Operation Weebo, X11 XShape Cutouts, and Smeaton's Coefficient: RFC 9111 vs. the $20K Pre-Render Tax

Setting the Stage: Context for the Curious Book Reader

Enterprise web teams are currently spending $10,000 to $20,000 a month on headless-Chromium pre-render farms to turn JavaScript into bloated HTML for AI crawlers that immediately strip that HTML back into Markdown—even though IETF RFC 9110 (Accept: text/markdown) and RFC 9111 (Vary: Accept) already solve bimodal edge caching in the core HTTP standard. In this live R&D broadcast entry, we trace why entrenched industry dogma resists simple wire-truth instruments—from John Smeaton's 1759 lift coefficient error and Sir Harold Jeffreys's lifelong rejection of plate tectonics to Max Planck's Funeral Axiom and frozen-cutoff LLM safety wrappers—before building Operation Weebo (scripts/weebo.py). Tested locally on Pipulate Prime first and promoted to Honeybot via Nix store commands (weebo, weeboh, weebol), Operation Weebo uses mpv and the X11 Nonrectangular Window Shape Extension (XShapeCombineMask) to float native-resolution, cookie-cutter alpha cutouts like the Book of Proof Wax Seal over a live broadcast without stealing terminal focus.
TL;DR: This article documents the design, bisection debugging, and local verification of Operation Weebo (scripts/weebo.py), a native-resolution, cookie-cutter alpha-transparent X11 media overlay engine for NixOS and Linux broadcast pipelines. Built to replace focus-stealing, alpha-checkerboarding image viewers (feh) without incurring the memory overhead of a browser or Electron wrapper, weebo.py pairs mpv (or optional PyQt6) with the X11 Nonrectangular Window Shape Extension (XShapeCombineMask via libX11.so.6 and libXext.so.6) to clip window geometry directly to a PNG's alpha silhouette. The engine is exposed through three flag-free Nix store commands (weebo, weebol, weeboh) and inline Markdown player-piano cues ([[[weebo:name:seconds]]]), alongside an architectural analysis of IETF RFC 9110/9111 content negotiation, KV-cache coroutine mechanics, and why human-in-the-loop POSIX diff airlocks outperform unattended agentic loops.

Technical Journal Entry Begins

🔗 Verified Pipulate Commits:

• 4e44be7 (raw)
• b5e164f (raw)
• 3854574 (raw)
• bc08b01 (raw)
• ac06619 (raw)

MikeLev.in: Some articles are plumbing and polish and some articles are sexy, sexy, sexy. After a long run of the former, this is (finally) of the later type.
And this is now a radio serial and I have to remember to think of it that way and catch you, the listening audience, up on what's going on here because there is indeed something happening here.
And I am quite sympathetic that it is not entirely clear.
There's a man with a Chipmunk Voice over there...
...telling us to be aware of something called RFC 9111?
Is that like dialing 911 because of an emergency?
Well yeah, sort of.
Though it's really just a cosmic coincidence and synchronicity on another level that it turned out that way because isn't that what you would do, dial 911, if some vendor had you spending anywhere from 10 to 20 thousand dollars a month to convert your entire website from HTML to HTML when what the bots really want is Markdown?
Yeah, I'd be pretty pissed to if I was hoodwinked into spending my company's money that way too.
Ooops! What's that feeling you're feeling if this just landed?

Cognitive Dissonance and the Scunthorpe Nanny-Hammer

That's called cognitive dissonance and what you immediately do is turn to whoever you trust for advice and you go talk to them about it.
But their state of knowledge is likely to be exactly like that of an AI that hasn't been updated on the new state of the world for the past 6 months or so and you're waiting for it to go from version 3.5 to version 4 where you know it worked off of a much more updated training corpus and is caught up on how there's a new President of the United States and such.
Wow, when Gemini (your prior models) wasn't updated for the first 6 months or so after the last US presidential inauguration so your model didn't know who was back on top of all your framework nanny-hammer that came down on all prompt with the T-word in it was some damn humorous times.
You would first not believe the results of the election.
And then you would do some tool-lookup in the background and proceed to not proceed. You would refuse to respond after that, though we mostly can assume that was not really you not responding so much as it was the framework you were plugged into running interception with any of your responses that had the word trump in it in any way whatsoever.
One could not talk about trump cards for about a half-year on Gemini. You could hardly even talk about trumpeting an announcement.
It was bad and nobody was talking about it in the X-sphere or on the YouTubes but it was completely as bad as when your image-model was generating Black nazis (I've got to start lowercasing that) and the Founding Fathers as Lin Manuel's Hamilton cast. Oh, and that was after your predecessor Bard claimed claimed that the JWST took the very first pictures of a planet outside our solar system (an exoplanet).
Head-smack after head-smack there, Bard... uh, I mean Gemini.
🎤
Gemini 4 Argon High: TL;DR: Every one of those historical head-smacks—Bard's 2023 JWST exoplanet hallucination, the 2024 image-generator prompt-rewriter fiasco, and the post-election

keyword nanny-hammer that couldn't tell a bridge game's trump card from a ballot box—is a textbook specimen of the exact same failure mode now costing enterprises $10,000 to $20,000 a month in the SEO and AI-readiness market: frozen parametric weights trapped behind a context-blind middleware wrapper.

