

Linux Kernel Programming

Bottom-Halves and Deferring Work

Pierre Olivier

Systems Software Research Group @ Virginia Tech

March 2, 2017

Outline

- 1 Bottom-halves: general information
- 2 Softirqs
- 3 Tasklets
- 4 Workqueues
- 5 Using the right bottom-half and misc. information
- 6 Additional sources of information

Outline

- 1 **Bottom-halves: general information**
- 2 Softirqs
- 3 Tasklets
- 4 Workqueues
- 5 Using the right bottom-half and misc. information
- 6 Additional sources of information

Bottom-halves: general information

Presentation

- ▶ **Top-halves (interrupt handlers) must be as quick as possible**
 - ▶ Because they interrupt kernel/user code
 - ▶ Affects performance
 - ▶ Because they run with one/all lines disabled
 - ▶ Processing network traffic should not prevent the kernel from receiving keystrokes
- ▶ **Top-halves run in interrupt context: they cannot block**
 - ▶ Limit what they can do
- ▶ When processing an interrupt, ***the less-critical part of the work is deferred to a bottom-half***
 - ▶ Runs *later* (regarding the moment the actual interrupt occurs)

Bottom-halves: general information

Which part of the job in which half, reason of being


- ▶ **Work is time sensitive?** → top-half
- ▶ **Work is related to the hardware?** → top-half
- ▶ **Work should not be interrupted by another/the same interrupt?** → top-half
- ▶ **Everything else:** → **bottom-half**

- ▶ Bottom-halves run *later*
 - ▶ They generally run very soon after the actual interrupt
 - ▶ Crucial point is to run with interrupts enabled

Bottom-halves: general information

A bit of history

- ▶ Old "**Bottom-Half**" (BH) mechanism
 - ▶ 32 of them, globally synchronized: only one running in the system at the same time
 - ▶ Performance bottleneck
- ▶ **Task queues**: queues of function pointers (handlers)
 - ▶ Handlers run at various time depending on which queue they are on
 - ▶ Not sufficient for performance critical subsystems (ex: networking)
- ▶ **BH and task queues were removed in 2.5**



Outline

- 1 Bottom-halves: general information
- 2 Softirqs**
- 3 Tasklets
- 4 Workqueues
- 5 Using the right bottom-half and misc. information
- 6 Additional sources of information

Softirqs

Implementing softirqs, softirq execution

▶ Softirqs

- ▶ Bottom-half with the best performance
- ▶ `kernel/softirq.c`
- ▶ Rarely used directly (tasklets instead)
 - ▶ However, tasklets are build upon softirqs

▶ `include/linux/
interrupt.h:`

```
1 struct softirq_action {  
2     void (*action)(struct softirq_action *);  
3 }
```

▶ `kernel/interrupt.c:`

```
1 static struct softirq_action  
2     softirq_vec[NR_SOFTIRQS];  
3 /* NR_SOFTIRQ is max 32 */
```

▶ Softirq handler:

```
1 void handler_name(struct softirq_action *);
```


Softirqs

Implementing softirqs, softirq execution (2)

- ▶ The kernel runs a softirq by executing the handler:

```
1 /* let's assume my_softirq is a struct softirq_action * */  
2 my_softirq->action(my_softirq);
```

- ▶ **Softirq execution**

- ▶ Once registered, a softirq must be *raised* to indicate it needs to execute
 - ▶ Generally done by the top-half handler
 - ▶ **Raised softirqs are checked and executed:**
 - ▶ In return from interrupt
 - ▶ In a kernel thread, `ksoftirqd`
 - ▶ In any code that explicitly checks for and runs raised softirqs (`do_softirq()`)
- ▶ Check and run is done in `__do_softirq()`
 - ▶ Goes over the softirq array (`softirq_vec`)
 - ▶ Executes the handlers of raised (pending) softirqs

Softirqs

Using softirqs: softirq indexes

- ▶ **Softirqs are declared statically at compile time**

- ▶ enum in linux/interrupt.h

```
1 enum {
2   HI_SOFTIRQ=0,      /* 0 */
3   TIMER_SOFTIRQ,    /* 1 */
4   NET_TX_SOFTIRQ,   /* 2 */
5   NET_RX_SOFTIRQ,   /* 3 */
6   BLOCK_SOFTIRQ,    /* 4 */
7   IRQ_POLL_SOFTIRQ, /* 5 */
8   TASKLET_SOFTIRQ,  /* 6 */
9   SCHED_SOFTIRQ,    /* 7 */
10  HRTIMER_SOFTIRQ   /* 8 */
11  RCU_SOFTIRQ,      /* 9 */
12  NR_SOFTIRQ
13 };
```

