

Appendix: List of events which Linux Kernel State Tracer records

Copyright (C) Hitachi, Ltd., 2002. All rights reserved.

Event type [hex]	Category	Mnemonic	Description of events	where to hook	filename	data recorded as "log_arg1"	data recorded as "log_arg2"	data recorded as "log_arg3"	data recorded as "log_arg4"	remarks	
01	Process management	PROCESS_CONTEXTSWITCH	Process context switching	schedule()	/kernel/sched.c	address of the task_struct of "prev"	address of the task_struct of "next"	prev_process state (value after switch)	data recorded as "log_arg4" prev_process count (value before switch)	Remarks from log_arg3, log_arg4, can determine why processes were switched	
02		PROCESS_WAKEUP	WAKEUP	try to wake up()		value of "p" in the function	synchronous				
03		PROCESS_SIGSEND	SENDING SIGNAL	send sig info()		value of "sig" in the function	value of "t" in the function	pointer to info (info)			
04		PROCESS_LTHREADEN	CREATING A KERNEL THREAD	creating a kernel thread		/arch/i386/kernel/process.c	value of "t" in the function	pointer to argument of kernel thread	flag		
10	Interrupts	INT_HARDWARE_ENTRY	hardware	entrance	do_IRQ()	value of "irq" in the function	interrupt status (status)				
12		INT_TASKLETH_ENTRY	software	entrance	tasklet_hi_action()	value of "t->func" in the function					
14		INT_TASKLET_ENTRY	software	entrance	tasklet_action()	value of "t->func" in the function					
16		INT_BH_ENTRY	software	entrance	bh_action()	value of "nr" in the function	address of action (bh_base)				
20	Exceptions	EXCEPTION_ENTRY	de								
			int3								
			overflow								
			bounds								
			invald_gp								
			double_fault								
			coprocessor_segment_overnrun								
			invald_TSS								
			segment_not_present								
			stack_segment	entrance	error_code			handler address (edi)	error code (esi)	exception occurred address (eip)	
			alignment_check								
			coprocessor_error								
			smiud_coprocessor_error								
			debug								
	general_protection										
	page_fault										
	machine_check										
	spurious_interrupt_bug										
	device_not_available										
	nmi										
21		EXCEPTION_EXIT	device_not_available	exit		handler address	the number of this exception				
			exceptions other than above two			handler address (edi)					
30	System calls	SYSCALL_ENTRY	entrance	beginning of system call()	/arch/i386/kernel/entry.S		the number of this system call			recording arguments of system calls is optional feature	
31		SYSCALL_EXIT	exit	ending of system call()	/arch/i386/kernel/entry.S		the number of this system call				
40	Filesystems	FS_DEVRW	device IO	creation of request for device	ll_rw_block()	/drivers/block/ll_rw_blk.c	buffer (bh)	READ/WRITE (rw)	num of blocks to transfer (nr)		
41		FS_DEVRND		completion of request for device	end_buffer_io_sync()	/fs/buffer.c	buffer (bh)	uptodate			
42		FS_BUFFERASY		buffer busy wait	wait_on_buffer()	/fs/buffer.c	buffer (bh)				
50		MEM_SWAPOUT	swap out	swap out	try to swap out()	/mm/vmstat.c	pointer to page swapped out (page)				
51		MEM_SWAPIN	swap in	swap in	do_swap_page()	/mm/memory.c	pointer to page swapped in (page)				
52		MEM_DO_NOPAGE	mem do nopage	mem do nopage	do_no_page()	/mm/memory.c	pointer to page allocated (new_page)				
53		MEM_DO_WPPAGE	mem do wppage	mem do wppage	do_wp_page()	/mm/memory.c	pointer to page (new_page)				
54		MEM_WAIT_PAGE	mem wait page	mem wait page	wait_on_page()	/mm/filemap.c	pointer to page (page)				
55		MEM_GET_FREEPAGE	mem get freepage	mem get freepage	get_free_page()	/mm/page_alloc.c	pointer to page (paddr)	type of page (gfp_mask)	the number of page (order)	call address	
56		MEM_GET_ZEROPAGE	mem get zeropage	mem get zeropage	get_zeroed_page()	/mm/page_alloc.c	pointer to page (address)	type of page (gfp_mask)	call address		
57	MEM_FREEPAGE	mem freepage	mem freepage	free_page()	/mm/page_alloc.c	pointer to (paddr)	the number of page (order)	call address			
58	MEM_VMALLOC	mem vmalloc	mem vmalloc	vmalloc()	/mm/vmalloc.c	address (addr)	size	call address			
59	MEM_VFREE	mem vfree	mem vfree	vfree()	/mm/vmalloc.c	address (addr)					
5a	MEM_CACHE_CREATE	mem cache create	mem cache create	kmem_cache_create()	/mm/slab.c	name	size	cachep			
5b	MEM_CACHE_ALLOC	mem cache alloc	mem cache alloc	kmem_cache_alloc()	/mm/slab.c	cachep	flags	obp	call address		
5c	MEM_MALLOCC	mem malloc	mem malloc	kmalloc()	/mm/slab.c	cachep	flags	obp	call address		
