Beyond Copy-and-Paste:
How Game Studios Are Reorganizing Around AI

Zimran Ahmed
zimran@wharton.upenn.edu
Senior Fellow Wharton Generative AI Labs
The Wharton School, University of Pennsylvania
April 10, 2026

Summary

The gaming industry has long used AI in its products and tools. This would suggest that the incorporation of Generative AI, the newest iteration, would be more straightforward than it's reported to be in many other industries, especially given its potential to transform the development process. But our research—20 interviews across game studios from AAA to indie—found a consistent pattern: organizations began by empowering individuals and automating existing, tractable tasks, but struggled when workflows became more complicated and spanned multiple teams. Progress came when individuals used AI to reach across domain boundaries, completing adjacent tasks themselves. The most dramatic results, however, came from studios designed around AI from day one, where small generalist teams replaced specialist silos and cycle times collapsed from months to weeks. The implication: don't stop after asking "what can we automate?" Also consider "how would we build if we designed around AI from day one?"

Introduction

At a recent European technology conference, an indie studio CEO demonstrated their feedback loop: a user reports a bug in Discord, a workflow automation bot creates a ticket in Linear, which Cursor responds to by generating a working prototype that fixes the issue and is ready to review minutes later. All with no human intervention. A programmer from a larger studio messaged afterward: "We still have a guy putting issues from Slack into Jira... and you have a whole pipeline using AI!"

This anecdote captures both the promise and the gap that exists in the gaming industry. Some studios have reimagined their workflows around AI, but many have not. To understand what separates them, we conducted 20 interviews with practitioners: product managers, engineers, and designers, across AAA, mobile, and indie game studios in the US and EU, all of which had at least some AI adoption. We spoke directly to learn what was actually happening on the ground three years after ChatGPT's launch.

The Sensible Starting Point

Most organizations started with giving individuals access to LLMs. This made sense, but had limits.

The earliest experiments included studios building their own private chatbots. One AAA studio built an internal tool based on GPT, and while it met legal and security requirements, it did not work well enough for users. The homegrown tools demonstrated the interest studios were taking in the new technology, however, and once foundation model developers (OpenAI, Google, etc.) offered enterprise tiers with sufficient data protection, they were adopted. We found that every existing studio began its Generative AI journey the same way: provisioning a secure LLM tool, such as ChatGPT Enterprise, and offering it to employees, usually without training, structured field experiments, or organizational assessments. They also turned on AI features in existing video conferencing tools (Zoom, Google Meet) and productivity software (Microsoft Office, Google Workspace), giving workers new tools to manage their own work.

We refer to these as "Copy-and-Paste AI" deployments because the user puts context into a chat window, gets a result, copies it and pastes it back into their work. The AI doesn't remember anything between sessions — each conversation starts from scratch — and it requires a human to shuttle information back and forth. These deployments exposed people to the technology, required no organizational change, and let individuals experiment within their existing workflows. This organic approach to AI adoption quickly resulted in the regular use and normalization of AI-generated notes and follow-up items, making them a part of existing standard operating procedures (SOPs).

Even so, adoption at this basic level wasn't always seamless. Some studios required employees to formally request access:

"They're actually trying not to give it to people. And you have to justify why you need it." — Product Manager, AAA studio

"I was clicking around and I found a link to it and I was like, I'm gonna request it. So I got it. And then I started spreading it to all the other product managers in my team." — Product Manager, AAA studio

Despite this friction, we found anecdotal productivity gains. People used AI to help with writing, particularly tailoring communications to different audiences, and with project management, where it made

tracking commitments and sharing meeting summaries easier. In fact, 7 of 20 interviewees went further, building persistent AI context systems—markdown files, learning documents, or custom GPTs with memory—that improved over time.