- ▶ Create an entry in this array to add a softirq to Linux
- ▶ Entries ranked by priority (HI_SOFTIRQ is the highest)
 - ▶ **This is the order in which the array is iterated for softirq execution**

Softirqs

Using softirqs (2)

- ▶ (Very) simplified version of `__do_softirq()`:

```
1 int i;
2
3 /* Iterate in priority order */
4 for(i = 0; i < NR_SOFTIRQ; i++) {
5     struct softirq_action *handler = softirq_vec[i];
6     int pending = is_pending(handler);
7
8     if(pending) {
9         handler->(action(handler));
10    }
11 }
```

Softirqs

Using softirqs: handler registration

- ▶ Handler registration done through `open_softirq()`:

```
1 open_softirq(SOFTIRQ_INDEX, softirq_handler);
2 /* real example (net/core/dev.c): */
3 open_softirq(NET_TX_SOFTIRQ, net_tx_action);
```

- ▶ **Softirqs run in interrupt context**
 - ▶ Cannot block/sleep
- ▶ When a softirq handler is running, **softirqs are disabled on the local core**
 - ▶ However softirqs can run concurrently on different cores
 - ▶ Including softirqs with the same index → **shared data must be protected against concurrent accesses** (generally use per-processor data)
 - ▶ **Good scalability vs tasklets**
 - ▶ If the same softirq is raised while its handler is running it can execute on another core
 - ▶ One tasklet cannot run concurrently on multiple cores

Softirqs

Using softirqs: raising a softirq

▶ `raise_softirq()`:

```
1 /* kernel/time/timer.c: */  
2 raise_softirq(TIMER_SOFTIRQ);
```

- ▶ Disables interrupts on the local core before marking the softirq as pending
 - ▶ With `local_irq_save()`
- ▶ Restores them to the previous state afterward
 - ▶ With `local_irq_restore()`
- ▶ Generally called from the interrupt handler (top-half)
- ▶ Optimization if interrupts are already off:

```
1 /* net/core/dev.c: */  
2 raise_softirq_irqoff(NET_TX_SOFTIRQ);
```

Softirqs

Softirq example

- ▶ Cannot be registered from module (softirq symbols not exported)
 - ▶ Adding a softirq must be done from inside the kernel code
- ▶ `include/linux/interrupt.h`:

```
1 diff -rc linux-4.10.1/include/linux/interrupt.h linux-4.10.1.modified/include/linux/interrupt.  
  h  
2 *** linux-4.10.1/include/linux/interrupt.h 2017-02-26 10:09:33.000000000 +0000  
3 --- linux-4.10.1.modified/include/linux/interrupt.h 2017-02-28 15:57:46.088406158 +0000  
4 *****  
5 *** 456,461 ****  
6 --- 456,462 ----  
7 SCHED_SOFTIRQ,  
8 HRTIMER_SOFTIRQ, /* Unused, but kept as tools rely on the  
9     numbering. Sigh! */  
10 + PIERRE_SOFTIRQ,  
11 RCU_SOFTIRQ,     /* Preferable RCU should always be the last softirq */  
12  
13 NR_SOFTIRQS
```

Softirqs

Softirq example (2)

► kernel/main.c:

```

13 diff -rc linux-4.10.1/init/main.c linux
    -4.10.1.modified/init/main.c
14 *** linux-4.10.1/init/main.c 2017-02-26
    10:09:33.000000000 +0000
15 --- linux-4.10.1.modified/init/main.c
    2017-02-28 16:18:02.672173235 +0000
16 *****
17 *** 89,94 ****
18 --- 89,100 ----
19 #include <asm/sections.h>
20 #include <asm/cacheflush.h>
21 +
22 +
23 + void pierre_softirq_handler(struct
    softirq_action * action)
24 + {
25 +     printk("Pierre softirq running!\n");
26 +     ;
27 + }
28 +
29 static int kernel_init(void *);
30 extern void init_IRQ(void);

```

```

31 *****
32 *** 669,674 ****
33 --- 675,686 ----
34
35     ftrace_init();
36
37 +
38 +
39 +     open_softirq(PIERRE_SOFTIRQ,
    pierre_softirq_handler);
40 +     printk("Raising Pierre softirq\n");
41 +     raise_softirq(PIERRE_SOFTIRQ);
42 +
43     /* Do the rest non-__init'ed, we're
    now alive */
44     rest_init();
45 }

