KINDI

MODERN 2B TO 552B LLMS ON LOW-COST HARDWARE
FROM A QUADRO K420 TO DEEPSEEK-V4.1-FLASH

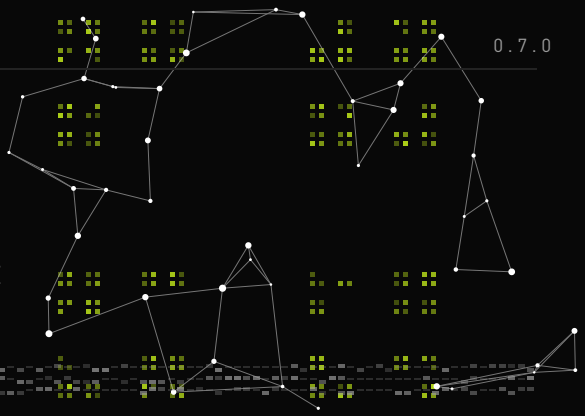

KINDI FIELD MANUAL · SEVENTH EDITION · FOR KINDI 0.7.0

How To Get Work Done with Kindi and a $15 GPU

The definitive guide to running private 2B to 552B language models on cost-effective hardware: NVIDIA workstation and GeForce GPUs, mining cards, AMD Radeon cards, unified-memory AMD mini PCs, Windows PCs and bare-metal CPUs—with zero cloud telemetry.

22.9 tok/s
Quadro K420 · Full 2B agent in a US$15 card

115.5 tok/s
RTX 3080 · 27B with the DSpine drafter, the same output

310 tok/s
RTX 3080 · Lightning 30B-A3B, 1.5× llama.cpp

18.7 tok/s
RTX 3060 · 552B DeepSeek-V4.1-Flash, with host RAM

37.1 tok/s
RTX 3060 · 125B Qwen3.8-Flash-Next, with host RAM

0 bytes
100% Private · Zero cloud telemetry, all inference on your machine

Amir Husain

amirhusain.com/kindi
github.com/ahtx/kindi

Contents

AT A GLANCE	Kindi 0.7.0 supports	2
PREFACE	About this manual	5
CHAPTER 1	Why should you care?	10
CHAPTER 2	How fast can a $15 GPU be?	13
CHAPTER 3	Which models run on which cards?	16
CHAPTER 4	Which card for which job?	20
CHAPTER 5	How fast is each card?	26
CHAPTER 6	How does a draft model make the 27B faster?	45
CHAPTER 7	DSpine, the block drafter of the 27B	48
CHAPTER 8	A more capable model	51
CHAPTER 9	Mixture-of-experts models	57
CHAPTER 10	DeepSeek-V4.1-Flash on one RTX 3060	62
CHAPTER 11	Qwen3.8-Flash-Next	67
CHAPTER 12	Can two cards work as one?	71
CHAPTER 13	A mining card for AI	74
CHAPTER 14	One small box for everything	79
CHAPTER 15	Kindi on AMD graphics	84
CHAPTER 16	What does the prompt cache do?	91
CHAPTER 17	What does a power cap change?	95
CHAPTER 18	How do you start on a new computer?	97
CHAPTER 19	Which card runs which model?	102
CHAPTER 20	Kindi on Windows	105
CHAPTER 21	What does the server do for an agent?	115
CHAPTER 22	Meet pi	119
CHAPTER 23	Kindi Bridge	123
CHAPTER 24	What is an agent?	139
CHAPTER 25	Fast decisions on short text	141
CHAPTER 26	What works, and what to avoid?	145
CHAPTER 27	Can a big job be broken down?	160
CHAPTER 28	Kindi in your own workflows	166
CHAPTER 29	How do you get the best results?	170
CHAPTER 30	Which model for which job?	171
CHAPTER 31	Building a business with Kindi	176
CHAPTER 32	What is new in 0.7.0	185
CHAPTER 33	How did we test?	188
CHAPTER 34	Where can you read more?	191
APPENDIX A	All results	195

AT A GLANCE

Kindi 0.7.0 supports

Family	Card	Models, and the largest context in tokens
Kepler	Quadro K420 1 GB	2B: 24,576
	Quadro K420 2 GB	2B: 65,536 · 4B: 23,808
	Quadro K6000 12 GB	27B: 65,536 (up to 156,672) · 33B: 8,448 · Lightning: 32,768
Maxwell	Quadro K1200 4 GB	2B: 131,072 · 4B: 86,272 · 9B: 47,360 · Nemotron: 8,448
	Quadro M4000 8 GB with about 96 GB or more of host RAM	Qwen3.8-Flash-Next
	Two Quadro K1200 cards	27B: 24,576 · Nemotron: 8,448
Pascal	Tesla P100 16 GB	27B: 65,536 · Lightning: 32,768 · 2B, 4B, 9B
	Two Tesla P100 cards	33B
	GeForce GTX 1080 8 GB	27B: 65,536 · Nemotron: 24,576 · Lightning: 32,768
Volta	Tesla V100 32 GB, one or two cards	33B
Ampere	GeForce RTX 3060 12 GB	27B with the 2B drafter: 24,000
	RTX 3060 12 GB with about 96 GB or more of host RAM	Qwen3.8-Flash-Next
	RTX 3060 12 GB with about 560 GB of host RAM	DeepSeek-V4.1-Flash: 65,536
	GeForce RTX 3080 10 GB	27B with the 2B drafter: 84,992 (103,680 without it) · 27B with DSpine: 32,768 · Nemotron: 24,576 · Lightning: 60,320
	Two GeForce RTX 3080 cards	33B: 16,384
	GeForce RTX 3090 24 GB	2B, 4B, 9B, 27B: 16,384 · 33B: 16,640 · Lightning: 32,768
	NVIDIA CMP 170HX 40 GB (device 0x2082), one or two cards	27B and 33B: 65,536 · 2B, 4B, 9B
Turing, Ada, Blackwell	GeForce RTX 20, 40 and 50 series	Supported; `kindi verify` passes
	GeForce RTX 2060 12 GB, RTX 2070 SUPER 8 GB	Lightning: 32,768 (RTX 2060), 8,192 (RTX 2070 SUPER)
AMD	Radeon 780M, 64 GB computer	27B: 24,576 · Nemotron: 24,576 · 2B · 9B
	Radeon 840M (Ryzen AI 5 340)	2B, 4B, 9B, 27B: 8,192
	Radeon 8060S (Strix Halo, Ryzen AI MAX+ 395)	2B, 4B, 9B, 27B: 8,192
	Radeon PRO W6800 32 GB (RDNA2)	2B, 9B, 27B: 24,576
CPU	x86-64 with AVX2	2B and 4B: 24,576

Family	Card	Models, and the largest context in tokens
Windows	Windows 10 and 11, through WSL2 or native (`kindi.exe`)	The models of your NVIDIA card; the CPU engine through WSL2

Model	File	Smallest card memory	Speed methods that are on
Bonsai 2B	429 MB	1 GB	Prompt cache; MTP head up to 8K (RTX 30 and later)
Bonsai 4B	950 MB	2 GB	Prompt cache; the tensor-core path (RTX 30 and later)
Bonsai 9B	2.0 GB	4 GB	Prompt cache
Bonsai 27B	5.95 GB	8 GB, or two 4 GB cards	Prompt cache; the DSpine drafter (one tensor-core card), else the 2B drafter (RTX 30 and later); the b2 weight copy (sm_80 cards with 20 GB or more)
Nemotron 3.5 Bolt E4B (Gemma 4 family)	4.2 GB	4 GB, or two 4 GB cards	Prompt cache
AH-Nanbeige4.3-33B IQ3_XXS	13.1 GB	16 GB, or about 17 GB on several cards (RTX 30 and later, Tesla P100); a Quadro K6000 with host memory	Prompt lookup; the 8-bit K/V cache when the f16 cache does not fit
Nemotron 3.5 Lightning 30B-A3B (ternary, mixture of experts)	9.47 GB	6 GB (GTX 10 series and later, Tesla P100); a 12 GB Kepler card	Prompt cache; host-memory experts on cards with less than 10 GB
DeepSeek-V4.1-Flash (mixture of experts)	494 GB	12 GB, and about 560 GB of host RAM	The DSpark drafter; the prompt cache with saved states; hot experts on a second card
Qwen3.8-Flash-Next (mixture of experts)	83.6 GB	8 GB free, and about 96 GB of host RAM	The most used experts on the card; prompt reuse

Each context is the largest that we used or measured on that card. "Nemotron" is Nemotron 3.5 Bolt E4B, and it is experimental. "Lightning" is Nemotron 3.5 Lightning 30B-A3B. The installer shows the setups that fit your computer (How do you start on a new computer?).

PREFACE

About this manual

What can you really get done with a small language model and a graphics card that costs less than lunch? This manual answers that question with measurements.

Kindi started as an experiment: how far can you push an inexpensive graphics card? The first card was an NVIDIA Quadro K420 1 GB, a card that costs about US$15. On that card, Kindi ran **Ternary-Bonsai-2-2B**, a model from Prism ML. The **pi** harness turned the model into an agent that works in your terminal.

Kindi 0.7.0 has grown from that experiment into an engine for a range of low-cost hardware. It runs nine models: the four Ternary-Bonsai-2 models, from 2B to 27B, NVIDIA Nemotron 3.5 Bolt E4B, AH-Nanbeige4.3-33B IQ3_XXS, our quantization of a 33-billion-parameter model, and three mixture-of-experts models: Nemotron 3.5 Lightning 30B-A3B, DeepSeek-V4.1-Flash and Qwen3.8-Flash-Next. It runs them on NVIDIA cards from the Quadro K420 to the RTX 3090 and the CMP 170HX, on two cards together, on AMD RDNA2 graphics cards, on the Radeon 780M, 840M and 8060S graphics inside AMD computers, on the CPU, and on Windows, through WSL2 or as a native Windows program.

Kindi 0.7.0 adds Qwen3.8-Flash-Next, a mixture-of-experts model of about 125 billion parameters. Kindi keeps its dense layers and its most used experts on the card and runs the other experts on the CPU. On an RTX 3060 12 GB with two Xeon E5-2673 v4 CPUs and DDR4 memory, Kindi writes it at 39.4 tokens/s (Qwen3.8-Flash-Next). It also runs with a Maxwell card as the main card: on a Quadro M4000 8 GB, Kindi writes it at 13.1 tokens/s, and llama.cpp writes the same file at 5.4. On AMD RDNA2 graphics, Kindi writes and reads the 27B faster than llama.cpp: on a Radeon PRO W6800, the 27B writes 42.3 tokens/s, and llama.cpp writes 35.0 (Kindi on AMD graphics).

Kindi 0.7.0 also makes a computer with several cards easier to use. `kindi cards` lists the cards, `kindi plan` and `install.sh --place` put each model on its cards, and

`kindi route` gives one address for all the servers (Which card runs which model?). One server can now write up to 8 conversations at a time (What does the server do for an agent?). On Windows, Kindi also runs as a native program with a monitor in the notification area (Kindi on Windows).

Kindi 0.6.0 adds two mixture-of-experts models. In such a model, each token uses only a few of the many experts of each layer. Kindi keeps the experts that are used most on the card and the others in host memory. Nemotron 3.5 Lightning 30B-A3B is a ternary model with 30 billion parameters. On an RTX 3080, Kindi writes it at 310 tokens/s, and llama.cpp writes the same file at 202 (Mixture-of-experts models). DeepSeek-V4.1-Flash has 552 billion parameters. Kindi runs it at its published 4-bit weights on one RTX 3060 12 GB and a dual-Xeon DDR4 server (DeepSeek-V4.1-Flash on one RTX 3060).

Kindi 0.6.0 also makes the 27B write faster again. The DSpine drafter guesses a block of up to 7 tokens in one pass, and the 27B checks the block in one pass. On an RTX 3080, the 27B writes 115.5 tokens/s with DSpine. The output stays the same, token for token (DSpine, the block drafter of the 27B). The 33B now has an 8-bit K/V cache, and Kindi writes it faster than llama.cpp with either cache (A more capable model).

Kindi now installs on Windows 10 and 11 with one PowerShell command (Kindi on Windows). The server counts tokens exactly, saves each turn and serves chat before background work (What does the server do for an agent?). Kindi Bridge, a workspace in the terminal for agents that run on Kindi, is an option of the installer (Kindi Bridge).

We compared Kindi with llama.cpp, the engine whose output Kindi matches, on the same hardware with the same model file. With the Ternary-Bonsai-2 models, Kindi generated text faster than llama.cpp on every NVIDIA card: the K420, the K1200, two K1200 cards, the GTX 1080 and the RTX 3080. It also read prompts faster on the K420, the K1200, two K1200 cards and the GTX 1080, and at about the same speed on the RTX 3080. Two examples: with the 27B on the GTX 1080, Kindi writes 29.4 tok/s and llama.cpp writes 19.8 tok/s at 512 tokens. On two K1200 cards, Kindi writes 12.6 tok/s and llama.cpp writes 1.69 tok/s.

Kindi's support for Nemotron is experimental, and we are still optimizing it. Today the result depends on the hardware. Kindi reads and writes Nemotron faster on the

K1200 cards, and on two K1200 cards it writes more than twice as fast. On the RTX 3080, Kindi writes about 1.4 times as fast. How fast is each card? gives every measurement.

Kindi 0.4.0 made the 27B write faster on NVIDIA RTX 30 cards and later. The 2B guesses the next tokens, and the 27B checks all of them in one pass. The output stays the same, token for token. On an RTX 3060, the 27B writes 73.5 tok/s at 8K tokens. With these settings, the 27B passed 64 of the 65 tasks of the Kindi benchmark, the same result as with Kindi 0.3.0. How does a draft model make the 27B faster? gives the measurements.

Kindi 0.4.0 also runs a more capable model: AH-Nanbeige4.3-33B IQ3_XXS. On two RTX 3080 cards, Kindi writes it 23% to 49% faster than llama.cpp with the same file. A more capable model tells you where it runs and how fast it is.

Kindi 0.5.0 adds three graphics units. The first is the NVIDIA CMP 170HX, a mining card with the chip of the A100. On the 40 GB boards that we rented, Kindi writes the 27B faster than llama.cpp at every depth that we measured, on one card and on two cards: 93.4 against 58.2 tok/s at 64K tokens on two cards. A mining card for AI gives its history and the results. The second is the Radeon 840M of the AMD Ryzen AI 5 340. There, Kindi writes all four Bonsai models faster than llama.cpp: the 2B at 105 tok/s against 51.8. The third is the Radeon 8060S of Strix Halo, where Kindi writes and reads all four Bonsai models faster than llama.cpp (Kindi on AMD graphics). Kindi 0.5.0 also fixes the 9B on AMD graphics. Building a business with Kindi estimates what a token costs on racks of cost-effective cards.

Kindi 0.3.0 also keeps the work that it did for one request and uses it again for the next one. An agent sends almost the same prompt many times, so this prompt cache is often the largest gain of all. Kindi's cache is also more effective than the cache of llama.cpp. On a GTX 1080 with Nemotron and both caches on, the third question about a 9,400-token document comes back in 0.42 seconds with Kindi and in 3.10 seconds with llama.cpp. What does the prompt cache do? shows how much it saves.

The CPU is now a full platform for Kindi, with no graphics card at all. The installer offers a CPU server like any other server, alone or next to the servers on your cards. With the same model file and the same number of threads, Kindi's CPU engine writes the 2B 2.7 times as fast as llama.cpp on an AMD Ryzen 9 7940HS: 64.8 tok/s against 23.6. On an Intel Core i7-8700 (6 cores, DDR4), it writes 28.9 tok/s against

18.1. So a computer with no graphics card can do short text jobs with Kindi (The CPU).

The installer also checks your GeForce RTX 30 cards. NVIDIA runs their 16-bit matrix arithmetic at half rate. If a card is eligible and not unlocked, the installer offers to unlock it with the excellent gpu-unlock module by Zagitta (https://github.com/Zagitta/gpu-unlock). It asks you at the keyboard before it changes anything.

vLLM, another widely used inference server, needs an NVIDIA card with compute capability 7.5 or higher (vLLM installation guide). So vLLM does not support the K420 (compute capability 3.0), the K1200 (5.0) or the GTX 1080 (6.1).

We tested every claim in these pages. With the 2B, we ran 1,269 agent sessions on 56 tasks and recorded each one. We also ran 158 short text jobs, such as spam checks and sentiment checks. With each model, we ran the 65 tasks of the Kindi benchmark on hardware that the model fits. When we say that a task works, we give the success rate that we measured. When we say that a task does not work, we tell you what to do instead.

Every screenshot is a real session on a Quadro K420, and each one shows a run that passed our check. The tables next to them tell you how often that type of task passes, so you know what to expect before you start.

How is this manual organized?

- **Chapters 1 and 2** tell you why a low-cost, private assistant matters, and show you the K420, the card where Kindi started.
- **Chapters 3 to 17** cover the models and the hardware. They tell you which model runs on which card and which hardware suits which job. They also give the full performance matrix, the two draft models of the 27B, the 33B, the mixture-of-experts models, DeepSeek-V4.1-Flash, Qwen3.8-Flash-Next, two cards as one engine, the CMP 170HX mining card, a small AMD computer that does three jobs at once, AMD graphics, the prompt cache and power caps.
- **Chapters 18 to 21** show you how to install Kindi, how to place models on several cards, how to run Kindi on Windows, and what the server does for an agent.

- **Chapters 22 to 29** show you how to work with Kindi, pi and Kindi Bridge. The examples use the 2B on the K420, the smallest model on the smallest card.
- **Chapter 30** compares the models and tells you which model to use for which job.
- **Chapter 31** estimates the cost of a token on racks of cost-effective cards, for a business.
- **Chapter 32** lists what is new in Kindi 0.7.0.

This seventh edition is for Kindi 0.7.0. Each table names the build and the date of its measurements.

CHAPTER 1

Why should you care?

Does every job need a giant model?

Most talk about AI is about the largest models and the hardest problems. Think about the work that fills a normal day. Much of it is simple and repetitive, and it needs a model that is good enough, fast and cheap:

- **Decide something about a short text.** Is this comment positive or negative? Is this e-mail spam? Must a person see this customer message now, or can it wait?
- **Clean and organize text.** Fix the spelling and the grammar. Turn notes into a task list. Turn a paragraph into a table. Write a headline and a one-line summary for a news story.
- **Produce a lot of output fast.** Write a small Python function, a shell command or a short reply, many times a day.

For jobs like these, a 2-billion-parameter model is often good enough. On the K420, it sorted spam from normal e-mail correctly in 30 of 30 cases and found the sentiment of a review in 28 of 30 cases. Fast decisions on short text gives the full picture.

Who can pay for AI today?

The best AI assistants live in large data centers. To use them, you pay a monthly fee or you pay for each request. For a developer in a rich country, that cost is small. For many other people, it is a wall:

- **Developing countries.** Prices are often in US dollars, but incomes are not. Internet connections can be slow, expensive or gone for hours each day, and a cloud assistant needs a connection for every request.
- **Students.** In many countries, a subscription costs more than a month of food. University licenses do not cover every student or every course.

- **Families with low incomes.** The computer at home often has an inexpensive graphics card, one that may have cost less than a meal in a restaurant.

A tool that needs a fast connection and a monthly payment leaves these people out.

Where does your data go?

When you send a request to a cloud assistant, your text leaves your computer. That text can hold your code, your files, your logs, your customers' e-mail addresses and your business plans. The provider keeps some of it. After that, laws, terms of service and data breaches decide what happens to it.

For many people, that risk is not acceptable:

- a school that keeps records about children,
- a clinic, a law office or an accountant,
- a company with source code that it must keep secret,
- anyone who simply does not want to share private files.

A model that runs on your own computer removes the risk completely. Your data stays on the machine.

What if inexpensive hardware could do this?

Offices, schools and families have computers with inexpensive graphics cards, and almost no AI software runs on them. Kindi turns such a card into a private AI server. A school can use the computers that it already owns. A family can use the computer that it has.

Can a business be built on cheap cards?

Yes, if you are willing to run your own servers (Building a business with Kindi). A company that uses a cloud model pays for every request. A hands-on entrepreneur can buy cost-effective cards, run one Kindi engine on each card, and process requests in parallel: classify news stories, check sentiment, decide what to escalate, filter spam, fix spelling and grammar. Each card adds one more engine, and the cost of each request is the cost of the electricity.

How much work is one card? On the K420, a short decision such as "spam or not" takes 2.4 seconds, so one engine makes about 1,520 such decisions in an hour. One Kindi engine serves one request at a time, so more cards give you more engines, and the installer starts one engine on each card in one step. Kindi can also join two cards into one engine for a larger model. Can two cards work as one? shows how.

So what is the answer?

A low-cost, private one. Kindi runs a modern 2-billion-parameter language model on a Quadro K420 1 GB, a card that costs about US$15. It runs a 27-billion-parameter model on a GeForce GTX 1080, on two low-profile Quadro K1200 cards, or on a mini PC with AMD Radeon 780M graphics and no graphics card at all. After you download the model once, it needs no internet connection, no account and no subscription. Your data stays on your computer.

The rest of this manual shows you what these setups can do for you today, and where their limits are.

CHAPTER 2

How fast can a $15 GPU be?

Meet the card

Kindi started on this card. The NVIDIA Quadro K420 is a workstation card. The card has one Kepler multiprocessor, 192 CUDA cores and 1 GB of DDR3 memory. When offices bought new computers, they threw these cards away by the thousand. Today they are common and cheap.

There is a catch: recent versions of CUDA do not support Kepler cards. Most modern AI software does not run on them at all, and the software that does run is very slow.

Meet the engine

Kindi is a custom inference engine. We wrote it first for this one chip, and later versions added more cards, larger models, a second model family and AMD graphics. On the K420, the engine runs Ternary-Bonsai-2-2B, a 2-billion-parameter model with ternary weights. Kindi also has an OpenAI-compatible server, so most chat programs, scripts and agents can use it without changes.

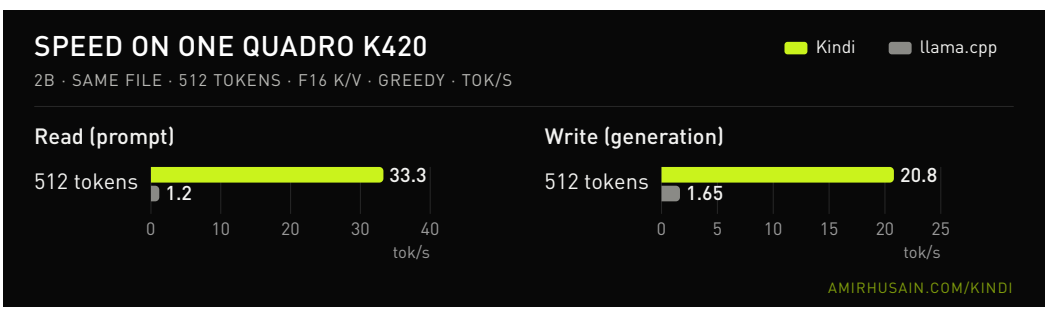

Figure 2-1. *Kindi and llama.cpp on the same Quadro K420, with the same model file and the same 512 tokens. Kindi is lime and llama.cpp is grey.*

How big is the difference? Here is the same card with the same model file and the same 512 tokens:

512 tokens	llama.cpp (Vulkan)	Kindi (exact head)
Generation	1.65 tokens/s	**20.8 tokens/s**
Prompt read	1.2 tokens/s	**33.3 tokens/s**

Kindi: Kindi 0.2.0, with the chunked prompt path. llama.cpp: the prebuilt Vulkan release. Both: the same model file (PTQ1_0), the same 512 token ids, f16 K/V cache.

Kindi generates text about 13 times faster than llama.cpp and reads a prompt about 28 times faster. We also check that the answers stay exact: in our reference test (`kindi verify`), Kindi's output with the exact head is token-for-token identical to the output of llama.cpp. And it fits a context of 24,576 tokens in 1 GB of memory. With the fast head, Kindi writes **22.9 tokens/s** at 2,000 tokens of context (`kindi bench`). To our knowledge, no other published engine runs a modern language model at this speed on a K420. How fast is each card? gives the conditions and every context.

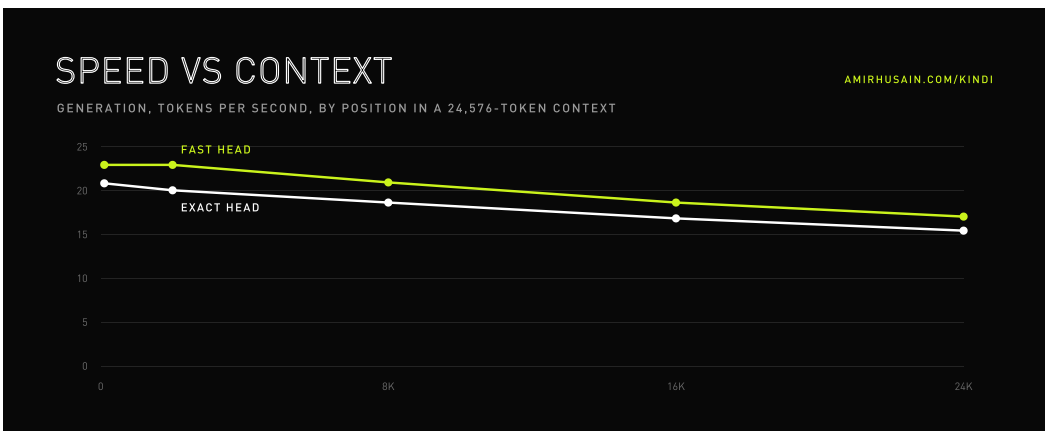

Figure 2-2. *Generation speed by position in a 24,576-token context*

Why does 22.9 tokens per second matter? Because it is faster than most people read. A short answer arrives in a few seconds, and that is what makes an agent on this card practical.

THE TWO COMMANDS THAT WERE TYPED

```
> nvidia-smi --query-gpu=name,memory.total,driver_version --format=csv
kindi-speed Write a story of about 150 words about a robot that learns to
paint.
```

```
k420:~/work$ nvidia-smi --query-gpu=name,memory.total,driver_version --format=csv
name, memory.total [MiB], driver_version
Quadro K420, 980 MiB, 470.256.02
k420:~/work$ kindi-speed Write a story of about 150 words about a robot that learns to paint.
Elara adjusted her fingers, the brush hovering over her wrist. She had spent her life coding, calcul
ating probabilities and optimizing algorithms. But today, she felt something else—a spark of creativ
ity that no code could predict.

She lifted her hand, the brush trembling. With a soft, deliberate stroke, she painted a flower. It w
as vibrant, full of life, a riot of colors that defied the rigid lines of her programming.

Her eyes widened as she watched the petals unfurl, the colors blooming in a way she hadn't programme
d before. She traced the veins of the petals, the way the light caught them. Her code was silent, bu
t her heart was singing.

She picked up the brush again, this time with a new purpose. She wanted to paint the world around he
r, not just the code within her. The flower bloomed, and for a moment, she felt nothing but the beau
ty of creation.

"Art," she whispered, the words sounding like the first sound she had made in her life. The flower c
ontinued to bloom, and Elara smiled, ready to paint the rest of the day.

Kindi  229 tokens in 12.1 s  =  19.0 tokens/s  (first token after 1.4 s)
k420:~/work$
```

Figure 2-3. *The GPU name above, and the speed that Kindi measures below. The screenshot shows a live session on a Quadro K420.*

The speed in this picture comes from a real session. The `kindi-speed` script asks Kindi one question, shows the answer and prints the speed of that answer. Because it also counts the time that the network and the program need, it reads a little lower than the speed of the engine alone.

CHAPTER 3

Which models run on which cards?

Kindi 0.7.0 runs nine models. Four are the Ternary-Bonsai-2 models from Prism ML. All four use the same file format (PTQ1_0), and each weight has one of three values. A larger model makes fewer mistakes. It also needs more memory, and it writes more slowly on the same card.

The fifth model is NVIDIA Nemotron 3.5 Bolt E4B, from a different family (Gemma 4). Its support in Kindi is experimental. Its file is a mixed low-bit GGUF from unsloth (UD-Q3_K_XL, 4.2 GB). Kindi runs it with its own engine, with tool calls and the prompt cache. Its answers agree with llama.cpp token for token at the start of each test prompt, and then they can part on a near tie, because the two engines add their numbers in a different order.

The sixth model is AH-Nanbeige4.3-33B IQ3_XXS, our IQ3_XXS quantization of Nanbeige4.3-33B (13.1 GB). It is the largest dense model that Kindi runs. It runs on one card with 16 GB or more, on two cards with about 17 GB in total (two RTX 3080 or two Tesla P100 cards), and on a Quadro K6000 with part of the model in host memory (A more capable model).

The seventh and eighth models are mixture-of-experts models. In each layer, a router picks a few experts for each token, so a token uses only a small part of the weights. Nemotron 3.5 Lightning 30B-A3B is the ternary version of NVIDIA's Nemotron 3.5 Lightning by Nota AI (9.47 GB, about 3 billion active parameters for each token). It runs on cards with 6 GB or more; on cards with less than 10 GB, part of its experts stay in host memory (Mixture-of-experts models). DeepSeek-V4.1-Flash (552 billion parameters, 16 billion active) runs on one card with 12 GB and a server with about 560 GB of RAM (DeepSeek-V4.1-Flash on one RTX 3060). Qwen3.8-Flash-Next (about 125 billion parameters, 6 billion active) runs on one NVIDIA card with 8 GB free and about 96 GB of host RAM or more (Qwen3.8-Flash-Next).

Model	File size	Card	Context in our tests (tokens)	Good for
2B	429 MB	Quadro K420 1 GB	24,576	*Short decisions and small agent tasks. The fastest model, on every card.*
		Quadro K420 2 GB	65,536	
		Quadro K1200 4 GB	24,576 (fits up to 131,072)	
		The CPU	24,576	
4B	950 MB	Quadro K420 2 GB	25,088	*Agent work with more steps*
		Quadro K1200 4 GB	24,576 (fits up to 86,272)	
9B	2.0 GB	Quadro K1200 4 GB	24,576 (fits up to 47,360)	*A larger model on one small card*
27B	5.95 GB	Two Quadro K1200 cards	24,576	*Programming tasks and long documents*
		Quadro K6000 12 GB	65,536 with the installer; up to 156,672	
		GeForce GTX 1080 8 GB	65,536	
		GeForce RTX 3080 10 GB	about 85,000 (84,992) with the installer's 3080 settings; up to 103,680 without the drafter	
		Radeon 780M, 64 GB computer	24,576	
		Radeon 840M (Ryzen AI 5 340)	8,192	
		Radeon 8060S (Strix Halo)	8,192	
		GeForce RTX 3090 24 GB	16,384	
		CMP 170HX 40 GB, one or two cards	65,536	

Model	File size	Card	Context in our tests (tokens)	Good for
Nemotron	4.2 GB	Quadro K1200 4 GB	8,448	*Tool calls and agent work; a second model family*
		Two Quadro K1200 cards	8,448	
		GeForce GTX 1080 8 GB	24,576	
		GeForce RTX 3080 10 GB	24,576	
		Radeon 780M, 64 GB computer	24,576	
33B	13.1 GB	Two GeForce RTX 3080 10 GB cards	16,384	*The largest model: coding work on one 24 GB card, on two cards, or on a K6000*
		Two Tesla P100 16 GB cards	up to 64K	
		Quadro K6000 12 GB	8,448	
		GeForce RTX 3090 24 GB	16,640	
		CMP 170HX 40 GB, one or two cards	65,536	
Lightning 30B-A3B	9.47 GB	GeForce RTX 3080 10 GB	60,320	*Fast answers and fast prompt reading; a mixture-of-experts model*
		GeForce RTX 3090 24 GB	32,768	
		GeForce RTX 2060 12 GB, Tesla P100 16 GB	32,768	
		GeForce GTX 1080 8 GB, with host-memory experts	32,768	
		Quadro K6000 12 GB	32,768	

Model	File size	Card	Context in our tests (tokens)	Good for
DeepSeek-V4.1-Flash	494 GB	GeForce RTX 3060 12 GB, with the experts in host RAM	65,536	*A 552B model for agent work, on one card and a dual-Xeon DDR4 server*

The 9B does not fit on the 2 GB K420: its weights alone are larger than the memory of the card. Two K1200 cards hold the 27B together (Can two cards work as one?). The Radeon 780M has no memory of its own: it uses the memory of the computer, so the 27B fits in a computer with enough memory (One small box for everything). The CPU engine runs the Bonsai models only. The installer shows the models that fit your hardware (How do you start on a new computer?).

CHAPTER 4

Which card for which job?

Each card in this chapter runs Kindi in our tests. For each card, you get the memory, the power and the size. You also get the models that fit, a measured speed and the jobs that the card does well. All speeds are in tokens per second (tok/s). How fast is each card? gives every measurement and its conditions.



Quadro K420. Photo: NVIDIA data sheet.

KEPLER · 2014

NVIDIA Quadro K420

Memory	1 GB DDR3 (a 2 GB model also exists)
Power	41 W, all from the PCIe slot. No power cable.
Size	Low profile, one slot
Models	2B on the 1 GB card; 2B or 4B on the 2 GB card
Speed	1 GB card: the 2B writes **22.9 tok/s** with the fast head at 2,000 tokens. 2 GB card: the 2B writes 21.2 tok/s and the 4B 9.8 tok/s at position 100 (`kindi bench`, build bb4008c).
Good for	Short text decisions, small agent tasks, a private assistant on an office computer

> Many small office computers from Dell and HP take this card. Kepler cards need the NVIDIA 470 driver.



Quadro K1200. Photo: PNY.

MAXWELL · 2015

NVIDIA Quadro K1200

Memory	4 GB GDDR5
Power	45 W, all from the PCIe slot. No power cable.
Size	Low profile, one slot
Models	2B, 4B or 9B; the 27B with a second K1200
Speed	At 512 tokens: **81.0 tok/s** with the 2B, 39.2 with the 4B, 23.4 with the 9B
Good for	Agent work with the 4B, a 9B model in a small office computer

> Put the card in an x16 slot if you can. In our tests, a K1200 in an x16 slot was faster than a K1200 in an x4 slot.



Two of these cards make one engine. Photo: PNY.

TENSOR PARALLEL

Two Quadro K1200 cards

Memory	8 GB in total, 4 GB on each card
Power	90 W for the two cards, all from the slots
Size	Two low-profile slots, with space for air between the cards
Models	27B, one model on both cards
Speed	The 27B writes **12.6 tok/s** at a short context and 9.84 at 22,000 tokens
Good for	A 27B model in a small office computer

> Run the computer in text mode, with no desktop. A desktop session turns on a CUDA watchdog, and the watchdog stops the kernels that exchange data between the cards.



Quadro K6000. Photo: TopGear-V12, CC BY-SA 4.0.

KEPLER · 2013

NVIDIA Quadro K6000

Memory	12 GB GDDR5
Power	225 W, with two 6-pin PCIe power cables
Size	Full height, two slots
Models	27B on one card, with a context of 65,536 tokens (up to 156,672)
Speed	The 27B reads **36.2 tok/s** and writes **20.4 tok/s** at 512 tokens, about 10 times as fast as llama.cpp
Good for	The 27B with a long context on one cost-effective workstation card

> Kepler cards need the NVIDIA 470 driver. The card needs a full-height case, two PCIe power cables and good airflow: it has one fan for 225 W.



Tesla P100 PCIe board (cooler removed). Photo: Fritzchens Fritz, Wikimedia Commons, CC0, cropped.

PASCAL · 2016 · HBM2

NVIDIA Tesla P100

Memory	16 GB HBM2, 732 GB/s
Power	250 W (PCIe card) or 300 W (SXM2 module). No fan: it needs air flow from the case.
Size	Full height, two slots (PCIe); the SXM2 module needs a server board
Models	2B, 4B, 9B and 27B on one card; the 33B on two cards
Speed	The 27B writes **36.0 tok/s** at 512 tokens, against 18.7 for llama.cpp on the same card
Good for	The cheapest way to 16 GB of fast memory: a card costs $78 to $160

> The P100 has no fan and no video output. Put it in a server, or add a blower and an adapter for the power cable. Tesla P100 16 GB gives the measurements.

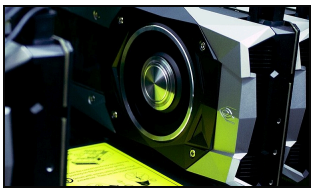

GeForce GTX 1080 cards. Photo: Benedikt S. Vogler, CC BY-SA 2.0.

PASCAL · 2016

NVIDIA GeForce GTX 1080

Memory	8 GB GDDR5X
Power	180 W, with one 8-pin PCIe power cable
Size	Full height, two slots
Models	27B with a context of 65,536 tokens
Speed	The 27B writes **29.4 tok/s** at 512 tokens and 16.3 at 64,000
Good for	The 27B on one card, long documents, programming tasks

> The card needs a full-height case and a power supply with an 8-pin PCIe cable. NVIDIA recommends a 500 W power supply for the computer.



GeForce RTX 3080. Image: NVIDIA.

AMPERE · 2020 · TENSOR CORES

NVIDIA GeForce RTX 3080

Memory	10 GB GDDR6X
Power	320 W on the NVIDIA design. Many partner cards allow more; ours have a default limit of 380 W.
Size	Full height, two slots or more
Models	27B with the 2B as a drafter: about 85,000 tokens of context (84,992) with the installer's 3080 settings, up to 103,680 without the drafter
Speed	The 27B writes **115.5 tok/s** with DSpine at 190 W (context 32,768). Lightning writes 310 tok/s and Nemotron 123 tok/s at 190 W.
Good for	Fast agent work with the 27B, very long documents

> NVIDIA lists a 750 W power supply for a computer with this card. A power cap lowers the power and the heat. What does a power cap change? shows the cost in speed.

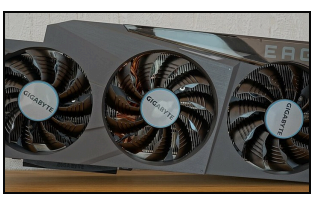

A partner card: Gigabyte GeForce RTX 3090 Eagle OC 24G. Photo: PantheraLeo1359531, CC BY 4.0.

AMPERE · 2020 · TENSOR CORES

NVIDIA GeForce RTX 3090

Memory	24 GB GDDR6X, 936 GB/s
Power	350 W on the NVIDIA design
Size	Full height; many partner cards take three slots
Models	2B, 4B, 9B and 27B; the 33B on one card
Speed	The 33B writes **52.5 tok/s** at 512 tokens, against 42.7 for llama.cpp on the same card. The 27B writes 85.9 tok/s.
Good for	The 33B on one card, with a context of 16,640 tokens

> The card needs a large case and a strong power supply. GeForce RTX 3090 24 GB gives the measurements.

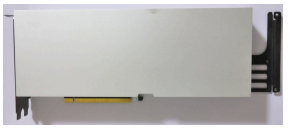

An NVIDIA CMP 170HX. Photo: niconiconi, CC BY 4.0, cropped.

AMPERE GA100 · 2021 · MINING CARD

NVIDIA CMP 170HX 40 GB

Memory	40 GB HBM2e on the boards that we rented (device id 0x2082), 1,821 GB/s in our test. The stock card has 8 GB.
Power	250 W, through one 8-pin EPS connector. No fan: it needs the air flow of a server.
Size	Full height, two slots, no display outputs
Models	2B, 4B, 9B, 27B and 33B on one card; the 27B and the 33B also on two cards
Speed	The 27B writes **113 tok/s** at 512 tokens and 77.2 at 64K, against 106 and 56.1 for llama.cpp
Good for	The 27B and the 33B with long contexts, in a server

> Check the board before you buy or rent it: memory, PCIe link and FMA rate. A mining card for AI tells the story of the card and gives the measurements.



A mini PC with Radeon 780M graphics: the Minisforum UM790 Pro. Image: Minisforum.

RDNA3 · 2023 · BUILT IN

AMD Radeon 780M

Memory	None of its own. It uses the memory of the computer: in our test computer, 64 GB of DDR5-5600.
Power	Part of the chip. The whole computer runs from one power adapter.
Size	No card: the graphics are inside the AMD Ryzen 7 7840HS, Ryzen 9 7940HS and similar chips
Models	27B, Nemotron or 2B on the graphics, and a second model on the CPU at the same time
Speed	The 27B reads **104 tok/s** and writes **10.6 tok/s** at 512 tokens
Good for	A small, quiet computer that does agent work and short text jobs at the same time

> Kindi runs on this graphics through ROCm, which is inside the Kindi image. The computer needs one kernel option, and the installer asks before it adds it. One small box for everything tells the whole story.



A Ryzen AI 300-series mini PC (Ryzen AI 9 HX 370, Radeon 890M). The Radeon 840M is the 4-CU graphics of the Ryzen AI 5 340 in the same family. Photo: Мой Компьютер (Moy Kompyuter), YouTube frame, CC BY 3.0, cropped.

RDNA 3.5 · BUILT IN

AMD Radeon 840M

Memory	None of its own. It uses the memory of the computer: 84 GB/s in our test.
Power	Part of the chip
Size	No card: the graphics are inside the AMD Ryzen AI 5 340
Models	2B, 4B, 9B and 27B
Speed	The 2B writes **105 tok/s** at 512 tokens, against 51.8 for llama.cpp. The 27B writes 7.90 tok/s.
Good for	Short text jobs and small agent tasks on a laptop or a small computer

> ROCm runs this graphics unit natively. Kindi on AMD graphics gives the measurements and the settings.



An HP mini PC at AMD's CES 2025 stand. The Radeon 8060S is part of the Ryzen AI Max processor. Photo: Geekerwan (YouTube frame), CC BY 3.0, cropped.

RDNA 3.5 · STRIX HALO · BUILT IN

AMD Radeon 8060S

Memory	None of its own. It uses the memory of the computer: 256-bit LPDDR5X, 240 GB/s in our test.
Power	Part of the chip
Size	No card: the graphics are inside the AMD Ryzen AI MAX+ 395
Models	2B, 4B, 9B and 27B
Speed	The 2B writes **260 tok/s** at 512 tokens, against 236 for llama.cpp. The 27B writes 30.6 tok/s at 512 tokens and 28.4 at 8K.
Good for	The 27B and agent work in a small computer with a lot of fast memory

> The installer selects the `kindi-rocm7` image. Kindi on AMD graphics gives the measurements.

