ENGINIOUS GROUP · RESEARCH PAPER
MAY 2026 · CREATIVE COMMONS CC BY 4.0

Metaverse Reloaded

The $83 Billion Detour

Eighty-three billion dollars bought one lesson:
the metaverse was never the headset.

On December 1, 2025, WebGPU shipped by default across every major browser. The other architecture is now infrastructurally feasible. This paper is about which one we are about to live in.

PIOTR KOCEL · ENGINIOUS GROUP S.A.

ABSTRACT

The conventional wisdom of 2025 is that the metaverse failed. Meta's Reality Labs has lost $83.6 billion over six years. Apple Vision Pro production halted in late 2025 with fewer than 90,000 units shipped. Global VR headset shipments fell 14% year on year, and the dedicated MR/VR category is on track for a 42.8% contraction. Gartner placed "metaverse" deep in the Trough of Disillusionment.

This paper argues that this reading is precisely wrong — not because the numbers are wrong, but because they describe only one of two architectural bets that the word "metaverse" silently conflated. Architecture A — immersion through dedicated head-mounted displays — was prosecuted at extraordinary expense and stalled. Architecture B — ubiquity through GPU-accelerated 3D in the browser — was infrastructurally blocked until December 1, 2025, when WebGPU shipped by default across Chrome, Edge, Firefox, and Safari simultaneously.

Architecture A reaches an installed base of roughly 22 million active devices. Architecture B reaches the roughly 5.5 billion people who open a browser. The ratio is 250 to one. Roblox alone now serves 151.5 million daily active users — a 70% year-on-year increase — through a proto-metaverse delivered without a headset. Three.js, the dominant browser-3D library, pulls 2.7 million downloads per week and is growing at ~40% annually. Figma has migrated its entire rendering pipeline to WebGPU. Gaussian Splatting renders photoreal volumetric capture directly in a browser tab.

The hype cycle predicts what happens next: Architecture A consolidates around its small, defensible niches; Architecture B begins the Slope of Enlightenment, then the Plateau of Productivity. The $83 billion detour is not the end of the metaverse. It is the end of the metaverse-as-headset. We are now entering the metaverse the way we always entered the internet: at a URL.

> We are not waiting for the metaverse. We are waiting for the headsets to stop blocking our view of it.

§ 01

The Trough Is Real

We begin with the obituary, because it is well-earned.

Meta's Reality Labs division, the corporate vehicle through which Mark Zuckerberg prosecuted the metaverse-as-headset thesis, has accumulated $83.55 billion in operating losses across the six fiscal years from 2020 through 2025. The annual pattern is unforgiving: $6.62 billion (2020), $10.19 billion (2021), $13.71 billion (2022), $16.12 billion (2023), $17.72 billion (2024), $19.19 billion (2025). Losses widen every year. Revenue does not. Reality Labs generated roughly $2.2 billion across all of 2025 — a spend-to-revenue ratio approaching nine to one.

In 2025 alone, Meta closed three of its first-party VR studios: Sanzaru Games (developer of Asgard's Wrath), Armature Studio, and Twisted Pixel (which had just shipped Deadpool VR). The closures were not framed as failures. The losses were not framed as a reconsideration. The narrative continues. The balance sheet does not lie.

Apple's entry was supposed to validate the category. It did the opposite. Vision Pro shipped roughly 390,000 units across 2024 against an internal target of three million. By Q4 2025, IDC estimates put the figure at 45,000 units for the holiday quarter. Apple's manufacturing partner Luxshare halted production. Marketing was cut by more than 95%. Apple Vision Pro is, at this writing, effectively a discontinued category-defining product.

FIGURE 1B

Apple Vision Pro: from category-definer to discontinued

Annual shipments, thousands of units. Source: IDC Worldwide Quarterly AR/VR Headset Tracker.

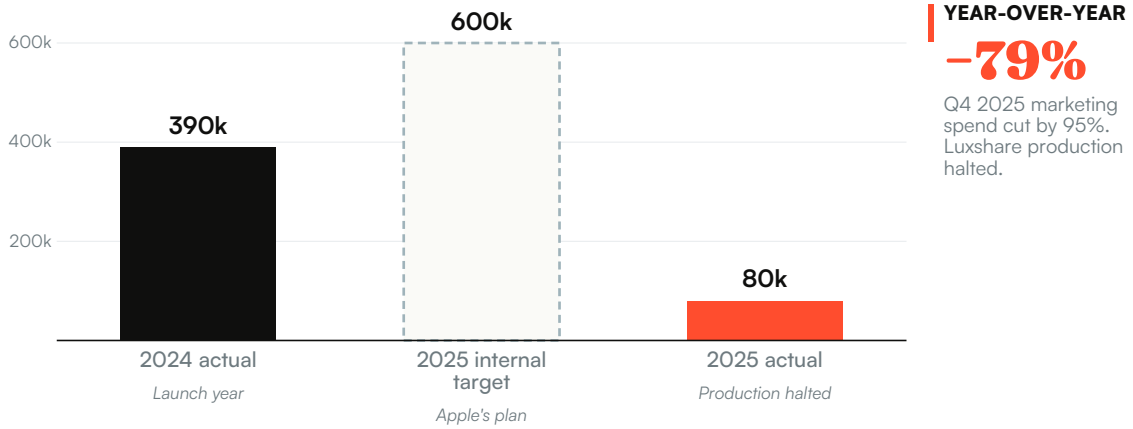

The broader market follows. Counterpoint Research reports the global VR headset market fell 14% year on year in 2025. IDC's Worldwide Quarterly AR/VR Headset Tracker forecasts the dedicated MR/VR category to contract 42.8% once final 2025 numbers are tallied. Meta Quest shipments fell from 5.6 million units in the first three quarters of 2024 to 1.7 million in the same period of 2025 — a 70% collapse for the category leader. The total installed base of active VR

headsets globally now stands at 21.9 million, down 8% year on year. That figure is roughly the population of metropolitan Beijing.

