
#include <limits.h>
#include <uct/api/uct.h>
#include <assert.h>
#include <inttypes.h>
typedef enum {
    FUNC_AM_SHORT,
    FUNC_AM_BCOPY,
    FUNC_AM_ZCOPY
} func_am_t;
typedef struct {
    int is_uct_desc;
} rcv_desc_t;
typedef struct {
    char                *server_name;
    uint16_t            server_port;
    sa_family_t         ai_family;
    func_am_t           func_am_type;
    const char          *dev_name;
    const char          *tl_name;
    long                test_strlen;
} cmd_args_t;
typedef struct {
    uct_iface_attr_t    iface_attr; /* Interface attributes: capabilities and limitations */
    uct_iface_h         iface;      /* Communication interface context */
    uct_md_attr_t       md_attr;    /* Memory domain attributes: capabilities and limitations */
    uct_md_h            md;         /* Memory domain */
    uct_worker_h        worker;     /* Workers represent allocated resources in a communication thread */
} iface_info_t;
/* Helper data type for am_short */
typedef struct {

```

```

    uint64_t      header;
    char          *payload;
    size_t       len;
} am_short_args_t;
/* Helper data type for am_bcopy */
typedef struct {
    char          *data;
    size_t       len;
} am_bcopy_args_t;
/* Helper data type for am_zcopy */
typedef struct {
    uct_completion_t uct_comp;
    uct_md_h         md;
    uct_mem_h        memh;
} zcopy_comp_t;
static void* desc_holder = NULL;
int print_err_usage(void);
static char *func_am_t_str(func_am_t func_am_type)
{
    switch (func_am_type) {
    case FUNC_AM_SHORT:
        return "uct_ep_am_short";
    case FUNC_AM_BCOPY:
        return "uct_ep_am_bcopy";
    case FUNC_AM_ZCOPY:
        return "uct_ep_am_zcopy";
    }
    return NULL;
}
static size_t func_am_max_size(func_am_t func_am_type,
                              const uct_iface_attr_t *attr)
{
    switch (func_am_type) {
    case FUNC_AM_SHORT:
        return attr->cap.am.max_short;
    case FUNC_AM_BCOPY:
        return attr->cap.am.max_bcopy;
    case FUNC_AM_ZCOPY:
        return attr->cap.am.max_zcopy;
    }
    return 0;
}
/* Helper function for am_short */
void am_short_params_pack(char *buf, size_t len, am_short_args_t *args)
{
    args->header = *(uint64_t *)buf;
    if (len > sizeof(args->header)) {
        args->payload = (buf + sizeof(args->header));
        args->len = len - sizeof(args->header);
    } else {
        args->payload = NULL;
        args->len = 0;
    }
}
ucs_status_t do_am_short(uct_iface_info_t *if_info, uct_ep_h ep, uint8_t id,
                        const cmd_args_t *cmd_args, char *buf)
{
    ucs_status_t status;
    am_short_args_t send_args;
    am_short_params_pack(buf, cmd_args->test_strlen, &send_args);
    do {
        /* Send active message to remote endpoint */
        status = uct_ep_am_short(ep, id, send_args.header, send_args.payload,
                                send_args.len);
        uct_worker_progress(if_info->worker);
    } while (status == UCS_ERR_NO_RESOURCE);
    return status;
}
/* Pack callback for am_bcopy */
size_t am_bcopy_data_pack_cb(void *dest, void *arg)
{
    am_bcopy_args_t *bc_args = arg;
    mem_type_memcpy(dest, bc_args->data, bc_args->len);
    return bc_args->len;
}
ucs_status_t do_am_bcopy(uct_iface_info_t *if_info, uct_ep_h ep, uint8_t id,
                        const cmd_args_t *cmd_args, char *buf)
{
    am_bcopy_args_t args;
    ssize_t len;
    args.data = buf;
    args.len = cmd_args->test_strlen;
    /* Send active message to remote endpoint */
    do {
        len = uct_ep_am_bcopy(ep, id, am_bcopy_data_pack_cb, &args, 0);
        uct_worker_progress(if_info->worker);
    } while (len == UCS_ERR_NO_RESOURCE);
}

```

```

    /* Negative len is an error code */
    return (len >= 0) ? UCS_OK : (ucs_status_t)len;
}
/* Completion callback for am_zcopy */
void zcopy_completion_cb(uct_completion_t *self)
{
    zcopy_comp_t *comp = (zcopy_comp_t *)self;
    assert((comp->uct_comp.count == 0) && (self->status == UCS_OK));
    if (comp->memh != UCT_MEM_HANDLE_NULL) {
        uct_md_mem_dereg(comp->md, comp->memh);
    }
    desc_holder = (void *)0xDEADBEEF;
}
ucs_status_t do_am_zcopy(iface_info_t *if_info, uct_ep_h ep, uint8_t id,
                        const cmd_args_t *cmd_args, char *buf)
{
    ucs_status_t status = UCS_OK;
    uct_mem_h memh;
    uct_iov_t iov;
    zcopy_comp_t comp;
    if (if_info->md_attr.cap.flags & UCT_MD_FLAG_NEED_MEMH) {
        status = uct_md_mem_reg(if_info->md, buf, cmd_args->test_strlen,
                                UCT_MD_MEM_ACCESS_RMA, &memh);
    } else {
        memh = UCT_MEM_HANDLE_NULL;
    }
    iov.buffer = buf;
    iov.length = cmd_args->test_strlen;
    iov.memh = memh;
    iov.stride = 0;
    iov.count = 1;
    comp.uct_comp.func = zcopy_completion_cb;
    comp.uct_comp.count = 1;
    comp.uct_comp.status = UCS_OK;
    comp.md = if_info->md;
    comp.memh = memh;
    if (status == UCS_OK) {
        do {
            status = uct_ep_am_zcopy(ep, id, NULL, 0, &iov, 1, 0,
                                     (uct_completion_t *)&comp);
            uct_worker_progress(if_info->worker);
        } while (status == UCS_ERR_NO_RESOURCE);
        if (status == UCS_INPROGRESS) {
            while (!desc_holder) {
                /* Explicitly progress outstanding active message request */
                uct_worker_progress(if_info->worker);
            }
            status = UCS_OK;
        }
    }
    return status;
}
static void print_strings(const char *label, const char *local_str,
                        const char *remote_str, size_t length)
{
    fprintf(stdout, "\n\n---- UCT TEST SUCCESS ----\n\n");
    fprintf(stdout, "[%s] %s sent %s (%# PRIu64 " bytes)", label, local_str,
            (length != 0) ? remote_str : "<none>", length);
    fprintf(stdout, "\n\n-----\n\n");
    fflush(stdout);
}
/* Callback to handle receive active message */
static ucs_status_t hello_world(void *arg, void *data, size_t length,
                                unsigned flags)
{
    func_am_t func_am_type = *(func_am_t *)arg;
    recv_desc_t *rdesc;
    print_strings("callback", func_am_t_str(func_am_type), data, length);
    if (flags & UCT_CB_PARAM_FLAG_DESC) {
        rdesc = (recv_desc_t *)data - 1;
        /* Hold descriptor to release later and return UCS_INPROGRESS */
        rdesc->is_uct_desc = 1;
        desc_holder = rdesc;
        return UCS_INPROGRESS;
    }
    /* We need to copy-out data and return UCS_OK if want to use the data
     * outside the callback */
    rdesc = malloc(sizeof(*rdesc) + length);
    CHKERR_ACTION(rdesc == NULL, "allocate memory\n", return UCS_ERR_NO_MEMORY);
    rdesc->is_uct_desc = 0;
    memcpy(rdesc + 1, data, length);
    desc_holder = rdesc;
    return UCS_OK;
}
/* Init the transport by its name */
static ucs_status_t init_iface(char *dev_name, char *tl_name,
                               func_am_t func_am_type,

```

```

        iface_info_t *iface_p)
{
    ucs_status_t      status;
    uct_iface_config_t *config; /* Defines interface configuration options */
    uct_iface_params_t params;
    params.field_mask = UCT_IFACE_PARAM_FIELD_OPEN_MODE |
                       UCT_IFACE_PARAM_FIELD_DEVICE |
                       UCT_IFACE_PARAM_FIELD_STATS_ROOT |
                       UCT_IFACE_PARAM_FIELD_RX_HEADROOM |
                       UCT_IFACE_PARAM_FIELD_CPU_MASK;
    params.open_mode = UCT_IFACE_OPEN_MODE_DEVICE;
    params.mode.device.tl_name = tl_name;
    params.mode.device.dev_name = dev_name;
    params.stats_root = NULL;
    params.rx_headroom = sizeof(recv_desc_t);
    UCS_CPU_ZERO(&params.cpu_mask);
    /* Read transport-specific interface configuration */
    status = uct_md_iface_config_read(iface_p->md, tl_name, NULL, NULL, &config);
    CHKERR_JUMP(UCS_OK != status, "setup iface_config", error_ret);
    /* Open communication interface */
    assert(iface_p->iface == NULL);
    status = uct_iface_open(iface_p->md, iface_p->worker, &params, config,
                           &iface_p->iface);
    uct_config_release(config);
    CHKERR_JUMP(UCS_OK != status, "open temporary interface", error_ret);
    /* Enable progress on the interface */
    uct_iface_progress_enable(iface_p->iface,
                             UCT_PROGRESS_SEND | UCT_PROGRESS_RECV);
    /* Get interface attributes */
    status = uct_iface_query(iface_p->iface, &iface_p->iface_attr);
    CHKERR_JUMP(UCS_OK != status, "query iface", error_iface);
    /* Check if current device and transport support required active messages */
    if ((func_am_type == FUNC_AM_SHORT) &&
        (iface_p->iface_attr.cap.flags & UCT_IFACE_FLAG_AM_SHORT)) {
        if (test_mem_type != UCS_MEMORY_TYPE_CUDA) {
            return UCS_OK;
        } else {
            fprintf(stderr, "AM short protocol doesn't support CUDA memory");
        }
    }
    if ((func_am_type == FUNC_AM_BCOPY) &&
        (iface_p->iface_attr.cap.flags & UCT_IFACE_FLAG_AM_BCOPY)) {
        return UCS_OK;
    }
    if ((func_am_type == FUNC_AM_ZCOPY) &&
        (iface_p->iface_attr.cap.flags & UCT_IFACE_FLAG_AM_ZCOPY)) {
        return UCS_OK;
    }
}
error_iface:
    uct_iface_close(iface_p->iface);
    iface_p->iface = NULL;
error_ret:
    return UCS_ERR_UNSUPPORTED;
}
/* Device and transport to be used are determined by minimum latency */
static ucs_status_t dev_tl_lookup(const cmd_args_t *cmd_args,
                                 iface_info_t *iface_p)
{
    uct_tl_resource_desc_t *tl_resources = NULL; /* Communication resource descriptor */
    unsigned num_tl_resources = 0; /* Number of transport resources resource objects
    created */
    uct_component_h *components;
    unsigned num_components;
    unsigned cmpt_index;
    uct_component_attr_t component_attr;
    unsigned md_index;
    unsigned tl_index;
    uct_md_config_t *md_config;
    ucs_status_t status;
    status = uct_query_components(&components, &num_components);
    CHKERR_JUMP(UCS_OK != status, "query for components", error_ret);
    for (cmpt_index = 0; cmpt_index < num_components; ++cmpt_index) {
        component_attr.field_mask = UCT_COMPONENT_ATTR_FIELD_MD_RESOURCE_COUNT;
        status = uct_component_query(components[cmpt_index], &component_attr);
        CHKERR_JUMP(UCS_OK != status, "query component attributes",
                   release_component_list);
        component_attr.field_mask = UCT_COMPONENT_ATTR_FIELD_MD_RESOURCES;
        component_attr.md_resources = alloca(sizeof(*component_attr.md_resources) *
                                             component_attr.md_resource_count);
        status = uct_component_query(components[cmpt_index], &component_attr);
        CHKERR_JUMP(UCS_OK != status, "query for memory domain resources",
                   release_component_list);
        iface_p->iface = NULL;
        /* Iterate through memory domain resources */
        for (md_index = 0; md_index < component_attr.md_resource_count; ++md_index) {
            status = uct_md_config_read(components[cmpt_index], NULL, NULL,
                                       &md_config);

```

```

CHKERR_JUMP(UCS_OK != status, "read MD config",
            release_component_list);
status = uct_md_open(components[cmpt_index],
                    component_attr.md_resources[md_index].md_name,
                    md_config, &iface_p->md);
uct_config_release(md_config);
CHKERR_JUMP(UCS_OK != status, "open memory domains",
            release_component_list);
status = uct_md_query(iface_p->md, &iface_p->md_attr);
CHKERR_JUMP(UCS_OK != status, "query iface",
            close_md);
status = uct_md_query_tl_resources(iface_p->md, &tl_resources,
                                   &num_tl_resources);
CHKERR_JUMP(UCS_OK != status, "query transport resources", close_md);
/* Go through each available transport and find the proper name */
for (tl_index = 0; tl_index < num_tl_resources; ++tl_index) {
    if (!strcmp(cmd_args->dev_name, tl_resources[tl_index].dev_name) &&
        !strcmp(cmd_args->tl_name, tl_resources[tl_index].tl_name)) {
        if ((cmd_args->func_am_type == FUNC_AM_ZCOPY) &&
            !(iface_p->md_attr.cap.reg_mem_types &
              UCS_BIT(test_mem_type))) {
            fprintf(stderr, "Unsupported memory type %s by "
                    "UCT_TL_RESOURCE_DESC_FMT" on %s MD\n",
                    ucs_memory_type_names[test_mem_type],
                    UCT_TL_RESOURCE_DESC_ARG(&tl_resources[tl_index]),
                    component_attr.md_resources[md_index].md_name);
            status = UCS_ERR_UNSUPPORTED;
            break;
        }
        status = init_iface(tl_resources[tl_index].dev_name,
                            tl_resources[tl_index].tl_name,
                            cmd_args->func_am_type, iface_p);
        if (status != UCS_OK) {
            break;
        }
        fprintf(stdout, "Using " UCT_TL_RESOURCE_DESC_FMT "\n",
                UCT_TL_RESOURCE_DESC_ARG(&tl_resources[tl_index]));
        goto release_tl_resources;
    }
}
release_tl_resources:
uct_release_tl_resource_list(tl_resources);
if ((status == UCS_OK) &&
    (tl_index < num_tl_resources)) {
    goto release_component_list;
}
tl_resources = NULL;
num_tl_resources = 0;
uct_md_close(iface_p->md);
}
fprintf(stderr, "No supported (dev/tl) found (%s/%s)\n",
        cmd_args->dev_name, cmd_args->tl_name);
status = UCS_ERR_UNSUPPORTED;
release_component_list:
uct_release_component_list(components);
error_ret:
return status;
close_md:
uct_md_close(iface_p->md);
goto release_component_list;
}
int print_err_usage()
{
    const char func_template[] = " -%c      Select \"%s\" function to send the message%s\n";
    fprintf(stderr, "Usage: uct_hello_world [parameters]\n");
    fprintf(stderr, "UCT hello world client/server example utility\n");
    fprintf(stderr, "\nParameters are:\n");
    fprintf(stderr, func_template, 'i', func_am_t_str(FUNC_AM_SHORT), " (default)");
    fprintf(stderr, func_template, 'b', func_am_t_str(FUNC_AM_BCOPY), "");
    fprintf(stderr, func_template, 'z', func_am_t_str(FUNC_AM_ZCOPY), "");
    fprintf(stderr, " -d      Select device name\n");
    fprintf(stderr, " -t      Select transport layer\n");
    fprintf(stderr, " -n <name> Set node name or IP address "
            "of the server (required for client and should be ignored "
            "for server)\n");
    print_common_help();
    fprintf(stderr, "\nExample:\n");
    fprintf(stderr, "  Server: uct_hello_world -d eth0 -t tcp\n");
    fprintf(stderr, "  Client: uct_hello_world -d eth0 -t tcp -n localhost\n");
    return UCS_ERR_UNSUPPORTED;
}
int parse_cmd(int argc, char * const argv[], cmd_args_t *args)
{
    int c = 0, idx = 0;
    assert(args);
    memset(args, 0, sizeof(*args));

```

```

/* Defaults */
args->server_port = 13337;
args->ai_family = AF_INET;
args->func_am_type = FUNC_AM_SHORT;
args->test_strlen = 16;
while ((c = getopt(argc, argv, "6ibzd:t:n:p:s:m:h")) != -1) {
    switch (c) {
        case 'i':
            args->func_am_type = FUNC_AM_SHORT;
            break;
        case 'b':
            args->func_am_type = FUNC_AM_BCOPY;
            break;
        case 'z':
            args->func_am_type = FUNC_AM_ZCOPY;
            break;
        case 'd':
            args->dev_name = optarg;
            break;
        case 't':
            args->tl_name = optarg;
            break;
        case 'n':
            args->server_name = optarg;
            break;
        case '6':
            args->ai_family = AF_INET6;
            break;
        case 'p':
            args->server_port = atoi(optarg);
            if (args->server_port <= 0) {
                fprintf(stderr, "Wrong server port number %d\n",
                    args->server_port);
                return UCS_ERR_UNSUPPORTED;
            }
            break;
        case 's':
            args->test_strlen = atol(optarg);
            if (args->test_strlen < 0) {
                fprintf(stderr, "Wrong string size %ld\n", args->test_strlen);
                return UCS_ERR_UNSUPPORTED;
            }
            break;
        case 'm':
            test_mem_type = parse_mem_type(optarg);
            if (test_mem_type == UCS_MEMORY_TYPE_LAST) {
                return UCS_ERR_UNSUPPORTED;
            }
            break;
        case 'h':
        default:
            return print_err_usage();
    }
}
fprintf(stderr, "INFO: UCT_HELLO_WORLD AM function = %s server = %s port = %d\n",
    func_am_t_str(args->func_am_type), args->server_name,
    args->server_port);
for (idx = optind; idx < argc; idx++) {
    fprintf(stderr, "WARNING: Non-option argument %s\n", argv[idx]);
}
if (args->dev_name == NULL) {
    fprintf(stderr, "WARNING: device is not set\n");
    return print_err_usage();
}
if (args->tl_name == NULL) {
    fprintf(stderr, "WARNING: transport layer is not set\n");
    return print_err_usage();
}
return UCS_OK;
}
/* The caller is responsible to free *rbuf */
int sendrcv(int sock, const void *sbuf, size_t slen, void **rbuf)
{
    int ret = 0;
    size_t rlen = 0;
    *rbuf = NULL;
    ret = send(sock, &slen, sizeof(slen), 0);
    if ((ret < 0) || (ret != sizeof(slen))) {
        fprintf(stderr, "failed to send buffer length\n");
        return -1;
    }
    ret = send(sock, sbuf, slen, 0);
    if (ret != (int)slen) {
        fprintf(stderr, "failed to send buffer, return value %d\n", ret);
        return -1;
    }
    ret = recv(sock, &rlen, sizeof(rlen), MSG_WAITALL);
}

```