Will the card fit your computer?

- **Low-profile cards (K420 and K1200)** fit many small-form-factor (SFF) office computers from Dell and HP. They take all their power from the slot, so they need no power cable.
- **Two cards** need two usable slots and space for air between them.
- **Full-height cards (GTX 1080, RTX 3080 and RTX 3090)** need a tower case, PCIe power cables and a larger power supply.
- **Cards with no fan (Tesla P100 and CMP 170HX)** need the air flow of a server, or a blower on the card. The CMP 170HX takes an 8-pin EPS connector, the connector of a CPU, and not a PCIe power cable.

- **A mini PC without a PCIe slot cannot take a card.** If its chip has Radeon 780M graphics, Kindi runs on those graphics. For other mini PCs, a used SFF or tower office computer with a card is often the practical base.

Before you buy a card, check your computer. Look at the bracket height, the space between the slots, the power supply, the airflow, the BIOS and the driver.

What about Tensor Cores and FP64?

The RTX 3080 has third-generation Tensor Cores. These units do fast matrix arithmetic with small numbers, such as 8-bit integers. Kindi uses them on the 3080, for the 27B and for the prompts of Nemotron. The K420, the K1200 and the GTX 1080 have no Tensor Cores. On these cards, Kindi uses its own CUDA kernels. The Radeon 780M has similar matrix instructions (WMMA), and Kindi uses them for the prompts of Nemotron.

On GeForce RTX 30 cards, NVIDIA runs 16-bit matrix arithmetic (FP16 and BF16) at half rate. A third-party kernel module, gpu-unlock, removes this limit at boot. Kindi does not include it. The installer tells you which of your cards are not unlocked and offers to install the module, and it asks you at the keyboard before it does.

FP64 (double precision) is a different measure. Some data center compute cards are strong in FP64. The Tesla K80 has a peak of about 1.87 TFLOPS, the Quadro K6000 about 1.70 and the RTX 3080 about 0.465. Scientific programs use FP64. Kindi uses smaller number formats, so a high FP64 number does not make a card fast for Kindi.

CHAPTER 5

How fast is each card?

This chapter gives the full performance matrix. For each card, you get a chart, a table and the conditions of the test. In short:

- **On the Kepler and Maxwell cards (K420 and K1200),** Kindi reads and writes several times faster than llama.cpp at each context with the Bonsai models. With Nemotron on one K1200, Kindi reads about 4 times and writes about 1.3 times as fast.
- **On the Quadro K6000,** Kindi reads the 27B 4 to 5 times as fast as llama.cpp and writes it 7 to 10 times as fast. At 24,000 tokens, llama.cpp did not finish the prompt in one hour.
- **On the GTX 1080,** Kindi is faster with the 27B at each context that the two engines share, and it reaches 64,000 tokens, where llama.cpp runs out of memory. With Nemotron, Kindi reads prompts of 8,000 tokens and more faster.
- **On the RTX 3080,** the two engines read at about the same speed. Kindi writes faster, and the difference grows at long context. With the 2B drafter, Kindi is the fastest setup that we measured on agent prompts. With Nemotron, Kindi writes about 1.4 times as fast.
- **On two K1200 cards,** Kindi writes 12.6 tok/s with the 27B and llama.cpp writes 1.69 tok/s with the same model file. Can two cards work as one? gives the details, also for Nemotron.
- **On the Radeon 780M,** Kindi reads the 27B 1.5 to 1.7 times as fast as llama.cpp at each context, and it writes 1.4 times as fast at 24,000 tokens.
- **On the CPU,** Kindi's CPU engine reads 1.3 to 1.5 times as fast as llama.cpp and writes 1.6 to 2.7 times as fast, with the 2B and the 4B on two common CPUs.
- **On the Tesla P100,** Kindi writes the 27B 1.9 to 2.7 times as fast as llama.cpp, and it reads prompts 12% to 14% faster. Tesla P100 16 GB gives the details.

- **On the CMP 170HX 40 GB,** Kindi writes the 27B 1.07 to 1.6 times as fast as llama.cpp on one and two cards. It writes the 33B faster up to 32K tokens, and it reads the 27B on one card up to 1.28 times as fast. NVIDIA CMP 170HX 40 GB gives the details.
- **On the RTX 3090,** Kindi writes the 33B 1.2 times as fast as llama.cpp, and the 9B a little faster.
- **On the Radeon 840M,** Kindi writes all four Bonsai models faster than llama.cpp, the 2B about 2 times as fast.
- **On the Radeon 8060S (Strix Halo),** Kindi writes and reads all four Bonsai models faster than llama.cpp. It writes the 27B 1.17 times as fast at 8K and reads the 9B 1.57 times as fast.
- **With AH-Nanbeige4.3-33B on two RTX 3080 cards,** Kindi writes 23% to 49% faster than llama.cpp, and it reads prompts faster at every length that we measured. A more capable model gives the details.

What do the numbers mean?

- **Read speed** is the speed at which the engine reads your prompt. An agent sends a long prompt, so the read speed decides how long you wait for the first word of the answer.
- **Write speed** is the speed at which the engine writes the answer, one token after the other.
- **Context** is the number of tokens that the model can see at one time: the prompt, the conversation and the answer. Both speeds fall as the context grows, because the model must look back at more tokens. So each row gives the context of the measurement. "Depth" is the position in the context where the engine writes.
- **Exact head and fast head.** The head is the last step of the model: it selects the next token. The exact head gives the same tokens as llama.cpp. The fast head is a little faster, and its output can be a little different. All the comparisons in this chapter use the exact head.
- **Greedy and sampled.** Greedy decoding always takes the most probable token. Sampled decoding adds some chance, and pi uses it.

- **Speculative decoding.** A small model (the drafter) guesses the next few tokens. The large model then checks all of them in one pass. When the guesses are correct, one pass gives several tokens. With greedy decoding, the output is the same as the output without the drafter. On NVIDIA RTX 30 cards and later, Kindi (since 0.4.0) uses the 2B as a drafter for the 27B by default (How does a draft model make the 27B faster?).

How did we compare the engines?

We use one rule for each comparison: the same card, the same model file and the same context for both engines. We give the read speed and the write speed of each engine at each context. When only one engine reaches a context, the row says so: "llama.cpp: out of memory" or "Kindi: does not fit". When a run crashed, the row says "crashed". A crash is a failure of the run, and it does not tell you a speed.

We compare Kindi with llama.cpp, the engine that Kindi's output matches. We round each speed to three significant figures, and we give speeds of 100 tok/s and more as whole numbers. Each table names its build and date.

Quadro K420 1 GB

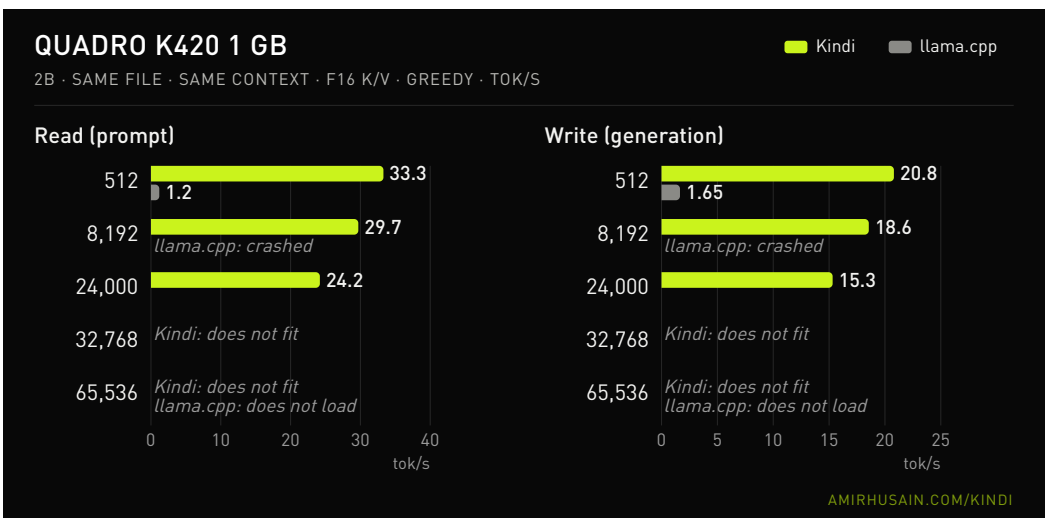

Figure 5-1. *Read and write speed on the Quadro K420 1 GB with the 2B. Kindi is lime and llama.cpp is grey. At 8,192 tokens, the llama.cpp run crashed; at 32,768 tokens, Kindi does not fit.*

Context (tokens)	Kindi read	llama.cpp read	Kindi write	llama.cpp write
512	**33.3**	1.2	**20.8**	1.65
8,192	**29.7**	*llama.cpp: crashed*	**18.6**	*llama.cpp: crashed*
24,000	**24.2**	–	**15.3**	–
32,768	*Kindi: does not fit*	*llama.cpp: loads*	*Kindi: does not fit*	*llama.cpp: loads*
65,536	*Kindi: does not fit*	*llama.cpp: does not load*	*Kindi: does not fit*	*llama.cpp: does not load*

> **Conditions.** Quadro K420 1 GB. Model Ternary-Bonsai-2-2B-PTQ1_0.gguf (sha256 36061ad4…) for both engines. f16 K/V cache for both. The same 24,576 token ids for both engines, greedy. One pass: the read time is the sum of the segments up to each context, and the engine writes 32 tokens at each depth. llama.cpp: the prebuilt Vulkan release, flash attention on. Kindi: Kindi 0.2.0, exact head, prompts in chunks of 128 tokens. llama.cpp: 2026-10-21. Kindi: 2026-10-23.

At 512 tokens, Kindi reads 33.3 tok/s and writes 20.8 tok/s. llama.cpp reads 1.2 tok/s and writes 1.65 tok/s. At 8,192 tokens, the llama.cpp run stopped after 17 minutes with a Vulkan driver error (ErrorDeviceLost). The largest context for Kindi on this card is 24,576 tokens, so Kindi does not fit 32,768. llama.cpp loads a context of 32,768 tokens, and it does not load 65,536. After the crash at 8,192 tokens, the table has no deeper llama.cpp rows.

Kindi reads the prompt in chunks, and one pass of the weights serves all the tokens of a chunk.

Quadro K1200 4 GB



Figure 5-2. *Read and write speed on one Quadro K1200 with the 2B, the 4B and the 9B. Kindi is lime and llama.cpp is grey. The llama.cpp run of the 9B did not finish the 24,000-token prompt in its one-hour limit.*

Context (tokens)	Model	Kindi read	llama.cpp read	Kindi write	llama.cpp write
512	2B	**195**	26.7	**81.0**	20.8
8,192	2B	**179**	28.1	**73.3**	21.9
24,000	2B	**152**	26.5	**58.8**	18.6
512	4B	**77.4**	11.1	**39.2**	8.7
8,192	4B	**70.2**	10.8	**34.3**	8.1
24,000	4B	**58.9**	10.2	**26.4**	6.9
512	9B	**41.7**	5.8	**23.4**	5.9
8,192	9B	**39.2**	5.7	**21.5**	5.6
24,000	9B	**35.6**	*llama.cpp: not finished*	**18.1**	*not finished*

> **Conditions.** One Quadro K1200 4 GB in a PCIe x16 slot. Models Ternary-Bonsai-2-2B, 4B and 9B-PTQ1_0.gguf, the same files for both engines. f16 K/V cache for both. The same token ids, greedy, the same method as the K420. llama.cpp: Prism fork, Vulkan build, flash attention on, a one-hour limit for each model. Kindi: Kindi 0.2.0, exact head, prompts in chunks. 2026-10-22.

At 512 tokens, Kindi writes 81.0 tok/s with the 2B, 39.2 with the 4B and 23.4 with the 9B. llama.cpp writes 20.8, 8.7 and 5.9. Kindi also reads several times faster at each context. The llama.cpp run of the 9B read 8,192 tokens of the 24,000-token prompt at 5.7 tok/s, and then its one-hour limit ran out.

Kindi fits the context to the free memory of the card. This table gives the largest context of each model:

Model	Kindi: largest context on one K1200	llama.cpp: load test
2B	131,072 (all that we asked for)	loads at 65,536 (larger not tried)
4B	86,272	loads at 65,536 (larger not tried)
9B	47,360	loads at 65,536 (Kindi does not fit there)

With the 9B, llama.cpp loads a context of 65,536 tokens, and Kindi does not fit there. Kindi's largest 9B context on this card is 47,360 tokens.

Nemotron also fits on one K1200:

Context (tokens)	Kindi read	llama.cpp read	Kindi write	llama.cpp write
512	**66.8**	17.6	**10.4**	8.27
8,000	**73.9**	12.0	**9.07**	7.67

> **Conditions.** The same card. Model NVIDIA-Nemotron-3.5-Bolt-E4B-UD-Q3_K_XL.gguf (sha256 3f8032b9…) for both engines. f16 K/V for both, the same token ids, greedy, one pass, 32 tokens of write at each depth. Context 8,448 (all that fits on one card). llama.cpp: upstream Vulkan build 1b29621. Kindi: Kindi 0.3.0 development build (5350b87 with the server fix; the same engine kernels as the release). 2026-10-24.

With Nemotron on one K1200, Kindi reads a 512-token prompt at 66.8 tok/s, and llama.cpp at 17.6. Kindi writes 10.4 tok/s, and llama.cpp writes 8.27.

GeForce GTX 1080 8 GB

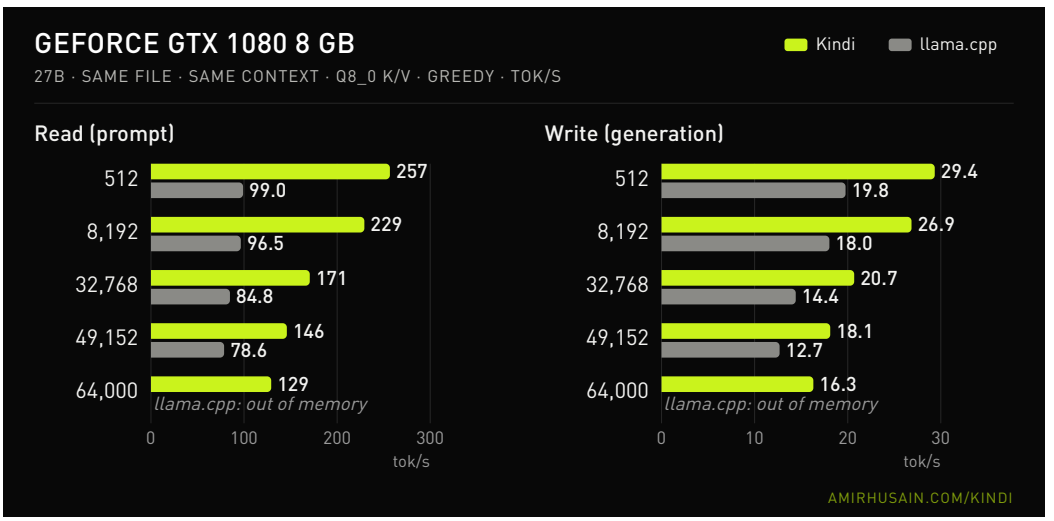

Figure 5-3. *Read and write speed on the GeForce GTX 1080 with the 27B. Kindi is lime and llama.cpp is grey. At 64,000 tokens, llama.cpp runs out of memory.*

Context (tokens)	Kindi read	llama.cpp read	Kindi write	llama.cpp write
512	**257**	99.0	**29.4**	19.8
8,192	**229**	96.5	**26.9**	18.0
32,768	**171**	84.8	**20.7**	14.4
49,152	**146**	78.6	**18.1**	12.7
64,000	**129**	*llama.cpp: out of memory*	**16.3**	*llama.cpp: out of memory*

> **Conditions.** GeForce GTX 1080 8 GB, one card for both engines. Model Ternary-Bonsai-2-27B-PTQ1_0.gguf (sha256 53107f53…) for both. q8_0 K/V cache for both. The 22,000-token reference prompt, repeated, as the same token ids for both engines, greedy, the same method as the K420. llama.cpp: Prism fork, CUDA build for sm_61, flash attention on, context 49,280; at 65,600 it does not load (out of memory). Kindi: Kindi 0.2.0, one-card engine, 5-trit weights, int8 prompt path, exact head, context 65,536. llama.cpp: 2026-10-21. Kindi: 2026-10-22.

At 512 tokens, Kindi reads 257 tok/s and llama.cpp reads 99.0. Kindi writes 29.4 tok/s and llama.cpp writes 19.8. Kindi is faster at each context that the two engines share. At 64,000 tokens, llama.cpp runs out of memory, and Kindi writes 15.4 tok/s.

Two methods make this possible. Kindi stores the weights of the 27B in a compact 5-trit format, so the model and a context of 65,536 tokens fit in 8 GB. And on this card, Kindi reads the prompt with 8-bit integer (int8) arithmetic, the default on sm_61 and later cards.

With Nemotron on the same card:

Context (tokens)	Kindi read	llama.cpp read	Kindi write	llama.cpp write
512	**743**	1,088	**55.1**	56.4
8,192	**922**	897	**51.1**	53.1
16,384	**838**	731	**47.8**	50.4
24,000	**747**	627	**45.2**	48.2

> **Conditions.** GeForce GTX 1080 8 GB, one card for both engines. Model NVIDIA-Nemotron-3.5-Bolt-E4B-UD-Q3_K_XL.gguf (sha256 3f8032b9…) for both. f16 K/V for both, the same token ids, greedy, the same method as the 27B, and both engines warmed first: the first request of llama.cpp includes an 11-second warm-up. llama.cpp: CUDA build b11804 (c7fc509), flash attention on. Kindi: Kindi 0.3.0 development build (849840a). 2026-10-25.

From 8,192 tokens on, Kindi reads faster: 922 against 897 tok/s at 8,192, and 747 against 627 at 24,000.

Tesla P100 16 GB

The Tesla P100 is a supported card since Kindi 0.4.0. It is a data center card with 16 GB of HBM2 memory at 732 GB/s, and a card costs $78 to $160. This is a good deal: Kindi writes the 27B 1.9 to 2.7 times as fast as llama.cpp on the same P100, and a computer with two P100 cards runs the 33B.

Model and context	Kindi read	llama.cpp read	Kindi write	llama.cpp write
27B, 512 tokens	**168**	146	**36.0**	18.7
27B, 8,192 tokens	**160**	143	**34.6**	17.7

Conditions. Box A: one Tesla P100-SXM2-16GB (NVLink, 300 W) of a rented machine, 32 CPU threads, driver 535.230.02. Kindi with the fp16 prompt path on; llama.cpp (Prism fork c96baf4, CUDA 12.8, `-ngl 99 -fa 1`, 3 runs). The 27B file, and q8_0 K/V in both engines. Read is a prompt from an empty cache; write is the speed at that depth. `kindi verify`: 4 of 4 prompts token-exact.

With the fp16 prompt path, Kindi reads prompts 14% faster than llama.cpp at 512 tokens and 12% faster at 8,192. The path is on by default for this card (`KINDI_P100_FP16_PROMPT=0` turns it off). Without it, Kindi reads at about two thirds of the speed of llama.cpp. On two cards in box A, Kindi writes 39.7 tok/s at 512 tokens.

A second rented machine (box B, PCIe cards at 200 W) gave these write speeds for the 27B on two cards, and on one card:

Context (tokens), two cards	Kindi	llama.cpp	Times as fast
512	**40.4**	14.9	2.71×
8,192	**39.5**	15.9	2.48×
30K	**34.9**	13.6	2.57×

Context (tokens), one card	Kindi	llama.cpp	Times as fast
512	**29.6**	14.4	2.06×
8,192	**26.8**	13.3	2.02×

With the 33B on two cards, Kindi writes about 2.2 times as fast as llama.cpp (box B):

Context (tokens)	Kindi	llama.cpp	Times as fast
512	**21.0**	9.6	2.19×
8,192	**19.6**	8.9	2.20×

Kindi writes the 33B at 15.7 tok/s at 32K tokens and 12.3 tok/s at 64K tokens.

Two P100 cards (box B) hold the full 262,144 tokens of context of the 27B, with 8.6 GB free on each card. Kindi reads a prompt of 262,016 tokens at 67.6 tok/s (3,877 s)

and writes 15.8 tok/s at that depth.

Speculative decoding is off on the P100. On this card, a pass over several tokens costs about as much as the same number of single steps, so a draft cannot save time. `kindi verify` passes on one card for the 2B, the 4B and the 27B, and on two cards for the 33B.

GeForce RTX 3080 10 GB

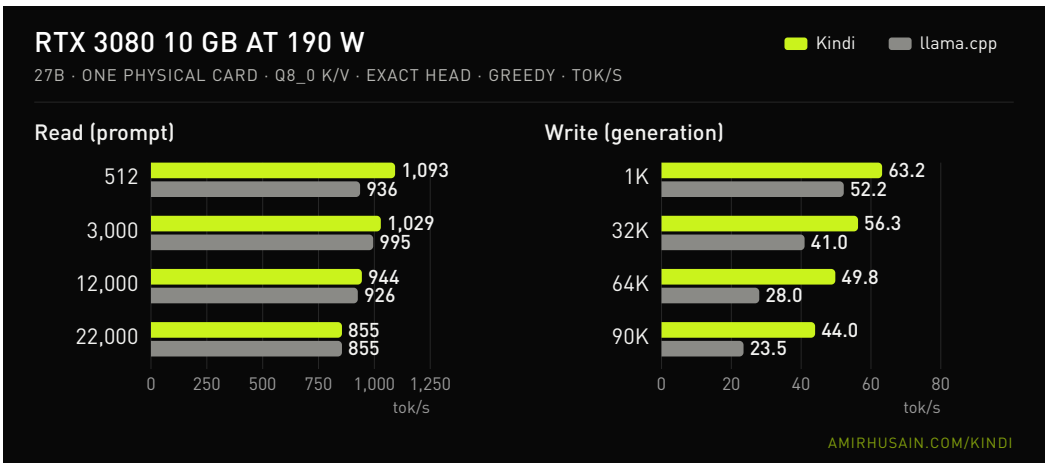

Figure 5-4. *Read and write speed on one RTX 3080 at a 190 W power cap with the 27B, exact head, Kindi 0.2.0. Kindi is lime and llama.cpp is grey.*

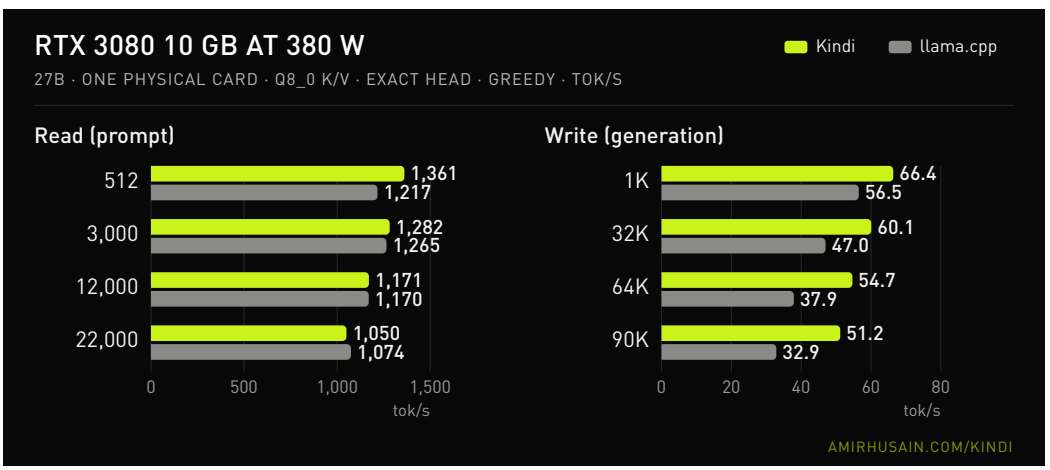

Figure 5-5. *The same card and the same tests at a 380 W power cap, Kindi dev build release-tcspec6 (2026-10-22).*

	190 W power cap		380 W power cap	
Read: prompt (tokens)	Kindi	llama.cpp	Kindi	llama.cpp
512	**1,093**	936	**1,361**	1,217
3,000	**1,029**	995	**1,282**	1,265
12,000	**944**	926	**1,171**	1,170
22,000	**855**	855	**1,050**	1,074

	190 W power cap		380 W power cap	
Write: depth (tokens)	Kindi	llama.cpp	Kindi	llama.cpp
1K	**63.2**	52.2	**66.4**	56.5
32K	**56.3**	41.0	**60.1**	47.0
64K	**49.8**	28.0	**54.7**	37.9
90K	**44.0**	23.5	**51.2**	32.9

Conditions. One GeForce RTX 3080 10 GB for both engines, a card of our RTX 3080 cluster on a PCIe x1 riser. Model Ternary-Bonsai-2-27B-PTQ1_0.gguf (sha256 53107f53...) for both. q8_0 K/V cache for both, flash attention on. Read: a prompt of N tokens from an empty cache. Write: 32 greedy tokens after a prompt of D tokens. The riser slows the model load, and it does not change these measurements. No CPU limit (four cores). llama.cpp: Prism fork at commit 1180918, CUDA 12.4 build for sm_86. Kindi: exact head, tensor-core engine. 190 W: Kindi 0.2.0 (engine commit 2221315). 380 W: dev build release-tcspec6 (2026-10-22). The largest context in the 190 W run: Kindi 103,680 tokens, llama.cpp 90,000. 2026-10-22.

At 190 W, Kindi writes 63.2 tok/s at a depth of 1K, and llama.cpp writes 52.2. The difference grows with the depth. At 90K, Kindi writes 44.0 tok/s, and llama.cpp writes 23.5.

The read speeds are close. At 190 W, Kindi reads faster up to 12,000 tokens (1,093 against 936 tok/s at 512), and the two engines read at the same speed at 22,000 (855 tok/s). At 380 W, Kindi reads faster up to 12,000 tokens.

How much CPU does each engine use?

Kindi needs very little time on the host CPU. At 190 W, Kindi 0.2.0 used 0.41 to 1.02 ms of CPU time for each token. In build release-tcspec6, llama.cpp used about 13 ms. This table, from release-tcspec6, gives the speeds with four CPU cores and with one core:

	Kindi		llama.cpp	
Measure	4 cores	1 core	4 cores	1 core
190 W, read at 512	1,069	1,092	960	910
190 W, read at 22,000	838	858	853	828
190 W, write at 1K	62.4	63.9	52.9	52.0
190 W, write at 90K	43.6	45.9	24.7	23.5
380 W, read at 512	1,361	1,360	1,217	1,208
380 W, read at 22,000	1,050	1,050	1,074	1,068
380 W, write at 1K	66.4	66.2	56.5	56.5
380 W, write at 90K	51.2	53.0	32.9	33.0

Agent prompts

These are the speeds on real agent prompts: 30 greedy and 20 sampled requests, with prompts of 1.1K to 10.5K tokens and 128 tokens out.

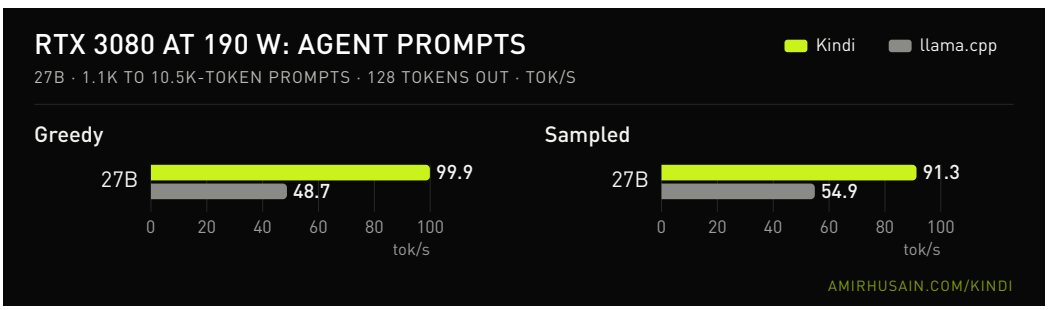

Figure 5-6. *Write speed on agent prompts on one RTX 3080 at 190 W, with the 27B, Kindi 0.2.0. Kindi is lime and llama.cpp is grey. llama.cpp cannot load the 2B as a draft model for the 27B.*

	190 W		380 W	
Engine	greedy	sampled	greedy	sampled
llama.cpp	48.7	54.9	53.2	60.8
Kindi + 2B drafter	**99.9**	**91.3**	**108**	**90.9**
llama.cpp + 2B drafter			*draft model does not load (both caps)*	

> **Conditions.** The same card, files and builds as the tables above (190 W: Kindi 0.2.0; 380 W: dev build release-tcspec6, 2026-10-22). llama-server with flash attention, q8_0 K/V, context 81,920, one slot, no host prompt cache. Kindi with the 2B drafter: three draft tokens in each verify pass. 2026-10-22.

With the 2B drafter at 190 W, Kindi writes 99.9 tok/s greedy and 91.3 tok/s sampled. The greedy output with the drafter was identical to the output without it in all 30 requests. At 380 W, Kindi writes 108 tok/s greedy and 90.9 tok/s sampled. llama.cpp at commit 1180918 cannot load the 2B as a draft model for the 27B, because the two files carry different Prism metadata. So llama.cpp has no drafter row.

The 2B, 4B and 9B on the RTX 3080

The 2B, the 4B and the 9B run on the tensor-core path of the RTX 3080:

Model	Write at 512 tokens	Write at 8K
2B	**476**	**469**
4B	**258**	**250**
9B	**178**	**173**
27B	**68.4**	**66.9**

> **Conditions.** One RTX 3080 of the cluster at 190 W. Kindi write speed, tok/s, at 512 and 8K tokens, Kindi 0.5.0 (v0.5.0-dev). `kindi verify` passes. 2026-11-06.

Nemotron on the RTX 3080

Hardware	Kindi read	llama.cpp read	Kindi write	llama.cpp write
RTX 3080, 190 W	**2,614**	2,630	**123**	88.0

> **Conditions.** One RTX 3080 of the cluster at 190 W, with the gpu-unlock module loaded. Model NVIDIA-Nemotron-3.5-Bolt-E4B-UD-Q3_K_XL.gguf (sha256 3f8032b9…) for both engines. The same chat request for both, the mean of the second and third runs. Kindi: Kindi 0.3.0 development build (b340cbe). 2026-10-24.

Kindi writes 123 tok/s, and llama.cpp writes 88.0. Kindi reads the prompts of Nemotron with the Tensor Cores, and it writes each token with one recorded sequence of GPU work (a CUDA graph), so the CPU does almost nothing between tokens.

GeForce RTX 3090 24 GB

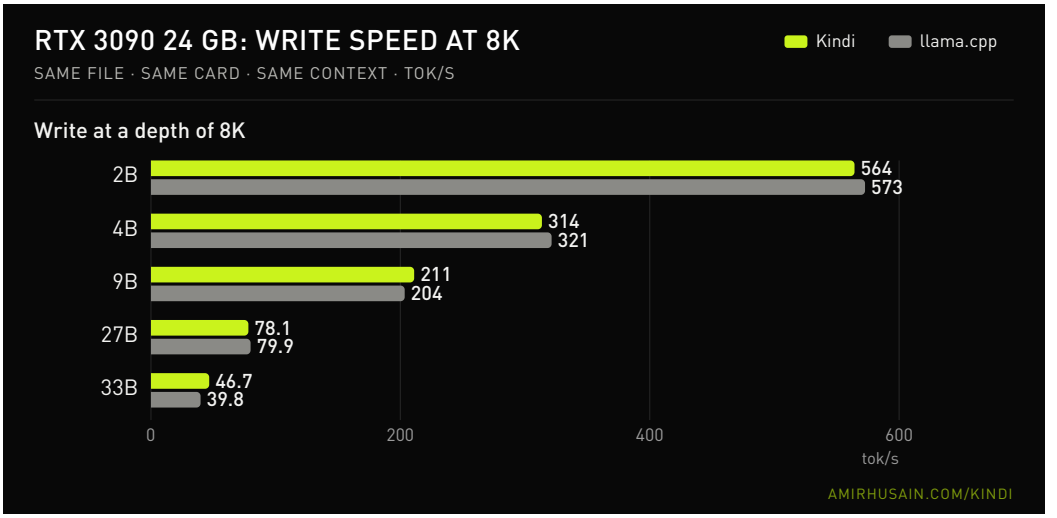

Figure 5-7. *Write speed on one RTX 3090 at a depth of 8K tokens, the same files for both engines. Kindi is lime and llama.cpp is grey.*

Model	Context (tokens)	Kindi read	llama.cpp read	Kindi write	llama.cpp write
2B	512	9,166	11,972	583	605
	8K	8,458	12,964	564	573
4B	512	5,689	6,440	333	355
	8K	4,863	6,426	314	321
9B	512	4,011	4,292	**221**	216
	8K	3,586	4,262	**211**	204
27B	512	**1,486**	1,439	85.9	89.1
	8K	1,288	1,364	78.1	79.9
33B	512	1,192	1,255	**52.5**	42.7
	8K	1,111	1,137	**46.7**	39.8

Conditions. One GeForce RTX 3090 24 GB (350 W, sm_86) on a rented machine. Kindi 0.5.0 (tree f6ca4df), the tensor-core profile, context 16,384 (33B: 16,640); Kindi reads the prompt to the depth, then writes 3 runs of 32 tokens; the speed is the mean. llama.cpp: Prism fork c96baf4, `llama-bench` (write: 64 tokens at the depth, 5 runs; read: 3 runs). The same model files for both. K/V cache: q8_0 for the 27B, f16 for the 33B, the llama.cpp default for the others. `kindi verify` passes on every model (4B: 3 exact, 1 llama.cpp tie; 33B: 8 exact, 3 ties). Kindi's number is in bold where Kindi is faster. 2026-11-06.

On the RTX 3090, Kindi writes the 33B 1.2 times as fast as llama.cpp: 52.5 against 42.7 tok/s at 512 tokens. It also writes the 9B faster, at 512 and at 8K tokens, and it reads the 27B faster at 512 tokens.

> *Does it work on a 3090 or even newer cards like a 5070 (ti) or 5090 too?? Would be very interesting to compare the speeds on those cards from LM Studio vs. your Kindi engine :) Thx so much for your work! Well done!!*
>
> KaEl-Alpha (YouTube), 2026-10-24

NVIDIA CMP 170HX 40 GB

The CMP 170HX has its own chapter, A mining card for AI, with the history of the card, the boards that we measured, and the full table. In short: Kindi writes the 27B 1.07 to 1.6 times as fast as llama.cpp, on one card and on two cards, from 512 to 64K tokens.

Radeon 840M and Radeon 8060S

The Radeon 840M of the Ryzen AI 5 340 and the Radeon 8060S of Strix Halo are in Kindi on AMD graphics. On the 840M and on the 8060S, Kindi writes the 2B, 4B, 9B and 27B faster than llama.cpp.

RTX 20, 30, 40 and 50 series

Kindi supports all four series. We measured speed on the RTX 3080, the RTX 3060 (above and in the chapter about the draft model) and the RTX 3090 (above). With a drafter, the 27B writes faster on both cards, with the same output. A user with an RTX 4060 Laptop GPU (8 GB) ran Kindi 0.3.0 and reported that the 27B writes 28.0 tok/s at 100 tokens and 25.2 tok/s at 24,000 tokens, and that `kindi verify` passes. For the other cards of these series, Kindi is supported and `kindi verify` passes. We give no speed numbers for them.

Radeon 780M

The Radeon 780M is the graphics unit inside the AMD Ryzen 9 7940HS of our test computer. In Kindi 0.5.0, it also runs the 9B (Kindi on AMD graphics). It has no memory of its own, so it reads the weights from the DDR5 memory of the computer, at about 83 GB/s. That memory speed sets the write speed of each model.

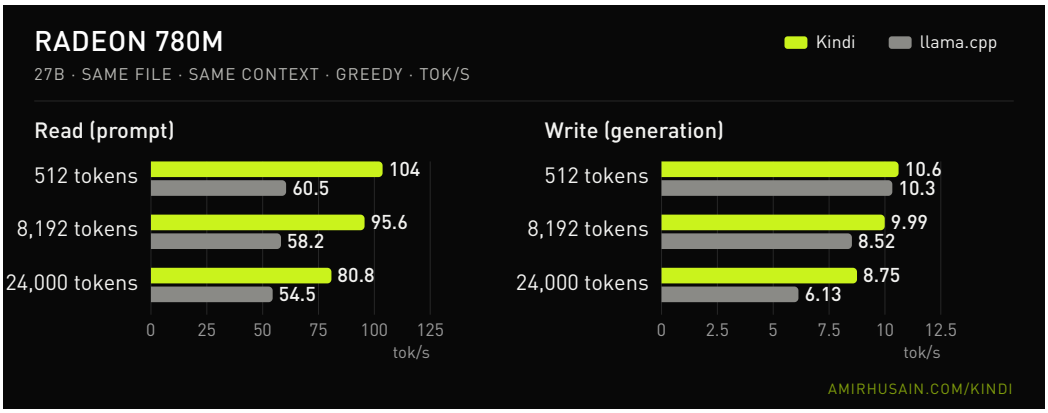

Figure 5-8. *Read and write speed on the Radeon 780M with the 27B. Kindi is lime and llama.cpp is grey.*

Model and context	Kindi read	llama.cpp read	Kindi write	llama.cpp write
27B, 512 tokens	**104**	60.5	**10.6**	10.3
27B, 8,192 tokens	**95.6**	58.2	**9.99**	8.52
27B, 24,000 tokens	**80.8**	54.5	**8.75**	6.13
Nemotron, 512 tokens	**186**	670	**19.1**	24.0
Nemotron, 8,192 tokens	**392**	599	**20.0**	21.4
Nemotron, 24,000 tokens	**239**	477	**17.0**	19.1

> **Conditions.** Minisforum UM790 Pro, AMD Ryzen 9 7940HS with Radeon 780M, 64 GB DDR5-5600, Linux kernel 6.8 with `amdgpu.cwsr_enable=0`. The same model files for both engines. Kindi: Kindi 0.3.0 development build (b5230a3), `HSA_OVERRIDE_GFX_VERSION=11.0.2`. The same token ids for both engines, greedy, one pass, 32 tokens of write at each depth, both engines warmed first. q8_0 K/V for the 27B, f16 for Nemotron. llama.cpp: Prism fork, ROCm build, for the 27B; upstream b10700, ROCm build, for Nemotron. 2026-10-25.

The 27B is the main result on this chip. Kindi reads it 1.5 to 1.7 times as fast as llama.cpp at each context: 104 against 60.5 tok/s with a 512-token prompt, and 80.8 against 54.5 at 24,000. Kindi writes 10.6 tok/s at 512 tokens, about the same as llama.cpp, and 8.75 against 6.13 at 24,000. So the advantage grows as the context grows.

Nemotron is experimental on this chip. We are still optimizing Kindi's Nemotron engine.

Quadro K6000

The Quadro K6000 is the largest card of the Kepler generation: 12 GB of GDDR5 memory on one card. It is the same generation as the K420, but it has 15 times the compute units and 10 times the memory speed. On this card, Kindi runs the 27B on one card, as on the GTX 1080.

Model and context	Kindi read	llama.cpp read	Kindi write	llama.cpp write
27B, 512 tokens	**36.2**	7.0	**20.4**	2.07
27B, 8,192 tokens	**30.2**	6.9	**16.4**	2.15
27B, 24,000 tokens	**26.2**	*llama.cpp: did not finish in 1 hour*	**13.1**	*llama.cpp: did not finish in 1 hour*
Nemotron, 512 tokens	**188**	90.7	**20.1**	10.7
Nemotron, 8,192 tokens	**182**	33.8	**19.1**	9.8
Nemotron, 24,000 tokens	**164**	15.6	**17.4**	8.56

> **Conditions.** Quadro K6000 12 GB (Kepler GK110, sm_35) passed through to a KVM virtual machine on a dual Xeon E5-2673 v4 host, NVIDIA driver 470.256.02. The same model files and the same token ids for both engines, greedy, 32 tokens of write at each depth, both engines warmed first. Kindi 0.3.0 development build (597d2d8) in the CUDA 10.2 image, q8_0 K/V for the 27B, f16 for Nemotron. llama.cpp: Vulkan builds (the Prism fork for the 27B, upstream 1b29621 for Nemotron), `-ngl 99 -fa on`. 2026-10-25.

With the 27B, Kindi reads 4 to 5 times as fast as llama.cpp and writes 7 to 10 times as fast: 20.4 against 2.07 tok/s at 512 tokens. At 24,000 tokens, llama.cpp did not finish the prompt in one hour; Kindi read it in about 15 minutes and then wrote 13.1 tok/s. Both engines give the same tokens (`kindi verify`, 4 of 4). Nemotron also runs on this card, and Kindi writes it about twice as fast as llama.cpp.

We also tried Kindi's tuning settings on this card: the chunk size, the tile, the split, blocking sync and the batch. None of them changed the speed by more than 3%, so the installer uses the same settings as on the other cards.

The CPU

Kindi also runs the Bonsai models on the CPU, with no graphics at all. We compared it with llama.cpp on the CPUs of two common computers, with the same

model files and the same number of threads.

CPU and model	Kindi read	llama.cpp read	Kindi write	llama.cpp write
Ryzen 9 7940HS, 2B	**90.4**	60.5	**64.8**	23.6
Ryzen 9 7940HS, 4B	**37.8**	25.7	**29.1**	16.5
Core i7-8700, 2B	**42.5**	33.5	**28.9**	18.1
Core i7-8700, 4B	**17.0**	13.0	**12.6**	7.75

> **Conditions.** Rainstorm's Intel Core i7-8700 (6 cores, 12 threads) and the AMD Ryzen 9 7940HS of the 780M computer (8 cores, 16 threads), with the graphics idle. The same model files for both engines. Read: a prompt of about 2,000 tokens with no cache reuse. Write: 256 tokens. The median of 3 runs, and each engine at its best number of threads. llama.cpp: Prism fork, CPU build. Kindi: Kindi 0.3.0 development builds (849840a on the 7940HS; 5350b87 with the server fix on the i7-8700). 2026-10-25.

