



# **Open Source in Industry: Linux tracing and debugging/ Trouble shooting of real-time Linux**


**Technical Heidelberg OSADL Talks, September 30, 2020, Online Session 2b**

**Latency fighting**

# Latency fighting case #1







# Culprit/victim table


Characteristics of the 20 highest latencies: System rack9slot3.osadl.org (updated Tue Mar 19, 2019 00:43:28)									
Delayed (victim)					Switcher (culprit)			Timestamp	CPU
PID	Prio	Total latency (µs)	T*(.W**) latency (µs)	Cmd	PID	Prio	Cmd		
9952	99	16981	0,16980	cyclctest	25744	-21	modprobe	22:34:26	4
9948	99	16848	0,16846	cyclctest	25744	-21	modprobe	22:34:26	0
9951	99	196	194,1	cyclctest	28093	-21	cat	22:27:59	3
13655	2	156	0,1	sleep1	3657	-21	runrttasks	22:51:09	1
28522	2	154	0,0	sleep6	0	-21	swapper/6	21:59:59	6
22371	2	137	0,0	sleep1	0	-21	swapper/1	22:59:58	1
9742	2	94	49,15	sleep2	0	-21	swapper/2	19:07:55	2
9640	2	90	49,34	sleep3	0	-21	swapper/3	19:06:28	3
9856	2	81	48,25	sleep4	0	-21	swapper/4	19:09:34	4
9800	2	73	53,13	sleep6	0	-21	swapper/6	19:08:48	6
9601	2	73	51,15	sleep0	0	-21	swapper/0	19:05:53	0
9866	2	72	51,14	sleep5	0	-21	swapper/5	19:09:43	5
9653	2	72	51,14	sleep7	0	-21	swapper/7	19:06:40	7
9629	2	72	50,15	sleep1	0	-21	swapper/1	19:06:18	1
9952	99	53	0,52	cyclctest	0	-21	swapper/4	21:14:56	4
9115	2	36	0,0	sleep1	0	-21	swapper/1	22:47:40	1
5127	2	36	0,0	sleep7	0	-21	swapper/7	00:10:03	7
11408	2	36	0,0	sleep5	0	-21	swapper/5	23:46:36	5
30046	2	34	0,1	sleep0	30048	-21	sh	22:37:29	0
3665	2	33	0,1	sleep1	0	-21	swapper/1	22:05:23	1

\*Timer \*\*Wakeup (latency=timer+wakeup+contextswitch)

# Culprit/victim table

Characteristics of the 20 highest latencies: System rack9slot3.osadl.org (updated Tue Mar 19, 2019 00:43:28)									
Delayed (victim)					Switcher (culprit)			Timestamp	CPU
PID	Prio	Total latency (µs)	T*(.W**) latency (µs)	Cmd	PID	Prio	Cmd		
9952	99	16981	0,16980	cyclicttest	25744	-21	modprobe	22:34:26	4
9948	99	16848	0,16846	cyclicttest	25744	-21	modprobe	22:34:26	0
9951	99	196	194,1	cyclicttest	28093	-21	cat	22:27:59	3
13655	2	156	0,1	sleep1	3657	-21	runrttasks	22:51:09	1
28522	2	154	0,0	sleep6	0	-21	swapper/6	21:59:59	6
22371	2	137	0,0	sleep1	0	-21	swapper/1	22:59:58	1
9742	2	94	49,15	sleep2	0	-21	swapper/2	19:07:55	2
9640	2	90	49,34	sleep3	0	-21	swapper/3	19:06:28	3
9856	2	81	48,25	sleep4	0	-21	swapper/4	19:09:34	4
9800	2	73	53,13	sleep6	0	-21	swapper/6	19:08:48	6
9601	2	73	51,15	sleep0	0	-21	swapper/0	19:05:53	0
9866	2	72	51,14	sleep5	0	-21	swapper/5	19:09:43	5
9653	2	72	51,14	sleep7	0	-21	swapper/7	19:06:40	7
9629	2	72	50,15	sleep1	0	-21	swapper/1	19:06:18	1
9952	99	53	0,52	cyclicttest	0	-21	swapper/4	21:14:56	4
9115	2	36	0,0	sleep1	0	-21	swapper/1	22:47:40	1
5127	2	36	0,0	sleep7	0	-21	swapper/7	00:10:03	7
11408	2	36	0,0	sleep5	0	-21	swapper/5	23:46:36	5
30046	2	34	0,1	sleep0	30048	-21	sh	22:37:29	0
3665	2	33	0,1	sleep1	0	-21	swapper/1	22:05:23	1

\*Timer \*\*Wakeup (latency=timer+wakeup+contextswitch)





# Let's have a look at dmesg

```
# dmesg | tail -11
```

```
[7216.176506] CPU3: Core temperature above threshold, cpu clock throttled (total events = 1)
[7216.176507] CPU7: Core temperature above threshold, cpu clock throttled (total events = 1)
[7216.176509] CPU5: Package temperature above threshold, cpu clock throttled (total events = 1)
[7216.176510] CPU1: Package temperature above threshold, cpu clock throttled (total events = 1)
[7216.176512] CPU6: Package temperature above threshold, cpu clock throttled (total events = 1)
[7216.176514] CPU4: Package temperature above threshold, cpu clock throttled (total events = 1)
[7216.176515] CPU2: Package temperature above threshold, cpu clock throttled (total events = 1)
[7216.176516] CPU0: Package temperature above threshold, cpu clock throttled (total events = 1)
[7216.176517] CPU7: Package temperature above threshold, cpu clock throttled (total events = 1)
[7216.176518] CPU3: Package temperature above threshold, cpu clock throttled (total events = 1)
[7216.176532] mce: [Hardware Error]: Machine check events logged
```

# At the time of the latency, a new module appeared


```
# lsmod | head
```

Module	Size	Used by
<b>msr</b>	<b>16384</b>	<b>16</b>
eeeprom	16384	0
nfsv4	409600	2
ppdev	16384	0
bnep	20480	2
lp	20480	0
rfcomm	36864	0
bluetooth	299008	10 bnep, rfcomm
rfkill	20480	2 bluetooth

# Countermeasures

- Quick fix
  - Load the *msr* module immediately after the system was booted
- Better
  - File a bug
- Even better
  - Fix it yourself
    - ✓ Analyze the init part of *msr*.
    - ✓ Search for busy loops or similar while preemption and/or interrupts are disabled.
    - ✓ Alternatively, use *cyclictest*'s breaktrace feature and function tracing to pinpoint the culprit.
    - ✓ Shorten the busy loop, allow intermittent interrupt execution and preemption.

