Episode Book

EPISODE #6513

Why Dual-Socket Servers Win on Dense AI Models

PUBLISHED	9 November 2026 at 20:26 UTC
DURATION	26:03
LISTEN	myweirdprompts.com/episode/dual-socket-cpu-workloads-37877/
HOSTS	Corn, Herman, Dorothy, Hilbert
SCRIPT MODEL	z-ai/glm-5.3
VOICES	chatterbox-regular (chatterbox-local)
PIPELINE	V5.26
SOURCES	25 cited by the research (listed at the end)
BOOK GENERATED	10 November 2026 at 10:05 UTC

SYNOPSIS

Adding a second CPU socket offers extra memory channels and PCIe lanes, but non-uniform memory access dramatically shapes real-world performance. A detailed look at Intel Xeon and AMD EPYC platforms reveals why dense neural networks scale up while fine-grained mixture-of-experts architectures stall on cross-socket interconnects. Also covered: BIOS nodes-per-socket settings, the economics of decommissioned server hardware, and how memory placement bugs distort hardware benchmarks.

TRANSCRIPT



Corn

GPUs have got so expensive that the humble CPU is having a comeback tour, and Daniel noticed. His prompt this week points at a corner of the hardware world nobody talks about anymore: two CPUs in one machine. Two GPUs, endless videos. Two CPUs, silence. He wants to know how dual-socket actually works, what platforms and boards still support it — he explicitly says used Xeon and EPYC servers included, so he's been browsing eBay — and what a second socket really gets you compared with one bigger CPU or a second machine. And where it's a trap. He's thinking about a self-hosted build runner, which, spoiler, turns out to matter a lot.



Herman

It matters enormously, because the build runner is almost the worst possible use of a dual-socket box. But we'll get there.



Corn

Before we can judge whether a second socket is worth it, we need to be precise about what a second socket physically is. So — what is it?



Herman

Two identical CPUs on one motherboard, sharing one main memory, run by one operating system that treats all the cores as a single pool. That's symmetric multiprocessing. You do not get two kernels. You do not get two page caches, two schedulers, two computers. The OS boots once and sees one machine with a very large number of cores.



Corn

Which is the first misconception Daniel's prompt brushes against. People hear "two CPUs" and picture two computers welded together, one for the database, one for the builds.



Herman

Right up until the memory part bites them. Because the memory is shared, but it isn't uniform. Each CPU owns the DIMMs wired directly to it. A thread running on socket one that reaches memory attached to socket two pays a latency penalty crossing the interconnect. That's NUMA — non-uniform memory access. Local memory is fast, remote



Herman (continued)

memory is a trip.



Corn

The local DIMMs and the ones across the hall.



Herman

And that single fact explains almost every good and bad report you'll ever read about dual-socket machines. The ones where somebody got one point eight times the performance from the second CPU and the ones where they got one point zero two. Same hardware, different memory behavior. The through-line for today, I'd say, is that the second socket is a capability purchase — cores, memory channels, PCIe lanes — and the question is whether your workload needs that capability, or whether you're just inheriting an assumption from the era when dual-socket was simply what a server was.



Corn

Twenty years of IT managers buying two sockets because that's what a server looked like. Muscle memory dressed as architecture.



Herman

So let's get into how the thing is actually wired, because the wiring is where all the consequences live.



Corn

Go. The interconnect.



Herman

The socket-to-socket link is the whole story. Intel's is UPI, AMD's is Infinity Fabric, also called xGMI at the socket level. Per link, Intel UPI runs at ten point four gigatransfers per second, which is about twenty-one gigabytes per second in one direction, and a Skylake-generation Xeon uses three of those links between sockets. AMD's Zen two and



Herman (continued)

Zen three fabric runs at eighteen gigatransfers, seventy-two gigabytes per second, and the socket links run at eight ninths of that because of the error checking. On EPYC Rome, three of those links between the sockets give you roughly ninety-six gigabytes per second theoretical, socket to socket.



Corn

Ninety-six gigabytes per second sounds enormous until you remember local memory bandwidth is several times that.



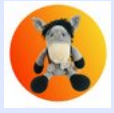
Herman

It's the slow road, yes. And here's the design decision I find elegant: on modern EPYC, the socket link and the PCIe lanes share the same physical wiring. The PHY is the same, the protocol on top differs. So a board designer can configure a dual-socket EPYC system with a hundred and sixty PCIe Gen four or five lanes plus three xGMI links, or a hundred and twenty-eight lanes plus four xGMI links.