"Of course I use ChatGPT a lot to get started: ‘can you outline what should the content be? Can you highlight the main reasons for doing this?’ And so I get a lot of input from people and machines." — CTO, mid-size studio

"The first time I talked with OpenAI I felt like I was finally talking to someone more senior than I was... it was like talking to a mentor and having him available for me 24/7." — Art Manager/Producer, startup studio

"One of the things that GPT really helps me with is writing OKRs. Because I hate writing OKRs. Like I know what my goals are, I just hate having to phrase it." — Product Manager, AAA studio

One product manager at a AAA studio described using GPT to "tweak [their] language depending on [the] audience" — putting things "in the lens of an exec or in the lens of someone who's an individual contributor."

"It's never contributed to final draft work. But it's always a good enough draft that I can use it to get me to the next step." — Senior Product Manager, AAA studio

These individual productivity improvements did not ladder up to broader organizational impact. Copy-and-Paste AI is fundamentally a siloed, individual practice that happens repeatedly across the organization. Some people saved prompts, building a library of personal resources, but this was idiosyncratic and disconnected from organizational upskilling efforts. Employees reported benefits, but the studios operated in the same way they had before. To capture the opportunity that AI offered, organizations realized they needed to connect and automate entire workflows. But this led to a second group of challenges.

Looking for Pilots

To go beyond the individual, studios looked for workflow automation projects they could pilot, which highlighted the challenges with organizational change.

The sensible next step was workflow automation: applying AI to an existing process to complete it faster and save money in the process. Yet this seemingly logical and direct approach faced challenges, not with the technology itself, but with the organizational coordination required.

The challenge was tacit knowledge. Multi-person workflows run on unwritten rules, institutional memory, and "how things really get done." (1) This knowledge is fragmented across many individuals within the organization, often leaders within specific teams or pipelines. The true process often differs from what's documented in SOPs. Before AI could automate reliably, someone had to extract, reconcile, and codify all of this, which is a difficult task.

There was also a human dimension. Tacit knowledge can be a form of job security — if no one else knows how to do something, then you become indispensable to the organization. So when organizations announced top-down pilots aimed at efficiency, some employees understandably grew cautious about

codifying their knowledge. Artists in particular felt concerned, not just about job stability, but about whether AI itself might diminish the status of their field and the quality of organizational output if “slop” became accepted as "art." Management tried to select pilots with sensitivity to these concerns, but the dynamic was pervasive.

"In the beginning there were several crisis meetings... they were just afraid for their jobs... they were afraid that instead of being the one actually performing or preparing the art, they were just being a prompter instead." — CFO, mobile studio

"The art director said 'I will support this, but then I will go and protest. In other places.'" — VP of Engineering, mobile studio

"We did hire a concept artist who was reluctant... now she cannot work without it. You need to make a weapon and then make 10 variations—you can get a machine to do the variations. She makes the original, and then she can make all variations and pick the nine best." — CTO, mid-size studio

Given these challenges, organizations looked for pilot projects that balanced three factors: important enough to justify the investment, feasible enough given the upfront codification work, and culturally sensitive enough to avoid triggering resistance within the organization. Examples included:

1. Accelerating concept art pipeline
    - Tuned an AI model to the studio style so it could be used to create concept art
    - This met with more resistance from the artists, as it was considered creative and required a high degree of craft, and needed more executive sponsorship.
2. Automating standard asset creation
    - Used AI to automate the production of updated standard art assets with the goal of reducing the time needed by 50%
    - This was broadly supported by the artists because there was little creativity involved and it was repetitive
3. Converting graphics assets from an old format to a new one
    - Traditional approaches had converted half the assets already but could not be applied to the remainder due to technical complexities. Generative AI could finish the job and in the process, actually learn how to do additional similar conversions in the future. Note this was dependent on having experts build the models to support this new pipeline.
    - There was no reported resistance from the art teams on this because the work was automated from the beginning. AI was simply a more flexible automation tool.