# Latency fighting case #2



# Culprit/victim table

Characteristics of the 20 highest latencies: System rack6slot2.osadl.org (updated Mon Mar 18, 2019 00:43:27)									
Delayed (victim)					Switcher (culprit)			Timestamp	CPU
PID	Prio	Total latency (µs)	T*(.W**) latency (µs)	Cmd	PID	Prio	Cmd		
5528	2	114	0,1	sleep1	21	-21	rcuc/1	19:31:47	1
18185	2	61	0,0	sleep1	0	-21	swapper/1	21:32:02	1
28116	2	43	8,8	sleep1	0	-21	swapper/1	19:05:18	1
28316	99	39	32,6	cyclicttest	2408	50	irq/16-enp2s0f0	20:36:47	1
28316	99	38	33,4	cyclicttest	2408	50	irq/16-enp2s0f0	20:41:45	1
28316	99	38	32,5	cyclicttest	2408	50	irq/16-enp2s0f0	00:11:50	1
28316	99	38	31,6	cyclicttest	2408	50	irq/16-enp2s0f0	22:54:36	1
28316	99	37	34,2	cyclicttest	2408	50	irq/16-enp2s0f0	20:13:03	1
28316	99	37	32,4	cyclicttest	2408	50	irq/16-enp2s0f0	23:43:28	1
28316	99	37	32,4	cyclicttest	2408	50	irq/16-enp2s0f0	21:06:47	1
28316	99	37	32,4	cyclicttest	2408	50	irq/16-enp2s0f0	20:56:53	1
28316	99	37	32,4	cyclicttest	2408	50	irq/16-enp2s0f0	19:26:50	1
28316	99	37	32,4	cyclicttest	2408	50	irq/16-enp2s0f0	19:12:46	1
28316	99	37	31,5	cyclicttest	2408	50	irq/16-enp2s0f0	21:52:05	1
28316	99	36	33,2	cyclicttest	2408	50	irq/16-enp2s0f0	22:20:44	1
28316	99	36	32,3	cyclicttest	2408	50	irq/16-enp2s0f0	23:41:50	1
28316	99	36	32,3	cyclicttest	2408	50	irq/16-enp2s0f0	23:26:47	1
28316	99	36	32,3	cyclicttest	2408	50	irq/16-enp2s0f0	21:45:53	1
28316	99	36	32,3	cyclicttest	2408	50	irq/16-enp2s0f0	21:35:53	1
28316	99	36	32,3	cyclicttest	2408	50	irq/16-enp2s0f0	20:03:54	1

\*Timer \*\*Wakeup (latency=timer+wakeup+contextswitch)

# Culprit/victim table

Characteristics of the 20 highest latencies: System rack6slot2.osadl.org (updated Mon Mar 18, 2019 00:43:27)									
Delayed (victim)				Switcher (culprit)				Timestamp	CPU
PID	Prio	Total latency (µs)	T*(.W**) latency (µs)	Cmd	PID	Prio	Cmd		
5528	2	114	0,1	sleep1	21	-21	rcuc/1	19:31:47	1
18185	2	61	0,0	sleep1	0	-21	swapper/1	21:32:02	1
28116	2	43	8,8	sleep1	0	-21	swapper/1	19:05:18	1
28316	99	39	32,6	cyclicttest	2408	50	irq/16-enp2s0f0	20:36:47	1
28316	99	38	33,4	cyclicttest	2408	50	irq/16-enp2s0f0	20:41:45	1
28316	99	38	32,5	cyclicttest	2408	50	irq/16-enp2s0f0	00:11:50	1
28316	99	38	31,6	cyclicttest	2408	50	irq/16-enp2s0f0	22:54:36	1
28316	99	37	34,2	cyclicttest	2408	50	irq/16-enp2s0f0	20:13:03	1
28316	99	37	32,4	cyclicttest	2408	50	irq/16-enp2s0f0	23:43:28	1
28316	99	37	32,4	cyclicttest	2408	50	irq/16-enp2s0f0	21:06:47	1
28316	99	37	32,4	cyclicttest	2408	50	irq/16-enp2s0f0	20:56:53	1
28316	99	37	32,4	cyclicttest	2408	50	irq/16-enp2s0f0	19:26:50	1
28316	99	37	32,4	cyclicttest	2408	50	irq/16-enp2s0f0	19:12:46	1
28316	99	37	31,5	cyclicttest	2408	50	irq/16-enp2s0f0	21:52:05	1
28316	99	36	33,2	cyclicttest	2408	50	irq/16-enp2s0f0	22:20:44	1
28316	99	36	32,3	cyclicttest	2408	50	irq/16-enp2s0f0	23:41:50	1
28316	99	36	32,3	cyclicttest	2408	50	irq/16-enp2s0f0	23:26:47	1
28316	99	36	32,3	cyclicttest	2408	50	irq/16-enp2s0f0	21:45:53	1
28316	99	36	32,3	cyclicttest	2408	50	irq/16-enp2s0f0	21:35:53	1
28316	99	36	32,3	cyclicttest	2408	50	irq/16-enp2s0f0	20:03:54	1


\*Timer \*\*Wakeup (latency=timer+wakeup+contextswitch)

# Countermeasures


- Quick fix
  - Uninstall Broadcom Limited NetLink BCM57785 Gigabit Ethernet PCIe network device and unload driver *tg3*.
  - Install another, for example USB, Ethernet controller and reconfigure the system to no longer use IRQ 16 handler of network device *enp2s0f0*.
- Better
  - File a bug
- Even better
  - Fix it yourself
    - ✓ Analyze the source code of *drivers/net/ethernet/broadcom/tg3.c*.
    - ✓ Search for busy loops or similar while preemption and/or interrupts are disabled.
    - ✓ Alternatively, use *cyclictest*'s breaktrace feature and function tracing to pinpoint the culprit.
    - ✓ Shorten the busy loop, allow intermittent interrupt execution and preemption.




# Before



After



# Latency fighting case #3



# Culprit/victim table

Characteristics of the 50 highest latencies: System rack9slot3.osadl.org (updated Thu Mar 21, 2019 12:43:28)									
Delayed (victim)					Switcher (culprit)			Timestamp	CPU
PID	Prio	Total latency (µs)	T*(,W**) latency (µs)	Cmd	PID	Prio	Cmd		
3053	99	67	65,1	cyclctest	914	-21	meminfo	11:35:15	1
3052	99	66	63,2	cyclctest	7422	-21	meminfo	10:15:13	0
3057	99	64	62,1	cyclctest	11143	-21	meminfo	08:30:12	5
3055	99	64	61,2	cyclctest	29980	-21	meminfo	10:35:14	3
3053	99	64	62,1	cyclctest	7460	-21	meminfo	07:15:13	1
3057	99	62	59,2	cyclctest	26677	-21	meminfo	07:55:14	5
3052	99	62	60,1	cyclctest	19536	-21	meminfo	07:40:17	0
3055	99	57	54,2	cyclctest	13529	-21	meminfo	08:35:13	3
3056	99	55	53,1	cyclctest	23084	-21	meminfo	08:55:16	4
3052	99	53	51,1	cyclctest	32180	-21	meminfo	09:40:13	0
3053	99	52	50,1	cyclctest	3954	-21	meminfo	08:15:15	1
3054	99	51	49,1	cyclctest	24304	-21	meminfo	07:50:16	2
3054	99	51	48,2	cyclctest	25223	-21	meminfo	12:25:15	2
3052	99	51	49,1	cyclctest	15098	-21	meminfo	09:25:12	0
3056	99	49	46,2	cyclctest	25476	-21	meminfo	09:00:14	4
3052	99	49	47,1	cyclctest	5060	-21	meminfo	11:10:15	0
3054	99	44	42,1	cyclctest	26094	-21	meminfo	11:00:12	2
3054	99	44	41,2	cyclctest	6368	-21	meminfo	08:20:12	2
3052	99	44	42,1	cyclctest	27854	-21	meminfo	09:05:12	0
3052	99	44	42,1	cyclctest	20412	-21	meminfo	10:55:13	0

\*Timer \*\*Wakeup (latency=timer+wakeup+contextswitch)