```

if ((ret != sizeof(rlen)) || (rlen > (SIZE_MAX / 2))) {
    fprintf(stderr,
            "failed to receive device address length, return value %d\n",
            ret);
    return -1;
}
*rbuf = calloc(1, rlen);
if (!*rbuf) {
    fprintf(stderr, "failed to allocate receive buffer\n");
    return -1;
}
ret = recv(sock, *rbuf, rlen, MSG_WAITALL);
if (ret != (int)rlen) {
    fprintf(stderr, "failed to receive device address, return value %d\n",
            ret);
    return -1;
}
return 0;
}
static void progress_worker(void *arg)
{
    uct_worker_progress((uct_worker_h)arg);
}
int main(int argc, char **argv)
{
    uct_device_addr_t *peer_dev = NULL;
    uct_iface_addr_t *peer_iface = NULL;
    uct_ep_addr_t *own_ep = NULL;
    uct_ep_addr_t *peer_ep = NULL;
    uint8_t id = 0;
    int oob_sock = -1; /* OOB connection socket */
    ucs_status_t status = UCS_OK; /* status codes for UCS */
    uct_device_addr_t *own_dev;
    uct_iface_addr_t *own_iface;
    uct_ep_h ep; /* Remote endpoint */
    ucs_async_context_t *async; /* Async event context manages
                                times and fd notifications */

    cmd_args_t cmd_args;
    iface_info_t if_info;
    uct_ep_params_t ep_params;
    int res;
    /* Parse the command line */
    if (parse_cmd(argc, argv, &cmd_args)) {
        status = UCS_ERR_INVALID_PARAM;
        goto out;
    }
    /* Initialize context
     * It is better to use different contexts for different workers */
    status = ucs_async_context_create(UCS_ASYNC_MODE_THREAD_SPINLOCK, &async);
    CHKERR_JUMP(UCS_OK != status, "init async context", out);
    /* Create a worker object */
    status = uct_worker_create(async, UCS_THREAD_MODE_SINGLE, &if_info.worker);
    CHKERR_JUMP(UCS_OK != status, "create worker", out_cleanup_async);
    /* Search for the desired transport */
    status = dev_tl_lookup(&cmd_args, &if_info);
    CHKERR_JUMP(UCS_OK != status, "find supported device and transport",
                out_destroy_worker);
    /* Set active message handler */
    status = uct_iface_set_am_handler(if_info.iface, id, hello_world,
                                    &cmd_args.func_am_type, 0);
    CHKERR_JUMP(UCS_OK != status, "set callback", out_destroy_iface);
    own_dev = (uct_device_addr_t*)calloc(1, if_info.iface_attr.device_addr_len);
    CHKERR_JUMP(NULL == own_dev, "allocate memory for dev addr",
                out_destroy_iface);
    own_iface = (uct_iface_addr_t*)calloc(1, if_info.iface_attr.iface_addr_len);
    CHKERR_JUMP(NULL == own_iface, "allocate memory for if addr",
                out_free_dev_addrs);
    oob_sock = connect_common(cmd_args.server_name, cmd_args.server_port,
                             cmd_args.ai_family);
    CHKERR_ACTION(oob_sock < 0, "OOB connect",
                  status = UCS_ERR_IO_ERROR; goto out_close_oob_sock);
    /* Get device address */
    if (if_info.iface_attr.device_addr_len > 0) {
        status = uct_iface_get_device_address(if_info.iface, own_dev);
        CHKERR_JUMP(UCS_OK != status, "get device address", out_free_if_addrs);
        res = sendrecv(oob_sock, own_dev, if_info.iface_attr.device_addr_len,
                      (void*)&peer_dev);
        CHKERR_ACTION(0 != res, "device exchange", status = UCS_ERR_NO_MESSAGE;
                      goto out_close_oob_sock);
    }
    /* Get interface address */
    if (if_info.iface_attr.iface_addr_len > 0) {
        status = uct_iface_get_address(if_info.iface, own_iface);
        CHKERR_JUMP(UCS_OK != status, "get interface address",
                    out_close_oob_sock);
        status = (ucs_status_t)sendrecv(oob_sock, own_iface, if_info.iface_attr.iface_addr_len,
                                       (void *)&peer_iface);
    }
}

```



```

    CHKERR_JUMP(0 != status, "ifaces exchange", out_close_oob_sock);
}
status = (ucs_status_t)uct_iface_is_reachable(if_info.iface, peer_dev,
                                             peer_iface);
CHKERR_JUMP(0 == status, "reach the peer", out_close_oob_sock);
ep_params.field_mask = UCT_EP_PARAM_FIELD_IFACE;
ep_params.iface      = if_info.iface;
if (if_info.iface_attr.cap.flags & UCT_IFACE_FLAG_CONNECT_TO_EP) {
    own_ep = (uct_ep_addr_t*)calloc(1, if_info.iface_attr.ep_addr_len);
    CHKERR_ACTION(NULL == own_ep, "allocate memory for ep addrs",
                  status = UCS_ERR_NO_MEMORY; goto out_close_oob_sock);
    /* Create new endpoint */
    status = uct_ep_create(&ep_params, &ep);
    CHKERR_JUMP(UCS_OK != status, "create endpoint", out_free_ep_adrs);
    /* Get endpoint address */
    status = uct_ep_get_address(ep, own_ep);
    CHKERR_JUMP(UCS_OK != status, "get endpoint address", out_free_ep);
    status = (ucs_status_t)sendrecv(oob_sock, own_ep, if_info.iface_attr.ep_addr_len,
                                   (void **)&peer_ep);
    CHKERR_JUMP(0 != status, "EPs exchange", out_free_ep);
    /* Connect endpoint to a remote endpoint */
    status = uct_ep_connect_to_ep(ep, peer_dev, peer_ep);
    if (barrier(oob_sock, progress_worker, if_info.worker)) {
        status = UCS_ERR_IO_ERROR;
        goto out_free_ep;
    }
} else if (if_info.iface_attr.cap.flags & UCT_IFACE_FLAG_CONNECT_TO_IFACE) {
    /* Create an endpoint which is connected to a remote interface */
    ep_params.field_mask |= UCT_EP_PARAM_FIELD_DEV_ADDR |
                           UCT_EP_PARAM_FIELD_IFACE_ADDR;
    ep_params.dev_addr    = peer_dev;
    ep_params.iface_addr  = peer_iface;
    status = uct_ep_create(&ep_params, &ep);
    CHKERR_JUMP(UCS_OK != status, "create endpoint", out_free_ep_adrs);
} else {
    status = UCS_ERR_UNSUPPORTED;
    goto out_free_ep_adrs;
}
if (cmd_args.test_strlen > func_am_max_size(cmd_args.func_am_type, &if_info.iface_attr)) {
    status = UCS_ERR_UNSUPPORTED;
    fprintf(stderr, "Test string is too long: %ld, max supported: %lu\n",
           cmd_args.test_strlen,
           func_am_max_size(cmd_args.func_am_type, &if_info.iface_attr));
    goto out_free_ep;
}
if (cmd_args.server_name) {
    char *str = (char *)mem_type_malloc(cmd_args.test_strlen);
    CHKERR_ACTION(str == NULL, "allocate memory",
                  status = UCS_ERR_NO_MEMORY; goto out_free_ep);
    res = generate_test_string(str, cmd_args.test_strlen);
    CHKERR_ACTION(res < 0, "generate test string",
                  status = UCS_ERR_NO_MEMORY; goto out_free_ep);
    /* Send active message to remote endpoint */
    if (cmd_args.func_am_type == FUNC_AM_SHORT) {
        status = do_am_short(&if_info, ep, id, &cmd_args, str);
    } else if (cmd_args.func_am_type == FUNC_AM_BCOPY) {
        status = do_am_bcopy(&if_info, ep, id, &cmd_args, str);
    } else if (cmd_args.func_am_type == FUNC_AM_ZCOPY) {
        status = do_am_zcopy(&if_info, ep, id, &cmd_args, str);
    }
    mem_type_free(str);
    CHKERR_JUMP(UCS_OK != status, "send active msg", out_free_ep);
} else {
    recv_desc_t *rdesc;
    while (desc_holder == NULL) {
        /* Explicitly progress any outstanding active message requests */
        uct_worker_progress(if_info.worker);
    }
    rdesc = desc_holder;
    print_strings("main", func_am_t_str(cmd_args.func_am_type),
                 (char *)rdesc + 1, cmd_args.test_strlen);
    if (rdesc->is_uct_desc) {
        /* Release descriptor because callback returns UCS_INPROGRESS */
        uct_iface_release_desc(rdesc);
    } else {
        free(rdesc);
    }
}
if (barrier(oob_sock, progress_worker, if_info.worker)) {
    status = UCS_ERR_IO_ERROR;
}
out_free_ep:
    uct_ep_destroy(ep);
out_free_ep_adrs:
    free(own_ep);
    free(peer_ep);
out_close_oob_sock:

```