5d	MEM_CACHE_FREE	mem cache free	mem cache free	kmem_cache_free()	/mm/slab.c	cachep	obp	call address			
5e	MEM_FREE	mem free	mem free	kfree()	/mm/slab.c	obp	call address				
60	Networking	NET_PKTSEND	sending packets	entrance	dev_queue_xmit()	/net/core/dev.c	skb				
61		NET_PKTSENDI	interrupt on sending packets	entrance	net_tx_action()	/net/core/dev.c	it				
62		NET_PKTRECV	receiving packets	entrance	netif_rx()	/net/core/dev.c	skb				
63		NET_PKTRECVI	interrupt on receiving packets	entrance	net_rx_action()	/net/core/dev.c	h				
64	NET_SOCKETIF	socket()	entrance	sys_socketcall	/net/socket.c	call	args		exit is recorded as exit of system call.		
70	SysV IPC	SYSV_IPC	IPC functions	entrance	sys_ipc()	/arch/i386/kernel/sys_i386.c	call/first	second/third	*ptr		
80		LK_SPINLOCK	spin lock	lock	spin_lock()		address where it was called	lock		inline	
81		LK_SPINTRYLOCK	spin lock	try lock (exit)	spin_trylock()		address where it was called	lock	return value	inline	
82		LK_SPINUNLOCK	spin lock	unlock	spin_unlock()		address where it was called	lock		inline	
83		LK_WRLCK	write lock	write lock	write_lock()	/include/asm-i386/spinlock.h	address where it was called	rwlock		inline	
84		LK_WRTWYLOCK	write try lock (exit)	write try lock (exit)	write_trylock()		address where it was called	rwlock	return value	inline	
85		LK_WRLUNCK	read/write lock	write unlock	write_unlock()		address where it was called	rwlock		define	
86		LK_RDLOCK	read lock	read lock	read_lock()		address where it was called	rwlock		inline	
87	LK_RDUNLOCK	read unlock	read unlock	read_unlock()		address where it was called	rwlock		define		
a0	Timer	TIMER_RUN	run timer list	run timer list		function address (fn)	argument for the function (data)				
a1		TIMER_ADD	add to timer list	add timer()		pointer to timer list (timer)	unexpired term (timer->expires)	function address (timer->function)	argument for the function (timer->argument)		
a2		TIMER_MOD	modify timer list	modify timer()	/kernel/timer.c	pointer to timer list (timer)	unexpired term (timer->expires)	function address (timer->function)	argument for the function (timer->argument)		
a3		TIMER_DEL	delete from timer list	del timer()		pointer to timer list (timer)	unexpired term (timer->expires)	function address (timer->function)	argument for the function (timer->argument)		
a4	TIMER_DEL_SYNC	delete from timer list with synchronous	del timer sync()		pointer to timer list (timer)	unexpired term (timer->expires)	function address (timer->function)	argument for the function (timer->argument)			
90	Others	O_PORTRN	io commands	port output		port address/byte width	value to output	address where it was called		inline	
91		O_PORTRIN	io commands	port input		port address/byte width	value to input	address where it was called		inline	
92		O_PANIC	panic	panic	panic()	/kernel/panic.c	address of argument	address where it was called			
93	O_PRINTK	printk	printk	printk()	/kernel/printk.c	address of argument	address where it was called				
100	LKST internal event	LKST_INIT	Progress of LKST initialization process	lkst_init_stage(0-1)()	/driver/lkst/lkst.c	initialization status					
108		LKST_MSET_XCHG	LKST switches the masksets	lkst_evhandlerprim_maskset_xchg_inlin	/driver/lkst/lkst.c	old maskset ID	new maskset ID	pointer to old maskset	pointer to new maskset	Recorded 2 times; before/after	
110		LKST_BUFF_SHIFT	LKST shifts the buffers	lkst_evhandlerprim_bufshift_inlin()	/driver/lkst/lkst.c	old buffer ID	new buffer ID	pointer to old buffer	pointer to new buffer	Recorded 2 times; before/after Used for automatically shifting buffer. If masked, LKST stops it.	
111		LKST_BUFF_OVERFLOW	overrun occurred in the current buffer.	lkst_evhandlerprim_entry_next()	/include/linux/lkst_private.h	pointer to the buffer					
119		LKST_SYNC_UID	Synchronization with UID	sys_uid(), set_user()	/kernel/timer.c, sys.c	UID		pointer to the process table		for compensation of dropped log data	
11a		LKST_SYNC_GID	Synchronization with GID	sys_gid()	/kernel/timer.c, sys.c	GID		pointer to the process table		for compensation of dropped log data	
11b	LKST_SYNC_PGID	Synchronization with PGID	sys_pgid(), sys_setsid()	/kernel/sys.c	PGID		pointer to the process table	session leader flag	for compensation of dropped log data		
11c	LKST_SYNC_TID	Synchronization with TID	sys_gettid()	/kernel/timer.c, sys.c	TID(pid)		pointer to the process table		for compensation of dropped log data		