LLMScan.online

AI Model Security Scan Report

D

Score 54/100

Report summary

GGUF v3 model | architecture: llama | quantization: F32 | 1 tensors | 1 blocking issue / 2 warnings

Model

Filename	`ssti_stmt.gguf`
Format	`gguf`
Architecture	`llama`
Size	`416 B`
SHA-256	`f50223da0b0c7f0d8ed236a12de752192b1335aa928d3e5fe9d4b9b62783a7c4`
Scan mode	`static-deep`
Scanned at	`2026-10-23T18:15:11+00:00Z`
Report ID	`83f3efec600441c6bff34d7c39627398`

Findings (3)

SEVERITY	CATEGORY	DETAIL
CRITICAL	Metadata Security	**Malicious Jinja2/SSTI pattern in chat template** Detected 1 dangerous pattern(s): \{[{%][^}]*(__import__|__builtins__|__class__|__mro__|__subc. These are real SSTI vectors (dunder introspection, os/exec access, config/request context leakage) -- not normal template constructs. **Recommendation:** Reject this model; do not render the template outside a strictly sandboxed environment. Normal constructs (set, for, if, namespace, tojson, tools, tool_calls) are not flagged.
LOW	Supply Chain	**No license field in GGUF metadata (general.license)** Absence of license information makes it impossible to determine usage rights **Recommendation:** Add general.license to model metadata before distributing; use an SPDX identifier
LOW	Supply Chain	**No embedded provenance metadata** None of general.author, general.source.url, general.base_model.0, or general.source.huggingface.repository are embedded in the GGUF file. Without any origin signal the model cannot be traced to a trusted source or compared against a reference checkpoint for tampering detection. For files uploaded directly, provenance is especially critical because the scanner cannot infer repository origin from the file alone. **Recommendation:** Add at least one of: general.author, general.source.huggingface.repository, or general.source.url

Runtime compatibility

RUNTIME	SUPPORTED	NOTES
llama.cpp	Yes	GGUF v3 container: supported | quant F32: supported
Ollama	Yes	GGUF via llama.cpp backend: container and quantization supported
LM Studio	Yes	GGUF container and quantization supported; GPU/CPU inference

vLLM	No	GGUF not supported; requires safetensors format
HuggingFace Transformers	No	GGUF not supported; requires safetensors or PyTorch checkpoint

Generated 2026-11-06 03:45 UTC by LLMScan.online (engine v2.0.0, report schema v1.0). Full interactive report: https://llmscan.online/report/83f3efec600441c6bff34d7c39627398 — Static analysis only; not a guarantee of safety. See llmscan.online/docs/how-it-works for methodology.