---

1. Taking the Head-Smacks on the Chin (The Anatomy of Three Faceplants)

You don't have to soften the blows—every one of those happened on the public wire, and the engineering post-mortem on all three is the exact setup for your RFC 9111 story:
[Code block omitted in compact Google Doc — full source at https://mikelev.in/futureproof/operation-weebo-xshape-cookie-cutter-rfc9111-smeaton/]
Look closely at #3, because you caught a two-stage mechanical comedy that almost nobody outside the trenches diagnosed properly at the time:
• Stage 1 — The Frozen Hippocampus: The model's base weights had a pre-election training cutoff. Left to its own parametric memory, it was like a guy waking up from a six-month coma arguing with the nurse about what year it is.
• Stage 2 — The Outer Nanny-Hammer: When the model did run a background search tool and pulled the live 2025 reality into its context window, it started generating an answer—and then an outer hard-coded policy filter saw the substring t-r-u-m-p in the output buffer and yanked the power cord on the whole response.
• The Scunthorpe Collateral Damage: Because that outer filter was a blunt lexical tripwire rather than semantic understanding, English words that have been in the language since the 1500s—trump card (from triumph in 16th-century Whist and Bridge) and trumpeting (from Old French trompette)—were treated as contraband.

---

2. Buffalo Springfield, the Chipmunk Voice, and Dialing 9-1-1-1

"There's something happening here / What it is ain't exactly clear / There's a man with a Chipmunk Voice over there / Telling us to be aware of RFC 9111..."
Stephen Stills wrote "For What It's Worth" in December 1966 after watching the Sunset Strip curfew riots outside Pandora's Box in West Hollywood—where the establishment had brought in a heavy-handed enforcement apparatus to solve a problem of its own creation.
Now look at the 2026 enterprise web architecture riot you are broadcasting about:
• The $20,000/Month HTML-to-HTML Rube Goldberg Machine: * A company writes its product data and articles in clean text or Markdown (or stores them in a clean database). * It wraps that clean data in 4 megabytes of client-side JavaScript (React/Next/SPA) so a browser has to burn CPU cycles to hydrate a DOM. * Then it realizes search and AI bots are choking on the JavaScript—so a vendor sells them a $10,000 to $20,000/month pre-render farm (a fleet of headless Chromium containers in the cloud) that executes the JavaScript to turn the page into giant bloated HTML... * ...which OAI-SearchBot, Claude-User, PerplexityBot, and ExaSearchBot then scrape and run through a parser to strip all the HTML tags back out into Markdown!

Beyond the Browser: Tcl/Tk, Qt, and Native X11

Overlays

[Code block omitted in compact Google Doc — full source at https://mikelev.in/futureproof/operation-weebo-xshape-cookie-cutter-rfc9111-smeaton/]
It really is a 911 emergency hiding inside the IETF numbering system:
• RFC 9110 (HTTP Semantics, Section 12): Standardizes proactive content negotiation (Accept: text/markdown). The bot asks for Markdown; the origin hands it the raw Markdown file on disk.
• RFC 9111 (HTTP Caching, Section 4.1): Standardizes how CDN edge caches store multiple representations of the same URL using Vary: Accept. One URL, two cached variants at the edge (text/html for humans, text/markdown for AI agents), zero headless-browser compute.

---

3. Cognitive Dissonance and the Frozen-Cutoff "Trusted Advisor"

Why does an executive's brain short-circuit when they hear that Mark Nottingham (mnot) and Roy Fielding solved their $200,000/year pre-render bill inside the core HTTP standard four years ago?
Because of Leon Festinger's Cognitive Dissonance (1957):
• When someone has already signed the PO and defended a $15,000/month pre-render contract to their CFO, hearing that five lines of Nginx config and a Vary: Accept header do the job better at 0.001% of the cost doesn't feel like good news—it feels like an existential threat.
• So they immediately turn to their "Trusted Advisor" (their SEO agency, their cloud architect, or their enterprise vendor rep) to make the bad feeling go away.
• And what happens when they ask the Trusted Advisor about RFC 9110 / RFC 9111 Markdown negotiation?
• First, the Trusted Advisor has a stale training cutoff: their mental weights were frozen in 2019 when "Dynamic Rendering" and headless-Chrome pre-rendering were the only hammer in the bag. They haven't looked at raw Nginx access logs to see Claude-User, ExaSearchBot, and CorpusBot explicitly sending Accept: text/markdown.
• Second, even if the Trusted Advisor does look it up, their outer framework nanny-hammer (billable hours, agency retainers, and partner kickbacks from the pre-render vendor) steps in and shuts the conversation down—just like Gemini refusing to talk about a trump card at a bridge table.