# Culprit/victim table

Characteristics of the 50 highest latencies: System rack9slot3.osadl.org (updated Thu Mar 21, 2019 12:43:28)									
Delayed (victim)					Switcher (culprit)			Timestamp	CPU
PID	Prio	Total latency (µs)	T*(.W**) latency (µs)	Cmd	PID	Prio	Cmd		
3053	99	67	65,1	cyclicttest	914	-21	meminfo	11:35:15	1
3052	99	66	63,2	cyclicttest	7422	-21	meminfo	10:15:13	0
3057	99	64	62,1	cyclicttest	11143	-21	meminfo	08:30:12	5
3055	99	64	61,2	cyclicttest	29980	-21	meminfo	10:35:14	3
3053	99	64	62,1	cyclicttest	7460	-21	meminfo	07:15:13	1
3057	99	62	59,2	cyclicttest	26677	-21	meminfo	07:55:14	5
3052	99	62	60,1	cyclicttest	19536	-21	meminfo	07:40:17	0
3055	99	57	54,2	cyclicttest	13529	-21	meminfo	08:35:13	3
3056	99	55	53,1	cyclicttest	23084	-21	meminfo	08:55:16	4
3052	99	53	51,1	cyclicttest	32180	-21	meminfo	09:40:13	0
3053	99	52	50,1	cyclicttest	3954	-21	meminfo	08:15:15	1
3054	99	51	49,1	cyclicttest	24304	-21	meminfo	07:50:16	2
3054	99	51	48,2	cyclicttest	25223	-21	meminfo	12:25:15	2
3052	99	51	49,1	cyclicttest	15098	-21	meminfo	09:25:12	0
3056	99	49	46,2	cyclicttest	25476	-21	meminfo	09:00:14	4
3052	99	49	47,1	cyclicttest	5060	-21	meminfo	11:10:15	0
3054	99	44	42,1	cyclicttest	26094	-21	meminfo	11:00:12	2
3054	99	44	41,2	cyclicttest	6368	-21	meminfo	08:20:12	2
3052	99	44	42,1	cyclicttest	27854	-21	meminfo	09:05:12	0
3052	99	44	42,1	cyclicttest	20412	-21	meminfo	10:55:13	0

\*Timer \*\*Wakeup (latency=timer+wakeup+contextswitch)


# Analysis

- The script *meminfo* is a *Munin* plugin that provides a detailed snapshot of memory allocation and usage.
- It may report an in-depth analysis of the *slub* allocator, if the related debug configuration is enabled. If so, *slub* data can be inspected through the */proc/slabinfo* interface.



Trouble shooting of real-time Linux – Latency fighting  
Technical Heidelberg OSADL Talks, September 30, 2020, Online Session 2b  
Open Source Automation Development Lab (OSADL), Heidelberg





# Countermeasure #1 (quick fix)

```
--- .config.old 2019-03-17 12:49:05.216088137 +0100
+++ .config 2019-03-17 13:08:18.013949211 +0100
@@ -5649,8 +5649,8 @@
# CONFIG_PAGE_POISONING is not set
# CONFIG_DEBUG_PAGE_REF is not set
# CONFIG_DEBUG_OBJECTS is not set
-CONFIG_SLUB_DEBUG_ON=y
-CONFIG_SLUB_STATS=y
+# CONFIG_SLUB_DEBUG_ON is not set
+# CONFIG_SLUB_STATS is not set
CONFIG_HAVE_DEBUG_KMEMLEAK=y
# CONFIG_DEBUG_KMEMLEAK is not set
# CONFIG_DEBUG_STACK_USAGE is not set
```

# Countermeasure #2 (provide patch)

Subject: Reading /proc/slabinfo may cause large latencies

From: Carsten Emde <C.Emde@osadl.org>

Date: Sat, 29 Jun 2013 15:28:27 +0100

Reading /proc/slabinfo may cause large latencies of up to several milliseconds. This is due to a mutex lock that spans over the entire readout period and, thus, may prevent a higher-priority process from allocating memory during this amount of time.

This patch shortens the mutex lock to the data accumulation of a single item. It still guarantees coherence of a single output line but not across items which probably is a minor disadvantage compared to the otherwise fatal effect on the system's real-time capabilities.

Signed-off-by: Carsten Emde <C.Emde@osadl.org>

```
--  
mm/slab_common.c | 4 +--  
1 file changed, 2 insertions(+), 2 deletions(-)
```

Index: linux-3.12.0-rt1/mm/slab\_common.c


```
=====
```

--- linux-3.12.0-rt1.orig/mm/slab\_common.c  
+++ linux-3.12.0-rt1/mm/slab\_common.c


```
@@ -543,7 +543,6 @@ static void *s_start(struct seq_file *m,  
{  
    loff_t n = *pos;  
  
-    mutex_lock(&slab_mutex);  
    if (!n)  
        print_slabinfo_header(m);  
  
@@ -557,7 +556,6 @@ void *slab_next(struct seq_file *m, void  
  
void slab_stop(struct seq_file *m, void *p)  
{  
-    mutex_unlock(&slab_mutex);  
}  
  
static void  
@@ -570,6 +568,7 @@ memcg_accumulate_slabinfo(struct kmem_ca  
    if (!is_root_cache(s))  
        return;  
  
+    mutex_lock(&slab_mutex);  
    for_each_memcg_cache_index(i) {  
        c = cache_from_memcg(s, i);  
        if (!c)
```

```
@@ -584,6 +583,7 @@ memcg_accumulate_slabinfo(struct kmem_ca  
        info->active_objs += sinfo.active_objs;  
        info->num_objs += sinfo.num_objs;  
    }  
+    mutex_unlock(&slab_mutex);  
}  
  
int cache_show(struct kmem_cache *s, struct seq_file *m)
```

# Before



After



# Culprit/victim table

Characteristics of the 20 highest latencies: System rack9slot3.osadl.org (updated Sun Mar 17, 2019 00:43:25)									
Delayed (victim)					Switcher (culprit)			Timestamp	CPU
PID	Prio	Total latency (µs)	T*(,W**) latency (µs)	Cmd	PID	Prio	Cmd		
26221	2	95	50,15	sleep1	0	-21	swapper/1	18:52:07	1
26102	2	92	49,35	sleep3	0	-21	swapper/3	18:50:24	3
26059	2	91	47,14	sleep7	0	-21	swapper/7	18:49:48	7
26206	2	90	48,35	sleep4	0	-21	swapper/4	18:51:54	4
26198	2	90	48,35	sleep5	0	-21	swapper/5	18:51:47	5
26032	2	90	47,14	sleep0	0	-21	swapper/0	18:49:25	0
25252	2	73	0,0	sleep0	0	-21	swapper/0	19:52:43	0
1094	2	72	0,0	sleep2	30	1	ktimersoftd/2	23:26:17	2
31196	2	71	0,1	sleep5	4203	-21	runrttasks	20:07:38	5
25955	2	71	50,14	sleep6	0	-21	swapper/6	18:48:19	6
25949	2	70	48,15	sleep2	0	-21	swapper/2	18:48:15	2
13130	2	64	0,0	sleep7	0	-21	swapper/7	23:37:03	7
27566	2	61	0,0	sleep6	0	-21	swapper/6	21:55:28	6
26321	99	27	0,11	cyclicttest	0	-21	swapper/3	21:42:11	3
26318	99	27	0,4	cyclicttest	0	-21	swapper/0	00:13:03	0
26318	99	27	0,21	cyclicttest	0	-21	swapper/0	21:00:39	0
26318	99	26	0,19	cyclicttest	0	-21	swapper/0	22:24:27	0
26321	99	25	0,11	cyclicttest	0	-21	swapper/3	22:10:52	3
26318	99	25	9,2	cyclicttest	0	-21	swapper/0	23:18:58	0
26318	99	25	8,2	cyclicttest	0	-21	swapper/0	23:54:25	0

\*Timer \*\*Wakeup (latency=timer+wakeup+contextswitch)


