




ZIKUU miniã®?è¼æ©?ã??ã ¥æ??

Description

Minisforumã®AI X1 Proã??æ??ã«ã ¥ã??ã¼ã·?ã·?ã??
64GB RAM/1TB SSDã®ã?¢ã??ã?«ã·§ã·?ã??
ã·?ã??ã??ZIKUU miniã®?è¼æ©?ã·?ã·?ã¼ã·?ã??ã·?ã·?ã¼ã·?ã??



- npm
- VSCodium
- neovim
- tmux
- Python 3
- python-is-python3
- ripgrep
- jq
- git
- cmake
- build-essential

llama.cpp

GitHub•?ã??ã? ¯ã?ã?¼ã?³

git clone https://github.com/ggml-org/llama.cpp.git

Vulkanç°èª•

vulkaninfo --summary ã•šã?•

AMD Radeon 890M Graphics (RADV GFX1150)
Mesa RADV
Vulkan 1.4

ã??ç°èª•ã??

ROCmã• ¯æ?ªã°?ã ¥ã??

ls /opt/rocm

No such file or directory

ã•¾ã•?Vulkan backendã??æ?iç?¨ã??

llama.cppã•®ã??ã?«ã??

ã??ã?«ã??ã•«¿ è|ã•ªã??ã•®ã??ã?ªã?³ã?¹ã??ã?¼ã?«ã??

sudo apt install libvulkan-dev glslc spirv-headers glslang-tools spirv-tools

ã??ã?«ã??

cmake -B build -DGGML_VULKAN=ON
cmake --build build -j

æ?“è«?ă??ă?³ă?•ă??ă?¼ă?™

Qwen3.5 9B Q4_K_M

å®?èj?:

```
./build/bin/llama-bench \  
-m ../models/Qwen3.5-9B-Q4_K_M.gguf \  
-ngl 999
```

GPU:

AMD Radeon 890M Graphics
RADV GFX1150
UMA: 1

çµ•æ??:

é ç?®	çµ•æ??
Model size	5.28 GiB
Params	8.95B
pp512	367.67 t/s
tg128	13.46 t/s

å¾¾è©±ç?“é??ă•ă?ă•ă•ă½¿ă?ă??ă•?ă?•ç??æ?•é??å!ă•™dGPUă??ă??ă?ă•ă??ă½?ă•?ă??

gpt-oss 20B Q4_K_M

å®?èj?:

```
./build/bin/llama-bench \  
-m ../models/gpt-oss-20b-Q4_K_M.gguf \  
-ngl 999
```

çµ•æ??:

é ç?®	çµ•æ??
Model size	10.81 GiB
Params	20.91B
pp512	401.93 t/s
tg128	27.68 t/s

20Bă®gpt-ossă®æ?¹ă•9B Qwen3.5ă??ă??é«?é??ă??

ç??æ?•é??å!:

Qwen3.5 9B	13.46 t/s
gpt-oss 20B	27.68 t/s

ç'2å?•ã??

MoEæ\$?é? ã"890M/Vulkan backendã"ã•®ç?,æ?§ã•?è?¬ã?å•¬è?½æ?§ã?é«?ã•?ã??

ã?;ã?¢ã?"ã•®æ??å??

890Mã«ã¬å°§ã•ã°?ç?"VRAMã??å?²ã??å½?ã•|ã•|ã•?ã•ã•?ç?¶æ ?ã??

Qwen3.5 9Bæ?"è«?æ??

å¹³å„æ??:

RAM used	4.5 GiB
VRAM used	695 MB
GTT used	105 MB

æ?"è«?æ??:

RAM used	9.5 GiB
VRAM used	1.84 GB
GTT used	5.28 GB

GGUF:

5.28 GiB

ã?¢ã??ã?«ã?µã?°ã"ã»ã¼å•?é?•ã?GTTã,è¼?ã•£ã|ã•?ã??ã??

gpt-oss 20Bæ?"è«?æ??

RAM used	15 GiB
RAM available	45 GiB
Swap	0
VRAM used	1.85 GB
GTT used	11.12 GB

GGUF:

10.81 GiB

ã•?ã??ã??ã?¢ã??ã?«æ?¬å½?ã•?ã»ã¼GTT / shared RAMå'ã,å±?é??ã•?ã??ã•|ã•?ã??ã??

ç¾æ??ç?'ã•§ã•®è©?ã¾i

AI X1 Proã¬ç¾ç?¶ã•

HX 370
12C / 24T
64GB RAM
Radeon 890M

2.5GbE 2
1TB NVMe

Ubuntu, 890M

890M

VRAM 2GBç GPU

890M

64GB UMAç GPU

890M

20B Q4ç GPU offloadç

RAM available 45 GiB
Swap 0

890M

890M

27.68 token/s

890M

890M

ZIKUU 890M

890M

- PostgreSQL
- MinIO
- Qdrant
- Gitea
- LLM Broker
- CPS/Nerve/Pivot

890M

890M

ã•?ã??ã•\$Alâ¼é?·ã?•Repository Explorerã?•Pivot Reasoning Engineã•ã•?ã•£ã•
?ã¼è©±ç³»ã?£ã??ã?ªã?±ã?¼ã?·ã?Sã?³ã?å¿«é•©ã•«ã½¿ã•?ã??ã•°ã?ã??ã??ã•@ã•²ã??ã•
«ä¹ã??ã??æ??æ??ã?•ã??ã?¯ã?£ã??ã?? = ZIKUU miniã?ã?æ?•ç«?ã?ã??ã?ã?ã?ã?ã?ã?ã?ã??
å¼ã?ã•ç¶?ã•èªæ?»ã?»ç ?ç©¶ã?»é??ç?°ã??ç¶?ã?ã??ã??

Category

1. æ?¥ã? ã•@æ´»å??

Tags

1. AI
2. LLM
3. Strix Point
4. ZIKUU Mini
5. ä¿é??å??å¿
6. å±±æ£ç??

Date Created

2026å¹8æ??24æ?¥

Author

kazuo-tsubaki

default watermark