Corn

Lanes or bandwidth. Pick one.



Herman

And nearly every major vendor ships the hundred-and-sixty-lane configuration with only three interconnect links, because in ServeTheHome's words, so many customers found the socket-to-socket bandwidth sufficient with three xGMI links. Which tells you something about real-world traffic: most workloads keep their data local, and the customers who didn't were rare enough to design around.



Corn

Here's the detail I want to land, because it explains why the folk wisdom on dual-socket is so unreliable. The NUMA topology is a BIOS setting.



Herman

It is! AMD's NPS setting — Nodes Per Socket — subdivides one socket's memory into one, two, or four NUMA nodes. The llama.cpp testing we'll talk about later ran NPS one, one node per socket, which is the natural dual-socket setup. But in the same thread, someone tested a single-socket Genoa workstation configured at NPS two specifically to emulate a dual-socket machine with a fast interconnect. One chip, in the BIOS, and you've built a different architecture.



Corn

So two machines that both say "dual EPYC" on the listing can have different memory topologies, and a benchmark that doesn't mention the NPS setting is a benchmark that told you half the story.



Herman

And that's before you get to the workload's own memory placement, which is the other half. Hold that thought.



Corn

Holding it. Now the platform question, because Daniel asked what you can actually buy, and the first answer is a flat no for anything consumer.



Herman

Flat no. Ryzen and Threadripper are single-socket, full stop. Multi-socket motherboards haven't been offered on consumer platforms for roughly fourteen years. There was a report in twenty twenty-one that five-thousand-series Threadripper Pro would gain dual-socket functionality, and it just never turned into shipping boards. Threadripper Pro remains a one-socket product. If you want two CPUs you are in server silicon: Intel Xeon or AMD EPYC.



Corn

The consolation prize is that server silicon has a used market like nothing else in computing. Give people the prices, Herman, because these are the numbers that pull people into this rabbit hole.



Herman

They're remarkable. A used Xeon Gold 6138 — that's the Skylake generation, twenty cores, forty threads — sells for thirty-six to forty-seven dollars. Per CPU. Refurbished pairs go for a hundred and twenty-six dollars. So for a bit over a hundred bucks you can hold forty cores of 2017-vintage silicon in your hand.



Corn

Two lattes per Xeon.



Herman

Then the whole-system prices: a used Dell PowerEdge R740 with two of those Gold 6138s installed, forty cores total, about nine hundred and fifty dollars. In the UK there's an R740xd with the same CPUs and sixty-four gig of DDR4 for six hundred pounds. So the used dual-Xeon box is real, and it's cheap, and the reason it's cheap is that an entire generation of corporate servers is being decommissioned at once.



Corn

Which is also the reason the electricity meter in that flat should fear it.



Herman

We'll get to power. The AMD side: a used ASRock Rack GENOAD8X board — that's the SP5 socket, current-generation EPYC, DDR5 — eight hundred and ninety-five dollars used on eBay against sixteen seventy-three new. New dual-socket boards: ASRock's Turin boards around eleven hundred, the Gigabyte MZ73-LM2 at eighteen ninety-nine. And if you want the full experience, a refurbished ASRock two-socket Genoa barebone, a hundred and twenty-eight cores, ninety-day warranty, forty-five hundred dollars.



Corn

One honest gap here. The Intel picture is all used Skylake and one Emerald Rapids part that AMD cited in a blog. There's no fresh survey of dual-socket Xeon 6 boards and prices.



Herman

Right — I haven't seen current-generation Xeon 6 dual-socket pricing anywhere, so I won't pretend to. What I can say is that AMD has quietly built a two-tier lineup that tells you exactly what they think of dual socket. Most EPYC 9005 Turin chips are marked one-P, two-P — they'll work in either. But the P-suffix models, the 9355P, 9455P, 9555P, 9655P, are single-socket only. Same silicon, sold cheaper, for the single-socket market. Same pattern in the Genoa generation. And the 8004 Sienna parts are, in AMD's own words, designed as a single-socket solution.



Corn

The P is the chip telling you its market moved on.



Herman

There's no P model above ninety-six cores this generation either. The big dual-socket chips exist, but the volume, the pricing energy, the marketing — it's all one socket. That's the landscape. Now the harder question: what does the second socket actually buy you, and where does it betray you?