# Culprit/victim table

Characteristics of the 20 highest latencies: System rack9slot3.osadl.org (updated Sun Mar 17, 2019 00:43:25)									
Delayed (victim)					Switcher (culprit)			Timestamp	CPU
PID	Prio	Total latency (µs)	T*(.W**) latency (µs)	Cmd	PID	Prio	Cmd		
26221	2	95	50,15	sleep1	0	-21	swapper/1	18:52:07	1
26102	2	92	49,35	sleep3	0	-21	swapper/3	18:50:24	3
26059	2	91	47,14	sleep7	0	-21	swapper/7	18:49:48	7
26206	2	90	48,35	sleep4	0	-21	swapper/4	18:51:54	4
26198	2	90	48,35	sleep5	0	-21	swapper/5	18:51:47	5
26032	2	90	47,14	sleep0	0	-21	swapper/0	18:49:25	0
25252	2	73	0,0	sleep0	0	-21	swapper/0	19:52:43	0
1094	2	72	0,0	sleep2	30	1	ktimersoftd/2	23:26:17	2
31196	2	71	0,1	sleep5	4203	-21	runrttasks	20:07:38	5
25955	2	71	50,14	sleep6	0	-21	swapper/6	18:48:19	6
25949	2	70	48,15	sleep2	0	-21	swapper/2	18:48:15	2
13130	2	64	0,0	sleep7	0	-21	swapper/7	23:37:03	7
27566	2	61	0,0	sleep6	0	-21	swapper/6	21:55:28	6
26321	99	27	0,11	cyclicttest	0	-21	swapper/3	21:42:11	3
26318	99	27	0,4	cyclicttest	0	-21	swapper/0	00:13:03	0
26318	99	27	0,21	cyclicttest	0	-21	swapper/0	21:00:39	0
26318	99	26	0,19	cyclicttest	0	-21	swapper/0	22:24:27	0
26321	99	25	0,11	cyclicttest	0	-21	swapper/3	22:10:52	3
26318	99	25	9,2	cyclicttest	0	-21	swapper/0	23:18:58	0
26318	99	25	8,2	cyclicttest	0	-21	swapper/0	23:54:25	0

\*Timer \*\*Wakeup (latency=timer+wakeup+contextswitch)

# Latency fighting case #4





# Culprit/victim table

Characteristics of the 20 highest latencies: System rack9slot5s.osadl.org (updated Sat Mar 16, 2019 12:46:13)									
Delayed (victim)					Switcher (culprit)			Timestamp	CPU
PID	Prio	Total latency (µs)	T*(.W**) latency (µs)	Cmd	PID	Prio	Cmd		
7787	99	346	343,0	cyclicttest	2337	-21	kerneloops	08:54:20	2
7791	99	309	307,0	cyclicttest	2337	-21	kerneloops	09:54:09	6
7791	99	299	299,0	cyclicttest	2337	-21	kerneloops	10:42:46	6
7787	99	299	297,0	cyclicttest	2337	-21	kerneloops	07:49:38	2
7791	99	296	296,0	cyclicttest	2337	-21	kerneloops	10:33:43	6
7789	99	296	294,0	cyclicttest	2337	-21	kerneloops	10:04:18	4
7790	99	294	294,0	cyclicttest	2337	-21	kerneloops	12:15:52	5
7790	99	294	292,0	cyclicttest	2337	-21	kerneloops	10:23:32	5
7787	99	294	293,0	cyclicttest	2337	-21	kerneloops	07:37:29	2
7790	99	293	293,0	cyclicttest	2337	-21	kerneloops	07:48:38	5
7790	99	292	292,0	cyclicttest	2337	-21	kerneloops	09:41:58	5
7790	99	292	290,0	cyclicttest	2337	-21	kerneloops	09:53:07	5
7792	99	291	291,0	cyclicttest	2337	-21	kerneloops	11:33:20	7
7792	99	291	289,0	cyclicttest	2337	-21	kerneloops	10:57:55	7
7790	99	291	289,0	cyclicttest	2337	-21	kerneloops	09:02:28	5
7789	99	291	291,0	cyclicttest	2337	-21	kerneloops	11:07:02	4
7788	99	291	291,0	cyclicttest	2337	-21	kerneloops	09:10:35	3
7790	99	290	290,0	cyclicttest	2337	-21	kerneloops	08:58:24	5
7792	99	289	289,0	cyclicttest	2337	-21	kerneloops	10:45:48	7
7792	99	289	289,0	cyclicttest	2337	-21	kerneloops	09:32:52	7

\*Timer \*\*Wakeup (latency=timer+wakeup+contextswitch)

# Culprit/victim table


Characteristics of the 20 highest latencies: System rack9slot5s.osadl.org (updated Sat Mar 16, 2019 12:46:13)									
Delayed (victim)					Switcher (culprit)			Timestamp	CPU
PID	Prio	Total latency (µs)	T*(.W**) latency (µs)	Cmd	PID	Prio	Cmd		
7787	99	346	343,0	cyclicttest	2337	-21	kerneloops	08:54:20	2
7791	99	309	307,0	cyclicttest	2337	-21	kerneloops	09:54:09	6
7791	99	299	299,0	cyclicttest	2337	-21	kerneloops	10:42:46	6
7787	99	299	297,0	cyclicttest	2337	-21	kerneloops	07:49:38	2
7791	99	296	296,0	cyclicttest	2337	-21	kerneloops	10:33:43	6
7789	99	296	294,0	cyclicttest	2337	-21	kerneloops	10:04:18	4
7790	99	294	294,0	cyclicttest	2337	-21	kerneloops	12:15:52	5
7790	99	294	292,0	cyclicttest	2337	-21	kerneloops	10:23:32	5
7787	99	294	293,0	cyclicttest	2337	-21	kerneloops	07:37:29	2
7790	99	293	293,0	cyclicttest	2337	-21	kerneloops	07:48:38	5
7790	99	292	292,0	cyclicttest	2337	-21	kerneloops	09:41:58	5
7790	99	292	290,0	cyclicttest	2337	-21	kerneloops	09:53:07	5
7792	99	291	291,0	cyclicttest	2337	-21	kerneloops	11:33:20	7
7792	99	291	289,0	cyclicttest	2337	-21	kerneloops	10:57:55	7
7790	99	291	289,0	cyclicttest	2337	-21	kerneloops	09:02:28	5
7789	99	291	291,0	cyclicttest	2337	-21	kerneloops	11:07:02	4
7788	99	291	291,0	cyclicttest	2337	-21	kerneloops	09:10:35	3
7790	99	290	290,0	cyclicttest	2337	-21	kerneloops	08:58:24	5
7792	99	289	289,0	cyclicttest	2337	-21	kerneloops	10:45:48	7
7792	99	289	289,0	cyclicttest	2337	-21	kerneloops	09:32:52	7

\*Timer \*\*Wakeup (latency=timer+wakeup+contextswitch)


# Countermeasures

- Quick fix
  - `# systemctl stop kerneloops`
  - `# systemctl disable kerneloops`
- Better
  - File a bug
- Even better
  - Fix it yourself and provide a patch