FIGURE 1

Meta Reality Labs operating losses

USD billions, fiscal years 2020 to 2025. Source: Meta Form 10-K filings.

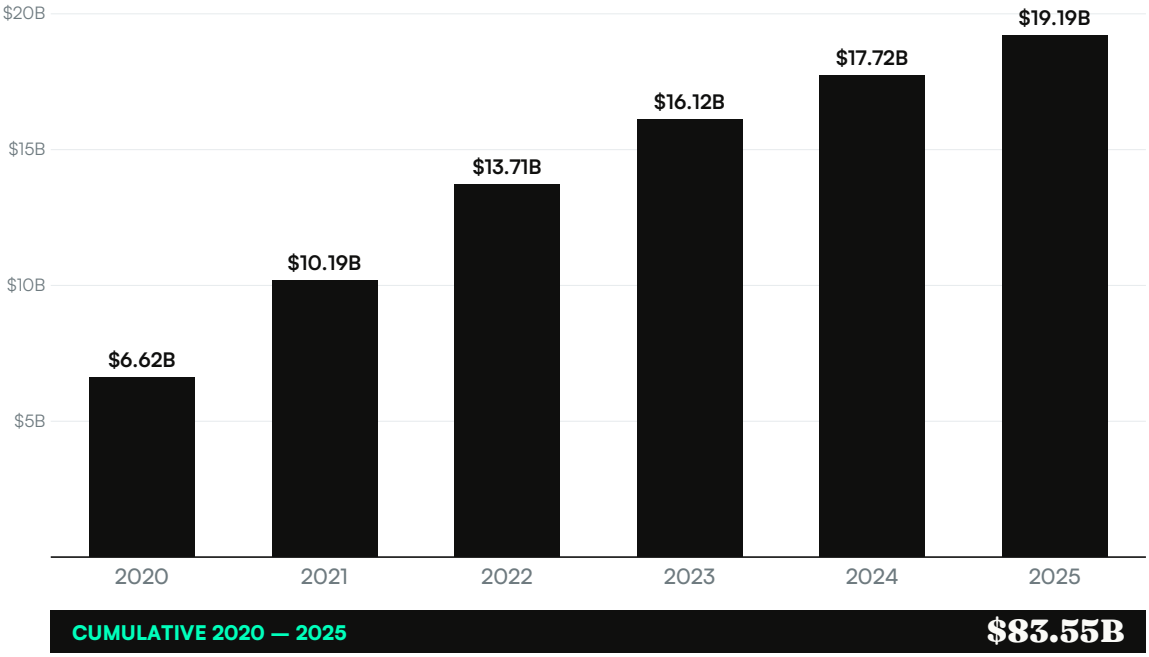

Gartner placed "metaverse" on its 2022 Hype Cycle for Emerging Technologies with a forecast horizon of more than ten years to the Plateau of Productivity. By 2023, sector analysts agreed the descent into the Trough of Disillusionment was occurring "quicker than expected." By 2025, the trough was not a forecast. It was the financial statement.

The trough is real. It is also incomplete. Because what failed was not the metaverse. What failed was a single architectural interpretation of it.

§ 02

But the Trough Was Architecture A Only

The press release of the metaverse — the one with the legless avatars, the Horizon Worlds screenshots, the Vision Pro spatial photos — was always selling a single architectural premise: that mass-market immersive 3D requires a dedicated head-mounted display. Call this Architecture A. Its binding constraint is friction. Hardware cost. Motion sickness. Battery weight. Social acceptability. Software install. Account creation. Headset onboarding. Every one of these constraints compounds. Each one reduces the addressable market by a multiplicative factor. The result is an installed base measured in the low tens of millions after a decade of effort and roughly $100 billion in combined industry capital expenditure.

Architecture B was always different. Its premise is that mass-market immersive 3D requires only a browser tab. Its binding constraint is not friction but graphics API capability — a problem of software, not hardware, that until very recently kept browser 3D bound to WebGL, an API designed for the GPU architectures of 2008 and shipped to consumers in 2011. While Architecture A was building hardware, Architecture B was waiting for the spec process to finish.

FIGURE 2

Architecture A vs. Architecture B

The two architectural bets conflated under the single word "metaverse."