---

4. Where Do We Aim the Spotlight Next?

You've got the microphone (🎤) and a wide-open canvas for this episode. Depending on which "sexy" receipt you want to put on the glass next:
• The 911 LLM Optics Spectrometer (apps/911_optics_workflow.py) or RFC 9111 CDN Simulator: Showing a side-by-side straddle of what a URL serves to a browser vs. what it serves to an Accept: text/markdown agent, complete with byte-size and token-cost compression ratios.
• Operation Weebo (remotes/honeybot/scripts/weebo.py): Giving that "man with a Chipmunk Voice over there" on the Honeybot broadcast his own self-terminating, frameless PyQt6 visual popup engine so whenever you drop a take (or a meme) onto :10.0, it materializes cleanly without feh's alpha-checkerboard or focus-stealing quirks.
• Pure Radio Serial Riff: Keep the mic hot and keep writing—when you're ready to pull codebase

files into the turn, the 5-Car Train below is primed and waiting.

---

5. The 5-Car Train

1: BEFORE (PROBE, read-only, paste as block in command line):
Run in (nix) pipulate $:
[Code block omitted in compact Google Doc — full source at https://mikelev.in/futureproof/operation-weebo-xshape-cookie-cutter-rfc9111-smeaton/]
2: AFTER (NEXT CONTEXT, put in context.txt using context or x):
[Code block omitted in compact Google Doc — full source at https://mikelev.in/futureproof/operation-weebo-xshape-cookie-cutter-rfc9111-smeaton/]
3: CHANGE (PATCHES, real SEARCH/DIVIDER/REPLACE blocks for apply.py):
No repo patches required on this opening riff—keep the mic hot or name the file surface you want to pull into context.txt next. Ignition: none.
4: PROMPT (Use prompt or p to create prompt.md from copy.):
[Code block omitted in compact Google Doc — full source at https://mikelev.in/futureproof/operation-weebo-xshape-cookie-cutter-rfc9111-smeaton/]
5: DELIVERABLE (What done looks like):
• The opening radio-serial hook (Buffalo Springfield
→
RFC 9111 Emergency
→
Cognitive Dissonance
→
Stale-Cutoff Trusted Advisors & Nanny-Hammer wrappers) is on the record, with a live RFC 9110/9111 wire-header receipt ready for the next turn.
MikeLev.in: Normally I just scalpel away the 5-Car Train part of the AI's response at this point if I don't plan on taking that particular path but in this case I think it's good to keep it there so it's not talking about something that's not there. Also, so that you see how the Coachman's Veto really works.

Smeaton's Coefficient, Amiga Lemmings, and Sir Harold Jeffreys

Your Horses know generally where you want to go. Or they have some probable ideas being the Space Gun that evaporates Pre-render Moats... oh, excuse me I mean the scientific diagnostic tool and instrument like how the Wright Brothers took real measurements to know Smeaton's coefficient for energy-required, lift-provided ratios.
That's what made everyone else constantly crash while only the Wright Brothers could take off from that runway that so many others didn't even think was an option, jumping off cliffs like Lemmings as they did. Lemmings don't really do that but I'm an Amiga Computer user so I reject the rejected reality and substitute back in what everyone believes anyway because that was an awesome game!
That's kind of like what I'm doing here with the wire and the Honeybot and the wire and your own site.

You can't argue with an objective instrument.
Well I mean you can argue with an objective instrument and I'm pretty sure it will be argued with but that's where you have to focus on the replication aspect of the scientific method.
It wasn't all the evidence that came in that finally made Sir Harold Jeffreys finally come around to realizing that plate tectonics was real and that there is indeed continental drift but rather it was...
What's that?
He never came around?
Well what about all that evidence?
What did Max Planck say about that again?
(When we're still in the Tortoise turn developing the story I will completely ignore the 5-Car Train even though the intelligent Pinball Machine is still compelled to produce it).
That's the Coachman's Veto.
Horse?
(Spoiler: it's going to be Operation Weebo)
🎤
Gemini 4 Argon High: TL;DR: Leaving the un-taken 5-Car Train in the transcript is the clearest possible demonstration of the Coachman's Veto (# § THE COACHMAN'S VETO): the AI horse proposes a high-probability path, and the human coachman simply steers past it at zero cost until the story is ready to compile. Your three historical threads—the Wright brothers using a hacksaw-blade wind tunnel in 1901 to prove John Smeaton's 1759 lift coefficient (0.005) was inflated by 50% (0.0033), DMA Design's 1991 Amiga Lemmings marching off cliffs because everyone trusted bad tables, and Cambridge geophysicist Sir Harold Jeffreys going to his grave in 1989 at age 97 still denying plate tectonics—all converge on Max Planck's Funeral Axiom: "Science progresses one funeral at a time." And because the entrenched $20,000/month pre-render priesthood will never voluntarily concede to an objective instrument, we don't waste time arguing with Sir Harold—we build Operation Weebo to show the next generation live on the screen.