If Art represented the department with the most reservations about AI, Engineering exhibited the least (although they too had concerns). In part, this was because the AI labs have focused on coding use cases and their use in this space was highly effective from the start. But it was also because engineering work already involves making tacit knowledge explicit: working with stakeholders to understand how a business process should work and encoding that logic. Many engineers were also accustomed to building systems to work across departmental boundaries. This existing pattern mapped more naturally to how AI needs to be set up for reliable output. In some ways, engineering pointed toward how other disciplines might work in the future.

"You must be precise. If you can't articulate the logic, the AI can't write the code." — Solo Developer

Reaching Across Boundaries

Progress came when people used AI to reach into adjacent domains. They understood the problem space well enough to direct the AI, and just needed help with technical execution.

Top-down pilots faced the twin challenges of extracting tacit knowledge and overcoming resistance. But we observed another pattern within legacy studios that may illuminate the path forward. Individuals used AI to complete tasks where, in the past, they would have had to rely on other teams. Unlike top-down pilots, these individuals already had the requisite tacit knowledge, they just lacked the technical skills to execute. We found them using AI to "fill in the gaps."

We saw this approach also being used for marginal jobs—tasks which were important at a local level, but struggled to get funded or prioritized organizationally. This enabled work to get done that wouldn't have happened otherwise, without threatening the status quo. Like pilot programs, these efforts spanned multiple departments, but were bottom-up and went beyond "Copy-and-Paste AI" to meaningfully improve project outcomes.

Examples:

- **PM → Data Science:** The product manager needed to run a number of ad hoc queries that weren't available in the dashboard, and the data tables themselves were messy and complicated. He wrote a markdown document that listed all of the idiosyncrasies in the tables he knew about, and used this to augment his requests to the LLM. Each time the LLM generated a bad result, he would update the markdown document with additional information, making what was tacit explicit, and the LLM would do better. Over several months, the query failure rate dropped from 70-80% to 5%. They were now scaling this company-wide.

"I basically thought, well, how did I ever write queries? I would query the database to figure out what the structure looks like and then I would plan my query. So I basically did the same thing with Claude... I went from having a fail rate of about 70 to 80% and now a fail rate of almost 5%." — Product Manager, mobile studio

- **Artist → Code:** An artist who had no experience with coding started doing so with AI. He was able to lean on his more technical, experienced teammates for help when needed, and translate his creative direction directly into the code, letting the AI handle syntax.

"[The designer] began coding in this past month using AI — he'd never really done it before, though he understood the basic logic." — Developer, Roblox platform

- **Finance → Everything:** The finance department would have to routinely ask other teams to manually send them current numbers so they could update their financial projections, and the firm had higher-priority projects for its data science teams. So they used AI to write queries that would pull the data automatically. Financial model updates went from 2-3 days of work to a push-button process that produced almost instant updates enabling a rapid and regular flow of accurate information to decision makers.

"We had the whole AI project... our business controller wanted to learn some more SQL and do some scripting. He used ChatGPT to generate a Google App Script that extracts data from BigQuery. We can just push the button and then our [financial] model is updated… Before, we needed to send that file to our [remote analytics team]. Then we need to wait for them to wake up due to the time difference... It took two to three days. Now we can press a button and it's ready." — CFO, mobile studio

- **Engineer → Artist:** The game needed some 2D art assets, but the studio did not have a designer in house. In the past they would have contracted out the work, but this time an engineer used AI to generate "good enough" assets themselves. Not only did this save money, but the scale of production (from dozens to hundreds) enabled new kinds of gameplay, and they extended that part of the game system to run with AI art.

"30 icons done the same day vs weeks with a contractor." — Developer, Roblox platform

- **Marketing → Art Direction:** The mobile marketing team was interested in trying different art styles to see which would resonate with their target audience. This was for concept testing, not production, so they used AI to generate 16 different style variations of their content and used that for their research. In the past, it would have required commissioning multiple concept artists for the tests alone. This is an example of why artists are concerned with AI making them redundant.