```
    close(oob_sock);
out_free_if_addrs:
    free(own_iface);
    free(peer_iface);
out_free_dev_addrs:
    free(own_dev);
    free(peer_dev);
out_destroy_iface:
    uct_iface_close(if_info.iface);
    uct_md_close(if_info.md);
out_destroy_worker:
    uct_worker_destroy(if_info.worker);
out_cleanup_async:
    ucs_async_context_destroy(async);
out:
    return (status == UCS_ERR_UNSUPPORTED) ? UCS_OK : status;
}
```

# Index

- completed\_cb
  - uct\_tag\_context, 248
- finish
  - UCP Data type routines, 132
- pack
  - UCP Data type routines, 132
- packed\_size
  - UCP Data type routines, 131
- priv
  - uct\_tag\_context, 249
- rndv\_cb
  - uct\_tag\_context, 248
- start\_pack
  - UCP Data type routines, 131
- start\_unpack
  - UCP Data type routines, 131
- tag\_consumed\_cb
  - uct\_tag\_context, 248
- UCP Application Context, 12
  - UCP\_ATTR\_FIELD\_MEMORY\_TYPES, 18
  - UCP\_ATTR\_FIELD\_NAME, 18
  - UCP\_ATTR\_FIELD\_REQUEST\_SIZE, 18
  - UCP\_ATTR\_FIELD\_THREAD\_MODE, 18
  - ucp\_cleanup, 20
  - ucp\_context\_attr\_field, 18
  - ucp\_context\_attr\_t, 16
  - ucp\_context\_h, 16
  - ucp\_context\_print\_info, 21
  - ucp\_context\_query, 21
  - UCP\_ENTITY\_NAME\_MAX, 16
  - ucp\_feature, 18
  - UCP\_FEATURE\_AM, 18
  - UCP\_FEATURE\_AMO32, 18
  - UCP\_FEATURE\_AMO64, 18
  - UCP\_FEATURE\_RMA, 18
  - UCP\_FEATURE\_STREAM, 18
  - UCP\_FEATURE\_TAG, 18
  - UCP\_FEATURE\_WAKEUP, 18
  - ucp\_get\_version, 19
  - ucp\_get\_version\_string, 19
  - ucp\_init, 19
  - ucp\_lib\_attr\_field, 18
  - UCP\_LIB\_ATTR\_FIELD\_MAX\_THREAD\_LEVEL, 18
  - ucp\_lib\_attr\_t, 16
  - ucp\_lib\_query, 19
  - UCP\_PARAM\_FIELD\_ESTIMATED\_NUM\_EPS, 17
  - UCP\_PARAM\_FIELD\_ESTIMATED\_NUM\_PPN, 18
  - UCP\_PARAM\_FIELD\_FEATURES, 17
  - UCP\_PARAM\_FIELD\_MT\_WORKERS\_SHARED, 17
  - UCP\_PARAM\_FIELD\_NAME, 18
  - UCP\_PARAM\_FIELD\_REQUEST\_CLEANUP, 17
  - UCP\_PARAM\_FIELD\_REQUEST\_INIT, 17
  - UCP\_PARAM\_FIELD\_REQUEST\_SIZE, 17
  - UCP\_PARAM\_FIELD\_TAG\_SENDER\_MASK, 17
  - ucp\_params\_field, 17
  - ucp\_request\_cleanup\_callback\_t, 17
  - ucp\_request\_init\_callback\_t, 17
  - ucp\_tag\_rcv\_info\_t, 16
- UCP Communication routines, 78
  - ucp\_am\_data\_release, 91
  - ucp\_am\_handler\_param\_field, 88
  - UCP\_AM\_HANDLER\_PARAM\_FIELD\_ARG, 88
  - UCP\_AM\_HANDLER\_PARAM\_FIELD\_CB, 88
  - UCP\_AM\_HANDLER\_PARAM\_FIELD\_FLAGS, 88
  - UCP\_AM\_HANDLER\_PARAM\_FIELD\_ID, 88
  - UCP\_AM\_RECV\_ATTR\_FIELD\_REPLY\_EP, 88
  - UCP\_AM\_RECV\_ATTR\_FLAG\_DATA, 88
  - UCP\_AM\_RECV\_ATTR\_FLAG\_RNDV, 88
  - ucp\_am\_recv\_attr\_t, 88
  - ucp\_am\_recv\_data\_nbx, 90
  - ucp\_am\_recv\_data\_nbx\_callback\_t, 85
  - ucp\_am\_send\_nb, 89
  - ucp\_am\_send\_nbx, 89
  - ucp\_atomic\_add32, 115
  - ucp\_atomic\_add64, 116
  - ucp\_atomic\_cswap32, 120
  - ucp\_atomic\_cswap64, 120
  - ucp\_atomic\_fadd32, 117
  - ucp\_atomic\_fadd64, 117
  - ucp\_atomic\_fetch\_nb, 109
  - UCP\_ATOMIC\_FETCH\_OP\_CSWAP, 86
  - UCP\_ATOMIC\_FETCH\_OP\_FADD, 86
  - UCP\_ATOMIC\_FETCH\_OP\_FAND, 86
  - UCP\_ATOMIC\_FETCH\_OP\_FOR, 86
  - UCP\_ATOMIC\_FETCH\_OP\_FXOR, 86
  - UCP\_ATOMIC\_FETCH\_OP\_LAST, 86
  - UCP\_ATOMIC\_FETCH\_OP\_SWAP, 86
  - ucp\_atomic\_fetch\_op\_t, 86
  - UCP\_ATOMIC\_OP\_ADD, 86
  - UCP\_ATOMIC\_OP\_AND, 86

- UCP\_ATOMIC\_OP\_CSWAP, 86
- UCP\_ATOMIC\_OP\_LAST, 87
- ucp\_atomic\_op\_nbx, 110
- UCP\_ATOMIC\_OP\_OR, 87
- UCP\_ATOMIC\_OP\_SWAP, 86
- ucp\_atomic\_op\_t, 86
- UCP\_ATOMIC\_OP\_XOR, 87
- ucp\_atomic\_post, 109
- UCP\_ATOMIC\_POST\_OP\_ADD, 86
- UCP\_ATOMIC\_POST\_OP\_AND, 86
- UCP\_ATOMIC\_POST\_OP\_LAST, 86
- UCP\_ATOMIC\_POST\_OP\_OR, 86
- ucp\_atomic\_post\_op\_t, 85
- UCP\_ATOMIC\_POST\_OP\_XOR, 86
- ucp\_atomic\_swap32, 118
- ucp\_atomic\_swap64, 119
- ucp\_datatype\_t, 82
- ucp\_err\_handler\_cb\_t, 83
- ucp\_err\_handler\_t, 84
- ucp\_get, 115
- ucp\_get\_nb, 107
- ucp\_get\_nbi, 106
- ucp\_get\_nbx, 108
- UCP\_OP\_ATTR\_FIELD\_CALLBACK, 87
- UCP\_OP\_ATTR\_FIELD\_DATATYPE, 87
- UCP\_OP\_ATTR\_FIELD\_FLAGS, 87
- UCP\_OP\_ATTR\_FIELD\_MEMH, 87
- UCP\_OP\_ATTR\_FIELD\_MEMORY\_TYPE, 87
- UCP\_OP\_ATTR\_FIELD\_RECV\_INFO, 87
- UCP\_OP\_ATTR\_FIELD\_REPLY\_BUFFER, 87
- UCP\_OP\_ATTR\_FIELD\_REQUEST, 87
- UCP\_OP\_ATTR\_FIELD\_USER\_DATA, 87
- UCP\_OP\_ATTR\_FLAG\_FAST\_CMPL, 87
- UCP\_OP\_ATTR\_FLAG\_FORCE\_IMM\_CMPL, 87
- UCP\_OP\_ATTR\_FLAG\_MULTI\_SEND, 87
- UCP\_OP\_ATTR\_FLAG\_NO\_IMM\_CMPL, 87
- ucp\_op\_attr\_t, 87
- ucp\_put, 114
- ucp\_put\_nb, 105
- ucp\_put\_nbi, 104
- ucp\_put\_nbx, 105
- ucp\_req\_attr\_field, 87
- ucp\_request\_alloc, 113
- UCP\_REQUEST\_ATTR\_FIELD\_INFO\_STRING, 88
- UCP\_REQUEST\_ATTR\_FIELD\_INFO\_STRING\_SIZE, 88
- UCP\_REQUEST\_ATTR\_FIELD\_MEM\_TYPE, 88
- UCP\_REQUEST\_ATTR\_FIELD\_STATUS, 88
- ucp\_request\_cancel, 112
- ucp\_request\_check\_status, 111
- ucp\_request\_free, 113
- ucp\_request\_is\_completed, 114
- ucp\_request\_query, 89
- ucp\_send\_callback\_t, 83
- ucp\_send\_nbx\_callback\_t, 83
- ucp\_stream\_data\_release, 113
- ucp\_stream\_recv\_callback\_t, 84
- ucp\_stream\_recv\_data\_nb, 99
- UCP\_STREAM\_RECV\_FLAG\_WAITALL, 87
- ucp\_stream\_recv\_flags\_t, 87
- ucp\_stream\_recv\_nb, 97
- ucp\_stream\_recv\_nbx, 98
- ucp\_stream\_recv\_nbx\_callback\_t, 84
- ucp\_stream\_recv\_request\_test, 112
- ucp\_stream\_send\_nb, 92
- ucp\_stream\_send\_nbx, 92
- ucp\_tag\_message\_h, 82
- ucp\_tag\_msg\_recv\_nb, 103
- ucp\_tag\_msg\_recv\_nbx, 103
- ucp\_tag\_probe\_nb, 102
- ucp\_tag\_recv\_callback\_t, 84
- ucp\_tag\_recv\_nb, 100
- ucp\_tag\_recv\_nbr, 100
- ucp\_tag\_recv\_nbx, 101
- ucp\_tag\_recv\_nbx\_callback\_t, 85
- ucp\_tag\_recv\_request\_test, 112
- ucp\_tag\_send\_nb, 93
- ucp\_tag\_send\_nbr, 94
- ucp\_tag\_send\_nbx, 96
- ucp\_tag\_send\_sync\_nb, 95
- ucp\_tag\_send\_sync\_nbx, 97
- ucp\_tag\_t, 82
- UCP Configuration, 122
  - ucp\_config\_modify, 125
  - ucp\_config\_print, 126
  - ucp\_config\_read, 125
  - ucp\_config\_release, 125
  - ucp\_config\_t, 124
  - ucp\_params\_t, 124
- UCP Data type routines, 127
  - finish, 132
  - pack, 132
  - packed\_size, 131
  - start\_pack, 131
  - start\_unpack, 131
  - UCP\_DATATYPE\_CLASS\_MASK, 130
  - UCP\_DATATYPE\_CONTIG, 130
  - UCP\_DATATYPE\_GENERIC, 130
  - UCP\_DATATYPE\_IOV, 130
  - UCP\_DATATYPE\_SHIFT, 130
  - UCP\_DATATYPE\_STRIDED, 130
  - ucp\_dt\_create\_generic, 130
  - ucp\_dt\_destroy, 130
  - ucp\_dt\_iov\_t, 129
  - ucp\_dt\_make\_contig, 128
  - ucp\_dt\_make\_iov, 129
  - ucp\_dt\_type, 129
  - ucp\_generic\_dt\_ops\_t, 129
  - unpack, 132
- UCP Endpoint, 62
  - ucp\_am\_callback\_t, 67
  - ucp\_am\_recv\_callback\_t, 67
  - UCP\_CB\_PARAM\_FLAG\_DATA, 70
  - ucp\_cb\_param\_flags, 70
  - ucp\_conn\_request\_h, 66

- ucp\_disconnect\_nb, 75
- ucp\_ep\_attr\_field, 71
- UCP\_EP\_ATTR\_FIELD\_LOCAL\_SOCKADDR, 71
- UCP\_EP\_ATTR\_FIELD\_NAME, 71
- UCP\_EP\_ATTR\_FIELD\_REMOTE\_SOCKADDR, 71
- ucp\_ep\_attr\_t, 66
- UCP\_EP\_CLOSE\_FLAG\_FORCE, 69
- ucp\_ep\_close\_flags\_t, 69
- ucp\_ep\_close\_mode, 69
- UCP\_EP\_CLOSE\_MODE\_FLUSH, 70
- UCP\_EP\_CLOSE\_MODE\_FORCE, 70
- ucp\_ep\_close\_nb, 72
- ucp\_ep\_close\_nbx, 72
- ucp\_ep\_create, 71
- ucp\_ep\_destroy, 75
- ucp\_ep\_evaluate\_perf, 74
- ucp\_ep\_flush, 76
- ucp\_ep\_flush\_nb, 73
- ucp\_ep\_flush\_nbx, 74
- ucp\_ep\_h, 66
- ucp\_ep\_modify\_nb, 76
- UCP\_EP\_PARAM\_FIELD\_CONN\_REQUEST, 68
- UCP\_EP\_PARAM\_FIELD\_ERR\_HANDLER, 68
- UCP\_EP\_PARAM\_FIELD\_ERR\_HANDLING\_MODE, 68
- UCP\_EP\_PARAM\_FIELD\_FLAGS, 68
- UCP\_EP\_PARAM\_FIELD\_LOCAL\_SOCK\_ADDR, 68
- UCP\_EP\_PARAM\_FIELD\_NAME, 68
- UCP\_EP\_PARAM\_FIELD\_REMOTE\_ADDRESS, 68
- UCP\_EP\_PARAM\_FIELD\_SOCK\_ADDR, 68
- UCP\_EP\_PARAM\_FIELD\_USER\_DATA, 68
- ucp\_ep\_params\_field, 68
- UCP\_EP\_PARAMS\_FLAGS\_CLIENT\_SERVER, 69
- ucp\_ep\_params\_flags\_field, 69
- UCP\_EP\_PARAMS\_FLAGS\_NO\_LOOPBACK, 69
- UCP\_EP\_PARAMS\_FLAGS\_SEND\_CLIENT\_ID, 69
- ucp\_ep\_params\_t, 68
- ucp\_ep\_perf\_attr\_field, 70
- UCP\_EP\_PERF\_ATTR\_FIELD\_ESTIMATED\_TIME, 70
- ucp\_ep\_perf\_attr\_field\_t, 66
- ucp\_ep\_perf\_param\_field, 70
- UCP\_EP\_PERF\_PARAM\_FIELD\_MESSAGE\_SIZE, 70
- ucp\_ep\_perf\_param\_field\_t, 66
- ucp\_ep\_print\_info, 73
- ucp\_ep\_query, 75
- UCP\_ERR\_HANDLING\_MODE\_NONE, 71
- UCP\_ERR\_HANDLING\_MODE\_PEER, 71
- ucp\_err\_handling\_mode\_t, 71
- ucp\_request\_release, 75
- ucp\_request\_test, 76
- ucp\_stream\_poll\_ep\_t, 66
- UCP Memory routines, 47
  - ucp\_ep\_rkey\_unpack, 56
  - UCP\_MADV\_NORMAL, 52
  - UCP\_MADV\_WILLNEED, 52
  - ucp\_mem\_advise, 51
  - ucp\_mem\_advise\_t, 49
  - ucp\_mem\_advise, 55
  - UCP\_MEM\_ADVISE\_PARAM\_FIELD\_ADDRESS, 51
  - UCP\_MEM\_ADVISE\_PARAM\_FIELD\_ADVICE, 51
  - UCP\_MEM\_ADVISE\_PARAM\_FIELD\_LENGTH, 51
  - ucp\_mem\_advise\_params\_field, 51
  - ucp\_mem\_advise\_params\_t, 50
  - ucp\_mem\_attr\_field, 52
  - UCP\_MEM\_ATTR\_FIELD\_ADDRESS, 52
  - UCP\_MEM\_ATTR\_FIELD\_LENGTH, 52
  - UCP\_MEM\_ATTR\_FIELD\_MEM\_TYPE, 52
  - ucp\_mem\_attr\_t, 50
  - ucp\_mem\_h, 50
  - ucp\_mem\_map, 52
  - UCP\_MEM\_MAP\_ALLOCATE, 51
  - UCP\_MEM\_MAP\_FIXED, 51
  - UCP\_MEM\_MAP\_NONBLOCK, 51
  - UCP\_MEM\_MAP\_PARAM\_FIELD\_ADDRESS, 50
  - UCP\_MEM\_MAP\_PARAM\_FIELD\_FLAGS, 50
  - UCP\_MEM\_MAP\_PARAM\_FIELD\_LENGTH, 50
  - UCP\_MEM\_MAP\_PARAM\_FIELD\_MEMORY\_TYPE, 50
  - UCP\_MEM\_MAP\_PARAM\_FIELD\_PROT, 50
  - ucp\_mem\_map\_params\_field, 50
  - ucp\_mem\_map\_params\_t, 49
  - UCP\_MEM\_MAP\_PROT\_LOCAL\_READ, 51
  - UCP\_MEM\_MAP\_PROT\_LOCAL\_WRITE, 51
  - UCP\_MEM\_MAP\_PROT\_REMOTE\_READ, 51
  - UCP\_MEM\_MAP\_PROT\_REMOTE\_WRITE, 51
  - ucp\_mem\_print\_info, 54
  - ucp\_mem\_query, 54
  - ucp\_mem\_unmap, 53
  - ucp\_rkey\_buffer\_release, 56
  - ucp\_rkey\_destroy, 57
  - ucp\_rkey\_h, 50
  - ucp\_rkey\_pack, 55
  - ucp\_rkey\_ptr, 57
- UCP Wake-up routines, 58
  - ucp\_worker\_arm, 60
  - ucp\_worker\_get\_efd, 58
  - ucp\_worker\_signal, 61
  - ucp\_worker\_wait, 59
  - ucp\_worker\_wait\_mem, 59
- UCP Worker, 22
  - ucp\_address\_t, 30
  - ucp\_am\_cb\_flags, 34
  - UCP\_AM\_FLAG\_PERSISTENT\_DATA, 35
  - UCP\_AM\_FLAG\_WHOLE\_MSG, 35
  - ucp\_am\_handler\_param\_t, 30
  - ucp\_am\_rcv\_param\_t, 30
  - UCP\_AM\_SEND\_FLAG\_EAGER, 35

- UCP\_AM\_SEND\_FLAG\_REPLY, 35
- UCP\_AM\_SEND\_FLAG\_RNDV, 35
- UCP\_AM\_SEND\_REPLY, 35
- ucp\_conn\_request\_attr\_field, 34
- UCP\_CONN\_REQUEST\_ATTR\_FIELD\_CLIENT\_ADDR, 34
- UCP\_CONN\_REQUEST\_ATTR\_FIELD\_CLIENT\_ID, 34
- ucp\_conn\_request\_attr\_t, 30
- ucp\_conn\_request\_query, 42
- ucp\_listener\_accept\_callback\_t, 31
- ucp\_listener\_accept\_handler\_t, 30
- ucp\_listener\_attr\_field, 34
- UCP\_LISTENER\_ATTR\_FIELD\_SOCKADDR, 34
- ucp\_listener\_attr\_t, 30
- ucp\_listener\_conn\_callback\_t, 31
- ucp\_listener\_conn\_handler\_t, 32
- ucp\_listener\_create, 41
- ucp\_listener\_destroy, 41
- ucp\_listener\_h, 30
- UCP\_LISTENER\_PARAM\_FIELD\_ACCEPT\_HANDLER, 33
- UCP\_LISTENER\_PARAM\_FIELD\_CONN\_HANDLER, 33
- UCP\_LISTENER\_PARAM\_FIELD\_SOCK\_ADDR, 33
- ucp\_listener\_params\_field, 33
- ucp\_listener\_params\_t, 30
- ucp\_listener\_query, 42
- ucp\_listener\_reject, 42
- ucp\_send\_am\_flags, 35
- ucp\_stream\_worker\_poll, 40
- UCP\_WAKEUP\_AMO, 35
- UCP\_WAKEUP\_EDGE, 36
- ucp\_wakeup\_event\_t, 32
- ucp\_wakeup\_event\_types, 35
- UCP\_WAKEUP\_RMA, 35
- UCP\_WAKEUP\_RX, 36
- UCP\_WAKEUP\_TAG\_RECV, 35
- UCP\_WAKEUP\_TAG\_SEND, 35
- UCP\_WAKEUP\_TX, 35
- ucp\_worker\_address\_attr\_field, 34
- UCP\_WORKER\_ADDRESS\_ATTR\_FIELD\_UID, 34
- ucp\_worker\_address\_attr\_t, 30
- UCP\_WORKER\_ADDRESS\_FLAG\_NET\_ONLY, 33
- ucp\_worker\_address\_flags\_t, 33
- ucp\_worker\_address\_query, 39
- ucp\_worker\_attr\_field, 33
- UCP\_WORKER\_ATTR\_FIELD\_ADDRESS, 34
- UCP\_WORKER\_ATTR\_FIELD\_ADDRESS\_FLAGS, 34
- UCP\_WORKER\_ATTR\_FIELD\_MAX\_AM\_HEADER, 34
- UCP\_WORKER\_ATTR\_FIELD\_MAX\_INFO\_STRING, 34
- UCP\_WORKER\_ATTR\_FIELD\_NAME, 34
- UCP\_WORKER\_ATTR\_FIELD\_THREAD\_MODE, 34
- ucp\_worker\_attr\_t, 29
- ucp\_worker\_create, 36
- ucp\_worker\_destroy, 36
- ucp\_worker\_fence, 44
- UCP\_WORKER\_FLAG\_IGNORE\_REQUEST\_LEAK, 33
- ucp\_worker\_flags\_t, 33
- ucp\_worker\_flush, 45
- ucp\_worker\_flush\_nb, 44
- ucp\_worker\_flush\_nbx, 45
- ucp\_worker\_get\_address, 38
- ucp\_worker\_h, 31
- UCP\_WORKER\_PARAM\_FIELD\_AM\_ALIGNMENT, 32
- UCP\_WORKER\_PARAM\_FIELD\_CLIENT\_ID, 32
- UCP\_WORKER\_PARAM\_FIELD\_CPU\_MASK, 32
- UCP\_WORKER\_PARAM\_FIELD\_EVENT\_FD, 32
- UCP\_WORKER\_PARAM\_FIELD\_EVENTS, 32
- UCP\_WORKER\_PARAM\_FIELD\_FLAGS, 32
- UCP\_WORKER\_PARAM\_FIELD\_NAME, 32
- UCP\_WORKER\_PARAM\_FIELD\_THREAD\_MODE, 32
- UCP\_WORKER\_PARAM\_FIELD\_USER\_DATA, 32
- ucp\_worker\_params\_field, 32
- ucp\_worker\_params\_t, 29
- ucp\_worker\_print\_info, 38
- ucp\_worker\_progress, 40
- ucp\_worker\_query, 38
- ucp\_worker\_release\_address, 39
- ucp\_worker\_set\_am\_handler, 43
- ucp\_worker\_set\_am\_rcv\_handler, 43
- ucp\_address\_t
  - UCP Worker, 30
- ucp\_am\_callback\_t
  - UCP Endpoint, 67
- ucp\_am\_cb\_flags
  - UCP Worker, 34
- ucp\_am\_data\_release
  - UCP Communication routines, 91
- UCP\_AM\_FLAG\_PERSISTENT\_DATA
  - UCP Worker, 35
- UCP\_AM\_FLAG\_WHOLE\_MSG
  - UCP Worker, 35
- ucp\_am\_handler\_param, 28
- ucp\_am\_handler\_param\_field
  - UCP Communication routines, 88
- UCP\_AM\_HANDLER\_PARAM\_FIELD\_ARG
  - UCP Communication routines, 88
- UCP\_AM\_HANDLER\_PARAM\_FIELD\_CB
  - UCP Communication routines, 88
- UCP\_AM\_HANDLER\_PARAM\_FIELD\_FLAGS
  - UCP Communication routines, 88
- UCP\_AM\_HANDLER\_PARAM\_FIELD\_ID
  - UCP Communication routines, 88
- ucp\_am\_handler\_param\_t
  - UCP Worker, 30

- UCP\_AM\_RECV\_ATTR\_FIELD\_REPLY\_EP
  - UCP Communication routines, 88
