



Event Chaining™ Enables Real-Time Systems to Respond to Multiple Real-Time Events More Efficiently

*Innovative function callback capability
permits responsiveness, while reducing overhead*

Introduction

Express Logic's ThreadX® RTOS provides several advanced technology features that can be beneficial during the development stage as well as during run-time. These features include real-time Event-Chaining™, Application Notification "Callback" Functions, and many others. We will investigate the Event Chaining and Notification Callback Function topics in this paper.

Event-Chaining

Event-Chaining is a technique that enables a single RTOS action based on the occurrence of independent events. This is particularly useful in activating an application thread that is suspended on two or more resources. For example, suppose a single thread is responsible for processing messages from 5 or more message queues, and must suspend when no messages are available. Such resources might be messages being awaited in one or more queues, a semaphore from one of several cooperating threads, or an event in an event flags group. In general, Event-Chaining results in fewer threads, less overhead, and smaller RAM requirements. It also provides a highly flexible mechanism to handle synchronization requirements of more complex systems. Implementing this technique is a three-step process as follows:

1. Register one or more notification callback functions. We'll explain notification callback functions below.
2. The event occurs, and the registered notification callback function is automatically invoked. Each such function typically contains a `tx_semaphore_put` service call, which increments a "gatekeeper" semaphore which communicates to a waiting thread that a particular event has occurred. However, many other service calls could be used.
3. A thread, suspended on the "gatekeeper" semaphore mentioned above, is activated. Getting this semaphore signifies that one of the events in question has occurred and the thread determines which, and then performs the actions appropriate for that event.

There are three types of Event-Chaining available:

1. Queue Event-Chaining
2. Semaphore Event Chaining
3. Event Flags Group Event Chaining

A typical use for Event-Chaining is to create a mechanism for a thread to suspend on two or more objects. For example, this technique can be used to permit a thread to suspend on any of the following situations:

- Suspend on a queue, a semaphore, and an event flags group
- Suspend on a queue or a semaphore
- Suspend on a queue or an event flags group
- Suspend on two queues
- Suspend on three queues
- Suspend on four queues

An important advantage of the Event-Chaining technique is that one or more threads waiting for an event to occur can be activated automatically when the event occurs. In general, this technique will reduce the number of threads needed to respond to an event and will reduce the associated resources and overhead required for processing systems of this nature.

In this paper, we will focus on Queue Event Chaining. The principles are the same across all three types, so the process described below for Queue Event Chaining can be replicated for either of the other two types.

Notification Callback Functions

Some applications may find it advantageous to be notified whenever a message is placed on a queue. ThreadX provides this ability through the *tx_queue_send_notify* service. This service registers the supplied application notification function with the specified queue. ThreadX will subsequently invoke this application notification function whenever a message is sent to the queue. The processing within the application notification function is determined by the application; however, it typically consists of resuming the appropriate thread for processing the new message.

For example, the *tx_queue_send_notify(&my_queue, queue_notify)* function registers a callback function (“queue_notify”) that would be called every time a message is sent to the specified queue (“my_queue”).

Queue Event-Chaining

Suppose a single thread is responsible for processing messages from five different queues and must also suspend when no messages are available. This is easily accomplished by registering an application notification function for each queue and introducing an additional counting semaphore. Specifically, the application notification function performs a *tx_semaphore_put* whenever it is called (the semaphore count represents the total number of messages in all five queues). The processing thread suspends on this semaphore via the *tx_semaphore_get* service. When the semaphore is available (in this case, when a message is available!), the processing thread is resumed. It then interrogates each queue for a message, processes the found message, and performs another *tx_semaphore_get* to wait for the next message. Accomplishing this without event-chaining is quite difficult and likely would require more threads and/or additional application code. As noted, implementing Event-Chaining is a multiple-step process.

Figure 1 contains a template that illustrates the components involved for Event-Chaining with a message queue.