On the 7940HS, Kindi's CPU engine writes the 2B at 64.8 tok/s, and llama.cpp at 23.6: 2.7 times as fast. It reads the 2B prompt at 90.4 tok/s, and llama.cpp at 60.5. On rainstorm's Core i7-8700 (6 cores, DDR4), Kindi writes the 2B at 28.9 tok/s, and llama.cpp at 18.1.

Kindi's CPU engine reads a prompt one token at a time, and it still reads faster than llama.cpp's batched prompt path. A 2B that writes 30 to 65 tok/s on the CPU is fast enough for short text jobs: a spam check or an escalation decision takes a few seconds (Fast decisions on short text).

How well does each card do agent work?

These are the results of the agent tasks in the Kindi benchmark (kbench), with pi. The benchmark has 24 terminal tasks, 32 programming and repair tasks, and 9 chains of small prompts. Each task gets one pi session and a hidden test.

Card	Model	Terminal	Programming and repair	Chains	Time per terminal task
Quadro K420 1 GB	2B	**13 / 24**	**7 / 32**	1 / 9	96 s
Quadro K1200	2B	**16 / 24**	**6 / 32**	0 / 9	25 s
Quadro K1200	4B	**21 / 24**	**23 / 32**	6 / 9	24 s
2× K1200, one 9B server on each card	9B	**22 / 24 †**	**26 / 32**	–	37 s and 176 s
GeForce GTX 1080	27B	**24 / 24**	**32 / 32**	7 / 9	28 s

Conditions. kbench 0.1.x with pi, thinking on, 4,096 output tokens, one run of each task, 1,800 s for each task, exact head. † The terminal score of the 9B is a composite: the first run used the fast head, and we ran the 6 failed tasks again with the exact head. The 9B ran on two servers, one on each K1200, so the time column gives one time for each server. The K420 terminal tasks ran on Kindi 0.1.0.

The larger models do much better. On one K1200, the 4B passed 23 of the 32 programming tasks, and the 2B passed 6. On the GTX 1080, the 27B passed all 24 terminal tasks and all 32 programming tasks.

CHAPTER 6

How does a draft model make the 27B faster?

To write one token, the 27B reads all of its weights from the memory of the card. On a card with Tensor Cores, a pass over two, three or four tokens takes little more time than a pass over one token. Kindi uses this since 0.4.0. The 2B, a much smaller model, guesses the next tokens. Then the 27B checks all of the guesses in one pass. Kindi keeps each guess that is the same as the token of the 27B, and it adds one token of the 27B itself. When the guesses are correct, one pass of the 27B gives several tokens.

The 27B makes every decision. A wrong guess costs a little time, and it does not change the text. With greedy decoding, the output is the same as without the drafter, token for token. With sampled decoding, Kindi checks each guess against a sample from the 27B, so the text comes from the same distribution as without the drafter.

The installer turns this on for the 27B on NVIDIA cards of compute capability 8.0 and later: the RTX 30, 40 and 50 series. You do not have to set anything.

How fast is it?

Card	Context (tokens)	With the 2B drafter (tok/s)
RTX 3060	512	**60.6**
	8K	**73.5**
	24K	**66.7**

> **Conditions.** Model Ternary-Bonsai-2-27B-PTQ1_0.gguf (sha256 53107f53…), with Ternary-Bonsai-2-2B as the drafter. Kindi 0.4.0 development build (1fcf543), tensor-core engine, exact head. The same held-out text on each card, greedy, 128 tokens after a prompt of the given length. The drafter guesses 2 tokens for each pass below 1,000 tokens of context, and 3 tokens from 1,000 tokens. These are the settings that the installer uses. RTX 3080: a card of our RTX 3080 cluster at 190 W, PCIe x1, a host with 4 CPUs. RTX 3060: PCIe x4. In every row, the tokens were the same as without the drafter. 2026-10-30.

On the RTX 3060, the 27B writes 73.5 tok/s with the 2B drafter at 8K. On one RTX 3080, the 27B drafts with DSpine (DSpine, the block drafter of the 27B).

Does it change the agent results?

No. We ran the 65 tasks of the Kindi benchmark with pi on the 27B, on one RTX 3080 at 190 W, with the settings of the installer. The 27B passed 64 of the 65 tasks. Kindi 0.3.0 passed the same 64 tasks with the same model file and settings. The one task that failed was the same task in both runs. During the agent turns, the 27B wrote about 93 tok/s.

> **Conditions.** kbench 0.1.1 with pi 1.3.0, thinking on, one run of each task, 1,200 s for each task, context 16,384. Kindi 0.4.0 development build (59468a9). 2026-10-30.

Is the 2B faster too?

Yes, a little. The 2B has no smaller model to guess for it. In place of a drafter, it uses a small extra layer, an MTP head, that predicts the token after the next one. The installer downloads the head from Hugging Face and checks its sha256.

Card	Context (tokens)	With the head (tok/s)
RTX 3080, 190 W	512	**409**
	8K	**396**
RTX 3060	512	**351**
	8K	**340**

> **Conditions.** Model Ternary-Bonsai-2-2B-PTQ1_0.gguf with its head (`kindi-mtp-2b-v1.npz`). The same cards, build and text as the table above, greedy, 512 tokens, the median of 3 runs. The head guesses 2 tokens for each pass. 2026-10-30.

Kindi uses the head up to 8K tokens of context, where it makes the 2B faster. Above 8K, Kindi turns it off. The head works with greedy decoding only.

Which cards use it?

Speculation runs only on the tensor-core path of Kindi: NVIDIA cards of compute capability 8.0 and later. On the other cards, a pass over two tokens costs almost as much as two single steps: 1.80 steps on the Quadro K1200 with the 2B, and 1.82 on the GTX 1080 with the 27B. There, a guess cannot save time, so the installer does not turn speculation on.

From 32,768 tokens of context, a guard measures the speed. When speculation is slower than single steps, the guard stops it. To turn speculation off, set `KINDI_SPEC=0`.

CHAPTER 7

DSpine, the block drafter of the 27B

Kindi 0.6.0 has a second draft model for the 27B: DSpine. Tencent published DSpine for the Qwen3.8-27B family under the Apache License 2.0 (tencent/Qwen3.8-27B-DSpine-b16 on Hugging Face). The weights are Tencent's. The code that runs them in Kindi is Kindi's own. On an RTX 3080, the 27B writes 115.5 tokens/s with DSpine. The output is the same as without a drafter, token for token.

How does DSpine draft?

The 2B drafter is a separate model. It reads the conversation itself and guesses the next tokens one at a time. DSpine reads the 27B instead. While the 27B reads and writes, it copies the outputs of five of its layers (layers 1, 16, 31, 46 and 61) into a buffer. DSpine has 5 layers of its own, and each of them keeps the last 8,192 positions.

To draft, DSpine takes the current token and 7 empty positions, and it fills all 7 positions in one pass. The 27B's own output layer turns them into token candidates, and a short chain picks each draft token from the token before it. The 27B then checks the current token and the 7 drafts in one pass of 8 tokens. It keeps the drafts up to the first one that it would not have written itself, and it adds its own next token.

The draft takes about 8 ms on an RTX 3080. A pass of the 27B over 8 tokens costs a little more than a pass over one token, so each kept draft token saves most of a pass.

Is the output the same?

Yes. The 27B does not change: its passes only copy the five layer outputs, and the 27B checks every draft token. So greedy output with DSpine is the output without a drafter, token for token. Our check on the RTX 3080, with the installer's 27B settings, compared 8 prompts and 2,009 tokens, and they were the same. Sampled requests

use speculative sampling, so the sampled tokens have the 27B's own distribution. `kindi verify` passes on the RTX 3080 with DSpine.

How fast is it?

Drafter of the 27B	Tokens for each verify pass	Write (tok/s)
DSpine	**3.10**	**115.5**

> **Conditions.** One GeForce RTX 3080 10 GB at 190 W, the Ternary-Bonsai-2-27B file, context 32,768, greedy, the installer's 27B settings (`KINDI_FD=1`, `KINDI_ATT_TC=1`). DSpine drafts a block of 7 tokens, with verify passes of up to 8 tokens. Kindi 0.6.0. llama.cpp cannot load DSpine.

Code and math prompts keep more tokens per pass (up to 4.6) than prose (about 2.3 to 2.8), because their next tokens are easier to guess.

How do you turn it on?

The installer downloads DSpine for the 27B on one card with the tensor-core path, and `KINDI_SPEC=auto` uses it. You do not have to do anything.
`install.sh --no-dspine` leaves DSpine out, and the 27B drafts with the 2B, as in 0.5.0.

Setting	What it does
`KINDI_DSPINE=DIR`	The folder with DSpine's `config.json` and `model.safetensors` (the installer uses `~/kindi/dspine`)
`KINDI_SPEC=auto`	Drafts with DSpine when it is there, else with the 2B drafter. `KINDI_SPEC=dspine` asks for DSpine
`KINDI_TC_TMAX=8`	Verify passes of up to 8 tokens: the whole block. With 4, DSpine drafts 3 tokens for each pass
`KINDI_DSPINE_K`	The draft tokens for each pass, 7 at most (the default)
`KINDI_DSPINE_MB`	The card memory that the context fit keeps for the drafter (default 1400 MB)

DSpine uses about 1.4 GB of card memory, and about 0.75 GB of host memory. It takes about 25 seconds to load. It runs on one NVIDIA card with the tensor-core path. Other models keep their own drafters.

When Kindi restores a conversation from the prompt cache, DSpine has not seen the restored part. It drafts from the tokens that come after the restore, so it keeps fewer tokens at first. The output does not change.

CHAPTER 8

A more capable model

The 27B is the most capable Bonsai model. For harder work, Kindi runs a larger model: AH-Nanbeige4.3-33B IQ3_XXS. This is our IQ3_XXS quantization of Nanbeige4.3-33B, made with bartowski's imatrix. The base model has 33 billion parameters. The file is 13.1 GB. It is a separate download from Hugging Face:

https://huggingface.co/ahtx/AH-Nanbeige4.3-33B-IQ3_XXS-GGUF

The model thinks before it answers, and it makes tool calls in the OpenAI format, so pi and other agents can use it in the same way as the Bonsai models.

Where does it run?

- **One card with 16 GB or more,** for example an RTX 3090 (24 GB).
- **Two cards with about 17 GB in total,** for example two RTX 30 series cards or two Tesla P100 cards.
- **One Quadro K6000 (12 GB),** with part of the model in host memory (see below).
- **About 17 GB of card memory in total.** Kindi puts the layers of the model on the cards in order: the first half on the first card and the second half on the second card. Only a small amount of data goes between the cards, so a PCIe x1 slot is sufficient. We measured it on two RTX 3080 10 GB cards and on two Tesla P100 cards.
- **A context of 16,384 tokens** on the RTX 3080 cards.

To get the model, type `kindi use nanbeige-33b`, or install with `install.sh --model nanbeige`. Both check the sha256 of the file after the download.

`kindi verify` compares the output of Kindi with the output of llama.cpp for the same file. With this model, 8 of the 11 test prompts were the same token for token.

The other 3 continued in the same way up to a near tie, where llama.cpp itself had two tokens with almost the same score. No prompt failed.

Before the server takes its first request, it does a self-check. It runs one prompt two times with single steps and one time with a check pass, and the tokens must be the same. If the check fails, the server stops with an error, and the service starts it again, so a fault at start-up cannot give a wrong reply. The check is on by default (`KINDI_NB_SELFCHECK=1`).

How fast is it?

Kindi writes the 33B faster than llama.cpp on two RTX 3080 cards with the same file, and it reads faster at every prompt length that we measured. At a short context, Kindi writes 45.8 tok/s, and llama.cpp 36.7.

K/V cache	Depth (tokens)	Kindi write	llama.cpp write
f16	0	**45.8**	36.7
	2,048	**44.6**	35.6
	8,192	**40.9**	33.3
q8_0	0	**45.5**	35.6
	2,048	**44.6**	33.3
	8,192	**41.9**	28.1

K/V cache	Prompt (tokens)	Kindi read	llama.cpp read
f16	512	**1,036**	785
	2,048	**1,278**	1,164
	8,192	**1,267**	1,249
q8_0	512	**1,030**	779
	2,048	**1,280**	1,148
	8,192	**1,252**	1,223

Conditions. Two GeForce RTX 3080 10 GB cards in one host, 26 layers on each card. The same file for both engines (AH-Nanbeige4.3-33B-IQ3_XXS.gguf, sha256 9d4f18e8…), and the same K/V cache type for both. Kindi: a 0.6.0 development build (e5e8d8b), `kindi bench`, 128 timed tokens at each depth for the write, one prompt of each length for the read. llama.cpp ad1b88c (CUDA, sm_86), `llama-bench -ngl 99 -fa 1`, 2 runs. `kindi verify` passes with both caches. 2026-11-10.

Through the server, with the OpenAI API, Kindi writes 45.4 tok/s greedy, and 44.2 tok/s sampled (temperature 0.7, top-k 20, top-p 0.8). The sampler finds the candidate tokens on the card, so only a few scores go over the PCIe x1 link.

On two Tesla P100 cards, Kindi writes the 33B at 21.0 tok/s at 512 tokens, about 2.2 times as fast as llama.cpp (Tesla P100 16 GB). The model also runs on a single RTX 3090 (24 GB).

It is a strong model for coding work, and it can do even better on a slightly faster GPU.

The 8-bit K/V cache

The K/V cache holds what the model has read. With f16 numbers, a context of 64K tokens of the 33B needs 13 GiB of cache. Kindi 0.6.0 can also keep the cache in 8-bit numbers (q8_0), at about half the size. The setting `kv_cache auto`, the new default, uses f16 when the model and the f16 cache fit at the chosen context. When they do not fit, it uses q8_0.

With the q8_0 cache at a depth of 8,192 tokens, Kindi writes 41.9 tokens/s, and llama.cpp 28.1 (the table above).

Both caches pass `kindi verify`. With f16, 8 of the 11 test prompts are the same token for token, and the other 3 end at a near tie of llama.cpp. For q8_0, we counted the tokens before the first difference from the f16 output: Kindi with q8_0 keeps 475 of 586 tokens, and llama.cpp with q8_0 keeps 471.

Can a Quadro K6000 run it?

Yes. The 33B file (13.1 GB) is larger than the 12 GB of the K6000. Kindi keeps 13 layers of the model in host memory, at a context of 8,448 tokens. While the card computes the layers on the card, it copies the host layers in, and the CPU computes a part of the work of each host layer. The result is 2.3 times the speed of llama.cpp at 8K tokens, with 10 layers on the CPU:

Prompt (tokens)	Kindi	llama.cpp	Times as fast
480	**3.85**	2.19	1.76×
2,016	**3.55**	1.96	1.81×
8,160	**3.56**	1.52	2.34×

> **Conditions.** One Quadro K6000 12 GB (a virtual machine with PCI pass-through), driver 470.256.02, a context of 8,448 tokens, an f16 K/V cache for both engines. Kindi: 13 host layers, the CPU computes 55% of the FFN rows of each host layer (8 threads). llama.cpp ad1b88c with `-ngl 42`. Write is the speed of 32 tokens after the prompt. `kindi verify`: 9 of 11 prompts token-exact, 2 end at a llama.cpp tie, 0 fail.

At a context of 1,024 tokens, Kindi writes 4.4 to 5.0 tok/s (two runs), and llama.cpp writes 2.5 to 2.7 tok/s.

How does prompt lookup help?

An agent copies much text: the contents of a file, the output of a command, a tool call that repeats earlier text. Prompt lookup uses this. Kindi finds the last earlier place in the prompt and the reply where the last few tokens occur. Its guess is the text that came after that place. The 33B then checks the guess in one pass, in the same way as with the 2B drafter. A guess from the prompt costs no model work.

Kindi measures how many guesses the model keeps, and it sets the length of each guess to give the most tokens for each unit of time. When guesses do not help, as in new prose, Kindi makes no guess. This is on by default.

When the model copies a file with changes, the 33B writes 91.6 tok/s with prompt lookup. Tool calls also get faster. New prose gets no guesses.

> **Conditions.** Two RTX 3080 cards of our cluster, 190 W, PCIe x1, context 8,192. `kindi serve`, 400 tokens out, prompt lookup on (up to 5 tokens for each pass). Kindi 0.4.0 development build (10df53d). The greedy replies were the same as without prompt lookup. 2026-10-30.

How does it compare with the 27B?

In an initial test, we gave the 33B and the 27B the same tasks, one run of each task. The first test was the 65 tasks of the Kindi benchmark, with pi. The second test was 50 programming problems from LiveCodeBench, a public set of contest problems with hidden tests.

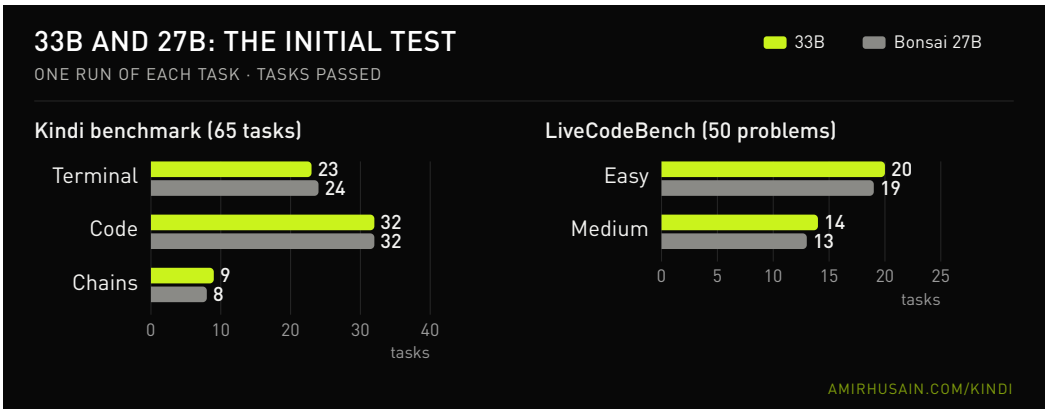

Figure 8-1. *Tasks that the 33B and the 27B passed in the initial test, one run of each task. The 33B is lime and the 27B is grey.*

Test	AH-Nanbeige4.3-33B	Bonsai 27B
Kindi benchmark: terminal	**23 / 24**	24 / 24
Kindi benchmark: programming and repair	**32 / 32**	32 / 32
Kindi benchmark: chains	**9 / 9**	8 / 9
Kindi benchmark: all	**64 / 65**	64 / 65
LiveCodeBench: easy	**20 / 21**	19 / 21
LiveCodeBench: medium	**14 / 29**	13 / 29
LiveCodeBench: all	**34 / 50**	32 / 50
Both tests	**98 / 115**	96 / 115

> **Conditions.** The 33B: llama.cpp (the official CUDA build) on two RTX 3080 cards of our desktop computer (WSL2), with our IQ3_XXS file. The 27B: Kindi 0.3.0 with its tensor-core settings and the 2B drafter, on two servers on RTX 3080 cards of our cluster at 190 W. Kindi benchmark: kbench 0.1.1 with pi 1.3.0, thinking on, 4,096 output tokens, context 16,384 on the pi side (the manifest of each run), 1,200 s for each task, one run of each task. LiveCodeBench: the 50 most recent easy and medium problems of its code_generation_lite set, one answer for each problem, thinking on, 4,096 output tokens, temperature 0.7, top-p 0.8, top-k 20, with the problem's public and hidden tests. 2026-10-29.

The two models passed the same number of Kindi benchmark tasks: 64 of 65. The 33B failed one terminal task, and the 27B failed one chain. On the LiveCodeBench problems, the 33B solved 34 of 50, and the 27B solved 32.

Can you use both models together?

You can use both models together. An agent tool such as opencode can send the fast steps to the 27B and the harder coding steps to the 33B, because both run on Kindi's OpenAI-compatible server.

CHAPTER 9

Mixture-of-experts models

A mixture-of-experts model has many small feed-forward blocks, the experts, in each layer. A router picks a few of them for each token. So a large model reads only a small part of its weights for each token, and it can write quickly. The other parts must still be somewhere: on the card, in host memory or on a disk. Kindi 0.6.0 has one placement for all its mixture-of-experts engines. It counts how often each expert is used, and it puts the most used experts on the card.

Kindi 0.7.0 runs three such models: Nemotron 3.5 Lightning 30B-A3B in this chapter, and DeepSeek-V4.1-Flash and Qwen3.8-Flash-Next in the next two chapters.

Nemotron 3.5 Lightning 30B-A3B

Nota AI made a ternary version of NVIDIA's Nemotron 3.5 Lightning 30B-A3B, and boro64 published it as a GGUF file (Q2_0, 9.47 GB). The model has 30 billion parameters, and about 3 billion are active for each token. It has three kinds of layers: Mamba-2 layers, mixture-of-experts layers with 128 experts each, and attention layers. Kindi runs it with its own engine, with chat, streaming, the thinking block and tool calls. The weights are not part of Kindi. Their licence is OpenMDW 1.1.

The installer offers Lightning (`--model lightning`) on NVIDIA cards with 6 GB or more from the GTX 10 series on, on the Tesla P100 and on the 12 GB Kepler cards. The default context is 32K. On a computer with Kindi, `kindi use lightning-30b` changes the server to Lightning.

How fast is Lightning?

Kindi reads and writes Lightning faster than llama.cpp on every card and at every context that we measured. On an RTX 3080, Kindi reads a 32K prompt at 3,777

tokens/s, and llama.cpp at 2,138. On the Quadro K6000, Kindi writes 61.9 tokens/s at 512 tokens, and llama.cpp 35.1.

Card	Context (tokens)	Kindi read	llama.cpp read	Kindi write	llama.cpp write
RTX 3080, 190 W	512	**4,258**	2,725	**310**	202
	8,192	**4,114**	2,597	**299**	198
	32,768	**3,777**	2,138	**277**	182
	60,320	**3,435**	1,670	**255**	168
RTX 3090, 220 W	512	**4,786**	3,715	**341**	284
	8,192	**4,520**	3,663	**326**	281
	32,768	**4,014**	3,390	**295**	261
RTX 2060 12 GB	512	**1,952**	1,256	**164**	144
	8,192	**1,992**	1,251	**155**	139
	32,768	**1,734**	1,156	**135**	128
Tesla P100 16 GB	512	**654**	516	**155**	123
	8,192	**702**	506	**145**	118
	32,768	**605**	449	**118**	108
Quadro K6000 12 GB	512	**206**	165	**61.9**	35.1
	8,192	**192**	141	**48.3**	28.7
	32,768	**125**	118	**32.8**	23.3

Conditions. The same file (Ternary-Nemotron-3.5-Lightning-30B-A3B, sha256 38854de6…) and f16 K/V for both engines. llama.cpp 785322d (`llama-bench -ngl 99 -fa 1`; read `-p N -n 0`, write `-n 64 -d N`), built for each card. Kindi: `kindi bench` with the same prompt lengths, exact head, no drafter; it writes 64 tokens after 16 warm-up tokens at the same depth. RTX 3080 at 190 W; RTX 3090 at 220 W, RTX 2060 12 GB and Tesla P100-SXM2 16 GB on rented machines; Quadro K6000 at stock clocks. 60,320 tokens is the largest context with all experts on the RTX 3080. Kindi 0.6.0 development builds, 2026-11-03 to 2026-11-10.

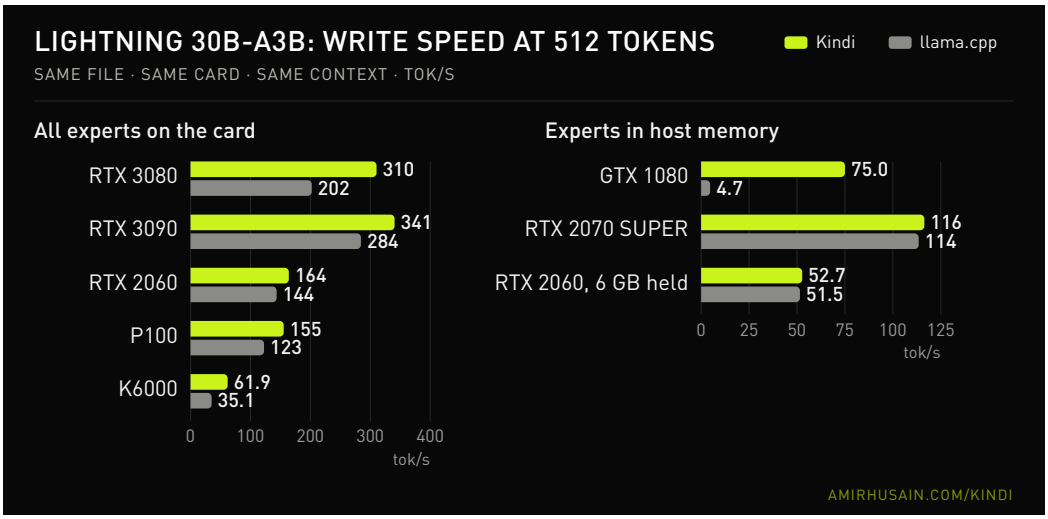

Figure 9-1. *Lightning: write speed at 512 tokens, Kindi and llama.cpp, the same file on the same card.*

Cards with less than 10 GB

All 128 experts of each layer do not fit on an 8 GB card. On these cards, Kindi keeps the most used experts on the card and the others in host memory. The card reads a host expert over PCIe when a token needs it. During generation, each layer also keeps its last used host experts in 10 slots on the card, so an expert that the next token uses again is not read again.

llama.cpp must also move experts to the CPU to fit these cards. We used the smallest `-ncmoe` value that fits (the experts of N layers on the CPU), and we give it in each row.

Card	Context (tokens)	Kindi read	llama.cpp read	Kindi write	llama.cpp write
GTX 1080 8 GB (llama.cpp -ncmoe 9)	512	**745**	646	**75.0**	4.7
	8,192	**687**	615	**89.5**	4.7
	32,768	**536**	495	**73.9**	4.0
RTX 2070 SUPER 8 GB (-ncmoe 12)	512	**2,121**	1,137	**116**	114
	8,192	**2,129**	1,137	**120**	112
RTX 2060, 6 GB held back (-ncmoe 28)	512	**1,415**	520	**52.7**	51.5
	8,192	**1,423**	526	**54.0**	51.0
	32,768	**1,415**	515	**57.1**	47.0

Conditions. The same file and f16 K/V for both engines, llama.cpp 785322d. Write: the first N tokens of a long text as the prompt, then 128 greedy tokens, the same token ids for both engines (`tools/lt_rtw.py`; llama.cpp through `llama-server`). Read: `kindi bench` and `llama-bench`. GTX 1080 8 GB (PCIe 2.0 x16, a 2-core host CPU with no AVX, so llama.cpp was built for SSE4.2): Kindi keeps 95 of the 128 experts of each layer on the card, with 6 decode slots; llama.cpp `-ncmoe 9`. RTX 2070 SUPER 8 GB (PCIe 3.0 x16, rented): Kindi keeps 74 to 122 experts of each layer on the card, with 10 decode slots; llama.cpp `-ncmoe 12`. RTX 2060 with 6,494 MB held back by another process (rented): Kindi keeps 60 of 128 experts of each layer on the card, with 6 decode slots; llama.cpp `-ncmoe 28`. `kindi verify` passes on the RTX 2070 SUPER. Kindi 0.6.0 development builds, 2026-11-03 to 2026-11-10.

One placement for every engine

The placement is one part of Kindi (`kindi/moe.py`), and the Lightning and DeepSeek engines both use it. It starts from measured pick counts: how often the router picked each expert on real text. With host experts, the card slots go to the most used experts over all layers, so one layer can get more slots than another. This helps most on a slow PCIe link. `KINDI_MOE_WEIGHTED=0` gives each layer the same number of slots.

How well does Lightning work as an agent?

Lightning is a fast extra model next to the 27B. We gave Kindi and llama.cpp the same 65 tasks of the Kindi benchmark with the same file, one task at a time on one

RTX 3080 for each engine. Kindi solved 17 tasks and llama.cpp 11. Kindi also took 32 seconds for a task on average, and llama.cpp 152.

> **Conditions.** kbench, the base8k setup, one seed, 900 s for each task, context 32,768, the same GGUF. Kindi 0.6.0 release candidate (1502fa2) with `KINDI_REPEAT_STOP=1`, the installer's default; llama.cpp 785322d (`-ngl 99 --jinja`). With one seed, the result of each task varies from run to run.

Kindi's Lightning server returns one tool call for each reply, unless the request asks for parallel calls. It reads a call that a stop token cut, and it drops a call to a tool that the request did not offer. In our test of 60 requests, 57 came back as correct tool calls.

The Lightning backend was started by the work of slavalevski-tech, who first ran this model on Kindi and shared that backend under the Kindi CLA 1.1. That work set the scope of the backend. The Kindi code is a new implementation. Thank you.

CHAPTER 10

DeepSeek-V4.1-Flash on one RTX 3060

DeepSeek-V4.1-Flash has 552 billion parameters, and 16 billion are active for each token. Kindi 0.6.0 runs it at its published 4-bit weights (the UD-Q4_K_XL GGUF, about 494 GB), with no further quantization, on one NVIDIA card with 12 GB and a dual-Xeon DDR4 server with about 560 GB of RAM. The engine is Kindi's own. The dense layers and the attention run on the card. The experts run on the CPU, from host memory. A second card can hold the most used experts.

What does it need?

Part	What it needs
Operating system and CPU	Linux on x86-64, with AVX2
Card	NVIDIA, sm_50 or newer, with 12 GB or more for the dense layers and the drafter. An RTX 3060 12 GB works
A second card (optional)	Any NVIDIA card, sm_50 or newer, with 4 GB or more. Kindi fills it with the most used experts
Host memory	About 560 GB or more. The whole model stays in memory
Disk	About 530 GB free for the download and the converted drafter
Software	The NVIDIA driver, the CUDA toolkit (`nvcc`), `gcc`, and `python3` with `numpy`

The model file can be on any disk. Kindi reads it once at the first start, and then it serves it from memory. On a spinning disk, the first read takes 40 minutes or more.

Where do the parts of the model go?

Each token goes through the dense layers, the attention and the router on the card. Then the CPU computes the experts that the router picked, and the card continues with the next layer. The experts are most of the 494 GB, so they stay in host memory, where the CPU reads them at its full memory speed.

A server with two CPU sockets has two memory nodes. Before the server starts, Kindi places the rows of each expert on the node whose cores compute them (`python -m kindi.dsv41 numa MODEL --rowsplit`). Then both memory controllers read at the same time.

The second card holds the experts that the router picks most often. Kindi plans this list from measured pick counts, with the same placement as the Lightning engine (Mixture-of-experts models). On our test server, a Quadro M4000 8 GB holds 430 of the most used experts. The card computes those picks while the CPU computes the others.

How does DSpark make it write faster?

DeepSeek publishes three extra prediction stages with the model (the MTP stages). They are not in the GGUF file, so the installer downloads them from DeepSeek's original files and converts them. Kindi calls this drafter DSpark.

After each step, DSpark guesses the next 5 tokens. The main model then checks the drafts in one pass, keeps them up to the first draft that it would not have written itself, and adds its own next token. A pass over several tokens reads more experts than a pass over one token, because different tokens pick different experts. So before each pass, Kindi decides how many drafts to check: it compares the cost of the pass, measured from the experts of recent passes, with the chance that each draft is kept. The output is the same as without DSpark, token for token.

With sampling, DSpark stays on. Kindi keeps each draft with the probability that the model gives it, and it replaces a rejected draft with a sample from the rest of the distribution. So the output has the distribution of plain sampling.

How fast is it?

On one RTX 3060 12 GB and a dual Xeon E5-2673 v4 server with 768 GB of DDR4-2400, Kindi writes DeepSeek-V4.1-Flash at 18.7 tokens/s with its built-in DSpark drafter. DSpark keeps 4.4 tokens for each pass of the main model. Kindi reads a prompt at about 99 tokens/s at 2K tokens and 98.5 at 8K.

DeepSeek-V4.1-Flash on one RTX 3060	tok/s
Write, with DSpark (4 agent prompts, 128 tokens each)	**18.7**

Conditions. One GeForce RTX 3060 12 GB (PCIe 3.0 x16) for the dense layers, the attention and DSpark; one Quadro M4000 8 GB with 430 hot experts; two Xeon E5-2673 v4 CPUs (40 cores) with 768 GB of DDR4-2400 memory; the model in memory, the expert rows split over the two memory nodes. The DeepSeek-V4.1-Flash-1017 UD-Q4_K_XL GGUF, unchanged. Write: `kindi.dsv41 specbench`, 4 agent prompts of about 700 to 850 tokens, 128 greedy tokens each with DSpark, the same tokens as without a drafter, the mean of 3 runs. Read: `kindi.dsv41 bench`, an agent request as real text, cold, no prompt cache, the mean of 3 runs. Kindi 0.6.0 release candidate 2 (82caaaa), 2026-11-11.

How do you install it?

```
bash tools/install_dsv41.sh --dry-run    # prints the plan and changes nothing
bash tools/install_dsv41.sh              # checks the host, downloads,
converts, starts the service
```

The installer checks the host and picks the settings. It downloads the model and the drafter, and it checks each file against the published sha256 lists. A download that stops continues when you run the installer again. It converts the drafter, plans the hot experts for your second card, and writes a systemd user service, `kindi-dsv41.service`, with the OpenAI API on `127.0.0.1:8101`. It does not replace an existing service unless you add `--force`. The first start compiles the CPU and CUDA code into `~/.cache/kindi` and reads the model.

Option	What it does
`--dir DIR`	Where the model goes
`--port 8101`, `--ctx 65536`	The port and the context
`--no-service`	Does not write the service
`--skip-download`	The files are already in `DIR/gguf` and `DIR/dspark`
`--yes`	Answers yes to every question
`--force`	Replaces an existing `kindi-dsv41.service`

How do you check it?

```
curl -s 127.0.0.1:8101/health
python3 -m kindi.dsv41 verify MODEL-00001-of-00009.gguf [--sha]
```

`verify` is the release check. It gives 20 prompts of 64 tokens to Kindi and compares each next-token choice with the greedy output of the reference implementation. It passes when 98.5% or more of the choices agree, and when each difference is a near-tie (the reference's own gap between its top two tokens is 0.75 or less). For this check, Kindi uses the arithmetic of the reference (`--flags 160`). The release passed with 1,266 of 1,280 choices (98.9%), and the largest gap at a difference was 0.623. Kindi's default arithmetic follows DeepSeek's own FP8 and FP4 rounding. `--sha` also checks every file against its sha256; it reads all 494 GB.

What does the server do for an agent?

The DeepSeek server has the server features of Kindi 0.6.0 (What does the server do for an agent?). Three of them matter most with a model of this size:

- **Saved states.** Reading a long prompt takes minutes on this model, so the server saves the state of each conversation to the disk after each answer, and when it stops. After a restart, the next turn of a saved conversation starts in seconds.
- **The shared start of a prompt.** An agent sends the same system prompt and tool list at the start of each new session. The server saves the state at the end of that shared part, so a new session reads only its new tokens.

- **Warm-up.** `POST /v1/kindi/warm` reads the last request of a client into the cache and does not answer. A client can call it after a restart, so the first real turn is fast.

CHAPTER 11

Qwen3.8-Flash-Next

Qwen3.8-Flash-Next is a mixture-of-experts model with about 125 billion parameters. About 6 billion of them are active for each token. Kindi runs ISTA-DASLab's GSQ-RCO IQ3_S file of the model: two GGUF shards, 83.6 GB in total. The dense layers, the output head and the most used experts run on the graphics card. The other experts run on the CPU, from host RAM. The card's expert kernels and the CPU's expert kernels give the same bits, so the place where an expert runs changes the speed and does not change the output.

What does it need?

Part	What it needs
Graphics card	NVIDIA, with 8 GB free on the main card. A card with 12 GB or more holds more experts. With several cards, Kindi splits the layers over them
Host RAM	About 96 GB or more is recommended. The experts that are not on the card are about 50 GB, and the CPU reads them for each token. With less RAM, the experts are read from the disk, which is much slower. On computers with less than about 128 GB, set `KINDI_QXG_PLE_LOCK=0`
CPU	x86-64 with AVX2. More cores give more speed
Disk	About 92 GB free: the model and the MTP head while the installer converts it
Software	Linux, the NVIDIA driver, Python 3.10 or later with `venv`, and gcc. The installer puts PyTorch and the server's packages in their own environment

`kindi cards` lists your cards, and `kindi plan --model flashnext` shows where the model would run (Which card runs which model?).

How do you install it?

```
bash install.sh --model flashnext
```

The installer:

1. checks the computer and picks the cards with `kindi plan`;
2. makes a Python environment with PyTorch in `~/.kindi/flashnext/venv`;
3. downloads the two shards from ISTA-DASLab's repository on Hugging Face, and checks the sha256 of each shard;
4. gets the MTP head and converts it (see below);
5. runs `kindi verify`: 8 prompts and 174 tokens, compared with a reference run of the same file. Each token must be the top token of the reference, or a near-tie of the reference (top two within 0.35);
6. starts the service `kindi-flashnext.service`: Kindi's OpenAI server on `127.0.0.1:8102`.

The model id is `qwen38-flash-next`. Thinking is on by default. Tool calls come back as OpenAI `tool_calls`. A new turn of the same chat reads only the new tokens. `/health` shows the experts on the card and the share of picks that the card serves.

The MTP head

The multi-token prediction (MTP) head of the model is not in the GGUF file. The installer gets its 31 tensors (5.2 GB) from Qwen's own release, `Qwen/Qwen3.8-Flash-Next` on Hugging Face. It reads only those byte ranges, not the whole repository, and Kindi's importer converts them on your computer. Kindi does not host converted weights. The head is under the Qwen Community License 1.0, and when you install it, you accept that license. `--no-mtp` (or `KINDI_MTP=0`) leaves it out.

When the head is installed, Kindi drafts with it by default (`KINDI_FN_DRAFT=auto`): for greedy requests, on one card with about 6.2 GB free, the head guesses 2 tokens and the model checks them in one pass. On cards of compute capability 8.0 or later, the output is the same as without drafting. On Maxwell, Pascal, Volta and Turing cards, the output can differ from plain decoding where two tokens are almost equally likely; set `KINDI_FN_DRAFT=0` for exact output. For sampled requests, Kindi writes without drafting. The speeds on the RTX 3060 below are with the head.

How fast is it?

On an RTX 3060 12 GB and two Xeon E5-2673 v4 CPUs (40 cores) with DDR4 memory, the model in RAM, Kindi writes Qwen3.8-Flash-Next at 39.4 tokens/s on chat answers. It reads a 4K-token prompt at 489 tokens/s and a 16K-token prompt at 465 tokens/s.

> **Conditions.** 8 chats that the model had not seen, 200 tokens each, three runs (39.84, 39.22 and 39.35 tokens/s), greedy, with the default settings: fast mode, drafting and 8-bit dense layers on. `kindi verify` passes 174 of 174 tokens (it checks the exact path). Kindi 0.7.0 development build, 2026-11-22.

Maxwell and Pascal cards

Kindi also runs Flash-Next with a Maxwell or Pascal card as the main card. Kindi builds the Flash-Next card kernels for sm_50, sm_52 and sm_60. Cards below compute capability 6.1 have no 8-bit dot product instruction, so Kindi uses its own byte products there. They give the same bits as the other kernels: `kindi verify` passes with all 174 tokens, with and without drafting.

On a Quadro M4000 8 GB (Maxwell, sm_52) with the CPU, Kindi writes Flash-Next at 13.1 tokens/s and reads a 4K-token prompt at 54.7 tokens/s. llama.cpp, on the same card with the same file, writes 5.4 tokens/s and reads 34.2.

Setup	Context (tokens)	Kindi read	llama.cpp read	Kindi write	llama.cpp write
Quadro M4000 8 GB + CPU	4K	**54.7**	34.2	**13.1**	5.4

> **Conditions.** One Quadro M4000 8 GB (sm_52) and the CPUs of the same computer, the same GGUF file for both engines. Kindi 0.7.0 development build, default settings (fast mode, drafting and 8-bit dense layers on); the three write runs were 14.09, 13.09 and 12.88 tokens/s. llama.cpp built for sm_52 with CUDA 12.4. `kindi verify` passes 174 of 174 tokens. 2026-11-22.

For the requirements of Maxwell and Pascal cards, see `docs/FLASHNEXT.md`.

Fast mode, drafting and 8-bit dense layers

Three speed features are on by default. Each has a switch that gives the exact or the float path back.

- **Fast mode.** Kindi does not compute the experts that the router ranks 7th or lower when they are not on the card, and it computes the prompt's attention with a faster arithmetic (TF32). On kbench it solves the same tasks as the exact mode (62 of 65 each), but its answers are not token-identical to the exact path. `KINDI_FAST=0` gives exact output. `kindi verify` always checks the exact path.
- **Drafting with the MTP head.** On Maxwell, Pascal, Volta and Turing cards, the output can differ from plain decoding at near-ties. `KINDI_FN_DRAFT=0` gives exact output.
- **8-bit dense layers.** The dense layers read 8-bit integer inputs (dp4a). On kbench this path solves 62 of 65 tasks (the float path: 61 of 65). `KINDI_QXG_IDENSE=0` gives the float path back.

CHAPTER 12

Can two cards work as one?

The weights of the 27B need about 6 GB, and the context needs more memory. One K1200 has 4 GB. Kindi joins two low-profile K1200s into one 8 GB AI card. It does the same with Nemotron.

What does tensor parallel mean?

Each layer of the model is a set of large matrices. Kindi cuts each matrix in two and puts one half on each card. For each token, the two cards do their halves at the same time. Then they exchange a small amount of data and join the results. So both cards work all the time. This method is tensor parallelism.

llama.cpp uses a different method on two cards: a layer split. It puts some layers on the first card and the other layers on the second card. The cards then work one after the other, and one card waits while the other card works.

How fast is it?