```

Outline

- 1 Bottom-halves: general information
- 2 Softirqs
- 3 Tasklets**
- 4 Workqueues
- 5 Using the right bottom-half and misc. information
- 6 Additional sources of information

Tasklets

Implementing tasklets

- ▶ **Tasklets** are implemented on top of softirqs
 - ▶ Same behavior, simpler interface, less locking rules
 - ▶ **The same tasklet cannot run concurrently on multiple cores**
 - ▶ Implemented in `HI_SOFTIRQ` and `TASKLET_SOFTIRQ` softirqs
- ▶ A tasklet is represented by a `tasklet_struct` (`include/linux/interrupt.h`):

```
1 struct tasklet_struct
2 {
3     struct tasklet_struct *next;
4     unsigned long state;
5     atomic_t count;
6     void (*func)(unsigned long);
7     unsigned long data;
8 }
```

- ▶ `func`: handler (`args`: `data`)
- ▶ `next`: linked list of tasklets
- ▶ `state`: enum (scheduled to run/running)
- ▶ `count`:
 - ▶ `!=0`: disabled, cannot run
 - ▶ `0`: enabled, can be marked pending

Tasklets

Implementing tasklets: scheduling tasklets

- ▶ *Scheduling* tasklets == *raising* softirqs
- ▶ Scheduled tasklets put in two per-processor linked lists:
 - ▶ `tasklet_hi_vec` (high priority)
 - ▶ `tasklet_vec`
- ▶ Scheduling a tasklet is done through `tasklet_schedule(t)` or `tasklet_hi_schedule(t)`:
 - 1 Check if the tasklet is already scheduled, if not, call `__tasklet_schedule(t)`
 - 2 Disable interrupts
 - 3 Add the tasklet to the corresponding linked list
 - 4 Raise `TASKLET_SOFTIRQ` or `HI_SOFTIRQ`
 - 5 Restore interrupts and return

Tasklets

Implementing tasklets: tasklets softirqs handlers

Tasklet softirqs handlers:

- ▶ `tasklet_action()` and `tasklet_hi_action()`
 - 1 Clear the list for the current CPU and iterate over its content, for each raised tasklet:
 - 1 Check `TASKLET_STATE_RUN` flag
 - 2 If the tasklet is not running, set that flag
 - 3 Check the `count` value (is it enabled?)
 - 4 Run the tasklet and clear the `TASKLET_STATE_RUN` flag

Tasklets

Using tasklets: creation

▶ Creation:

```
1 /* Statically: */
2 DECLARE_TASKLET(tasklet_name, handler_name, handler_arguments);
3 DECLARE_TASKLET_DISABLED(tasklet_name, handler_name, handler_arguments);
4
5 /* Dynamically */
6 tasklet_ptr = kcalloc(sizeof(struct tasklet_struct), GFP_KERNEL);
7 tasklet_init(tasklet_ptr, handler_name, handler_arguments);
```

- ▶ Difference between `DECLARE_TASKLET()` and `DECLARE_TASKLET_DISABLED()`:
 - ▶ count field of the initialized struct `task_struct`
 - ▶ 1 for *enabled*, 0 for *disabled* (will not run even if it is scheduled to)

Tasklets

Using tasklets: handler

▶ Handler

▶ Prototype:

```
1 void handler_name(unsigned long data);
```

- ▶ Like softirqs, **tasklets cannot sleep** (run in interrupt context)
 - ▶ No use of blocking semaphores, no `kmallo` with `GFP_KERNEL`, etc.
- ▶ **Tasklets run with interrupts enabled**
 - ▶ Shared data with an interrupt handler? *disable interrupts/get a (non-sleeping) lock*
 - ▶ Two different tasklets can run concurrently on different cores: *need locking if a tasklet shares data with another tasklet or a softirq*

Tasklets

Using tasklets: scheduling the tasklet to run, enabling/disabling a tasklet

- ▶ To mark the tasklet as pending: `tasklet_schedule()`

```
1 tasklet_schedule(&tasklet_name); /* tasklet_name being of type struct tasklet */
```

- ▶ Runs one in a near future on the same CPU it is schedule from
 - ▶ Independently of the number of calls to `tasklet_schedule()`