# Before



After



# Culprit/victim table

Characteristics of the 20 highest latencies: System rack9slot5s.osadl.org (updated Sun Mar 17, 2019 00:46:14)									
Delayed (victim)					Switcher (culprit)			Timestamp	CPU
PID	Prio	Total latency (µs)	T*(W**) latency (µs)	Cmd	PID	Prio	Cmd		
31028	2	41	0,0	sleep6	0	-21	swapper/6	00:38:07	6
31398	2	31	0,0	sleep2	0	-21	swapper/2	23:39:26	2
8339	99	27	27,0	cyclicttest	0	-21	swapper/2	22:05:45	2
8339	99	26	26,0	cyclicttest	0	-21	swapper/2	20:55:05	2
8339	99	25	21,2	cyclicttest	19866	-21	munin-node	21:31:14	2
8337	99	24	24,0	cyclicttest	0	-21	swapper/0	21:44:53	0
29650	2	24	10,0	sleep6	0	-21	swapper/6	22:38:00	6
8341	99	23	10,6	cyclicttest	0	-21	swapper/4	20:55:03	4
8344	99	22	10,5	cyclicttest	0	-21	swapper/7	21:45:03	7
8339	99	22	10,3	cyclicttest	0	-21	swapper/2	19:35:03	2
4357	2	22	10,1	sleep1	0	-21	swapper/1	23:14:45	1
8343	99	21	10,5	cyclicttest	0	-21	swapper/6	20:25:03	6
8342	99	21	21,0	cyclicttest	0	-21	swapper/5	21:54:42	5
8340	99	21	9,8	cyclicttest	41	-21	ksoftirqd/3	23:05:04	3
792	2	21	10,0	sleep5	0	-21	swapper/5	21:41:29	5
3929	2	21	10,0	sleep3	0	-21	swapper/3	00:13:16	3
12612	2	21	10,0	sleep1	0	-21	swapper/1	21:23:17	1
9635	2	20	10,0	sleep0	0	-21	swapper/0	20:25:04	0
9025	2	20	10,0	sleep4	0	-21	swapper/4	21:20:07	4
8710	2	20	10,0	sleep6	0	-21	swapper/6	22:19:18	6

\*Timer \*\*Wakeup (latency=timer+wakeup+contextswitch)


# Culprit/victim table

Characteristics of the 20 highest latencies: System rack9slot5s.osadl.org (updated Sun Mar 17, 2019 00:46:14)									
Delayed (victim)					Switcher (culprit)			Timestamp	CPU
PID	Prio	Total latency (µs)	T*(,W**) latency (µs)	Cmd	PID	Prio	Cmd		
31028	2	41	0,0	sleep6	0	-21	swapper/6	00:38:07	6
31398	2	31	0,0	sleep2	0	-21	swapper/2	23:39:26	2
8339	99	27	27,0	cyclicttest	0	-21	swapper/2	22:05:45	2
8339	99	26	26,0	cyclicttest	0	-21	swapper/2	20:55:05	2
8339	99	25	21,2	cyclicttest	19866	-21	munin-node	21:31:14	2
8337	99	24	24,0	cyclicttest	0	-21	swapper/0	21:44:53	0
29650	2	24	10,0	sleep6	0	-21	swapper/6	22:38:00	6
8341	99	23	10,6	cyclicttest	0	-21	swapper/4	20:55:03	4
8344	99	22	10,5	cyclicttest	0	-21	swapper/7	21:45:03	7
8339	99	22	10,3	cyclicttest	0	-21	swapper/2	19:35:03	2
4357	2	22	10,1	sleep1	0	-21	swapper/1	23:14:45	1
8343	99	21	10,5	cyclicttest	0	-21	swapper/6	20:25:03	6
8342	99	21	21,0	cyclicttest	0	-21	swapper/5	21:54:42	5
8340	99	21	9,8	cyclicttest	41	-21	ksoftirqd/3	23:05:04	3
792	2	21	10,0	sleep5	0	-21	swapper/5	21:41:29	5
3929	2	21	10,0	sleep3	0	-21	swapper/3	00:13:16	3
12612	2	21	10,0	sleep1	0	-21	swapper/1	21:23:17	1
9635	2	20	10,0	sleep0	0	-21	swapper/0	20:25:04	0
9025	2	20	10,0	sleep4	0	-21	swapper/4	21:20:07	4
8710	2	20	10,0	sleep6	0	-21	swapper/6	22:19:18	6

\*Timer \*\*Wakeup (latency=timer+wakeup+contextswitch)



# Latency fighting case #5



# Culprit/victim table

Characteristics of the 20 highest latencies: System rack6slot6.osadl.org (updated Sun Mar 17, 2019 12:43:25)									
Delayed (victim)					Switcher (culprit)			Timestamp	CPU
PID	Prio	Total latency (µs)	T*(.W**) latency (µs)	Cmd	PID	Prio	Cmd		
6409	99	2612	2597,13	cyclicttest	1802	-21	hald	03:46:17	0
6409	99	682	666,14	cyclicttest	728	-21	sensors	04:06:29	0
6409	99	673	653,18	cyclicttest	15095	-21	sensors	05:56:26	0
6409	99	655	8,618	cyclicttest	26624	-21	ssh	05:27:37	0
6409	99	652	11,6	cyclicttest	19597	-21	sh	03:47:27	0
6409	99	650	7,620	cyclicttest	0	-21	swapper	05:14:53	0
6409	99	649	7,640	cyclicttest	0	-21	swapper	05:46:42	0
6409	99	647	7,5	cyclicttest	0	-21	swapper	02:44:19	0
6409	99	645	7,636	cyclicttest	0	-21	swapper	06:28:54	0
6409	99	644	7,5	cyclicttest	0	-21	swapper	04:47:15	0
6409	99	644	7,4	cyclicttest	0	-21	swapper	02:35:16	0
6409	99	643	7,615	cyclicttest	0	-21	swapper	01:34:38	0
6409	99	643	617,23	cyclicttest	1406	50	irq/9-eth0	05:01:19	0
6409	99	641	7,632	cyclicttest	0	-21	swapper	06:14:00	0
6409	99	640	7,631	cyclicttest	0	-21	swapper	06:23:23	0
6409	99	640	7,631	cyclicttest	0	-21	swapper	05:58:46	0
6409	99	640	7,631	cyclicttest	0	-21	swapper	04:15:46	0
6409	99	640	7,630	cyclicttest	21458	-21	expr	03:06:15	0
6409	99	639	7,5	cyclicttest	0	-21	swapper	04:24:08	0
6409	99	639	619,18	cyclicttest	30600	-21	sh	05:33:08	0

\*Timer \*\*Wakeup (latency=timer+wakeup+contextswitch)

# Culprit/victim table

Characteristics of the 20 highest latencies: System rack6slot6.osadl.org (updated Sun Mar 17, 2019 12:43:25)									
Delayed (victim)					Switcher (culprit)			Timestamp	CPU
PID	Prio	Total latency (µs)	T*(.W**) latency (µs)	Cmd	PID	Prio	Cmd		
6409	99	2612	2597,13	<i>cyclictest</i>	1802	-21	<i>hald</i>	03:46:17	0
6409	99	682	666,14	<i>cyclictest</i>	728	-21	<i>sensors</i>	04:06:29	0
6409	99	673	653,18	<i>cyclictest</i>	15095	-21	<i>sensors</i>	05:56:26	0
6409	99	655	8,618	<i>cyclictest</i>	26624	-21	<i>ssh</i>	05:27:37	0
6409	99	652	11,6	<i>cyclictest</i>	19597	-21	<i>sh</i>	03:47:27	0
6409	99	650	7,620	<i>cyclictest</i>	0	-21	<i>swapper</i>	05:14:53	0
6409	99	649	7,640	<i>cyclictest</i>	0	-21	<i>swapper</i>	05:46:42	0
6409	99	647	7,5	<i>cyclictest</i>	0	-21	<i>swapper</i>	02:44:19	0
6409	99	645	7,636	<i>cyclictest</i>	0	-21	<i>swapper</i>	06:28:54	0
6409	99	644	7,5	<i>cyclictest</i>	0	-21	<i>swapper</i>	04:47:15	0
6409	99	644	7,4	<i>cyclictest</i>	0	-21	<i>swapper</i>	02:35:16	0
6409	99	643	7,615	<i>cyclictest</i>	0	-21	<i>swapper</i>	01:34:38	0
6409	99	643	617,23	<i>cyclictest</i>	1406	50	<i>irq/9-eth0</i>	05:01:19	0
6409	99	641	7,632	<i>cyclictest</i>	0	-21	<i>swapper</i>	06:14:00	0
6409	99	640	7,631	<i>cyclictest</i>	0	-21	<i>swapper</i>	06:23:23	0
6409	99	640	7,631	<i>cyclictest</i>	0	-21	<i>swapper</i>	05:58:46	0
6409	99	640	7,631	<i>cyclictest</i>	0	-21	<i>swapper</i>	04:15:46	0
6409	99	640	7,630	<i>cyclictest</i>	21458	-21	<i>expr</i>	03:06:15	0
6409	99	639	7,5	<i>cyclictest</i>	0	-21	<i>swapper</i>	04:24:08	0
6409	99	639	619,18	<i>cyclictest</i>	30600	-21	<i>sh</i>	05:33:08	0

