Episode Book

EPISODE #6188

Every Way Linux Relieves Memory Pressure

PUBLISHED	21 October 2026 at 17:21 UTC
DURATION	22:30
LISTEN	myweirdprompts.com/episode/linux-memory-pressure-zram-zswap-68660/
HOSTS	Corn, Herman, Larry, Hilbert
SCRIPT MODEL	anthropic/claude-sonnet-5.5
VOICES	chatterbox-regular (chatterbox-local)
PIPELINE	V5.12
BOOK GENERATED	22 October 2026 at 07:46 UTC

SYNOPSIS

Daniel's 16GB laptop runs six Claude Code sessions and stays alive, while his 64GB desktop barely notices. What's the difference? This episode walks the entire ladder of memory pressure relief in Linux — from dropping file cache for free, through zram and zswap compression, swap files, the swappiness dial, page-cluster, and multi-generational page tracking, down to the out of memory killers at the bottom. Along the way: why you shouldn't run zram and zswap together, why swappiness above 100 isn't a typo, why page-cluster should be zero on zram, and why used memory isn't pressure — stall time is. If you've ever watched a machine grind instead of die, this is the episode that explains what the kernel was actually doing.

TRANSCRIPT



Corn

Two gigabytes available out of thirteen, and the thing still scrolls. Six Claude Code sessions open, Chrome doing whatever Chrome does, and it just keeps going. That's what Daniel's asking about. His laptop is a sixteen gigabyte machine under real pressure that somehow stays alive, and his desktop has sixty-four gigabytes and rarely notices anything. He wants every mechanism Linux has for relieving that pressure laid out in one go, and then a verdict on what he should be running.



Herman

There's an order to it, Corn, and it's the order in which the kernel runs out of things it's willing to sacrifice. Zram, zswap, swap files, the swappiness dial, the killers. They look like a pile of unrelated tricks, but they're a ladder.



Corn

A pecking order of sacrifices. Fine. Start at the top.



Herman

Rung one is the file cache. Every file you've read recently sits in RAM, and the kernel can throw any of it away for free, because the copy on disk is the real one. No writing, no compressing, it just forgets. That's why a Linux box always looks like it's using all its memory. Most of it is cache that's perfectly happy to leave.



Corn

Free is the right price.



Herman

It stops being free when the cache you drop is the stuff you're about to read again. And then there's memory that has no copy on disk at all, the anonymous pages. Your browser's heap, a Python process, the guts of a Claude session. The kernel can't forget those. It has to put them somewhere, and every other rung is a different answer to where. At the very bottom of the ladder is the out of memory killer, and people picture a little judge deciding who deserves to live. It is not a person. It is a process, it runs a scoring routine, and it does not negotiate.

L

Larry

This is a public service announcement. If your computer has ever felt the squeezing hand of an out of memory killer, you are not alone, and you are not safe. Every day, decent machines suffocate because their owners never learned about swap. So listen carefully. The Swap Paw is a plush toy shaped like a hard drive that you kneel on and whisper to, and it will personally absorb all the pressure in your house. It holds sixteen gigabytes of anxiety, sixty four gigabytes of regret, and one hundred percent of your neighbors' disapproval. The Swap Paw. Because somebody has to carry this, and it cannot legally be your laptop.



Corn

Which is why you want every rung above it to be cheap, so the killer never gets its chance. What's rung two?



Herman

Compression. Take the anonymous pages the kernel wants out of the way and squeeze them while they're still in RAM. There are two ways to do it, and they get confused constantly. Zram is a device. The kernel creates a block device that lives in memory, you format it as swap, and anything swapped to it is compressed on the way in. Zswap is a cache. It sits in front of a real swap device on disk, compresses pages into a pool in RAM first, and only when that pool fills does it push the oldest ones out to the disk behind it.



Corn

Device versus cache in front of a device. So one doesn't need a disk and the other is useless without one.



Herman

That's why you don't run both. If your swap is a zram device, zswap would compress pages into its pool and then you'd be compressing them a second time into zram. The usual advice, and I've mostly seen it from people running these in practice rather than from the kernel documentation, is pick one. Daniel's laptop has zswap off and zram on, which is the coherent choice.



Corn

When would I pick zswap, then? If it's just a cache, it sounds like the lesser option.



Herman

When you want a disk swap anyway and you'd like the hot pages kept compressed in front of it. The pool isn't preallocated, so it grows as needed, and when it fills, the least recently used pages get written back to the disk. There's a catch with the default setup, though. Once the pool is full, zswap refuses new pages until usage falls back to a threshold, so a bad moment can send pages straight to the slow device. Zram on the other hand is simple. It's a device, it's there, and it's fast. For a machine that wants the compression to be the main event, I'd lean toward zram, and Daniel's choice reflects that.



