

LOCI: AI-Driven Observability & Performance Intelligence for Compiled Binaries

```
1 /**  
2  * amdgpu_gart_dummy_page_fini - free dummy page used by  
3  the driver  
4  */  
5  * @adev: amdgpu_device pointer  
6  */  
7  * Frees the dummy page used by the driver (all asics).  
8  */  
9  void amdgpu_gart_dummy_page_fini(struct amdgpu_device  
10 *adev)  
11 {  
12     if (!adev->dummy_page_addr)  
13         return;  
14     dma_unmap_page(&adev->pdev->dev,  
15 adev->dummy_page_addr, PAGE_SIZE,  
16 DMA_BIDIRECTIONAL);  
17     adev->dummy_page_addr = 0;  
18 }
```

Shift Observability Left

LOCI, Line-of-Code Intelligence platform, transforms observability and shift-left approach by extracting deep performance insights from compiled binary files, without requiring source code.

Traditional static analysis and observability tools fail to detect performance issues in compiled BIN files due to missing execution context, hardware interactions, and real-time software behavior analysis.

LOCI bridges this gap by modeling compiled binaries with real-world execution data, enabling early detection of performance degradations, power inefficiencies, and test coverage prediction, after project build.

How LOCI Works

LOCI leverages Aurora Labs' proprietary vertical LLM, known as Large Code Language Model (LCLM), that is specifically designed for compiled binaries. Unlike general-purpose Large Language Models (LLMs), LCLM delivers superior, efficient, and accurate binary analysis and detection of software behavior changes, offering deep contextual insights into system-wide impacts - without the need of source code.

LOCI has accumulated extensive performance intelligence from real-world execution data, enabling shift observability left and allowing it to:

- Predict performance impacts & detect degradations across software versions
- Identify power-hungry functions & optimize resource efficiency
- Provide test coverage references to enhance testing & validation

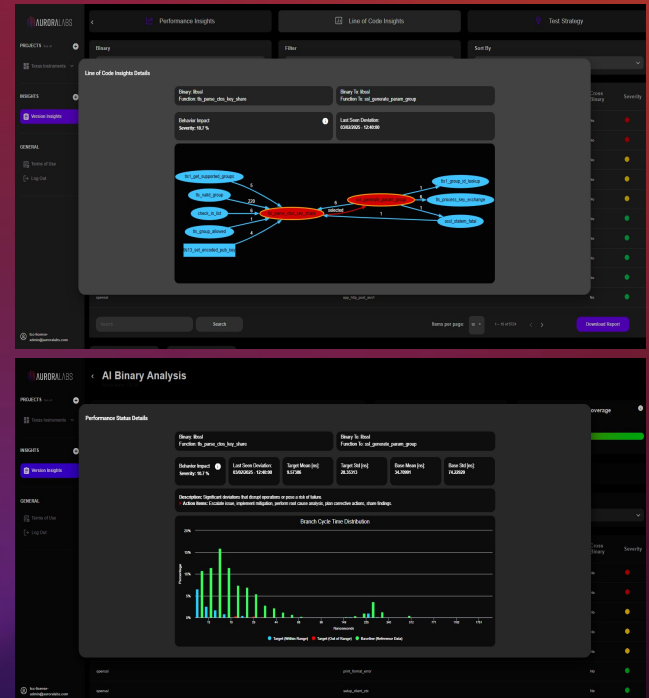
Converge static and dynamic analysis for early-stage optimizations

Unlock R&D Productivity & Reduce Time to Resolution

LOCI dramatically reduces Mean Time to Resolution (MTTR) and enhances key performance indicators (KPIs) by shifting performance analysis left in the development cycle.

Key Benefits:

- Detects issues immediately after build
- Rapid root cause analysis
- Minimizes costly late-stage performance issues
- Faster performance debugging & fixes
- Power efficiency gains
- Enhance testing optimization



Seamless CI/CD Integration for Early Optimization

LOCI integrates directly into CI/CD pipelines, providing real-time feedback post-build and eliminating the need for extensive post-deployment profiling. This ensures continuous monitoring, early anomaly detection, and reduced rollbacks/hotfixes, leading to a more stable and efficient release software cycle.

Why Aurora Labs?

- ML & LLM tuning expertise since 2017
- Proven track record of using ML/LLM to optimize embedded & HPC systems
- 100+ Patents in AI, Embedded Systems & Observability
- ISO 27001, ISO 21434, ASPICE L2, ASIL B, and ISO 26262 Certified
- Trusted by Global Industry Leaders: LG, VW, AWS, ST Micro, Infineon, NTT Data and Toyota - TTC

Get Started **Today!**