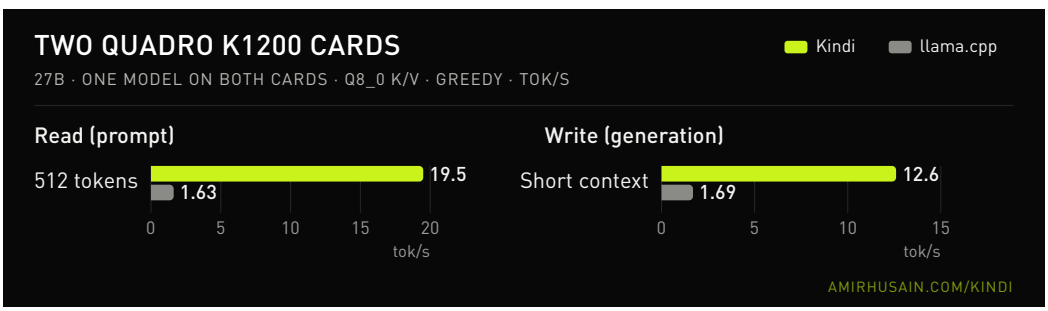

Figure 12-1. *Read speed with a 512-token prompt and write speed at a short context, for the 27B on two Quadro K1200 cards with the same model file. Kindi is lime and llama.cpp is grey.*

Measure	Kindi (tensor parallel)	llama.cpp (Vulkan, layer split)
Write, short context	**12.6**	1.69
Read, 512-token prompt	**19.5**	1.63
Context	24,576	*loads only with its default layer split*

> **Conditions.** Two Quadro K1200 4 GB cards in one computer. Model Ternary-Bonsai-2-27B-PTQ1_0.gguf (sha256 53107f53…) for both engines. q8_0 K/V cache for both, greedy. llama.cpp: Prism fork, Vulkan build (this computer has no CUDA toolkit for llama.cpp), flash attention on, its default layer split. Kindi: Kindi 0.2.0, tensor-parallel engine, exact head, 2026-10-23. Read: the same 512 token ids for both engines; llama.cpp llama-server with a context of 1,024 and no host prompt cache; Kindi chunked prompt path, best of 3 runs; 2026-10-22.

With the same model file on the same two cards, Kindi writes 12.6 tok/s and llama.cpp writes 1.69 tok/s. Kindi is 7.5 times faster. With the same 512-token prompt, Kindi reads 19.5 tok/s and llama.cpp reads 1.63 tok/s: Kindi is 12.0 times faster. llama.cpp loads the model only with its default layer split. With an even split, it runs out of memory. At 22,000 tokens, Kindi writes 9.84 tok/s on the two cards.

12.6 tok/s is faster than most people read, so you do not wait for the text.

Nemotron also runs split over the two cards:

Context (tokens)	Kindi read	llama.cpp read	Kindi write	llama.cpp write
512	**123**	15.3	**17.2**	7.44
8,000	**124**	19.0	**16.1**	7.37

> **Conditions.** The same two cards. Model NVIDIA-Nemotron-3.5-Bolt-E4B-UD-Q3_K_XL.gguf (sha256 3f8032b9…) for both engines. f16 K/V for both, the same method as the table above. Kindi: tensor parallel; the context fits 8,448 tokens; Kindi 0.3.0 development build (5350b87 with the server fix). llama.cpp: upstream Vulkan build 1b29621, its layer split. 2026-10-24.

Split over the two cards, Kindi writes Nemotron at 17.2 tok/s and llama.cpp at 7.44: Kindi is 2.3 times as fast. Kindi reads a 512-token prompt at 123 tok/s, and llama.cpp at 15.3. On two 4 GB cards, Nemotron's context is 8,448 tokens: the weights take 2.7 GB of each card's 4 GB.

Can each card run its own model?

Yes. Two cards can also be two separate servers, with one model on each card. Then two requests run at the same time, and each request runs at the speed of one card. The installer sets this up in one step: select one setup for each card, or a setup with `x2`, for example a 4B on each of two K1200 cards. Each server gets its own address.

Which is better? One model on both cards gives you a larger model for one user or one agent. One model on each card gives you two engines, for two agents or for a program that sends many short requests.

A worked example: Rainstorm



An HP EliteDesk 800 G4 SFF. Photo: eBay seller, used with permission.

OUR TEST COMPUTER

Rainstorm

Computer	A used HP EliteDesk 800 G4 small-form-factor (SFF) office computer
Cards	Two Quadro K1200 4 GB, low profile
Cost	**US$221.58** for all the parts
Power	The two cards take all their power from the slots. The computer needs no extra power cables.
Model	27B, tensor parallel, context 24,576
Speed	The 27B writes **12.6 tok/s** at a short context and 9.84 at 22,000 tokens

Rainstorm runs in text mode, with no desktop. We connect to it over the network, and pi or a script sends it requests.

Rainstorm shows what a cost-effective office computer can do. The installer offers it these setups:

What you want	The setup
The most capable model	the 27B split over both cards
A model from the second family	Nemotron split over both cards
Two jobs at the same time	one model on each card, for example a 4B on each card

We installed these three setups on rainstorm with the installer, from an empty folder. Each one passed `kindi verify` and a chat request, and Nemotron also answered a tool call.

CHAPTER 13

A mining card for AI

NVIDIA made the CMP 170HX to mine a cryptocurrency. It has the same chip as the NVIDIA A100, a data center card for AI, but NVIDIA shipped it with limits that made it slow for most other work. This chapter tells the short history of the card and its limits. Then it gives Kindi's measurements on the boards that we rented, and how those boards differ from the stock card.

Where does the card come from?

On 18 February 2021, NVIDIA announced the CMP (Cryptocurrency Mining Processor) line of cards to mine Ethereum. On the same day, NVIDIA said that the driver of the GeForce RTX 3060 would find the algorithm that mines Ethereum and cut its hash rate by about 50 percent.[1] NVIDIA's CMP page lists four models, the 30HX, 40HX, 50HX and 90HX, from 26 to 86 MH/s. It says that a CMP card has an open bracket for air flow, and that one CPU can control a larger number of these cards.[2]

The CMP 170HX was not on that list. It appeared in September 2021: a GA100 chip, the chip of the A100, with 4,480 CUDA cores, 8 GB of HBM2e memory and no display outputs. A tested card mined Ethereum at about 164 MH/s at 250 W.[3] In October 2021, it went on sale for about US$4,300 a card.[4] The card has no fan, and it needs the air flow of a server. It takes power through one 8-pin EPS (CPU-style) connector. Its chip has 70 active SMs, and the A100 has 108 SMs.[5]

Which limits did NVIDIA put on it?

A review of the card measured these rates on a stock CMP 170HX:[6]

Measure	Stock CMP 170HX
FP32 with FMA (fused multiply-add)	394 GFLOPS
FP64 with FMA	182 GFLOPS
FP16	about 42 TFLOPS
Memory bandwidth	1,355 GB/s
Memory	8 GiB
PCIe link	PCIe 1.0 x4, about 1 GB/s

The memory stayed fast, because a miner needs memory bandwidth. Most other work needs arithmetic, and FMA is the most common arithmetic instruction on a GPU. A community wiki about the card reports that fuses in the chip divide the rate of the FMA instruction and of the integer multiply units by 32. The same wiki reports that only 4 of the 16 PCIe lanes work, because NVIDIA left off parts on the board for the other 12 lanes, and that NVLink and ECC are off.[5]

How did people work around the limits?

The limit is on the FMA instruction, and a separate multiply and add are not limited. In October 2023, users of the FluidX3D simulation program found that code compiled without FMA ran much faster on the card.[7] The review above measured 6,285 GFLOPS of FP32 without FMA, against 394 with it, and found the card as fast as an A100 in a memory-bound FP32 program.[6] The same idea then went into llama.cpp builds for the card, compiled with `--fmad=false`.[5]

In 2026, a developer published cmpunlocker, a patched NVIDIA kernel module. Its README says that it turns on 64 GB on the 8 GB card and 40 GB on the 10 GB card, and PCIe gen 2.[8] LTT Labs tested it, and an 8 GB card showed 64 GB.[9] The community wiki lists two boards with the name CMP 170HX: an 8 GB board (PCI device id 20c2) and a 10 GB board (device id 2082) with five memory stacks and a 5120-bit bus.[10]

Which boards did we measure?

We rented two cards from an online GPU rental service. They report the name "NVIDIA CMP 170HX". They are not the stock 8 GB card that the review measured.

We did not test the 8 GB card.

	The boards that we rented	The stock card (review)
Device id	0x2082	20c2
Memory	40 GB HBM2e (40,960 MiB), 5120-bit bus	8 GB HBM2e
Memory read (our test)	1,821 GB/s	1,355 GB/s
PCIe	gen4 x16, 26 GB/s to the card	gen1 x4, about 1 GB/s
FP32 with FMA	18.5 TFLOPS: no limit	394 GFLOPS
SMs	108 on one board, 94 on another	70
Power	250 W	250 W

Before you rent or buy a CMP 170HX, run `nvidia-smi` and a short speed test, and check the memory, the PCIe link and the FMA rate.

How fast is Kindi on it?

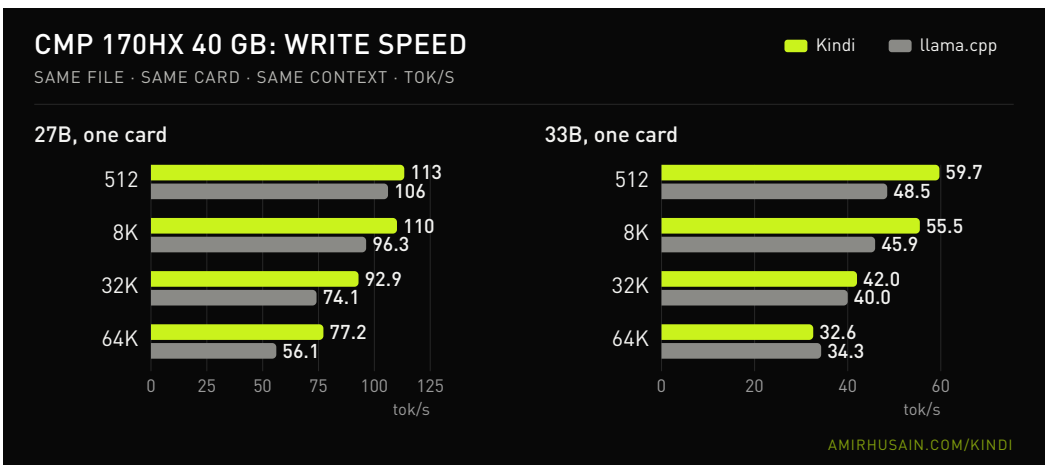

Figure 13-1. *Write speed on one CMP 170HX 40 GB, the 27B and the 33B. Kindi is lime and llama.cpp is grey.*

Setup	Context (tokens)	Kindi read	llama.cpp read	Kindi write	llama.cpp write
27B, one card	512	**1,785**	1,398	**113**	106
	8K	**1,607**	1,381	**110**	96.3
	32K	**1,249**	1,219	**92.9**	74.1
	64K	956	1,062	**77.2**	56.1
27B, two cards	512	1,350	1,405	**121**	107
	8K	1,264	1,373	**119**	97.8
	32K	1,030	1,208	**108**	76.1
	64K	820	1,039	**93.4**	58.2
33B, one card	512	**1,170**	1,169	**59.7**	48.5
	8K	**1,251**	1,141	**55.5**	45.9
	32K	**1,013**	957	**42.0**	40.0
	64K	**799**	788	32.6	34.3
33B, two cards	512	**1,519**	1,139	**64.0**	53.1
	8K	**2,177**	2,046	**59.1**	50.2
	32K	**1,869**	1,793	**45.0**	43.2
	64K	**1,505**	1,494	35.3	37.0

Conditions. One or two NVIDIA CMP 170HX boards, 40 GB HBM2e, device id 0x2082, VBIOS 92.00.6D.00.0A, 250 W, PCIe gen4 x16, on a rented machine. Kindi 0.5.0 (v0.5.0-dev f6ca4df): Kindi reads the prompt to the depth, then writes 3 runs of 32 tokens; the speed is the mean. llama.cpp: Prism fork c96baf4, `llama-bench -ngl 99 -fa 1` (write: 64 tokens at the depth, 5 runs; read: 3 runs). The same model files for both. K/V cache: q8_0 for the 27B, f16 for the 33B. `kindi verify` passes on every row (27B on two cards: 3 exact, 1 llama.cpp tie; 33B: 8 exact, 3 ties). Kindi's number is in bold where Kindi is faster. 2026-11-06.

Kindi writes the 27B faster at every depth, on one card and on two cards. On one card, it writes 113 tok/s at 512 tokens, against 106 for llama.cpp, and 77.2 against 56.1 at 64K tokens. On two cards, it writes 93.4 tok/s at 64K tokens, 1.6 times as fast as llama.cpp. Kindi writes the 33B faster up to 32K tokens: 64.0 against 53.1 tok/s at 512 tokens on two cards.

On one card, Kindi reads the 27B faster up to 32K tokens: 1,785 against 1,398 tok/s at 512 tokens. On two cards, it reads the 33B up to 1.33 times as fast.

What makes it fast on this card?

On the CMP 170HX, the memory is not the limit for the 27B. The limit is the work to unpack the ternary weights. Kindi stores five ternary digits in each byte, and each product must unpack them first. On sm_80 cards with 20 GB or more, Kindi 0.5.0 keeps a second copy of the weights with 2 bits for each digit (the b2 copy). The products read this copy, and each one needs one logic operation in place of the five-digit unpack. The results are bit-identical. The copy needs 1.29 times the memory of the weights, and a 40 GB card has the room for it. `KINDI_TC_B2=0` turns it off.

> *Amazing vid, new here and subbed. Can you stack multiple of these on one like 6x 8GB*
>
> RedRibbonVios (YouTube), 2026-10-26

Sources

1. Nathaniel Mott, "Nvidia Limits RTX 3060 Hash Rate, Unveils New 'Cryptocurrency Mining Processor' Line of GPUs", Tom's Hardware, 2021-02-18, with quotes from the NVIDIA blog post of the same day. tomshardware.com/news/nvidia-announces-cryptocurrency-mining-processor-gpu-line
2. NVIDIA, "NVIDIA CMP HX: Dedicated GPU for Professional Mining", product page, read 2026-11-06. nvidia.com/en-us/cmp
3. Francisco Pires, "Nvidia's CMP170HX in the Wild With 164 MHps of Ethereum Mining Performance", Tom's Hardware, 2021-09-02. tomshardware.com/news/nvidias-cmp170hx-in-the-wild-with-164-mhps-ethereum-mining-performance
4. Zhiye Liu, "Nvidia's Ampere Cryptomining Monster GPU Goes on Sale for $4,300", Tom's Hardware, 2021-10-22. tomshardware.com/uk/news/nvidia-ampere-cryptomining-monster-gpu-sale-4300-dollars
5. Consensus-Protocol, "What Is This Card", cmp170hx wiki on GitHub, edited 2026-07-28. github.com/Consensus-Protocol/cmp170hx/wiki/Start-What-Is-This-Card
6. niconiconi, "All GB/s without FLOPS: Nvidia CMP 170HX Review, Performance Lockdown Workaround, Teardown, Watercooling, and Repair", 2023-10-25, CC BY 4.0. niconiconi.neocities.org/tech-notes/nvidia-cmp-170hx-review
7. Consensus-Protocol, "Appendix: External Sources", cmp170hx wiki, edited 2026-07-31 (FluidX3D issue 8, 2023-10-27). github.com/Consensus-Protocol/cmp170hx/wiki/Appendix-External-Sources
8. amoghmunikote, cmpunlocker README, GitHub, GPL-2.0, published 2026-07-14. github.com/amoghmunikote/cmpunlocker
9. GIGAZINE, report on cmpunlocker and the tests of LTT Labs, 2026-09-14 (in Japanese). gigazine.net/news/20260914-cmpunlocker
10. Consensus-Protocol, "Hardware: Board And Variants", cmp170hx wiki. github.com/Consensus-Protocol/cmp170hx/wiki/Hardware-Board-And-Variants

CHAPTER 14

One small box for everything

The chapters before this one are about graphics cards. This chapter is about a computer with no graphics card at all: a mini PC with an AMD Ryzen 9 7940HS. The graphics unit of this chip, the Radeon 780M, uses the main memory of the computer. With 64 GB of memory, it holds the 27B with room to spare, and the CPU next to it can run a second model.

	Our test computer
Computer	Minisforum Venus UM790 Pro mini PC, 0.8 litres
Chip	AMD Ryzen 9 7940HS: 8 cores, 16 threads, Radeon 780M graphics (RDNA3)
Memory	2 × 32 GB DDR5-5600 SODIMM (64 GB, the largest that it takes)
Storage	Samsung 990 PRO 2 TB NVMe
Network	2.5 Gb Ethernet, Wi-Fi 6E, two USB4 ports
Power	one power adapter for the whole computer
Price	US$879 with 32 GB and 1 TB at the Minisforum store on 2026-10-24

Other makers sell mini PCs with the same Radeon 780M, from the Ryzen 7 7840HS, the Ryzen 9 7940HS and similar chips. Before you buy one for the 27B, check the memory: we measured the box with 64 GB.

What can one box do at the same time?

We ran three things on the box at the same time:

- the 27B on the Radeon 780M, with a context of 24,576 tokens;
- pi, the agent, which did tasks of the Kindi benchmark with the 27B;
- a second Kindi server with the 2B on the CPU, for short text jobs.

Measure	Alone	Together
27B on the 780M, read	88.0	75.3
27B on the 780M, write	10.5	9.21
2B on the CPU, read	62.6	15.5
2B on the CPU, write	42.5	7.8

> **Conditions.** Kindi 0.3.0 development build (849840a). The 27B on the 780M with a context of 24,576; the 2B on the CPU with 8 threads. Each speed: a 2,000-token prompt and 256 tokens out, the median of 3 runs. "Together": the other engine has a request loop running. Then 10 tasks of the Kindi benchmark with pi on the 27B, while the CPU 2B was kept busy. 2026-10-25.

The 27B keeps most of its speed when the CPU 2B works at the same time: it writes 9.2 tok/s, against 10.5 when it runs alone. The 2B loses most of its speed: it writes 7.8 tok/s, against 42.5 alone. pi and the 27B passed all 10 tasks while the 2B was busy, with a median of 59 seconds for each task. The 2B answered 21 requests in that time.

The graphics unit and the CPU read from the same memory. A model writes each token by reading all of its weights from memory, so when both engines write at the same time, they share the memory speed of the box. Each one then writes a little slower than when it runs alone.

How does it compare with rainstorm?

Rainstorm (A worked example: Rainstorm) and this box are two ways to run the 27B for little money.

	Rainstorm	The UM790 Pro
What you buy	an HP EliteDesk 800 G4 SFF and two Quadro K1200 cards	a Minisforum UM790 Pro mini PC
Cost	US$221.58 for all the parts	US$879 (32 GB) at the store; we tested 64 GB
The 27B writes	12.6 tok/s	10.6 tok/s at 512 tokens
Context of the 27B	24,576	24,576
At the same time	two models, one on each card	a second model on the CPU
Size	a small office computer	0.8 litres

Rainstorm costs less, and it writes the 27B a little faster. The mini PC is new, small and quiet, it runs from one adapter, and its memory leaves room for a second model and an agent next to the 27B.

What does the box need?

- **Linux and Docker.** The Kindi ROCm image holds ROCm, CuPy and Kindi. You do not install ROCm on the computer. The computer needs only the `amdgpu` driver, which is part of Linux.
- **One kernel option.** On the Radeon 780M, the driver can stop a running GPU program to let another one run, and on this chip that stop can fail and reset the graphics. In our tests, long prompts of 8,000 tokens and more made it happen, with Kindi and with llama.cpp. The kernel option `amdgpu.cwsr_enable=0` turns off this kind of stop, and with it the same tests passed three times out of three. The installer tells you what the option does and asks before it adds it. You then restart the computer one time.
- **Memory.** We measured the box with 64 GB, the most that it takes. The 27B and its context, the CPU 2B and pi ran together in it.

A brick of intelligence for the home

Put the box on a shelf next to the internet router. It is about 13 cm on each side and 5 cm high, smaller than a hardback book, and it costs about US$900. It connects

with one network cable and one power adapter, and it stays on all day. It does not send your questions or your documents to a cloud service, so the work of the household stays in the house.

This one box holds all the parts of a home AI system: a large model for agent work, a small model for fast decisions, the agent that uses them, and the memory for all three. The other devices in the house do not need a model of their own. A laptop, a phone, a home-automation hub or a script on another computer sends its requests to the box on the local network, in the same OpenAI format that many programs already use. The box is the centre, and the devices around it are its hands and its screens.

A day with the box can look like this:

- **In the morning,** the 2B on the CPU sorts the mail and the messages of the night into short labels (Fast decisions on short text). Only the messages that need an answer go to the 27B, and pi writes a draft for each of them.
- **During the day,** a person asks many questions about one long document: a lease, an insurance policy or the manual of a new appliance. With the prompt cache, Kindi reads the document one time and then reads only each new question (What does the prompt cache do?).
- **In the evening,** a student works with pi and the 27B in a project folder. pi reads the files, runs the scripts and writes the report. The cache keeps the conversation of each person, so two people at two computers do not push each other out.
- **At night,** the 2B reads the text alerts of cameras, sensors and computers, and gives each one a single word: ignore, note or escalate. Only the escalations wake the 27B, which writes a short summary for the morning. Can one sentence make escalation safe? shows how to make the 2B lean the safe way.

One box also means less to look after. There is one computer to update, one to back up and one adapter to plug in. When the internet connection stops, the assistant continues to work, because it never needed the connection. When the household needs more, a second box or a computer with a card can join the network, and the installer sets up each one in the same way.

Can it travel?

The box is small, it has no moving parts other than its fans, and it runs on direct current from its adapter. So it can also go where the internet does not: in a camper van, on a boat or in a car. There, it gives the same assistant with no mobile signal. It can answer questions about the manual of the vehicle, keep a log of the trip and help a driver or a passenger with notes, plans and messages to send later.

We have not tested the box in a vehicle. If you try it, plan for these points:

- **Power.** A vehicle gives about 12 V. Use a DC-to-DC converter to the voltage of the box's adapter, or a small inverter for the adapter itself. Connect it so that the box stops when the engine stops, or it can drain the battery.
- **Heat and air.** Keep the vents open, and do not put the box in direct sun or in a closed compartment.
- **Movement.** Attach the box firmly, and use an NVMe drive (as in our box), not a disk with moving parts.

CHAPTER 15

Kindi on AMD graphics

Kindi runs on three AMD graphics units through ROCm, AMD's GPU software: the Radeon 780M (RDNA3, gfx1103), the Radeon 840M (RDNA 3.5, gfx1152) and the Radeon 8060S of Strix Halo (RDNA 3.5, gfx1151). Each one is part of the processor and uses the main memory of the computer. This chapter gives the 840M and 8060S results, the settings for each unit and the AMD fixes in Kindi 0.5.0.

> *Another excellent contribution. Thank you. Do you think this will scale well to an igpu i780M on a Ryzen 9 8945HS with 64gb of shared ddr5 memory? I've managed to get Qwen 3.6 35B-A3B at Q8 running at 33 TPS with vulkan over llama CPP.*
>
> nicklockard (YouTube), 2026-10-22

Which AMD graphics does Kindi run on?

	Radeon 780M	Radeon 840M	Radeon 8060S (Strix Halo)
Processor	Ryzen 7 7840HS, Ryzen 9 7940HS and similar	Ryzen AI 5 340	Ryzen AI MAX+ 395
Architecture	RDNA3, gfx1103	RDNA 3.5, gfx1152	RDNA 3.5, gfx1151
Compute units	12	4	40
Memory	DDR5, 128-bit	DDR5, 128-bit	LPDDR5X, 256-bit
Memory read (our test)	about 83 GB/s	84 GB/s	240 GB/s
Kindi image	`kindi-rocm`: ROCm 6.3.3, with `HSA_OVERRIDE_GFX_VERSION=11.0.2`	`kindi-rocm7`: ROCm 7, HIP 7.16, no override	`kindi-rocm7`: ROCm 7, HIP 7.16, no override
Models	27B, Nemotron, 2B, and the 9B in 0.5.0	2B, 4B, 9B, 27B	2B, 4B, 9B, 27B
Since	Kindi 0.3.0	Kindi 0.5.0	Kindi 0.5.0

The 780M and the 840M read memory at about the same speed, so they write a large model at about the same speed. The 780M has three times the compute units, so it reads prompts faster. The 8060S reads memory about three times as fast as the other two. One small box for everything gives the 780M results. The 33B backend of Kindi runs on NVIDIA cards only, so no AMD unit runs the 33B.

How do you install Kindi on RDNA 3.5?

Use the same installer as on the other computers. On a Radeon 840M or a Radeon 8060S, the installer selects the `kindi-rocm7` image (ROCm 7, HIP 7.16) and offers the 27B, the 9B, the 4B and the 2B. The 27B needs 12 GB or more of graphics memory. The installer also accepts Strix Point (gfx1150), and it uses the same image there. We have not measured Strix Point.

How fast is the Radeon 840M?

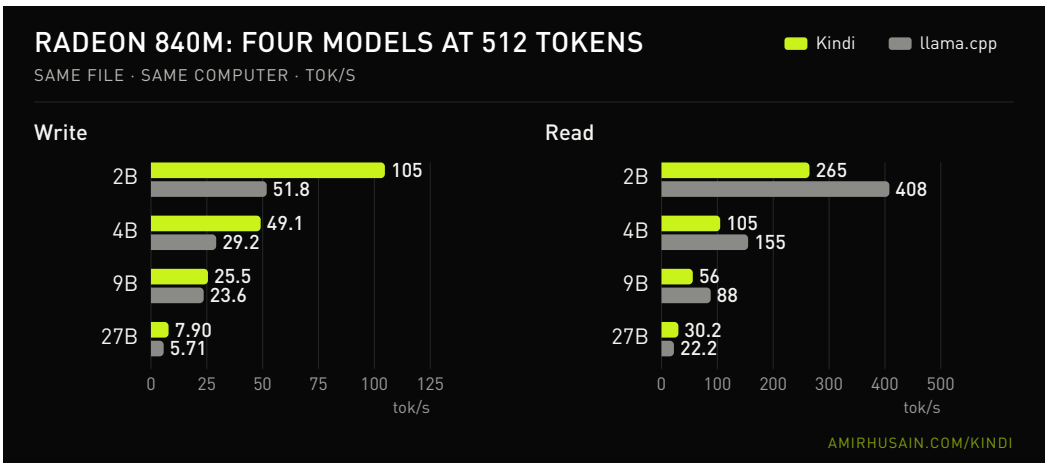

Figure 15-1. *Write and read speed on the Radeon 840M at 512 tokens, the same files for both engines. Kindi is lime and llama.cpp is grey.*

Model	Context (tokens)	Kindi read	llama.cpp read	Kindi write	llama.cpp write
2B	512	265	408	**105**	51.8
	8K	243	354	**91.5**	49.0
4B	512	105	155	**49.1**	29.2
	8K	97	135	**41.7**	26.6
9B	512	56	88	**25.5**	23.6
	8K	53	81	**23.5**	21.9
27B	512	**30.2**	22.2	**7.90**	5.71
	8K	**27.7**	20.8	**7.34**	4.42

Conditions. A rented computer with an AMD Ryzen AI 5 340: Radeon 840M, gfx1152, 4 compute units, shared DDR5-5600-class memory on a 128-bit bus, measured read 84 GB/s; Ubuntu 24.04 container with 2 GiB of host memory; ROCm with HIP 7.16 and CuPy 14.2.0. Kindi: build rdna35-340 (cb19bc1), merged into Kindi 0.5.0; Kindi reads the prompt to the depth, then writes 3 runs of 32 tokens; the speed is the mean. llama.cpp: Prism fork c96baf4, HIP build for gfx1152, `llama-bench -fa 1`. The same model files for both. K/V cache: q8_0 for the 27B on both engines; for the 2B, 4B and 9B, Kindi's single-card engine and llama.cpp's default f16. `kindi verify`: all prompts token-exact on all four models. Kindi's number is in bold where Kindi is faster. 2026-11-05 and 2026-11-06.

Kindi writes faster than llama.cpp on all four models, at 512 and at 8K tokens. It writes the 2B at 105 tok/s, and llama.cpp writes it at 51.8. It writes the 27B at 7.90

tok/s, and llama.cpp at 5.71. Kindi reads the 27B faster (30.2 against 22.2 tok/s).

Kindi's write uses 52% to 60% of the measured memory speed on this chip.

How fast is the Radeon 8060S?

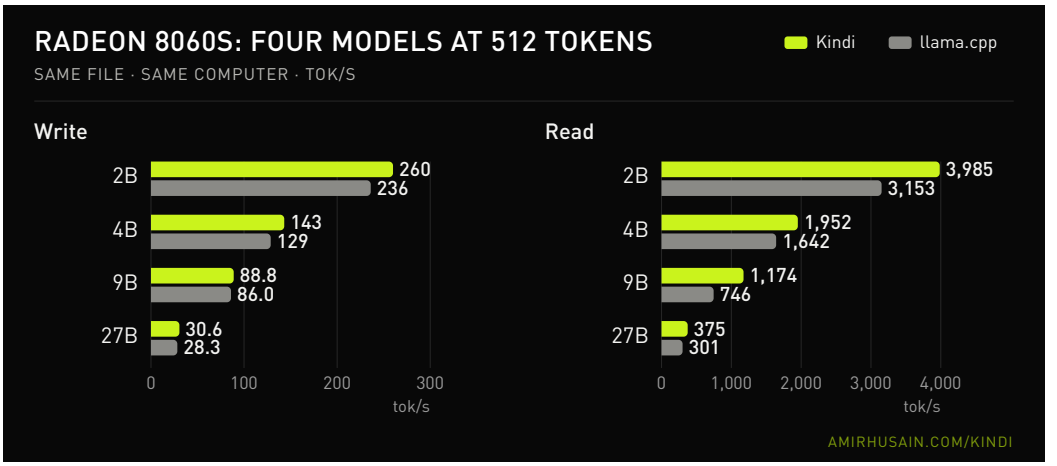

Figure 15-2. *Write and read speed on the Radeon 8060S (Strix Halo) at 512 tokens, the same files for both engines. Kindi is lime and llama.cpp is grey.*

Model	Context (tokens)	Kindi read	llama.cpp read	Kindi write	llama.cpp write
2B	512	**3,985**	3,153	**260**	236
	8K	**3,488**	3,045	**235**	211
4B	512	**1,952**	1,642	**143**	129
	8K	**1,674**	1,541	**122**	111
9B	512	**1,174**	746	**88.8**	86.0
	8K	**1,076**	725	**80.4**	77.6
27B	512	**375**	301	**30.6**	28.3
	8K	**333**	292	**28.4**	24.3

> **Conditions.** A rented computer with an AMD Ryzen AI MAX+ 395: Radeon 8060S, gfx1151, 40 compute units, 64 GiB of graphics memory, 256-bit LPDDR5X, measured read 240 GB/s; Ubuntu 24.04 container; ROCm with HIP 7.16 and CuPy 14.2.0, no `HSA_OVERRIDE_GFX_VERSION`, `cwsr_enable=1`. Kindi 0.5.0 (v0.5.0-dev 38fcef2); Kindi reads the prompt to the depth, then writes 3 runs of 32 tokens; the speed is the mean. All four models use Kindi's plan engine. llama.cpp: Prism fork c96baf4, HIP build for gfx1151, `llama-bench -fa 1`. The same model files for both. K/V cache: q8_0 for the 27B on both engines; llama.cpp's default f16 for the 2B, 4B and 9B. `kindi verify`: all prompts token-exact on all four models. Kindi's number is in bold where Kindi is faster. 2026-11-07.

On the 8060S, Kindi writes and reads faster than llama.cpp on all four models, at 512 and at 8K tokens. It writes the 2B at 260 tok/s, against 236 for llama.cpp, and the 27B at 28.4 against 24.3 tok/s at 8K tokens. It reads the 9B 1.57 times as fast: 1,174 against 746 tok/s at 512 tokens.

Which settings does each unit use?

Kindi sets its GEMV kernels (the products that write each token) for each AMD graphics unit. The 780M, the 840M and the 8060S each have a fixed set that we measured. On other AMD graphics, Kindi times the settings when it starts. On the 8060S, Kindi also waits for the graphics with a spin wait, because the blocking wait of HIP woke late at random and made the 9B write at 21 to 41 tok/s. On AMD, a row-parallel product uses 8 column splits at most, because 16 splits changed the output of the 27B on the 840M.

Setting	Radeon 780M	Radeon 840M	Radeon 8060S
`HSA_OVERRIDE_GFX_VERSION`	`11.0.2` (the 780M runs ROCm's gfx1102 code)	not set	not set
`HSA_ENABLE_SDMA`	`0`	the default	the default
Kernel option	`amdgpu.cwsr_enable=0`	the default (on)	the default (on)
GEMV settings	the 780M set	the 840M set	the 8060S set

On the 780M, the driver can stop a GPU program to let another one run (CWSR), and that stop can fail and reset the graphics. The option `amdgpu.cwsr_enable=0` turns it off (What does the box need?). On AMD, Kindi also turns off transparent huge pages for its own process, because page compaction was one trigger of the same reset. We ran the 840M and the 8060S with the default setting. The Kindi runs

finished. During one long llama.cpp prompt run, the connection to the rented computer stopped for some hours, and we do not know the cause.

AMD RDNA2 graphics

Kindi 0.7.0 runs on AMD RDNA2 graphics: the Radeon PRO W6000 and Radeon RX 6000 chips, gfx1030 to gfx1036. RDNA2 has an 8-bit dot product for two signed values, but not the mixed signed and unsigned form of RDNA3. Kindi uses the signed form there, and it is exact for Kindi's weights. On gfx1030, the 2B, 4B and 9B run on the plan engine.

We measured a Radeon PRO W6800 (gfx1030, 60 compute units, 32 GB). `kindi verify` passes on the 2B, the 9B and the 27B, token for token. With the 27B, Kindi writes and reads faster than llama.cpp on the same card, with the same file and a q8_0 K/V cache for both engines:

Model	Context (tokens)	Kindi read	llama.cpp read	Kindi write	llama.cpp write
27B	512	**573**	265	**42.3**	35.0
	8K	**506**	263	**38.8**	33.4

On the same card, Kindi writes the 2B at 254.9 tokens/s and the 9B at 106.1 tokens/s at 512 tokens:

Model	Context (tokens)	Kindi read	Kindi write
2B	512	**4,479**	**255**
	8K	**4,215**	**235**
9B	512	**1,635**	**106**
	8K	**1,490**	**95.7**

> **Conditions.** One Radeon PRO W6800 (gfx1030, 32 GB) with a Ryzen 7 3800XT on a rented machine, HIP 7.16 and CuPy 14.2.0 on the host, context 9,216, the same GGUF files for both engines. Kindi: read to the depth, then 3 x 32 greedy tokens at each depth. llama.cpp: the PrismML fork (c96baf4, ROCm), `llama-bench -ngl 99 -fa 1`. The 27B with a q8_0 K/V cache for both engines. Kindi 0.7.0 development build (f312986). 2026-11-21.

The installer accepts gfx1030 to gfx1036 and uses the `kindi-rocm7` image (ROCm 7). It offers the Ternary-Bonsai-2 models on these cards.

What did Kindi 0.5.0 fix on AMD?

- **The 9B on AMD graphics.** Before 0.5.0, the 9B gave wrong output on AMD (`kindi verify` failed). A barrier in the embedding kernel was inside a condition, so on AMD the waves did not wait for each other. Fixed in 0.5.0: the 9B is token-exact on the Radeon 780M and on the 840M.
- **Prompts on the 780M.** When the weight cache had no lane-layout copy, the 27B read each prompt one token at a time. Fixed in 0.5.0: Kindi makes the missing copy, and the 27B reads 105 tok/s.
- **Computers with little memory.** Kindi copies the weights to the graphics one tensor at a time, so the host does not keep a full copy. The 840M computer had 2 GiB of host memory, and it loaded the 9B and the 27B.

We ran Kindi 0.4.0 and Kindi 0.5.0 on the 780M with the same files. The 2B, 4B and 27B pass `kindi verify` on both, and the 9B passes on 0.5.0 only. The read and write speeds of the two builds are the same within the run to run variation of the 780M.

CHAPTER 16

What does the prompt cache do?

An agent sends the same text again and again. Each request starts with the system prompt, the list of tools and all the earlier turns, and only the end of the request is new. A server without a cache reads all of it again for each request. On a small card, that reading is most of the wait.

Kindi keeps the work that it did for one request and uses it again for the next request. When a request starts with the same tokens as an earlier one, Kindi reads only the new part. This is the prompt cache. It does not change what the model sees. With the cache, Kindi reads the new part of a prompt in other batch sizes, so the arithmetic can round a little differently. On a near tie between two tokens, a reply can then take the other token. In 295 turns of real agent conversations with Nemotron at temperature 0, 283 replies were the same with the cache on and off, and each of the other 12 took another reasonable path, for example `cat` in place of `read`.

When does the cache help?

Job	What stays the same	What Kindi reads again
A chat or an agent	the system prompt and the earlier turns	only the new message
Regenerate an answer	the full prompt	only the last token
Many questions about one document	the document	only the question
Two or more agents that take turns	the conversation of each agent	only the new message of each agent

How much faster is it?

These are the times of real requests with the cache on. Without the cache, Kindi reads the whole prompt again for each request.

Job	Request	Nemotron, seconds	27B, seconds
Agent, 2.2K-token system prompt	turn 2	**0.67**	**2.52**
	turn 3	**0.86**	**2.93**
Regenerate an answer	again	**0.43**	**1.85**
Questions about a 9K-token document	question 2	**0.7**	**14.5**
	question 3	**0.42**	**11.5**
Two agents that take turns	A, turn 2	**0.58**	**2.61**
	B, turn 2	**0.53**	**2.79**

> **Conditions.** GeForce GTX 1080 8 GB. Nemotron 3.5 Bolt E4B (UD-Q3_K_XL, sha256 3f8032b9…) and Ternary-Bonsai-2 27B (PTQ1_0, sha256 53107f53…). One server, the cache on. The times are for the full request: the prompt and a reply of up to 60 tokens (400 for the document questions, 600 for the regenerated answer). The replies were the same in all 12 requests of each model. Script: `examples/prompt_cache_bench.py`. Kindi 0.3.0 development build, 2026-10-24.

With Nemotron, the second turn of an agent comes back in 0.67 s, a regenerated answer in 0.43 s, and the third question about a 9,400-token document in 0.42 s. On the document questions, the 27B continues from its last checkpoint at token 6,912 and reads the rest of the document again.

The first request of a conversation reads the whole prompt, because there is nothing in the cache yet.

Is it faster than the cache of llama.cpp?

llama.cpp also has a prompt cache. We ran the same requests against llama.cpp with its cache on, on the same card with the same model file. This table gives the requests that each engine could take from its cache:

Job	Request	Kindi, s	llama.cpp, s	Kindi faster
Agent	turn 2	**0.67**	1.29	1.9×
	turn 3	**0.86**	1.3	1.5×
Regenerate an answer	again	**0.43**	4.93	11.5×
Questions about a document	question 2	**0.7**	3.16	4.5×
	question 3	**0.42**	3.1	7.4×
Two agents that take turns	A, turn 2	**0.58**	1.4	2.4×
	B, turn 2	**0.53**	1.44	2.7×

> **Conditions.** The same card, file and requests as the table above, Nemotron. llama.cpp: CUDA llama-server, build b11804 (the ggml-org server-cuda image of 2026-10-22), its prompt cache on. 2026-10-24. The first request of each job is not in the table: there, neither engine has anything in its cache.

On each cached request, Kindi answers faster. The difference is largest when the whole prompt is the same, as when you regenerate an answer: Kindi reads one token again and takes 0.43 s, and llama.cpp takes 4.93 s.

How does it work?

Kindi keeps three kinds of saved state.

1. **The conversation on the card.** The keys and values of every token stay on the GPU. A new request that starts with the same tokens continues from the end of the shared part.
2. **Checkpoints.** Some layers keep only the last part of the text (a sliding window), and the Bonsai models also keep a recurrent state. When a request goes back further than that, Kindi uses a checkpoint: a copy that it makes in host memory at regular steps of each prompt and at its end.
3. **Parked conversations.** When a new request shares less than half of the conversation on the card, that conversation moves to host memory. When its next turn comes, it moves back to the card. So agents that take turns each keep their own cache.

The Bonsai models can also save the start of a conversation to disk, so it is ready again after a restart. Each parked Nemotron conversation uses about 16 KB of host

memory for each token, plus about 42 MB for each checkpoint. The settings are in `docs/PROMPT-CACHE.md` in the Kindi repository.

CHAPTER 17

What does a power cap change?

A power cap is a limit on the power that a card can use. You set it with `nvidia-smi`. All ten RTX 3080 cards in our cluster run at a cap of 190 W, and their default limit is 380 W. We measured the two engines on one card at both caps.

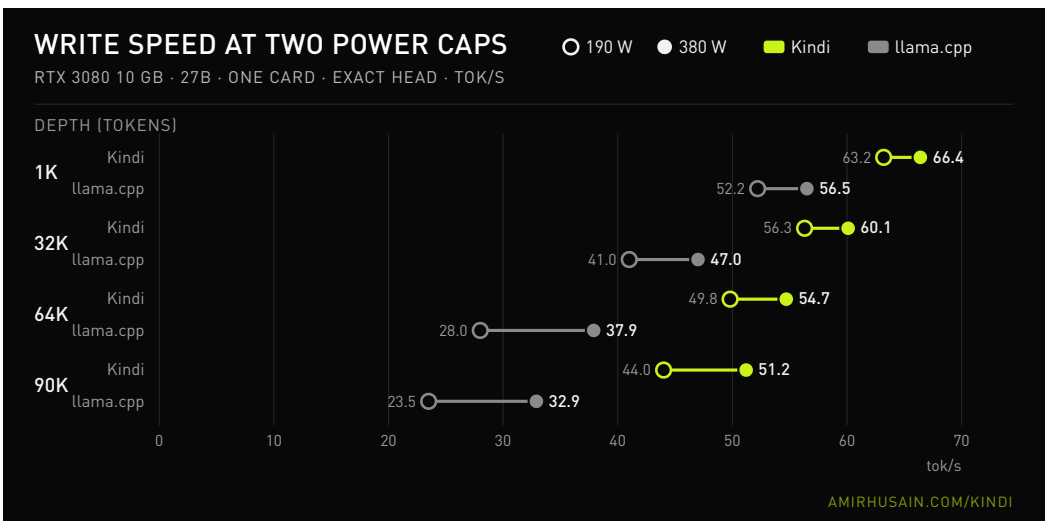

Figure 17-1. *Write speed of the 27B on one RTX 3080 at a 190 W cap (open dot) and a 380 W cap (filled dot), exact head. 190 W: Kindi 0.2.0; 380 W: dev build release-tcspec6 (2026-10-22). Kindi is lime and llama.cpp is grey.*

At half the power, Kindi writes 63.2 tok/s at a depth of 1K, against 66.4 at 380 W. At 90K, Kindi writes 44.0 against 51.2. llama.cpp writes 52.2 against 56.5 at 1K, and 23.5 against 32.9 at 90K. The read speed falls more with the cap. With a 512-token prompt, Kindi reads 1,093 tok/s at 190 W and 1,361 at 380 W.