"They produced like 16 different styles... for marketing to figure out what do people react to, what do they like? And we found a winner, which is the game they're now developing." — CTO, mid-size studio

What made some of these efforts particularly powerful was how the AI improved over time. We call this pattern Read/Write AI, in contrast to the Copy-and-Paste AI approach described earlier. By saving prompts and continually adding to text files that captured what they knew about the system, the employee gave their AI memory. The PM → Data Science example illustrates this: each time the LLM generated a bad query, he updated his markdown document, making what was tacit explicit. Over months, the failure rate dropped as the knowledge compounded.

Even though we saw these approaches being taken at the margin, the efficiency they enable points toward how AI may be utilized in the future. Processes which currently extend across multiple departments of specialists with numerous handoffs can now be collapsed and managed by a smaller number of generalist employees using AI to fill technical and talent gaps. One can imagine this digital documentation, continually updated and maintained, becoming a shared substrate on which the organization runs.

That future began to come into view when we looked at new studios which had been built around AI from the beginning.

What AI-First Looks Like

The most dramatic results came from organizations designed around AI from day one.

What individuals discovered at the margin, AI-first studios built into their DNA from the beginning. Unlike incumbents experimenting with this new technology, these new studios had built it into the core of what they did. They relied on AI to accelerate development speed while reducing costs so they could operate

more efficiently, meet milestones, and raise capital to complete and launch their game. Some went further and wanted to deliver AI-enabled game experiences that hadn't been possible before.

These entrants did not need to worry about existing org structure, or layers of difficult-to-document tacit knowledge. With necessity being the mother of invention, they selected employees who were comfortable with the technology and interested in finding new ways to build.

The knowledge which used to be tacit existed as documentation from the start. Designs, playbooks, and procedures were captured in markdown files and used by LLMs to complete tasks. For example, the game design document, which detailed the pillars of the core game itself, was referenced when asking an LLM to build a loop or generate art assets. These documents were updated, sometimes by the AI, to keep them complete and current. This was Read/Write AI operating at an organizational level.

We saw small teams capable of quickly executing complete workflows which used to take multiple teams of specialists, with many handoffs, much longer. People wore multiple hats and used AI to execute tasks that spanned disciplines: product management, data science, engineering, QA, design. We call this increase in tempo and reduction in steps "pipeline collapse."

For game studios, iteration speed is the name of the game. Anything that speeds up the process of moving from idea to design to prototype to playtest helps teams identify whether something new works before investing heavily. AI enabled teams to brainstorm quickly, try different approaches, and iterate on how a game loop felt. Studios could "discover the fun" faster than before.

Team sizes in AI-first studios: One AI-first platform/engine company operated with ~22 people. Another AI-first game studio ran with just 5 technical team members, all generalists. A third studio developed detailed design documents with just two people, to guide an external developer. We even found a single developer who had built an entire, fully featured game by himself using AI.

We did not see consistent organizational structures emerging from these examples. These were small, scrappy startups still figuring things out. But if generalists using AI can span technical disciplines, one can imagine the studio's organizing principle shifting from domain (Engineering, Design, QA) to business outcome (Growth, Monetization, Retention). In gaming, teams might form around user acquisition or live ops; in banking, around loans or new accounts. It is too early to say how these forms will mature, but the direction is suggestive.

Before AI	**After AI**	**Compression**
4 months (vertical slice)	4 weeks	4x
Weeks (30 icons)	Same day	10-20x
Weeks (dev->QA cycle)	1 day	5-10x

"[Thanks to AI a] vertical slice that we spent four weeks building will look like something that normally would've taken us four months." — CPO, early-stage indie studio

"Instead of having a pipeline that has defined stages, what happens now is that you have multiple cycles that are constantly running all the time." — CEO, startup studio

"AI now rewards the generalist — even though I do believe that depth is better. Teams become smaller because they have extremely powerful people that can cross domains." — CEO, startup studio

"Myself and our COO, we are doing a lot of things which, even two years ago we would've been hiring people for or contracting for—like complex content creation, even things like legal review." — CEO, startup studio

Only 3 of 20 studios we interviewed achieved team-wide AI workflow adoption. All were AI-first from founding, built their processes explicitly on top of AI systems, and selected for employees who were culturally open to this technology. This greenfield development gave us a glimpse of what the endpoint may be for organizations that build themselves around AI.