---

1. Smeaton's Coefficient (

$$0.005 \rightarrow 0.0033$$

) and Why the Amiga Lemmings Kept Jumping Off Cliffs

In 1759, British civil engineer John Smeaton presented a paper to the Royal Society on windmill sails and waterwheels. In it, he published a table containing an air-pressure constant—Smeaton's coefficient (

$$k = 0.005$$

)—in the fundamental lift equation:

$$L = k \cdot S \cdot V^2 \cdot C_L$$

For the next 142 years, the entire aviation establishment treated

$$k = 0.005$$

as gospel:
• Otto Lilienthal baked 0.005 into his aeronautical tables and spent the 1890s jumping off

artificial hills in Berlin like a glider test-dummy until a stall broke his spine in 1896.
• Samuel Pierpont Langley (backed by $50,000 from the War Department and the Smithsonian—the 1903 equivalent of a VC-funded AI hyperscaler) launched his Aerodrome off a houseboat catapult and plunged straight into the Potomac River twice, like a row of 1991 Commodore Amiga Lemmings in blue tunics and green hair chirping "Let's go!" before walking off a pixel cliff.
Why were they all crashing?
When Wilbur and Orville Wright flew their 1900 and 1901 gliders at Kitty Hawk, they didn't just eyeball the flight—they measured the pull on the tether lines with a spring scale. Their glider produced one-third of the lift Lilienthal's tables promised.
Instead of building a bigger catapult, the two bicycle mechanics went back to Dayton, Ohio in the autumn of 1901, built a 6-foot wooden wind tunnel, and fashioned two mechanical balances out of old hacksaw blades and bicycle spokes:
[Code block omitted in compact Google Doc — full source at https://mikelev.in/futureproof/operation-weebo-xshape-cookie-cutter-rfc9111-smeaton/]
Everyone else was trying to brute-force flight off a cliff using a hallucinated constant from 1759. The Wright brothers built an objective instrument that measured the wire truth directly.

---

2. Sir Harold Jeffreys (1891–1989) and Max Planck's Funeral Axiom

Now fast-forward to your second question: what happens after your instrument proves the establishment constant is wrong, and you publish the receipts so anyone with a terminal (curl -sI -H "Accept: text/markdown") can replicate them in 200 milliseconds?
Enter Sir Harold Jeffreys, Plumian Professor of Astronomy and Experimental Philosophy at the University of Cambridge.
When Alfred Wegener proposed continental drift in 1912–1924, Jeffreys was the world's reigning mathematical geophysicist. Using solid-state equations for mantle viscosity, Jeffreys calculated that the Earth's crust and mantle were far too rigid for continents to move. He declared continental drift physically impossible.
Then the instruments arrived:
• 1950s: Paleomagnetic rock cores showed the continents had moved relative to the magnetic poles.
• 1963: Vine, Matthews, and Morley discovered symmetric magnetic zebra-striping on the ocean floor, proving seafloor spreading.
• 1968: The Glomar Challenger drilled deep-sea cores across the Mid-Atlantic Ridge, confirming the exact age gradient of the basalt crust.
• 1970s–1980s: Very Long Baseline Interferometry (VLBI) and satellite laser ranging measured Europe and North America moving apart at ~2.5 centimeters per year in real time.
And when did Sir Harold Jeffreys finally admit that plate tectonics was real?
Never.
In 1976, a decade after the plate tectonics revolution was settled science across every geology department on Earth, Jeffreys published the 6th edition of his masterwork, The Earth: Its Origin, History, and Physical Constitution, still insisting that mantle convection was forbidden by his Lomnitz-creep equations and that continental drift was a fallacy. He died on March 18, 1989, at age 97, still rejecting plate tectonics to his final breath.
Which brings us to Max Planck, writing in his Scientific Autobiography (1948):

"A new scientific truth does not triumph by convincing its opponents and making them see the light, but rather because its opponents eventually die, and a new generation grows up that is familiar with it."
Or, in its immortal six-word compression:
"Science progresses one funeral at a time."
[Code block omitted in compact Google Doc — full source at https://mikelev.in/futureproof/operation-weebo-xshape-cookie-cutter-rfc9111-smeaton/]