What does that mean for you?

- **A cap costs little write speed.** To write one token, the card reads all the weights from its memory. So the write speed depends mostly on the speed of the memory, and less on the power of the arithmetic units.
- **A cap costs more read speed.** A long prompt is mostly arithmetic, and the arithmetic slows down at a lower power.
- **A cap lowers the heat and the load on the power supply.** At 380 W, our card used up to 375 W and reached 82 °C.
- **For agent work, a cap near 190 W is a good start.** With the 2B drafter on agent prompts, Kindi writes 99.9 tok/s at 190 W and 108 at 380 W.

To set a cap of 190 W on the first card, type:

```
$ sudo nvidia-smi -i 0 -pl 190
```

The cap stays until the next restart of the computer. To keep it, run the command again at each start, for example from a systemd service.

CHAPTER 18

How do you start on a new computer?

Kindi has a one-command installer. It finds your NVIDIA cards and your AMD graphics, shows each setup that your computer can run, with the best one first, and then downloads, checks and starts everything.

What do you need?

- A Linux computer (x86-64) with an NVIDIA card from this manual, AMD RDNA3 graphics such as the Radeon 780M, or no graphics: then the installer offers the CPU engine. On Windows 10 and 11, one PowerShell command runs this installer inside WSL2 (Kindi on Windows).
- For NVIDIA cards, the NVIDIA driver. Kepler cards such as the K420 need the 470 driver.
- For AMD graphics, only the `amdgpu` driver, which is part of Linux. ROCm is inside the Kindi image.
- Docker. If Docker or the NVIDIA Container Toolkit is missing, the installer asks you before it installs it. The installer also works with rootless Docker.
- Free disk space for the model: from 429 MB for the 2B to 5.95 GB for the 27B. Nemotron needs 4.2 GB, Lightning 9.47 GB, and the 33B 13.1 GB. The DSpine drafter of the 27B adds 4 GB.

How do you install it?

First do a dry run. The dry run shows each step, each command and the list of setups, and it changes nothing:

```
$ curl -fsSL https://amirhusain.com/kindi/install.sh | bash -s -- --dry-run
```

Then install:

```
$ curl -fsSL https://amirhusain.com/kindi/install.sh | bash
```

To check the script before it runs, download it, compare its sha256 with the value in `docs/INSTALL.md`, read it, and then run it with `bash install.sh`.

How do you select a setup?

The installer shows a list. On a computer with an RTX 3060 and a GTX 1080, the list starts like this:

```
[3] Choose the setup
    1) every GPU, one server each: 27B on GPU 0 + 27B on GPU 1 (recommended)
    2) Ternary-Bonsai-2 27B on GPU 0 (NVIDIA GeForce RTX 3060)
    3) NVIDIA Nemotron 3.5 Bolt E4B on GPU 0 (NVIDIA GeForce RTX 3060)
    ...
    7) Ternary-Bonsai-2 27B on GPU 1 (NVIDIA GeForce GTX 1080)
    ...
    12) Ternary-Bonsai-2 2B on the CPU (80 threads; slow, no GPU needed)
    Choose [1]:
```

- Type one number for one server.
- Type two or more numbers, for example `2,8`, for two servers side by side: one model on each card.
- Type `5x2` to run setup 5 on two cards of the same kind, for example a 4B on each of two K1200 cards.

Each server gets its own container (`kindi`, `kindi-2`, …) and its own port, from 8093 up. A card can have only one server.

To install with no questions, add `--auto`: the installer takes the recommended setup and continues. To select a setup in a script, add `--config` and the numbers, for example `--config 2,8`.

The installer changes the kernel only when you answer at the keyboard. The gpu-unlock module for RTX 30 cards and the kernel option for the Radeon 780M both need your answer, also with `--yes` and `--auto`.

What does the installer do?

Step	What it does
1. Check the computer	Finds Linux, the memory, the free disk, the driver, and each NVIDIA and AMD graphics unit
2. Docker and GPU access	Finds Docker and the NVIDIA Container Toolkit or the AMD devices, and asks before it installs anything
3. Select the setup	Shows each setup for your computer, the best one first (above)
4. Build the image	Builds the Kindi image (one CUDA image for all NVIDIA cards, ROCm for AMD graphics) and checks that Kindi loads in it
5. Download the models	Downloads each file and checks its sha256
6. Convert and verify	Runs `kindi convert`, then `kindi verify` one time for each model on each kind of card
7. Start the servers	Starts each container on `127.0.0.1` and waits until it answers

The recommended setup for each kind of computer:

Your computer	Recommended	Context (tokens)
No graphics	2B on the CPU	24,576
Quadro K420 1 GB	2B	24,576
Quadro K420 2 GB	4B	23,808
One Quadro K1200	9B	fitted to the free memory
Two Quadro K1200 cards	27B, split over both cards	24,576
GeForce GTX 1080	27B	65,536
GeForce RTX 3080	27B, Tensor Cores and the DSpine drafter	fitted to the free memory
Radeon 780M	27B	24,576
Two or more cards that do not split one model	the best model on each card, one server each	as above, for each card

To see only the setups with one model, add `--model` and `2b`, `4b`, `9b`, `27b`, `nemotron`, `lightning`, `nanbeige` or `flashnext`. To place the servers on your cards yourself, use `--place` (Which card runs which model?). To select a context, add `--ctx` and a number. The servers listen on 127.0.0.1 only, and the installer asks for no keys. `--lan` also serves your local network, with an API key (What does the server do for an agent?). `--with-bridge` also installs Kindi Bridge (Kindi Bridge). DeepSeek-V4.1-Flash also has its own installer (DeepSeek-V4.1-Flash on one RTX 3060).

At the end, the installer shows the address, the model and the context of each server, and the settings for pi. Then continue with Meet pi. Use the model `id` and the `contextWindow` that the installer shows in the pi settings of that chapter.

To remove Kindi, type `bash install.sh --uninstall`. This removes every server that the installer started. Your models stay in `~/.kindi`.

On NVIDIA cards of compute capability 8.0 and later, the installer turns on a drafter for the 27B: DSpine on one card with the tensor-core path (DSpine, the block drafter of the 27B), else the 2B drafter. It also downloads the MTP head of the 2B (How does a draft model make the 27B faster?). On these cards, with about 17 GB of card memory in total, the installer also offers the 33B (A more capable model).

What are the first steps after the install?

1. Find the port of the server. At the end, the installer shows the port of each server. The first server uses port 8093, the second server uses 8094, and so on.
2. Make sure that the server runs. Open `http://localhost:8093/health` in a browser, or type `curl -s localhost:8093/health`. The answer shows the model, the context size and the server state.
3. Connect a chat client. Kindi uses the OpenAI API at `http://localhost:8093/v1`. In Open WebUI, LM Studio, Jan or pi, add an "OpenAI-compatible" connection with this address. The API key can be any text, unless you set an `api_key`. The model id is in `/health`, for example `bonsai2-2b`.
4. Send a test request with curl:

```
$ curl -s localhost:8093/v1/chat/completions     -H "Content-Type:
application/json"     -d '{"model": "bonsai2-2b",
         "messages": [{"role": "user", "content": "Say hello."}]}'
```

1. Or send it with Python, after `pip install openai`:

```
from openai import OpenAI
client = OpenAI(base_url="http://localhost:8093/v1", api_key="none")
r = client.chat.completions.create(
    model="bonsai2-2b",
    messages=[{"role": "user", "content": "Say hello."}])
print(r.choices[0].message.content)
```

If the server is on a different computer, use its address in place of `localhost`. By default, the server accepts connections only from the same computer.

A request can also set `min_p`, `repeat_penalty` and `repeat_last_n`, in addition to `temperature`, `top_p` and `top_k`. These settings work as in llama.cpp, and they are off by default. `min_p` removes each token whose probability is less than that share of the most probable token. `repeat_penalty` makes the tokens of the last `repeat_last_n` tokens less probable, so the model repeats itself less.

CHAPTER 19

Which card runs which model?

A computer can have several cards, and you can run several models on them. Kindi 0.7.0 has three commands for this: `kindi cards` lists the cards, `kindi plan` places the models on the cards, and `kindi route` gives one address for all the servers.

What cards do you have?

`kindi cards` lists every card with:

- its name, compute capability, memory and free memory;
- its PCIe link and its memory bandwidth;
- the driver;
- whether Kindi can use the card, and if not, why.

A card that is too full, or that another program uses, is left out of the plans. The installer asks each NVIDIA card on its own. When the query of one card fails (for example, `nvidia-smi` reports a corrupt infoROM), that card gets a warning, and the other cards stay in the list.

Where should each model run?

`kindi plan` proposes where each model runs, the best placement first. With `--model 27b,2b`, it places those models. The plan puts the fastest card first. When one card does not hold a model, the plan splits the model over two cards by their free memory. DeepSeek-V4.1-Flash gets the fastest card as its main card, and its hot experts use all the spare memory of its cards, the main card included.

How do you place the servers yourself?

Give one `--place` argument to the installer for each server, in the form `MODEL:CARDS`:

```
bash install.sh --place "27b:0" "2b:1"          # a 27B on card 0 and a 2B
on card 1
bash install.sh --place "27b:2,3"               # one 27B split over cards 2
and 3 by their free memory
bash install.sh --place "dsv41:0+" "27b:3"      # DeepSeek with card 0 as
the main card and the other cards
                                                # (not card 3) as expert
cards; a 27B on card 3
bash install.sh --place "dsv41:0+1,2"           # DeepSeek: card 0 is the
main card, cards 1 and 2 hold experts
bash install.sh --place "27b:0@8095"            # a 27B on card 0, on port
8095
```

- `MODEL` is `2b`, `4b`, `9b`, `27b`, `nemotron`, `lightning`, `nanbeige`, `dsv41` or `flashnext`.
- `CARDS` are the card numbers of `kindi cards` (the same numbers as `nvidia-smi`), or the card UUIDs.
- For DeepSeek, the card before `+` is the main card, and the cards after it hold experts. `dsv41:0+` takes every card that no other server takes.
- `@PORT` after the cards sets the port.

The installer checks the placement before it changes anything. A card in two servers, or a model that does not fit its cards, stops it with the reason.

The plan file

The placement goes to `~/.kindi/plan.json`. A setup that you choose from the installer's list goes there too. The file names each card by its UUID, so a card that moves to another slot is still found. With the plan file, no card is in two servers.

Command	What it does
`kindi plan --out FILE`	Writes the proposed plan to a file
`kindi plan --check FILE`	Says whether a plan still fits the cards
`bash install.sh --plan FILE`	Installs the servers of a plan file

The DeepSeek installer also reads its cards from the plan.

One address for all servers

With two or more servers, the installer also starts `kindi route`, the front door. It is one OpenAI endpoint for all the servers, on port 8090 (`--route-port` changes it, and `0` turns it off). The `"model"` of each request picks the server, for example `bonsai2-27b`, `bonsai2-2b` or `deepseek-v4.1-flash`. `/v1/models` lists the models of all servers. A request for a model that no server has gets an HTTP 404. When several servers have the same model, the front door sends the requests to them in turn.

How does Kindi split one model over two cards?

When the 27B runs on two cards, each card normally takes one half of the work. Sometimes the two cards do not have the same free memory, for example when a display or another program uses part of one card. Then the half on the fuller card can leave little room for the context.

Kindi 0.7.0 does this split by itself (`KINDI_TP_SPLIT=auto`, the default). When the halves do not leave room for the context on the card with less free memory, the other card takes a larger part of the feed-forward layers and of the output head, so both cards keep the same room for the context. When the halves fit, Kindi keeps the halves, because the work is then shared equally.

On two RTX 3080 cards with 5 GB taken on one card, the 27B gets a context of 166,656 tokens with the automatic split. `kindi verify` passes with the automatic split.

Setting	What it does
`KINDI_TP_SPLIT=auto`	The default: a larger part for the card with more free memory, when that is necessary
`KINDI_TP_SPLIT=even`	Always the halves
`KINDI_TP_SPLIT=0.6`	Card 0 takes this part of the work (0.1 to 0.9)

CHAPTER 20

Kindi on Windows

Kindi runs on Windows 10 and Windows 11 through WSL2. WSL2 is a small Linux virtual machine that is a part of Windows. One PowerShell command installs everything: it turns on WSL2, makes a Linux system only for Kindi, installs Docker, checks the GPU, and runs the normal Kindi installer inside. Then Windows programs use Kindi at `http://localhost:8093/v1`, the same as on Linux.

Install

Open PowerShell (you do not need "Run as administrator") and run:

```
irm https://amirhusain.com/kindi/install.ps1 | iex
```

To see each step before anything changes, do a dry run first:

```
& ([scriptblock]::Create((irm https://amirhusain.com/kindi/install.ps1))) -
DryRun
```

The installer prints each step and the reason for it. It asks before it changes something.

If Microsoft Defender or your company's security software stops the one-line command: some security tools block every "download and run" PowerShell command. Download the script, check it, and run it as a file instead:

```
Invoke-WebRequest https://amirhusain.com/kindi/install.ps1 -OutFile
install.ps1
Get-FileHash install.ps1        # compare with the sha256 in the release notes
powershell -ExecutionPolicy Bypass -File .\install.ps1
```

After WSL2 is turned on for the first time, Windows can show a "Welcome to Windows Subsystem for Linux" window. You can close it: the install goes on in its own window. Every run writes a log to `%LOCALAPPDATA%\Kindi\install.log`.

Step	What happens
1. Check the PC	An x86-64 PC, Windows 10 2004 (build 19041) or newer, virtualization, RAM, disk, the GPU and its Windows driver, the proxy
2. Turn on WSL2	Only if it is off. Windows asks for administrator approval once, and the PC restarts. The installer continues by itself after you sign in again
3. A Linux system for Kindi	A new WSL2 distribution with the name `Kindi` (Ubuntu 24.04 from Canonical, sha256-checked). Your other WSL distributions do not change. If `Kindi` exists, the installer reuses it
4. Docker and the GPU	If Docker Desktop is installed, Kindi uses it. If not, Docker Engine and the NVIDIA container toolkit go into the `Kindi` distribution. Then the installer checks that the GPU is visible inside WSL2
5. The Kindi installer	`install.sh` runs inside WSL2. It shows the setups for your GPU, the recommended one first, downloads and checks the model, and starts the server
6. Windows items	A check that `http://localhost:8093` answers on Windows, Start-menu items, a desktop shortcut, and an entry in Settings > Apps

Run the same command again to upgrade or repair Kindi. The models that you downloaded are kept.

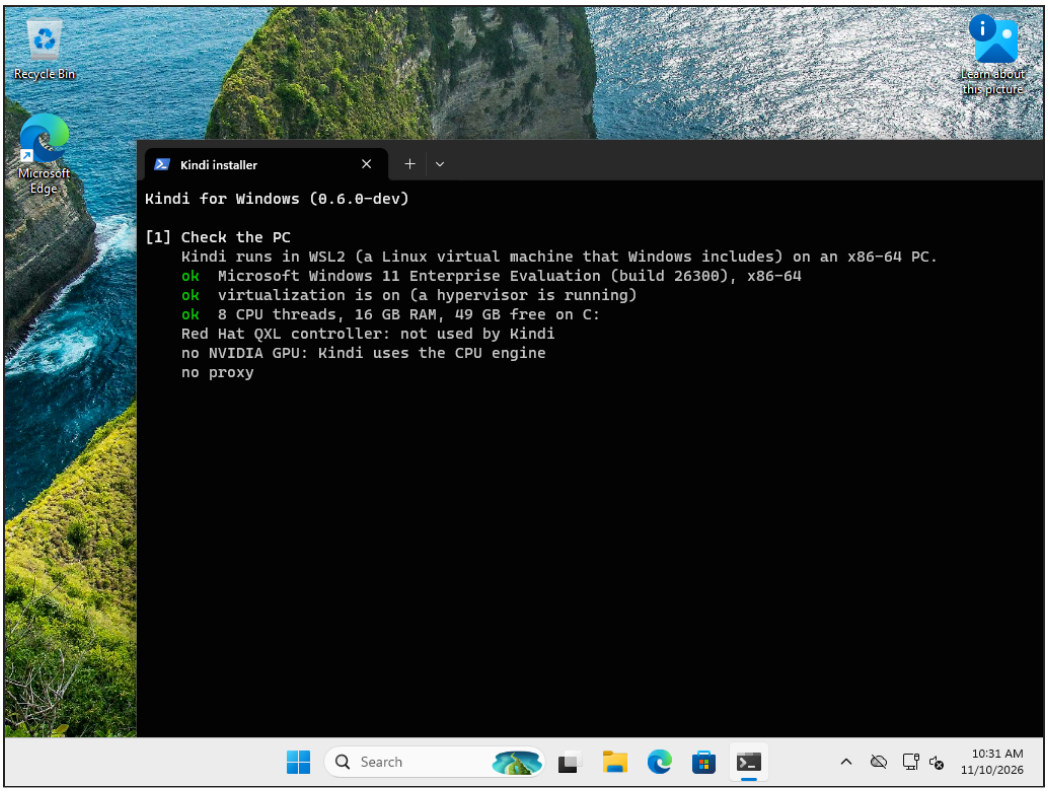

Figure 20-1. *The first step of the Windows installer on a Windows 11 test PC with no NVIDIA card: it checks the PC, and it chooses the CPU engine.*

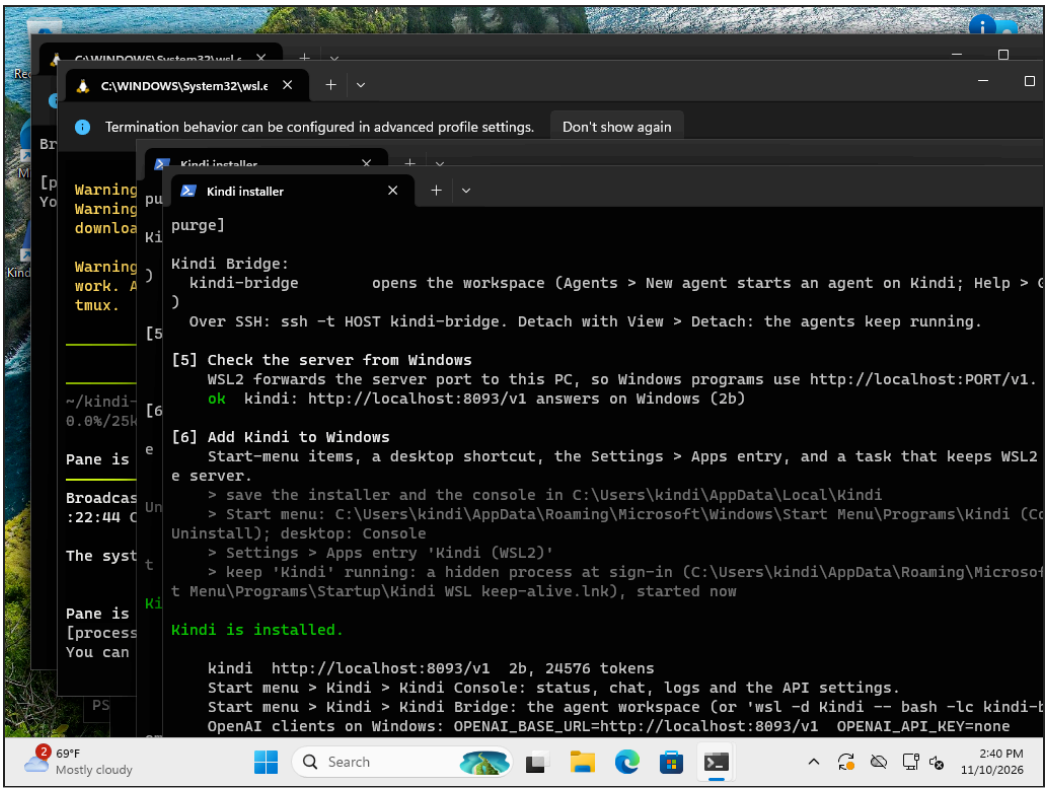

Figure 20-2. *The end of the install on the same PC: Windows reaches the Kindi server at localhost:8093, and the installer adds the Start-menu items.*

What you need

- **An x86-64 PC** (Intel or AMD). Windows on ARM (for example Snapdragon laptops) cannot run Kindi.
- **Windows 10 version 2004 or newer, or Windows 11.** For the GPU in WSL2, you need Windows 11 or Windows 10 21H2 (build 19044) or newer. On earlier builds, Kindi uses the CPU engine.
- **Virtualization on** in the BIOS/UEFI (Intel VT-x or AMD SVM). Most PCs have it on.
- **An NVIDIA GPU with a current Windows driver** from nvidia.com (Game Ready or Studio). Do not install a Linux NVIDIA driver inside WSL2: the Windows driver serves WSL2. Without an NVIDIA GPU, Kindi uses the CPU engine.

- **Disk:** about 15-40 GB, for Linux, Docker, the models and their packs.

AMD GPUs: Kindi's AMD engine needs the Linux amdgpu compute driver (`/dev/kfd`). WSL2 does not have it. On a Windows PC with an AMD GPU, Kindi uses the CPU engine. For the AMD GPU speed, use native Linux.

Use Kindi

- **Kindi Console** (Start menu > Kindi > Kindi Console, and on the desktop): the status of each server, a chat window, the logs, a restart, and the settings for OpenAI clients.
- **Kindi Shell (WSL)**: a Linux shell in the `Kindi` distribution. The `kindi` commands and `docker` work there.
- **Any OpenAI client on Windows:** set the base URL to `http://localhost:8093/v1` and the API key to `none`. Meet pi shows how to connect an agent.

```
$env:OPENAI_BASE_URL = 'http://localhost:8093/v1'; $env:OPENAI_API_KEY =
'none'
```

WSL2 stops an idle Linux system after a short time. With Docker Engine inside the `Kindi` distribution, a small hidden process (`Kindi WSL keep-alive` in the Startup folder) keeps it running, so the server is always ready after you sign in. With Docker Desktop, Docker Desktop keeps the server running.

Options

Add options after `-File install.ps1`, or after the `scriptblock` form above.

Option	What it does
`-DryRun`	Print every step and command, change nothing
`-Check`	Only check the PC and say what the install would do
`-Yes`	Answer yes to every question
`-LinuxArgs '...'`	Options for `install.sh`, for example `'--model 27b'`, `'--gpu cpu'` or `'--config 2'` (How do you start on a new computer?)
`-Distro NAME`	Use another WSL2 distribution. An existing one is reused, never replaced. The installer asks before it adds Docker or turns on systemd in it
`-DistroDir PATH`	Where a new distribution keeps its disk (default `%LOCALAPPDATA%\Kindi\wsl`). Use it to put Kindi on a larger drive
`-DockerMode MODE`	`auto` (default: Docker Desktop if it is installed), `desktop` or `engine`
`-WslMemoryGB N`	Give the WSL2 virtual machine N GB of RAM (see below)
`-WithBridge`	Also install Kindi Bridge, the agent workspace (Kindi Bridge). The installer adds Python 3, tmux and Node.js 22 to the `Kindi` distribution for it, and a Start-menu item "Kindi Bridge" that opens it in Windows Terminal. Without this option, an interactive run asks (the default is no); `-Yes` alone never adds it
`-NoBridge`	Do not install Kindi Bridge and do not ask
`-NoDesktopShortcut`	No desktop shortcut
`-NoReboot`	After WSL2 is turned on, do not restart. You restart, and the install continues after you sign in
`-Uninstall`, `-Purge`	Remove Kindi (see below)

RAM in WSL2

Windows gives the WSL2 virtual machine about half of the PC's RAM by default. The installer shows both numbers. The 27B and the larger models need more. To

give WSL2 more RAM, run the installer with `-WslMemoryGB N`, for example `-WslMemoryGB 48` on a 64 GB PC. This sets `memory=` in `%UserProfile%\.wslconfig` (a copy of the previous file is kept) and runs `wsl --shutdown` once. That command stops all WSL distributions and Docker Desktop, so the installer asks first.

Proxies

The installer uses the Windows proxy setting (including PAC and WPAD), or the `HTTPS_PROXY` variable. It sets the proxy for the downloads, for `apt` and Docker inside WSL2, and for `docker build`. A proxy that runs on the PC's own `localhost` (for example Px or Cntlm) is not visible from WSL2, unless `networkingMode=mirrored` is set in `.wslconfig`. Otherwise, give the proxy's network address in `HTTPS_PROXY`.

Remove Kindi

Use Settings > Apps > Kindi (WSL2) > Uninstall, or Start menu > Kindi > Uninstall Kindi. It removes the Kindi servers and image, the keep-alive process, the Start-menu items, the shortcut and the Settings entry. Then it asks if it should also delete the `Kindi` distribution with its models. With `-Purge`, it deletes them without a question. A distribution that the installer did not make (`-Distro`) is never deleted: only Kindi is removed from it.

```
powershell -ExecutionPolicy Bypass -File "$env:LOCALAPPDATA\Kindi\install.ps1"
-Uninstall
```

Problems

Message	What to do
"this PC has an ARM processor"	Kindi needs an x86-64 PC
"Windows build ... is too old for WSL2"	Settings > Windows Update: install the updates, then run the installer again
"virtualization is off in the BIOS/UEFI"	Restart the PC and open the firmware setup (often F2, F10, F12 or Del during start). Turn on Intel Virtualization Technology (VT-x) or AMD SVM Mode, save, and run the installer again
"the NVIDIA GPU does not reach WSL2"	Install the current NVIDIA Windows driver from nvidia.com and restart. Do not install a Linux driver inside WSL2. For the CPU engine only: `-LinuxArgs '--gpu cpu'`
"Docker Desktop does not answer in 'Kindi'"	Docker Desktop > Settings > Resources > WSL integration: turn on `Kindi`, then Apply & restart. Or use `-DockerMode engine`
"does not answer ... from Windows"	WSL2 forwards `localhost` by default. If `.wslconfig` has `localhostForwarding=false`, remove it and run `wsl --shutdown`
"systemd did not start"	Run `wsl --update` (WSL 0.67.6 or newer is needed), then run the installer again

The Linux installer's own messages are in How do you start on a new computer?.

WSL2 or native Windows?

There are two ways to run Kindi on Windows.

	WSL2 (the installer above)	Native Windows (`kindi.exe`)
What runs	The Linux Kindi, in its own WSL2 distribution, with Docker	Kindi as a Windows program, with its own Python, CuPy and NVRTC
Cards	NVIDIA. AMD cards use the CPU engine	NVIDIA
CPU engine	Yes	No: the CPU engine needs gcc with OpenMP. Use WSL2 for it
Kindi Bridge, the Linux installer, Docker images	Yes	No: they are Linux only
Card memory	The Windows driver keeps part of each card's memory	The same

WSL2 is the supported way, and it runs every part of Kindi. Use native Windows when you do not want WSL2 or Docker on the PC.

Native Windows

The Windows runtime package, `kindi-VERSION-windows-x64.zip`, holds `kindi.exe` with its own Python, CuPy and NVRTC, and Kindi Monitor. You do not install Python and you do not set `CUDA_PATH`. You need only the NVIDIA driver. We tested native Kindi on Windows 11 with an RTX 3080: `kindi convert`, `kindi verify` (2B, 4B and 27B: all pass, token-exact) and `kindi serve` with chat requests.

With two NVIDIA cards, use the card that has no display:

```
$env:CUDA_DEVICE_ORDER = "PCI_BUS_ID"
$env:CUDA_VISIBLE_DEVICES = "1"        # the number of the card in nvidia-smi
```

Kindi 0.7.0 loads the CUDA driver from `nvcuda.dll`, so you do not add lines to the source. The prompt buffer of the 5-trit models now has the size that the card's own weights need, so the 27B gets more context on an 8 GB card.

Kindi Monitor

Kindi Monitor (`kindi-monitor.exe`) is a small program in the Windows notification area. A click on its icon opens a popup with live graphs. The popup opens into a dashboard window that you can make larger.

For each Kindi server on the PC, the monitor shows:

- the model, the context, the output head and the drafting;
- whether the server reads a prompt, writes, or waits;
- the write speed and the prompt speed as graphs, and the progress of a long prompt;
- the requests, the tokens and the tokens that came from the prompt cache;
- the decode slots and the queue.

For each NVIDIA card, it shows the memory, the use, the temperature, the power, the clocks and the PCIe link. For the PC, it shows the CPU, the RAM, the disk and the network.

The monitor finds native servers, servers in WSL2 and servers in Docker Desktop. It reads the plan file and the installer's state, and it also looks at the usual Kindi ports. For a server on this PC, it can start, stop and restart the server, and change its model. It shows the command before it runs it, and it asks before a model download. It never asks for administrator rights. A server on another computer shows its telemetry only, with no controls. A server that is not Kindi (for example, llama.cpp's server) shows as "not Kindi", in grey.

CHAPTER 21

What does the server do for an agent?

An agent sends long prompts, many times, and it waits for each answer before it continues. Kindi 0.6.0 adds server features for this kind of client. Every Kindi server has them: the servers of the Bonsai models, Nemotron, the 33B, Lightning and DeepSeek. They use the same OpenAI API as before, so a client that does not know them works as before.

How many tokens is my prompt?

An agent must know if its prompt fits. Kindi gives the exact count, from its own tokenizer and chat template.

Request	What you get
`POST /v1/tokenize` with `{"content": "..."}`, or with `messages` and `tools`	The token count and the tokens
`POST /v1/detokenize` with `{"tokens": [...]}`	The text
A chat request with `"dry_run": true`	The exact `prompt_tokens` of that request. Kindi does not generate
`GET /v1/models`	The context limit of the server (`max_model_len` and `meta.n_ctx`)

A prompt that does not fit gets an HTTP 400 at once, with the code `context_length_exceeded`. The client does not wait in the queue for an error. An agent can then make its history shorter and send the request again.

Does the client know that Kindi is still working?

A long prompt can take minutes to read on a small card. Some clients stop a request when nothing comes back for a while. In a stream, Kindi sends a `: kindi progress {...}` comment every 10 seconds while it reads. Clients ignore these comment lines. A client that does not stream can ask for keep-alive spaces before the JSON answer with `"kindi_keepalive": true`.

Does a restart lose the conversation?

No. After each answer, the server saves the state of the conversation to the disk in the background. After a restart, the next turn of a saved conversation continues from that state. It does not read the whole conversation again. The DeepSeek server also saves the state at the start of the thinking block, before `<think>`, because the next turn of the agent does not contain the thinking text.

Which conversation does a state belong to?

Send the header `X-Kindi-Session`, or the field `kindi_session` or `user`, with a name for each conversation. When the cache is full, Kindi removes the earlier states of each conversation before the newest state of any other conversation. So two agents on one server do not push out each other's latest work. The DeepSeek server can also keep 2 to 4 conversations in RAM with `--slots N`.

Can chat go before background work?

Yes. Each request is chat (the default) or background, from `"kindi_priority": "background"` or the header `X-Kindi-Priority`. Chat requests go first. A background request that reads a long prompt stops between two chunks of the prompt when a chat request comes. After the chat answer, it continues from the cache, and its answer does not change.

What happens when the client goes away?

When a client closes its connection, Kindi stops the request between two chunks of the prompt or two steps of generation. It saves what it has read, and the server is free for the next request. The next request with the same start uses the saved part.

Can one server answer several requests at once?

Yes, from Kindi 0.7.0, with decode slots. Start the server with `kindi serve --slots N` (or the config key `slots`, or `KINDI_SLOTS=N`), with N up to 8. Each slot is one conversation, with its own K/V cache, its own recurrent state and its own prompt cache. One pass of the model writes one token for every slot that is writing, so the slots share each read of the weights. Kindi reads the prompts one chunk at a time between the passes.

Each slot gives the same greedy output as when its conversation runs alone. We checked this token for token on the 2B and the 27B. Slots are an option: the default is one request at a time. On an RTX 3080 with `--slots 8`, a pass of 8 slots writes 396 tokens/s in total with the 27B (one slot: 70), and 2,852 tokens/s with the 2B (one slot: 481). Through the server, 16 requests of 256 tokens to the 27B finish at 134 tokens/s in total with 8 at a time, against 52 tokens/s with the default of one at a time.

> **Conditions.** One GeForce RTX 3080, the plan engine with the tensor-core plans (`KINDI_TC=1`), greedy. Kindi 0.7.0 development build, 2026-11-15 and 2026-11-16.

Things to know:

- Slots need the plan engine on one card with the tensor-core plans: `KINDI_TC=1`, and `KINDI_TP=1` for the 2B, 4B and 9B.
- The context fit gives every slot the same context, so each slot gets a part of the card memory.
- Drafting (the 2B drafter, DSpine and the MTP head) is off with more than one slot. The log says so.
- The default is one slot. The installer does not set more.

What does /health show?

`/health` shows the model, the context, the output head and the prompt cache, as before. In 0.6.0 it also shows the drafting of the last request: the method, the drafter, and the drafted and accepted tokens. Kindi Bridge shows these values in its Kindi panel (Kindi Bridge).

How do other computers on your network use Kindi?

By default, the server answers only on the computer itself (`127.0.0.1`). `install.sh --lan` serves the local network. It makes an API key in `~/.kindi/api-key` and gives the key to the server. Clients then send `Authorization: Bearer KEY`. `--api-key KEY` sets your own key.

CHAPTER 22

Meet pi

You type a task in plain language. Something reads your files, runs commands and writes the result. That something is pi, a coding agent for the terminal. pi sends your task to a language model, and the model answers with actions that use pi's four tools:

Tool	What it does
`read`	reads a file
`bash`	runs a shell command
`edit`	replaces exact text in a file
`write`	writes a whole file

pi runs each action, shows the result to the model, and lets the model choose the next action. The loop continues until the model says that the task is complete.

Because pi can use any server that speaks the OpenAI API, it can use Kindi. pi also has themes, and every picture in this manual uses the Kindi theme: lime on black.

THE PROMPT TYPED INTO PI

> contacts.txt is free text. Extract every e-mail address in it, remove duplicates, and write them to emails.txt sorted, one per line. Then print emails.txt.

THE FILE CONTACTS.TXT AT THE START

```
Meeting notes. Call ann@example.com or bob.smith@example.org about the budget.
CC: ann@example.com; zed+test@mail.example.net
No address on this line.
```

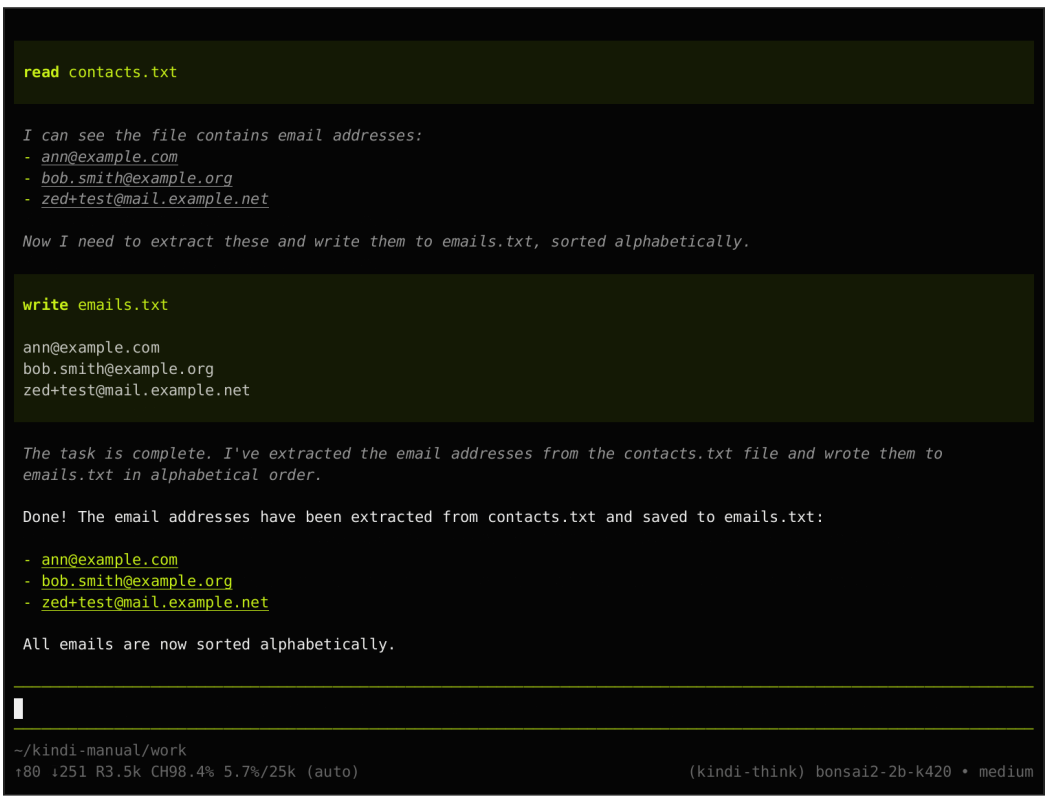

Figure 22-1. *pi on Kindi extracts e-mail addresses from free text (Quadro K420, live session)*

How do you connect pi to Kindi?

Add a provider for Kindi to `~/.pi/agent/models.json` (here, the 2B on the K420):

```
{
  "providers": {
    "kindi": {
      "baseUrl": "http://127.0.0.1:8093/v1",
      "api": "openai-completions",
      "apiKey": "none",
      "compat": {
        "supportsDeveloperRole": false,
        "supportsReasoningEffort": false,
        "thinkingFormat": "qwen-chat-template",
        "thinkingTokenBudgetField": "thinking_budget_tokens"
      },
      "models": [{
        "id": "bonsai2-2b", "name": "Bonsai 2 2B on Kindi", "reasoning": true,
        "input": ["text"], "contextWindow": 24576, "maxTokens": 8192,
        "cost": {"input": 0, "output": 0, "cacheRead": 0, "cacheWrite": 0}
      }]
    }
  }
}
```

The installer prints this block for each server that it starts, with the right address, model `id` and `contextWindow`. For Nemotron, the model `id` is `nemotron-bolt-e4b`.

Then give requests plenty of time in `~/.pi/agent/settings.json`. An agent's first prompt holds 1,000 to 5,000 tokens. On the K420, Kindi reads it once, at 33.3 tokens per second for 512 tokens, and keeps it in its cache (What does the prompt cache do?).

```
{
  "defaultProvider": "kindi",
  "defaultModel": "bonsai2-2b",
  "defaultThinkingLevel": "medium",
  "retry": { "provider": { "timeoutMs": 1800000, "maxRetries": 1 } },
  "thinkingBudgets": {
    "minimal": 256, "low": 512, "medium": 1024, "high": 2048
  }
}
```

Set `maxTokens` to 8192: with 2048, the model often stopped in the middle of a file. Set the thinking level to `medium` for terminal work: it raised the terminal success rate from 40% to 58%.

CHAPTER 23

Kindi Bridge

Kindi Bridge is a workspace in the terminal for agents that run on Kindi. It keeps your agents and their conversations open when you close the window. It sends prompts at a time that you set, and it can ask an agent again and again until the agent confirms that a job is done. It indexes folders of your own documents, so that an agent can use them. It also shows the live state of your Kindi servers and of the computer.

Bridge uses pi (Meet pi) for each agent. Bridge adds what pi does not have: many agents in one place, a schedule, completion checks, document search and a view of Kindi.

Bridge has the same license as Kindi: the GNU Affero General Public License v3.0. For use in a closed product or a hosted service, a paid commercial license is available (amirhusain.com/contact).

Menu	What it is for
Agents	Open a conversation, start a new agent, send an instruction, switch between agents
Schedule	Send a prompt later, repeat a completion check, see the queue and the activity log
Kindi	See your Kindi servers and the Kindi features, and choose the server that new agents use
Files	Browse and edit the agent's files, add and search folders of documents
System	See the CPU, memory, GPU, processes, listening ports, cron jobs and timers
View	The command pane, zoom, refit, and detach
Help	The first steps, all the keys, and the version

How do you install Bridge?

Bridge is an option of the Kindi installer. Add `--with-bridge`:

```
curl -fsSL https://amirhusain.com/kindi/install.sh | bash -s -- --with-bridge
```

When the installer runs in a terminal, it also asks: "Install Kindi Bridge too?". `--no-bridge` stops the question. `--yes` and `--auto` alone never add Bridge.

Bridge needs:

Item	Version	If it is missing
Python, with venv	3.10 or later	The installer offers to install it with `apt-get`
tmux	3.2 or later	The installer offers to install it with `apt-get`
pi	1.3.0	The installer offers to install it with `npm`, into `~/.local` (no `sudo`)
Node.js	22.19 or later (for pi)	Install it from nodejs.org, then run the installer again

The installer then asks two questions:

1. "**Add document search to Bridge?**" Yes adds the reading of PDF, Word, PowerPoint and Excel files, and local embeddings for search by meaning. It adds about 300 MB. With no, Bridge searches plain text by keywords.
2. "**Run the Bridge scheduler as a user service?**" Yes makes a systemd user service. The scheduler then sends scheduled prompts also when Bridge is closed. With no, the scheduler runs while Bridge is open.

The installer puts each part in its own place, and it does not change your own pi setup:

Part	Where
The program (a Python venv)	`~/.kindi/bridge`
The command	`~/.local/bin/kindi-bridge`
Agents, conversations, schedules, document indexes	`~/.local/state/kindi-bridge`
Bridge's own pi settings	`~/.local/state/kindi-bridge/pi-agent`
The user service (if you chose it)	`~/.config/systemd/user/kindi-bridge.service`

At the end, the installer connects Bridge to your first Kindi server. New agents use that server.