- UCP\_AM\_RECV\_ATTR\_FLAG\_DATA
  - UCP Communication routines, 88
- UCP\_AM\_RECV\_ATTR\_FLAG\_RNDV
  - UCP Communication routines, 88
- ucp\_am\_recv\_attr\_t
  - UCP Communication routines, 88
- ucp\_am\_recv\_callback\_t
  - UCP Endpoint, 67
- ucp\_am\_recv\_data\_nbx
  - UCP Communication routines, 90
- ucp\_am\_recv\_data\_nbx\_callback\_t
  - UCP Communication routines, 85
- ucp\_am\_recv\_param, 28
- ucp\_am\_recv\_param\_t
  - UCP Worker, 30
- UCP\_AM\_SEND\_FLAG\_EAGER
  - UCP Worker, 35
- UCP\_AM\_SEND\_FLAG\_REPLY
  - UCP Worker, 35
- UCP\_AM\_SEND\_FLAG\_RNDV
  - UCP Worker, 35
- ucp\_am\_send\_nb
  - UCP Communication routines, 89
- ucp\_am\_send\_nbx
  - UCP Communication routines, 89
- UCP\_AM\_SEND\_REPLY
  - UCP Worker, 35
- ucp\_atomic\_add32
  - UCP Communication routines, 115
- ucp\_atomic\_add64
  - UCP Communication routines, 116
- ucp\_atomic\_cswap32
  - UCP Communication routines, 120
- ucp\_atomic\_cswap64
  - UCP Communication routines, 120
- ucp\_atomic\_fadd32
  - UCP Communication routines, 117
- ucp\_atomic\_fadd64
  - UCP Communication routines, 117
- ucp\_atomic\_fetch\_nb
  - UCP Communication routines, 109
- UCP\_ATOMIC\_FETCH\_OP\_CSWAP
  - UCP Communication routines, 86
- UCP\_ATOMIC\_FETCH\_OP\_FADD
  - UCP Communication routines, 86
- UCP\_ATOMIC\_FETCH\_OP\_FAND
  - UCP Communication routines, 86
- UCP\_ATOMIC\_FETCH\_OP\_FOR
  - UCP Communication routines, 86
- UCP\_ATOMIC\_FETCH\_OP\_FXOR
  - UCP Communication routines, 86
- UCP\_ATOMIC\_FETCH\_OP\_LAST
  - UCP Communication routines, 86
- UCP\_ATOMIC\_FETCH\_OP\_SWAP
  - UCP Communication routines, 86
- ucp\_atomic\_fetch\_op\_t
  - UCP Communication routines, 86
- UCP Communication routines, 86
- UCP\_ATOMIC\_OP\_ADD
  - UCP Communication routines, 86
- UCP\_ATOMIC\_OP\_AND
  - UCP Communication routines, 86
- UCP\_ATOMIC\_OP\_CSWAP
  - UCP Communication routines, 86
- UCP\_ATOMIC\_OP\_LAST
  - UCP Communication routines, 87
- ucp\_atomic\_op\_nbx
  - UCP Communication routines, 110
- UCP\_ATOMIC\_OP\_OR
  - UCP Communication routines, 87
- UCP\_ATOMIC\_OP\_SWAP
  - UCP Communication routines, 86
- ucp\_atomic\_op\_t
  - UCP Communication routines, 86
- UCP\_ATOMIC\_OP\_XOR
  - UCP Communication routines, 87
- ucp\_atomic\_post
  - UCP Communication routines, 109
- UCP\_ATOMIC\_POST\_OP\_ADD
  - UCP Communication routines, 86
- UCP\_ATOMIC\_POST\_OP\_AND
  - UCP Communication routines, 86
- UCP\_ATOMIC\_POST\_OP\_LAST
  - UCP Communication routines, 86
- UCP\_ATOMIC\_POST\_OP\_OR
  - UCP Communication routines, 86
- ucp\_atomic\_post\_op\_t
  - UCP Communication routines, 85
- UCP\_ATOMIC\_POST\_OP\_XOR
  - UCP Communication routines, 86
- ucp\_atomic\_swap32
  - UCP Communication routines, 118
- ucp\_atomic\_swap64
  - UCP Communication routines, 119
- UCP\_ATTR\_FIELD\_MEMORY\_TYPES
  - UCP Application Context, 18
- UCP\_ATTR\_FIELD\_NAME
  - UCP Application Context, 18
- UCP\_ATTR\_FIELD\_REQUEST\_SIZE
  - UCP Application Context, 18
- UCP\_ATTR\_FIELD\_THREAD\_MODE
  - UCP Application Context, 18
- UCP\_CB\_PARAM\_FLAG\_DATA
  - UCP Endpoint, 70
- ucp\_cb\_param\_flags
  - UCP Endpoint, 70
- ucp\_cleanup
  - UCP Application Context, 20
- ucp\_config\_modify
  - UCP Configuration, 125
- ucp\_config\_print
  - UCP Configuration, 126
- ucp\_config\_read
  - UCP Configuration, 125
- ucp\_config\_release

- UCP Configuration, 125
- ucp\_config\_t
  - UCP Configuration, 124
- ucp\_conn\_request\_attr, 27
- ucp\_conn\_request\_attr\_field
  - UCP Worker, 34
- UCP\_CONN\_REQUEST\_ATTR\_FIELD\_CLIENT\_ADDR
  - UCP Worker, 34
- UCP\_CONN\_REQUEST\_ATTR\_FIELD\_CLIENT\_ID
  - UCP Worker, 34
- ucp\_conn\_request\_attr\_t
  - UCP Worker, 30
- ucp\_conn\_request\_h
  - UCP Endpoint, 66
- ucp\_conn\_request\_query
  - UCP Worker, 42
- ucp\_context\_attr, 13
- ucp\_context\_attr\_field
  - UCP Application Context, 18
- ucp\_context\_attr\_t
  - UCP Application Context, 16
- ucp\_context\_h
  - UCP Application Context, 16
- ucp\_context\_print\_info
  - UCP Application Context, 21
- ucp\_context\_query
  - UCP Application Context, 21
- UCP\_DATATYPE\_CLASS\_MASK
  - UCP Data type routines, 130
- UCP\_DATATYPE\_CONTIG
  - UCP Data type routines, 130
- UCP\_DATATYPE\_GENERIC
  - UCP Data type routines, 130
- UCP\_DATATYPE\_IOV
  - UCP Data type routines, 130
- UCP\_DATATYPE\_SHIFT
  - UCP Data type routines, 130
- UCP\_DATATYPE\_STRIDED
  - UCP Data type routines, 130
- ucp\_datatype\_t
  - UCP Communication routines, 82
- ucp\_disconnect\_nb
  - UCP Endpoint, 75
- ucp\_dt\_create\_generic
  - UCP Data type routines, 130
- ucp\_dt\_destroy
  - UCP Data type routines, 130
- ucp\_dt\_iov, 128
- ucp\_dt\_iov\_t
  - UCP Data type routines, 129
- ucp\_dt\_make\_contig
  - UCP Data type routines, 128
- ucp\_dt\_make\_iov
  - UCP Data type routines, 129
- ucp\_dt\_type
  - UCP Data type routines, 129
- UCP\_ENTITY\_NAME\_MAX
  - UCP Application Context, 16
- ucp\_ep\_attr, 64
- ucp\_ep\_attr\_field
  - UCP Endpoint, 71
- UCP\_EP\_ATTR\_FIELD\_LOCAL\_SOCKADDR
  - UCP Endpoint, 71
- UCP\_EP\_ATTR\_FIELD\_NAME
  - UCP Endpoint, 71
- UCP\_EP\_ATTR\_FIELD\_REMOTE\_SOCKADDR
  - UCP Endpoint, 71
- ucp\_ep\_attr\_t
  - UCP Endpoint, 66
- UCP\_EP\_CLOSE\_FLAG\_FORCE
  - UCP Endpoint, 69
- ucp\_ep\_close\_flags\_t
  - UCP Endpoint, 69
- ucp\_ep\_close\_mode
  - UCP Endpoint, 69
- UCP\_EP\_CLOSE\_MODE\_FLUSH
  - UCP Endpoint, 70
- UCP\_EP\_CLOSE\_MODE\_FORCE
  - UCP Endpoint, 70
- ucp\_ep\_close\_nb
  - UCP Endpoint, 72
- ucp\_ep\_close\_nbx
  - UCP Endpoint, 72
- ucp\_ep\_create
  - UCP Endpoint, 71
- ucp\_ep\_destroy
  - UCP Endpoint, 75
- ucp\_ep\_evaluate\_perf
  - UCP Endpoint, 74
- ucp\_ep\_evaluate\_perf\_attr\_t, 64
- ucp\_ep\_evaluate\_perf\_param\_t, 64
- ucp\_ep\_flush
  - UCP Endpoint, 76
- ucp\_ep\_flush\_nb
  - UCP Endpoint, 73
- ucp\_ep\_flush\_nbx
  - UCP Endpoint, 74
- ucp\_ep\_h
  - UCP Endpoint, 66
- ucp\_ep\_modify\_nb
  - UCP Endpoint, 76
- UCP\_EP\_PARAM\_FIELD\_CONN\_REQUEST
  - UCP Endpoint, 68
- UCP\_EP\_PARAM\_FIELD\_ERR\_HANDLER
  - UCP Endpoint, 68
- UCP\_EP\_PARAM\_FIELD\_ERR\_HANDLING\_MODE
  - UCP Endpoint, 68
- UCP\_EP\_PARAM\_FIELD\_FLAGS
  - UCP Endpoint, 68
- UCP\_EP\_PARAM\_FIELD\_LOCAL\_SOCK\_ADDR
  - UCP Endpoint, 68
- UCP\_EP\_PARAM\_FIELD\_NAME
  - UCP Endpoint, 68
- UCP\_EP\_PARAM\_FIELD\_REMOTE\_ADDRESS
  - UCP Endpoint, 68
- UCP\_EP\_PARAM\_FIELD\_SOCK\_ADDR



- UCP Endpoint, 68
- UCP\_EP\_PARAM\_FIELD\_USER\_DATA
  - UCP Endpoint, 68
- ucp\_ep\_params, 65
- ucp\_ep\_params\_field
  - UCP Endpoint, 68
- UCP\_EP\_PARAMS\_FLAGS\_CLIENT\_SERVER
  - UCP Endpoint, 69
- ucp\_ep\_params\_flags\_field
  - UCP Endpoint, 69
- UCP\_EP\_PARAMS\_FLAGS\_NO\_LOOPBACK
  - UCP Endpoint, 69
- UCP\_EP\_PARAMS\_FLAGS\_SEND\_CLIENT\_ID
  - UCP Endpoint, 69
- ucp\_ep\_params\_t
  - UCP Endpoint, 68
- ucp\_ep\_perf\_attr\_field
  - UCP Endpoint, 70
- UCP\_EP\_PERF\_ATTR\_FIELD\_ESTIMATED\_TIME
  - UCP Endpoint, 70
- ucp\_ep\_perf\_attr\_field\_t
  - UCP Endpoint, 66
- ucp\_ep\_perf\_param\_field
  - UCP Endpoint, 70
- UCP\_EP\_PERF\_PARAM\_FIELD\_MESSAGE\_SIZE
  - UCP Endpoint, 70
- ucp\_ep\_perf\_param\_field\_t
  - UCP Endpoint, 66
- ucp\_ep\_print\_info
  - UCP Endpoint, 73
- ucp\_ep\_query
  - UCP Endpoint, 75
- ucp\_ep\_rkey\_unpack
  - UCP Memory routines, 56
- ucp\_err\_handler, 82
- ucp\_err\_handler\_cb\_t
  - UCP Communication routines, 83
- ucp\_err\_handler\_t
  - UCP Communication routines, 84
- UCP\_ERR\_HANDLING\_MODE\_NONE
  - UCP Endpoint, 71
- UCP\_ERR\_HANDLING\_MODE\_PEER
  - UCP Endpoint, 71
- ucp\_err\_handling\_mode\_t
  - UCP Endpoint, 71
- ucp\_feature
  - UCP Application Context, 18
- UCP\_FEATURE\_AM
  - UCP Application Context, 18
- UCP\_FEATURE\_AMO32
  - UCP Application Context, 18
- UCP\_FEATURE\_AMO64
  - UCP Application Context, 18
- UCP\_FEATURE\_RMA
  - UCP Application Context, 18
- UCP\_FEATURE\_STREAM
  - UCP Application Context, 18
- UCP\_FEATURE\_TAG
  - UCP Application Context, 18
- UCP Application Context, 18
- UCP\_FEATURE\_WAKEUP
  - UCP Application Context, 18
- ucp\_generic\_dt\_ops, 247
- ucp\_generic\_dt\_ops\_t
  - UCP Data type routines, 129
- ucp\_get
  - UCP Communication routines, 115
- ucp\_get\_nb
  - UCP Communication routines, 107
- ucp\_get\_nbi
  - UCP Communication routines, 106
- ucp\_get\_nbx
  - UCP Communication routines, 108
- ucp\_get\_version
  - UCP Application Context, 19
- ucp\_get\_version\_string
  - UCP Application Context, 19
- ucp\_init
  - UCP Application Context, 19
- ucp\_lib\_attr, 13
- ucp\_lib\_attr\_field
  - UCP Application Context, 18
- UCP\_LIB\_ATTR\_FIELD\_MAX\_THREAD\_LEVEL
  - UCP Application Context, 18
- ucp\_lib\_attr\_t
  - UCP Application Context, 16
- ucp\_lib\_query
  - UCP Application Context, 19
- ucp\_listener\_accept\_callback\_t
  - UCP Worker, 31
- ucp\_listener\_accept\_handler, 29
- ucp\_listener\_accept\_handler\_t
  - UCP Worker, 30
- ucp\_listener\_attr, 27
- ucp\_listener\_attr\_field
  - UCP Worker, 34
- UCP\_LISTENER\_ATTR\_FIELD\_SOCKADDR
  - UCP Worker, 34
- ucp\_listener\_attr\_t
  - UCP Worker, 30
- ucp\_listener\_conn\_callback\_t
  - UCP Worker, 31
- ucp\_listener\_conn\_handler, 29
- ucp\_listener\_conn\_handler\_t
  - UCP Worker, 32
- ucp\_listener\_create
  - UCP Worker, 41
- ucp\_listener\_destroy
  - UCP Worker, 41
- ucp\_listener\_h
  - UCP Worker, 30
- UCP\_LISTENER\_PARAM\_FIELD\_ACCEPT\_HANDLER
  - UCP Worker, 33
- UCP\_LISTENER\_PARAM\_FIELD\_CONN\_HANDLER
  - UCP Worker, 33
- UCP\_LISTENER\_PARAM\_FIELD\_SOCK\_ADDR
  - UCP Worker, 33

- ucp\_listener\_params, [27](#)
- ucp\_listener\_params\_field
  - UCP Worker, [33](#)
- ucp\_listener\_params\_t
  - UCP Worker, [30](#)
- ucp\_listener\_query
  - UCP Worker, [42](#)
- ucp\_listener\_reject
  - UCP Worker, [42](#)
- UCP\_MADV\_NORMAL
  - UCP Memory routines, [52](#)
- UCP\_MADV\_WILLNEED
  - UCP Memory routines, [52](#)
- ucp\_mem\_advice
  - UCP Memory routines, [51](#)
- ucp\_mem\_advice\_t
  - UCP Memory routines, [49](#)
- ucp\_mem\_advise
  - UCP Memory routines, [55](#)
- UCP\_MEM\_ADVISE\_PARAM\_FIELD\_ADDRESS
  - UCP Memory routines, [51](#)
- UCP\_MEM\_ADVISE\_PARAM\_FIELD\_ADVICE
  - UCP Memory routines, [51](#)
- UCP\_MEM\_ADVISE\_PARAM\_FIELD\_LENGTH
  - UCP Memory routines, [51](#)
- ucp\_mem\_advise\_params, [49](#)
- ucp\_mem\_advise\_params\_field
  - UCP Memory routines, [51](#)
- ucp\_mem\_advise\_params\_t
  - UCP Memory routines, [50](#)
- ucp\_mem\_attr, [49](#)
- ucp\_mem\_attr\_field
  - UCP Memory routines, [52](#)
- UCP\_MEM\_ATTR\_FIELD\_ADDRESS
  - UCP Memory routines, [52](#)
- UCP\_MEM\_ATTR\_FIELD\_LENGTH
  - UCP Memory routines, [52](#)
- UCP\_MEM\_ATTR\_FIELD\_MEM\_TYPE
  - UCP Memory routines, [52](#)
- ucp\_mem\_attr\_t
  - UCP Memory routines, [50](#)
- ucp\_mem\_h
  - UCP Memory routines, [50](#)
- ucp\_mem\_map
  - UCP Memory routines, [52](#)
- UCP\_MEM\_MAP\_ALLOCATE
  - UCP Memory routines, [51](#)
- UCP\_MEM\_MAP\_FIXED
  - UCP Memory routines, [51](#)
- UCP\_MEM\_MAP\_NONBLOCK
  - UCP Memory routines, [51](#)
- UCP\_MEM\_MAP\_PARAM\_FIELD\_ADDRESS
  - UCP Memory routines, [50](#)
- UCP\_MEM\_MAP\_PARAM\_FIELD\_FLAGS
  - UCP Memory routines, [50](#)
- UCP\_MEM\_MAP\_PARAM\_FIELD\_LENGTH
  - UCP Memory routines, [50](#)
- UCP\_MEM\_MAP\_PARAM\_FIELD\_MEMORY\_TYPE
  - UCP Memory routines, [50](#)
- UCP\_MEM\_MAP\_PARAM\_FIELD\_PROT
  - UCP Memory routines, [50](#)
- ucp\_mem\_map\_params, [48](#)
- ucp\_mem\_map\_params\_field
  - UCP Memory routines, [50](#)
- ucp\_mem\_map\_params\_t
  - UCP Memory routines, [49](#)
- UCP\_MEM\_MAP\_PROT\_LOCAL\_READ
  - UCP Memory routines, [51](#)
- UCP\_MEM\_MAP\_PROT\_LOCAL\_WRITE
  - UCP Memory routines, [51](#)
- UCP\_MEM\_MAP\_PROT\_REMOTE\_READ
  - UCP Memory routines, [51](#)
- UCP\_MEM\_MAP\_PROT\_REMOTE\_WRITE
  - UCP Memory routines, [51](#)
- ucp\_mem\_print\_info
  - UCP Memory routines, [54](#)
- ucp\_mem\_query
  - UCP Memory routines, [54](#)
- ucp\_mem\_unmap
  - UCP Memory routines, [53](#)
- UCP\_OP\_ATTR\_FIELD\_CALLBACK
  - UCP Communication routines, [87](#)
- UCP\_OP\_ATTR\_FIELD\_DATATYPE
  - UCP Communication routines, [87](#)
- UCP\_OP\_ATTR\_FIELD\_FLAGS
  - UCP Communication routines, [87](#)
- UCP\_OP\_ATTR\_FIELD\_MEMH
  - UCP Communication routines, [87](#)
- UCP\_OP\_ATTR\_FIELD\_MEMORY\_TYPE
  - UCP Communication routines, [87](#)
- UCP\_OP\_ATTR\_FIELD\_RECV\_INFO
  - UCP Communication routines, [87](#)
- UCP\_OP\_ATTR\_FIELD\_REPLY\_BUFFER
  - UCP Communication routines, [87](#)
- UCP\_OP\_ATTR\_FIELD\_REQUEST
  - UCP Communication routines, [87](#)
- UCP\_OP\_ATTR\_FIELD\_USER\_DATA
  - UCP Communication routines, [87](#)
- UCP\_OP\_ATTR\_FLAG\_FAST\_CMPL
  - UCP Communication routines, [87](#)
- UCP\_OP\_ATTR\_FLAG\_FORCE\_IMM\_CMPL
  - UCP Communication routines, [87](#)
- UCP\_OP\_ATTR\_FLAG\_MULTI\_SEND
  - UCP Communication routines, [87](#)
- UCP\_OP\_ATTR\_FLAG\_NO\_IMM\_CMPL
  - UCP Communication routines, [87](#)
- ucp\_op\_attr\_t
  - UCP Communication routines, [87](#)
- UCP\_PARAM\_FIELD\_ESTIMATED\_NUM\_EPS
  - UCP Application Context, [17](#)
- UCP\_PARAM\_FIELD\_ESTIMATED\_NUM\_PPN
  - UCP Application Context, [18](#)
- UCP\_PARAM\_FIELD\_FEATURES
  - UCP Application Context, [17](#)
- UCP\_PARAM\_FIELD\_MT\_WORKERS\_SHARED
  - UCP Application Context, [17](#)

- UCP\_PARAM\_FIELD\_NAME
  - UCP Application Context, [18](#)
- UCP\_PARAM\_FIELD\_REQUEST\_CLEANUP
  - UCP Application Context, [17](#)
- UCP\_PARAM\_FIELD\_REQUEST\_INIT
  - UCP Application Context, [17](#)
- UCP\_PARAM\_FIELD\_REQUEST\_SIZE
  - UCP Application Context, [17](#)
- UCP\_PARAM\_FIELD\_TAG\_SENDER\_MASK
  - UCP Application Context, [17](#)