<p>1. Initialization</p> <pre>TX_QUEUE my_queue; TX_SEMAPHORE gatekeeper; ULONG my_message[4]; tx_queue_send_notify (&my_queue, queue_notify); void queue_notify (TX_QUEUE *my_queue);</pre>	<pre>/* The queue, semaphore, and message declarations, the registration of the notification callback function, and the prototype for the notification callback function are usually placed in the tx_application_define function, which is part of the initialization process */</pre>
<p>2a. Event Occurrence</p> <pre>tx_queue_send (&my_queue, my_message, TX_NO_WAIT);</pre>	<pre>/* A message is sent to the queue somewhere in the application. Whenever a message is sent to this queue, the notification callback function is automatically invoked, thus causing the semaphore gatekeeper to be incremented. */</pre>
<p>2b. Notification Callback Function Called</p> <pre>void queue_notify (TX_QUEUE *my_queue) { tx_semaphore_put (&gatekeeper); }</pre>	<pre>/* Notification callback function to increment the "gatekeeper" semaphore is called whenever a message has been sent to my_queue */</pre>
<p>3. Thread Activation</p> <pre>tx_semaphore_get (&gatekeeper, TX_WAIT_FOREVER);</pre>	<pre>/* Somewhere in the application, a thread suspends on semaphore gatekeeper, which is equivalent to waiting for a message to appear on the queue */</pre>

Figure 1. Template for Event-Chaining with a message queue

Sample System Using Event-Chaining

We will now study a complete sample system that uses Event-Chaining. The system is characterized in Figure 2.

All the thread suspension examples in previous chapters involved one thread waiting on one object, such as a mutex, a counting semaphore, an event flags group, or a message queue. In this sample system, we have 2 threads waiting on multiple objects. Specifically, threads wait for a message to appear on either *queue_1* or *queue_2*.

Speedy_thread has priority 5 and *slow_thread* has priority 15. We will use Event-Chaining to automatically increment the counting semaphore named “gatekeeper” whenever a message is sent to either *queue_1* or *queue_2*. We use two application timers to send messages to *queue_1* or *queue_2* at periodic time intervals and the threads wait for a message to appear.

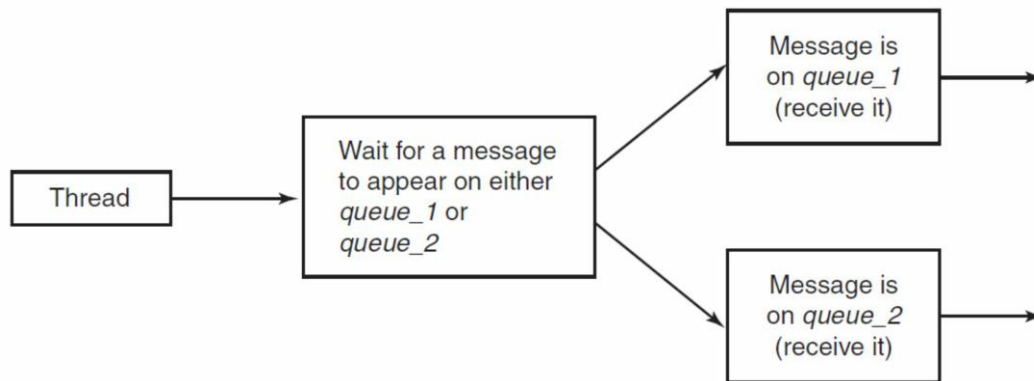


Figure 2. Multiple object suspension problem

Figure 3 contains a description of the two activities for *speedy_thread*.

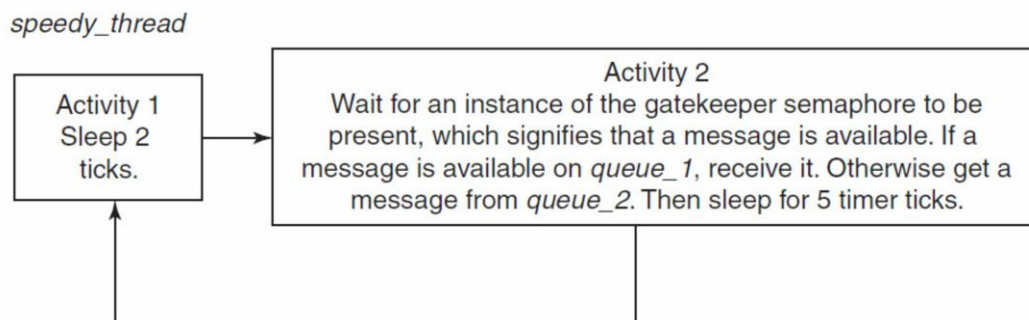


Figure 3. *speedy_thread* activities

Figure contains a description of the two activities for *slow_thread*.