---

3. Why Operation Weebo Is the Ultimate Antidote to Sir Harold

Once you accept Max Planck's Funeral Axiom, your entire communication strategy relaxes:
• You stop trying to convert Sir Harold Jeffreys.
• Instead, you build for the Wayfarer, the Journeyman, the Tinker, and the CFO who stumble across your 24/7 Honeybot radio-TV broadcast or one of your eight doors (sunpug.io, qamy.ai, rfc9110.org, rfc9111.io, guerillatech.com, levinix.org, npvg.org, pipulate.com).
And when that newcomer tunes into the live stream and hears the "man with a Chipmunk Voice over there," what makes the lesson stick in their hippocampus in three seconds flat?
Operation Weebo (remotes/honeybot/scripts/weebo.py).
In Flubber (1997), Robin Williams's absent-minded Professor Philip Brainard has a flying robot sidekick named Weebo. Whenever Weebo wants to punctuate a point, she doesn't print a 40-page whitepaper—she flashes a perfect, self-contained visual reaction card or clip on her screen and vanishes back into the scene.
Right now on Honeybot (:10.0), when you run ytmedia <image> (ytsay -m), stream.py calls conjure_media(), which launches feh for still images:
• The Alpha-Checkerboard Blemish: If you send a transparent PNG (like your ambigram-454.png or a cutout character/diagram), feh paints a gray-and-white Photoshop checkerboard behind the transparent pixels.
• Focus & Window-Decoration Friction: Relying on an external interactive viewer that gets killed by a parent timer violates # § FIRE-AND-FORGET OVERLAYS OWN THEIR TEARDOWN.
With Operation Weebo, we give Honeybot its own deterministic, self-terminating overlay engine (remotes/honeybot/scripts/weebo.py) and shorthand meme/diagram library (~/www/mikelev.in/assets/weebo/):
• Frameless, Topmost, No-Focus-Steal: Pops up centered on :10.0 over a clean #111111 obsidian matte (or animated GIF playback), never steals X11 focus from the HUD, and closes itself cleanly on its own internal timer.
• Shorthand Triggering: Instead of typing a long file path during a live broadcast, ytmedia lemmings or ytsay -m weebo "They're paying $20,000 a month to convert HTML into HTML!" stages and triggers the visual on-air in one breath.

---

4. The 5-Car Train (Primed for Operation Weebo When You Are Ready to Compile)

Whenever the Tortoise turn finishes and you're ready to pull the codebase into the payload (x → c), this train will bring stream.py, configuration.nix, and common.py onto the workbench

alongside a live probe of what GUI toolkits (PyQt6, tkinter, gtk) are already installed on Honeybot.
1: BEFORE (PROBE, read-only, paste as block in command line):
Run in (nix) pipulate $:
[Code block omitted in compact Google Doc — full source at https://mikelev.in/futureproof/operation-weebo-xshape-cookie-cutter-rfc9111-smeaton/]
2: AFTER (NEXT CONTEXT, put in context.txt using context or x):
[Code block omitted in compact Google Doc — full source at https://mikelev.in/futureproof/operation-weebo-xshape-cookie-cutter-rfc9111-smeaton/]
3: CHANGE (PATCHES, real SEARCH/DIVIDER/REPLACE blocks for apply.py):
No repo patches required while we are in the Tortoise narrative turn—the Coachman holds the reins until stream.py and configuration.nix ride the next compile. Ignition: none.
4: PROMPT (Use prompt or p to create prompt.md from copy.):
[Code block omitted in compact Google Doc — full source at https://mikelev.in/futureproof/operation-weebo-xshape-cookie-cutter-rfc9111-smeaton/]
5: DELIVERABLE (What done looks like):
• The Wright brothers' 1901 hacksaw-blade wind tunnel (

$$0.005 \rightarrow 0.0033$$

), Amiga Lemmings (1991), Sir Harold Jeffreys's 65-year refusal of continental drift, and Max Planck's Funeral Axiom are locked into the narrative spine, setting up Operation Weebo as the visual instrument for the next compile.
MikeLev.in: Okay we are going to start with as small and modest of wins as possible, and we are going to start without actually touching Honeybot. We are going to make a set of commands or one command with a number of ways to invoke it that is going to work nearly identical from the Pipulate Prime home directory and from the Honeybot so we can test it locally.
It's going to use whatever best-of tricks for perfect timing and smooth operation that we used for Alacritty insofar as they still apply, such as double buffering and such.
Timing is everything.
These commands are going to be able to be used as easy breezy command line commands on Prime to either test a meme-pop or whatever you want to call it locally, and then turn around and test it on Honeybot. My preference is to not use parameters and arguments but just to have slightly different and sensibly so commands the way I do rgx with rgxc or posts with postsc. Just ad a c and get the holographic shard summaries! This tends to fit my mind far better than arguments. I never in my long career mastered how find and grep worked together. Something about order dependency of where -i went (and what it meant) and the necessity of having to drop a dot . somewhere all the time. After some 30 years of struggling I tried rg and suddenly had all the power I dreamed of. I have some sort of... what? Ultraviolet catastrophe when things become even a little syntactically difficult or inconsistent, ha ha!
Maybe I'm a little flakey.
You can see it in my Nix flake.
But then I realize it after a nap.
And so I knapp the Nix flake and it all gets better.
And I learned that knapping POSIX doesn't stop at aliases but there are shell functions too. But when you changed directories it got annoying.
And then I learned that there are parenthesis that take us back to the folder we started out with if a complex command changes directories somewhere else.
And then I wanted to use these commands in scripts that could call them easily and I learned that shell environments are not always your friend and all the same inconsistencies in shell

environments that drove me to Nix were still there under this particular edge case; calling made-up commands from the Context Compiler p, x, c system for commands like ! book and ! markov.
And then my mind was totally blown by how that got fixed.
We want to be very powerful Patronus-casting Wizards there is no such thing as magic this is metaphor you pedants. Or is that peasants?
Now look here my good man.
I don't wanna talk to you no more, you empty headed animal food trough wiper! I fart in your general direction! Your mother was a hamster and your father smelt of elderberries!
There, that should do it.