	ARCHITECTURE A the headset bet	ARCHITECTURE B the browser bet
DOMINANT PREMISE	Immersion through dedicated head-mounted displays	Ubiquity through GPU-accelerated 3D in a browser tab
BINDING CONSTRAINT	Hardware friction: cost, ergonomics, install	Graphics API capability — resolved Dec 1, 2025
INSTALLED BASE	≈22 million active VR headsets	≈5.5 billion browser-connected users
DISTRIBUTION MODEL	App store + physical hardware purchase	URL — no install, no account, no download
ACTIVE USERS (ENGAGED)	Tens of millions	600—700 million MAU across proto-metaverse
CAPITAL ALLOCATED	≈$100B (Meta + Apple + others)	≈$0 in dedicated capex — W3C standards work
POSITION ON HYPE CYCLE	Trough of Disillusionment	Slope of Enlightenment (entry)

The numbers under Architecture B were always there. They were simply miscategorised as "games" rather than "metaverse," because they did not arrive inside a headset. Consider Roblox. In its Q3 2025 shareholder letter, the company reported 151.5 million daily active users, up 70% year on year. Hours engaged: 39.6 billion in the quarter, up 91%. Peak concurrency: more than 45 million simultaneous users on August 23, 2025. Roblox is delivered overwhelmingly through browsers and lightweight native clients — no headset required.

Fortnite reports 236 million monthly active users. Aggregated across the top platforms, the proto-metaverse — persistent, shared, real-time 3D worlds — serves 600 to 700 million people every month. That figure is more than thirty times the total installed base of all VR headsets ever sold and still active.

Consider e-commerce. Shopify's own published data shows that merchants who deploy 3D product viewers experience a 94% average increase in conversion rate. Return rates fall by approximately 40%. These are not theoretical projections — they are measured outcomes from production deployments in furniture, apparel, footwear, and consumer electronics. The 3D

commerce revolution is being delivered through WebGL, in browser tabs, to customers who will never own a headset.

Consider professional tools. Figma — a $20 billion design platform — migrated its entire rendering pipeline from WebGL to WebGPU in 2025. Not for graphics fidelity, but because WebGL's draw-call overhead was throttling collaborative work at scale. The single largest production deployment of browser-based GPU rendering in the world is not a metaverse company. It is a SaaS tool used by twenty million people who never noticed the migration.

The trough is real. It is also architecturally specific. While Architecture A spent $83 billion learning that consumers do not want to wear goggles, Architecture B quietly accumulated 600 million users, billions of engaged hours, and double-digit-billion-dollar SaaS deployments — all without a single headset.

> While Architecture A spent $83 billion learning that consumers do not want to wear goggles, Architecture B quietly accumulated 600 million users.

§ 03

December 1, 2025: The Catalyst

Architecture B was, until very recently, throttled by a single technical constraint: the browser graphics API. WebGL — the only cross-browser GPU interface available to web developers from 2011 through 2025 — is an OpenGL ES 2.0 wrapper. It was designed for the mobile GPUs of the late 2000s. It assumes a single-threaded command queue. It requires expensive translation layers when running on the underlying APIs of modern systems: Direct3D 12 on Windows, Metal on Apple silicon, Vulkan elsewhere. It supports no compute shaders. It has no mechanism for general-purpose GPU work. It is, in 2026, structurally obsolete.

WebGPU is its successor. It is not "WebGL 2.5." It is a ground-up redesign of browser graphics around the architectural assumptions of modern GPUs — explicit pipelines, compute shaders, multi-threaded command buffers, near one-to-one mapping to Metal, D3D12, and Vulkan. It exposes general-purpose GPU compute to the browser for the first time. It enables on-device machine learning inference at approximately 80% of native speed (as measured by WebLLM). It enables real-time physics simulation with hundreds of thousands of particles. It enables Gaussian Splatting renderers — photoreal volumetric capture — directly in a browser tab on a mobile phone.

The spec process began in 2017. Chrome and Edge shipped WebGPU by default in version 113, in May 2023. Then nothing happened for two years. Apple, Mozilla, and Linux remained the gap. Production deployment was impossible because no developer could assume their users would have the API.

FIGURE 4

WebGPU rollout across major browsers

Two and a half years from first ship to universal default.

That gap closed in the autumn of 2025. Safari 26 (September 2025) brought WebGPU to macOS Tahoe, iOS 26, iPadOS 26, and visionOS 26. Firefox 141 (July 2025) shipped it on Windows. Firefox 145 (November 2025) extended it to macOS on Apple silicon. On November 25, 2025, Google's Chrome for Developers account announced the milestone. On December 1, 2025, the formal cross-browser announcement was published on web.dev. Coverage is now approximately 70% of the global browser market and growing. Linux on Firefox stable and Firefox Android remain in progress for 2026.

The performance implications are not incremental. Babylon.js — one of the major WebGL/WebGPU frameworks — reports that its Snapshot Rendering mode, which leverages WebGPU's Render Bundles, executes approximately ten times faster than the equivalent WebGL path. A test moving a particle simulation from CPU JavaScript to a WebGPU compute shader took it from 8 frames per second to 60. ChartGPU, a WebGPU-native data visualization library, renders one million data points at 60fps — a workload that brings Canvas2D or WebGL charting to its knees. Nexara Labs, a commercial AR e-commerce platform, runs real-time facial AI at 58fps on iPhone 15 in Safari, up from 12fps under WebGL/JavaScript; they report a 40% conversion uplift and serve three million users in production.