- ▶ **Disabling/enabling a tasklet:**

```
1 tasklet_disable(&tasklet_name);  
2 tasklet_enable(&tasklet_name);
```

- ▶ Set the `count` member of the `struct tasklet`
- ▶ `tasklet_disable()` blocks until any potential running handler finishes
 - ▶ `tasklet_disable_nosync()` in order not to wait, probably unsafe
 - ▶ `tasklet_enable()` must be called after declaring a tasklet through `DECLARE_TASKLET_DISABLED()`

Tasklets

ksoftirqd

- ▶ System can be flooded by softirqs (and tasklets)
 - ▶ Scenarios with high interrupts arrival frequency
 - ▶ Plus softirqs can reactivate themselves
- ▶ **Takes a lot of CPU time from userspace processes**
 - ▶ **But softirqs still need to be processed promptly**
- ▶ A solution is to defer *reactivated* softirqs until the next time softirq run
 - ▶ Generally at the next interrupt occurrence
 - ▶ Sub-optimal on an idle system
- ▶ **Solution:** defer reactivated softirqs into per-cpu, low priority kernel threads: `ksoftirqd`
 - ▶ No starvation of CPU time for user space processes in case of highly frequent interrupts
 - ▶ `ksoftirqd` priority is nice 19
 - ▶ On an idle system, `ksoftirqd` is scheduled quickly

Tasklets

Tasklet example

► Tasklet example in a module:

```
1 #include <linux/module.h>
2 #include <linux/kernel.h>
3 #include <linux/init.h>
4 #include <linux/interrupt.h>
5
6 void my_tasklet_handler(unsigned long data
7 );
8 static DECLARE_TASKLET(my_tasklet,
9     my_tasklet_handler, 0);
10 void my_tasklet_handler(unsigned long data
11 )
12 {
13     printk("tasklet executing.\n");
14 }
```

```
14 static int __init my_mod_init(void)
15 {
16     printk("Entering module.\n");
17     printk("Scheduling tasklet.\n");
18
19     tasklet_schedule(&my_tasklet);
20
21     return 0;
22 }
23
24 static void __exit my_mod_exit(void)
25 {
26     tasklet_disable(&my_tasklet);
27     printk(KERN_INFO "Exiting module.\n");
28 }
29
30 module_init(my_mod_init);
31 module_exit(my_mod_exit);
```


Outline

- 1 Bottom-halves: general information
- 2 Softirqs
- 3 Tasklets
- 4 Workqueues**
- 5 Using the right bottom-half and misc. information
- 6 Additional sources of information

Workqueues

Presentation

- ▶ **Workqueues run in *process context***
 - ▶ They are a schedulable entity, so **workqueues can sleep**
 - ▶ Needed when the bottom-half need to allocate a lot of memory, sleep on a lock (semaphore/mutex), perform blocking I/O
 - ▶ In these situations softirqs and tasklets cannot be used → **use workqueues**
- ▶ Work deferred into work queues is handled by **kernel threads**:
 - ▶ Default, per-cpu ones: `kworker/n` (before: `events/n`) where `n` is the CPU id
 - ▶ Syntax: `kworker/<cpu>[unbound]:<id><priority>`
 - ▶ Workqueues users can also create their own threads
 - ▶ Better performance & lighten the load on default threads

Workqueues

Work queues implementation: `worker_pool` and `work_struct`

▶ `kernel/workqueue.c`:

```
1 struct worker_pool {
2     spinlock_t    lock;
3     int           cpu;
4     int           node;
5     int           id;
6     unsigned int  flags;
7     /* list of work_struct: */
8     struct list_head worklist;
9     /* number of associated worker threads: */
10    int           nr_workers;
11
12
13    /* ... */
14 };
```

```
1 struct work_struct {
2     atomic_long_t  data;
3     struct list_head entry;
4     work_func_t    func;
5     #ifdef CONFIG_LOCKDEP
6     struct lockdep_map lockdep_map;
7     #endif
8 };
```

▶ `include/linux/workqueue.h`:

```
1 typedef void (*work_func_t)(struct work_struct *work);
```

Workqueues

Work queues implementation: worker thread function

- ▶ Worker threads execute the `worker_thread()` function (`kernel/workqueue.c`)
- ▶ Infinite loop doing the following:
 - 1 Check if there is some work to do in the current pool
 - 2 If so, execute all the `work_struct` objects pending in the pool `worklist` by calling **`process_scheduled_works()`**
 - ▶ Call the `work_struct` function pointer `func`, passing a pointer to the `work_struct` itself as a parameter
 - ▶ `work_struct` objects removed from the list after being processed
 - 3 Go to sleep
 - ▶ Worker threads are awoken next time some work is inserted in the workqueue