\*Timer \*\*Wakeup (latency=timer+wakeup+contextswitch)

Output of "hwlatdetect --duration=30m" on 03/17/2019 at 03:45:

# tracer: hwlat

#

# entries-in-buffer/entries-written: 89/89 #P:1

#

#

#

#

#

#

#

#

#

#

#

```

          -----> irqsoft
          /_-----> need-resched
          || /_-----> need-resched_lazy
          ||| /_-----> hardirq/softirq
          |||| /_-----> preempt-depth
          ||||| /_-----> preempt-lazy-depth
          ||||| /_-----> migrate-disable
          ||||| /_-----> delay
TASK-PID  CPU#  |||||  TIMESTAMP  FUNCTION
<...>-16302 [000] d..... 4877473.400363: #1    inner/outer(us): 614/0    ts:1552790720.782088594
<...>-16302 [000] d..... 4877479.517988: #2    inner/outer(us): 0/612    ts:1552790730.901208053
<...>-16302 [000] d..... 4877484.911291: #3    inner/outer(us): 0/608    ts:1552790740.981327040
<...>-16302 [000] d..... 4877490.348623: #4    inner/outer(us): 0/613    ts:1552790751.061446017
<...>-16302 [000] d..... 4877495.760501: #5    inner/outer(us): 0/607    ts:1552790761.141564983
<...>-16302 [000] d..... 4877501.162081: #6    inner/outer(us): 0/611    ts:1552790771.221683938
<...>-16302 [000] d..... 4877506.570170: #7    inner/outer(us): 0/607    ts:1552790781.301802883
<...>-16302 [000] d..... 4877511.961434: #8    inner/outer(us): 0/608    ts:1552790791.381921818
<...>-16302 [000] d..... 4877587.256690: #9    inner/outer(us): 0/609    ts:1552790931.581574965
<...>-16302 [000] d..... 4877592.666464: #10   inner/outer(us): 0/613    ts:1552790941.677693936
<...>-16302 [000] d..... 4877598.068932: #11   inner/outer(us): 574/0    ts:1552790951.765812804
<...>-16302 [000] d..... 4877603.480017: #12   inner/outer(us): 0/608    ts:1552790961.861931755
<...>-16302 [000] d..... 4877608.870198: #13   inner/outer(us): 612/0    ts:1552790971.942050509
<...>-16302 [000] d..... 4877614.278978: #14   inner/outer(us): 607/0    ts:1552790982.021169240
<...>-16302 [000] d..... 4877619.688860: #15   inner/outer(us): 0/614    ts:1552790992.109288068
<...>-16302 [000] d..... 4877625.125405: #16   inner/outer(us): 0/607    ts:1552791002.197406885
<...>-16302 [000] d..... 4877693.388482: #17   inner/outer(us): 614/0    ts:1552791122.517823263
<...>-16302 [000] d..... 4877698.787595: #18   inner/outer(us): 0/606    ts:1552791132.605941955
<...>-16302 [000] d..... 4877704.205377: #19   inner/outer(us): 610/0    ts:1552791142.702060731
<...>-16302 [000] d..... 4877709.602221: #20   inner/outer(us): 0/616    ts:1552791152.790179404
<...>-16302 [000] d..... 4877715.012943: #21   inner/outer(us): 0/609    ts:1552791162.869297962
<...>-16302 [000] d..... 4877720.422802: #22   inner/outer(us): 0/607    ts:1552791172.957416616
<...>-16302 [000] d..... 4877725.967553: #23   inner/outer(us): 0/613    ts:1552791183.053535355
<...>-16302 [000] d..... 4877731.361344: #24   inner/outer(us): 609/0    ts:1552791193.133653896
<...>-16302 [000] d..... 4877736.768349: #25   inner/outer(us): 611/0    ts:1552791203.213772428
<...>-16302 [000] d..... 4877742.162564: #26   inner/outer(us): 0/607    ts:1552791213.293890950
<...>-16302 [000] d..... 4877747.568916: #27   inner/outer(us): 608/0    ts:1552791223.374009462
<...>-16302 [000] d..... 4877752.971246: #28   inner/outer(us): 611/0    ts:1552791233.454127966
<...>-16302 [000] d..... 4877758.393812: #29   inner/outer(us): 608/0    ts:1552791243.533246448
```

Trouble shooting of real-time Linux – Latency fighting

Technical Heidelberg OSADL Talks, September 30, 2020, Online Session 2b

Open Source Automation Development Lab (OSADL), Heidelberg

Output of "hwlatdetect --duration=30m" on 03/17/2019 at 03:45:

# tracer: hwlat

#

# entries-in-buffer/entries-written: 89/89 #P:1

#

#

#

#

#

#

#

#

#

#

#

#

#

#

#

#

#

#

#

#

#

#

#

#

#

#

#

#

#

#

#

#

#

#

#

#

#

#

#

#

#

#

#

#

#

#

#

#

```

          -----> irqsoft
          /-----> need-resched
          || /-----> need-resched_lazy
          || /-----> hardirq/softirq
          ||| /-----> preempt-depth
          |||| /-----> preempt-lazy-depth
          ||||| /-----> migrate-disable
          ||||| /-----> delay
          ||||| /----->
TASK-PID CPU# ||||| TIMESTAMP FUNCTION
<...>-16302 [000] d..... 4877473.400363: #1 inner/outer(us): 614/0 ts:1552790720.782088594
<...>-16302 [000] d..... 4877479.517988: #2 inner/outer(us): 0/612 ts:1552790730.901208053
<...>-16302 [000] d..... 4877484.911291: #3 inner/outer(us): 0/608 ts:1552790740.981327040
<...>-16302 [000] d..... 4877490.348623: #4 inner/outer(us): 0/613 ts:1552790751.061446017
<...>-16302 [000] d..... 4877495.760501: #5 inner/outer(us): 0/607 ts:1552790761.141564983
<...>-16302 [000] d..... 4877501.162081: #6 inner/outer(us): 0/611 ts:1552790771.221683938
<...>-16302 [000] d..... 4877506.570170: #7 inner/outer(us): 0/607 ts:1552790781.301802883
<...>-16302 [000] d..... 4877511.961434: #8 inner/outer(us): 0/608 ts:1552790791.381921818
<...>-16302 [000] d..... 4877587.256690: #9 inner/outer(us): 0/609 ts:1552790931.581574965
<...>-16302 [000] d..... 4877592.666464: #10 inner/outer(us): 0/613 ts:1552790941.677693936
<...>-16302 [000] d..... 4877598.068932: #11 inner/outer(us): 574/0 ts:1552790951.765812804
<...>-16302 [000] d..... 4877603.480017: #12 inner/outer(us): 0/608 ts:1552790961.861931755
<...>-16302 [000] d..... 4877608.870198: #13 inner/outer(us): 612/0 ts:1552790971.942050509
<...>-16302 [000] d..... 4877614.278978: #14 inner/outer(us): 607/0 ts:1552790982.021169240
<...>-16302 [000] d..... 4877619.688860: #15 inner/outer(us): 0/614 ts:1552790992.109288068
<...>-16302 [000] d..... 4877625.125405: #16 inner/outer(us): 0/607 ts:1552791002.197406885
<...>-16302 [000] d..... 4877693.388482: #17 inner/outer(us): 614/0 ts:1552791122.517823263
<...>-16302 [000] d..... 4877698.787595: #18 inner/outer(us): 0/606 ts:1552791132.605941955
<...>-16302 [000] d..... 4877704.205377: #19 inner/outer(us): 610/0 ts:1552791142.702060731
<...>-16302 [000] d..... 4877709.602221: #20 inner/outer(us): 0/616 ts:1552791152.790179404
<...>-16302 [000] d..... 4877715.012943: #21 inner/outer(us): 0/609 ts:1552791162.869297962
<...>-16302 [000] d..... 4877720.422802: #22 inner/outer(us): 0/607 ts:1552791172.957416616
<...>-16302 [000] d..... 4877725.967553: #23 inner/outer(us): 0/613 ts:1552791183.053535355
<...>-16302 [000] d..... 4877731.361344: #24 inner/outer(us): 609/0 ts:1552791193.133653896
<...>-16302 [000] d..... 4877736.768349: #25 inner/outer(us): 611/0 ts:1552791203.213772428
<...>-16302 [000] d..... 4877742.162564: #26 inner/outer(us): 0/607 ts:1552791213.293890950
<...>-16302 [000] d..... 4877747.568916: #27 inner/outer(us): 608/0 ts:1552791223.374009462
<...>-16302 [000] d..... 4877752.971246: #28 inner/outer(us): 611/0 ts:1552791233.454127966
<...>-16302 [000] d..... 4877758.393812: #29 inner/outer(us): 608/0 ts:1552791243.533246448
```