The libraries that matter have already moved. Three.js, the dominant browser-3D framework with 2.7 million weekly downloads, shipped its WebGPU renderer in release r171 (September 2025) with automatic WebGL 2 fallback — zero-config migration for existing codebases. Babylon.js has supported WebGPU since version 5.0 (May 2022). PlayCanvas, Unity, Unreal Engine, ONNX Runtime, Transformers.js: all support WebGPU as of the rollout window. The framework question is settled. The deployment question is settled. What remains is adoption.

December 1, 2025 will be remembered the way May 27, 1995 is remembered by people who were there: the day a foundational web standard achieved enough cross-browser coverage to become deployable. Architecture B was infrastructurally blocked. It is no longer.

§ 04

Why Pixel Streaming Was Always a Detour

Between the death of Unity Web Player (2015) and the maturation of WebGPU (2025), the official answer to "how do I run AAA-quality 3D in a browser" had no good answer. Epic Games dropped HTML5 export from Unreal Engine at version 4.24 in 2019, citing low adoption and technical complexity. For UE5 specifically, Epic's only sanctioned path to the browser was, and remains, Pixel Streaming.

Pixel Streaming is a clever idea executed under unforgiving constraints. The 3D application renders on a GPU-equipped server in a data centre. Each frame is captured, H.264-encoded, and shipped to the client browser over WebRTC. User inputs are captured by the client and forwarded upstream. The client browser becomes, effectively, a thin video terminal.

The economics do not work outside of high-value enterprise use cases. A G4dn instance on AWS — the cost-effective sweet spot for UE5 pixel streaming — runs in the range of $0.50 to $1.50 per hour per concurrent user. Each session consumes 5—20 GB per hour of egress bandwidth at 1080p and 60fps. Server provisioning is one-to-one with concurrent users: every additional viewer demands an additional GPU. A consumer product with 10,000 simultaneous users requires 10,000 GPUs available on standby in geographically distributed data centres. The unit economics for a free-to-play web experience are absurd.

The latency physics are worse than the economics. Pixel streaming adds, at best, 30—80 milliseconds of round-trip delay between user input and visible response — a hard floor set by encode time, network propagation, and decode time. For passive product tours and architectural walkthroughs this is tolerable. For competitive games, virtual try-on, or any input-sensitive interactive experience, it is not. Google Stadia closed in January 2023 having proven the technology while failing the business model. Amazon Luna has been progressively reduced. Cloud gaming, as a consumer category, survives at the margins through GeForce NOW (which streams the user's already-purchased games) and Xbox Cloud (which is a bundled feature of an existing subscription).

WebGPU eliminates the trade-off. Wonder Interactive, a small studio that received an Epic MegaGrant in 2024, built a native WebGPU Rendering Hardware Interface for Unreal Engine 5. Their SimplyStream platform now exports UE5 projects to WebAssembly + WebGPU and runs them client-side at near-native performance — no GPU server farm, no input latency, no per-user egress cost. The economics flip from "GPU per user" to "GPU on user's device." The provider serves static .wasm and asset files from a CDN. The unit economics become the unit economics of any web application.

This is the structural break. Pixel streaming was the answer when the browser had no GPU. The browser now has a GPU. Pixel streaming will persist in legitimate niches — secure remote workstations, very-large CAD models, enterprise pilots with strict IP protection requirements — but as the default path for consumer-facing 3D it is now obsolete. Architecture B's distribution model is the one the web has always used: a URL, a download of code and assets, and execution on the client.

§ 05

Why Headsets Were Always a Detour

The case against the headset-first metaverse is not technological. Quest 3 is a competent piece of hardware. Vision Pro is a remarkable one. The case is structural, and it was visible in the early shipment numbers if anyone had wanted to see it.

The internet itself is the comparison. In 2001, after the first dot-com crash, the consensus was that consumer internet adoption would be gated by hardware penetration — the PC in the household. Mass internet adoption arrived not via PCs but via a different device class entirely: the smartphone, in 2007 onward. Crucially, the killer applications of the smartphone era did not require app installation in the early years. They were websites. Google Maps, Gmail, Facebook, YouTube, Twitter — all reached their first hundreds of millions of users through the mobile web. Native apps came after the audience was established, not before.

FIGURE 5

Addressable reach: Architecture A vs Architecture B

Active VR headsets (Omdia, 2024) vs browser-connected users (ITU, 2025).

If Architecture A's reach were drawn to scale relative to Architecture B's, it would be approximately the width of the period at the end of this sentence.

The metaverse-as-headset thesis inverted this lesson. It proposed that adoption would follow hardware penetration. It demanded that consumers buy a $300 to $3,500 device, charge it, set up an account, complete an onboarding, agree to motion-sickness disclaimers, and then access content. Every one of those steps is a funnel stage. Each stage drops attrition. The compounded drop-off explains the 22 million active devices.

The cost of switching to Architecture B is, by contrast, zero. There is no purchase. There is no install. There is no account. There is a URL. The user clicks it. The 3D world loads in the same tab they were already using. They are inside the metaverse without having decided to enter it.

The numerical contrast is severe. Active VR headsets globally: 21.9 million. Active monthly internet users globally: approximately 5.5 billion. Ratio: 250 to one. iOS 26 — released to a global

installed base of more than one billion iPhones over 2025 — included WebGPU by default. With a single OS update, Apple silently made one billion handsets capable of running Three.js-class 3D experiences with hardware-accelerated rendering. No headset purchase. No marketing campaign. No press release.