Workqueues

Using work queues

▶ Creating / destroying a **workqueue**:

```
1 struct workqueue_struct *create_workqueue(char *name);
2 void destroy_workqueue(struct workqueue_struct *);
```

▶ Creating a **work entity**:

```
1 /* statically: */
2 DECLARE_WORK(work, func);
3 DECLARE_DELAYED_WORK(work, func);
4 DECLARE_DEFERRABLE_WORK(work, func);
```

```
1 /* dynamically */
2 INIT_WORK(work_ptr, func);
3 INIT_DELAYED_WORK(work_ptr, func);
4 INIT_DEFERRABLE_WORK(work_ptr, func);
```

- ▶ **work** is the name of the initialized `work_struct`
- ▶ `work_ptr` is a pointer to an allocated `work_struct`
- ▶ `func` is the name of the handler function
- ▶ **DELAYED** is related to work item which execution can be delayed by a given time after they are scheduled to run
- ▶ **DEFERRED** indicates low priority work items

▶ Handler prototype:

```
1 void handler(struct work_struct *work);
```

Workqueues

Using work queues (2)

▶ Creating a work entity, example:

```
1 static void my_handler(struct work_struct *work)
2 {
3     /* ... */
4 }
5
6 /* Static creation: */
7 DECLARE_WORK(work_item, my_handler);
8
9 static int __init my_init_function(void)
10 {
11     /* dynamic creation: */
12     struct work_struct * work_item2;
13
14     work_item2 = kmalloc(sizeof(struct work_struct), GFP_KERNEL);
15     INIT_WORK(work_item2, my_handler);
16 }
```

Workqueues

Using work queues (3)

- ▶ Enqueueing a work on a specific, previously created work queue:

```
1 int queue_work (struct workqueue_struct *wq, struct work_struct *work);
2 int queue_work_on (int cpu, struct workqueue_struct *wq, struct work_struct *work);
3 int queue_delayed_work (struct workqueue_struct *wq, struct delayed_work *work,
4     unsigned long delay);
4 int queue_delayed_work_on (int cpu, struct workqueue_struct *wq, struct delayed_work *
    work), unsigned long delay);
```

- ▶ Enqueue on the default kernel threads (kworkers):

```
1 int schedule_work (struct work_struct *):
2 int schedule_work_on (int cpu, struct work_struct *):
3 int scheduled_delayed_work (struct delayed_work *, unsigned long delay);
4 int scheduled_delayed_work_on (int cpu, struct delayed_work *, unsigned long delay);
```

- ▶ **Work executes next time a worker thread is awoken**
- ▶ Can schedule on a specific core ("_on" functions)
- ▶ Can delay execution by a given number of timer ticks (delay)



Workqueues

Using work queues: work function parameters

- ▶ **How to pass parameters to the work item handler?**
 - ▶ Not done directly like in tasklets, *but the `work_struct` is passed as a parameter to its own handler execution*
 - ▶ Solution: **put the `work_struct` in a data structure**
 - ▶ Parameter as another member of the data structure
 - ▶ Use `container_of()` from inside the handler

```
1 struct work_item {
2     struct work_struct ws;
3     int parameter;
4 };
5
6 static void handler( struct work_struct *work)
7 {
8     struct work_item *wi = (struct work_item *)container_of(work, struct work_item, ws);
9     int parameter = wi->parameter;
10 }
```

- ▶ In that case (`work_struct` as the first field of the containing data structure), one can also do:

```
1 struct work_item *wi = (struct work_item *)work;
```


Workqueues

Using work queues: utility functions

- ▶ To ensure work finishes its execution

```
1 /* flush a specific work_struct */
2 int flush_work(struct work_struct *work);
3 /* flushes a specific workqueue: */
4 void flush_workqueue(struct workqueue_struct *);
5 /* flush the default workqueue (kworkers): */
6 void flush_scheduled_work(void);
```

- ▶ Canceling scheduled work:

```
1 int cancel_work_sync(struct work_struct *work);
2 int cancel_delayed_work_sync(struct delayed_work *dwork);
```

- ▶ Checking if work is pending:

```
1 work_pending(work); /* work is struct work_struct */
2 delayed_work_pending(work); /* same thing */
```