- ucp\_params, [122](#)
- ucp\_params\_field
  - UCP Application Context, [17](#)
- ucp\_params\_t
  - UCP Configuration, [124](#)
- ucp\_put
  - UCP Communication routines, [114](#)
- ucp\_put\_nb
  - UCP Communication routines, [105](#)
- ucp\_put\_nbi
  - UCP Communication routines, [104](#)
- ucp\_put\_nbx
  - UCP Communication routines, [105](#)
- ucp\_req\_attr\_field
  - UCP Communication routines, [87](#)
- ucp\_request\_alloc
  - UCP Communication routines, [113](#)
- UCP\_REQUEST\_ATTR\_FIELD\_INFO\_STRING
  - UCP Communication routines, [88](#)
- UCP\_REQUEST\_ATTR\_FIELD\_INFO\_STRING\_SIZE
  - UCP Communication routines, [88](#)
- UCP\_REQUEST\_ATTR\_FIELD\_MEM\_TYPE
  - UCP Communication routines, [88](#)
- UCP\_REQUEST\_ATTR\_FIELD\_STATUS
  - UCP Communication routines, [88](#)
- ucp\_request\_attr\_t, [82](#)
- ucp\_request\_cancel
  - UCP Communication routines, [112](#)
- ucp\_request\_check\_status
  - UCP Communication routines, [111](#)
- ucp\_request\_cleanup\_callback\_t
  - UCP Application Context, [17](#)
- ucp\_request\_free
  - UCP Communication routines, [113](#)
- ucp\_request\_init\_callback\_t
  - UCP Application Context, [17](#)
- ucp\_request\_is\_completed
  - UCP Communication routines, [114](#)
- ucp\_request\_param\_t, [14](#)
- ucp\_request\_param\_t.cb, [15](#)
- ucp\_request\_param\_t.recv\_info, [16](#)
- ucp\_request\_query
  - UCP Communication routines, [89](#)
- ucp\_request\_release
  - UCP Endpoint, [75](#)
- ucp\_request\_test
  - UCP Endpoint, [76](#)
- ucp\_rkey\_buffer\_release
  - UCP Memory routines, [56](#)
- ucp\_rkey\_destroy
  - UCP Memory routines, [57](#)
- ucp\_rkey\_h
  - UCP Memory routines, [50](#)
- ucp\_rkey\_pack
  - UCP Memory routines, [55](#)
- ucp\_rkey\_ptr
  - UCP Memory routines, [57](#)
- ucp\_send\_am\_flags
  - UCP Worker, [35](#)
- ucp\_send\_callback\_t
  - UCP Communication routines, [83](#)
- ucp\_send\_nbx\_callback\_t
  - UCP Communication routines, [83](#)
- ucp\_stream\_data\_release
  - UCP Communication routines, [113](#)
- ucp\_stream\_poll\_ep, [64](#)
- ucp\_stream\_poll\_ep\_t
  - UCP Endpoint, [66](#)
- ucp\_stream\_rcv\_callback\_t
  - UCP Communication routines, [84](#)
- ucp\_stream\_rcv\_data\_nb
  - UCP Communication routines, [99](#)
- UCP\_STREAM\_RECV\_FLAG\_WAITALL
  - UCP Communication routines, [87](#)
- ucp\_stream\_rcv\_flags\_t
  - UCP Communication routines, [87](#)
- ucp\_stream\_rcv\_nb
  - UCP Communication routines, [97](#)
- ucp\_stream\_rcv\_nbx
  - UCP Communication routines, [98](#)
- ucp\_stream\_rcv\_nbx\_callback\_t
  - UCP Communication routines, [84](#)
- ucp\_stream\_rcv\_request\_test
  - UCP Communication routines, [112](#)
- ucp\_stream\_send\_nb
  - UCP Communication routines, [92](#)
- ucp\_stream\_send\_nbx
  - UCP Communication routines, [92](#)
- ucp\_stream\_worker\_poll
  - UCP Worker, [40](#)
- ucp\_tag\_message\_h
  - UCP Communication routines, [82](#)
- ucp\_tag\_msg\_rcv\_nb
  - UCP Communication routines, [103](#)
- ucp\_tag\_msg\_rcv\_nbx
  - UCP Communication routines, [103](#)
- ucp\_tag\_probe\_nb
  - UCP Communication routines, [102](#)
- ucp\_tag\_rcv\_callback\_t
  - UCP Communication routines, [84](#)
- ucp\_tag\_rcv\_info, [14](#)
- ucp\_tag\_rcv\_info\_t
  - UCP Application Context, [16](#)
- ucp\_tag\_rcv\_nb
  - UCP Communication routines, [100](#)
- ucp\_tag\_rcv\_nbr

- UCP Communication routines, 100
- ucp\_tag\_recv\_nbx
  - UCP Communication routines, 101
- ucp\_tag\_recv\_nbx\_callback\_t
  - UCP Communication routines, 85
- ucp\_tag\_recv\_request\_test
  - UCP Communication routines, 112
- ucp\_tag\_send\_nb
  - UCP Communication routines, 93
- ucp\_tag\_send\_nbr
  - UCP Communication routines, 94
- ucp\_tag\_send\_nbx
  - UCP Communication routines, 96
- ucp\_tag\_send\_sync\_nb
  - UCP Communication routines, 95
- ucp\_tag\_send\_sync\_nbx
  - UCP Communication routines, 97
- ucp\_tag\_t
  - UCP Communication routines, 82
- UCP\_WAKEUP\_AMO
  - UCP Worker, 35
- UCP\_WAKEUP\_EDGE
  - UCP Worker, 36
- ucp\_wakeup\_event\_t
  - UCP Worker, 32
- ucp\_wakeup\_event\_types
  - UCP Worker, 35
- UCP\_WAKEUP\_RMA
  - UCP Worker, 35
- UCP\_WAKEUP\_RX
  - UCP Worker, 36
- UCP\_WAKEUP\_TAG\_RECV
  - UCP Worker, 35
- UCP\_WAKEUP\_TAG\_SEND
  - UCP Worker, 35
- UCP\_WAKEUP\_TX
  - UCP Worker, 35
- ucp\_worker\_address\_attr, 26
- ucp\_worker\_address\_attr\_field
  - UCP Worker, 34
- UCP\_WORKER\_ADDRESS\_ATTR\_FIELD\_UID
  - UCP Worker, 34
- ucp\_worker\_address\_attr\_t
  - UCP Worker, 30
- UCP\_WORKER\_ADDRESS\_FLAG\_NET\_ONLY
  - UCP Worker, 33
- ucp\_worker\_address\_flags\_t
  - UCP Worker, 33
- ucp\_worker\_address\_query
  - UCP Worker, 39
- ucp\_worker\_arm
  - UCP Wake-up routines, 60
- ucp\_worker\_attr, 25
- ucp\_worker\_attr\_field
  - UCP Worker, 33
- UCP\_WORKER\_ATTR\_FIELD\_ADDRESS
  - UCP Worker, 34
- UCP\_WORKER\_ATTR\_FIELD\_ADDRESS\_FLAGS
  - UCP Worker, 34
- UCP Worker, 34
- UCP\_WORKER\_ATTR\_FIELD\_MAX\_AM\_HEADER
  - UCP Worker, 34
- UCP\_WORKER\_ATTR\_FIELD\_MAX\_INFO\_STRING
  - UCP Worker, 34
- UCP\_WORKER\_ATTR\_FIELD\_NAME
  - UCP Worker, 34
- UCP\_WORKER\_ATTR\_FIELD\_THREAD\_MODE
  - UCP Worker, 34
- ucp\_worker\_attr\_t
  - UCP Worker, 29
- ucp\_worker\_create
  - UCP Worker, 36
- ucp\_worker\_destroy
  - UCP Worker, 36
- ucp\_worker\_fence
  - UCP Worker, 44
- UCP\_WORKER\_FLAG\_IGNORE\_REQUEST\_LEAK
  - UCP Worker, 33
- ucp\_worker\_flags\_t
  - UCP Worker, 33
- ucp\_worker\_flush
  - UCP Worker, 45
- ucp\_worker\_flush\_nb
  - UCP Worker, 44
- ucp\_worker\_flush\_nbx
  - UCP Worker, 45
- ucp\_worker\_get\_address
  - UCP Worker, 38
- ucp\_worker\_get\_efd
  - UCP Wake-up routines, 58
- ucp\_worker\_h
  - UCP Worker, 31
- UCP\_WORKER\_PARAM\_FIELD\_AM\_ALIGNMENT
  - UCP Worker, 32
- UCP\_WORKER\_PARAM\_FIELD\_CLIENT\_ID
  - UCP Worker, 32
- UCP\_WORKER\_PARAM\_FIELD\_CPU\_MASK
  - UCP Worker, 32
- UCP\_WORKER\_PARAM\_FIELD\_EVENT\_FD
  - UCP Worker, 32
- UCP\_WORKER\_PARAM\_FIELD\_EVENTS
  - UCP Worker, 32
- UCP\_WORKER\_PARAM\_FIELD\_FLAGS
  - UCP Worker, 32
- UCP\_WORKER\_PARAM\_FIELD\_NAME
  - UCP Worker, 32
- UCP\_WORKER\_PARAM\_FIELD\_THREAD\_MODE
  - UCP Worker, 32
- UCP\_WORKER\_PARAM\_FIELD\_USER\_DATA
  - UCP Worker, 32
- ucp\_worker\_params, 25
- ucp\_worker\_params\_field
  - UCP Worker, 32
- ucp\_worker\_params\_t
  - UCP Worker, 29
- ucp\_worker\_print\_info
  - UCP Worker, 38

- ucp\_worker\_progress
  - UCP Worker, 40
- ucp\_worker\_query
  - UCP Worker, 38
- ucp\_worker\_release\_address
  - UCP Worker, 39
- ucp\_worker\_set\_am\_handler
  - UCP Worker, 43
- ucp\_worker\_set\_am\_recv\_handler
  - UCP Worker, 43
- ucp\_worker\_signal
  - UCP Wake-up routines, 61
- ucp\_worker\_wait
  - UCP Wake-up routines, 59
- ucp\_worker\_wait\_mem
  - UCP Wake-up routines, 59
- UCS Communication Resource, 239
  - ucs\_async\_add\_timer, 244
  - ucs\_async\_context\_create, 245
  - ucs\_async\_context\_destroy, 245
  - ucs\_async\_event\_cb\_t, 240
  - ucs\_async\_modify\_handler, 245
  - ucs\_async\_poll, 246
  - ucs\_async\_remove\_handler, 244
  - ucs\_async\_set\_event\_handler, 243
  - UCS\_CALLBACKQ\_FLAG\_FAST, 241
  - UCS\_CALLBACKQ\_FLAG\_ONESHOT, 241
  - ucs\_callbackq\_flags, 241
  - UCS\_ERR\_ALREADY\_EXISTS, 242
  - UCS\_ERR\_BUFFER\_TOO\_SMALL, 242
  - UCS\_ERR\_BUSY, 242
  - UCS\_ERR\_CANCELED, 242
  - UCS\_ERR\_CONNECTION\_RESET, 242
  - UCS\_ERR\_ENDPOINT\_TIMEOUT, 242
  - UCS\_ERR\_EXCEEDS\_LIMIT, 242
  - UCS\_ERR\_FIRST\_ENDPOINT\_FAILURE, 242
  - UCS\_ERR\_FIRST\_LINK\_FAILURE, 242
  - UCS\_ERR\_INVALID\_ADDR, 242
  - UCS\_ERR\_INVALID\_PARAM, 242
  - UCS\_ERR\_IO\_ERROR, 242
  - UCS\_ERR\_LAST, 242
  - UCS\_ERR\_LAST\_ENDPOINT\_FAILURE, 242
  - UCS\_ERR\_LAST\_LINK\_FAILURE, 242
  - UCS\_ERR\_MESSAGE\_TRUNCATED, 242
  - UCS\_ERR\_NO\_DEVICE, 242
  - UCS\_ERR\_NO\_ELEM, 242
  - UCS\_ERR\_NO\_MEMORY, 242
  - UCS\_ERR\_NO\_MESSAGE, 242
  - UCS\_ERR\_NO\_PROGRESS, 242
  - UCS\_ERR\_NO\_RESOURCE, 242
  - UCS\_ERR\_NOT\_CONNECTED, 242
  - UCS\_ERR\_NOT\_IMPLEMENTED, 242
  - UCS\_ERR\_OUT\_OF\_RANGE, 242
  - UCS\_ERR\_REJECTED, 242
  - UCS\_ERR\_SHMEM\_SEGMENT, 242
  - UCS\_ERR\_SOME\_CONNECTS\_FAILED, 242
  - UCS\_ERR\_TIMED\_OUT, 242
  - UCS\_ERR\_UNREACHABLE, 242
  - UCS\_ERR\_UNSUPPORTED, 242
  - UCS\_INPROGRESS, 242
  - ucs\_memory\_type, 241
  - UCS\_MEMORY\_TYPE\_CUDA, 241
  - UCS\_MEMORY\_TYPE\_CUDA\_MANAGED, 241
  - UCS\_MEMORY\_TYPE\_HOST, 241
  - UCS\_MEMORY\_TYPE\_LAST, 241
  - UCS\_MEMORY\_TYPE\_ROCM, 241
  - UCS\_MEMORY\_TYPE\_ROCM\_MANAGED, 241
  - ucs\_memory\_type\_t, 240
  - UCS\_MEMORY\_TYPE\_UNKNOWN, 242
  - UCS\_OK, 242
  - ucs\_sock\_addr\_t, 240
  - ucs\_status\_ptr\_t, 241
  - ucs\_status\_t, 242
  - UCS\_THREAD\_MODE\_LAST, 243
  - UCS\_THREAD\_MODE\_MULTI, 243
  - UCS\_THREAD\_MODE\_SERIALIZED, 243
  - UCS\_THREAD\_MODE\_SINGLE, 243
  - ucs\_thread\_mode\_t, 243
  - ucs\_time\_t, 241
- ucs\_async\_add\_timer
  - UCS Communication Resource, 244
- ucs\_async\_context\_create
  - UCS Communication Resource, 245
- ucs\_async\_context\_destroy
  - UCS Communication Resource, 245
- ucs\_async\_event\_cb\_t
  - UCS Communication Resource, 240
- ucs\_async\_modify\_handler
  - UCS Communication Resource, 245
- ucs\_async\_poll
  - UCS Communication Resource, 246
- ucs\_async\_remove\_handler
  - UCS Communication Resource, 244
- ucs\_async\_set\_event\_handler
  - UCS Communication Resource, 243
- UCS\_CALLBACKQ\_FLAG\_FAST
  - UCS Communication Resource, 241
- UCS\_CALLBACKQ\_FLAG\_ONESHOT
  - UCS Communication Resource, 241
- ucs\_callbackq\_flags
  - UCS Communication Resource, 241
- UCS\_ERR\_ALREADY\_EXISTS
  - UCS Communication Resource, 242
- UCS\_ERR\_BUFFER\_TOO\_SMALL
  - UCS Communication Resource, 242
- UCS\_ERR\_BUSY
  - UCS Communication Resource, 242
- UCS\_ERR\_CANCELED
  - UCS Communication Resource, 242
- UCS\_ERR\_CONNECTION\_RESET
  - UCS Communication Resource, 242
- UCS\_ERR\_ENDPOINT\_TIMEOUT
  - UCS Communication Resource, 242
- UCS\_ERR\_EXCEEDS\_LIMIT
  - UCS Communication Resource, 242
- UCS\_ERR\_FIRST\_ENDPOINT\_FAILURE
  - UCS Communication Resource, 242

- UCS Communication Resource, [242](#)
- UCS\_ERR\_FIRST\_LINK\_FAILURE
  - UCS Communication Resource, [242](#)
- UCS\_ERR\_INVALID\_ADDR
  - UCS Communication Resource, [242](#)
- UCS\_ERR\_INVALID\_PARAM
  - UCS Communication Resource, [242](#)
- UCS\_ERR\_IO\_ERROR
  - UCS Communication Resource, [242](#)
- UCS\_ERR\_LAST
  - UCS Communication Resource, [242](#)
- UCS\_ERR\_LAST\_ENDPOINT\_FAILURE
  - UCS Communication Resource, [242](#)
- UCS\_ERR\_LAST\_LINK\_FAILURE
  - UCS Communication Resource, [242](#)
- UCS\_ERR\_MESSAGE\_TRUNCATED
  - UCS Communication Resource, [242](#)
- UCS\_ERR\_NO\_DEVICE
  - UCS Communication Resource, [242](#)
- UCS\_ERR\_NO\_ELEM
  - UCS Communication Resource, [242](#)
- UCS\_ERR\_NO\_MEMORY
  - UCS Communication Resource, [242](#)
- UCS\_ERR\_NO\_MESSAGE
  - UCS Communication Resource, [242](#)
- UCS\_ERR\_NO\_PROGRESS
  - UCS Communication Resource, [242](#)
- UCS\_ERR\_NO\_RESOURCE
  - UCS Communication Resource, [242](#)
- UCS\_ERR\_NOT\_CONNECTED
  - UCS Communication Resource, [242](#)
- UCS\_ERR\_NOT\_IMPLEMENTED
  - UCS Communication Resource, [242](#)
- UCS\_ERR\_OUT\_OF\_RANGE
  - UCS Communication Resource, [242](#)
- UCS\_ERR\_REJECTED
  - UCS Communication Resource, [242](#)
- UCS\_ERR\_SHMEM\_SEGMENT
  - UCS Communication Resource, [242](#)
- UCS\_ERR\_SOME\_CONNECTS\_FAILED
  - UCS Communication Resource, [242](#)
- UCS\_ERR\_TIMED\_OUT
  - UCS Communication Resource, [242](#)
- UCS\_ERR\_UNREACHABLE
  - UCS Communication Resource, [242](#)
- UCS\_ERR\_UNSUPPORTED
  - UCS Communication Resource, [242](#)
- UCS\_INPROGRESS
  - UCS Communication Resource, [242](#)
- ucs\_memory\_type
  - UCS Communication Resource, [241](#)
- UCS\_MEMORY\_TYPE\_CUDA
  - UCS Communication Resource, [241](#)
- UCS\_MEMORY\_TYPE\_CUDA\_MANAGED
  - UCS Communication Resource, [241](#)
- UCS\_MEMORY\_TYPE\_HOST
  - UCS Communication Resource, [241](#)
- UCS\_MEMORY\_TYPE\_LAST
  - UCS Communication Resource, [241](#)
- UCS\_MEMORY\_TYPE\_ROCM
  - UCS Communication Resource, [241](#)
- UCS\_MEMORY\_TYPE\_ROCM\_MANAGED
  - UCS Communication Resource, [241](#)
- ucs\_memory\_type\_t
  - UCS Communication Resource, [240](#)
- UCS\_MEMORY\_TYPE\_UNKNOWN
  - UCS Communication Resource, [242](#)
- UCS\_OK
  - UCS Communication Resource, [242](#)
- ucs\_sock\_addr, [240](#)
- ucs\_sock\_addr\_t
  - UCS Communication Resource, [240](#)
- ucs\_status\_ptr\_t
  - UCS Communication Resource, [241](#)
- ucs\_status\_t
  - UCS Communication Resource, [242](#)
- UCS\_THREAD\_MODE\_LAST
  - UCS Communication Resource, [243](#)
- UCS\_THREAD\_MODE\_MULTI
  - UCS Communication Resource, [243](#)
- UCS\_THREAD\_MODE\_SERIALIZED
  - UCS Communication Resource, [243](#)
- UCS\_THREAD\_MODE\_SINGLE
  - UCS Communication Resource, [243](#)
- ucs\_thread\_mode\_t
  - UCS Communication Resource, [243](#)
- ucs\_time\_t
  - UCS Communication Resource, [241](#)
- UCT Active messages, [195](#)
  - uct\_am\_callback\_t, [195](#)
  - uct\_am\_trace\_type, [197](#)
  - UCT\_AM\_TRACE\_TYPE\_LAST, [197](#)
  - UCT\_AM\_TRACE\_TYPE\_RECV, [197](#)
  - UCT\_AM\_TRACE\_TYPE\_RECV\_DROP, [197](#)
  - UCT\_AM\_TRACE\_TYPE\_SEND, [197](#)
  - UCT\_AM\_TRACE\_TYPE\_SEND\_DROP, [197](#)
  - uct\_am\_tracer\_t, [196](#)
  - uct\_ep\_am\_bcopy, [199](#)
  - uct\_ep\_am\_short, [198](#)
  - uct\_ep\_am\_short\_iov, [199](#)
  - uct\_ep\_am\_zcopy, [199](#)
  - uct\_iface\_release\_desc, [198](#)
  - uct\_iface\_set\_am\_handler, [197](#)
  - uct\_iface\_set\_am\_tracer, [198](#)
  - uct\_msg\_flags, [196](#)