The same pattern repeats across the device taxonomy. Every Chromebook in every school district can render WebGPU Three.js. Every mid-range Android phone with a Qualcomm or ARM GPU and OS version 12 or higher. Every integrated-graphics laptop. Every desktop. Every console (browsers on PlayStation 5 and Xbox Series X/S support WebGL today; WebGPU integration is expected). The browser is the most ubiquitous GPU runtime in human history. It is now a competent one.

Headsets will continue to find product-market fit in defensible niches. Industrial training, where the embodied learning gains are 4x over classroom methods. Surgical simulation and medical visualisation. Architectural design review. Specialist gaming for a passionate enthusiast base. None of these niches require, or benefit from, the consumer mass market that Meta and Apple spent ninety billion dollars chasing. The mass metaverse will not be wearing goggles. It will be in a browser tab.

> Active VR headsets: 21.9 million. Active browser users: 5.5 billion. The ratio is 250 to one.

§ 06

Signals from the Slope of Enlightenment

The Gartner Hype Cycle predicts that technologies emerging from the Trough of Disillusionment do so quietly, through accumulating production deployments rather than through new press releases. The early signals are visible in build statistics, library adoption, content tooling, and job market data — not in trade press headlines. By that test, Architecture B is well into the Slope of Enlightenment as of early 2026.

Three.js downloads. The dominant browser-3D library now ships 2.7 million npm downloads per week. Growth is approximately 40% year on year and has not plateaued. The library has 112,000 GitHub stars and is detected in production on 3,483 distinct active company domains in commercial crawl data, including Ford, Bombardier, Morningstar, and a long tail of e-commerce, agency, and SaaS deployments. Three.js outpaces its nearest competitor by roughly 270 to one in download volume.

Job market. Listings specifying Three.js or WebGL skills grew 25% in 2025. WebGPU-specific listings are a small but rapidly growing category. This is not the labour-market footprint of a dead

technology.

FIGURE 3

Roblox daily active users by quarter, 2025

Millions of DAU, with year-over-year growth. Source: Roblox Form 8-K filings.

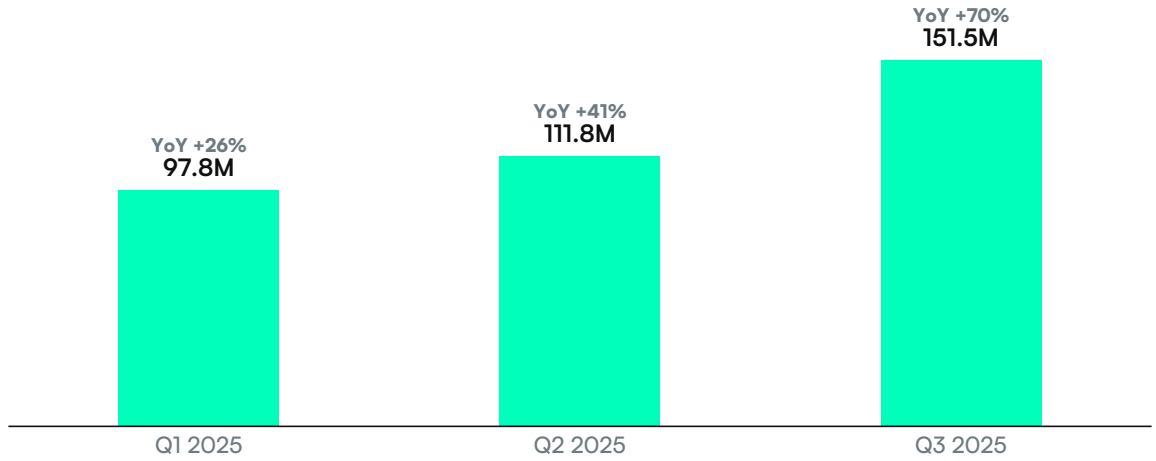

Hours engaged Q3 2025: 39.6 billion (+91% YoY). Peak concurrency Aug 23, 2025: 45M+ users.

Content pipeline. The 2021 metaverse hit a content wall: every world required bespoke 3D asset creation by trained artists at industry rates. That wall is gone. Tripo AI, building on the Stability AI TripoSR research release, generates a usable 3D model in approximately 0.5 seconds on an A100 GPU. Meshy reports more than 10 million users for its text-to-3D and image-to-3D pipelines. Luma AI, Polycam, and others have brought photogrammetry and Gaussian Splatting capture to consumer smartphones. The cost curve for 3D asset creation has fallen by approximately two orders of magnitude in three years.

Capture pipeline. 3D Gaussian Splatting, introduced at SIGGRAPH 2023, has matured into a production volumetric capture format with native browser renderers in WebGL and WebGPU, native plugins for Unity and Unreal, and official extension support in both OpenUSD and Khronos glTF (2025—2026). A splat scene captured from a phone video can be embedded in a webpage as trivially as a JPEG.

FIGURE 7

Weekly npm downloads, browser-3D libraries

Thousands of downloads per week, 2026. Source: npm registry / TechnologyChecker.io.

Three.js outpulls its nearest competitor by approximately 270× — a winner-take-most dynamic typical of late-stage standardisation. The framework question is no longer open.