- ▶ Destroying a workqueue:

```
1 void destroy_workqueue(struct workqueue_struct *);
```

Workqueues

Workqueue module example

```
1 #include <linux/module.h>
2 #include <linux/kernel.h>
3 #include <linux/init.h>
4 #include <linux/slab.h>
5 #include <linux/workqueue.h>
6
7 struct work_item {
8     struct work_struct ws;
9     int parameter;
10 };
11
12 struct work_item *wi, *wi2;
13 struct workqueue_struct *my_wq;
14
15 static void handler( struct work_struct *
16     work)
17 {
18     int parameter = ((struct work_item *)
19         container_of(work, struct work_item,
20         ws))->parameter;
21     printk("doing some work ... \n");
22     printk("parameter is: %d \n", parameter);
23 }
```

```
21 static int __init my_mod_init(void)
22 {
23     printk("Entering module. \n");
24
25     my_wq = create_workqueue("pierre_wq");
26     wi = kmalloc(sizeof(struct work_item),
27         GFP_KERNEL);
28     wi2 = kmalloc(sizeof(struct work_item),
29         GFP_KERNEL);
30
31     INIT_WORK(&wi->ws, handler);
32     wi->parameter = 42;
33     INIT_WORK(&wi2->ws, handler);
34     wi2->parameter = -42;
35
36     schedule_work(&wi->ws);
37     queue_work(my_wq, &wi2->ws);
38
39     return 0;
40 }
```

Workqueues

Workqueue module example (2)

```
39 static void __exit my_mod_exit(void)
40 {
41     flush_scheduled_work();
42     flush_workqueue(my_wq);
43     kfree(wi);
44     kfree(wi2);
45     destroy_workqueue(my_wq);
46     printk(KERN_INFO "Exiting module.\n");
47 }
48
49 module_init(my_mod_init);
50 module_exit(my_mod_exit);
51 MODULE_LICENSE("GPL");
```

Outline

- 1 Bottom-halves: general information
- 2 Softirqs
- 3 Tasklets
- 4 Workqueues
- 5 Using the right bottom-half and misc. information**
- 6 Additional sources of information

Using the right bottom-half and misc. information

Using the right bottom-half

- ▶ Work has a lot of potential for parallelism and/or is very performance critical
 - ▶ **Softirqs**
 - ▶ Can run concurrently on multiple cores (need to take care of concurrency issues)
 - ▶ High performance: generally run very quickly after marked pending
 - ▶ No need for parallelism but performance still important
 - ▶ **Tasklets**
 - ▶ Two tasklets cannot run concurrently on several cores
 - ▶ Performance close to softirqs (depends on the tasklet type considered)
 - ▶ Need to run in process context
 - ▶ **Workqueues**
 - ▶ Can sleep
 - ▶ Less performance, need to wait to be scheduled to perform work

Using the right bottom-half and misc. information

Locking between bottom-halves

▶ **Softirqs**

- ▶ Need intra-softirq (same softirq) locking (thread-safe)
- ▶ Need inter-softirq (different softirqs) locking in case of shared data between them

▶ **Tasklets**

- ▶ No need to protect against concurrent accesses from the same tasklet
- ▶ Two different tasklets sharing data need proper locking (inter-tasklets locking)

▶ **Workqueues**

- ▶ Intra- and inter-workqueue locking needed

▶ **All of these generally run with interrupts enabled**

- ▶ In case of shared data with an interrupt handler (top-half), disabling interrupts + locking needed

Using the right bottom-half and misc. information

Disabling bottom-half processing

- ▶ To disable / enable softirqs and tasklets on the local core:

```
1 void local_bh_disable();  
2 void local_bh_enable();
```

- ▶ Can nest
 - ▶ Need to call `local_bh_enable()` as much times as `local_bh_disable()` was called to re-enable interrupts
- ▶ These calls do not disable workqueues processing
 - ▶ Not a problem as these run in process context and do not happen asynchronously like softirqs/tasklets

Outline

- 1 Bottom-halves: general information
- 2 Softirqs
- 3 Tasklets
- 4 Workqueues
- 5 Using the right bottom-half and misc. information
- 6 Additional sources of information**

Additional sources of information

Disabling bottom-half processing

- ▶ Numerous details about the internals of Linux interrupt management (top/bottom half): <https://0xax.gitbooks.io/linux-insides/content/interrupts/> (Linux 3.18)
- ▶ Workqueues documentation in Linux sources: `Documentation/workqueues.txt`