



Contribution ID: 262

Type: not specified

Automating Linux Kernel Documentation for Safety-Critical Compliance through Large Language Models

Thursday 11 December 2025 17:45 (45 minutes)

ABSTRACT

The lack of standardized documentation for the Linux kernel poses a barrier to its adoption in safety-critical industries such as aerospace, where compliance with standards like DO-178C is required. We explored the use of locally trained Large Language Models (LLMs) to automatically generate compliant documentation for kernel modules and tools. As a case study, we applied this approach to the Linux kernel's ftrace utility and evaluated four LLM families—Meta Llama, StarCoder2, Mistral Devstral, and Google Gemma—across documentation validity, Graphical Processing Unit (GPU) utilization, and throughput. Results show that Mistral Devstral produced the most accurate and standards-aligned documentation and demonstrated that LLMs can provide an effective method for bridging the gap between open-source software and regulated environments, enabling safer and broader integration of the Linux kernel into aerospace and other compliance-driven domains.

Reason Behind Effort

The broader goal is not to propose an immediate solution, but to present empirical results that raise questions for the community: What criteria should kernel-generated documentation meet? Can LLMs be integrated into existing toolchains (e.g., Sphinx, SPDX) to support compliance goals? What processes would allow reproducibility, traceability, and expert validation in a way that certification authorities might accept?

Primary author: STANLEY, Justin (VES LLC)

Co-authors: STENSLAND, Grant (VES LLC); RICKS, Morgan (VES LLC); HILLIARD, Toby (VES LLC); ICE, Tom (VES LLC)

Presenters: STENSLAND, Grant (VES LLC); STANLEY, Justin (VES LLC); RICKS, Morgan (VES LLC); HILLIARD, Toby (VES LLC); ICE, Tom (VES LLC)

Session Classification: LPC Refereed Track

Track Classification: LPC Refereed Track