slow_thread

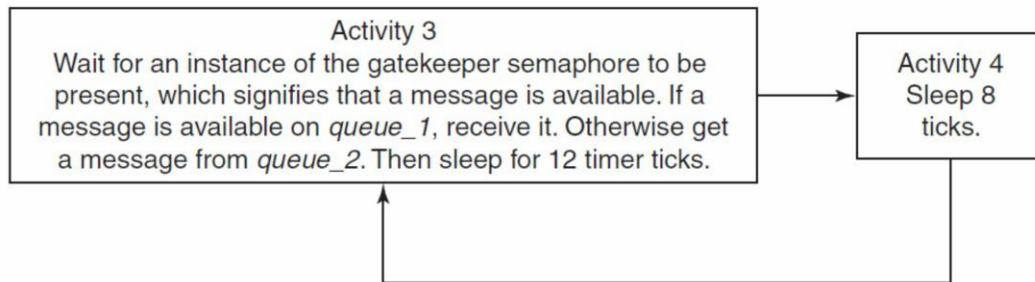


Figure 4. *slow_thread* activities

Listing for *sample_system.c*

The sample system named *sample_system.c* appears below; line numbers have been added for easy reference.

```

000 /* sample_system.c
001
002     Create two threads, one byte pool, two message queues, three timers, and
003     one counting semaphore. This is an example of multiple object suspension
004     using Event-Chaining, i.e., speedy_thread and slow_thread wait for a
005     message to appear on either of two queues */
006
007
008     /*****
009     /*      Declarations, Definitions, and Prototypes      */
010     /*****
011
012     #include "tx_api.h"
013     #include <stdio.h>
014
015     #define STACK_SIZE          1024
016     #define BYTE_POOL_SIZE     9120
017     #define NUMBER_OF_MESSAGES 100
018     #define MESSAGE_SIZE       TX_1_ULONG
019     #define QUEUE_SIZE         MESSAGE_SIZE*sizeof(ULONG)*NUMBER_OF_MESSAGES
020
021
022     /* Define the ThreadX object control blocks... */
023
024     TX_THREAD speedy_thread; /* higher priority thread */
025     TX_THREAD slow_thread; /* lower priority thread */
026
027     TX_BYTE_POOL my_byte_pool; /* byte pool for stacks and queues */
028     TX_SEMAPHORE gatekeeper; /* indicate how many objects available */
029
030     TX_QUEUE queue_1; /* queue for multiple object suspension */
031     TX_QUEUE queue_2; /* queue for multiple object suspension */
032
  
```

```

033 TX_TIMER stats_timer; /* generate statistics at intervals */
034 TX_TIMER queue_timer_1; /* send message to queue_1 at intervals */
035 TX_TIMER queue_timer_2; /* send message to queue_2 at intervals */
036
037 /* Variables needed to get info about the message queue */
038 CHAR *info_queue_name;
039 TX_THREAD *first_suspended;
040 TX_QUEUE *next_queue;
041 ULONG enqueued_1=0, enqueued_2=0, suspended_count=0, available_storage=0;
042
043 /* Define the variables used in the sample application... */
044 ULONG speedy_thread_counter=0, total_speedy_time=0;
045 ULONG slow_thread_counter=0, total_slow_time=0;
046 ULONG send_message_1[TX_1_ULONG]={0X0}, send_message_2[TX_1_ULONG]={0X0};
047 ULONG receive_message_1[TX_1_ULONG], receive_message_2[TX_1_ULONG];
048
049 /* speedy_thread and slow_thread entry function prototypes */
050 void speedy_thread_entry(ULONG thread_input);
051 void slow_thread_entry(ULONG thread_input);
052
053 /* timer entry function prototypes */
054 void queue_timer_1_entry(ULONG thread_input);
055 void queue_timer_2_entry(ULONG thread_input);
056 void print_stats(ULONG);
057
058 /* event notification function prototypes used for Event-Chaining */
059 void queue_1_send_notify(TX_QUEUE *queue_1_ptr);
060 void queue_2_send_notify(TX_QUEUE *queue_2_ptr);
061
062
063 /*****
064 /*          Main Entry Point          */
065 /*****
066
067 /* Define main entry point. */
068
069 int main()
070 {
071     /* Enter the ThreadX kernel. */
072     tx_kernel_enter();
073 }
074
075
076 /*****
077 /*          Application Definitions          */
078 /*****
079
080
081 /* Define what the initial system looks like. */
082
083 void tx_application_define(void *first_unused_memory)
084 {
085
086     CHAR *speedy_stack_ptr;

