Contribution ID: **134** Type: **not specified**

Integral Atomic Stack Switching for IST Exceptions

Friday, 20 September 2024 15:50 (20 minutes)

The x86_64 exception handling relies on complex indirect system structures such as the IDT and TSS. This process can sometimes involve complicated stack switching, which further complicates the situation when it comes to ring changes, syscall gaps, unblockable reentrant IST exceptions, the increasing number of Coco-introduced IST exception types, the nesting of the IST exceptions, and so forth, along with the necessity for accurate switching of GSBASE, CR3 or other bits related to mitigations.

The dancing of the IST stacks represents a major challenge; the NMI stack-switching had led to CVEs and the current more cumbersome and burdensome #VC stack-dancing adds more strain. The varied approaches used by different exceptions exacerbate the issues, making them more entrenched.

In this session, we introduce a new Integral Atomic Stack Switching mechanism. This mechanism aims to consolidate the diverse segregated stack-switching processes and handle all the essential event-handling states in a unified, atomic manner. We will explore the current problems, outline the design of the mechanism, examine how it addresses the issues, and discuss other potential derived benefits, such as enabling shadow stacks in the kernel.

Primary author: JIANGSHAN, Lai (AntGroup)

Presenter: JIANGSHAN, Lai (AntGroup)

Session Classification: x86 MC

Track Classification: x86 Microconference