Trouble shooting of real-time Linux – Latency fighting

Technical Heidelberg OSADL Talks, September 30, 2020, Online Session 2b

Open Source Automation Development Lab (OSADL), Heidelberg

```

<...>-16302 [000] d..... 4877804.630851: #30 inner/outer(us): 608/0 ts:1552791323.597187292
<...>-16302 [000] d..... 4877810.339991: #31 inner/outer(us): 0/617 ts:1552791333.733306362
<...>-16302 [000] d..... 4877815.749037: #32 inner/outer(us): 0/613 ts:1552791343.813424765
<...>-16302 [000] d..... 4877821.153488: #33 inner/outer(us): 0/610 ts:1552791353.901543253
<...>-16302 [000] d..... 4877826.582131: #34 inner/outer(us): 0/611 ts:1552791364.013662014
<...>-16302 [000] d..... 4877831.975550: #35 inner/outer(us): 0/605 ts:1552791374.093780390
<...>-16302 [000] d..... 4877837.380887: #36 inner/outer(us): 607/0 ts:1552791384.173898756
<...>-16302 [000] d..... 4877842.779632: #37 inner/outer(us): 0/612 ts:1552791394.262017208
<...>-16302 [000] d..... 4877923.496708: #38 inner/outer(us): 0/612 ts:1552791544.541780701
<...>-16302 [000] d..... 4877928.894798: #39 inner/outer(us): 610/0 ts:1552791554.629899013
<...>-16302 [000] d..... 4877934.300006: #40 inner/outer(us): 616/0 ts:1552791564.718017316
<...>-16302 [000] d..... 4877939.697218: #41 inner/outer(us): 0/608 ts:1552791574.806135611
<...>-16302 [000] d..... 4877945.106890: #42 inner/outer(us): 612/0 ts:1552791584.885253792
<...>-16302 [000] d..... 4877950.511634: #43 inner/outer(us): 611/0 ts:1552791594.973372070
<...>-16302 [000] d..... 4877956.425184: #44 inner/outer(us): 611/0 ts:1552791605.093490714
<...>-16302 [000] d..... 4878029.559309: #45 inner/outer(us): 0/610 ts:1552791735.446018845
<...>-16302 [000] d..... 4878034.969977: #46 inner/outer(us): 0/595 ts:1552791745.526137006
<...>-16302 [000] d..... 4878040.364004: #47 inner/outer(us): 0/614 ts:1552791755.605255147
<...>-16302 [000] d..... 4878045.769635: #48 inner/outer(us): 0/606 ts:1552791765.685373291
<...>-16302 [000] d..... 4878051.193046: #49 inner/outer(us): 0/608 ts:1552791775.773491521
<...>-16302 [000] d..... 4878056.738801: #50 inner/outer(us): 610/0 ts:1552791785.885610024
<...>-16302 [000] d..... 4878062.132220: #51 inner/outer(us): 610/0 ts:1552791795.965728144
<...>-16302 [000] d..... 4878067.541579: #52 inner/outer(us): 0/613 ts:1552791806.053846349
<...>-16302 [000] d..... 4878072.937369: #53 inner/outer(us): 0/611 ts:1552791816.141964546
<...>-16302 [000] d..... 4878078.345181: #54 inner/outer(us): 0/611 ts:1552791826.222082641
<...>-16302 [000] d..... 4878083.751705: #55 inner/outer(us): 0/606 ts:1552791836.309200810
<...>-16302 [000] d..... 4878089.166895: #56 inner/outer(us): 0/615 ts:1552791846.389318889
<...>-16302 [000] d..... 4878094.560043: #57 inner/outer(us): 0/611 ts:1552791856.469436959
<...>-16302 [000] d..... 4878099.968019: #58 inner/outer(us): 0/612 ts:1552791866.549555021
<...>-16302 [000] d..... 4878146.515204: #59 inner/outer(us): 0/613 ts:1552791946.445490511


```

```

<...>-16302 [000] d..... 4878151.921186: #60 inner/outer(us): 611/0 ts:1552791956.525608502
<...>-16302 [000] d..... 4878157.340646: #61 inner/outer(us): 0/606 ts:1552791966.613726579
<...>-16302 [000] d..... 4878162.734473: #62 inner/outer(us): 607/0 ts:1552791976.693844554
<...>-16302 [000] d..... 4878168.143763: #63 inner/outer(us): 0/599 ts:1552791986.781962615
<...>-16302 [000] d..... 4878173.541694: #64 inner/outer(us): 0/612 ts:1552791996.870080668
<...>-16302 [000] d..... 4878178.950005: #65 inner/outer(us): 611/0 ts:1552792006.957198702
<...>-16302 [000] d..... 4878184.370113: #66 inner/outer(us): 0/616 ts:1552792017.045316740
<...>-16302 [000] d..... 4878189.795701: #67 inner/outer(us): 0/612 ts:1552792027.149434956
<...>-16302 [000] d..... 4878195.188437: #68 inner/outer(us): 617/0 ts:1552792037.229552885
<...>-16302 [000] d..... 4878200.590301: #69 inner/outer(us): 0/610 ts:1552792047.309670805
<...>-16302 [000] d..... 4878205.991221: #70 inner/outer(us): 0/612 ts:1552792057.389788718
<...>-16302 [000] d..... 4878211.396720: #71 inner/outer(us): 0/596 ts:1552792067.469906623
<...>-16302 [000] d..... 4878216.798039: #72 inner/outer(us): 0/609 ts:1552792077.550024520
<...>-16302 [000] d..... 4878295.026662: #73 inner/outer(us): 0/607 ts:1552792217.877665008
<...>-16302 [000] d..... 4878301.310293: #74 inner/outer(us): 0/601 ts:1552792228.069784102
<...>-16302 [000] d..... 4878306.721636: #75 inner/outer(us): 0/609 ts:1552792238.165902066
<...>-16302 [000] d..... 4878312.146302: #76 inner/outer(us): 610/0 ts:1552792248.246019836
<...>-16302 [000] d..... 4878317.542008: #77 inner/outer(us): 0/609 ts:1552792258.326137598
<...>-16302 [000] d..... 4878322.971584: #78 inner/outer(us): 0/615 ts:1552792268.405255342
<...>-16302 [000] d..... 4878328.367607: #79 inner/outer(us): 616/0 ts:1552792278.493373183
<...>-16302 [000] d..... 4878333.761087: #80 inner/outer(us): 0/614 ts:1552792288.573490923
<...>-16302 [000] d..... 4878403.657367: #81 inner/outer(us): 607/0 ts:1552792418.718010425
<...>-16302 [000] d..... 4878409.071895: #82 inner/outer(us): 607/0 ts:1552792428.798128066
<...>-16302 [000] d..... 4878414.462985: #83 inner/outer(us): 0/600 ts:1552792438.877245688
<...>-16302 [000] d..... 4878419.897597: #84 inner/outer(us): 617/0 ts:1552792448.973363502
<...>-16302 [000] d..... 4878425.295041: #85 inner/outer(us): 0/608 ts:1552792459.053481121
<...>-16302 [000] d..... 4878430.685993: #86 inner/outer(us): 0/616 ts:1552792469.133598734
<...>-16302 [000] d..... 4878436.092827: #87 inner/outer(us): 0/613 ts:1552792479.213716340
<...>-16302 [000] d..... 4878441.501407: #88 inner/outer(us): 0/612 ts:1552792489.293833938
<...>-16302 [000] d..... 4878446.898398: #89 inner/outer(us): 612/0 ts:1552792499.373951530