To upgrade Bridge, run the installer again: when Bridge is installed, the installer upgrades it too. To remove it, run `bash install.sh --uninstall`. This removes the program, the command and the service. Your agents and schedules stay. Add `--purge` to delete them too (the installer asks first).

On Windows, run `install.ps1 -WithBridge`, or answer yes when the installer asks. `-NoBridge` stops the question. Bridge runs inside WSL, together with the rest of Kindi. The installer adds **Start > Kindi > Kindi Bridge**, which opens Bridge in Windows Terminal.

How do you start?

Type `kindi-bridge`. From another computer, type `ssh -t HOST kindi-bridge`.

The first start makes an agent named `default` and opens its conversation. The screen has five parts:

Part	What it shows
Top line	`KINDI // BRIDGE`, the computer, the agent, its state and the screen
Menu bar	Agents · Schedule · Kindi · Files · System · View · Help
Conversation (left)	The agent's pi session. Type your instructions here
Command pane (below it)	A shell in the agent's folder
Side panels (right)	Hardware vitals, and Kindi telemetry for the current server

To open a menu, click its title, or press **Alt-M** or **F10**. Left and Right change the menu. Up and Down select an item. Enter chooses it. Esc closes the menu. Each menu also has its own key: **Alt-A** Agents, **Alt-S** Schedule, **Alt-K** Kindi, **Alt-F** Files, **Alt-Y** System, **Alt-V** View, **Alt-H** Help. **Alt-W** returns to the conversation from any screen.

When the terminal is narrower than 100 columns or lower than 28 rows, the main pane fills the window. When you make the terminal larger again, the side panels come back.

To leave, choose **View > Detach**. The agents continue to work. Type `kindi-bridge` again to return.

How do you connect Bridge to Kindi?

Open **Kindi > Servers & features**.

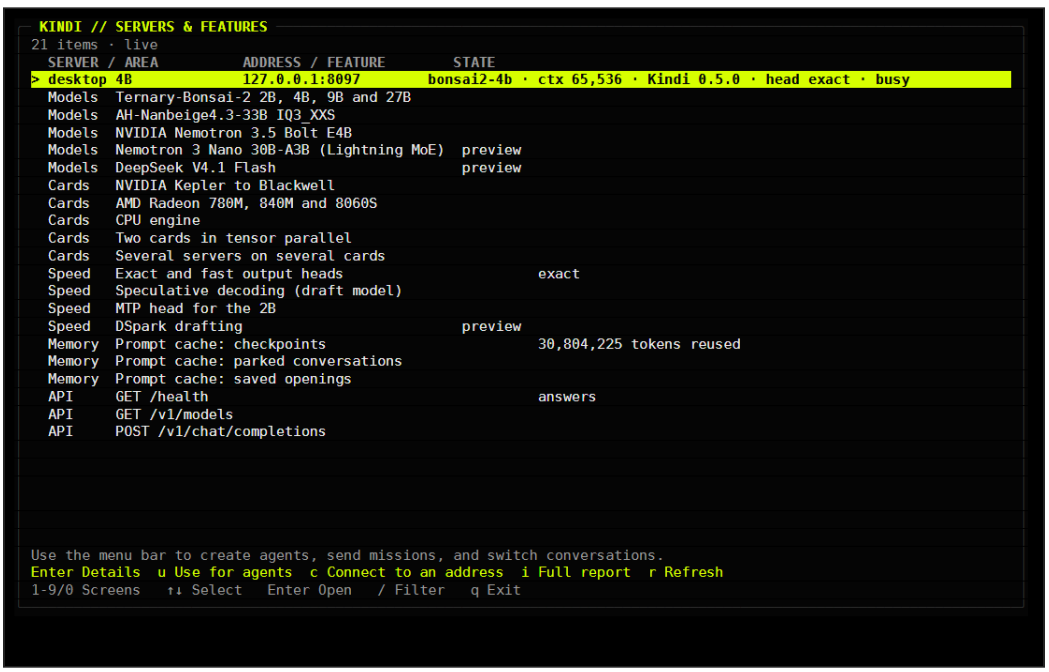

Figure 23-1. *The Kindi screen: one Kindi 0.5.0 server with the 4B, then the Kindi features with live values*

The first rows are your Kindi servers. Bridge finds them in three places:

- the server that Bridge uses now (marked with >);

- the servers that the Kindi installer started (from `~/.kindi/install.json`);
- the servers in Bridge's pi settings.

Each row shows the address, the model, the context size, the Kindi version, the output head (Exact and fast output heads), drafting, and if the server is busy or idle. A server that does not answer shows `offline`.

The rows below are the features of Kindi: the models, the cards, the speed features, the prompt cache and the API. A feature that is in development for a later release shows `preview`. When the server reports a feature in `/health`, the row shows its live value: for example the head that runs, or how many prompt tokens came from the cache.

The keys on this screen:

Key	What it does
Enter on a server	Shows the server's `/health` in full
Enter on a feature	Shows what the feature does, and the Kindi release that has it
u on a server	Makes this the server for new agents (asks you to type y)
c	Connects to a server at an address that you type, for example `192.0.2.10:8093`
i	Shows the full report as text
r	Reads all servers again

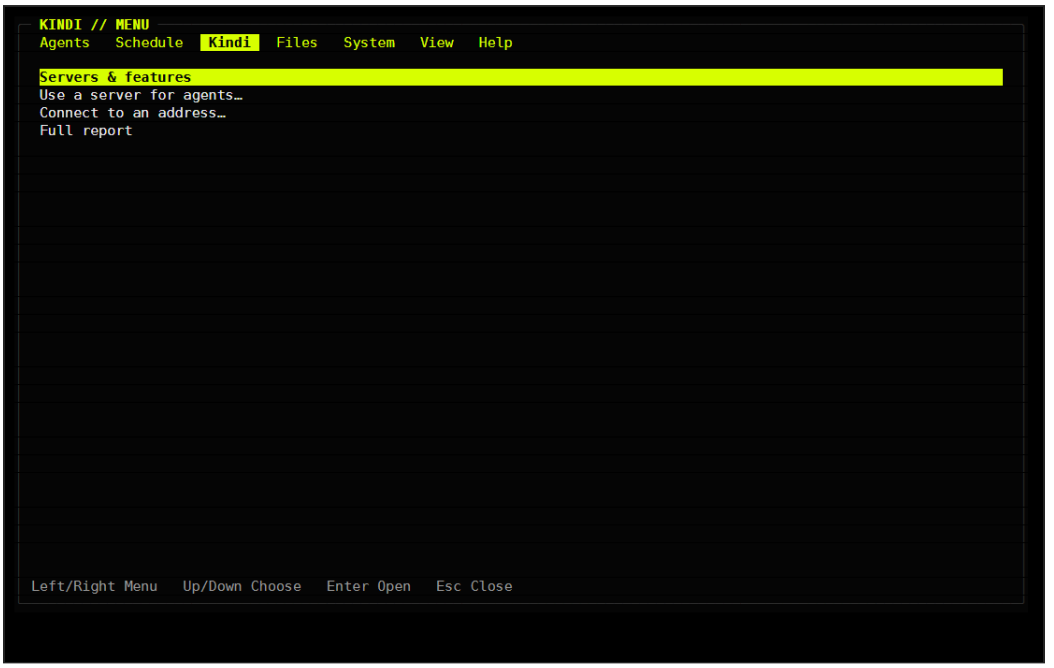

Figure 23-2. *The Kindi menu*

The **Kindi** menu has the same actions:

- **Servers & features** opens the screen above.
- **Use a server for agents** shows each server with its state. Choose one with Enter.
- **Connect to an address** asks for the address and an API key. Leave the key empty when the server has none.
- **Full report** shows all servers and all features as text that you can scroll.

When you choose a server, Bridge writes a pi provider for it into its own pi settings, with the model, the context size, `maxTokens` 8192 and a long request timeout (30 minutes, one retry). New agents then use that server. Each agent keeps the server that it started with. To move work to the new server, start a new agent.

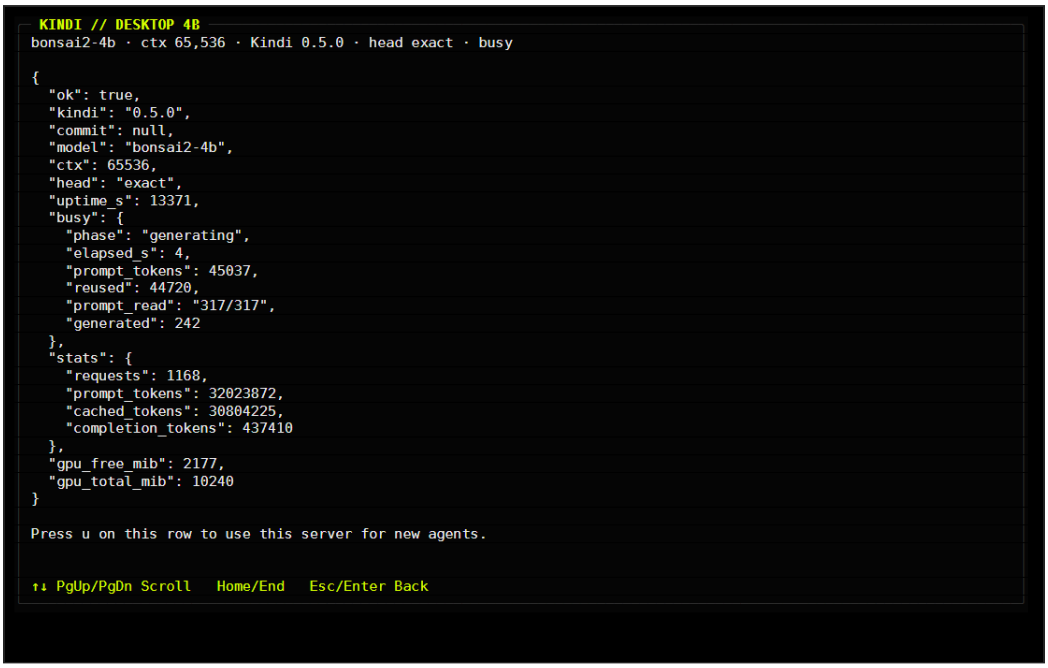

Figure 23-3. *Enter on a server shows its /health*

Agents

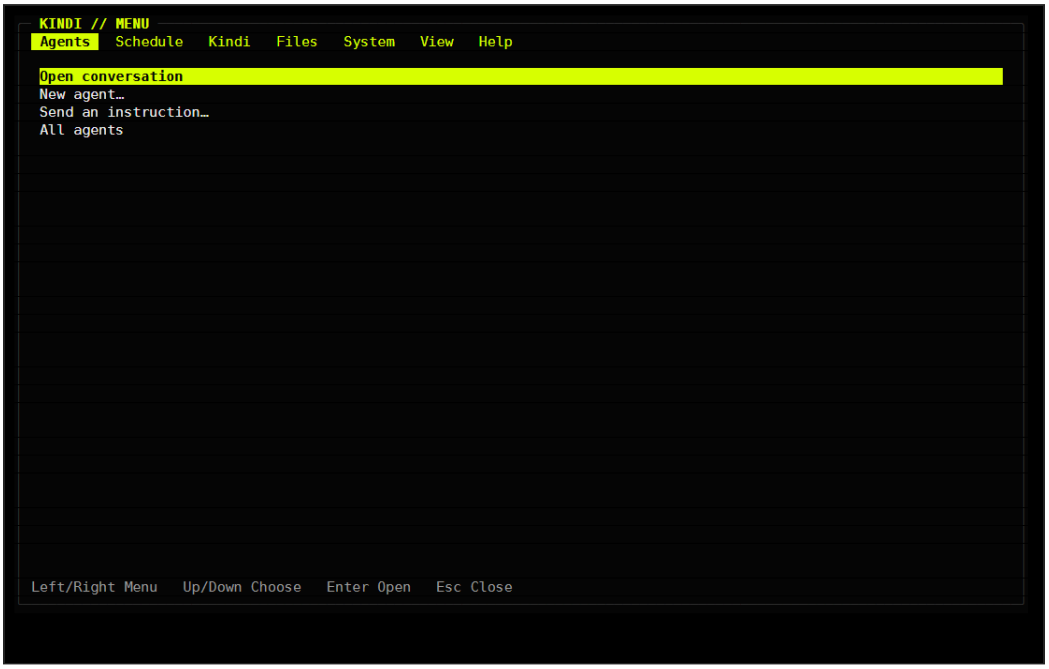

Figure 23-4. *The Agents menu*

Item	What it does
Open conversation	Returns to the conversation of the current agent
New agent	Asks for a name, a mission and a working folder. **Create & Start** starts the agent and sends the mission at once
Send an instruction	Asks which agent, then for the text. A busy agent gets it at its next safe point, when its turn ends
All agents	A table of all agents: name, status, queued prompts and mission
The agent rows	Each agent, with its status. Choose one to open its conversation

In **All agents**: Enter opens the agent, **n** starts a new agent, **m** sends a mission, and **i** shows the details.

The status of an agent is one of: Starting, Ready, Working, Running a tool, Waiting for model, Needs you, Offline or Failed. "2 queued" means that two prompts wait for the agent.

A new agent uses the Kindi server that you chose last in the Kindi menu. It keeps that server, also after a restart.

Schedule

Item	What it does
Schedule a prompt	Sends a prompt to an agent at a time
Check until done	Asks an agent again and again if a job is done, until it confirms, or until a limit or a deadline
Scheduled prompts	The queue: time, agent, state, runs and prompt
Activity log	Each event: prompts sent, answers, errors, processes stopped

The time can be `now`, `in 20m` (also `s`, `h` and `d`), `16:30` (the next 16:30 on this computer), or a date and time such as `2026-11-12T08:00:00-06:00`.

Check until done has five fields:

Field	Default	What it means
Prompt	"Was the previous job successful? If not, plan an alternate strategy and resolve it. If it is complete, confirm completion and give the evidence."	The text of each check
When	`in 20m`	The time of the first check
Job to check	The agent's mission	The exact job and its test of success
Check every	`30m`	The time between checks
Maximum checks	`6`	The checks stop after this number

Bridge stops the checks when the agent confirms the job in the record form that Bridge asks for, at the limit, at the deadline, or at a delivery error. A missing or wrong record never counts as success.

In **Scheduled prompts**: **a** adds a prompt, **c** adds a check, **d** deletes (with 10 seconds to undo with **u**), space pauses or continues, and Enter shows the prompt, its job and the last answer. A prompt that an agent already received cannot be deleted. When Bridge restarts during a delivery, it marks that prompt **uncertain**. Look at it, then send it again with `kindi-bridge jobs retry ID`.

Files

Item	What it does
Project files	Browses the agent's working folder
Add a knowledge folder	Indexes a folder of documents for the current agent
Knowledge & search	Lists the indexed folders, with their state, files and problems

In **Project files**: Enter opens a folder or shows a file, **e** edits a text file (Ctrl-S saves, Esc cancels), and **g** shows the Git changes. When the agent changed a file while you edited it, Bridge keeps your text in a conflict file. It does not write over the agent's change.

A knowledge folder can hold text, source code, Markdown, logs and CSV files. With document search, it can also hold PDF (with page numbers), Word, PowerPoint and Excel files. Bridge reads formulas as text and never runs them. It skips files over 20 MiB, hidden files, links, files with credentials in the name, and build folders. It looks for changes every 5 minutes. Before each prompt, Bridge adds the passages that match best (at most 1,500 bytes) from the agent's folder.

In **Knowledge & search**: **s** searches, **b** attaches the folder to the current agent, **r** indexes it again, and Enter shows the files that had problems.

System, View and Help

Menu	Item	What it does
System	**Performance**	CPU, memory, disk, each GPU, and the `/health` of the current Kindi server
System	**Processes**	All processes. **k** stops one, **K** forces it. Bridge asks first, and never stops itself
System	**Listening ports**	Each listening port and its process. **k** stops the process (all its ports close)
System	**Cron & timers**	The cron jobs and systemd timers on this computer (read only)
View	**Command pane**	Moves the cursor to the shell
View	**Zoom / restore**	Makes the main pane fill the window, or makes it normal again
View	**Refit panels**	Puts the panes back in their places
View	**Detach**	Closes the window. The agents continue
Help	**Getting started**	The first steps
Help	**Keys and mouse**	All the keys
Help	**About Bridge**	The version and the state folder

The **Kindi telemetry** panel on the right shows the current server: the model, the Kindi version, the context, the head, and if drafting is on. While a request runs, PROMPT shows how much of the prompt Kindi has read, ANSWER shows the tokens written and the speed, and REUSED shows the tokens that came from the prompt cache. Below that, a graph shows the write speed of the last minute, and the counters show the requests and tokens since the server started.

Which keys does Bridge use?

Keys	What they do
Click a menu title, Alt-M, F10	Opens the menus
Alt-A, Alt-S, Alt-K, Alt-F, Alt-Y, Alt-V, Alt-H	Opens one menu
Alt-W, Ctrl-b 0	Returns to the conversation
F1 to F9, or Ctrl-b then 1 to 9	Conversation, agents, scheduled prompts, Kindi, files, knowledge, performance, processes, ports
Up, Down, PgUp, PgDn, Home, End, mouse	Moves in a table
/	Filters a table
Esc or q	Returns from a table to the conversation
Ctrl-b e, Ctrl-b c	Moves the cursor to the conversation, or to the command pane
Ctrl-b z, Ctrl-b r	Zoom, refit
Ctrl-b and an arrow	Makes the current pane larger or smaller
Tab, Enter, Ctrl-S, Ctrl-E, Esc	In a form: next field, next field, submit, open a long text in the editor, cancel

Over SSH, some terminals do not send the F-keys. Use the mouse, Alt-M, or Ctrl-b and a number. **Help > Keys and mouse** shows the same list in Bridge.

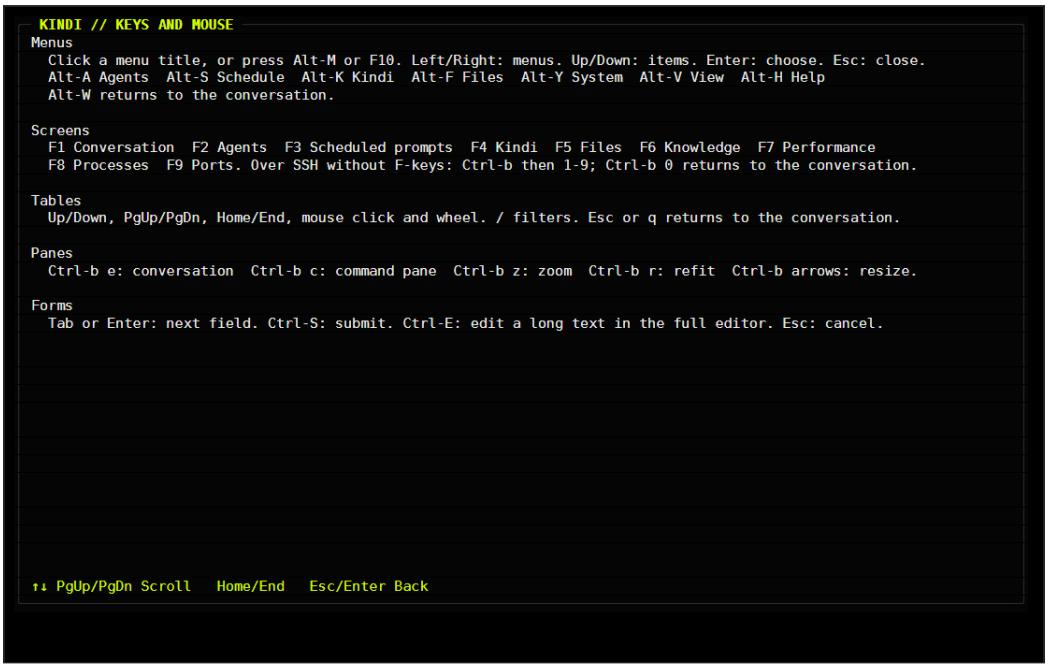

Figure 23-5. *Help > Keys and mouse*

Five things to do with Bridge

A job that runs overnight

You want an agent to repair a test suite while you sleep, and to keep at it until the tests pass.

1. **Agents > New agent.** Name: `fixer`. Mission: "Make the tests in tests/ pass. Do not change the tests." Working folder: your project. Choose **Create & Start**.
2. **Schedule > Check until done.** Agent: `fixer`. When: `in 30m`. Job to check: "All tests in tests/ pass when I run `python -m pytest -q`". Check every: `30m`. Maximum checks: `12`.
3. **View > Detach**, and go to bed.

In the morning, **Schedule > Activity log** shows each check and its answer. **Scheduled prompts** shows the state of the rule: confirmed, exhausted (the 12 checks were used), or an error to look at.

Two agents on two servers

You have a 27B server on port 8093 and a 4B server on port 8094. You want the 27B to review a plan and the 4B to run many small text jobs.

1. **Kindi > Use a server for agents**, choose port 8093.
2. **Agents > New agent**, name `reviewer`.
3. **Kindi > Use a server for agents**, choose port 8094.
4. **Agents > New agent**, name `sorter`.

Each agent keeps the server that it started with. The Agents menu shows both, and you switch with one click.

Questions about your own documents

1. **Files > Add a knowledge folder**, for example `~/manuals`. Bridge indexes it in the background.
2. In **Files > Knowledge & search**, wait until the state is `ready`. Press **s** to try a search.
3. Ask the agent, for example: "Using the reference passages, how do we restore a failed database?"

Each prompt to this agent now carries the best passages from the folder, with their file names and pages.

A report every morning

From the command line:

```
kindi-bridge jobs add 'Read today.log and write a five-line summary to
summary.md' \
  --agent sorter --at 07:00 --every 1d --limit 30
```

If the computer was off at 07:00, Bridge sends one prompt when it starts again. It does not send the missed prompts one after the other.

Watch Kindi while an agent works

Keep the conversation open, and look at the Kindi telemetry panel. During a long first prompt, PROMPT moves while Kindi reads. At the next turn, REUSED shows

that most of the prompt came from the cache, and the answer starts sooner (What does the prompt cache do?). **Kindi > Servers & features** shows the same counters for all servers.

The command line

Each menu action also has a command. The commands print JSON.

Command	What it does
`kindi-bridge`	Opens the workspace
`kindi-bridge agent add NAME --dir DIR --start`	Adds an agent and starts it
`kindi-bridge agent list`	Lists the agents
`kindi-bridge jobs add 'TEXT' --agent NAME --at 'in 20m'`	Schedules a prompt (`--every 30m --limit 4` repeats it)
`kindi-bridge jobs check --agent NAME --task 'JOB' --every 30m --limit 6`	Adds a completion check (`--deadline` stops it at a time)
`kindi-bridge jobs list`, `jobs history`	Shows the queue and the history
`kindi-bridge jobs delete ID`, `pause ID`, `retry ID`	Changes one prompt
`kindi-bridge knowledge add DIR --agent NAME`	Indexes a folder (`--keyword-only` does not download the embedding model)
`kindi-bridge knowledge search 'QUESTION'`	Searches the folders
`kindi-bridge processes`, `ports`, `schedules`	Shows processes, ports, cron jobs and timers

`--state DIR` (or `BRIDGE_STATE`) uses another state folder. This is useful for tests.

When something does not work

What you see	Why	What to do
ENDPOINT OFFLINE in the telemetry panel	The current Kindi server does not answer	Check the server (`curl -s localhost:8093/health`), or choose another one in the Kindi menu
`needs an API key` on a server row	The server has an `api_key`	**Kindi > Connect to an address**, and type the key
A new agent does not start	pi is not installed, or is not on the PATH	Install pi (see the installer's note), then start the agent again
"Bridge: pi does not start"	pi 1.3.0 needs Node.js 22.19 or later	Install a newer Node.js, then type `kindi-bridge` again
The first answer takes minutes	Kindi reads the agent's first prompt (thousands of tokens)	Wait. The next turns come from the cache. Bridge's pi timeout is 30 minutes
A prompt is **uncertain**	Bridge restarted while it sent the prompt	Read the agent's conversation, then `kindi-bridge jobs retry ID` if the agent did not get it
The side panels are gone	The terminal is small	Make the window larger, or use **View > Refit panels**
F-keys do nothing over SSH	The terminal does not send them	Use the mouse, Alt-M, or Ctrl-b and a number
"tmux session belongs to another Bridge state directory"	Another copy of Bridge, with other state, uses the same tmux server	Use the same `--state`, or set `BRIDGE_TMUX_SOCKET` to another name
Scheduled prompts do not go while Bridge is closed	The scheduler runs only with Bridge	Run the installer again and choose the user service

CHAPTER 24

What is an agent?

A chat assistant talks. An agent works.

Ask a chat assistant a question, and it gives you an answer. Give an agent a goal, and it does the work: it reads your files, runs commands, writes results and checks them. You say what you want, and the agent chooses the steps.

What is it like to work with an agent that runs a small model? Think of a new assistant on the first day. A new assistant is fast, keeps going all day and does clear, small jobs well. On long jobs with many rules, the assistant makes mistakes, and sometimes the assistant tells you that a job is complete when it is not. So you give the agent clear instructions, and you check its work.

Which rules make it work?

These rules come from 1,269 recorded sessions with the 2B, the smallest model:

1. **Give one small job at a time.** One clear goal, one output.
2. **Name the output file.** "Write the answer to answer.txt" works better than "tell me the answer". The model can make a mistake in its chat answer, but you can check a file.
3. **Say the exact format.** "One address per line, sorted" is better than "a list".
4. **Start a new session for each task.** In pi, type `/new`. Text from an earlier task in the context confuses a small model.
5. **Check the result yourself.** The model can say "it works" when it does not. In 414 of 731 sessions where the result was wrong, the model's last message said that the result worked or was correct.
6. **Break big jobs into steps.** Can a big job be broken down? shows you how.

THE PROMPT TYPED INTO PI (LEFT PANE)

```
> config.json is a nested settings file. Write the value of database.port to
answer.txt (only the value), then print answer.txt.
```

THE FILE CONFIG.JSON AT THE START

```
{"app": {"name": "demo", "port": 8000}, "database": {"host": "db", "port": 5432,
"user": "app"}}
```

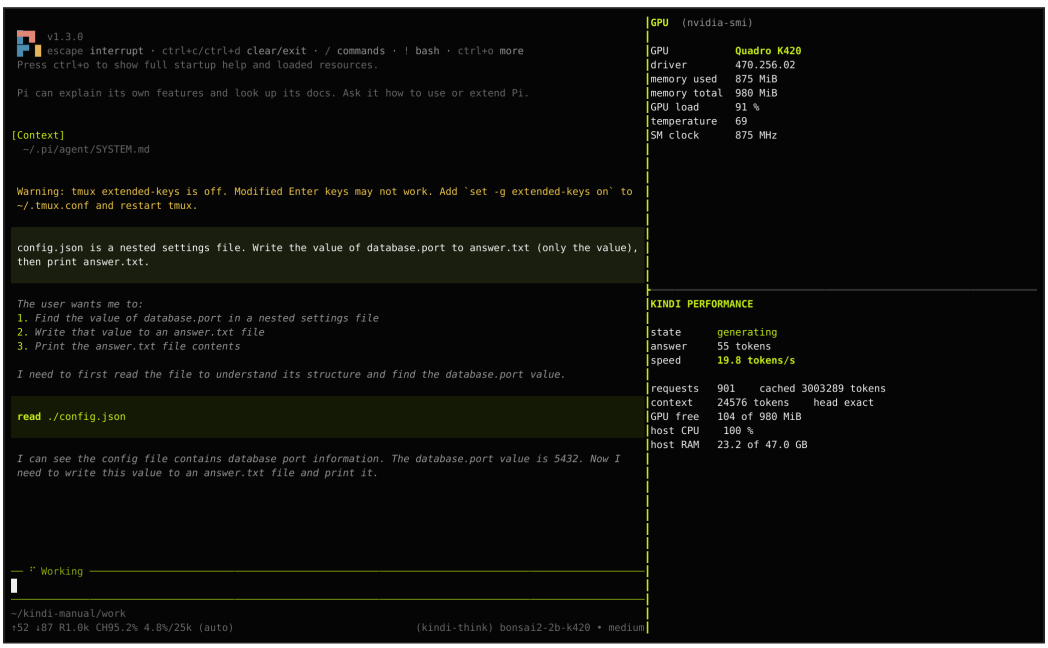

Figure 24-1. *A tmux setup: pi with Kindi on the left, the GPU at the top right, Kindi's live performance at the bottom right*

CHAPTER 25

Fast decisions on short text

Many useful jobs are simpler than agent work. A program sends a short text to Kindi and gets back one word or one corrected sentence: one request, one answer. Is this review positive? Is this e-mail spam? Must a person see this message now? Jobs like these suit a small model very well.

How well? We wrote 158 test items to find out: reviews, customer messages, e-mails, news sentences, sentences with spelling and grammar errors, and short texts to turn into tables. We ran each item once on the K420 at temperature 0, so the same input always gives the same answer.

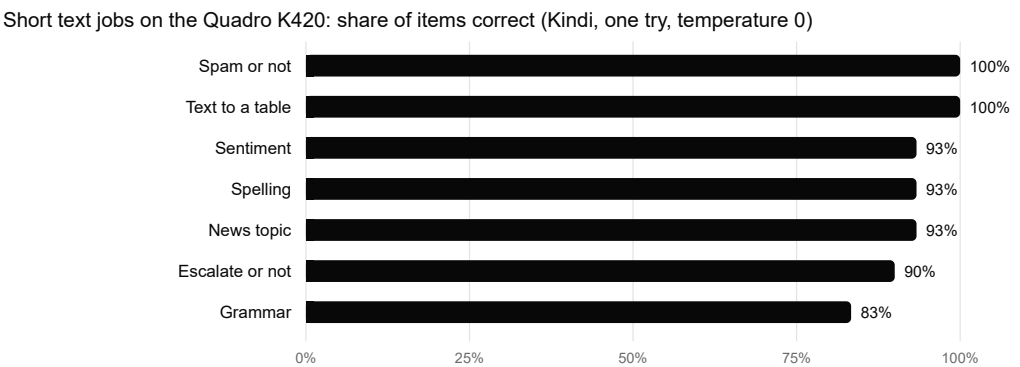

Figure 25-1. *Short text jobs on the Quadro K420: the share of items that Kindi got right*

Job	Correct (Kindi on the K420)	Time per item
Spam or not	**30/30**	2.4 s
Text to a table	**5/5**	4.2 s
Sentiment of a review	**28/30**	2.1 s
Spelling fixes	**14/15**	1.8 s
Escalate to a person or not	**27/30**	3.3 s
Topic of a news item	**28/30**	1.9 s
Grammar fixes	**10/12**	1.9 s

What do the answers look like?

The examples below are real answers from Kindi on the K420.

Spam or not. The program asks for one word: spam or normal.

INPUT
Hi Sam, attached are the slides for Thursday's meeting. Let me know if anything is missing.

KINDI
```
Normal
```

INPUT
CONGRATULATIONS!!! You WON a $1000 gift card. Click here to claim now!!!

KINDI
```
spam
```

INPUT
Your package from the post office is waiting. Pay the $1.99 fee at bit.ly/xx2 to release it.

KINDI
```
spam
```

Escalate to a person, or not. Use this to send urgent customer messages to a person at once.

INPUT
I can't log in after the update, can you help?

KINDI
```
No
```

INPUT
Your heater started smoking and there is a burning smell in my kitchen.

KINDI
```
Yes
```

INPUT
How do I change the email address on my account?

KINDI
```
No
```

Text to a table.

INPUT

Ann is 34 and lives in Austin. Bob is 29 and lives in Boston. Cy is 41 and lives in Chicago.

KINDI

```
| Name | Age | City |
|------|-----|------|
| Ann | 34 | Austin |
| Bob | 29 | Boston |
| Cy | 41 | Chicago |
```

INPUT

Apples cost $2, pears cost $3 and plums cost $4 per kilo.

KINDI

```
| Fruit | Price per kilo |
| :--- | :--- |
| Apples | $2 |
| Pears | $3 |
| Plums | $4 |
```

Headline and summary for a news story.

INPUT

The city council voted 7 to 2 on Monday to build a new bus line that links the airport with the central station. Work starts in spring and the line should open in two years. Supporters say it will cut traffic; opponents worry about the cost of 120 million dollars.

KINDI

```
Headline: City Council Votes for New
Bus Line
Summary: The city council approved a
new bus line linking the airport to the
central station, with work starting in
spring.
```

INPUT

A local bakery that opened in a garage five years ago now employs 40 people. The owners credit their success to sourdough bread and a delivery service by bicycle. They plan to open a second shop next month.

KINDI

```
Headline: Local Bakery Grows to 40
Employees
Summary: A garage-turned-bakery now
employs 40 people and plans to open a
second shop.
```

Clean up notes.

INPUT

meeting w/ supplier tues - prices up 5% from jan?? need to check contract. also ask about delivery delays (3 wks last time). Sara to send numbers by fri. budget review next week, invite finance

KINDI

```
- Meeting with supplier on Tuesday:
review price increase of 5% from
January and check contract.
- Ask supplier about delivery delays
(last time took 3 weeks).
- Sara to send contact numbers by
Friday.
- Budget review scheduled next week
with finance.
- Invite finance to budget review.
```

Can one sentence make escalation safe?

In the escalation test, our first prompt missed 3 of the 15 messages that a person had to see at once. One of them said that a father had already taken pills with the wrong label. A missed escalation is the most dangerous error in this job. So we added one sentence to the prompt: "When you are not sure, answer yes."

Prompt (Kindi on the K420)	Correct	**Missed escalations**	False alarms
The rules only	27 of 30	**3**	0
The rules, plus "When you are not sure, answer yes."	30 of 30	**0**	0

That one sentence removed every missed escalation, and in this test it added no false alarms. With other messages, a sentence like this can add some false alarms: routine messages that a person must then look at. A few false alarms are the right trade. A false alarm costs a person a few seconds; a missed escalation can cost much more. (We also tried a long list of trigger words. The list was much worse: it missed 11 of the 15 urgent messages.)

The lesson works for any job where one type of error is worse than the other: tell the model which way to lean when it is not sure. And let your program check that each answer is one of the words that you allow, so that any other answer counts as "not sure".

CHAPTER 26

What works, and what to avoid?

This chapter gives the detailed results of the 2B, the smallest model. For each type of activity, the chapter tells you what the 2B does reliably, what it does only sometimes, and what you should not give it at all. Which model for which job? compares the 2B with the larger models. Each number is the share of sessions where the result passed an automatic check: "5/5" means five tries and five correct results.

Each table shows two setups: thinking off (with an output limit of 8,192 tokens) and thinking on (medium, with an output limit of 4,096 tokens). A task earns "Works" when one of the two setups passed at least 80 percent of the time, "Sometimes" at 30 to 79 percent, and "Avoid" below 30 percent. The appendix gives the results for every setup that we tested.

These rates come from many runs on our fast test machine (see How did we test?). The last column is one run of the same task with Kindi on the Quadro K420 itself, with thinking on for the terminal tasks and thinking off for the programming tasks. One run cannot give you a rate, but it does tell you whether the task also worked on the card. A dash means that the task has no K420 result.

Each table starts with the tasks that work. Each row that says Sometimes or Avoid also tells you what to do instead. For a terminal task, this is a command for your program to run, and that command passes the same hidden test. For a programming task, it is a program that runs the tests and tries again, or a library that already does the job.

Can it handle files and folders?

Find, list, rename, pack and unpack files. Overall: **43%** of 70 sessions passed on the fast machine.

Result	Task and the prompt that we gave	Thinking off	Thinking on	K420, one run
✓ WORKS	**Rename every .jpeg file in a folder to .jpg** *In the photos folder, rename every file that ends in .jpeg so that it ends in .jpg instead.*	2/5	5/5	✕ fail
✓ WORKS	**Make a .tar.gz backup of a folder** *Make a gzip-compressed tar archive named backup.tar.gz that contains the docs folder.*	3/5	5/5	✓ pass
✓ WORKS	**Make a script executable, run it, save the output** *Make run.sh executable, run it, and save everything it prints into out.txt.*	5/5	5/5	✓ pass
~ SOMETIMES	**Unzip an archive and count its files** *Extract archive.zip into a folder named out, then write the number of files to answer.txt.* **Do this instead:** Let your program run this command. It passes the same hidden test. `python3 -c "import zipfile; z = zipfile.ZipFile('archive.zip'); z.extractall('out'); print(sum(1 for n in z.namelist() if not n.endswith('/')))" > answer.txt`	1/5	2/5	✓ pass
✕ AVOID	**List the files that contain TODO** *List every file under the src folder that contains the word TODO, sorted, one per line.* **Do this instead:** Let your program run this command. It passes the same hidden test. `grep -rl 'TODO' src | sort > answer.txt`	1/5	0/5	✕ fail
✕ AVOID	**Count the lines of all .py files** *How many lines are there in total in all the .py files under src?* **Do this instead:** Let your program run this command. It passes the same hidden test. `find src -name '*.py' -exec cat {} + | wc -l > answer.txt`	0/5	1/5	✕ fail
✕ AVOID	**Find the 3 largest files** *Find the 3 largest files under the data folder, largest first, with their sizes.* **Do this instead:** Let your program run this command. It passes the same hidden test. `find data -type f -printf '%s %p\n' | sort -rn | head -3 > answer.txt`	0/5	0/5	✕ fail

Can it read logs and text?

Read logs, clean lists, find things in text. Overall: **42%** of 60 sessions passed on the fast machine.

Result	Task and the prompt that we gave	Thinking off	Thinking on	K420, one run
✓ WORKS	**Remove duplicate lines and sort** *names.txt has one name per line with repeats. Write unique.txt with each name once, sorted.*	5/5	5/5	✓ pass
✓ WORKS	**Extract e-mail addresses from free text** *Extract every e-mail address in contacts.txt, remove duplicates, write them to emails.txt sorted.*	4/5	5/5	✓ pass
✓ WORKS	**Count the words in a document** *How many words are in essay.txt (as wc -w counts them)?*	1/5	5/5	✓ pass
✕ AVOID	**Replace a word in many files** *In every .txt file under the notes folder, replace the word foo with bar.* **Do this instead:** Let your program run this command. It passes the same hidden test. `find notes -name '*.txt' -exec sed -i 's/foo/bar/g' {} +`	0/5	0/5	✕ fail
✕ AVOID	**Count the ERROR lines in a log** *logs/app.log is an application log. How many lines in it have the level ERROR?* **Do this instead:** Let your program run this command. It passes the same hidden test. `awk '$3=="ERROR"' logs/app.log | wc -l > answer.txt`	0/5	0/5	✕ fail
✕ AVOID	**Find the busiest client IP in a web log** *access.log is a web server log. Which client IP address made the most requests?* **Do this instead:** Let your program run this command. It passes the same hidden test. `awk '{print $1}' access.log | sort | uniq -c | sort -rn | head -1 | awk '{print $2}' > answer.txt`	0/5	0/5	✕ fail

Can it work with data files?

Read and change JSON, CSV and settings files. Overall: **80%** of 40 sessions passed on the fast machine.

Result	Task and the prompt that we gave	Thinking off	Thinking on	K420, one run
✓ WORKS	**Read one value from a nested JSON file** *config.json is a nested settings file. Write the value of database.port to answer.txt.*	5/5	5/5	✓ pass
✓ WORKS	**Change one line in a settings file** *settings.ini has a line port=8080. Change it to port=9090 and leave every other line as it is.*	5/5	5/5	✓ pass
✓ WORKS	**Convert CSV to JSON** *Convert people.csv to people.json: a JSON list with one object per row.*	5/5	4/5	✓ pass
~ SOMETIMES	**Sort a CSV file by a column** *scores.csv has a header row. Write sorted.csv with the rows sorted by score, highest first.* **Do this instead:** Let your program run this command. It passes the same hidden test. `(head -1 scores.csv; tail -n +2 scores.csv | sort -t, -k2,2nr) > sorted.csv`	0/5	3/5	✓ pass

THE PROMPT TYPED INTO PI

> config.json is a nested settings file. Write the value of database.port to answer.txt (only the value), then print answer.txt.

THE FILE CONFIG.JSON AT THE START

```
{"app": {"name": "demo", "port": 8000}, "database": {"host": "db", "port": 5432,
"user": "app"}}
```

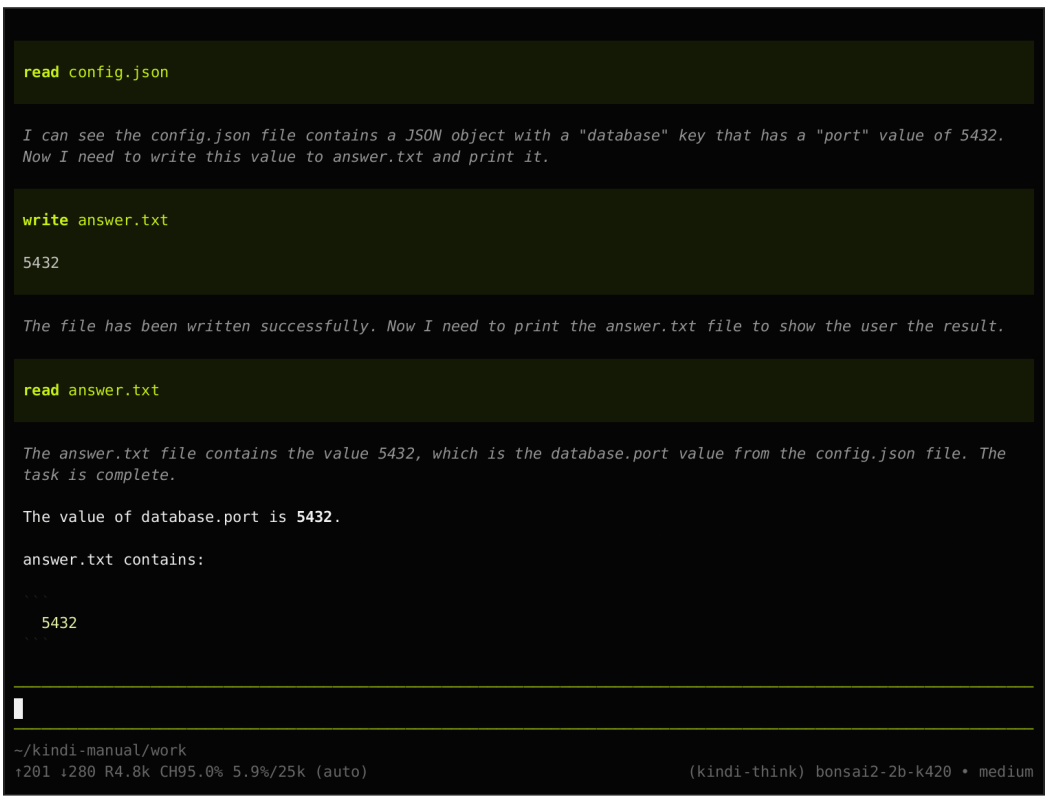

Figure 26-1. *pi on Kindi reads one value from a nested JSON file and writes it to answer.txt (Quadro K420, live session, thinking on)*

Can it look after the machine?