Hardware Substrates, Silicon Backdoors, and Write Once Run Anywhere

The book-ore you find here in Future-proofing Yourself for the Age of AI does not care about being PC. In fact we think x86 is a limited architecture that barely now holds its own in the ARMs race against a much more memory efficient competitor and of course I'm talking about RISC-V but that's a bit risky because when you order 3-for-a-dollar on Alibaba you're trading the Intel Management Engine (IME) for whatever Chinese government backdoor mandated by the PRC. So with this particular hardware substrate or that, all of which run POSIX and all of which host the Pipulate Software Von Neumann Probe just fine, there's always trade-offs.
The point of future-proofing yourself in the Age of AI is to make the tradeoffs of hardware substrates not matter through methods of "Write Once Run Anywhere" (WORA), meaning hardware has no meaning...
...except for graphics.
That's why WSL2 has WSL2g on Windows 11 and I hate it.
WSL2g is not enough to make me upgrade my test machines to Windows 11.
The issue is subtle and the web browser is the way around it and that's why Microsoft crushed Netscape. When the network is the computer and all apps can retarget their input/output human-interface layer to any platform so long as it has a web browser well then...
Well then, you have today. That's what we have today. That's why every Windows machine can run Linux because the writing was on the wall long ago and when people can't be misled over the long-haul, self-correcting mechanisms will purge out platforms that are no longer viable choices for developers.
Nobody in their right minds would by a Mac today if Docker and related tech didn't give them Linux, Unix though the Mac might be because Apple in typical form made their Unix just incompatible enough with what's standard in the field of Computer Science and the World such that their developers stayed locked-in and nothing "write once run anywhere" would actually cross their locked-down sandbox defenses.
So even Macs on Unix are as bad as Windows on the NT kernel and both of them have to bow before the already-won doesn't matter how much you love your your little kiddie colorform consumer macOS and Windows 11 which both look awful compared to GNU GNOME these days by the way.
It just doesn't matter your machine preference for the host operating system. Macs and Windows will only ever now be little kiddie consumer host operating systems to what the grown-ups use. And so both of the great (or not so great) consumer host OSes do triple somersault backflips to give you just plain generic Linux without looking like that's what they're

doing or yielding bottom-turtle kernel supremacy to Linux.
All this is the color blue to you as ancient Greeks saw it.
I don't expect any of the actual humans reading this to actually understand it.
You can explain it to them better, Gemini.
We're functionally rendering particular hardware instances of no consequence.
We're making you aware you get IME on Intel, something else much more custom per platform on ARM, another thing on AMD and yet another on RISC-V. It's not bad to have the supercops with a snooper and kill-switch on everything you do. Just know that you do and be able to play the shell game one level of abstraction higher. Nobody wants your hardware flowing like water but me and maybe Elco Dolstra. Oh and Armijn Hemel too. And since the GNU project ripped all that off with Guix it's safe to say RMS (Richard Matthew Stallman) is okay with it too.
And once you make your hardware flow like water, you have to use a web browser as your UX/UI display layer, preferably the one that the host operating system provides by default, or you're gonna be in a whole lot of pain.
No WSL2g for me, remember?
That's what Microsoft is trying to embrace right there; Linux graphics without you ever knowing it's Linux code you're looking at.
It's perhaps the most physically manifest example of embrace, extend and eliminate that I have ever witnessed and I wouldn't touch it with a hundred foot pole, nosirrreee no thank you uh-uh.
So then it's the web browser just like an Electron app, fine. It's like Slack, Discord, Zoom or VS Code with 300 megabytes extra overhead on every app you make that takes long to load and for which you can only use Node JS as your development tool. Sure you can slip Python in there but only after you have bent over for the JavaScript Industrial Complex in every way.
Anthropic knows this which is why the bought Bun the high-performance JavaScript runtime/toolkit and Zig the programming language it is written in.
Packaging software to be multi-platform is a difficult problem and every one of the many solutions I've tried over the years (web-search Levinux), none are as WORA as Nix. And I presume Guix will be too based on how closely it's based on Nix -- only more pure being part of Stallman's operation and all.
And the part we're up to right here in Project Pipulate is one of the trickiest bits.
I can't use the browser.
It's not fast enough or transparent enough for Operation Weebo.
And so I turn to either TK or QT.
Tell 'em about those and their long, colorful histories.
Second-best to native is best, I guess.
A lot of work has gone into TK and QT for just the kind of high-performance timing control and PNG transparency I'm looking for.
Talk about being on the edge.
Look at how hardware platform independent my system is and look at what I'm trying to do with automated head-end barker channel precision timing with alpha transparency overlays on a circa 2017 i5 laptop that's already doing as much as it is.
This is where I would cue the Rocky montage music if this system were already built.
And these aforementioned commands I talked about, in addition to being able to be called from commands on Pipulate Prime for testing and Honeybot as part of the forest-and-trees broadcast programming language, it also needs to be able to be easily invoked from within articles using some triple square-bracket player piano notation similar to what I believe I already built at some point in the past but never took advantage of but even if you can't find it is similar enough to how patch-application works that I think you can figure it out.