What AI Cannot Do

AI was a poor substitute when the process itself *was important; it excelled when what mattered most was the* output*.*

While studios told us that AI was helpful at generating code, graphics, queries, etc., we also heard that it could not substitute for activities that focused on aligning teams, such as strategy. These processes required human-to-human collaboration, discussion, and interaction. Examples included:

1. **Strategic Planning** — Leadership needed input from their teams, and to work together with them to agree on the direction and objectives of the firm.
2. **Team alignment** — Teams needed to understand the priorities and tradeoffs so they could collaborate to pursue the shared goal.
3. **Building Culture** — Establishing how a company worked, what it valued, how it communicated, etc., required people to express their perspectives, understand, and agree to a process.

"It's not about the plan itself. It's the thinking and reflection that happens in doing the plan." — CEO, startup studio

“The process [of building a roadmap] is what’s more important. The output is important too as a reference point. [but] The process is important, discovery with the leads, lots of research goes into it, understanding the trend analysis, getting perspectives, etc. Getting buy-in.” — Studio Director, startup studio

"The team needs to have ownership [over the plan]. Otherwise they are not going to be engaged with it." — Art Manager/Producer, startup studio

Consider a strategic planning exercise as an example, where a team was tasked with reviewing company metrics, analyzing industry trends, and getting input from stakeholders to produce a list of recommended investments for the next period. The formal output of this work may need to change after it has been "finalized," but the process itself remains valuable since it brings people together and had them collaborate to reach agreement. Those alignment dynamics cannot be short-circuited. If an AI had

produced the same strategy document, but the team was not part of the process and had not had that shared experience, they may not commit to the direction.

These human-driven coordination processes took knowledge that was tacit — assumptions, priorities, concerns — and worked through it so it could be documented in an explicit plan with metrics, goals, and guardrails. In the speculative future organization we described earlier, where small teams organize around business outcomes, this human coordination would be what keeps the different business outcome teams aligned. Leadership defines trade-offs, allocates resources, and sets goals; outcome teams execute against them and return to leadership when they need guidance or encounter information that changes the picture. The human-generated output becomes the input for the AI to include in the shared substrate of its digital documentation, updating it as needed after the next planning cycle.

Setting goals, defining priorities, and aligning a team around them—these are high-agency processes that humans need to drive. They require relationships, coordination, and teamwork that AI cannot replicate.

The Path Forward

Don't stop after asking "what can we automate?" The bigger opportunities lie with "how would we build if we designed around AI from day one?"

Taking a snapshot of the current industry, we saw sensible top-down efforts by incumbents to empower their people with tools and identify pilots to drive process-level efficiency. We also saw bottom-up adoption, where people increased their personal productivity, and several instances of Read/Write AI, where knowledge compounded over time.

As promising as these approaches seem, they leave organizations a long way from where AI-first entrants begin. Incumbents face complicated systems running on tacit knowledge, employees who are understandably cautious about making their expertise explicit, and the management challenge of transitioning from specialist departments to generalist teams. Nevertheless, the productivity potential is real. Read/Write AI, already emerging in software development, is applicable to any function with a well-documented playbook of robust SOPs.

Entrants face a different challenge. They may be efficient at production, but making a hit title is hard. Incumbents can rely on existing audiences, whether it's through ongoing relationships and live services (Fortnite), well-known IP (Call of Duty) or licensed properties (Marvel Unlimited/NetEase). Entrants can make great games, but must still find their audience and distribution channel, areas where AI has had less impact. As costs fall, we expect many more studios and titles. But they may struggle to find a market, and whether this translates to higher industry revenue is unclear.