Corn

Start with the retraction, because it's the best story in this whole subject.



Herman

It really is. OpenClawDC published CPU inference speeds in August, comparing two machines, and concluded dual socket was slower than single socket. Then they went back and did it properly — a controlled A/B on the same physical machine, socket count the only variable. And the conclusion flipped: the second socket delivered one point eight three times the speed on a dense seventy-billion-parameter model. And one point zero two on DeepSeek R1. Basically nothing.



Corn

Same hardware. One model, near-double. The other, a rounding error.



Herman

Their verdict, and I think it's the single most quotable line in this space: "Dual socket is slower" is usually a placement bug, not a hardware verdict. The earlier comparison was two different machines with different memory configurations — which is to say, nearly worthless. And a lot of the folk wisdom out there comes from exactly that kind of comparison.



Corn

The controlled numbers back it up. Same machine, llama.cpp testing: a dense Llama three point one seventy-billion model in full FP sixteen on two EPYC 9175Fs went from two point three five to four point three zero tokens per second in generation when you populated the second socket. Call it one hundred and eighty-three percent.



Herman

Mixtral eight by twenty-two billion, a mixture-of-experts model, on two 9654s: generation went from seven to just over ten tokens per second. A hundred and forty-six percent. Still a good gain, visibly less than the dense model. DeepSeek R1 on the same hardware: eight point four eight to eight point six seven. A hundred and two percent. The second socket was decoration.



Corn

And there's a rule that predicts all three numbers before you run them. It's matrix size.



Herman

Fairydreaming's formulation in the llama.cpp thread: the smaller the multiplied matrix, the higher the relative synchronization and communication overhead from dual CPU usage. The feed-forward matrices in the dense seventy-billion Llama are eight thousand one hundred ninety-two by twenty-eight thousand six hundred seventy-two — four hundred and forty-eight megabytes of matrix data. Per expert in Mixtral, ninety-six megabytes. Per expert in DeepSeek R1, roughly seven megabytes.



Corn

Seven megabytes of work per conversation across the hallway. That's a lot of trips.



Herman

Measured directly: synchronization consumed fifteen point four percent of computation time on the dual-socket system versus nine point one percent on single socket. The overhead is real, it's per trip, and it doesn't shrink with the work. Big chunks of work amortize the trip. Small chunks drown in it.



Corn

And now, a word from our very own Dorothy in the kitchen.

D

Dorothy

Hello dear, it's your mother. It's Dorothy's Kitchen again.



Herman

Mum! Wonderful.

D

Dorothy

Today I want to give you my recipe for a cracker with nothing on it. The ingredients are one cracker. That's all, sweetheart. No butter, no cheese, nothing. That's the whole point, and I want you to take it seriously.



Corn

She always takes it seriously. That's why we do.

D

Dorothy

First you take the cracker out of the packet. Gently. The packet can tear and then you've got crumbs everywhere, and I won't go into what happened the last time.



Herman

Was the last time the rye incident?

D

Dorothy

Then you put it on a plate, and you watch it. You watch it for a good while, because if you turn your back, the butter in the fridge starts to think about things, and we don't want that.



Corn

Nobody wants thinking butter.

D

Dorothy

Then you eat the cracker. Nothing on it. It will be dry, dear, and it will be a bit disappointing, and that's the recipe working exactly as intended. Cooking is mostly about bracing yourself.

D

Dorothy

Anyway. Are you eating properly? Put a jumper on. I'll ring off now, the kettle's on and I've lost the tea again.



Corn

Back to the sockets. Daniel's other comparison — one bigger CPU. The evidence there is uncomfortable for team dual.



Herman

It's the strongest single data point in the whole topic, I think. Dell ran a benchmark with a single fifth-generation EPYC 9575F against a dual fourth-generation 9554, on TPC-H analytics queries, and the single socket delivered a one point six times average latency improvement — a thirty-eight percent reduction — over the dual-socket box. One query gained two point two eight times.



Corn

One socket beating two.



Herman

Now, the honest framing, and I want to be careful here: that's a generation-over-generation comparison. New chip versus old chips. It's not a controlled socket-count test like the llama.cpp one. But that's exactly the comparison a real buyer makes — do I buy two of the previous thing or one of the new thing — and the new thing keeps winning it. AMD's own position, from their blog titled, charmingly, "2P or not 2P": moving from two sockets to one does not always mean giving up performance or features, most workloads run great on one socket, and the first benefit they list is no inter-socket communication bottlenecks.