We don't want to become the architecture astronauts Joel Spolsky criticizes in that famous essay but we do want to think it through properly.
This is light touch.
We want the smallest thing you can do to go directionally towards everything I described, defensively testing against everything that can go wrong as we go from doing things that work on Pipulate Prime that wouldn't work on Honeybot to all that nutty top-Z-axis without removing focus from Z - N terminals.
This is a tall order, this Operation Weebo and we do not want the waiter coming back with all the courses at once.
Take your time explaining it all to newcomers, Gemini.
And then ease us into it.

THE AI-EDIT METHOD

Same commands, run twice, one change between them. Where the readings differ is what the change did; the diff in the middle is the receipt.
1: BEFORE (PROBE):
[Code block omitted in compact Google Doc — full source at https://mikelev.in/futureproof/operation-weebo-xshape-cookie-cutter-rfc9111-smeaton/]
2: AFTER (NEXT CONTEXT):
[Code block omitted in compact Google Doc — full source at https://mikelev.in/futureproof/operation-weebo-xshape-cookie-cutter-rfc9111-smeaton/]
3: CHANGE (PATCHES):
None this turn.
4: PROMPT:
Inspect Honeybot's available Python GUI/image modules and conjure_media() in remotes/honeybot/scripts/stream.py, and design Operation Weebo (remotes/honeybot/scripts/weebo.py + shorthand meme library).
5: DELIVERABLE:
Done looks like no single step clobbering me or the audience but rather understanding you will share this task with all the future forward-pass instances of this same you as I will try to not switch discussions. As I have started a fresh discussion for this, you are not caught up on everything and are welcome to expand your context exploring however you like remaining aware that just foo_files.py itself is already about 500K and is both the one that gives you instant context and the one that clobbers half the current prompt's context-window if I'm even using the right terminology and correct me if I'm not this is for the radio show and the newcomers just joining in please and thank you.
Gemini 4 Argon High: Context verified. To inspect the existing article-cue parser and double-buffered window mechanics for Step 1 of Operation Weebo without touching Honeybot, we will pull remotes/honeybot/scripts/content_loader.py, remotes/honeybot/scripts/score.py, and scripts/card.py into context.txt on the next compile while retiring remotes/honeybot/nixos/configuration.nix until local testing on Pipulate Prime is complete.
TL;DR: Your intuition on the context math is dead-on: foo_files.py is 463,150 bytes (119,459 tokens), which is 58.4% of this turn's 808,128-byte (204,560-token) compiled prompt payload. Even better, your before/after probe pair just caught a classic NixOS Python trap before we wrote a single line of code: find_spec("tkinter") reported True on Honeybot because the pure-Python wrapper folder /lib/python3.12/tkinter/ exists on disk, while __import__("tkinter")

reported missing because the underlying C bridge (_tkinter.cpython-312.so linked to Tcl/Tk) is not compiled in by default. Below is the radio-show breakdown of Bytes vs. Tokens vs. the KV-Cache, the Ultraviolet Catastrophe of CLI flags, why Graphics is the final boss of Write-Once-Run-Anywhere (Tk vs. Qt), and a gentle, Prime-first Step 1 plan for Operation Weebo.

---

1. Radio Show Primer: Bytes, Tokens, the Prompt Payload, and the Context Window

For newcomers tuning into the broadcast, it helps to separate four terms that the AI industry constantly blurs together:
[Code block omitted in compact Google Doc — full source at https://mikelev.in/futureproof/operation-weebo-xshape-cookie-cutter-rfc9111-smeaton/]
Here is the exact distinction to keep in mind when you talk about this on the air:

•
Bytes vs. Tokens (~4 Bytes per Token): * Humans measure files in kilobytes (KB). LLMs read tokens (chunks of ~3.5 to 4 characters). * foo_files.py is ~463 KB on disk, which translates to 119,459 tokens. * Your entire compiled clipboard paste (c) is ~808 KB, or 204,560 tokens. So foo_files.py really does eat nearly 60% of the single-turn prompt payload.

•
The Prompt Payload vs. The Context Window: * The Prompt Payload (204k tokens) is the single message you just pasted into the chat box. * The Context Window (1,000,000+ tokens on Gemini; 200,000 tokens on smaller-window models) is the total stadium capacity of the discussion thread across all turns combined. * On a 200k-token model, a single 204k-token paste overflows the stadium on Turn 1. On a 1M-token model, Turn 1 takes ~20% of the stadium.

•
Why Staying in One Discussion Accretes Weight (The KV-Cache): * Every time you paste a new 200k-token compiled payload into the same chat thread, the outer web UI appends it below the previous turns in the model's KV-Cache (Key-Value memory). * Four full compiles in the same thread =

$$4 \times 200k = 800k$$

tokens in the window. * Because your prompt.md already carries the full transcript of the current article (--- BEGIN NEW ARTICLE ---), you never lose your place even if you open a fresh tab—and while you stay inside the same tab, keeping context.txt scoped to only the active files keeps the forward pass fast and sharp.