  - UCT\_SEND\_FLAG\_PEER\_CHECK, [197](#)
  - UCT\_SEND\_FLAG\_SIGNALED, [197](#)
- UCT Atomic operations, [204](#)
  - uct\_ep\_atomic32\_fetch, [205](#)
  - uct\_ep\_atomic32\_post, [204](#)
  - uct\_ep\_atomic64\_fetch, [205](#)
  - uct\_ep\_atomic64\_post, [205](#)
  - uct\_ep\_atomic\_cswap32, [204](#)
  - uct\_ep\_atomic\_cswap64, [204](#)
- UCT client-server operations, [214](#)
  - uct\_cm\_attr\_field, [223](#)

- UCT\_CM\_ATTR\_FIELD\_MAX\_CONN\_PRIV, 223
- uct\_cm\_client\_ep\_conn\_notify, 230
- uct\_cm\_close, 229
- uct\_cm\_config\_read, 229
- uct\_cm\_ep\_client\_connect\_args\_field, 226
- UCT\_CM\_EP\_CLIENT\_CONNECT\_ARGS\_FIELD\_REMOTE\_DATA, 226
- UCT\_CM\_EP\_CLIENT\_CONNECT\_ARGS\_FIELD\_STATUS, 226
- uct\_cm\_ep\_client\_connect\_args\_t, 221
- uct\_cm\_ep\_client\_connect\_callback\_t, 222
- uct\_cm\_ep\_priv\_data\_pack\_args\_field, 224
- UCT\_CM\_EP\_PRIV\_DATA\_PACK\_ARGS\_FIELD\_DEVICE\_NAME, 225
- uct\_cm\_ep\_priv\_data\_pack\_args\_t, 220
- uct\_cm\_ep\_priv\_data\_pack\_callback\_t, 222
- uct\_cm\_ep\_resolve\_args\_field, 225
- UCT\_CM\_EP\_RESOLVE\_ARGS\_FIELD\_DEV\_NAME, 225
- UCT\_CM\_EP\_RESOLVE\_ARGS\_FIELD\_STATUS, 225
- uct\_cm\_ep\_resolve\_args\_t, 220
- uct\_cm\_ep\_resolve\_callback\_t, 223
- uct\_cm\_ep\_server\_conn\_notify\_args\_field, 226
- UCT\_CM\_EP\_SERVER\_CONN\_NOTIFY\_ARGS\_FIELD\_STATUS, 226
- uct\_cm\_ep\_server\_conn\_notify\_args\_t, 221
- uct\_cm\_ep\_server\_conn\_notify\_callback\_t, 221
- uct\_cm\_listener\_conn\_request\_args\_field, 225
- UCT\_CM\_LISTENER\_CONN\_REQUEST\_ARGS\_FIELD\_CLIENT\_ADDR, 226
- UCT\_CM\_LISTENER\_CONN\_REQUEST\_ARGS\_FIELD\_CONN\_REQUEST, 226
- UCT\_CM\_LISTENER\_CONN\_REQUEST\_ARGS\_FIELD\_DEV\_NAME, 226
- UCT\_CM\_LISTENER\_CONN\_REQUEST\_ARGS\_FIELD\_REMOTE\_DATA, 226
- uct\_cm\_listener\_conn\_request\_args\_t, 220
- uct\_cm\_listener\_conn\_request\_callback\_t, 221
- uct\_cm\_open, 228
- uct\_cm\_query, 229
- uct\_cm\_remote\_data\_field, 225
- UCT\_CM\_REMOTE\_DATA\_FIELD\_CONN\_PRIV\_DATA, 225
- UCT\_CM\_REMOTE\_DATA\_FIELD\_CONN\_PRIV\_DATA\_LENGTH, 225
- UCT\_CM\_REMOTE\_DATA\_FIELD\_DEV\_ADDR, 225
- UCT\_CM\_REMOTE\_DATA\_FIELD\_DEV\_ADDR\_LENGTH, 225
- uct\_cm\_remote\_data\_t, 220
- uct\_ep\_connect, 227
- UCT\_EP\_CONNECT\_PARAM\_FIELD\_PRIVATE\_DATA, 224
- UCT\_EP\_CONNECT\_PARAM\_FIELD\_PRIVATE\_DATA\_LENGTH, 224
- uct\_ep\_connect\_params\_field, 224
- uct\_ep\_disconnect, 228
- uct\_ep\_disconnect\_cb\_t, 222
- uct\_iface\_accept, 227
- uct\_iface\_reject, 227
- uct\_listener\_attr\_field, 224
- UCT\_LISTENER\_ATTR\_FIELD\_SOCKADDR, 224
- uct\_listener\_create, 230
- uct\_listener\_destroy, 230
- UCT\_LISTENER\_PARAM\_FIELD\_BACKLOG, 224
- UCT\_LISTENER\_PARAM\_FIELD\_CONN\_REQUEST\_CB, 224
- UCT\_LISTENER\_PARAM\_FIELD\_USER\_DATA, 224
- uct\_listener\_params\_field, 224
- uct\_listener\_query, 231
- uct\_listener\_reject, 231
- uct\_sockaddr\_conn\_request\_callback\_t, 221
- UCT Communication Context, 176
- UCT\_ALLOC\_METHOD\_DEFAULT, 177
- UCT\_ALLOC\_METHOD\_HEAP, 176
- UCT\_ALLOC\_METHOD\_HUGE, 177
- UCT\_ALLOC\_METHOD\_LAST, 177
- UCT\_ALLOC\_METHOD\_MD, 176
- UCT\_ALLOC\_METHOD\_MMAP, 176
- uct\_alloc\_method\_t, 176
- STATUS\_ALLOC\_METHOD\_THP, 176
- uct\_config\_get, 178
- uct\_config\_modify, 179
- uct\_worker\_create, 177
- uct\_worker\_destroy, 177
- uct\_worker\_progress, 179
- uct\_worker\_progress\_register\_safe, 177
- uct\_worker\_progress\_unregister\_safe, 178
- UCT Communication Resource, 136
- DevName\_t, 152
- uct\_async\_event\_cb\_t, 156
- REMOTE\_FLAG\_ASYNC, 159
- UCT\_CB\_FLAG\_RESERVED, 159
- uct\_cb\_flags, 158
- UCT\_CB\_PARAM\_FLAG\_DESC, 161
- UCT\_CB\_PARAM\_FLAG\_FIRST, 161
- UCT\_CB\_PARAM\_FLAG\_MORE, 161
- uct\_cb\_param\_flags, 161
- uct\_cm\_attr\_t, 153
- uct\_cm\_config\_t, 151
- ENCTH, 153
- uct\_cm\_t, 153
- uct\_completion\_callback\_t, 155
- uct\_completion\_t, 152
- uct\_completion\_update\_status, 175
- uct\_component\_attr\_field, 157
- UCT\_COMPONENT\_ATTR\_FIELD\_FLAGS, 157
- UCT\_COMPONENT\_ATTR\_FIELD\_MD\_RESOURCE\_COUNT, 157
- UCT\_COMPONENT\_ATTR\_FIELD\_MD\_RESOURCES, 157
- UCT\_COMPONENT\_ATTR\_FIELD\_NAME, 157
- uct\_component\_attr\_t, 150
- UCT\_COMPONENT\_FLAG\_CM, 157

- uct\_component\_h, 150
- uct\_component\_query, 162
- uct\_config\_release, 165
- uct\_conn\_request\_h, 154
- uct\_device\_addr\_t, 153
- UCT\_DEVICE\_TYPE\_ACC, 157
- UCT\_DEVICE\_TYPE\_LAST, 157
- UCT\_DEVICE\_TYPE\_NET, 157
- UCT\_DEVICE\_TYPE\_SELF, 157
- UCT\_DEVICE\_TYPE\_SHM, 157
- uct\_device\_type\_t, 157
- uct\_ep\_addr\_t, 153
- uct\_ep\_attr\_t, 152
- uct\_ep\_check, 168
- uct\_ep\_connect\_params\_t, 153
- uct\_ep\_connect\_to\_ep, 171
- uct\_ep\_create, 170
- uct\_ep\_destroy, 170
- uct\_ep\_fence, 173
- uct\_ep\_flush, 173
- uct\_ep\_get\_address, 170
- uct\_ep\_h, 151
- UCT\_EP\_PARAM\_FIELD\_CM, 160
- UCT\_EP\_PARAM\_FIELD\_CM\_RESOLVE\_CB, 160
- UCT\_EP\_PARAM\_FIELD\_CONN\_REQUEST, 160
- UCT\_EP\_PARAM\_FIELD\_DEV\_ADDR, 160
- UCT\_EP\_PARAM\_FIELD\_IFACE, 160
- UCT\_EP\_PARAM\_FIELD\_IFACE\_ADDR, 160
- UCT\_EP\_PARAM\_FIELD\_LOCAL\_SOCKADDR, 160
- UCT\_EP\_PARAM\_FIELD\_PATH\_INDEX, 160
- UCT\_EP\_PARAM\_FIELD\_PRIV\_DATA, 160
- UCT\_EP\_PARAM\_FIELD\_PRIV\_DATA\_LENGTH, 160
- UCT\_EP\_PARAM\_FIELD\_SOCKADDR, 160
- UCT\_EP\_PARAM\_FIELD\_SOCKADDR\_CB\_FLAGS, 160
- UCT\_EP\_PARAM\_FIELD\_SOCKADDR\_CONNECT\_CB\_CLIENT, 160
- UCT\_EP\_PARAM\_FIELD\_SOCKADDR\_DISCONNECT\_CB\_SERVER, 160
- UCT\_EP\_PARAM\_FIELD\_SOCKADDR\_NOTIFY\_CB\_SERVER, 160
- UCT\_EP\_PARAM\_FIELD\_SOCKADDR\_PACK\_CB, 160
- UCT\_EP\_PARAM\_FIELD\_USER\_DATA, 160
- uct\_ep\_params\_field, 160
- uct\_ep\_params\_t, 153
- uct\_ep\_pending\_add, 172
- uct\_ep\_pending\_purge, 173
- uct\_error\_handler\_t, 155
- UCT\_EVENT\_RECV, 158
- UCT\_EVENT\_RECV\_SIG, 158
- UCT\_EVENT\_SEND\_COMP, 158
- UCT\_FLUSH\_FLAG\_CANCEL, 158
- UCT\_FLUSH\_FLAG\_LOCAL, 158
- uct\_flush\_flags, 158
- uct\_iface\_addr\_t, 153
- uct\_iface\_attr\_t, 152
- uct\_iface\_close, 166
- uct\_iface\_config\_t, 151
- uct\_iface\_event\_arm, 168
- uct\_iface\_event\_fd\_get, 168
- uct\_iface\_event\_types, 158
- uct\_iface\_feature, 160
- UCT\_IFACE\_FEATURE\_AM, 161
- UCT\_IFACE\_FEATURE\_AMO32, 161
- UCT\_IFACE\_FEATURE\_AMO64, 161
- UCT\_IFACE\_FEATURE\_GET, 161
- UCT\_IFACE\_FEATURE\_LAST, 161
- UCT\_IFACE\_FEATURE\_PUT, 161
- UCT\_IFACE\_FEATURE\_TAG, 161
- uct\_iface\_fence, 172
- uct\_iface\_flush, 171
- uct\_iface\_get\_address, 167
- uct\_iface\_get\_device\_address, 166
- uct\_iface\_h, 150
- uct\_iface\_is\_reachable, 167
- uct\_iface\_mem\_alloc, 169
- uct\_iface\_mem\_free, 169
- uct\_iface\_open, 165
- uct\_iface\_open\_mode, 159
- UCT\_IFACE\_OPEN\_MODE\_DEVICE, 159
- UCT\_IFACE\_OPEN\_MODE\_SOCKADDR\_CLIENT, 159
- UCT\_IFACE\_OPEN\_MODE\_SOCKADDR\_SERVER, 159
- UCT\_IFACE\_PARAM\_FIELD\_AM\_ALIGN\_OFFSET, 160
- UCT\_IFACE\_PARAM\_FIELD\_AM\_ALIGNMENT, 160
- UCT\_IFACE\_PARAM\_FIELD\_ASYNC\_EVENT\_ARG, 160
- UCT\_IFACE\_PARAM\_FIELD\_ASYNC\_EVENT\_CB, 160
- UCT\_IFACE\_PARAM\_FIELD\_CPU\_MASK, 159
- UCT\_IFACE\_PARAM\_FIELD\_DEVICE, 159
- UCT\_IFACE\_PARAM\_FIELD\_ERR\_HANDLER, 159
- UCT\_IFACE\_PARAM\_FIELD\_ERR\_HANDLER\_ARG, 159
- UCT\_IFACE\_PARAM\_FIELD\_ERR\_HANDLER\_FLAGS, 159
- UCT\_IFACE\_PARAM\_FIELD\_FEATURES, 160
- UCT\_IFACE\_PARAM\_FIELD\_HW\_TM\_EAGER\_ARG, 159
- UCT\_IFACE\_PARAM\_FIELD\_HW\_TM\_EAGER\_CB, 160
- UCT\_IFACE\_PARAM\_FIELD\_HW\_TM\_RNDV\_ARG, 160
- UCT\_IFACE\_PARAM\_FIELD\_HW\_TM\_RNDV\_CB, 160
- UCT\_IFACE\_PARAM\_FIELD\_KEEPLIVE\_INTERVAL, 160
- UCT\_IFACE\_PARAM\_FIELD\_OPEN\_MODE, 159



- UCT\_IFACE\_PARAM\_FIELD\_RX\_HEADROOM, 159
- UCT\_IFACE\_PARAM\_FIELD\_SOCKADDR, 159
- UCT\_IFACE\_PARAM\_FIELD\_STATS\_ROOT, 159
- uct\_iface\_params\_field, 159
- uct\_iface\_params\_t, 152
- uct\_iface\_progress, 175
- uct\_iface\_progress\_disable, 174
- uct\_iface\_progress\_enable, 174
- uct\_iface\_query, 166
- uct\_iov\_t, 154
- uct\_listener\_attr\_t, 153
- uct\_listener\_h, 154
- uct\_listener\_params\_t, 154
- uct\_md\_attr\_t, 152
- uct\_md\_close, 163
- uct\_md\_config\_t, 151
- uct\_md\_h, 151
- uct\_md\_iface\_config\_read, 164
- uct\_md\_open, 163
- uct\_md\_ops\_t, 151
- uct\_md\_query\_tl\_resources, 163
- uct\_md\_resource\_desc\_t, 150
- uct\_md\_t, 152
- uct\_mem\_h, 151
- uct\_pack\_callback\_t, 156
- uct\_pending\_callback\_t, 155
- uct\_pending\_purge\_callback\_t, 156
- uct\_pending\_req\_t, 152
- UCT\_PROGRESS\_RECV, 158
- UCT\_PROGRESS\_SEND, 158
- UCT\_PROGRESS\_THREAD\_SAFE, 158
- uct\_progress\_types, 158
- uct\_query\_components, 161
- uct\_release\_component\_list, 162
- uct\_release\_tl\_resource\_list, 164
- uct\_rkey\_ctx\_h, 151
- uct\_rkey\_t, 151
- uct\_tag\_context\_t, 154
- UCT\_TAG\_RECV\_CB\_INLINE\_DATA, 161
- uct\_tag\_t, 154
- uct\_tl\_resource\_desc\_t, 150
- uct\_unpack\_callback\_t, 156
- uct\_worker\_cb\_id\_t, 154
- uct\_worker\_h, 152
- UCT interface for asynchronous event capabilities, 237
  - UCT\_IFACE\_FLAG\_EVENT\_ASYNC\_CB, 237
  - UCT\_IFACE\_FLAG\_EVENT\_FD, 237
  - UCT\_IFACE\_FLAG\_EVENT\_RECV, 237
  - UCT\_IFACE\_FLAG\_EVENT\_RECV\_SIG, 237
  - UCT\_IFACE\_FLAG\_EVENT\_SEND\_COMP, 237
- UCT interface operations and capabilities, 232
  - UCT\_IFACE\_FLAG\_AM\_BCOPY, 233
  - UCT\_IFACE\_FLAG\_AM\_DUP, 235
  - UCT\_IFACE\_FLAG\_AM\_SHORT, 232
  - UCT\_IFACE\_FLAG\_AM\_ZCOPY, 233
  - UCT\_IFACE\_FLAG\_ATOMIC\_CPU, 234
  - UCT\_IFACE\_FLAG\_ATOMIC\_DEVICE, 234
  - UCT\_IFACE\_FLAG\_CB\_ASYNC, 236
  - UCT\_IFACE\_FLAG\_CB\_SYNC, 236
  - UCT\_IFACE\_FLAG\_CONNECT\_TO\_EP, 235
  - UCT\_IFACE\_FLAG\_CONNECT\_TO\_IFACE, 235
  - UCT\_IFACE\_FLAG\_CONNECT\_TO\_SOCKADDR, 235
  - UCT\_IFACE\_FLAG\_EP\_CHECK, 235
  - UCT\_IFACE\_FLAG\_EP\_KEEPALIVE, 236
  - UCT\_IFACE\_FLAG\_ERRHANDLE\_AM\_ID, 234
  - UCT\_IFACE\_FLAG\_ERRHANDLE\_BCOPY\_BUF, 234
  - UCT\_IFACE\_FLAG\_ERRHANDLE\_BCOPY\_LEN, 235
  - UCT\_IFACE\_FLAG\_ERRHANDLE\_PEER\_FAILURE, 235
  - UCT\_IFACE\_FLAG\_ERRHANDLE\_REMOTE\_MEM, 234
  - UCT\_IFACE\_FLAG\_ERRHANDLE\_SHORT\_BUF, 234
  - UCT\_IFACE\_FLAG\_ERRHANDLE\_ZCOPY\_BUF, 234
  - UCT\_IFACE\_FLAG\_GET\_BCOPY, 234
  - UCT\_IFACE\_FLAG\_GET\_SHORT, 233
  - UCT\_IFACE\_FLAG\_GET\_ZCOPY, 234
  - UCT\_IFACE\_FLAG\_PENDING, 233
  - UCT\_IFACE\_FLAG\_PUT\_BCOPY, 233
  - UCT\_IFACE\_FLAG\_PUT\_SHORT, 233
  - UCT\_IFACE\_FLAG\_PUT\_ZCOPY, 233
  - UCT\_IFACE\_FLAG\_TAG\_EAGER\_BCOPY, 236
  - UCT\_IFACE\_FLAG\_TAG\_EAGER\_SHORT, 236
  - UCT\_IFACE\_FLAG\_TAG\_EAGER\_ZCOPY, 236
  - UCT\_IFACE\_FLAG\_TAG\_RNDV\_ZCOPY, 236
- UCT Memory Domain, 181
  - uct\_allocated\_memory\_t, 185
  - UCT\_MADV\_NORMAL, 187
  - UCT\_MADV\_WILLNEED, 187
  - uct\_md\_config\_read, 192
  - uct\_md\_detect\_memory\_type, 190
  - UCT\_MD\_FLAG\_ADVISE, 186
  - UCT\_MD\_FLAG\_ALLOC, 186
  - UCT\_MD\_FLAG\_FIXED, 186
  - UCT\_MD\_FLAG\_INVALIDATE, 186
  - UCT\_MD\_FLAG\_NEED\_MEMH, 186
  - UCT\_MD\_FLAG\_NEED\_RKEY, 186
  - UCT\_MD\_FLAG\_REG, 186
  - UCT\_MD\_FLAG\_RKEY\_PTR, 186
  - UCT\_MD\_FLAG\_SOCKADDR, 186
  - uct\_md\_is\_sockaddr\_accessible, 192
  - UCT\_MD\_MEM\_ACCESS\_ALL, 187
  - UCT\_MD\_MEM\_ACCESS\_LOCAL\_READ, 187
  - UCT\_MD\_MEM\_ACCESS\_LOCAL\_WRITE, 187
  - UCT\_MD\_MEM\_ACCESS\_REMOTE\_ATOMIC, 187
  - UCT\_MD\_MEM\_ACCESS\_REMOTE\_GET, 187
  - UCT\_MD\_MEM\_ACCESS\_REMOTE\_PUT, 187
  - UCT\_MD\_MEM\_ACCESS\_RMA, 187
  - uct\_md\_mem\_advise, 189
  - uct\_md\_mem\_attr\_field, 187

- UCT\_MD\_MEM\_ATTR\_FIELD\_ALLOC\_LENGTH, 187
- UCT\_MD\_MEM\_ATTR\_FIELD\_BASE\_ADDRESS, 187
- UCT\_MD\_MEM\_ATTR\_FIELD\_MEM\_TYPE, 187
- UCT\_MD\_MEM\_ATTR\_FIELD\_SYS\_DEV, 187
- uct\_md\_mem\_attr\_field\_t, 185
- uct\_md\_mem\_attr\_t, 185
- uct\_md\_mem\_dereg, 190
- UCT\_MD\_MEM\_FLAG\_FIXED, 187
- UCT\_MD\_MEM\_FLAG\_HIDE\_ERRORS, 187
- UCT\_MD\_MEM\_FLAG\_LOCK, 187
- UCT\_MD\_MEM\_FLAG\_NONBLOCK, 187
- uct\_md\_mem\_flags, 186
- uct\_md\_mem\_query, 188
- uct\_md\_mem\_reg, 189
- uct\_md\_mkey\_pack, 193
- uct\_md\_query, 188
- uct\_mem\_advice\_t, 187
- uct\_mem\_alloc, 190
- UCT\_MEM\_ALLOC\_PARAM\_FIELD\_ADDRESS, 188
- UCT\_MEM\_ALLOC\_PARAM\_FIELD\_FLAGS, 188
- UCT\_MEM\_ALLOC\_PARAM\_FIELD\_MDS, 188
- UCT\_MEM\_ALLOC\_PARAM\_FIELD\_MEM\_TYPE, 188
- UCT\_MEM\_ALLOC\_PARAM\_FIELD\_NAME, 188
- uct\_mem\_alloc\_params\_field\_t, 187
- uct\_mem\_free, 192
- uct\_rkey\_bundle\_t, 186
- uct\_rkey\_ptr, 194
- uct\_rkey\_release, 194
- uct\_rkey\_unpack, 193
- UCT\_SOCKADDR\_ACC\_LOCAL, 186
- UCT\_SOCKADDR\_ACC\_REMOTE, 186
- uct\_sockaddr\_accessibility\_t, 186
- UCT Remote memory access operations, 201
  - uct\_ep\_get\_bcopy, 202
  - uct\_ep\_get\_short, 202
  - uct\_ep\_get\_zcopy, 202
  - uct\_ep\_put\_bcopy, 201
  - uct\_ep\_put\_short, 201
  - uct\_ep\_put\_zcopy, 201
- UCT Tag matching operations, 206
  - uct\_ep\_tag\_eager\_bcopy, 209
  - uct\_ep\_tag\_eager\_short, 208
  - uct\_ep\_tag\_eager\_zcopy, 209
  - uct\_ep\_tag\_rndv\_cancel, 211
  - uct\_ep\_tag\_rndv\_request, 211
  - uct\_ep\_tag\_rndv\_zcopy, 210
  - uct\_iface\_tag\_recv\_cancel, 212
  - uct\_iface\_tag\_recv\_zcopy, 212
  - uct\_tag\_unexp\_eager\_cb\_t, 206
  - uct\_tag\_unexp\_rndv\_cb\_t, 207
- UCT\_ALLOC\_METHOD\_DEFAULT
  - UCT Communication Context, 177
- UCT\_ALLOC\_METHOD\_HEAP
  - UCT Communication Context, 176
- UCT\_ALLOC\_METHOD\_HUGE
  - UCT Communication Context, 177
- UCT\_ALLOC\_METHOD\_LAST
  - UCT Communication Context, 177
- UCT\_ALLOC\_METHOD\_MD
  - UCT Communication Context, 176
- UCT\_ALLOC\_METHOD\_MMAP
  - UCT Communication Context, 176
- uct\_alloc\_method\_t
  - UCT Communication Context, 176
- UCT\_ALLOC\_METHOD\_THP
  - UCT Communication Context, 176
- uct\_allocated\_memory, 184
- uct\_allocated\_memory\_t
  - UCT Memory Domain, 185