```

```

087 CHAR *slow_stack_ptr;
088 CHAR *queue_1_ptr;
089 CHAR *queue_2_ptr;
090
091 /* Create a byte memory pool from which to allocate the thread stacks. */
092 tx_byte_pool_create(&my_byte_pool, "my_byte_pool",
093     first_unused_memory, BYTE_POOL_SIZE);
094
095 /* Create threads, queues, the semaphore, timers, and register functions
096     for Event-Chaining */
097
098 /* Allocate the stack for speedy_thread. */
099 tx_byte_allocate(&my_byte_pool, (VOID **) &speedy_stack_ptr, STACK_SIZE,
100     TX_NO_WAIT);
101
102 /* Create speedy_thread. */
103 tx_thread_create(&speedy_thread, "speedy_thread", speedy_thread_entry, 0,
104     speedy_stack_ptr, STACK_SIZE, 5, 5, TX_NO_TIME_SLICE,
105     TX_AUTO_START);
106
107 /* Allocate the stack for slow_thread. */
108 tx_byte_allocate(&my_byte_pool, (VOID **) &slow_stack_ptr, STACK_SIZE,
109     TX_NO_WAIT);
110
111 /* Create slow_thread */
112 tx_thread_create(&slow_thread, "slow_thread", slow_thread_entry, 1,
113     slow_stack_ptr, STACK_SIZE, 15, 15, TX_NO_TIME_SLICE,
114     TX_AUTO_START);
115
116 /* Create the message queues used by both threads. */
117 tx_byte_allocate(&my_byte_pool, (VOID **) &queue_1_ptr,
118     QUEUE_SIZE, TX_NO_WAIT);
119
120 tx_queue_create (&queue_1, "queue_1", MESSAGE_SIZE,
121     Queue_1_ptr, QUEUE_SIZE);
122
123 tx_byte_allocate(&my_byte_pool, (VOID **) &queue_2_ptr,
124     QUEUE_SIZE, TX_NO_WAIT);
125
126 tx_queue_create (&queue_2, "queue_2", MESSAGE_SIZE,
127     Queue_2_ptr, QUEUE_SIZE);
128
129 /* Create the gatekeeper semaphore that counts the available objects */
130 tx_semaphore_create (&gatekeeper, "gatekeeper", 0);
131
132 /* Create and activate the stats timer */
133 tx_timer_create (&stats_timer, "stats_timer", print_stats,
134     0x1234, 500, 500, TX_AUTO_ACTIVATE);
135
136 /* Create and activate the timer to send messages to queue_1 */
137 tx_timer_create (&queue_timer_1, "queue_timer", queue_timer_1_entry,
138     0x1234, 12, 12, TX_AUTO_ACTIVATE);
139
140 /* Create and activate the timer to send messages to queue_2 */