Guide for incumbents:

Incumbents may not transform their entire organization, but can create spaces where "how would we build this from scratch" gets asked.

1. **Codify tacit knowledge**
   If you can't articulate it, AI can't do it

2. **Enable boundary-crossing**
   Find people who can reach into adjacent domains
3. **Think in terms of tasks, not roles**
   Jobs are bundles of tasks, and as AI can help with some, the bundle (and therefore the job) changes
4. **Try Read/Write AI**
   Move beyond Copy-and-Paste; let AI read and write to build context
5. **Run end-to-end experiments**
   Small teams going all-AI for a defined project instead of just a workflow within a larger project. For example, a strategic sample (or "vertical slice") of the final game.
6. **Don't forget the human**
   Setting strategy, aligning teams, and building culture require human coordination. The process matters as much as the output.

Guide for AI-first entrants:

Entrants may be dramatically more productive at making games, but when production is easy, if there is no IP or established customer base, how can a studio differentiate itself and find an audience? While we do not attempt to address this marketing question here, we do note its importance.

1. **Move fast**
   AI is a disruptive technology that is difficult for incumbents to adopt. Use your speed advantage and "discover the fun" much faster, and at lower cost, than the established studios
2. **Distribution is the challenge**
   AI has solved production but has not yet solved distribution, or created a new distribution platform. Focus as much on reaching and developing audiences as on building the game.
3. **Consider IP partnerships**
   If production becomes commoditized, marketing complements such as IP or distribution channels become more valuable. There may be a partnership to be struck with IP holders or platforms who want to innovate and move fast.
4. **Watch for new platforms**
   Existing distribution platforms, which aggregate gamers, are wary of AI-generated content at volume as it may overwhelm their curation ability and lower the quality of their service. New distribution platforms, with more sophisticated, algorithmic distribution mechanisms, may be able to absorb and sift through much higher quantities of content submissions and so may have less problem accepting AI games.
5. **Invest in distribution innovation**
   Some studios are applying AI to distribution itself, not just production. This may be what unlocks a new audience eager for novel kinds of generative AI play experiences.

What the studio of the future might look like

Based on the patterns we observed, and 15 years' experience in the games industry, we can speculate what the AI-first studio of the future might look like, whether that emerges from a successful new entrant or an incumbent that’s re-invented itself around this new technology.

In the future, traditional pipelines are upended. A handful of designers focus on the concept, aesthetics, and direction. The production happens via a “software factory” model, where the game design document is the primary input to the AI, which then generates the title itself. Hand-crafted code or art is a rare occurrence, if it happens at all. Small teams of generalists organize around business outcomes (growth, retention, monetization) rather than technical domains. Production velocity increases by an order of magnitude, not just in early prototyping, where teams can "discover the fun" faster, but by leveraging generative AI through finaling and live operations as well.

If the studio has an existing audience, this new AI-first content appears as new game modes, levels, events, and other live operations that keep their players engaged with novel experiences in a context they are familiar with. The trend we observe with “black hole” games (massive franchise titles that have such a hold on players that it’s nearly impossible for a new game to pull them away (6)) becomes more pronounced as it’s easier than ever to keep them fresh at low cost.

If the studio needs to build a new audience, it finds novel channel partners who have IP, brands, or other paths to AI-friendly consumers outside of traditional distribution channels. The storefronts that dominate game publishing today (Apple’s App Store, itch.io, Valve’s Steam) are not designed to absorb a massive quantity of AI-generated content. They may also have an audience base who prefers more “traditionally” produced titles anyway, where no AI assistance can be perceived. The studio of the future understands that content discovery has moved to social media and finds ways to align their content with streamers, influencers, and other creators who are looking for new ways to keep their audience entertained.