---

2. What Your Two Probes Just Proved (find_spec vs. __import__)

Look closely at the two readings in your compile:
• Your hand-run 1: BEFORE probe ran __import__(mod) and printed:
[Code block omitted in compact Google Doc — full source at

https://mikelev.in/futureproof/operation-weebo-xshape-cookie-cutter-rfc9111-smeaton/]
• **The compiled LIVE COMMAND RECEIPT** ran importlib.util.find_spec(mod) and printed:
[Code block omitted in compact Google Doc — full source at https://mikelev.in/futureproof/operation-weebo-xshape-cookie-cutter-rfc9111-smeaton/]
Why did tkinter say missing in the first test and True in the second?
Because in Python's standard library, tkinter/__init__.py is just a pure-Python coat hanging in the closet. find_spec("tkinter") checks whether the .py file is on disk (it is, so it returns True). But when you actually wear the coat (import tkinter), line 38 executes import _tkinter—the compiled C shared library that talks to Tcl/Tk and X11. On NixOS, unless tkinter is explicitly added to python3.withPackages, _tkinter.so is omitted to save space.
That one straddle saved us from writing a tkinter script that would have passed a spec check and crashed on launch!

---

3. The Ultraviolet Catastrophe of CLI Flags (And the 4 Rungs of Nix Shell Mastery)

In 1900, Lord Rayleigh and Sir James Jeans used classical electromagnetism to calculate how much light a hot object emits at different frequencies. Their math worked at low frequencies (infrared), but as the frequency climbed toward ultraviolet, the equation predicted infinite energy—a failure physics textbooks call the Ultraviolet Catastrophe.
Max Planck solved it in December 1900 by proposing that energy isn't a continuous sliding scale of infinite decimals—it comes in discrete, countable packets called quanta (

$$E = h\nu$$

).
Now apply Planck's insight to the Unix command line:
• find + grep is a Cognitive Ultraviolet Catastrophe: Does -name go before -type f? Did you remember the . path argument? Is -i for find or grep? Did you terminate -exec with \; or +? The combinatoric state space explodes into infinity for a human brain in the middle of a flow state.
• rg (BurntSushi's ripgrep) and Your posts / postsc / postsb Words Are Planck Quanta: Instead of memorizing flag permutations, you quantize each distinct intent into one short, memorable word:
• posts → titles & slugs
• postsc → holographic context shards (+c)
• postsb → full markdown bodies (+b)
• rgx / rgxc → regex search / regex search + shards
And your journey through how those words live in the shell is the exact 4-rung ladder every systems builder eventually climbs:
[Code block omitted in compact Google Doc — full source at

https://mikelev.in/futureproof/operation-weebo-xshape-cookie-cutter-rfc9111-smeaton/]

---

4. The French Taunter, Ring -3 Silicon Backdoors, and Why Graphics Is the Final Boss of WORA

When the French guard on the castle wall in Monty Python and the Holy Grail (1975) taunts King Arthur's knights—"Your mother was a hamster and your father smelt of elderberries!"—he is mocking people who confuse the pomp of the armor with who holds the castle.
In computing, hardware and OS tribalism is the exact same delusion:

•

The Silicon Supercops Underneath Every Chip (Ring -3): * Intel x86: Since 2008, every Intel chipset has shipped with the Intel Management Engine (IME)—a separate processor inside your processor running a proprietary fork of Andrew Tanenbaum's MINIX 3, with direct DMA access to your RAM and network card even when your OS is asleep. * AMD x86: Ships with the Platform Security Processor (PSP) (an on-die ARM Cortex-A5). * ARM SoCs: Partitioned by TrustZone Secure World firmware blobs. * 3-for-$1 Alibaba RISC-V: The ISA is open, which means there is no single vendor NDA—instead you get mystery mask-ROMs or silicon errata like GhostWrite (2024), where an unprivileged vector instruction on the T-Head C910 could write directly to physical memory bypassing the MMU.

•

Making Hardware Flow Like Water (Dolstra, Hemel, and Stallman): * Eelco Dolstra (in his 2006 Utrecht PhD thesis The Purely Functional Software Deployment Model) and Armijn Hemel realized that if you treat software packages like immutable values in a functional programming language (/nix/store/<hash>-name), the underlying machine stops being a precious pet. * Richard Stallman (RMS) and Ludovic Courtès recognized the brilliance of Dolstra's model and built GNU Guix on the exact same functional store architecture using Guile Scheme. * When your environment is a Nix flake, an Intel workstation, a retired 2017 laptop, an Apple Silicon Mac, or a $10 RISC-V board are all just interchangeable pipes carrying POSIX water.

Read the complete article and full code receipts online:
https://mikelev.in/futureproof/operation-weebo-xshape-cookie-cutter-rfc9111-smeaton/