```

```

141     tx_timer_create (&queue_timer_2, "queue_timer", queue_timer_2_entry,
142                     0x1234, 9, 9, TX_AUTO_ACTIVATE);
143
144     /* Register the function to increment the gatekeeper semaphore when a
145        message is sent to queue_1 */
146     tx_queue_send_notify(&queue_1, queue_1_send_notify);
147
148     /* Register the function to increment the gatekeeper semaphore when a
149        message is sent to queue_2 */
150     tx_queue_send_notify(&queue_2, queue_1_send_notify);
151 }
152
153
154 /*****
155  */
156 /*****
157
158
159 /* Entry function definition of speedy_thread
160    it has a higher priority than slow_thread */
161
162 void speedy_thread_entry(ULONG thread_input)
163 {
164
165     ULONG start_time, cycle_time=0, current_time=0;
166     UINT status;
167
168     /* This is the higher priority speedy_thread */
169
170     while(1)
171     {
172         /* Get the starting time for this cycle */
173         start_time = tx_time_get();
174
175         /* Activity 1: 2 ticks. */
176         tx_thread_sleep(2);
177
178         /* Activity 2: 5 ticks. */
179         /* wait for a message to appear on either one of the two queues */
180         tx_semaphore_get (&gatekeeper, TX_WAIT_FOREVER);
181
182         /* Determine whether a message queue_1 or queue_2 is available */
183         status = tx_queue_receive (&queue_1, receive_message_1, TX_NO_WAIT);
184
185         if (status == TX_SUCCESS)
186             ; /* A message on queue_1 has been found-process */
187         else
188             /* Receive a message from queue_2 */
189             tx_queue_receive (&queue_2, receive_message_2, TX_WAIT_FOREVER);
190
191         tx_thread_sleep(5);
192
193         /* Increment the thread counter and get timing info */
194         speedy_thread_counter++;

```



```

195     current_time = tx_time_get();
196     cycle_time = current_time-start_time;
197     total_speedy_time = total_speedy_time + cycle_time;
198 }
199 }
200
201 /*****
202
203 /* Entry function definition of slow_thread
204    it has a lower priority than speedy_thread */
205
206 void slow_thread_entry(ULONG thread_input)
207 {
208
209     ULONG start_time, current_time=0, cycle_time=0;
210     UINT status;
211
212
213     while(1)
214     {
215         /* Get the starting time for this cycle */
216         start_time=tx_time_get();
217
218         /* Activity 3-sleep 12 ticks. */
219         /* wait for a message to appear on either one of the two queues */
220         tx_semaphore_get (&gatekeeper, TX_WAIT_FOREVER);
221
222         /* Determine whether a message queue_1 or queue_2 is available */
223         status = tx_queue_receive (&queue_1, receive_message_1, TX_NO_WAIT);
224
225         if (status == TX_SUCCESS)
226             ; /* A message on queue_1 has been found-process */
227         else
228             /* Receive a message from queue_2 */
229             tx_queue_receive (&queue_2, receive_message_2, TX_WAIT_FOREVER);
230
231         tx_thread_sleep(12);
232
233
234         /* Activity 4: 8 ticks. */
235         tx_thread_sleep(8);
236
237         /* Increment the thread counter and get timing info */
238         slow_thread_counter++;
239
240         current_time = tx_time_get();
241         cycle_time = current_time-start_time;
242         total_slow_time = total_slow_time + cycle_time;
243     }
244 }
245
246 /*****
247 /* print statistics at specified times */
248 Void print_stats (ULONG inval)

```

```

249 {
250     ULONG current_time, avg_slow_time, avg_speedy_time;
251
252     If ((speedy_thread_counter>0) && (slow_thread_counter>0))
253     {
254         current_time = tx_time_get();
255         avg_slow_time = total_slow_time/slow_thread_counter;
256         avg_speedy_time = total_speedy_time/speedy_thread_counter;
257         tx_queue_info_get (&queue_1, &info_queue_name, &enqueued_1,
258             &available_storage, &first_suspended,
259             &suspended_count, &next_queue);
260         tx_queue_info_get (&queue_2, &info_queue_name, &enqueued_2,
261             &available_storage, &first_suspended,
262             &suspended_count, &next_queue);
263         printf("\nEvent-Chaining: 2 threads waiting for 2 queues\n\n");
264         printf(" Current Time:           %lu\n", current_time);
265         printf(" speedy_thread counter: %lu\n", speedy_thread_counter);
266         printf(" speedy_thread avg time: %lu\n", avg_speedy_time);
267         printf(" slow_thread counter:   %lu\n", slow_thread_counter);
268         printf(" slow_thread avg time:  %lu\n", avg_slow_time);
269         printf(" total # queue_1 messages sent:
270 %lu\n", send_message_1[TX_1_ULONG-1]);
271         printf(" total # queue_2 messages sent:
272 %lu\n", send_message_2[TX_1_ULONG-1]);
273         printf(" current # messages in queue_1:
274 %lu\n", enqueued_1);
275         printf(" current # messages in queue_2: %lu\n\n", enqueued_2);
276     }
277     else printf("Bypassing print_stats function, Current Time: %lu\n",
278         tx_time_get());
279 }
280
281
282
283 /*****
284 /* Send a message to queue_1 at specified times */
285 void queue_timer_1_entry (ULONG inval)
286 {
287
288     /* Send a message to queue_1 using the multiple object suspension approach
289     */
290     /* The gatekeeper semaphore keeps track of how many objects are available
291     via the notification function */
292     send_message_1[TX_1_ULONG-1]++;
293     tx_queue_send (&queue_1, send_message_1, TX_NO_WAIT);
294 }
295
296 /*****
297 /* Send a message to the queue at specified times */
298 void queue_timer_2_entry (ULONG inval)
299 {
300