The market has the potential to become larger and more competitive, with AI lowering the technical barriers to entry, enabling more studios to enter the field and significantly increasing the production of new titles. Smaller teams will be able to take bigger creative risks and create new genres, just as YouTube and TikTok have enabled a vast expansion of who makes videos, what they are used for, and how audiences find them. The studio of the future may also find ways to use generative AI to create new *game mechanics*, ways to play that were not possible with older forms of AI (such as matchmaking algorithms to pair players by skill, procedural generation to cheaply create vast worlds, etc.)

Previous attempts at infinite procedural worlds showed that infinite content does not guarantee infinite fun. Whether generative AI changes that equation, or hits the same ceiling, is perhaps the most important open question in the industry.

What we need to figure out

Our research surfaced more questions than it answered. These are the ones that matter most.

While we saw glimpses of what the future AI-first studio might look like, much remains uncertain. The following questions point toward areas for future research.

- **Which technical gaps matter most?** AI has made rapid progress in many technical domains, but others have proven more stubborn. In gaming, for example, 3D art remains a challenge because these assets must meet specific technical requirements to function correctly, not just look convincing. Might there be other technical challenges, important to other industries, that generative AI cannot overcome?

- **Can artists be empowered by AI?** Artists have always faced a challenging labor market, but with the arrival of generative AI, that challenge has become an existential threat. In gaming, this tension is particularly visible, as art is often the most beloved part of a game, and gamers feel a strong emotional connection with the characters and their human creators. Perhaps this is why AI-generated output feels so threatening. Past technical innovations changed the shape of the creative workforce, but also brought new ways of making art. Digital painting and 3D modeling ultimately empowered artists rather than replacing them. Will generative AI follow the same pattern, and how can we ensure that it furthers human creativity rather than displacing it?

  "It's just like photos. Kodak doesn't exist anymore because they didn't believe in digital pictures. AI is coming, and it's coming hard, but it's coming to help us. I see artists doing a lot of code now, focused on performance. Before, you'd need a developer on the team just to ensure the art assets fit. But since we can use AI for coding, we can do it ourselves." — Art Manager/Producer, startup studio

  "We have zero intention of replacing humans. AI is supposed to be a copilot. Think about what mundane, repetitive tasks you could get rid of, so that you can use your time being creative." — VP of Engineering, mobile studio

  "When it comes to evaluating creative, be really careful. [The AI] can just tell you if it looks like something successful, they can't actually tell you if it's any good." — CPO, early-stage indie studio

- **Will AI shrink teams, or grow them?** The assumption that AI will reduce headcount may be premature. If AI makes each person dramatically more productive, studios may choose to scale output rather than cut staff. Cheaper content pipelines let studios entertain established audiences with more content and invest more in other parts of the game. This leans into the industry trend of "black hole" titles, where gamers remain loyal because the game keeps updating itself through live operations: 72% of time spent on PC, PlayStation, and Xbox worldwide goes to titles that are 6+ years old (2). The same could play out in other industries where increased productivity is a change to grow and take share, not trim costs.

  "We have not seen a retraction in personnel as a result of any AI initiatives. If anything, a lot of the AI work is being encouraged because it's getting difficult to find people." — Senior Technical PM, AAA studio

  "I could very easily see a future in which absolutely nobody loses their job. [A major franchise] has a thousand people making ten maps a game, and they want to be able to make a hundred. That scaling curve — everyone can still have their job. Now their job might be different." — Senior Technical PM, AAA studio

- **Is the generalist model permanent?** This paper has highlighted how AI rewards generalists — individuals who can reach across domain boundaries and wear multiple hats. But historically, as industries scale, they see returns to specialization. There may be a transitional period where generalists thrive because flexibility matters most, but ultimately, as companies complete their transformation, the returns to deep expertise will return. Will the small, generalist teams we

observed in AI-first studios remain the model, or are they a feature of this early, fluid moment? How might this dynamic play out across other industries?