Standards. OpenUSD, the universal scene description format originated at Pixar and now governed by the Alliance for OpenUSD, has continued to consolidate as the lingua franca for interoperable 3D scenes across creative tools, game engines, and now browsers. glTF 2.0 from Khronos remains the delivery format of choice for the web. Both are stable, both are royalty-free, both have wide engine support. The standards problem is solved.

Brand activations. The case studies are accumulating quietly. Nike's Air Max Day digital activations. Dior's AR campaigns. Major automotive OEMs deploying browser-based 3D configurators that displace showroom visits. Architectural firms presenting BIM-derived walkthroughs to clients via URL. Museums offering volumetric exhibition tours. Each of these is a single corporate decision; collectively they are reshaping the browser into the default surface for premium product narrative.

Each signal in isolation is unremarkable. The pattern in aggregate is not. This is what the Slope of Enlightenment looks like before the Plateau of Productivity — quiet, distributed, unmarketed, structural.

§ 07

The Gartner Parallel: Internet 2000–2005

The closest historical analogue to the current metaverse trajectory is not VR's previous cycle in the 1990s. It is the consumer internet itself, in the period bookended by the dot-com peak of March 2000 and the publication of Tim O'Reilly's "What Is Web 2.0?" essay in September 2005.

The dot-com bust was severe. The NASDAQ Composite fell 78% peak to trough. Pets.com, Webvan, eToys, Boo.com — the marquee failures of the era — became case studies in misallocated capital. Trade-press consensus by 2002 was that the internet was overhyped and that the next great consumer shift would not arrive for a decade.

FIGURE 6

The Hype Cycle, with Architecture A and B located

Architecture A is in the trough. Architecture B never had a peak — it bypassed the cycle and entered the slope quietly.