```

```

301     /* Send a message to queue_2 using the multiple object suspension approach */
302     /* The gatekeeper semaphore keeps track of how many objects are available
303        via the notification function */
304     send_message_2[TX_1_ULONG--1]++;
305     tx_queue_send (&queue_2, send_message_2, TX_NO_WAIT);
306
307 }
308
309 /*****
310  /* Notification function to increment gatekeeper semaphore
311     whenever a message has been sent to queue_1 */
312 void queue_1_send_notify(TX_QUEUE *queue_ptr_1)
313 {
314     tx_semaphore_put (&gatekeeper);
315 }
316
317 /*****
318  /* Notification function to increment gatekeeper semaphore
319     whenever a message has been sent to queue_2 */
320 void queue_2_send_notify(TX_QUEUE *queue_ptr_2)
321 {
322     Tx_semaphore_put (&gatekeeper);
323 }

```

END Example code.

Figure contains several comments about this listing, using the line numbers as references.

Lines	Comments
024 through 035	Declaration of system resources including threads, byte pool, semaphore, queues, and timers
037 through 047	Declaration of variables used in the system including parameters for the queue info get services
049 through 060	Declaration of prototypes for thread entry functions, timer entry function, and event notification functions
116 through 127	Creation of the two queues used for multiple object suspension
129 and 130	Creation of the gatekeeper semaphore used for Event-Chaining
132 through 142	Creation of the timer for display statistics at periodic intervals, and creation of the two timers to send messages to the queues at various intervals
144 through 150	Registration of the two functions that increment the gatekeeper semaphore whenever messages are sent to the queues
159 through 199	Entry function for Speedy Thread; lines 178 through 191 contain the implementation of Activity 2
203 through 244	Entry function for Slow Thread; lines 218 through 231 contain the implementation of Activity 3
247 through 276	Entry function for timer print stats, which includes calculating average cycle time, number of times through each cycle, and info get for the two queues
281 through 304	Entry functions for timers to send messages to queue_1 and queue_2 at periodic intervals
307 through 320	Entry functions for the notification callback functions; these functions increment semaphore gatekeeper whenever a message is send to either queue_1 or queue_2; these functions are essential to the Event-Chaining technique

Figure 5. Comments about sample system listing

Following is some sample output for this system after it has executed for 500 timer ticks, using information obtained from the *tx_queue_info_get* service:

Event-Chaining: 2 threads waiting for 2 queues

Current Time:	500
<i>speedy_thread</i> counter:	69
<i>speedy_thread</i> avg time:	7
<i>slow_thread</i> counter:	24
<i>slow_thread</i> avg time:	20
total # <i>queue_1</i> messages sent:	41
total # <i>queue_2</i> messages sent:	55
current # messages in <i>queue_1</i> :	0
current # messages in <i>queue_2</i> :	1

Conclusion

Event-Chaining is one technique that uses notification callback functions to reduce the number of threads required to manage responses to multiple events in a real-time system. For more information about Event-Chaining, Callback Functions, or any of the other advanced technology features of Express Logic's ThreadX RTOS, please send an email to: info@expresslogic.com, or call 1-888-THREADX.