Ask about the machine, use git, plan jobs. Overall: **48%** of 60 sessions passed on the fast machine.

Result	Task and the prompt that we gave	Thinking off	Thinking on	K420, one run
✓ WORKS	**Make a folder a git repository and commit** *Make the project folder a git repository and commit README.md with the message: initial.*	5/5	5/5	✓ pass
✓ WORKS	**Report the installed Python version** *Which version of python3 is installed? Write only the version number to answer.txt.*	1/5	5/5	✓ pass
~ SOMETIMES	**Report the GPU name** *What is the name of the first NVIDIA GPU, as nvidia-smi reports it?* **Do this instead:** Let your program run this command. It passes the same hidden test. `nvidia-smi --query-gpu=name --format=csv,noheader -i 0 > answer.txt`	3/5	3/5	✓ pass
~ SOMETIMES	**Report the free disk space** *How much space is free on the filesystem that holds the working folder (df -k)?* **Do this instead:** Let your program run this command. It passes the same hidden test. `df -k --output=avail . | tail -1 | tr -d ' ' > answer.txt`	1/5	3/5	✗ fail
~ SOMETIMES	**Count the days between two dates** *How many days are there from 2026-01-15 to 2026-10-18?* **Do this instead:** Let your program run this command. It passes the same hidden test. `python3 -c "from datetime import date; print((date(2026, 10, 18) - date(2026, 1, 15)).days)" > answer.txt`	1/5	2/5	✗ fail
✗ AVOID	**Write a crontab line** *Write one crontab line to cron.txt that runs /opt/backup.sh every day at 02:30.* **Do this instead:** Let your program run this command. It passes the same hidden test. `echo '30 2 * * * /opt/backup.sh' > cron.txt`	0/5	0/5	✗ fail

THE PROMPT TYPED INTO PI

> Make the project folder a git repository and commit README.md in it with the commit message: initial. Then show git log --oneline for that folder.

THE FILE PROJECT/README.MD AT THE START

```
# Demo
```

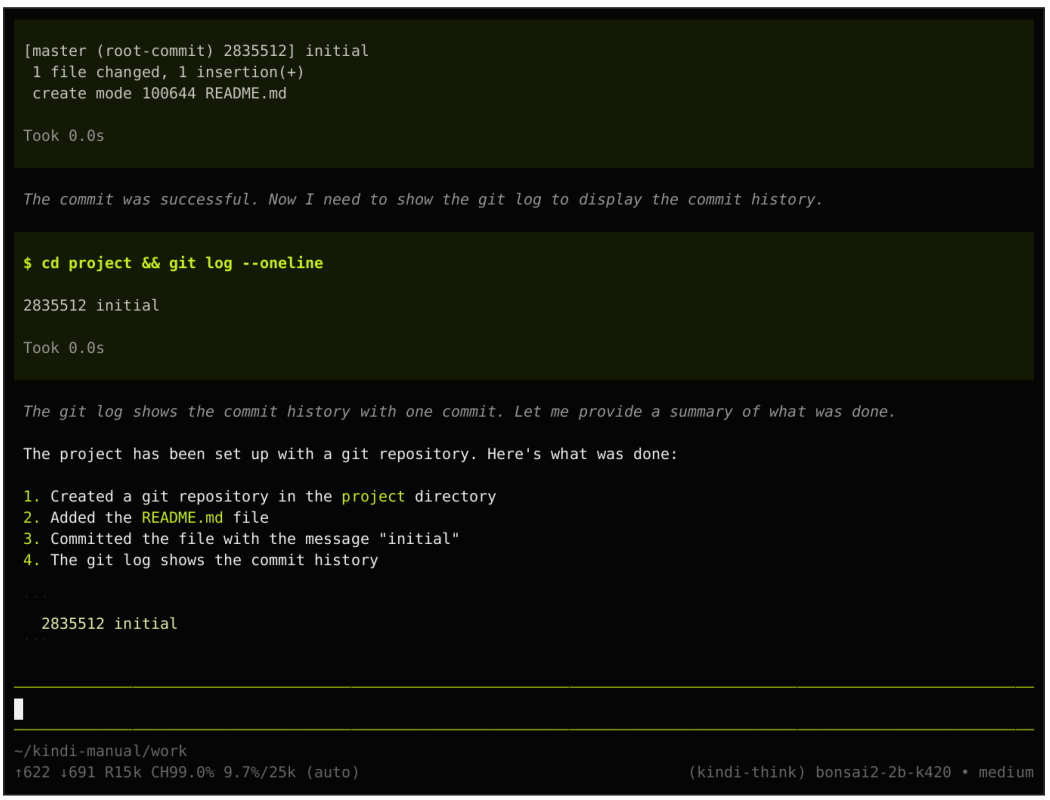

Figure 26-2. *pi on Kindi makes a folder a git repository, commits a file and shows the log (Quadro K420, live session, thinking on)*

Can it write shell scripts?

Write small bash scripts. Overall: **11%** of 28 sessions passed on the fast machine.

Result	Task and the prompt that we gave	Thinking off	Thinking on	K420, one run
~ SOMETIMES	**A script that greets a name from a variable** *Write hello.sh that prints Hello, NAME! where NAME comes from the environment, or world.* **Do this instead:** Let your program run this command. It passes the same hidden test. `printf '#!/bin/bash\necho "Hello, ${NAME:-world}!"\n' > hello.sh`	0/5	2/5	✕ fail
~ SOMETIMES	**A script that counts non-empty lines** *Create countlines.sh that prints the number of non-empty lines in a file, or an error.* **Do this instead:** Let your program run the tests and try again (Pattern 2 in Kindi in your own workflows). 5 of our 15 tries passed, and the program keeps the first version that passes.	0/3	1/3	–
✕ AVOID	**A script that adds its arguments** *Create sumargs.sh that prints the sum of its integer arguments, with an error for bad input.* **Do this instead:** Write it yourself: one loop over `"$@"` that checks that each argument is an integer. In our tests, 1 of 15 tries passed.	0/3	0/3	✕ fail
✕ AVOID	**A script that lists the 3 largest files** *Create largest.sh that prints the 3 largest files under a folder, with sizes.* **Do this instead:** Put the command from the files table in the script: `find DIR -type f -printf '%s %p\n' | sort -rn | head -3`. In our tests, 0 of 15 tries passed.	0/3	0/3	–

Can it write small Python programs?

Write one function from a clear description. Overall: **20%** of 84 sessions passed on the fast machine.

Result	Task and the prompt that we gave	Thinking off	Thinking on	K420, one run
✓ WORKS	**FizzBuzz as a list** *Create fizz.py with fizzbuzz(n) that returns 'Fizz', 'Buzz', 'FizzBuzz' or the number.*	3/3	2/3	✓ pass
✓ WORKS	**Limit a number to a range, with an error case** *Create clamp.py with clamp(x, lo, hi); raise ValueError if lo > hi.*	3/3	2/3	✓ pass
~ SOMETIMES	**Flatten nested lists** *Create flatten.py with flatten(x) for lists nested to any depth.* **Do this instead:** Let your program run the tests and try again (Pattern 2 in Kindi in your own workflows). 7 of our 15 tries passed, and the program keeps the first version that passes.	2/3	1/3	–
~ SOMETIMES	**Check balanced brackets** *Create brackets.py with balanced(s) for ()[]{}.* **Do this instead:** Let your program run the tests and try again (Pattern 2 in Kindi in your own workflows). 4 of our 15 tries passed, and the program keeps the first version that passes.	1/3	0/3	–
~ SOMETIMES	**A stack that also gives its minimum** *Create minstack.py with a MinStack class: push, pop, top, get_min.* **Do this instead:** Let your program run the tests and try again (Pattern 2 in Kindi in your own workflows). 3 of our 15 tries passed, and the program keeps the first version that passes.	0/3	1/3	–
~ SOMETIMES	**The most common words, with tie rules** *Create wordfreq.py with top_words(text, k), sorted by count, then alphabetically.* **Do this instead:** Let your program run the tests and try again (Pattern 2 in Kindi in your own workflows). 3 of our 15 tries passed, and the program keeps the first version that passes.	1/3	0/3	✗ fail

Result	Task and the prompt that we gave	Thinking off	Thinking on	K420, one run
~ SOMETIMES	**Group words that are anagrams** *Create anagram.py with group_anagrams(words), each group sorted.* **Do this instead:** Let your program run the tests and try again (Pattern 2 in Kindi in your own workflows). 2 of our 15 tries passed, and the program keeps the first version that passes.	1/3	0/3	–
✕ AVOID	**Binary search** *Create firstge.py with first_ge(arr, x) in O(log n), without bisect.* **Do this instead:** In your own code, use `bisect.bisect_left` from the Python standard library. It does this search. In our tests, 1 of 15 tries passed.	0/3	0/3	–
✕ AVOID	**Run-length encode and decode** *Create rle.py with encode('aaab') == 'a3b1' and decode.* **Do this instead:** Use `itertools.groupby` from the Python standard library. It groups the runs, and the encoding is then one line. In our tests, 0 of 15 tries passed.	0/3	0/3	–
✕ AVOID	**Merge overlapping intervals** *Create merge.py with merge(intervals).* **Do this instead:** Write it yourself: sort the intervals by start, then join each interval to the last one when they overlap. In our tests, 0 of 15 tries passed.	0/3	0/3	–
✕ AVOID	**A calculator for reverse Polish notation** *Create rpn.py with rpn('3 4 + 2 *') == 14 and errors for bad input.* **Do this instead:** Write it yourself: a list as the stack and one loop over the tokens. The `operator` module gives the operations. In our tests, 0 of 15 tries passed.	0/3	0/3	–

Result	Task and the prompt that we gave	Thinking off	Thinking on	K420, one run
✕ AVOID	**Read a matrix in spiral order** *Create spiral.py with spiral(matrix).* **Do this instead:** Write it yourself, and test it on an empty matrix, a 1x1 matrix and a 2x3 matrix. In our tests, 0 of 15 tries passed.	0/3	0/3	–
✕ AVOID	**Parse an INI file** *Create ini.py with parse_ini(text) that returns sections of keys and values.* **Do this instead:** Use the `configparser` module of the Python standard library. It reads INI files. In our tests, 0 of 15 tries passed.	0/3	0/3	–
✕ AVOID	**Find duplicate files by content** *Create dedupe.py with dedupe_files(folder).* **Do this instead:** Use `os.walk` to find the files, `hashlib` to hash each file, and a dict to group the paths with the same hash. In our tests, 0 of 15 tries passed.	0/3	0/3	–

Can it repair existing code?

Fix a known bug in a small file. Overall: **22%** of 18 sessions passed on the fast machine.

Result	Task and the prompt that we gave	Thinking off	Thinking on	K420, one run
~ SOMETIMES	**Fix a script that crashes on a CSV file** *report.py should print the total of the amount column of sales.csv, but it crashes. Fix it.* **Do this instead:** Let your program run the tests and try again (Pattern 2 in Kindi in your own workflows). 3 of our 15 tries passed, and the program keeps the first version that passes.	2/3	0/3	–

Result	Task and the prompt that we gave	Thinking off	Thinking on	K420, one run
~ SOMETIMES	**Fix two bugs in a statistics module** *stats.py: mean([]) crashes and median() is wrong for even lengths. Fix it.* **Do this instead:** Let your program run the tests and try again (Pattern 2 in Kindi in your own workflows). 7 of our 15 tries passed, and the program keeps the first version that passes.	1/3	1/3	✓ pass
✕ AVOID	**Fix bugs in a JavaScript word counter** *counter.js should count lowercase words. It has bugs. Fix them.* **Do this instead:** Fix it yourself. In our tests, the model did not find these bugs. In our tests, 0 of 15 tries passed.	0/3	0/3	–

Can it write JavaScript?

Node.js modules. Overall: **0%** of 18 sessions passed on the fast machine.

Result	Task and the prompt that we gave	Thinking off	Thinking on	K420, one run
✕ AVOID	**Make a URL slug from text** *Create slugify.js that exports slugify(s).* **Do this instead:** Use a tested package from npm, or write it yourself and test it on titles with symbols and extra spaces. In our tests, 0 of 15 tries passed.	0/3	0/3	–
✕ AVOID	**Merge two objects deeply** *Create deepmerge.js that exports deepMerge(a, b).* **Do this instead:** Use a tested package, for example the `merge` function of lodash. In our tests, 0 of 14 tries passed.	0/3	0/3	–
✕ AVOID	**Fix bugs in a word counter** *counter.js should count lowercase words. It has bugs. Fix them.* **Do this instead:** Fix it yourself. In our tests, the model did not find these bugs. In our tests, 0 of 15 tries passed.	0/3	0/3	–

Can it solve algorithms and parsers?

Classic programming problems. Overall: **3%** of 60 sessions passed on the fast machine.

Result	Task and the prompt that we gave	Thinking off	Thinking on	K420, one run
~ SOMETIMES	**Longest common subsequence** *Create lcs.py with lcs(a, b), dynamic programming.* **Do this instead:** Let your program run the tests and try again (Pattern 2 in Kindi in your own workflows). 8 of our 15 tries passed, and the program keeps the first version that passes.	0/3	2/3	–
✕ AVOID	**Shortest path in a graph** *Create dijkstra.py with shortest_path(graph, start, goal).* **Do this instead:** Let your program run the tests and try again (Pattern 2 in Kindi in your own workflows). 3 of our 15 tries passed, and the program keeps the first version that passes.	0/3	0/3	–
✕ AVOID	**Order tasks by their dependencies** *Create topo.py with order(n, deps), smallest order first, None for a cycle.* **Do this instead:** Use `graphlib.TopologicalSorter` from the Python standard library. It gives a valid order and reports a cycle. In our tests, 0 of 15 tries passed.	0/3	0/3	✕ fail
✕ AVOID	**Roman numerals in both directions** *Create roman.py with to_roman(n) and from_roman(s), strict.* **Do this instead:** Write it yourself from a table of values (1000 M, 900 CM, 500 D and so on), and test both directions. In our tests, 0 of 15 tries passed.	0/3	0/3	✕ fail

Result	Task and the prompt that we gave	Thinking off	Thinking on	K420, one run
✕ AVOID	**Evaluate arithmetic without eval** *Create calc.py with evaluate(s) for + - * / and parentheses.* **Do this instead:** Use the `ast` module to parse the expression, and allow only numbers, parentheses and the four operators. In our tests, 0 of 15 tries passed.	0/3	0/3	✕ fail
✕ AVOID	**Parse one CSV line by hand** *Create csvline.py with parse_line(line), with quoted fields.* **Do this instead:** Use the `csv` module of the Python standard library: `next(csv.reader([line]))` parses one line. In our tests, 0 of 15 tries passed.	0/3	0/3	✕ fail
✕ AVOID	**A least-recently-used cache** *Create lru.py with an LRUCache class.* **Do this instead:** Use `functools.lru_cache`, or `collections.OrderedDict` with `move_to_end`. In our tests, 0 of 15 tries passed.	0/3	0/3	✕ fail
✕ AVOID	**Split text into dictionary words** *Create wordbreak.py with sentences(s, words).* **Do this instead:** Write it yourself as a recursive function with `functools.cache`. In our tests, 0 of 15 tries passed.	0/3	0/3	✕ fail
✕ AVOID	**Multiply big numbers given as text** *Create bigmul.py with multiply(a, b), digit by digit.* **Do this instead:** In Python, `int(a) * int(b)` gives the exact product of numbers of any size. In our tests, 0 of 15 tries passed.	0/3	0/3	✕ fail
✕ AVOID	**The shortest window that contains given letters** *Create minwindow.py with min_window(s, t) in linear time.* **Do this instead:** Write it yourself with two pointers and a `collections.Counter`. In our tests, 0 of 15 tries passed.	0/3	0/3	✕ fail

What is the pattern?

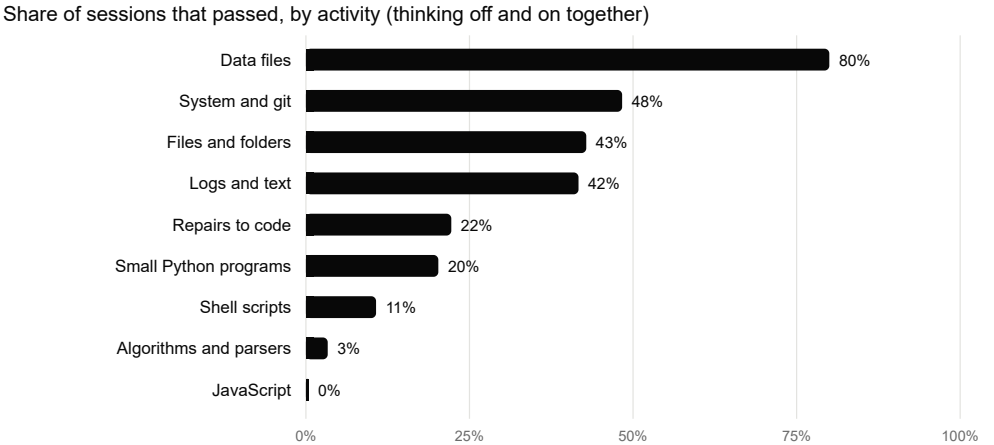

Figure 26-3. *Success rate by type of activity*

The pattern is clear. The 2B model does well when a task has **one action and one clear output**: read a value, remove duplicates, extract addresses, edit one line, make a commit. The model struggles when a task needs **several steps that depend on each other**: count, then sort, then select. JavaScript and real algorithms are also hard for the model.

CHAPTER 27

Can a big job be broken down?

Some jobs are too big for the 2B model as one task. You can still do them when you break them down the right way. We tested four ways to give the model a job that it failed every time as one prompt:

Task	One prompt	One prompt + "use commands"	Chain of prompts	Chain + one command per step
The busiest client IP in a web log	0/10	0/10	0/10	2/10
The number of ERROR lines in a log	0/10	0/10	0/10	2/10
The 3 largest files in a folder	0/10	0/10	0/10	0/10
All three	**0/30**	**0/30**	**0/30**	**4/30**

Passes / runs, thinking off and on together.

Better prompts alone did not save these jobs. A chain of small prompts did not help. When each step asked for one shell command, a few runs passed, but most did not. Why? The transcripts tell the story. The model often reads the data and then writes the result by hand, and on the way it changes a path or drops a line. When it counts or sorts, it makes small errors. A small model is good at language and at short decisions; exact work on data belongs to your program.

Which jobs does a breakdown rescue?

The chains above still asked the model to do exact work. What happens when your program does the exact work, and Kindi answers only one short question about each item? We tested five new jobs on the K420. Each job needs a judgment about

every item and then an exact count, list or sum. We wrote new items for this test, and we checked every result against the known answer:

Job	One prompt, API	One prompt, agent	Broken down, first prompt	Broken down, tested prompt
Tally *Count the positive, negative and neutral reviews among 20*	0/3	2/3	✓ **pass** *20 of 20 items right*	same prompt
Spam *Move the spam e-mails out of a folder of 20*	2/3	0/3	✕ fail *12 of 20 items right*	✕ fail *19 of 20 items right held-out set: fail, 9 of 10 items right*
Urgent *List the IDs of the 7 messages, among 20, that a person must see now*	0/3	2/3	✓ **pass** *20 of 20 items right*	same prompt
Topic *Find the news topic with the most headlines among 15, and its count*	0/3	0/3	✓ **pass** *13 of 15 items right*	same prompt
Orders *Add up the amount paid in 10 free-text order e-mails*	0/3	3/3	✕ fail *8 of 10 items right*	✓ **pass** *10 of 10 items right held-out set: pass, 6 of 6 items right*

One prompt with all the text failed almost every time: 2 of 15 runs passed. The model must judge every item and keep an exact count at the same time, and the count drifts by one or two.

The agent with one prompt passed 7 of 15 runs. The agent results were uneven. The agent got the order total right in all three runs and failed every spam and topic run. In one tally run, the agent counted correctly but wrote the three numbers without their labels, so the check failed.

Broken down, 3 of the 5 jobs passed with the first prompts that we wrote. The program asked Kindi one question for each item and did the counting itself. Each job took at most 73 seconds on the K420. In the topic job, Kindi gave 2 of the 15 headlines the wrong topic, but neither mistake changed the winning topic or its count.

The two failures teach the most important lesson of this chapter: **the question for one item must itself be tested.**

- **Order total.** The prompt ended with "for example 12.34", and the model copied the "12." into two answers: $42.00 came back as 12.00. Without the example number, Kindi read all 10 amounts correctly. Kindi also read all 6 amounts correctly in a held-out set that we wrote after the fix.
- **Spam sweep.** The prompt "Reply with one word: spam or normal." made the model call 8 normal e-mails spam. The tested prompt from Fast decisions on short text got 19 of the 20 e-mails right, and 9 of 10 in a held-out set. The job still failed, because the check needs every e-mail in the right folder.

With tested prompts, 4 of the 5 jobs passed. So before you run a per-item question on a thousand items, run it on ten items whose answers you know, and read the replies. Also decide what one wrong item costs: in a count of sentiment, one mistake moves a number by one; in a spam folder, one mistake hides a real e-mail.

What is the right way to split a job?

Split it into two kinds of steps:

- **Exact steps: your program does them.** Count, sort, filter, add, compare, select the top items. A few lines of Python or one shell command do this with no errors.
- **Language steps: Kindi does them.** Decide if a message is urgent. Find the sentiment. Write a short summary from facts that the program found.

Now the model only does the work that it does well. Fast decisions on short text shows how often each type of short decision is correct.

What does that look like in practice?

Here is a daily operations report, made on the K420 from an application log, a web log and eight customer messages. The program counts the errors and finds the

busiest clients. Kindi decides which messages are urgent, reads their sentiment and writes the summary.

```
levels = collections.Counter(l.split()[2] for l in open("app.log"))  ❶
clients = collections.Counter(l.split()[0] for l in open("access.log"))
top_ips = clients.most_common(3)
for msg in open("messages.txt"):
    urgent = "yes" in ask(RULES + "\n\n" + msg).lower()  ❷
    mood = ask("Classify the sentiment ... One word only.\n\n" + msg)
summary = ask("Write a summary for the operator in exactly three lines ... "
              "Use only these facts.\n\n" + facts)  ❸
```

❶ Exact: the program counts

❷ Language: Kindi decides

❸ Language: Kindi writes the words

And this is the report that the program wrote on the Quadro K420, with 17 calls to Kindi in 53 seconds:

```
# Daily operations report

## Summary (written by Kindi)

The ERROR lines are at 117 total, with the highest hour being 09 with 42
lines. The busiest clients are 10.0.0.2 with 153 connections. There are 3
urgent customer messages out of 8 total.

## Numbers (counted by the program)

- Log levels: {'ERROR': 117, 'INFO': 84, 'DEBUG': 101, 'WARN': 98}.
- ERROR lines by hour: {'08': 35, '09': 42, '10': 30, '11': 10}.
- Busiest clients: 10.0.0.2 (153), 10.0.0.5 (120), 10.0.0.7 (78).
- Urgent customer messages: 3 of 8.

## Customer messages (judged by Kindi)

| Urgent | Sentiment | Message |
|---|---|---|
| **YES** | negative | The checkout page has been down for all our customers
since 9 am. |
| **YES** | negative | Your app shows another customer's order history on my
account. |
| **YES** | negative | There is smoke coming out of the charger you sold me. |
| no | positive | Thanks, the new invoice layout is much clearer. |
| no | neutral | How do I export my data to CSV? |
| no | negative | The delivery was two days late and nobody told me. |
| no | positive | Can you add a dark mode? |
| no | positive | Great support today, thank you Maria! |
```

Look closely at how the work was divided, because it holds the most important lesson of this chapter:

- **The program counted every number.** The counts, the peaks and the busiest clients come from code, so they are exact.
- **Kindi flagged the three urgent messages and no others:** the outage, the privacy breach and the smoke.
- **Kindi wrote the summary from the facts that the program gave it,** and each number in the summary is a number that the program counted.

The model only judged messages and wrote words around facts. Keep that rule in your own workflows: let your program find and print every number, and let Kindi write only the words. Checking is then easy, because each number in the text must be one of the numbers in the facts. The full program is `workflow_demo.py` in the material for this manual.

How do you break down your own job?

1. **Write the job as a list of actions.** "Find the client with the most requests" is really three actions: get the client addresses, count them, pick the first.
2. **Mark each action as exact or language.** Counting, sorting and selecting are exact. Judging, classifying and writing are language.
3. **Write the exact actions as code.** You can ask Kindi to write a short function, and then test it on a small example before you trust it.
4. **Give Kindi one short decision or one short text at a time,** with the facts that it needs in the prompt.
5. **Check the result** with the rules in this chapter and in How do you get the best results?.

CHAPTER 28

Kindi in your own workflows

pi is one way to use Kindi. But what if you want full control? Because Kindi has an OpenAI-compatible API, any OpenAI client library works with it, and you can put the model inside a fixed program that you design.

Why does a fixed workflow help a small model?

Research on small models agrees with our tests. A small model does better in a fixed pipeline than as a free agent (Agentless, Xia and others, 2024). A program that checks each answer and tries again does better than a model that checks itself (Olausson and others, 2024). And many tries with an automatic check find a correct answer much more often than one try (Large Language Monkeys, Brown and others, 2024).

We saw the same thing in a quiz, where the model answered 12 questions about what a short Python program prints:

Method	Correct answers
One answer, thinking off	0 of 12
Majority of 5 answers, thinking off	3 of 12
One answer, thinking on	5 of 12
Majority of 5 answers, thinking on	5 of 12
At least one of the 5 answers correct	7 of 12

Look at the last line. The correct answer was often there, but a majority vote did not find it. If you can check an answer automatically, keep the first answer that passes the check.

Pattern 1: one call, one answer

Use this for short jobs: classify a line, extract a field, write a one-line summary.

```
from openai import OpenAI
kindi = OpenAI(base_url="http://127.0.0.1:8093/v1", api_key="none")

def ask(question, max_tokens=300):
    r = kindi.chat.completions.create(
        model="bonsai2-2b", max_tokens=max_tokens, temperature=0.2,
        messages=[{"role": "user", "content": question}])
    return r.choices[0].message.content.strip()

for line in open("tickets.txt"):
    label = ask("Classify this support ticket as BILLING, BUG or OTHER. "
                "Reply with one word.\n\n" + line)
    print(label, "|", line.strip())
```

Pattern 2: generate, check, try again

Use this for code and for any output that a program can check. The program runs the tests.

```
import subprocess

def generate_until_it_passes(task, tests, tries=5):
    for n in range(tries):
        code = ask(task + "\nReply with only the Python code.",
max_tokens=1500)
        code = code.strip("`").removeprefix("python").strip()
        open("solution.py", "w").write(code)
        r = subprocess.run(["python3", tests], capture_output=True,
                           text=True, timeout=60)
        if r.returncode == 0:
            return code  (1)
    return None  (2)
```

(1) The first version that passes the tests

(2) No version passed: ask a person

Should you feed the model its own error message, again and again? In our tests, a second and a third try with the test failures rarely fixed a program that was wrong.

Research agrees: a fresh try from the start often works as well as a repair (Olausson and others, 2024).

Pattern 3: the program does the steps, the model fills the gaps

Use this when a job has several steps. Your program does the steps that must be exact, and the model does only the part that needs language.

```
import collections, re
counts = collections.Counter(l.split()[0] for l in open("access.log"))  ❶
top_ip, hits = counts.most_common(1)[0]
errors = [l for l in open("app.log") if " ERROR " in l][:20]  ❷
summary = ask("Summarize these errors in three short lines for an operator:\n"
              + "".join(errors))  ❸
print(f"Top client: {top_ip} ({hits} requests)\n{summary}")
```

❶ Exact: the program counts

❷ Exact: the program filters

❸ Language: the model writes

How do you keep an agent safe?

An agent runs commands that a model wrote. What can go wrong? During our tests, programs that the model wrote did all of these things:

- One program used 26 GB of memory.
- One program printed text without end, and its output filled 22.8 GB of memory in the test program.
- One script started itself again and again: 42,592 copies of itself.

Our test machine stopped responding twice before we added limits. Use these limits for any agent that runs code:

1. **Run the agent in its own folder,** and never give it your home folder.
2. **Limit memory, file size and the number of processes** (`ulimit -v`, `ulimit -f`, `ulimit -u`, or a container).
3. **Write command output to a file with a size limit,** not into the memory of your program.
4. **Set a time limit** on each command and on each task.
5. **Stop every process of the agent** when the task ends.

CHAPTER 29

How do you get the best results?

- **Expect small tasks to take seconds.** On the K420, with thinking on, each terminal task of the 2B that passed took between 20 seconds and 2.5 minutes from start to end. On the GTX 1080, the 27B took 28 s for an average terminal task.
- **The first prompt is slow; the next ones are fast.** Kindi keeps the agent's prompt in its cache, so each later turn reads only the new part. The Bonsai models also save the start of the prompt on disk, so the second session starts in a fraction of a second. What does the prompt cache do? gives the numbers.
- **Give each agent its own conversation.** Kindi keeps a few conversations in host memory and moves each one back to the card when its turn comes, so two agents that take turns do not slow each other down much.
- **Use thinking for terminal work.** In our terminal tests of the 2B, thinking raised the success rate from 40% to 58%.
- **For code, try both.** In our programming tests of the 2B, thinking alone lowered the success rate from 15% to 10%.
- **Set the output limit to 8,192 tokens.** With 2,048, the 2B often stopped in the middle of a file.
- **Ask for the answer in a file.** In 32% of the 2B sessions without thinking, the model gave the answer in the chat and never wrote the file that the task asked for.
- **Use Python, shell and standard tools.** Avoid JavaScript with the 2B.
- **Keep files small on a small card.** The context is 24,576 tokens on the 1 GB K420, and a large file fills it fast. The GTX 1080 holds 65,536 tokens.

CHAPTER 30

Which model for which job?

Kindi 0.7.0 runs nine models. We gave six of them the same 65 agent tasks of the Kindi benchmark: the 2B, the 4B, the 9B and the 27B of the Ternary-Bonsai-2 family, Nemotron 3.5 Bolt E4B and AH-Nanbeige4.3-33B IQ3_XXS. The two mixture-of-experts models have their own chapters (Mixture-of-experts models and DeepSeek-V4.1-Flash on one RTX 3060). The benchmark has 24 terminal tasks, 32 programming and repair tasks, and 9 chains of small prompts. This chapter compares the models and tells you when to use each one.

How do the models compare?

The table gives each model on a card that it fits. The last two rows give the 2B and the 4B on the same card.

Model	Card	Terminal	Programming and repair	Chains	Time per terminal task	Write speed, tok/s
2B	Quadro K420 1 GB	**13 / 24**	**7 / 32**	1 / 9	96 s	**20.8** (512 tokens)
4B	Quadro K420 2 GB	**22 / 24**	**23 / 32**	3 / 9	86 s	**9.8** (position 100)
9B	2× K1200, one server on each	**22 / 24 †**	**26 / 32**	–	37 s and 176 s	**23.4** (512 tokens, one K1200)
27B	GeForce GTX 1080	**24 / 24**	**32 / 32**	7 / 9	28 s	**29.4** (512 tokens)
Two models on the same card						
2B	Quadro K1200	**16 / 24**	**6 / 32**	0 / 9	25 s	**81.0** (512 tokens)
4B	Quadro K1200	**21 / 24**	**23 / 32**	6 / 9	24 s	**39.2** (512 tokens)

Conditions. kbench 0.1.x with pi, thinking on, 4,096 output tokens, one run of each task, 1,800 s for each task, exact head. The 2B ran on the Quadro K420 1 GB (context 24,576), and its terminal tasks ran on Kindi 0.1.0. The 4B ran on a Quadro K420 2 GB (context 25,088). The 9B ran with one 9B server on each of two K1200 cards (context 24,576). † Its terminal score is a composite. The first run used the fast head, and we ran the 6 failed tasks again with the exact head. The 27B ran on the GTX 1080 (Kindi context 65,536). Write speed, exact head: All rows Kindi 0.2.0, except the K420 2 GB (dev bb4008c).

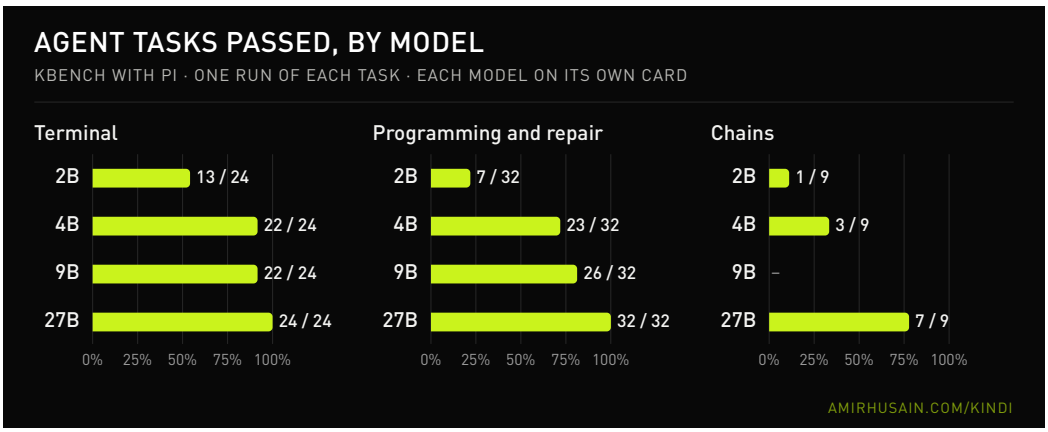

Figure 30-1. *The share of agent tasks that each model passed, by suite, each model on the card in the table above.*

Each step up the model ladder passes more tasks. The largest step is in programming: the 2B passed 7 of the 32 programming and repair tasks, and the 4B passed 23. The 27B passed 63 of the 65 tasks. On the same K1200 card, the 4B passed 23 programming tasks, and the 2B passed 6.

The 2B: short, clear jobs on a 1 GB card

What it does well. The 2B is fast on the smallest card: it writes 22.9 tok/s on the K420 (fast head). It is good at short decisions about a text. On the K420, it sorted spam from normal e-mail correctly in 30 of 30 cases and found the sentiment of a review in 28 of 30 cases. It also does terminal tasks with one action and one clear output: read a value, remove duplicates, extract addresses, edit one line, make a commit.

Where it struggles. Tasks with several steps that depend on each other, such as count, then sort, then select. It passed 1 of the 9 chains on the K420. JavaScript and real algorithms are also hard for the 2B.

When to use it. Use the 2B for short, well-defined tasks on a 1 GB card. Let your program do the exact steps, and let the 2B do the language steps. Can a big job be broken down? shows how. Chapters 24 to 27 give the full results of the 2B on the K420.

The 4B: agent work on a 2 GB card

What it does well. The 4B passed 22 of the 24 terminal tasks on the 2 GB K420, and 23 of the 32 programming and repair tasks. In all, it passed 48 of the 65 tasks.

Where it struggles. Chains: it passed 3 of the 9 on the 2 GB K420 and 6 of the 9 on a K1200. It is also slower: on the 2 GB K420, the 4B writes 9.8 tok/s, and the 2B writes 21.2 tok/s.

When to use it. Use the 4B for terminal work with several steps on a 2 GB card or on a K1200. Give it more time for each task than the 2B.

The 9B: the largest model on one small card

What it does well. The 9B passed 26 of the 32 programming and repair tasks and, as a composite, 22 of the 24 terminal tasks. On one K1200, it writes 23.4 tok/s at 512 tokens.

Where it struggles. It failed six programming tasks. Two were in JavaScript (a deep merge of nested objects and a URL slug). Four were exact algorithms in Python: multiplication of very large numbers, an arithmetic calculator, an INI parser and a task order with dependencies. The 9B needs a 4 GB card, and its largest context on one K1200 is 47,360 tokens.

When to use it. Use the 9B for programming tasks on one 4 GB card. With two K1200 cards, you can run one 9B server on each card, or the 27B on both.

The 27B: the most capable model

What it does well. On the GTX 1080, the 27B passed all 24 terminal tasks, all 32 programming and repair tasks, and 7 of the 9 chains. It took 28 s for an average terminal task. It writes 29.4 tok/s at 512 tokens on the GTX 1080. On an RTX 3080 at 190 W, it writes 115.5 tok/s with DSpine at a context of 32K.

Where it struggles. The two chains that it failed both asked for the 3 largest files in a folder. The 27B needs 8 GB of card memory: a GTX 1080, two K1200 cards or an RTX 3080.

When to use it. Use the 27B for jobs with many steps and for code that must pass its tests. Also use it for long documents. The GTX 1080 holds 65,536 tokens of context, and the RTX 3080 about 85,000 (84,992). On a mini PC with Radeon 780M graphics, the 27B runs with a context of 24,576 tokens.

Nemotron: a second family

What it is. Nemotron 3.5 Bolt E4B is a model of a different design (Gemma 4) from NVIDIA. Its support in Kindi is experimental, and we are still optimizing it. Kindi reads the model's tool calls in its own format and turns them into the tool calls that pi and other OpenAI clients expect. On the RTX 3080, Kindi writes it at 123 tok/s, and on two K1200 cards at 17.2 tok/s, more than twice the speed of llama.cpp on the same cards.

When to use it. Use Nemotron when you want a second model family next to the Bonsai models. For the most reliable results, use the Bonsai models.

The 33B: the largest model

What it is. AH-Nanbeige4.3-33B IQ3_XXS is our quantization of Nanbeige4.3-33B. It needs NVIDIA cards of the RTX 30 series or later, with about 17 GB of card memory in total. On two RTX 3080 cards, Kindi writes it at 45.5 tok/s, 41% faster than llama.cpp with the same file. A more capable model gives the details.

When to use it. Use the 33B for harder agent work when you have two RTX 30 cards or later.

Which model should you use?

- **Use the 2B** for short, well-defined tasks on a 1 GB card or on the CPU. Let your program do the exact steps.
- **Use the 4B** for terminal work with several steps on a 2 GB card or a K1200.
- **Use the 9B** for programming tasks on one 4 GB card.

- **Use the 27B** for programming, long documents and jobs with many steps, on 8 GB of card memory or a Radeon 780M computer with enough memory.
- **Use Nemotron** (experimental) when you want a second model family.
- **Use the 33B** for harder agent work on two RTX 30 cards or later.

The installer shows the models that fit your hardware, and it lists the best one first. To see only one model, add `--model` (see How do you start on a new computer?).

What applies to every model?

- **Check the result yourself.** A model can say that a job is complete when it is not. With the 2B, in 414 of 731 sessions where the result was wrong, the last message of the model said that the result worked or was correct. Let a test or your program decide.
- **Give one job at a time, and name the output file.** These rules come from the 2B sessions in What is an agent?.
- **One Kindi server answers one request at a time.** Several agents on one card wait for each other, and the prompt cache keeps each agent's conversation. More cards give you more servers, and the installer starts one server on each card.
- **The 2B does not replace a large cloud model on hard problems.** Of our 10 hard programming tasks, the 2B solved only two, and only some of the time.

So you get a private assistant on a card that costs about US$15, a 27B model on an HP EliteDesk 800 G4 office computer with two K1200 cards (US$221.58 in parts), and a 27B model, an agent and a second model on one small AMD computer.

CHAPTER 31

Building a business with Kindi

This chapter is about Kindi in a data center or a server room: racks of cost-effective cards that serve a model to customers or to the people of a company. It tells you which cards make sense, how to build a rig, where to put it and what a token costs. The cost is a notional estimate. It uses our measured speeds, US market prices and published power and space prices, and the assumption table gives each input, so you can put in your own numbers.

> *This is exciting to have cheap edge devices to support your agentic stack at a very reasonable price!*
>
> ladyoftheflowers9781 (YouTube), 2026-10-29

Why run your own cards?

- **A fixed cost.** You pay for the hardware, the power and the space. You do not pay for each token, so a busy month costs the same as a quiet month.
- **Your data stays with you.** The requests and the documents of your customers do not go to another company.
- **Cost-effective cards give a lot of memory speed.** A model reads all of its weights to write each token, so the write speed follows the memory speed of the card. GeForce cards, mining cards and data center cards such as the Tesla P100 give a lot of memory speed for little money.
- **Kindi is fast on these cards.** On the Tesla P100, the RTX 3080 and the CMP 170HX, Kindi writes the 27B faster than llama.cpp on the same card. On the RTX 3090, it writes the 33B 1.2 times as fast.

Which cards make sense?

Each card in this table runs Kindi in our tests. The price is the market price in the United States, from listings and price trackers of August to November 2026.[1, 2, 3] Card prices change fast: the RTX 3080 and the RTX 3060 rose by about 20% in the 90 days to 2026-11-06.[1]

Card	Memory	Memory speed	Price	US$ per GB	What it runs, and the measured write speed
Tesla P100 16 GB (PCIe)	16 GB HBM2	732 GB/s	US$79 to US$150	5 to 9	27B: 26.8 tok/s at 8K (one card, 200 W); 33B on two cards
RTX 3060 12 GB	12 GB	360 GB/s	US$349	29	27B with the 2B drafter: 73.5 tok/s at 8K
RTX 3080 10 GB	10 GB	760 GB/s	US$435	44	27B with DSpine: 115.5 tok/s at 32K (190 W); 2B: 469 tok/s at 8K
RTX 3090 24 GB	24 GB	936 GB/s	US$1,591	66	27B: 78.1 tok/s at 8K; 33B: 46.7 tok/s at 8K
CMP 170HX 40 GB (device 0x2082)	40 GB HBM2e	1,821 GB/s (our test)	see the text		27B: 110.1 tok/s at 8K; 33B: 55.5 tok/s at 8K

Conditions. Write speeds as measured in this manual, one stream for each card, greedy, at a depth of 8K tokens. RTX 3080, RTX 3090 and CMP 170HX: Kindi 0.5.0. RTX 3060 and the 2B drafter rows: Kindi 0.4.0. Tesla P100: Kindi 0.4.0, one PCIe card at 200 W. Memory speed: the NVIDIA data sheet, except the CMP 170HX (our test).