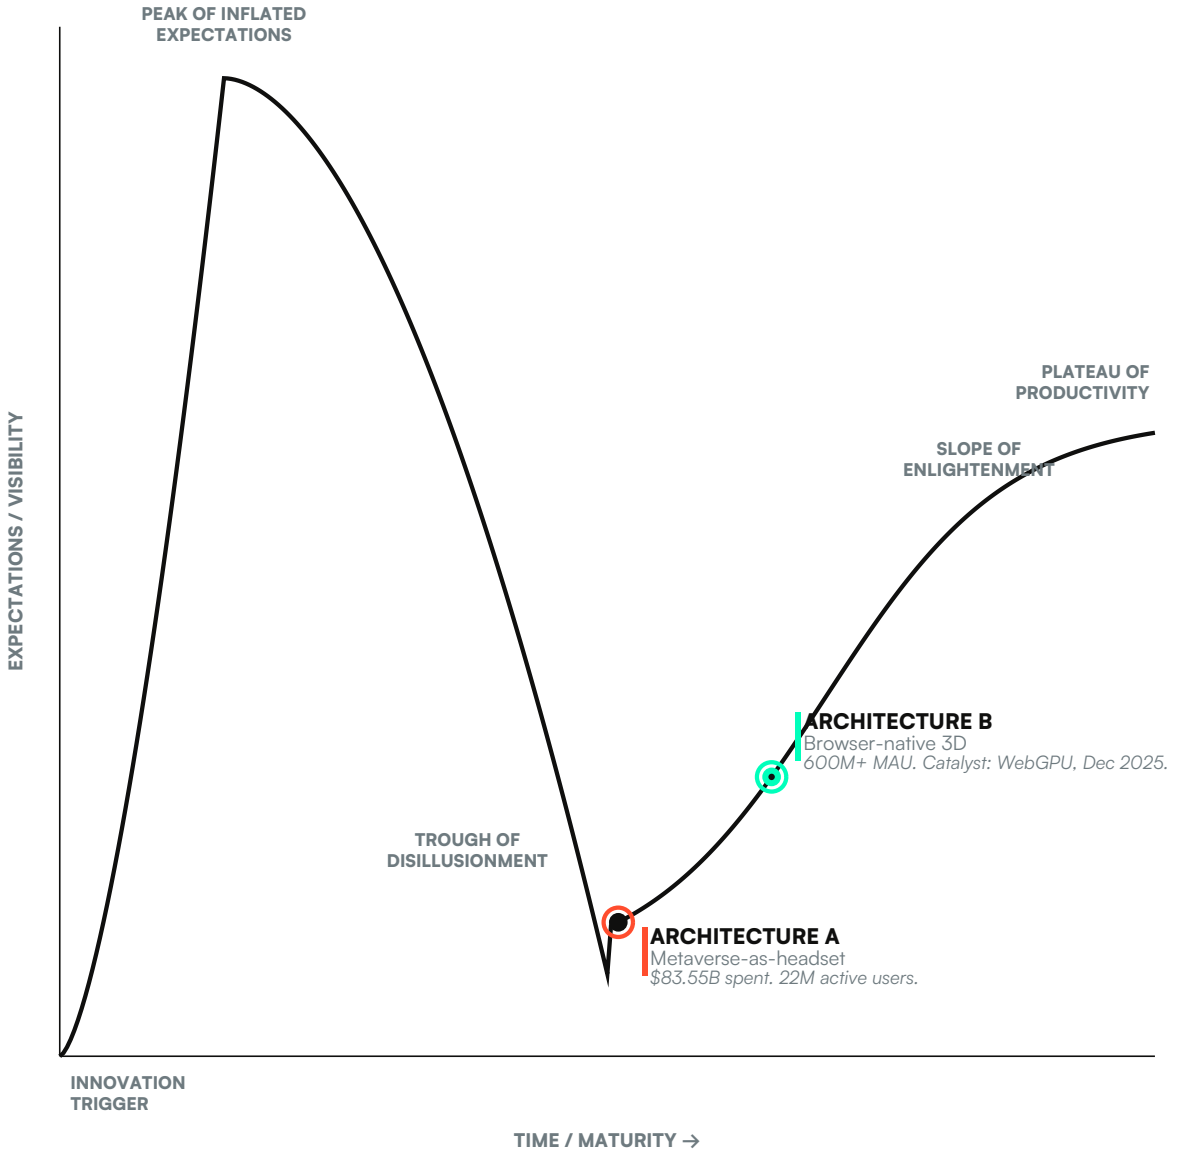

The data underneath told a different story. A University of Maryland / UC San Diego study of more than 700 venture-funded internet companies from the bubble era found a five-year survival rate of 48% — comparable to the early auto industry. eBay, Amazon, Google, PayPal, Salesforce — companies that survived the trough — built compounding businesses through 2001—2004 while the press was looking elsewhere. Broadband penetration in US households went from 5% in 2000 to 30% in 2005. The infrastructure for the second wave was being laid quietly under the trough.

When the consensus shifted back to "internet is real," it shifted quickly. From 2005 to 2010, social networking (Facebook), video (YouTube), mobile (iPhone), and cloud (AWS) all moved from speculative to dominant. The dot-com survivors and the second-wave founders captured an order of magnitude more value in the 2005—2015 decade than the bubble-era cohort ever did. The trough was not the end. It was the filter.

The mapping to the metaverse is structural, not metaphorical. Roblox (founded 2004, IPO 2021) is the eBay of this cycle: the proto-metaverse company that survived through both the bubble and the trough by running an actual business with actual users. Epic Games — Unreal Engine, Fortnite, MetaHumans, RealityScan — is the Amazon: the platform that uses the trough to entrench. Unity, despite its corporate troubles, remains structurally important. Adobe, Autodesk, and Pixar's Alliance for OpenUSD are the picks-and-shovels suppliers. The casualties — Meta's Horizon Worlds, Decentraland, The Sandbox in its bubble-era form, Magic Leap — are the Pets.coms.

The catalyst for the second wave of the metaverse — the equivalent of broadband penetration crossing the 30% threshold — is the cross-browser ubiquity of WebGPU. Architecture B becomes infrastructurally feasible at 70% browser coverage, the threshold it crossed in late 2025. This is the same threshold that, for the consumer internet, made Web 2.0 possible.

The implication for capital allocation is direct. The next decade of metaverse value capture will accrue not to the companies that built headsets, but to the companies that build the second-wave applications, infrastructure, and content on top of the new browser GPU. Most of those companies do not yet exist. Some of them, like the eBay and Amazon of 2002, have been running quietly under the trough for years.

§ 08

Architecture B in Practice

Architecture B is not theoretical. It is a stack. The components are mature, well-documented, and in production use today. A complete picture of Architecture B as it stands in mid-2026:

Runtime layer. The user's browser. Chrome, Edge, Firefox, Safari, all shipping WebGPU. WebAssembly with SIMD and threads enabled for high-performance code. WebRTC for real-time multiplayer transport. WebXR for optional AR/VR overlay on supported devices (without requiring it). WebCodecs for hardware-accelerated video. The Project Fugu APIs for file system, clipboard, peripherals.

Rendering layer. Three.js (2.7M weekly downloads) for the broad market, Babylon.js for enterprise applications, PlayCanvas for production games, native WebGPU for specialised workloads. For Unreal Engine projects, Wonder Interactive's SimplyStream provides UE5-to-WebGPU compilation. Unity's WebGPU export path is stabilising. The framework war is over: Three.js won the breadth game, the others occupy specialised niches.

Asset pipeline. glTF 2.0 as the delivery format. OpenUSD for source scene description and tool interchange. Gaussian Splatting for volumetric photoreal capture, supported natively in both. AI-generated meshes from Tripo, Meshy, Rodin for parametric asset creation at scale. Procedural generation through PCG (Unreal) or Houdini for systematic worlds.

AI and behaviour layer. WebLLM and Transformers.js for browser-native language model inference at 80% of native performance. ONNX Runtime Web for general ML inference. WebGPU compute shaders for custom GPU workloads. For agentic NPCs and conversational characters, embedded inference paired with cloud-served large-context reasoning. A category of AI-character platforms — Soul Machines, Convai, Inworld, and others, with new entrants arriving quarterly — has begun to converge on this hybrid local-cloud pattern as the production architecture for persistent, voice-driven browser characters.

Multiplayer and persistence layer. WebSocket and WebRTC data channels for transport. Conventional cloud backends (Firebase, Supabase, custom Node/Go services) for persistence. The data layer is solved. The novel work is in the synchronisation models for spatial state.

Distribution layer. A URL. A CDN. SEO. The same growth and distribution mechanics that govern every other web application. No app store gate. No platform tax. No 30% commission. No platform-specific build. Cross-device by default.