```






# Countermeasure

- Run the following process:

```
int timeout = bios_fixed = success = 0;

install_timeout_handler(&timeout);
while (!bios_fixed && !timeout)
    return_hardware_to_vendor(&bios_fixed);
if (timeout) {
    ask_osadl();
    docs = osadl_mkdocs(qafarm);
    success = osadl_convince_vendor(docs);
    if (!success)
        remove_feature_from_hardware(rt);
}
```

# Latency fighting case #6



# Culprit/victim table

Characteristics of the 20 highest latencies: System rackbslot4s.osadl.org (updated Sun Mar 17, 2019 00:43:30)									
Delayed (victim)					Switcher (culprit)			Timestamp	CPU
PID	Prio	Total latency (µs)	T*(,W**) latency (µs)	Cmd	PID	Prio	Cmd		
27884	99	308	305,2	cyclictest	25156	-21	kworker/u8:1	22:20:02	2
27885	99	261	245,12	cyclictest	1724	-21	gnome-shell	22:03:00	3
11961	99	261	247,11	cyclictest	1724	-21	gnome-shell	09:01:59	0
11962	99	254	249,3	cyclictest	1724	-21	gnome-shell	10:35:58	1
27884	99	253	249,2	cyclictest	1724	-21	gnome-shell	23:59:59	2
27884	99	245	241,2	cyclictest	1724	-21	gnome-shell	23:43:00	2
27884	99	243	237,3	cyclictest	1724	-21	gnome-shell	23:51:00	2
27884	99	235	222,10	cyclictest	1724	-21	gnome-shell	23:34:59	2
27882	99	225	212,10	cyclictest	1724	-21	gnome-shell	00:15:00	0
27884	99	220	215,3	cyclictest	1724	-21	gnome-shell	23:49:00	2
27882	99	204	190,11	cyclictest	1724	-21	gnome-shell	00:03:00	0
27882	99	198	182,13	cyclictest	1724	-21	gnome-shell	00:09:00	0
27884	99	167	163,2	cyclictest	1724	-21	gnome-shell	23:37:00	2
27885	99	166	161,2	cyclictest	1724	-21	gnome-shell	23:35:59	3
27885	99	124	119,2	cyclictest	1724	-21	gnome-shell	23:55:00	3
27885	99	85	81,2	cyclictest	1724	-21	gnome-shell	23:33:59	3
27884	99	84	81,1	cyclictest	1724	-21	gnome-shell	00:00:59	2
27885	99	55	24,4	cyclictest	41	-21	ksoftirqd/3	00:15:07	3
27885	99	54	29,5	cyclictest	41	-21	ksoftirqd/3	00:00:00	3
27885	99	53	28,8	cyclictest	41	-21	ksoftirqd/3	23:45:14	3

\*Timer \*\*Wakeup (latency=timer+wakeup+contextswitch)

# Culprit/victim table


Characteristics of the 20 highest latencies: System rackbslot4s.osadl.org (updated Sun Mar 17, 2019 00:43:30)									
Delayed (victim)					Switcher (culprit)			Timestamp	CPU
PID	Prio	Total latency (µs)	T*(,W**) latency (µs)	Cmd	PID	Prio	Cmd		
27884	99	308	305,2	cyclctest	25156	-21	kworker/u8:1	22:20:02	2
27885	99	261	245,12	cyclctest	1724	-21	gnome-shell	22:03:00	3
11961	99	261	247,11	cyclctest	1724	-21	gnome-shell	09:01:59	0
11962	99	254	249,3	cyclctest	1724	-21	gnome-shell	10:35:58	1
27884	99	253	249,2	cyclctest	1724	-21	gnome-shell	23:59:59	2
27884	99	245	241,2	cyclctest	1724	-21	gnome-shell	23:43:00	2
27884	99	243	237,3	cyclctest	1724	-21	gnome-shell	23:51:00	2
27884	99	235	222,10	cyclctest	1724	-21	gnome-shell	23:34:59	2
27882	99	225	212,10	cyclctest	1724	-21	gnome-shell	00:15:00	0
27884	99	220	215,3	cyclctest	1724	-21	gnome-shell	23:49:00	2
27882	99	204	190,11	cyclctest	1724	-21	gnome-shell	00:03:00	0
27882	99	198	182,13	cyclctest	1724	-21	gnome-shell	00:09:00	0
27884	99	167	163,2	cyclctest	1724	-21	gnome-shell	23:37:00	2
27885	99	166	161,2	cyclctest	1724	-21	gnome-shell	23:35:59	3
27885	99	124	119,2	cyclctest	1724	-21	gnome-shell	23:55:00	3
27885	99	85	81,2	cyclctest	1724	-21	gnome-shell	23:33:59	3
27884	99	84	81,1	cyclctest	1724	-21	gnome-shell	00:00:59	2
27885	99	55	24,4	cyclctest	41	-21	ksoftirqd/3	00:15:07	3
27885	99	54	29,5	cyclctest	41	-21	ksoftirqd/3	00:00:00	3
27885	99	53	28,8	cyclctest	41	-21	ksoftirqd/3	23:45:14	3

\*Timer \*\*Wakeup (latency=timer+wakeup+contextswitch)

# Countermeasures


- Quick fix for testing
  - `# systemctl isolate multi-user`
  - `# systemctl set-default multi-user`
- Fix with considerable impact on graphics performance
  - Additional kernel parameter: `nomodeset`
  - Use frame buffer graphics
- Better
  - There is nothing better
- Even better
  - No, you can't fix it yourself.

# Before







After




After





# SIMs hidden behind other latency




# Latency fighting case #7

























# Countermeasures

- Quick fix for testing
  - `# systemctl isolate multi-user`
  - `# systemctl set-default multi-user`
- Fix with considerable impact on graphics performance
  - Additional kernel parameter: `nomodeset`
  - Use frame buffer graphics
- Better
  - There is nothing better
- Even better
  - No, you can't fix it yourself.


# Before



After




Before  
and  
after




Trouble shooting of real-time Linux – Latency fighting  
 Technical Heidelberg OSADL Talks, September 30, 2020, Online Session 2b  
 Open Source Automation Development Lab (OSADL), Heidelberg



Before  
and  
after



# Before



After