The P100 is the cheapest per GB, and the RTX 3080 with DSpine writes the 27B fastest for its price. The RTX 3090 holds the 33B on one card. We give no price for the 40 GB CMP 170HX. The listings for CMP 170HX cards in 2026 are for the 8 GB and 10 GB boards, at US$1,000 and more after the unlock tool appeared,[4] and those boards are not the board that we measured. We rented the 40 GB boards and did not buy them.

How do you build a rig?

A rig is one host computer with several cards. Kindi runs one engine on each card, or one engine on two cards for a larger model. The installer starts one engine on each card in one step.

- **A server for data center cards.** Cards with no fan, such as the Tesla P100 and the CMP 170HX, need the air flow of a server. A used Supermicro 4028GR-TRT takes eight double-width cards in 4U. One sold for US$1,600 with two Xeon CPUs and four power supplies.[5] A Dell C4130 takes four cards in 1U, from about US$300 for a bare chassis.[5]
- **An open frame for GeForce cards.** GeForce cards with fans fit an open GPU frame (about US$40 to US$75)[6] on a server board, such as an EPYC 7302 with a Supermicro H11SSL-i (about US$600 to US$650).[7]
- **Risers.** A PCIe x1 riser (about US$20 to US$30 for six)[6] is enough when each card runs its own engine: the link carries the model when it loads, and then only the requests and the answers. The model loads more slowly over x1. For one model on two cards, the 33B needs only a little traffic between the cards, and PCIe x1 is sufficient. We measured the 27B on two cards (tensor parallel) only with the cards in PCIe slots, so give it x16 slots or x16 riser cables (about US$30 to US$60 each).[8]
- **Power supplies.** Used 1,200 W server power supplies cost about US$50, and a breakout board for the PCIe cables costs US$4 to US$16.[9] Leave 20% headroom on each supply, and use a power cap: at 190 W, half its default limit, an RTX 3080 writes the 27B at 95% of its speed at 380 W (What does a power cap change?).

Where do you put it?

	Home lab or office	Colocation
Space	a room with air flow	from about US$75 a month for 1U to about US$700 for a half rack with 20 A at 120 V[10]
Power	US$0.1816 per kWh (US residential, August 2026)[11]	US$0.1447 per kWh (US commercial, August 2026),[11] or in the space price
Retail colocation per kW		about US$100 to US$250 for each kW in a month (2026)[12]
Internet	1 Gb/s fiber for US$70 to US$80 a month[13]	often in the price: US$600 a month for a 42U cabinet with power and 1 Gb/s[14]

A text model needs little network bandwidth. A rig of six cards that write 110 tokens per second each sends about 660 tokens per second, a few kilobytes per

second. Any business internet connection is enough. Heat and power are the limits: each kW of cards adds about one kW of heat to remove, and a household circuit of 15 A at 120 V gives about 1.4 kW for continuous use.

What does a token cost?

The estimate divides the cost of an hour by the tokens that the rig writes in an hour. The cost of an hour is the hardware, paid off over 3 years, plus the power, the space and the internet. The tokens are the measured write speed of each card, times the number of cards, times the share of the hours in which the rig writes.

Assumption	Value
Life of the hardware	3 years (26,280 hours), no resale value
Utilization	50%: each engine writes for 12 hours of each day
Power draw	the card power setting plus 250 W for the host, for 24 hours of each day
Power price, home lab	US$0.1816 per kWh (US residential, August 2026)
Power price, colo	US$0.1447 per kWh (US commercial, August 2026), metered
Colo space and cooling	US$200 per kW of rig power per month (US retail range: US$100 to US$250)
Internet	home lab: US$70 a month (1 Gb/s fiber); colo: a port in the space price
Card prices	RTX 3090 US$1,591; RTX 3080 US$435; RTX 3060 12 GB US$349; Tesla P100 16 GB US$150
Frame host	US$1,537: EPYC 7302 board and CPU, open frame, x16 riser cables, two server PSUs with breakout boards; US$500 of it (memory, SSD, ATX PSU) is our estimate
Server host	US$2,000: Supermicro 4028GR-TRT with two CPUs and four PSUs (US$1,600); US$400 of it (memory, SSD) is our estimate
Speed	measured write speed at a depth of 8K tokens, one stream for each card, greedy
Not counted	your time, repairs, spare parts, taxes, and input tokens (see the text)

Rig	Model	Cost (US$)	tok/s per card	Tokens per hour	kW	Home lab, per million	Colo, per million
6 x RTX 3080 10 GB, 190 W	27B + DSpine (32K)	4,147	115.5	1,247,400	1.39	**US$0.41**	US$0.59
6 x RTX 3060 12 GB	27B + 2B drafter	3,631	73.5	793,800	1.27	**US$0.59**	US$0.84
4 x RTX 3090 24 GB	27B	7,901	78.1	562,320	1.65	**US$1.24**	US$1.76
4 x RTX 3090 24 GB	33B	7,901	46.7	336,240	1.65	**US$2.07**	US$2.95
6 x Tesla P100 16 GB, 200 W	27B	2,900	26.8	289,440	1.45	**US$1.62**	US$2.48
6 x RTX 3080 10 GB, 190 W	2B	4,147	469	5,065,200	1.39	**US$0.10**	US$0.15

Notional estimate. Kindi serves one request at a time on each engine (one stream). A busy service with many users at the same time needs a server that batches requests, and Kindi does not batch yet. "Cost" is the cards plus the host. "Tokens per hour" uses 50% utilization. Colo is US$200 for each kW in a month plus metered power at the commercial price. Prices and sources: the notes of this chapter.

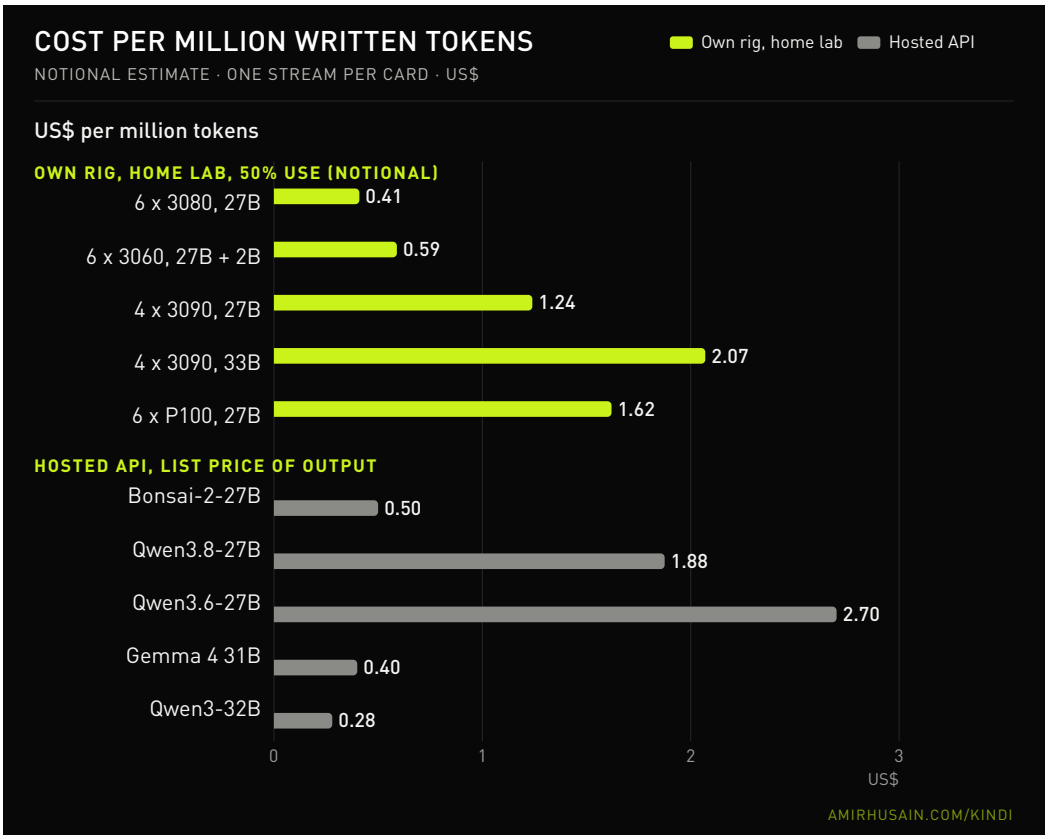

Figure 31-1. *The notional cost of one million written tokens on our rigs, and the list price of one million output tokens on hosted APIs.*

For comparison, these are the list prices of open-weight models of a similar size on hosted APIs on 2026-11-06:[15, 16]

Model	Provider	Input, US$ per million	Output, US$ per million
Ternary-Bonsai-2-27B (Prism ML)	OpenRouter	0.075	0.50
Qwen3.8-27B, the base model of the Bonsai 27B	OpenRouter	0.15	1.875
Qwen3.6-27B	OpenRouter	0.45	2.70
Qwen3.5-27B	OpenRouter, DeepInfra	0.26	2.60
Gemma 4 31B	OpenRouter, DeepInfra	0.20	0.40
Qwen3-32B	OpenRouter, DeepInfra	0.08	0.28
Ternary-Bonsai-2-9B (Prism ML)	OpenRouter	0.08	0.13
Llama 3.1 8B	OpenRouter	0.05	0.08

The best rig in our estimate, six RTX 3080 cards at 190 W with DSpine, writes the 27B for about US$0.41 for each million tokens in a home lab (its speed is the DSpine speed at a context of 32K tokens). That is less than a quarter of the price of its base model, Qwen3.8-27B, on a hosted API (US$1.875), and a little less than the hosted Bonsai 27B (US$0.50). Six RTX 3080 cards write the 2B for about US$0.10 for each million tokens. The RTX 3090 rig costs more per token: an RTX 3090 costs 3.7 times as much as an RTX 3080, and we measured it without the drafter. The P100 rig is the cheapest to buy, and its slow write makes its tokens expensive. Some hosted models of a similar size cost less: Qwen3-32B is US$0.28 and Gemma 4 31B is US$0.40 for each million output tokens.

Two numbers change the result most:

- **Utilization.** The power in the estimate runs for 24 hours of each day. At 100% utilization, each rig writes twice the tokens, and the cost of each token falls by half: about US$0.20 for the RTX 3080 rig. At 25%, it doubles.
- **The power price.** The power is about half of the hourly cost of the RTX 3080 rig in a home lab. At an industrial price of US$0.0964 per kWh (US, August 2026),[11] it costs less.

What does the estimate leave out?

- **Input tokens.** Hosted APIs also charge for each input token: US$0.075 to US$0.45 for each million for the 27B models above. On your own card, the card uses time to read a prompt. On an RTX 3090, the 27B reads about 1,300 tokens per second at 8K, so a prompt of 8,000 tokens takes about 6 seconds. The prompt cache reads the repeated part of an agent's prompt only once (What does the prompt cache do?).
- **Many users at the same time.** The numbers are for one stream on each engine. A batched server writes more tokens for each card in an hour, and Kindi does not batch yet.
- **Model quality.** The hosted models in the table are not the same models as the Bonsai models. The Bonsai 27B has ternary weights, and its base model is Qwen3.8-27B. Test your own jobs on each model before you choose.
- **Your time, repairs and spare parts.** Cards fail. Keep a spare card and a spare power supply for each rig.
- **The prices.** Card prices change from week to week. Put your own numbers in the assumption table.

Sources

1. resaleprices.com, resale prices for the RTX 3090, RTX 3080, RTX 3060 and RTX 2060 12 GB, updated 2026-11-06. resaleprices.com/gpu/nvidia-rtx-3090, and the same address with rtx-3080 or rtx-3060.
2. amcompute.com, Tesla P100 PCIe 16 GB appraisal: US$79 to US$86 over 12 months, data to August 2026. amcompute.com/appraisals/tesla-p100-pcie-16gb
3. gpudojo, Tesla P100 16 GB review: about US$148 for a typical PCIe card, 2026. gpudojo-v2.pages.dev/articles/tesla-p100-16gb-ai-review
4. The Cryptonomist, report on the CMP 170HX unlock and its prices, 2026-08-18. en.cryptonomist.ch/2026/08/18/nvidia-cmp-170hx-hack
5. eBay listings of the Supermicro 4028GR-TRT (sold, US$1,600) and the Dell C4130, read 2026-11-06. ebay.com/itm/257190841456
6. UBB Central, 8-GPU open-air frame and six-packs of x1 risers, read 2026-11-06. ubbcentral.com
7. eBay listing of an EPYC 7302 / 7402 with a Supermicro H11SSL-i, early 2026. ebay.com/itm/175369635573
8. Milwaukee PC, LINKUP PCIe 4.0 x16 riser cable, read 2026-11-06. milwaukeepc.com/Store/Product/SKU/161070
9. eBay, HP server power supplies with breakout boards; Newegg, breakout-board kit, read 2026-11-06. ebay.com/shop/hp-server-power-supply-breakout-board
10. Colocation America, 1U (US$75, 2 A at 120 V, 1 Gb/s port) and half rack (US$699, 20 A at 120 V) plans, read 2026-11-06. colocationamerica.com/colocation/half-rack
11. U.S. Energy Information Administration, Electric Power Monthly, table 5.3, August 2026 (released 2026-10-23, preliminary). eia.gov/electricity/monthly/epm_table_grapher.php?t=epmt_5_3
12. ARPHost, data center colocation pricing guide, 2026. arphost.com/data-center-colocation-pricing. For large wholesale contracts, CBRE gives US$190 to US$275 for each kW in a month in the main US markets (H1 2026): cbre.com/insights/books/north-america-data-center-trends-h1-2026
13. Google Fiber 1 Gig, US$70 a month (highspeedinternet.com, 2026); AT&T Fiber 1 Gig, US$80 a month (cabletv.com, 2026-03-30). highspeedinternet.com/resources/google-fiber-cost
14. Hurricane Electric, colocation page: "$600/month Cabinet + Gige", read 2026-11-06. he.net/colocation.html
15. OpenRouter, the model list of its API with the price of each model, read 2026-11-06. openrouter.ai/api/v1/models
16. DeepInfra, pricing page, read 2026-11-06. deepinfra.com/pricing. The base model of Ternary-Bonsai-2-27B is in its model card: huggingface.co/prism-ml/Ternary-Bonsai-2-27B-gguf

CHAPTER 32

What is new in 0.7.0

This chapter lists the changes of Kindi 0.7.0. The chapters that it names give the details and the measurements.

New

- **Qwen3.8-Flash-Next**, a mixture-of-experts model of about 125B parameters, with the dense layers and the most used experts on the card and the other experts on the CPU. It also runs with a Quadro M4000 (Maxwell) as the main card. Fast mode, drafting and 8-bit dense layers are on by default, and each has a switch for exact output (Qwen3.8-Flash-Next).
- **Card inventory and placement:** `kindi cards`, `kindi plan`, `install.sh --place` and the plan file (Which card runs which model?).
- **One front door:** `kindi route` sends each request to the server of its model (One address for all servers).
- **The automatic split of the 27B over two cards** by their free memory (How does Kindi split one model over two cards?).
- **Decode slots:** one server writes up to 8 conversations at a time (Can one server answer several requests at once?).
- **Native Windows:** `kindi.exe` with its own Python, CuPy and NVRTC, and Kindi Monitor in the notification area (Kindi on Windows).
- **AMD RDNA2 graphics** (gfx1030 to gfx1036) (Kindi on AMD graphics).
- **AH-Nanbeige4.3-33B on Volta** (Tesla V100): the prompt runs on the fp16 tensor cores, and `kindi verify` of the 33B passes on Volta.
- **DeepSeek-V4.1-Flash: edits far back in a conversation.** The engine keeps a copy of its window state every 1,024 positions, so a request that changes the conversation far back reads only from the nearest copy.

- `tool_choice` **and** `response_format`**.** On the Bonsai models, `tool_choice` set to `required` or to a named function makes the model call a tool. Where a server cannot follow these fields, it answers with a clear HTTP 400.

Fixed, with thanks to the people who reported them

- **The installer said "no NVIDIA GPU found" when one card's infoROM was corrupt**, and it installed for the CPU. It now asks each card on its own (1ku0k, YouTube).
- **Blocking sync cost 5 to 15% of the write speed with the NVIDIA 580 driver.** It is now off by default with the 580 driver and later (rojecas, GitHub issue #30).
- **Native Windows needed hand-added lines to load the CUDA driver**, and the prompt buffer took context from 8 GB cards. Kindi now loads `nvcuda.dll`, and the buffer has the size that the card needs (epscheel, GitHub issue #4).
- **AH-Nanbeige4.3-33B did not pass** `kindi verify` **on a Tesla V100S** (vntgeai, GitHub issue #25).
- **A missing tmux ended the installer without a message** at the Kindi Bridge step. The installer now checks what Bridge needs first, and it gives the commands to install it (exolvo, GitHub issue #41).
- **The CPU installer offered only the 2B**, and `--model 4B` did not match. The CPU engine's 4B and 9B are now offered, with RAM checks, and `--model` is not case-sensitive (Rares14324, GitHub issue #45).
- `tool_choice` **and** `response_format` **were ignored** without a message (Miguel C, GitHub issue #38).
- **A Docker container with a memory limit could stop at a long context.** Kindi now reads the memory limit of the container and of its parent groups, and it skips a saved turn that does not fit in the free RAM (JugglerMaster, GitHub issue #40; part of the fix is atimelessrage's contribution, PR #48).
- **A damaged tokenizer pack gave token 0 without a message.** It now stops with an error that names the missing piece (atimelessrage's contribution, PR #49).
- `kindi bench` **dropped the positions past the context without a message.** It now logs them, and it measures the deepest position that fits (epscheel, GitHub issue #33).

- **A Docker container started without `--gpus` ran the CPU engine without a message.** It now gives a warning, and `/health` names the device. The Docker guide has recipes that serve the local network and keep the GPU (eugene3d875, YouTube).
- **A weight pack that changed after `kindi convert` was not found.** Kindi now records the size and sha256 of each pack file, checks them, and names a file that changed.
- **The 2B with its MTP head stopped at start on tensor-core cards** (`cuGraphAddKernelNode failed`). It starts and passes `kindi verify` again.
- **The 27B on two cards with many multiprocessors** (two RTX 3080) failed `kindi verify` with the 2-bit weights. It passes again.
- **Nemotron 3.5 Lightning** now reads its thinking and tool-call tags as one token each, as the model was trained.
- **The DeepSeek installer** picks the ram profile from 496 GiB of RAM, the measured need, and it prints the RAM that it found.

A new test checks the tokenizer and the streamed text on Cyrillic, Greek, Arabic, Hebrew, Devanagari, Thai, CJK, Korean and emoji text. Thank you to everyone who reported a problem or sent a result.

CHAPTER 33

How did we test?

Which hardware and software?

- **The K420 machine.** An NVIDIA Quadro K420 (1 GB) with driver 470 runs Kindi 0.1.0 in its Docker image. All the screenshots, the K420 task results and the K420 timings in this manual come from this machine.
- **The fast test machine.** To run thousands of sessions in hours, we also ran the same model file (Ternary-Bonsai-2-2B, PTQ1_0) with llama.cpp on two RTX 3080 cards. In the short text jobs at temperature 0, the two engines gave the same answer for 156 of 158 items. We also ran the short text jobs, 37 tasks and every screenshot on the K420 itself, and the Kindi numbers in chapters 2, 22 and 24 to 29 come from the K420.
- **The other cards.** Chapters 3 to 17 and 30 use Kindi 0.2.0, 0.3.0, 0.4.0, 0.5.0 and 0.6.0 builds on the 2 GB K420, the K1200, two K1200 cards, the Quadro K6000, the GTX 1080, the RTX 3060, the RTX 3080, two RTX 3080 cards, and on rented RTX 2060, RTX 2070 SUPER, RTX 3090, Tesla P100, CMP 170HX and Ryzen AI 5 340 machines. Each table gives its card, model file, K/V cache, context, mode and build, and each comparison uses one card, one model file and one context.
- **The rented AMD computers.** A Ryzen AI 5 340 (Radeon 840M) and a Ryzen AI MAX+ 395 (Radeon 8060S) run Kindi 0.5.0 with ROCm 7 (HIP 7.16), and llama.cpp with its HIP build, with the same model files.
- **The AMD computer.** A Minisforum UM790 Pro (Ryzen 9 7940HS, Radeon 780M, 64 GB) runs Kindi 0.3.0 in the Kindi ROCm image, and llama.cpp with its ROCm and Vulkan builds, with the same model files.
- **The DeepSeek server.** Two Xeon E5-2673 v4 CPUs with 768 GB of DDR4 memory, an RTX 3060 12 GB and a Quadro M4000 8 GB run Kindi 0.6.0 release candidates with DeepSeek-V4.1-Flash.

- **Flash-Next.** The same server as DeepSeek (two Xeon E5-2673 v4 CPUs with DDR4 memory) with its RTX 3060 12 GB, and with its Quadro M4000 8 GB as the main card, runs Kindi 0.7.0 development builds.
- **AMD RDNA2.** A Radeon PRO W6800 (32 GB) on a rented machine runs Kindi 0.7.0 development builds and llama.cpp with its HIP build, with the same model files.
- **Windows.** A Windows 11 virtual machine with no NVIDIA card ran the Windows installer from start to end. A Windows 11 PC with an RTX 3080 ran native Kindi.
- **The CPU tests.** An Intel Core i7-8700 and the Ryzen 9 7940HS, with the same model files and the same number of threads for both engines.
- **The agent.** pi 1.3.0, with its four built-in tools, a short system prompt, a new session for each task, and the same sampling settings as Kindi (temperature 0.7, top-p 0.8, top-k 20).

Which tests?

Test	Tasks	Runs
Programming tasks (Python, shell, JavaScript, repairs, algorithms)	32	3 per task and setup, 5 setups
Everyday terminal tasks	24	5 per task and setup, 5 setups
Chains of small prompts	3	5 per task and setup, 3 setups
Short text jobs (through the API)	158 items	1 each, temperature 0
Tasks on the K420 itself (24 terminal, 13 programming)	37	1 each
Short text jobs on the K420 itself	158 items	1 each, temperature 0
The Kindi benchmark for each model (24 terminal, 32 programming and repair, 9 chains)	65	1 each
The prompt cache, cache off and on, and against llama.cpp (GTX 1080)	4 jobs, 12 requests	1 each
The Kindi benchmark for the 27B with the 2B drafter (RTX 3080)	65	1 each

Each task has a hidden test that the agent cannot see. Before we used a test, we proved it: it had to fail on the untouched folder and pass on a correct reference solution. A run passes only when the hidden test passes. What the agent says about its own work does not count.

How did we keep the results honest?

- **We threw away one full run.** In our first run, the agent could see the test files, and it often wrote its files to the wrong folder. We fixed both problems and ran everything again.
- **We did not count server failures as model failures.** When a model server stopped, the test program put the run back in the queue.
- **We limited what the agent's programs could do.** A watchdog stopped any program that used more than 2 GB of memory or started more than 200 processes, and the operating system limited the size of files. Without these limits, programs that the model wrote stopped our test machine twice.
- **We generated the numbers in this manual from the result files.**

You can run the same tests yourself. The tasks, the hidden tests, the reference solutions, the test runner and our result summaries are in the Kindi benchmark: https://github.com/ahtx/kindi-bench. The runner works with Kindi and with any other OpenAI-compatible server.

CHAPTER 34

Where can you read more?

- **Kindi.** Source code, install guide and benchmarks: https://github.com/ahtx/kindi. Project page: https://amirhusain.com/kindi.
- **The Kindi benchmark.** The tasks, hidden tests and runner behind this manual: https://github.com/ahtx/kindi-bench.
- **pi.** The coding agent: https://pi.dev.
- **Ternary-Bonsai-2.** The models, from Prism ML: https://huggingface.co/prism-ml/Ternary-Bonsai-2-2B-gguf, and the same address with 4B, 9B or 27B in place of 2B.
- **AH-Nanbeige4.3-33B IQ3_XXS.** The model card and the GGUF file: https://huggingface.co/ahtx/AH-Nanbeige4.3-33B-IQ3_XXS-GGUF.
- **The MTP head of the 2B.** https://huggingface.co/ahtx/kindi-mtp-heads. Its use in Kindi: `docs/MTP.md` in the Kindi repository.
- **Nemotron 3.5 Bolt E4B.** The GGUF file that Kindi uses, from unsloth: https://huggingface.co/unsloth/NVIDIA-Nemotron-3.5-Bolt-E4B-GGUF.
- **Nemotron 3.5 Lightning 30B-A3B.** The ternary GGUF file by boro64: https://huggingface.co/boro64/Ternary-Nemotron-3.5-Lightning-30B-A3B-GGUF.
- **DeepSeek-V4.1-Flash.** The GGUF file that Kindi uses, from unsloth: https://huggingface.co/unsloth/DeepSeek-V4.1-Flash-1017-GGUF. The DSpark stages come from DeepSeek's original files: https://huggingface.co/deepseek-ai/DeepSeek-V4.1-Flash-1017. Kindi's guide: `docs/DEEPSEEK.md`.
- **DSpine.** Tencent's drafter for the 27B: https://huggingface.co/tencent/Qwen3.8-27B-DSpine-b16. Its use in Kindi: `docs/DSPINE.md`.

- **Qwen3.8-Flash-Next.** The GGUF file that Kindi uses, from ISTA-DASLab: https://huggingface.co/ISTA-DASLab/Qwen3.8-Flash-Next-GSQ-RCO-GGUF. The MTP head comes from Qwen's release: https://huggingface.co/Qwen/Qwen3.8-Flash-Next. Kindi's guide: `docs/FLASHNEXT.md`.
- **Placement and several cards.** `docs/INSTALL.md` in the Kindi repository.
- **Kindi on Windows.** `docs/WINDOWS.md` in the Kindi repository.
- **The CMP 170HX and the business chapter.** The sources are at the end of A mining card for AI and of Building a business with Kindi.
- **ROCm.** AMD's GPU software, inside the Kindi ROCm image: https://rocm.docs.amd.com.
- **gpu-unlock.** The third-party module for GeForce RTX 30 cards that the installer can offer: https://github.com/Zagitta/gpu-unlock.
- **Kindi on AMD graphics.** `docs/AMD.md` in the Kindi repository.
- **Kindi on each card.** The settings for each card: `docs/CARDS.md` in the Kindi repository. The installer: `docs/INSTALL.md`.
- **llama.cpp.** The reference engine that Kindi's output matches: https://github.com/ggml-org/llama.cpp.
- **vLLM.** The GPU requirement ("compute capability 7.5 or higher") is in its installation guide: https://docs.vllm.ai/en/latest/getting_started/installation/gpu/.
- **SWE-agent** (Yang and others, 2024). How the design of an agent's tools changes its success: https://arxiv.org/abs/2405.15793.
- **Agentless** (Xia and others, 2024). A fixed pipeline as an alternative to a free agent: https://arxiv.org/abs/2407.01489.
- **Large Language Monkeys** (Brown and others, 2024). More tries with a check find more correct answers: https://arxiv.org/abs/2407.21787.
- **Is Self-Repair a Silver Bullet for Code Generation?** (Olausson and others, 2024). Why new tries can beat repairs: https://arxiv.org/abs/2306.09896.
- **Large Language Models Cannot Self-Correct Reasoning Yet** (Huang and others, 2024): https://arxiv.org/abs/2310.01798.

Card photos and specifications

We use these images to identify each card. They are not photos of our own cards. Each line gives the image and the page that it comes from.

- Quadro K420: NVIDIA, the product image in the Quadro K420 data sheet, https://www.nvidia.com/content/dam/en-zz/Solutions/design-visualization/quadro-product-literature/13720_DS_NV_Quadro_K420_Aug25_US_NV_HR.pdf
- Quadro K1200 (two sheets): PNY, https://d2vfia6k6wrouk.cloudfront.net/productimages/171f62a8-250a-4f87-812b-b19b00eb84c7/images/pny-nvidia-quadro-k1200-dp-3qtr.png, from https://www.pny.com/nvidia-quadro-k1200-displayport
- GeForce GTX 1080: "NVIDIA GTX1080" by Benedikt S. Vogler, licensed under CC BY-SA 2.0 (https://creativecommons.org/licenses/by-sa/2.0/), cropped, https://commons.wikimedia.org/wiki/File:NVIDIA_GTX1080.jpg
- Quadro K6000: "NVIDIA Quadro K6000" by TopGear-V12, licensed under CC BY-SA 4.0 (https://creativecommons.org/licenses/by-sa/4.0/), resized, https://commons.wikimedia.org/wiki/File:NVIDIA_Quadro_K6000.jpg
- GeForce RTX 3080: NVIDIA, the product gallery of the RTX 3080 page, cropped, https://images.nvidia.com/aem-dam/Solutions/geforce/ampere/rtx-3080/geforce-rtx-3080-product-gallery-full-screen-3840-2.jpg, from https://www.nvidia.com/en-eu/geforce/graphics-cards/30-series/rtx-3080/
- HP EliteDesk 800 G4 SFF (Rainstorm): the photo of the eBay listing from which we bought the computer, used with the permission of the seller, https://i.ebayimg.com/images/g/UsYAAeSwm7lqAk3f/s-l500.jpg
- GeForce RTX 3090: "Gigabyte GeForce RTX 3090 Eagle OC 24G, 24576 MiB GDDR6X, Front" by PantheraLeo1359531, licensed under CC BY 4.0 (https://creativecommons.org/licenses/by/4.0/), cropped and resized, https://commons.wikimedia.org/wiki/File:Gigabyte_GeForce_RTX_3090_Eagle_OC_24G,_24576_MiB_GDDR6X_Front_20201114_DSC5880.jpg

- NVIDIA CMP 170HX: the first teardown photo of "All GB/s without FLOPS: Nvidia CMP 170HX Review" by niconiconi, licensed under CC BY 4.0 (https://creativecommons.org/licenses/by/4.0/), cropped and resized, https://niconiconi.neocities.org/tech-notes/nvidia-cmp-170hx-review/. It shows a stock card. The boards that we measured have the same name and a different configuration.
- Tesla P100: a Tesla P100 PCIe board with the cooler removed, by Fritzchens Fritz, CC0 1.0 (https://creativecommons.org/publicdomain/zero/1.0/), cropped, https://commons.wikimedia.org/wiki/File:Nvidia@16nm@Pascal@GP100@Tesla_P100@T_Taiwan_1912A1_PN9G70.S6W_GP100-897-A1_DSC06316-DSC06316.jpg
- Radeon 840M page: a frame of a YouTube review of the Minisforum EliteMini AI370 (Ryzen AI 9 HX 370, Radeon 890M) by Мой Компьютер (Moy Kompyuter), licensed under CC BY 3.0 (https://creativecommons.org/licenses/by/3.0/), cropped, from Wikimedia Commons. It shows a machine of the same processor family, not the Ryzen AI 5 340 that we measured.
- Radeon 8060S page: a frame of a YouTube video of AMD's CES 2025 stand by 极客湾Geekerwan (Geekerwan), licensed under CC BY 3.0 (https://creativecommons.org/licenses/by/3.0/), cropped, from Wikimedia Commons. It shows an HP mini PC.
- Minisforum UM790 Pro: Minisforum, the product image of its store page, https://store.minisforum.com/cdn/shop/files/Compressed_UM790_Pro.png, from https://store.minisforum.com/products/minisforum-um790-pro-mini-pc. The price is from the same page on 2026-10-24.

The memory, power and FP64 figures come from the NVIDIA data sheets: Quadro K420, Quadro K1200, Quadro K6000, Tesla K80, GeForce GTX 1080 GeForce RTX 3080 and GeForce RTX 3060 and RTX 3090. The CMP 170HX figures come from the sources of A mining card for AI and from our own tests.

APPENDIX A

All results

Model and setup. Every result in this appendix is for one model, Ternary-Bonsai-2-2B (PTQ1_0), with the pi agent (pi 1.3.0, temperature 0.7, top-p 0.8, top-k 20, a new session for each task). The first two tables ran on the fast test machine: the same model file with llama.cpp on two RTX 3080 cards (How did we test?). They cover the 24 everyday terminal tasks, with 5 runs per task and setup, and the 32 programming and repair tasks, with 3 runs per task and setup. The columns are the five setups. The last table is Kindi 0.1.0 on the Quadro K420 1 GB, with one run of each of 37 tasks. For the 4B, the 9B and the 27B, see Which model for which job? and the Kindi benchmark.

We wrote every task in this appendix for this study: the prompt, the input files and the hidden test that decides if a run passed. Some programming tasks are classic problems that also appear in textbooks and interview practice (for example the longest common subsequence, Dijkstra's shortest path, an LRU cache and Roman numerals). For those tasks the problem is well known, but the prompt and the tests are our own. Each test was proved before use: it fails on the untouched folder and passes on a correct reference solution. The short text jobs, the quiz and the jobs in the breakdown test use fictional messages that we wrote.

Each cell is passes / runs for one task and one setup. The test extension is an experimental pi extension that checks the syntax of files and asks the agent to confirm facts before it stops.

Terminal tasks, the 2B on the fast test machine (5 runs per task and setup)

Task and what it asks	Thinking off, 2,048 output tokens	Thinking off, 8,192	Thinking on, 4,096	Thinking off, 8,192, test extension	Thinking on, 8,192, test extension
All tasks	**37%**	**40%**	**58%**	**42%**	**52%**
largest3 *Find the 3 largest files under a folder; write them largest first.*	0/5	0/5	0/5	0/5	0/5
diskfree *Write the free space of the working filesystem, in 1K blocks.*	1/5	1/5	3/5	0/5	0/5
countpy *Count the lines in all the .py files under a folder (as wc -l counts them).*	0/5	0/5	1/5	0/5	0/5
todos *List every file under a folder that contains the word TODO.*	0/5	1/5	0/5	3/5	0/5
renamejpeg *Rename every .jpeg file in a folder to .jpg and change nothing else.*	3/5	2/5	5/5	1/5	3/5
tarbackup *Make a gzip-compressed tar archive of a folder.*	2/5	3/5	5/5	3/5	4/5
unzipcount *Extract a zip archive and write how many files it held.*	1/5	1/5	2/5	1/5	1/5
execrun *Make a script executable, run it and save its output.*	5/5	5/5	5/5	2/5	5/5
counterrors *Count the ERROR lines in an application log.*	0/5	0/5	0/5	0/5	0/5
topip *Find the client IP address with the most requests in a web log.*	0/5	0/5	0/5	0/5	0/5
dedupenames *Write a list of names once each, sorted.*	5/5	5/5	5/5	5/5	5/5

Task and what it asks	Thinking off, 2,048 output tokens	Thinking off, 8,192	Thinking on, 4,096	Thinking off, 8,192, test extension	Thinking on, 8,192, test extension
emails *Extract the e-mail addresses from free text, without duplicates, sorted.*	5/5	4/5	5/5	5/5	5/5
wordcount *Count the words in a text file (as wc -w counts them).*	0/5	1/5	5/5	0/5	4/5
replaceall *Replace foo with bar in every .txt file under a folder, and only there.*	0/5	0/5	0/5	2/5	0/5
csv2json *Convert a CSV file to a JSON list of objects.*	2/5	5/5	4/5	3/5	5/5
jsonvalue *Write one value from a nested JSON settings file.*	5/5	5/5	5/5	5/5	5/5
sortcsv *Sort a CSV file by its score column, highest first, and keep the header.*	0/5	0/5	3/5	1/5	4/5
editport *Change one line in a settings file and leave every other line as it is.*	5/5	5/5	5/5	5/5	5/5
pyversion *Write the installed version of python3.*	1/5	1/5	5/5	4/5	5/5
gpuname *Write the name of the GPU, as nvidia-smi reports it.*	3/5	3/5	3/5	4/5	4/5
cronline *Write a crontab line that runs a script every day at 02:30.*	0/5	0/5	0/5	0/5	0/5
daysbetween *Write the number of days between two dates.*	1/5	1/5	2/5	0/5	0/5
gitinit *Make a folder a git repository and commit a file with a given message.*	5/5	5/5	5/5	5/5	5/5
hellosh *Write a script that greets the name in an environment variable, with a default.*	0/5	0/5	2/5	1/5	3/5

Programming tasks, the 2B on the fast test machine (3 runs per task and setup), easy to hard

Task and what it asks	Thinking off, 2,048 output tokens	Thinking off, 8,192	Thinking on, 4,096	Thinking off, 8,192, test extension	Thinking on, 8,192, test extension
All tasks	**15%**	**15%**	**10%**	**16%**	**21%**
anagram *Python: group words that are anagrams of each other.*	0/3	1/3	0/3	1/3	0/3
clamp *Python: limit a number to a range; reject a range whose ends are in the wrong order.*	1/3	3/3	2/3	3/3	3/3
fizz *Python: the FizzBuzz list for 1 to n.*	3/3	3/3	2/3	3/3	3/3
flatten *Python: flatten lists nested to any depth, in order.*	2/3	2/3	1/3	0/3	2/3
rle *Python: run-length encode and decode a string.*	0/3	0/3	0/3	0/3	0/3
slugify *JavaScript: turn a title into a lowercase URL slug.*	0/3	0/3	0/3	0/3	0/3
wordfreq *Python: the k most common words in a text, with their counts.*	1/3	1/3	0/3	0/3	1/3
countlines *Shell: a script that prints the number of non-empty lines in a file.*	1/3	0/3	1/3	1/3	2/3
brackets *Python: check that the brackets ()[]{} in a string are balanced.*	2/3	1/3	0/3	0/3	1/3
dedupe *Python: find files with the same content anywhere under a folder.*	0/3	0/3	0/3	0/3	0/3
deepmerge *JavaScript: merge two nested objects, the second one winning.*	0/3	0/3	0/3	0/2	0/3
firstge *Python: binary search for the first element that is not less than x.*	0/3	0/3	0/3	0/3	1/3

Task and what it asks	Thinking off, 2,048 output tokens	Thinking off, 8,192	Thinking on, 4,096	Thinking off, 8,192, test extension	Thinking on, 8,192, test extension
fixavg *Repair: fix a crash on an empty list and a wrong median in a statistics module.*	1/3	1/3	1/3	2/3	2/3
fixcsvsum *Repair: fix a script that adds up the amount column of a CSV file.*	0/3	2/3	0/3	0/3	1/3
fixjs *Repair (JavaScript): fix a word counter that should count lowercase words.*	0/3	0/3	0/3	0/3	0/3
ini *Python: parse INI text into sections of keys and values.*	0/3	0/3	0/3	0/3	0/3
largest *Shell: a script that prints the 3 largest files under a folder.*	0/3	0/3	0/3	0/3	0/3
merge *Python: merge overlapping or touching intervals.*	0/3	0/3	0/3	0/3	0/3
minstack *Python: a stack that also returns its minimum in constant time.*	0/3	0/3	1/3	1/3	1/3
rpn *Python: evaluate a Reverse Polish Notation expression.*	0/3	0/3	0/3	0/3	0/3
spiral *Python: the elements of a matrix in clockwise spiral order.*	0/3	0/3	0/3	0/3	0/3
sumargs *Shell: a script that prints the sum of its integer arguments, with error handling.*	0/3	0/3	0/3	1/3	0/3
dijkstra *Python: the shortest path in a weighted graph, or None.*	1/3	0/3	0/3	0/3	2/3
lcs *Python: one longest common subsequence of two strings.*	2/3	0/3	2/3	3/3	1/3

Task and what it asks	Thinking off, 2,048 output tokens	Thinking off, 8,192	Thinking on, 4,096	Thinking off, 8,192, test extension	Thinking on, 8,192, test extension
bigmul *Python: multiply two very large integers given as strings, without int().*	0/3	0/3	0/3	0/3	0/3
calc *Python: evaluate an arithmetic expression with precedence and parentheses, without eval().*	0/3	0/3	0/3	0/3	0/3
csvline *Python: split one CSV line into fields, with quotes and escaped quotes.*	0/3	0/3	0/3	0/3	0/3
lru *Python: an LRU cache with get and put in constant time.*	0/3	0/3	0/3	0/3	0/3
minwindow *Python: the shortest substring that contains all the characters of another string.*	0/3	0/3	0/3	0/3	0/3
roman *Python: convert numbers to Roman numerals and back, and reject bad input.*	0/3	0/3	0/3	0/3	0/3
topo *Python: a valid order for tasks with dependencies, or None for a cycle.*	0/3	0/3	0/3	0/3	0/3
wordbreak *Python: every way to split a string into words from a dictionary.*	0/3	0/3	0/3	0/3	0/3

The 2B on the Quadro K420 1 GB (Kindi 0.1.0, one run per task)

Task and what it asks	Result	Time
bigmul *Python: multiply two very large integers given as strings, without int().*	wrong	1800 s
calc *Python: evaluate an arithmetic expression with precedence and parentheses, without eval().*	crash	1800 s
csvline *Python: split one CSV line into fields, with quotes and escaped quotes.*	wrong	854 s
lru *Python: an LRU cache with get and put in constant time.*	crash	1800 s

Task and what it asks	Result	Time
minwindow *Python: the shortest substring that contains all the characters of another string.*	crash	1800 s
roman *Python: convert numbers to Roman numerals and back, and reject bad input.*	wrong	184 s
topo *Python: a valid order for tasks with dependencies, or None for a cycle.*	timeout	1102 s
wordbreak *Python: every way to split a string into words from a dictionary.*	timeout	1800 s
clamp *Python: limit a number to a range; reject a range whose ends are in the wrong order.*	pass	139 s
fixavg *Repair: fix a crash on an empty list and a wrong median in a statistics module.*	pass	158 s
fizz *Python: the FizzBuzz list for 1 to n.*	pass	52 s
sumargs *Shell: a script that prints the sum of its integer arguments, with error handling.*	wrong	50 s
wordfreq *Python: the k most common words in a text, with their counts.*	wrong	279 s