Corn

Zswap is the lawyer. Zram is the plumber.



Herman

I would not have said that, but I'm not going to stop you.



Corn

What does zram hold on his laptop?



Herman

About eight point six gigabytes of data, squeezed into four point six of actual RAM. That's close to one point nine to one. The kernel documentation assumes about two to one, so he's right on the expected curve, and it also says a device much bigger than twice your RAM has little point. His is ten and a half gigabytes on a thirteen gigabyte machine, so it's nowhere near a ceiling.



Corn

So he's paying four and a half gigs of rent to avoid owing eight and a half. That's a better rate than most banks.



Herman

The rent is the thing to watch, because the compressed pages still occupy real memory. The algorithm is the lever on that rent. Zram supports a handful, the lzo variants, L Z four and Zstandard. The usual claim is that L Z four is faster and Zstandard squeezes harder, and I haven't seen a clean benchmark that settles it on the same data. His laptop uses Zstandard, his desktop uses one of the lzo variants.



Corn

I'd have guessed the reverse. Squeeze hardest where RAM is the scarce thing, and the laptop is the one short of RAM. Oh. That's what he did.



Herman

It is, and it's the right bias. A few microseconds of decompression disappears next to what a disk costs. Speaking of which, Corn, what does your gut say a swap fault costs on a fast NVMe drive?



Corn

Fifty microseconds? A hundred?



Herman

On his laptop, about twenty-four, for a four kilobyte read where the process sits waiting for one page. Zram is quicker still, because it's a memory copy and a decompress. So the disk is a second tier. It's for pages that are cold, the stuff nobody has touched in an hour. And the laptop is set up like that. An eight gigabyte swap file on the NVMe, sitting behind the zram, and barely touched. A hundred and forty-five megabytes when we looked.



Corn

Does zram ever grow its own disk behind it? I half remember something about writeback.



Herman

It can, if the kernel was built with it. You point zram at a backing partition, not a file, and it can push idle or incompressible pages out to it. It also has a recompression feature that squeezes idle pages a second time with a stronger algorithm. Neither is on his laptop, as far as I know. He gets the same effect more simply, with a normal swap file at lower priority.



Corn

Simple is the right word. Fewer things to explain when it breaks. Why does the kernel fill zram first, though? Nobody told it which was faster.



Herman

Priority. Swap devices carry a number, and the kernel uses the highest first. Zram sits at one hundred, the file sits at minus one. Zram fills, then the file takes the overflow. Give them equal priority and the kernel would round-robin between them, which sends pages to the slow place for no reason.



Corn

Minus one. Somebody was being polite about the disk's prospects.



Herman

Fedora's choice is worth a look, as a contrast. When they made zram the default, they sized it at half of RAM, capped at four gigabytes, and created no disk swap at all. That was deliberately conservative, a default that's safe for the vast majority of people. Daniel's laptop is more aggressive than that, and it's the right amount of aggressive for a machine that's actually under pressure.



Corn

A default tuned for a laptop that's mostly browsing. His is running six AI sessions and a Python job. Different animal.



Herman

Now swappiness, because this is the dial everyone misunderstands. People read it as a percentage, how full RAM has to be before swapping begins. It isn't. The kernel documentation calls it the rough relative cost of swapping versus filesystem paging, on a scale of zero to two hundred. At one hundred they're treated as equal. Below that, the kernel would rather drop file cache. Above it, the kernel would rather swap anonymous pages out and keep the cache.



Corn

And the laptop sits at one eighty.



Herman

Which tells the kernel that swapping is cheaper than rereading files, and when swap is a memory copy into zram, that's simply true. The documentation says values above a hundred can be considered for in-memory swap like zram. The range only went up to two hundred around kernel five point eight, so a lot of old tuning guides cap out at a hundred.



Corn

So an old guide would look at his laptop and call it a typo.



Herman

It isn't a typo. Next, page-cluster. The default is three, which means when the kernel reads one page back from swap it reads eight. That made sense on spinning disks, where the seek was the expensive part and the neighbours were nearly free. Zram has no seek. Reading eight pages to use one is just wasted decompression. Zero is right, and his laptop has zero.



Corn

The kernel kept a habit from the disk era and never dropped it.