- **How will audiences respond?** Most players may be neutral toward AI-assisted production, but a vocal, passionate minority is negative. High-profile protests have made some studios avoid AI altogether or keep usage secret (3, 4, 5). What drives these emotions? Are there enough such players to support an "artisanal gaming" niche? Might cultural differences support the growth of AI gaming in some markets, such as Asia, while limiting it in others, such as Europe? Will other kinds of content, such as music, journalism, movies, or books, see the same sort of audience rejection or will the response be more measured?

  "PC players are spoiled... game companies are scared of having a [****]storm. They can just start something on Reddit and it can explode and it can be very expensive." — CTO, mid-size studio

- **What happens when everyone can make games?** This paper has focused on how AI transforms production. But if AI makes it dramatically easier to create content, what impact does that have on distribution? How does all this new content find its audience?

  Previous technology innovations answered this by opening new markets: Facebook brought casual players, mobile brought light console experiences to developing markets, and Roblox brought young people to the Metaverse. Generative AI has not done this (yet). The games it produces compete on existing platforms for existing players — platforms built on the assumption that content is scarce and expensive.

  "Unity is 20 years old, Roblox 19, Minecraft 18, Epic Games 35. They were not set up in the age of the internet and AI — they were set up with 'content is hard to create' in mind." — CEO, startup studio

  Now that this assumption is breaking down, these platforms face a curation problem: how do you surface quality when submissions increase by an order of magnitude? Whether a new distribution channel emerges — one designed for abundance rather than scarcity — may determine whether generative AI expands the gaming market or simply makes it more competitive. The same question will arise in any industry where AI dramatically lowers the cost of production without a corresponding innovation in how that output reaches a market.

Conclusion

The programmer mentioned at the start of this article had access to all the same technology as the indie studio CEO. The difference was in how that technology had been rigorously applied, without the assumptions that govern traditional game development, to create something very new. To understand what this may mean at your studio, or in your industry, challenge and empower a small team of generalists to make a new product, using AI at each step, and see what they create. What tacit knowledge gets documented? Which workflows get collapsed? What roles emerge when five people have to do the job of fifty? The results will go beyond faster pipelines and reveal what is possible within your organization.

References

1. Michael Polanyi, The Tacit Dimension (Chicago: University of Chicago Press, 1966).
2. Matthew Ball, "The State of Video Gaming in 2026," Epyllion, February 16, 2026.
3. Luke Plunkett, "I Don't Think Larian's CEO Understands How Art Is Made," Aftermath, December 16, 2025, https://aftermath.site/larian-ceo-ai-art-baldurs-gate-3-divinity/.
4. Jamie Hore, "New Postal Spinoff Axed Immediately as Publisher Is 'Overwhelmed' by Gen-AI Backlash," PCGamesN, December 5, 2025, https://www.pcgamesn.com/postal-bullet-paradise/canceled-ai-use.
5. Elie Gould, "Nexon's CEO Defends the Use of Generative AI in Game Development," PC Gamer, November 11, 2025, https://www.pcgamer.com/games/third-person-shooter/every-game-company-is-now-using-ai-nexons-ceo-defends-the-use-of-generative-ai-in-game-development-in-the-wake-of-arc-raiders-controversy/.
6. Matthew Ball, "The State of Video Gaming in 2025," Epyllion, 2025.

About This Research

Zimran Ahmed, who led the work, is a Senior Fellow at Wharton Generative AI Labs and gaming industry veteran (link). He writes about technology on his Substack.

We conducted 20 structured interviews (60-90 minutes each) across AAA, mid-size, and indie game studios in the US and EU. Roles included product managers, CTOs, VPs of Engineering, CFOs, game designers, art managers, technical artists, and developers. Interviews were recorded, transcribed, and analyzed for themes. We did not include novel interactive experiences, such as real-money casinos, or Consumer AI/Assistance apps, such as ChatGPT or virtual boyfriend/girlfriend experiences.

We want to thank the following companies for their support: Simcluster, UXIF Studios, Pet Fighters Studio, Neural Nonsense, and Nilo Technologies.