- uct\_am\_callback\_t
  - UCT Active messages, 195
- uct\_am\_trace\_type
  - UCT Active messages, 197
- UCT\_AM\_TRACE\_TYPE\_LAST
  - UCT Active messages, 197
- UCT\_AM\_TRACE\_TYPE\_RECV
  - UCT Active messages, 197
- UCT\_AM\_TRACE\_TYPE\_RECV\_DROP
  - UCT Active messages, 197
- UCT\_AM\_TRACE\_TYPE\_SEND
  - UCT Active messages, 197
- UCT\_AM\_TRACE\_TYPE\_SEND\_DROP
  - UCT Active messages, 197
- uct\_am\_trace\_type\_t
  - UCT Communication Resource, 152
- uct\_am\_tracer\_t
  - UCT Active messages, 196
- uct\_async\_event\_cb\_t
  - UCT Communication Resource, 156
- UCT\_CB\_FLAG\_ASYNC
  - UCT Communication Resource, 159
- UCT\_CB\_FLAG\_RESERVED
  - UCT Communication Resource, 159
- uct\_cb\_flags
  - UCT Communication Resource, 158
- UCT\_CB\_PARAM\_FLAG\_DESC
  - UCT Communication Resource, 161
- UCT\_CB\_PARAM\_FLAG\_FIRST
  - UCT Communication Resource, 161
- UCT\_CB\_PARAM\_FLAG\_MORE
  - UCT Communication Resource, 161
- uct\_cb\_param\_flags
  - UCT Communication Resource, 161
- uct\_cm\_attr, 217
- uct\_cm\_attr\_field
  - UCT client-server operations, 223
- UCT\_CM\_ATTR\_FIELD\_MAX\_CONN\_PRIV
  - UCT client-server operations, 223
- uct\_cm\_attr\_t
  - UCT Communication Resource, 153
- uct\_cm\_client\_ep\_conn\_notify
  - UCT client-server operations, 230

- uct\_cm\_close
  - UCT client-server operations, [229](#)
- uct\_cm\_config\_read
  - UCT client-server operations, [229](#)
- uct\_cm\_config\_t
  - UCT Communication Resource, [151](#)
- uct\_cm\_ep\_client\_connect\_args, [219](#)
- uct\_cm\_ep\_client\_connect\_args\_field
  - UCT client-server operations, [226](#)
- UCT\_CM\_EP\_CLIENT\_CONNECT\_ARGS\_FIELD\_REMOTE\_DATA
  - UCT client-server operations, [226](#)
- UCT\_CM\_EP\_CLIENT\_CONNECT\_ARGS\_FIELD\_STATUS
  - UCT client-server operations, [226](#)
- uct\_cm\_ep\_client\_connect\_args\_t
  - UCT client-server operations, [221](#)
- uct\_cm\_ep\_client\_connect\_callback\_t
  - UCT client-server operations, [222](#)
- uct\_cm\_ep\_priv\_data\_pack\_args, [218](#)
- uct\_cm\_ep\_priv\_data\_pack\_args\_field
  - UCT client-server operations, [224](#)
- UCT\_CM\_EP\_PRIV\_DATA\_PACK\_ARGS\_FIELD\_DEVICE\_NAME
  - UCT client-server operations, [225](#)
- uct\_cm\_ep\_priv\_data\_pack\_args\_t
  - UCT client-server operations, [220](#)
- uct\_cm\_ep\_priv\_data\_pack\_callback\_t
  - UCT client-server operations, [222](#)
- uct\_cm\_ep\_resolve\_args, [218](#)
- uct\_cm\_ep\_resolve\_args\_field
  - UCT client-server operations, [225](#)
- UCT\_CM\_EP\_RESOLVE\_ARGS\_FIELD\_DEV\_NAME
  - UCT client-server operations, [225](#)
- UCT\_CM\_EP\_RESOLVE\_ARGS\_FIELD\_STATUS
  - UCT client-server operations, [225](#)
- uct\_cm\_ep\_resolve\_args\_t
  - UCT client-server operations, [220](#)
- uct\_cm\_ep\_resolve\_callback\_t
  - UCT client-server operations, [223](#)
- uct\_cm\_ep\_server\_conn\_notify\_args, [220](#)
- uct\_cm\_ep\_server\_conn\_notify\_args\_field
  - UCT client-server operations, [226](#)
- UCT\_CM\_EP\_SERVER\_CONN\_NOTIFY\_ARGS\_FIELD\_STATUS
  - UCT client-server operations, [226](#)
- uct\_cm\_ep\_server\_conn\_notify\_args\_t
  - UCT client-server operations, [221](#)
- uct\_cm\_ep\_server\_conn\_notify\_callback\_t
  - UCT client-server operations, [221](#)
- uct\_cm\_h
  - UCT Communication Resource, [153](#)
- uct\_cm\_listener\_conn\_request\_args, [219](#)
- uct\_cm\_listener\_conn\_request\_args\_field
  - UCT client-server operations, [225](#)
- UCT\_CM\_LISTENER\_CONN\_REQUEST\_ARGS\_FIELD\_CLIENT\_ADDR
  - UCT Communication Context, [178](#)
- UCT\_CM\_LISTENER\_CONN\_REQUEST\_ARGS\_FIELD\_CONN\_REQUEST
  - UCT Communication Context, [179](#)
- UCT\_CM\_LISTENER\_CONN\_REQUEST\_ARGS\_FIELD\_DEV\_NAME
  - UCT Communication Resource, [165](#)
- UCT\_CM\_LISTENER\_CONN\_REQUEST\_ARGS\_FIELD\_REMOTE\_DATA
  - UCT Communication Resource, [154](#)
- UCT client-server operations, [226](#)
- uct\_cm\_listener\_conn\_request\_args\_t
  - UCT client-server operations, [220](#)
- uct\_cm\_listener\_conn\_request\_callback\_t
  - UCT client-server operations, [221](#)
- uct\_cm\_open
  - UCT client-server operations, [228](#)
- uct\_cm\_query
  - UCT client-server operations, [229](#)
- uct\_cm\_remote\_data, [218](#)
- uct\_cm\_remote\_data\_field
  - UCT client-server operations, [225](#)
- UCT\_CM\_REMOTE\_DATA\_FIELD\_CONN\_PRIV\_DATA
  - UCT client-server operations, [225](#)
- UCT\_CM\_REMOTE\_DATA\_FIELD\_CONN\_PRIV\_DATA\_LENGTH
  - UCT client-server operations, [225](#)
- UCT\_CM\_REMOTE\_DATA\_FIELD\_DEV\_ADDR
  - UCT client-server operations, [225](#)
- UCT\_CM\_REMOTE\_DATA\_FIELD\_DEV\_ADDR\_LENGTH
  - UCT client-server operations, [225](#)
- uct\_cm\_remote\_data\_t
  - UCT client-server operations, [220](#)
- uct\_cm\_t
  - UCT Communication Resource, [153](#)
- uct\_completion, [148](#)
- uct\_completion\_callback\_t
  - UCT Communication Resource, [155](#)
- uct\_completion\_t
  - UCT Communication Resource, [152](#)
- uct\_completion\_update\_status
  - UCT Communication Resource, [175](#)
- uct\_component\_attr, [141](#)
- uct\_component\_attr\_field
  - UCT Communication Resource, [157](#)
- UCT\_COMPONENT\_ATTR\_FIELD\_FLAGS
  - UCT Communication Resource, [157](#)
- UCT\_COMPONENT\_ATTR\_FIELD\_MD\_RESOURCE\_COUNT
  - UCT Communication Resource, [157](#)
- UCT\_COMPONENT\_ATTR\_FIELD\_MD\_RESOURCES
  - UCT Communication Resource, [157](#)
- UCT\_COMPONENT\_ATTR\_FIELD\_NAME
  - UCT Communication Resource, [157](#)
- uct\_component\_attr\_t
  - UCT Communication Resource, [150](#)
- UCT\_COMPONENT\_FLAG\_CM
  - UCT Communication Resource, [157](#)
- uct\_component\_h
  - UCT Communication Resource, [150](#)
- uct\_component\_query
  - UCT Communication Resource, [162](#)
- uct\_config\_get
  - UCT Communication Context, [178](#)
- uct\_config\_modify
  - UCT Communication Context, [179](#)
- uct\_config\_release
  - UCT Communication Resource, [165](#)
- uct\_conn\_request\_h
  - UCT Communication Resource, [154](#)

- uct\_device\_addr\_t
  - UCT Communication Resource, [153](#)
- UCT\_DEVICE\_TYPE\_ACC
  - UCT Communication Resource, [157](#)
- UCT\_DEVICE\_TYPE\_LAST
  - UCT Communication Resource, [157](#)
- UCT\_DEVICE\_TYPE\_NET
  - UCT Communication Resource, [157](#)
- UCT\_DEVICE\_TYPE\_SELF
  - UCT Communication Resource, [157](#)
- UCT\_DEVICE\_TYPE\_SHM
  - UCT Communication Resource, [157](#)
- uct\_device\_type\_t
  - UCT Communication Resource, [157](#)
- uct\_ep\_addr\_t
  - UCT Communication Resource, [153](#)
- uct\_ep\_am\_bcopy
  - UCT Active messages, [199](#)
- uct\_ep\_am\_short
  - UCT Active messages, [198](#)
- uct\_ep\_am\_short\_iov
  - UCT Active messages, [199](#)
- uct\_ep\_am\_zcopy
  - UCT Active messages, [199](#)
- uct\_ep\_atomic32\_fetch
  - UCT Atomic operations, [205](#)
- uct\_ep\_atomic32\_post
  - UCT Atomic operations, [204](#)
- uct\_ep\_atomic64\_fetch
  - UCT Atomic operations, [205](#)
- uct\_ep\_atomic64\_post
  - UCT Atomic operations, [205](#)
- uct\_ep\_atomic\_cswap32
  - UCT Atomic operations, [204](#)
- uct\_ep\_atomic\_cswap64
  - UCT Atomic operations, [204](#)
- uct\_ep\_attr\_t
  - UCT Communication Resource, [152](#)
- uct\_ep\_check
  - UCT Communication Resource, [168](#)
- uct\_ep\_connect
  - UCT client-server operations, [227](#)
- UCT\_EP\_CONNECT\_PARAM\_FIELD\_PRIVATE\_DATA
  - UCT client-server operations, [224](#)
- UCT\_EP\_CONNECT\_PARAM\_FIELD\_PRIVATE\_DATA\_LENGTH
  - UCT client-server operations, [224](#)
- uct\_ep\_connect\_params, [217](#)
- uct\_ep\_connect\_params\_field
  - UCT client-server operations, [224](#)
- uct\_ep\_connect\_params\_t
  - UCT Communication Resource, [153](#)
- uct\_ep\_connect\_to\_ep
  - UCT Communication Resource, [171](#)
- uct\_ep\_create
  - UCT Communication Resource, [170](#)
- uct\_ep\_destroy
  - UCT Communication Resource, [170](#)
- uct\_ep\_disconnect
  - UCT client-server operations, [228](#)
- uct\_ep\_disconnect\_cb\_t
  - UCT client-server operations, [222](#)
- uct\_ep\_fence
  - UCT Communication Resource, [173](#)
- uct\_ep\_flush
  - UCT Communication Resource, [173](#)
- uct\_ep\_get\_address
  - UCT Communication Resource, [170](#)
- uct\_ep\_get\_bcopy
  - UCT Remote memory access operations, [202](#)
- uct\_ep\_get\_short
  - UCT Remote memory access operations, [202](#)
- uct\_ep\_get\_zcopy
  - UCT Remote memory access operations, [202](#)
- uct\_ep\_h
  - UCT Communication Resource, [151](#)
- UCT\_EP\_PARAM\_FIELD\_CM
  - UCT Communication Resource, [160](#)
- UCT\_EP\_PARAM\_FIELD\_CM\_RESOLVE\_CB
  - UCT Communication Resource, [160](#)
- UCT\_EP\_PARAM\_FIELD\_CONN\_REQUEST
  - UCT Communication Resource, [160](#)
- UCT\_EP\_PARAM\_FIELD\_DEV\_ADDR
  - UCT Communication Resource, [160](#)
- UCT\_EP\_PARAM\_FIELD\_IFACE
  - UCT Communication Resource, [160](#)
- UCT\_EP\_PARAM\_FIELD\_IFACE\_ADDR
  - UCT Communication Resource, [160](#)
- UCT\_EP\_PARAM\_FIELD\_LOCAL\_SOCKADDR
  - UCT Communication Resource, [160](#)
- UCT\_EP\_PARAM\_FIELD\_PATH\_INDEX
  - UCT Communication Resource, [160](#)
- UCT\_EP\_PARAM\_FIELD\_PRIV\_DATA
  - UCT Communication Resource, [160](#)
- UCT\_EP\_PARAM\_FIELD\_PRIV\_DATA\_LENGTH
  - UCT Communication Resource, [160](#)
- UCT\_EP\_PARAM\_FIELD\_SOCKADDR
  - UCT Communication Resource, [160](#)
- UCT\_EP\_PARAM\_FIELD\_SOCKADDR\_CB\_FLAGS
  - UCT Communication Resource, [160](#)
- UCT\_EP\_PARAM\_FIELD\_SOCKADDR\_CONNECT\_CB\_CLIENT
  - UCT Communication Resource, [160](#)
- UCT\_EP\_PARAM\_FIELD\_SOCKADDR\_DISCONNECT\_CB
  - UCT Communication Resource, [160](#)
- UCT\_EP\_PARAM\_FIELD\_SOCKADDR\_NOTIFY\_CB\_SERVER
  - UCT Communication Resource, [160](#)
- UCT\_EP\_PARAM\_FIELD\_SOCKADDR\_PACK\_CB
  - UCT Communication Resource, [160](#)
- UCT\_EP\_PARAM\_FIELD\_USER\_DATA
  - UCT Communication Resource, [160](#)
- uct\_ep\_params, [146](#)
- uct\_ep\_params\_field
  - UCT Communication Resource, [160](#)
- uct\_ep\_params\_t
  - UCT Communication Resource, [153](#)
- uct\_ep\_pending\_add
  - UCT Communication Resource, [172](#)

- uct\_ep\_pending\_purge
  - UCT Communication Resource, [173](#)
- uct\_ep\_put\_bcopy
  - UCT Remote memory access operations, [201](#)
- uct\_ep\_put\_short
  - UCT Remote memory access operations, [201](#)
- uct\_ep\_put\_zcopy
  - UCT Remote memory access operations, [201](#)
- uct\_ep\_tag\_eager\_bcopy
  - UCT Tag matching operations, [209](#)
- uct\_ep\_tag\_eager\_short
  - UCT Tag matching operations, [208](#)
- uct\_ep\_tag\_eager\_zcopy
  - UCT Tag matching operations, [209](#)
- uct\_ep\_tag\_rndv\_cancel
  - UCT Tag matching operations, [211](#)
- uct\_ep\_tag\_rndv\_request
  - UCT Tag matching operations, [211](#)
- uct\_ep\_tag\_rndv\_zcopy
  - UCT Tag matching operations, [210](#)
- uct\_error\_handler\_t
  - UCT Communication Resource, [155](#)
- UCT\_EVENT\_RECV
  - UCT Communication Resource, [158](#)
- UCT\_EVENT\_RECV\_SIG
  - UCT Communication Resource, [158](#)
- UCT\_EVENT\_SEND\_COMP
  - UCT Communication Resource, [158](#)
- UCT\_FLUSH\_FLAG\_CANCEL
  - UCT Communication Resource, [158](#)
- UCT\_FLUSH\_FLAG\_LOCAL
  - UCT Communication Resource, [158](#)
- uct\_flush\_flags
  - UCT Communication Resource, [158](#)
- uct\_iface\_accept
  - UCT client-server operations, [227](#)
- uct\_iface\_addr\_t
  - UCT Communication Resource, [153](#)
- uct\_iface\_attr, [142](#)
- uct\_iface\_attr.cap, [142](#)
- uct\_iface\_attr.cap.am, [143](#)
- uct\_iface\_attr.cap.atomic32, [144](#)
- uct\_iface\_attr.cap.atomic64, [144](#)
- uct\_iface\_attr.cap.get, [143](#)
- uct\_iface\_attr.cap.put, [143](#)
- uct\_iface\_attr.cap.tag, [143](#)
- uct\_iface\_attr.cap.tag.eager, [144](#)
- uct\_iface\_attr.cap.tag.recv, [144](#)
- uct\_iface\_attr.cap.tag.rndv, [144](#)
- uct\_iface\_attr\_t
  - UCT Communication Resource, [152](#)
- uct\_iface\_close
  - UCT Communication Resource, [166](#)
- uct\_iface\_config\_t
  - UCT Communication Resource, [151](#)
- uct\_iface\_event\_arm
  - UCT Communication Resource, [168](#)
- uct\_iface\_event\_fd\_get
  - UCT Communication Resource, [168](#)
- uct\_iface\_event\_types
  - UCT Communication Resource, [158](#)
- uct\_iface\_feature
  - UCT Communication Resource, [160](#)
- UCT\_IFACE\_FEATURE\_AM
  - UCT Communication Resource, [161](#)
- UCT\_IFACE\_FEATURE\_AMO32
  - UCT Communication Resource, [161](#)
- UCT\_IFACE\_FEATURE\_AMO64
  - UCT Communication Resource, [161](#)
- UCT\_IFACE\_FEATURE\_GET
  - UCT Communication Resource, [161](#)
- UCT\_IFACE\_FEATURE\_LAST
  - UCT Communication Resource, [161](#)
- UCT\_IFACE\_FEATURE\_PUT
  - UCT Communication Resource, [161](#)
- UCT\_IFACE\_FEATURE\_TAG
  - UCT Communication Resource, [161](#)
- uct\_iface\_fence
  - UCT Communication Resource, [172](#)
- UCT\_IFACE\_FLAG\_AM\_BCOPY
  - UCT interface operations and capabilities, [233](#)
- UCT\_IFACE\_FLAG\_AM\_DUP
  - UCT interface operations and capabilities, [235](#)
- UCT\_IFACE\_FLAG\_AM\_SHORT
  - UCT interface operations and capabilities, [232](#)
- UCT\_IFACE\_FLAG\_AM\_ZCOPY
  - UCT interface operations and capabilities, [233](#)
- UCT\_IFACE\_FLAG\_ATOMIC\_CPU
  - UCT interface operations and capabilities, [234](#)
- UCT\_IFACE\_FLAG\_ATOMIC\_DEVICE
  - UCT interface operations and capabilities, [234](#)
- UCT\_IFACE\_FLAG\_CB\_ASYNC
  - UCT interface operations and capabilities, [236](#)
- UCT\_IFACE\_FLAG\_CB\_SYNC
  - UCT interface operations and capabilities, [236](#)
- UCT\_IFACE\_FLAG\_CONNECT\_TO\_EP
  - UCT interface operations and capabilities, [235](#)
- UCT\_IFACE\_FLAG\_CONNECT\_TO\_IFACE
  - UCT interface operations and capabilities, [235](#)