The stack is not exotic. Every layer of it is in production use somewhere. What is new is that the layers now compose into experiences that, in 2024, would have been technically impossible without an Unreal Engine native build, a headset, or a pixel-streaming server farm. The threshold of what can ship in a tab has moved by an order of magnitude in twelve months.

§ 09

Implications

For brands and retailers. The browser is now a competent 3D runtime on the device every customer already owns. The marginal cost of a 3D product viewer, configurator, or virtual showroom has fallen by an order of magnitude in three years; the conversion uplift on production deployments is documented at 94% with return-rate reductions of 40%. The question is no longer whether to deploy. It is which experiences to prioritise and which technical partner to use. The customers of Architecture A are not coming. The customers of Architecture B are already there.

For agencies and creative studios. The skill curve has shortened. Three.js with React Three Fiber is now within reach of any team with a competent senior front-end engineer; the artistic talent that previously sat in game studios is now deployable in a browser context. Agencies that move their senior 3D talent from headset-targeted projects to browser-targeted projects in 2026 will capture a disproportionate share of the next wave of brand work. Agencies that wait for Vision Pro to recover will not.

For platforms. Roblox, Epic, Unity, and the new entrants will divide the consumer browser-metaverse market. Adobe, Autodesk, and the OpenUSD alliance will provide the production tooling. The interesting strategic question is which non-game platforms — design tools (Figma), creative tools (Canva), commerce (Shopify), communication (Zoom) — incorporate ambient real-time 3D as a feature within the next eighteen months. The technical barrier has dropped to a level where the answer is: all of them, if they want to remain relevant.

For investors. The capital cycle pattern of dot-com bust to Web 2.0 recovery suggests that the firms which will capture disproportionate value in the 2027—2032 window are being founded now. They are not building headsets. They are not building pixel-streaming infrastructure. They are building browser-native 3D applications, AI-driven content pipelines, spatial collaboration tools, and the picks-and-shovels infrastructure (asset hosting, format conversion, optimisation, observability) that the new stack will need at scale. The 22 million active VR headsets are the ceiling for one architecture; the 5.5 billion browser users are the floor for the other.

For policymakers and institutional buyers. European sovereignty in the immersive sector does not require a European Meta. It requires a European ability to deploy and govern browser-native 3D infrastructure — the rendering stack, the AI character runtimes, the volumetric capture pipeline, the standards participation in W3C, Khronos, and OpenUSD. This is a more achievable strategic posture than European VR hardware competition. The browser is, by design, a sovereignty-friendly platform.

For everyone who has been told, repeatedly, that the metaverse failed. The metaverse did not fail. What failed was a specific architectural prosecution of it, at a specific cost ($83.55 billion), in a specific time window (2020—2025). The thing the headsets were trying to reach is now reachable

through a different door, which has been propped open since December 1, 2025, and which billions of people pass through every day without thinking about it. The next decade of 3D — for commerce, for entertainment, for collaboration, for learning, for representation of physical space — happens through that door.

The infrastructure was being built quietly all along. The door is now open.

REFERENCES

Apple (2025): Safari 26 Release Notes. webkit.org.

Counterpoint Research (2026): Global VR Headset Market Tracker.

Figma Engineering Blog (2025): "Figma rendering: Powered by WebGPU." figma.com/blog.

Gartner (2022, 2023, 2024, 2025): Hype Cycle for Emerging Technologies; Hype Cycle for AI.

Google Chrome / web.dev (2025): "WebGPU is now supported in major browsers." web.dev/blog.

IDC (2025—2026): Worldwide Quarterly AR/VR Headset Tracker.

Inria et al. (2023): "3D Gaussian Splatting for Real-Time Radiance Field Rendering." SIGGRAPH 2023.

Meta Platforms, Inc. (2021—2026): Form 10-K and Q1—Q4 earnings releases, Reality Labs segment.

Mozilla (2025): Firefox 141, 145 Release Notes.

Omdia (2024): "VR Headset Sales Decline 10%." Omdia Research.

Roblox Corporation (2025): Q1, Q2, Q3 2025 Shareholder Letters (Form 8-K).

Shopify (2024): "3D Ecommerce and the New Era of Retail." shopify.com/blog.

Stability AI (2024): TripoSR research release. stability.ai.

TechnologyChecker.io (2026): Three.js domain adoption and npm analytics.

W3C GPU for the Web Working Group (2017—2025): WebGPU Specification.

Wonder Interactive (2024—2026): SimplyStream platform announcement; UE5 WebGPU RHI.

Utsubo (2026): "What's New in Three.js (2026): WebGPU, New Workflows & Beyond."

COLOPHON

Author: Piotr Kocel, Prezes Zarządu, Enginious Group S.A. Contact: piotr@enginious.tech. The author would like to acknowledge the research and engineering teams at Enginious Group for source material, fact-checking, and technical review.

Published by Enginious Group S.A., ul. Łąkowa 29, 90-554 Łódź, Poland. KRS 0001007157 • NIP 7262678729 • enginious.tech

Citation: Kocel, P. (2026). Metaverse Reloaded — The $83 Billion Detour. Enginious Group Research, May 2026.

Typeset in Satoshi (Indian Type Foundry, Fontshare) and Fraunces (Undercase Type). Rendered with ReportLab.