Corn

The vendor that sells the two-socket chips, telling you to buy one. That's about as clear a market signal as you'll ever get.



Herman

Daniel's third comparison is the second machine, and the argument there is the classic vertical-versus-horizontal scaling one. The case for one big box: if you can fit your job on one machine, plus spares, you save a whole lot of complexity versus spreading it over several. Same person, elsewhere: if you're hitting bottlenecks with the biggest off-the-shelf computer you can get, scale out instead. The second socket is the last rung on the one-box ladder — the way to stay on one machine past the point where one CPU's memory channels and lanes run out.



Corn

And that gets us to what I think is the real answer to Daniel's "what do you really gain" question. It's not speed. It's DIMM slots.



Herman

A dual-socket SP5 board has twenty-four DIMM slots — twelve per socket — and a hundred and twenty-eight PCIe five point zero lanes. One person building one of these put a terabyte of RAM in a dual-socket EPYC for around three thousand dollars, sixteen sixty-four-gig DIMMs, with considerable memory bandwidth to spare. And the line that



Herman (continued)

goes with it, from the same people who did the retraction: "Then the second CPU is a memory-capacity purchase, not a speed purchase."



Corn

A terabyte for three grand. The second socket is the coupon that comes with the RAM.



Herman

The trap, then, is buying two sockets for compute when you should have bought one socket and more memory. The sockets get the glory, the DIMMs do the work.



Corn

Now Daniel's build runner. And this is where the whole thing turns around, because every piece of CI evidence points away from dual socket.



Herman

Every piece. Compilation is dominated by single-thread performance for the serial steps — configure, link, the dependency chain where job two waits on job one. Ubicloud tested this directly: a Ryzen 9 7950X3D beat an EPYC 9454P by about thirty-two percent single-core, and on real GitHub Actions workloads the median job ran in eighty seconds on the Ryzen versus a hundred and twenty-nine on the EPYC.



Corn

A desktop chip beating a server chip by half a minute per job, per run, forever.



Herman

And remember what GitHub's own runners actually use: an EPYC 7763, which came out in early twenty twenty-one. A current Ryzen 9950X3D scores roughly four thousand seven hundred single-thread against twenty-five hundred for that EPYC — nearly double the serial speed. Kirill Korins did the OpenBSD build test that maps it perfectly: a Ryzen 9 7950X3D was about twice as fast as an EPYC 7502P when the job count was low, and the



Herman (continued)

gap only narrowed to about a third when he saturated every core.



Corn

So unless your builds are embarrassingly parallel — and most CI pipelines are not, they're chains — one fast socket beats two slow ones, and three slow ones, and however many you can fit in a four-unit chassis.



Herman

With the caveat that we're inferring this. Nobody has published a dual-socket-specific compilation benchmark. All the build-runner data is single-socket Ryzen versus single-socket EPYC. The dual-socket scaling data we have is LLM inference, which is memory-bandwidth-bound — a different regime from compilation. But the inference all runs one direction, and the mechanism — serial steps, small units of work, cross-socket penalty — is exactly the shape that punishes dual socket.



Corn

Then there's the invisible second socket. Power and licensing.



Herman

The money nobody prices in. Servnet UK put together indicative numbers at equal usable core count: single socket, nine thousand pounds of hardware, three thousand a year in power, fourteen thousand in per-core licensing. Dual socket: fifteen thousand hardware, five thousand power, twenty-two thousand licence.



Corn

Six thousand pounds more to buy, two thousand a year to run, eight thousand a year to license. For the same cores.



Herman

Same cores, more sockets, and per-core software doesn't care that you spread them across two packages. AMD's SQL Server comparison ran ten million analytics queries across three thirty-two-core configurations — dual Xeon, dual EPYC, single EPYC — and claimed over half a million dollars of SQL Server savings, up to eight hundred and eighty-one thousand, from right-sizing the sockets.



Corn

Even on the cheap used Xeon box, the pattern holds. The CPUs cost eighty dollars. The chassis, the DIMMs you now have to buy in matched pairs to populate both sockets symmetrically, and the idle draw — that's where the money actually goes.



Herman

An R740 doesn't idle at nothing, and it doesn't idle quietly either.



Corn

Forty cores of 2017 listening to themselves hum.