Herman

There's one more piece in the machinery, which I'll keep short. The laptop has the multi-generational page tracker on, which arrived in kernel six point one. Instead of two lists, recent and not recent, the kernel sorts memory into numbered generations by age, so it's better at telling a page that's really cold from one that's just quiet. There's an optional setting, off by default, that protects the last second or so of working set from eviction, and if it can't be protected the kernel fires the out of memory killer instead of grinding.



Corn

Grinding being the thing where the laptop isn't dead, it's just refusing to answer. Which is worse. Dead I can reboot.



Herman

Right, and that raises the question of how anyone knows the difference between tight and grinding. Your fans can tell you before the kernel does.



Corn

My fans have never once filed a report.



Herman

The kernel does have one. Pressure stall information. It tracks the share of time that some tasks, or all non-idle tasks, are stalled waiting on memory, with averages over ten, sixty and three hundred seconds. A machine can be at ninety-five percent used with zero stall and be perfectly healthy. Another can sit at eighty with stalls climbing and be in trouble. Used memory isn't pressure. Stall time is pressure.



Corn

So that's the trigger to trust. Not when RAM is full, but when everything is stuck waiting for it.



Herman

And the killers are the last rung. The kernel has its own, the process I mentioned, and it tends to act late, sometimes after the machine has been miserable for a good while. So people run a userspace one that acts sooner. Earlyoom watches available memory and swap and kills the biggest offender when both run low. Systemd oomd reads the pressure numbers, works per control group, and kills a whole group.



Corn

We've done the calibration on those before. Today I only care how they sit next to swap.



Herman

They sit at the end of it. Systemd oomd's own manual strongly recommends swap be enabled, because swap is what buys it time to notice pressure. A machine with no swap goes from fine to frozen in seconds. A machine with a compressed layer and a disk layer slides down slowly, and a slow slide is something a killer can react to.



Corn

Here's the part I noticed on Daniel's laptop. Both are running. Earlyoom and systemd oomd.



Herman

I was going to raise that. Two killers with different triggers, one watching free memory and one watching stalls. I haven't found any source showing them actively fighting, and I wouldn't claim they conflict. Nobody I know of has documented a kill order between them. But it is redundancy, and on a bad day you don't know which one fires first. Fedora's packaging discussion said both can be installed and only one should run, and Fedora's own plan replaced earlyoom with systemd oomd. So the advice is pick one.



Corn

Which one?



Herman

You tell me. You're the one who likes the old simple tool.



Corn

Earlyoom. Fewer opinions. It looks at two numbers, kills the biggest thing, done. On a laptop with six Claude sessions, the biggest thing will be whichever one is working hardest, and I'd want that to hurt.



Herman

Systemd oomd can exempt chosen units, so you can shield the important ones. Ubuntu switched it on by default for the desktop in twenty-two oh four, and people reported it killing browsers and editors, which is the downside of a pressure trigger being less predictable than a threshold. Reasonable people go either way. What matters is one.



Corn

And the desktop?



Herman

Earlyoom only, so that's sorted. The desktop's oddity is different. Sixty-one gigabytes of RAM, a thirty point eight gigabyte zram device, and swappiness set to one.



Corn

One. Not a hundred and eighty. One.



Herman

At one, the kernel overwhelmingly prefers to drop file cache. The documentation I read doesn't spell out exactly what it does under real pressure, so that part is my inference. In practice the zram becomes a last-resort net. Unused zram costs about a tenth of a percent of its size in RAM, so around thirty megabytes here. It's cheap to leave sitting there.



Corn

Is it a problem, though? It's idle at ten gigs used, fifty-one free.



Herman

Not a demonstrated one. Pressure is zero. If a spike came, the desktop would reclaim cache before touching zram, so it would feel less relief than a high-swappiness setup. That's the whole cost.



Corn

So the net is rigged to catch you after you've hit the ground.



Herman

After you've hit the cache, which is softer ground. On a machine that idle, defensible.



Corn

Hibernation, then. I assume none of this helps there.



Herman

It's the last piece, and it's a trap. Zram lives in RAM, and RAM doesn't survive power off. To hibernate, the kernel writes memory to real disk swap, and the image has to fit on a single device. One report I saw had zram-only swap failing with a not enough swap space error. The laptop has an eight gigabyte file and thirteen of RAM, so hibernation is unlikely to work as it stands, though I haven't tried it on that machine. A partition has the edge for resume, since the setting names one device, and a file means an extra step that most



Herman (continued)

people would rather skip than...



