

TCB safety

Tuesday 25 August 2020 07:40 (20 minutes)

Thread Control Block (TCB) is a data structure in the Linux kernel which contains thread-specific information needed to manage it.

The Thread Control Block acts as a library of information about the threads in the system.

TCB is being manipulated by the kernel constantly, while the thread is being executed and while it is switched out.

Assuring the integrity of the TCB is critical to achieve safe thread life cycle management in Linux.

As part of making TCB management safe, several tasks will need to be performed:

1. Analysis of the TCB
 - What kind of information is stored in TCB. For example:
 - o All flags (unless there is a specific justification for an exception).
 - o Namespace Cgroups information
 - o Signal handlers
 - o MMU list
 - o Security fields
 - o Dependencies on LSMs (e.g., `in_execve` or `brk_randomized`)
 - o `stack_canary`
 - o `seccomp` related data
 - o stack pointer
 - o parent pointer
 - o child/sibling lists
 - o PI data structures
 - o RT mutexes
 - o `futex` list
 - o NUMA balancing fields
 - o `tlbflush_unmap_batch` data
 - What is the criticality of this information to the thread execution (Categorization to critical/non critical, etc..). For example:
 - o Parent pointer
 - o Signal handlers
 - Identify the safety critical part(s) of the TCB
1. Analysis of the possible failure modes
 - What possible faults might be caused by the kernel, that will influence the TCB. For example:
 - o Altering of data during context switch out
 - o Corruption of data while thread is not running (e.g. due to bit flip)
 - 2. Propose solutions for protecting the TCB –Examples:
 - Kernel configurations on kernel space code –Protect the kernel space code and data by using kernel self protection mechanisms (e.g., enable `CONFIG_HARDENED_USERCOPY`, or disable `CONFIG_DEVMEM`)
 - CRC the safety critical data after switch out
 - Allocate RO block and store immutable safety critical data in that block

I agree to abide by the anti-harassment policy

I agree

Primary authors: Dr COPPERMAN, Elana (Mobileye); Mr DAVIDOVICH, Rafi (Mobileye)

Presenters: Dr COPPERMAN, Elana (Mobileye); Mr DAVIDOVICH, Rafi (Mobileye)

Session Classification: Kernel Dependability & Assurance MC

Track Classification: Kernel Dependability & Assurance MC