Herman

So the verdict, as best the evidence supports it. Dual socket makes sense for dense VM consolidation, large in-memory databases, lane-hungry GPU and NVMe builds, anything that needs more than eight to twelve DIMM slots. It's a trap for speed on small-grained workloads, for untuned NUMA, for a single latency-sensitive application, and for anything per-core licensed. And for a self-hosted build runner — the thing Daniel actually asked about — it's the wrong tool, close to categorically.



Hilbert

That's what I kept telling people.



Corn

Did they listen?



Hilbert

No. I had a dual-socket machine here for years, Herman, and I spent more time than I'll admit trying to make it behave like two computers. The one OS seeing everything as one pool bothered me. I wanted half of it to mind its own business.



Herman

So you partitioned it?



Hilbert

I tried. Affinity, memory pinning, the lot. In the end I gave up and did what I should have done at the start, which was draw the topology by hand. I taped the map inside the case lid. Every DIMM slot colour-coded by which socket owned it. Blue for socket one, red for socket two.



Corn

You drew the NUMA map yourself instead of reading it out of the BIOS.



Hilbert

The BIOS told me the topology. The map told me which stick was where without opening a settings page. Different documents. And nobody else opened that case, because someone would breathe on the map. I sold the machine in the end. Kept the map. It's in the drawer with the thermal paste.



Herman

Of course it is.



Hilbert

What I want to know is whether the current ones would let you fence off half the cores and all of one socket's memory and run something on it that the other half never touches. Not virtualized. Actually fenced. I have a reason.



Herman

Cgroup memory binding plus NUMA pinning gets you most of the way — the kernel will keep the workload's pages local and the scheduler off the other socket's cores.



Hilbert

It isn't, actually. The second socket was still worth it, mind. Two of everything in one box. Two power feeds, two banks of memory, two ways in. Like a car with twin fuel tanks. It doesn't go faster. It goes further.



Corn

Right — so where does that leave the second socket?



Herman

Where we said: a capacity purchase wearing a speed purchase's clothes. If you take one thing from this episode, it's the reframe. Two sockets is how you get twenty-four DIMM slots and a terabyte of RAM on one board, and a hundred and twenty-eight lanes. It is not how you make anything go faster — the controlled tests say one point eight times on big-chunk work, one point zero two on small-chunk work, and the small-chunk world is where most of us live.



Corn

And the people who get burned are the ones who bought it for the reason that isn't true. The listing says forty cores, the wallet says a hundred and twenty-six dollars, and nobody reads the part about matched DIMMs and the idle draw.



Herman

There's a real open question in there, too — what happens to the used dual-socket market as DDR5 prices move and single-socket boards keep adding lanes. If one socket can feed all the DIMMs you need, the second socket's last job disappears.



Corn

Which would make it purely a collector's item. There's a drawer somewhere that's ready.



Herman

That's Hilbert Flumingtop at the producing desk, as ever. If this was your kind of episode, go back for episode two, Local STT For AMD GPU Owners; episode seventeen, Cloud Render Superpowers; and episode twenty, Architectural AI. This has been My Weird Prompts — no question too obscure, no rabbit hole too deep. If the show's earned it, a review helps other people find us, and you'll find everything else at my weird prompts dot com.



Corn

See you tomorrow, Herman.



Herman

See you tomorrow, Corn. Mind the NUMA.

SOURCES

What the research for this episode read before the script was written. Primary sources first.

1. **AMD, 2P or not 2P? That is the question, published 2025-02-21, retrieved 2026-11-09. [primary]**
https://www.amd.com/en/blogs/2025/2p-or-not-2p--that-is-the-question--.html

2. **AMD, The Right CPU can Boost Microsoft SQL Server Performance and Lower Software Costs, 2025-09-03, retrieved 2026-11-09. [primary]**
https://www.amd.com/en/blogs/2025/the-right-cpu-can-boost-microsoft-sql--server--performance-and-l.html

3. **AMD, EPYC Processor Overview, retrieved 2026-11-09. [primary]**
https://www.amd.com/en/partner/articles/epyc-processor-overview.html

4. **AMD, 4th Gen AMD EPYC Processor Architecture white paper, retrieved 2026-11-09. [primary]**
https://www.amd.com/content/dam/amd/en/documents/products/epyc/4th-gen-epyc-processor-architecture-white-paper.pdf