Hilbert

...be talked into by an aunt. Ninety-one euros, up front, for a swap partition in my Aunt Rosalind's garage. I want it back. It is plywood, two metres by one, put up this month. She wrote me a receipt. Number four hundred and seventeen, and the first one she's ever written that wasn't on the back of a bus ticket. I have the receipt. The plywood is not the part I kept. The plywood belongs to a man called Ignatius Boyle in Tallaght, who has not been told it was sold.



Corn

A partition is a wall that splits a space in two. The garage qualifies.



Hilbert

It doesn't reach the ceiling, Corn. That matters. The arrangement is, whatever I can't fit in my flat goes in the left half. Whatever she can't fit in her house goes in the right. When the left fills up, we swap halves. Thursdays. By ladder. At the moment the left half holds four hundred and twelve umbrella handles with no umbrellas attached, and a tuba that belongs to the council. The right half holds her lodger, Mr Prendergast, and a goat called Commissioner who has not been told he's a goat.



Herman

A fixed Thursday rotation is at least a policy. The kernel has no ladder and no calendar, so it has to guess what's cold from how recently each page was touched.



Hilbert

She says the goat does the guessing. The receipt is witnessed by a heron.



Corn

A live heron?



Hilbert

Stuffed. It came with the garage, along with a fridge nobody is allowed to open before four. Herman, you keep saying page, and you say it as if a page were a small thing. I've had four hundred and twelve umbrella handles moved up a ladder one at a time. A page is not a small thing.



Herman

In bytes it's four thousand and ninety-six, which is lighter than an umbrella handle. Though I take the point that at that scale, the volume is the problem.



Hilbert

She invoiced me for the right half as well. Thirty-eight euros for space nobody has used. It has been empty for eleven days. Mr Prendergast sleeps in the left, on the tuba, by arrangement with the council. Nobody has explained that either. You said the kernel drops the cache first because it can read it back from disk, Corn. I went and looked at the left half after that. I could read back most of it.



Corn

That's a lot of umbrella handles to have a spare copy of.



Hilbert

I stacked pallets at a paper mill once, I know what four hundred and twelve of anything weighs. Of the ninety-one euros, nineteen is a keyhole fee, for a lock on a door that is also not there. I asked for the nineteen back and she offered fourteen laundrette tokens. I took the tokens. They are in a drawer with the receipt, and the nearest laundrette closed in the spring. She won't take the tuba out. It's on the paperwork as a deposit, listed ahead of me. The council has written twice and she has answered both letters with a drawing of a ladder. The handbook for the plywood says to oil the hinges. There are no hinges. She oils them anyway, Thursdays, before the swap, with a feather from the heron. When the right half is finally used, she says, it will be used by Commissioner, and I am to bring a coat.



Herman

Eleven days of an empty reserve and she still bills for it. An idle reserve does cost something small, just not nothing. Thirty-eight euros is a lot to pay for a reserve that has never been touched, though. The kernel charges about a tenth of a percent.



Hilbert

She does the work properly. I'll say that for her.



Corn

So what do we actually tell Daniel to run? Laptop first. The setup is good. Zram on Zstandard, swappiness one eighty, page-cluster zero, the swap file behind it at lower priority. I'd change very little.



Herman

One change: pick a single killer. If it were my machine I'd keep earlyoom and switch systemd oomd off, because the second one adds ambiguity rather than protection. And if hibernation matters to him, he needs a real swap area at least as large as the RAM it has to hold, which an eight gigabyte file isn't. If he never hibernates, ignore that.



Corn

Desktop, my instinct is leave it alone. It's idle, the net is cheap, nothing's choking.



Herman

I'd go a notch further and raise the swappiness to something moderate, say sixty. That's a mild preference of mine, not a repair. It lets the zram act as a first buffer in a sudden spike for almost no cost, instead of a net you only reach after the cache has gone. And the single-killer rule is already satisfied there.



Corn

Sixty. You've picked a number between my laziness and your enthusiasm.



Herman

That's what a compromise is.



Corn

The misconception I hear most is that any swap use means your machine has run out of memory.



Herman

It's the reverse on a machine like the laptop. Swap holding cold anonymous pages means the kernel parked what nobody was using so the cache could stay. Running out is swap full and the stall numbers climbing, which is a different state altogether.



Corn

Which leaves the thing worth keeping on both machines: a small reserve you rarely touch costs a fraction of its size, and that beats a killer arriving without warning. Eleven empty days and thirty-eight euros is a bad rate, but the principle holds. Let the compressed layer keep working, let one killer stand guard.



Herman

Thanks to our producer, Hilbert Flumingtop, for keeping the desk running.



Corn

This has been My Weird Prompts. If you want more, it's at my weird prompts dot com.



Herman

We'll be back soon.