- UCT\_IFACE\_FLAG\_CONNECT\_TO\_SOCKETADDR
  - UCT interface operations and capabilities, [235](#)
- UCT\_IFACE\_FLAG\_EP\_CHECK
  - UCT interface operations and capabilities, [235](#)
- UCT\_IFACE\_FLAG\_EP\_KEEPALIVE
  - UCT interface operations and capabilities, [236](#)
- UCT\_IFACE\_FLAG\_ERRHANDLE\_AM\_ID
  - UCT interface operations and capabilities, [234](#)
- UCT\_IFACE\_FLAG\_ERRHANDLE\_BCOPY\_BUF
  - UCT interface operations and capabilities, [234](#)
- UCT\_IFACE\_FLAG\_ERRHANDLE\_BCOPY\_LEN
  - UCT interface operations and capabilities, [235](#)
- UCT\_IFACE\_FLAG\_ERRHANDLE\_PEER\_FAILURE
  - UCT interface operations and capabilities, [235](#)
- UCT\_IFACE\_FLAG\_ERRHANDLE\_REMOTE\_MEM
  - UCT interface operations and capabilities, [234](#)
- UCT\_IFACE\_FLAG\_ERRHANDLE\_SHORT\_BUF

- UCT interface operations and capabilities, [234](#)
- UCT\_IFACE\_FLAG\_ERRHANDLE\_ZCOPY\_BUF
  - UCT interface operations and capabilities, [234](#)
- UCT\_IFACE\_FLAG\_EVENT\_ASYNC\_CB
  - UCT interface for asynchronous event capabilities, [237](#)
- UCT\_IFACE\_FLAG\_EVENT\_FD
  - UCT interface for asynchronous event capabilities, [237](#)
- UCT\_IFACE\_FLAG\_EVENT\_RECV
  - UCT interface for asynchronous event capabilities, [237](#)
- UCT\_IFACE\_FLAG\_EVENT\_RECV\_SIG
  - UCT interface for asynchronous event capabilities, [237](#)
- UCT\_IFACE\_FLAG\_EVENT\_SEND\_COMP
  - UCT interface for asynchronous event capabilities, [237](#)
- UCT\_IFACE\_FLAG\_GET\_BCOPY
  - UCT interface operations and capabilities, [234](#)
- UCT\_IFACE\_FLAG\_GET\_SHORT
  - UCT interface operations and capabilities, [233](#)
- UCT\_IFACE\_FLAG\_GET\_ZCOPY
  - UCT interface operations and capabilities, [234](#)
- UCT\_IFACE\_FLAG\_PENDING
  - UCT interface operations and capabilities, [233](#)
- UCT\_IFACE\_FLAG\_PUT\_BCOPY
  - UCT interface operations and capabilities, [233](#)
- UCT\_IFACE\_FLAG\_PUT\_SHORT
  - UCT interface operations and capabilities, [233](#)
- UCT\_IFACE\_FLAG\_PUT\_ZCOPY
  - UCT interface operations and capabilities, [233](#)
- UCT\_IFACE\_FLAG\_TAG\_EAGER\_BCOPY
  - UCT interface operations and capabilities, [236](#)
- UCT\_IFACE\_FLAG\_TAG\_EAGER\_SHORT
  - UCT interface operations and capabilities, [236](#)
- UCT\_IFACE\_FLAG\_TAG\_EAGER\_ZCOPY
  - UCT interface operations and capabilities, [236](#)
- UCT\_IFACE\_FLAG\_TAG\_RNDV\_ZCOPY
  - UCT interface operations and capabilities, [236](#)
- uct\_iface\_flush
  - UCT Communication Resource, [171](#)
- uct\_iface\_get\_address
  - UCT Communication Resource, [167](#)
- uct\_iface\_get\_device\_address
  - UCT Communication Resource, [166](#)
- uct\_iface\_h
  - UCT Communication Resource, [150](#)
- uct\_iface\_is\_reachable
  - UCT Communication Resource, [167](#)
- uct\_iface\_mem\_alloc
  - UCT Communication Resource, [169](#)
- uct\_iface\_mem\_free
  - UCT Communication Resource, [169](#)
- uct\_iface\_open
  - UCT Communication Resource, [165](#)
- uct\_iface\_open\_mode
  - UCT Communication Resource, [159](#)
- UCT\_IFACE\_OPEN\_MODE\_DEVICE
  - UCT Communication Resource, [159](#)
- UCT\_IFACE\_OPEN\_MODE\_SOCKADDR\_CLIENT
  - UCT Communication Resource, [159](#)
- UCT\_IFACE\_OPEN\_MODE\_SOCKADDR\_SERVER
  - UCT Communication Resource, [159](#)
- UCT\_IFACE\_PARAM\_FIELD\_AM\_ALIGN\_OFFSET
  - UCT Communication Resource, [160](#)
- UCT\_IFACE\_PARAM\_FIELD\_AM\_ALIGNMENT
  - UCT Communication Resource, [160](#)
- UCT\_IFACE\_PARAM\_FIELD\_ASYNC\_EVENT\_ARG
  - UCT Communication Resource, [160](#)
- UCT\_IFACE\_PARAM\_FIELD\_ASYNC\_EVENT\_CB
  - UCT Communication Resource, [160](#)
- UCT\_IFACE\_PARAM\_FIELD\_CPU\_MASK
  - UCT Communication Resource, [159](#)
- UCT\_IFACE\_PARAM\_FIELD\_DEVICE
  - UCT Communication Resource, [159](#)
- UCT\_IFACE\_PARAM\_FIELD\_ERR\_HANDLER
  - UCT Communication Resource, [159](#)
- UCT\_IFACE\_PARAM\_FIELD\_ERR\_HANDLER\_ARG
  - UCT Communication Resource, [159](#)
- UCT\_IFACE\_PARAM\_FIELD\_ERR\_HANDLER\_FLAGS
  - UCT Communication Resource, [159](#)
- UCT\_IFACE\_PARAM\_FIELD\_FEATURES
  - UCT Communication Resource, [160](#)
- UCT\_IFACE\_PARAM\_FIELD\_HW\_TM\_EAGER\_ARG
  - UCT Communication Resource, [159](#)
- UCT\_IFACE\_PARAM\_FIELD\_HW\_TM\_EAGER\_CB
  - UCT Communication Resource, [160](#)
- UCT\_IFACE\_PARAM\_FIELD\_HW\_TM\_RNDV\_ARG
  - UCT Communication Resource, [160](#)
- UCT\_IFACE\_PARAM\_FIELD\_HW\_TM\_RNDV\_CB
  - UCT Communication Resource, [160](#)
- UCT\_IFACE\_PARAM\_FIELD\_KEEPALIVE\_INTERVAL
  - UCT Communication Resource, [160](#)
- UCT\_IFACE\_PARAM\_FIELD\_OPEN\_MODE
  - UCT Communication Resource, [159](#)
- UCT\_IFACE\_PARAM\_FIELD\_RX\_HEADROOM
  - UCT Communication Resource, [159](#)
- UCT\_IFACE\_PARAM\_FIELD\_SOCKADDR
  - UCT Communication Resource, [159](#)
- UCT\_IFACE\_PARAM\_FIELD\_STATS\_ROOT
  - UCT Communication Resource, [159](#)
- uct\_iface\_params, [144](#)
- uct\_iface\_params.mode, [145](#)
- uct\_iface\_params.mode.device, [146](#)
- uct\_iface\_params.mode.sockaddr, [146](#)
- uct\_iface\_params\_field
  - UCT Communication Resource, [159](#)
- uct\_iface\_params\_t
  - UCT Communication Resource, [152](#)
- uct\_iface\_progress
  - UCT Communication Resource, [175](#)
- uct\_iface\_progress\_disable
  - UCT Communication Resource, [174](#)
- uct\_iface\_progress\_enable
  - UCT Communication Resource, [174](#)

- uct\_iface\_query
  - UCT Communication Resource, [166](#)
- uct\_iface\_reject
  - UCT client-server operations, [227](#)
- uct\_iface\_release\_desc
  - UCT Active messages, [198](#)
- uct\_iface\_set\_am\_handler
  - UCT Active messages, [197](#)
- uct\_iface\_set\_am\_tracer
  - UCT Active messages, [198](#)
- uct\_iface\_tag\_recv\_cancel
  - UCT Tag matching operations, [212](#)
- uct\_iface\_tag\_recv\_zcopy
  - UCT Tag matching operations, [212](#)
- uct\_iov, [149](#)
- uct\_iov\_t
  - UCT Communication Resource, [154](#)
- uct\_listener\_attr, [217](#)
- uct\_listener\_attr\_field
  - UCT client-server operations, [224](#)
- UCT\_LISTENER\_ATTR\_FIELD\_SOCKADDR
  - UCT client-server operations, [224](#)
- uct\_listener\_attr\_t
  - UCT Communication Resource, [153](#)
- uct\_listener\_create
  - UCT client-server operations, [230](#)
- uct\_listener\_destroy
  - UCT client-server operations, [230](#)
- uct\_listener\_h
  - UCT Communication Resource, [154](#)
- UCT\_LISTENER\_PARAM\_FIELD\_BACKLOG
  - UCT client-server operations, [224](#)
- UCT\_LISTENER\_PARAM\_FIELD\_CONN\_REQUEST\_CB
  - UCT client-server operations, [224](#)
- UCT\_LISTENER\_PARAM\_FIELD\_USER\_DATA
  - UCT client-server operations, [224](#)
- uct\_listener\_params, [217](#)
- uct\_listener\_params\_field
  - UCT client-server operations, [224](#)
- uct\_listener\_params\_t
  - UCT Communication Resource, [154](#)
- uct\_listener\_query
  - UCT client-server operations, [231](#)
- uct\_listener\_reject
  - UCT client-server operations, [231](#)
- UCT\_MADV\_NORMAL
  - UCT Memory Domain, [187](#)
- UCT\_MADV\_WILLNEED
  - UCT Memory Domain, [187](#)
- uct\_md\_attr, [183](#)
- uct\_md\_attr.cap, [183](#)
- uct\_md\_attr\_t
  - UCT Communication Resource, [152](#)
- uct\_md\_close
  - UCT Communication Resource, [163](#)
- uct\_md\_config\_read
  - UCT Memory Domain, [192](#)
- uct\_md\_config\_t
  - UCT Communication Resource, [151](#)
- uct\_md\_detect\_memory\_type
  - UCT Memory Domain, [190](#)
- UCT\_MD\_FLAG\_ADVISE
  - UCT Memory Domain, [186](#)
- UCT\_MD\_FLAG\_ALLOC
  - UCT Memory Domain, [186](#)
- UCT\_MD\_FLAG\_FIXED
  - UCT Memory Domain, [186](#)
- UCT\_MD\_FLAG\_INVALIDATE
  - UCT Memory Domain, [186](#)
- UCT\_MD\_FLAG\_NEED\_MEMH
  - UCT Memory Domain, [186](#)
- UCT\_MD\_FLAG\_NEED\_RKEY
  - UCT Memory Domain, [186](#)
- UCT\_MD\_FLAG\_REG
  - UCT Memory Domain, [186](#)
- UCT\_MD\_FLAG\_RKEY\_PTR
  - UCT Memory Domain, [186](#)
- UCT\_MD\_FLAG\_SOCKADDR
  - UCT Memory Domain, [186](#)
- uct\_md\_h
  - UCT Communication Resource, [151](#)
- uct\_md\_iface\_config\_read
  - UCT Communication Resource, [164](#)
- uct\_md\_is\_sockaddr\_accessible
  - UCT Memory Domain, [192](#)
- UCT\_MD\_MEM\_ACCESS\_ALL
  - UCT Memory Domain, [187](#)
- UCT\_MD\_MEM\_ACCESS\_LOCAL\_READ
  - UCT Memory Domain, [187](#)
- UCT\_MD\_MEM\_ACCESS\_LOCAL\_WRITE
  - UCT Memory Domain, [187](#)
- UCT\_MD\_MEM\_ACCESS\_REMOTE\_ATOMIC
  - UCT Memory Domain, [187](#)
- UCT\_MD\_MEM\_ACCESS\_REMOTE\_GET
  - UCT Memory Domain, [187](#)
- UCT\_MD\_MEM\_ACCESS\_REMOTE\_PUT
  - UCT Memory Domain, [187](#)
- UCT\_MD\_MEM\_ACCESS\_RMA
  - UCT Memory Domain, [187](#)
- uct\_md\_mem\_advise
  - UCT Memory Domain, [189](#)
- uct\_md\_mem\_attr, [183](#)
- uct\_md\_mem\_attr\_field
  - UCT Memory Domain, [187](#)
- UCT\_MD\_MEM\_ATTR\_FIELD\_ALLOC\_LENGTH
  - UCT Memory Domain, [187](#)
- UCT\_MD\_MEM\_ATTR\_FIELD\_BASE\_ADDRESS
  - UCT Memory Domain, [187](#)
- UCT\_MD\_MEM\_ATTR\_FIELD\_MEM\_TYPE
  - UCT Memory Domain, [187](#)
- UCT\_MD\_MEM\_ATTR\_FIELD\_SYS\_DEV
  - UCT Memory Domain, [187](#)
- uct\_md\_mem\_attr\_field\_t
  - UCT Memory Domain, [185](#)
- uct\_md\_mem\_attr\_t
  - UCT Memory Domain, [185](#)

- uct\_md\_mem\_dereg
  - UCT Memory Domain, [190](#)
- UCT\_MD\_MEM\_FLAG\_FIXED
  - UCT Memory Domain, [187](#)
- UCT\_MD\_MEM\_FLAG\_HIDE\_ERRORS
  - UCT Memory Domain, [187](#)
- UCT\_MD\_MEM\_FLAG\_LOCK
  - UCT Memory Domain, [187](#)
- UCT\_MD\_MEM\_FLAG\_NONBLOCK
  - UCT Memory Domain, [187](#)
- uct\_md\_mem\_flags
  - UCT Memory Domain, [186](#)
- uct\_md\_mem\_query
  - UCT Memory Domain, [188](#)
- uct\_md\_mem\_reg
  - UCT Memory Domain, [189](#)
- uct\_md\_mkey\_pack
  - UCT Memory Domain, [193](#)
- uct\_md\_open
  - UCT Communication Resource, [163](#)
- uct\_md\_ops\_t
  - UCT Communication Resource, [151](#)
- uct\_md\_query
  - UCT Memory Domain, [188](#)
- uct\_md\_query\_tl\_resources
  - UCT Communication Resource, [163](#)
- uct\_md\_resource\_desc, [140](#)
- uct\_md\_resource\_desc\_t
  - UCT Communication Resource, [150](#)
- uct\_md\_t
  - UCT Communication Resource, [152](#)
- uct\_mem\_advice\_t
  - UCT Memory Domain, [187](#)
- uct\_mem\_alloc
  - UCT Memory Domain, [190](#)
- UCT\_MEM\_ALLOC\_PARAM\_FIELD\_ADDRESS
  - UCT Memory Domain, [188](#)
- UCT\_MEM\_ALLOC\_PARAM\_FIELD\_FLAGS
  - UCT Memory Domain, [188](#)
- UCT\_MEM\_ALLOC\_PARAM\_FIELD\_MDS
  - UCT Memory Domain, [188](#)
- UCT\_MEM\_ALLOC\_PARAM\_FIELD\_MEM\_TYPE
  - UCT Memory Domain, [188](#)
- UCT\_MEM\_ALLOC\_PARAM\_FIELD\_NAME
  - UCT Memory Domain, [188](#)
- uct\_mem\_alloc\_params\_field\_t
  - UCT Memory Domain, [187](#)
- uct\_mem\_alloc\_params\_t, [184](#)
- uct\_mem\_alloc\_params\_t.mds, [185](#)
- uct\_mem\_free
  - UCT Memory Domain, [192](#)
- uct\_mem\_h
  - UCT Communication Resource, [151](#)
- uct\_msg\_flags
  - UCT Active messages, [196](#)
- uct\_pack\_callback\_t
  - UCT Communication Resource, [156](#)
- uct\_pending\_callback\_t
  - UCT Communication Resource, [155](#)
- uct\_pending\_purge\_callback\_t
  - UCT Communication Resource, [156](#)
- uct\_pending\_req, [149](#)
- uct\_pending\_req\_t
  - UCT Communication Resource, [152](#)
- UCT\_PROGRESS\_RECV
  - UCT Communication Resource, [158](#)
- UCT\_PROGRESS\_SEND
  - UCT Communication Resource, [158](#)
- UCT\_PROGRESS\_THREAD\_SAFE
  - UCT Communication Resource, [158](#)
- uct\_progress\_types
  - UCT Communication Resource, [158](#)
- uct\_query\_components
  - UCT Communication Resource, [161](#)
- uct\_release\_component\_list
  - UCT Communication Resource, [162](#)
- uct\_release\_tl\_resource\_list
  - UCT Communication Resource, [164](#)
- uct\_rkey\_bundle, [184](#)
- uct\_rkey\_bundle\_t
  - UCT Memory Domain, [186](#)
- uct\_rkey\_ctx\_h
  - UCT Communication Resource, [151](#)
- uct\_rkey\_ptr
  - UCT Memory Domain, [194](#)
- uct\_rkey\_release
  - UCT Memory Domain, [194](#)
- uct\_rkey\_t
  - UCT Communication Resource, [151](#)
- uct\_rkey\_unpack
  - UCT Memory Domain, [193](#)
- UCT\_SEND\_FLAG\_PEER\_CHECK
  - UCT Active messages, [197](#)
- UCT\_SEND\_FLAG\_SIGNALED
  - UCT Active messages, [197](#)
- UCT\_SOCKADDR\_ACC\_LOCAL
  - UCT Memory Domain, [186](#)
- UCT\_SOCKADDR\_ACC\_REMOTE
  - UCT Memory Domain, [186](#)
- uct\_sockaddr\_accessibility\_t
  - UCT Memory Domain, [186](#)
- uct\_sockaddr\_conn\_request\_callback\_t
  - UCT client-server operations, [221](#)
- uct\_tag\_context, [247](#)
  - completed\_cb, [248](#)
  - priv, [249](#)
  - rndv\_cb, [248](#)
  - tag\_consumed\_cb, [248](#)
- uct\_tag\_context\_t
  - UCT Communication Resource, [154](#)
- UCT\_TAG\_RECV\_CB\_INLINE\_DATA
  - UCT Communication Resource, [161](#)
- uct\_tag\_t
  - UCT Communication Resource, [154](#)
- uct\_tag\_unexp\_eager\_cb\_t
  - UCT Tag matching operations, [206](#)



- uct\_tag\_unexp\_rndv\_cb\_t
  - UCT Tag matching operations, [207](#)
- uct\_tl\_resource\_desc, [141](#)
- uct\_tl\_resource\_desc\_t
  - UCT Communication Resource, [150](#)
- uct\_unpack\_callback\_t
  - UCT Communication Resource, [156](#)
- uct\_worker\_cb\_id\_t
  - UCT Communication Resource, [154](#)
- uct\_worker\_create
  - UCT Communication Context, [177](#)
- uct\_worker\_destroy
  - UCT Communication Context, [177](#)
- uct\_worker\_h
  - UCT Communication Resource, [152](#)
- uct\_worker\_progress
  - UCT Communication Context, [179](#)
- uct\_worker\_progress\_register\_safe
  - UCT Communication Context, [177](#)
- uct\_worker\_progress\_unregister\_safe
  - UCT Communication Context, [178](#)
- Unified Communication Protocol (UCP) API, [11](#)
- Unified Communication Services (UCS) API, [238](#)
- Unified Communication Transport (UCT) API, [135](#)
- unpack
  - UCP Data type routines, [132](#)