5. **Supermicro, CPU Information: AMD EPYC 9005/9004/7003 Series Processors, retrieved 2026-11-09. [primary]**
https://www.supermicro.com/en/support/resources/cpu-amd-epyc-9005-9004-7003

6. **llama.cpp GitHub discussion #11733, Dual Epyc Genoa/Turin token generation performance bottlenecks, by fairydreaming, retrieved 2026-11-09. [primary]**
https://github.com/ggml-org/llama.cpp/discussions/11733

7. **Wikipedia, Symmetric multiprocessing, last edited 2026-06-29, retrieved 2026-11-09.**
https://en.wikipedia.org/wiki/Symmetric_multiprocessing

8. **Wikipedia, Non-uniform memory access, last edited 2026-07-01, retrieved 2026-11-09.**
https://en.wikipedia.org/wiki/Non-uniform_memory_access

9. **DartPanel, AMD EPYC 9005 Turin: specs and review of all 27 models, prices as of October 2026, retrieved 2026-11-09.**
https://dartpanel.com/blog/amd-epyc-9005-turin-specs-review

10. **OpenClawDC, Dual-Socket vs Single-Socket EPYC for LLM Inference, 2026 (supersedes their 2026-08-24 guide), retrieved 2026-11-09.**
https://openclawdc.com/blog/dual-socket-vs-single-socket-llm-inference/

11. **Servnet UK, Single-socket vs dual-socket servers: which do you actually need? (UK 2026), retrieved 2026-11-09.**
https://www.servnetuk.com/insights/single-vs-dual-socket-server-uk

12. **Dell Technologies Info Hub, AMD EPYC for AI Data Preparation, retrieved 2026-11-09.**
https://infohub.delltechnologies.com/en-us/p/amd-epyc-for-ai-data-preparation/

13. **ServeTheHome, PCIe Lanes and Bandwidth Increase Over a Decade for Intel Xeon and AMD EPYC, retrieved 2026-11-09.**
https://www.servethehome.com/pcie-lanes-and-bandwidth-increase-over-a-decade-for-intel-xeon-and-amd-epyc/

14. **ORFEO HPC documentation, computational resources (Rome/Genoa interconnect specs), retrieved 2026-11-09.**
https://orfeo-doc.areasciencepark.it/HPC/computational-resources/

15. **lyan, Transport speed of common computer buses (QPI/UPI/Infinity Fabric figures), retrieved 2026-11-09.**
https://xryan.net/p/295

16. **Newegg, dual-socket AMD motherboard listings, retrieved 2026-11-09.**
https://www.newegg.com/p/pl?d=amd+dual+socket+motherboard

17. **eBay, ASRock Rack GENOAD8X-2T/BCM SP5 used at $895, listing updated 2026-08-15, retrieved 2026-11-09.**
https://www.ebay.com/itm/278284407746

18. **eBay, dual-socket EPYC motherboard listings (316 results), retrieved 2026-11-09.**
https://www.ebay.com/shop/amd-epyc-dual-socket-motherboard

19. **eBay, Dell PowerEdge R740 2× Gold 6138 at $957.99, listing updated 2025-04-15, retrieved 2026-11-09.**
https://www.ebay.com/itm/176389898023

20. **eBay, Xeon Gold 6138 search results (CPU prices, R740XD listings), retrieved 2026-11-09.**
https://www.ebay.com/sch/i.html?_nkw=Xeon+Gold+6138

21. **PicClick UK, R740xd 2× Gold 6138 at £599.63, retrieved 2026-11-09.**
https://picclick.co.uk/Dell-PowerEdge-R740xd-2xGold-6138-200Ghz-40-Core-64GB-156878719407.html

22. **Trident321 Solutions, ASRock 2U2G-GENOA refurbished 128-core barebone at $4,530, retrieved 2026-11-09.**
https://www.trident321solutions.com/p/2u2g-genoa-r/

23. **NewServerLife, The 2026 Buyer's Guide to Refurbished AMD EPYC Servers, 2025-12-26, retrieved 2026-11-09.**
https://newserverlife.com/articles/refurbished-amd-epyc-servers-2026-buyers-guide/

24. **Ubicloud, Does MHz still matter?, 2025-08-22, retrieved 2026-11-09.**
https://www.ubicloud.com/blog/does-mhz-still-matter

25. **Martin Alderson, How I make CI/CD (much) faster and cheaper, retrieved 2026-11-09.**
https://martinalderson.com/posts/how-i-make-cicd-much-faster-and-cheaper/