

The Orchestration Layer's Ballmer Problem

What Higgsfield's Unlimited
Access Really Bought


July 2026

AUTHOR:



Emma Rowand

Research Analyst



About Activant

Activant is a research-led global investment firm that partners with high-growth companies. Since 2015, we have invested in category-defining businesses during their most critical phases of growth, partnering with founders who have won their initial battles and are ready for the next challenge.

Our approach pairs deep, proprietary research with patient, flexible capital and hands-on operational partnership. We work alongside founders and leadership teams to refine strategy, strengthen operations, and accelerate sustainable growth.

Activant Research is dedicated to uncovering the most exciting emerging technologies, sectors, and companies we believe will shape the future. Our research-driven perspective informs everything we do, helping us invest at meaningful inflection points and support founders in building enduring, category-leading businesses.

You can find out more about Activant and our research at <https://activantcapital.com/>.

Last week, [Higgsfield](#), the generative media platform for creating cinematic videos and images, gave away 24 hours of unlimited access to all its top models for free to new users!¹ Given the brutal compute economics of video generation, handing out unlimited free access looks close to reckless. And it came just a week after Higgsfield put \$100,000 on the table for a contest to get people building apps on the platform.² Two expensive giveaways in the space of a week, both pointed at the same goal: get people to build.

In [Is Video AI Having Its GPT-3 Moment?](#), we argued that video models diverge on creative character rather than converge on capability, so no single model wins and durable value concentrates at the orchestration layer.³ If that is right, the race is not to own the best model but to own the place where people actually create. Higgsfield's free trial, its contest, and a wave of moves like them are that thesis in motion. They are also its first real stress test.

Developers, Developers, Developers

Twenty-six years ago, Microsoft CEO, Steve Ballmer bounced across a stage chanting one word until he was hoarse: developers.⁴ He was defending the most valuable structural fact the company owned. A developer who wrote to the Win32 API produced software that ran only on Windows, meaning every application made Windows more useful and every user who came for it found leaving harder than before. The developer was the moat precisely because the developer's output was **non-portable and bound to the platform**. That is the innovation-platform flywheel that later built the App Store and Google's suite, and the transaction-platform version that built Uber and Airbnb: *complements you do not build, network effects you do not buy*.



In this new era of development, lock-in is much lower. Software can be built anywhere, which means the thing a developer builds no longer binds them to the platform they built it on. Higgsfield aggregates more than thirty models under a single subscription starting at \$15 per month, and now ships its own Model Context Protocol (MCP) server, built on the open, portable, cross-host standard that Anthropic authored and OpenAI later adopted.^{5,6,7} Win32 was designed to make a developer's work stick to one platform. MCP is designed to let it slide off.

The chant was about lock-in, and highly technical developers were the only mechanism for it. Higgsfield's [Supercomputer](#) removes the requirement, turning a plain-language brief into a finished app, website, or game, so the builder needs no technical background at all. Now anyone can build anywhere, so the lock-in no longer lives in the code, which means that in this era, a durable advantage has to be built a different way.

The Playbook: Wrapper to Model

If you build an AI product on top of someone else's model and own no model yourself, you're essentially a wrapper, and your value lives in the workflow and the user experience. **That's a head start, not a moat**, because the model will eventually saturate the workflow. [Cursor](#) was the best place for AI-assisted coding right up until Claude got good enough that simply being a nice coding tool stopped mattering. But the workflow was never the endgame. It was how Cursor built its **data advantage**. Every diff a developer accepted or rejected, every comment, and every passing test became training data that fed [Composer](#), a model now sitting on the frontier of cost to performance.⁸ The data, not the menu of models, is what carried Cursor to roughly \$4 billion in annualized revenue and, last month, a \$60 billion all-stock deal to be acquired by SpaceX, which wanted the workflow, the users, and above all the coding data to train its own models.^{9,10}

[ElevenLabs](#) ran the same play in reverse. It started with the model and built the platform on top, launching in 2022 with a text-to-speech engine whose voice quality was the entire product.¹¹ But a standalone model commoditizes fast, so the company climbed the stack: dubbing, music, [ElevenAgents](#) for conversational voice, and [ElevenCreative](#), a studio product that orchestrates third-party video models. What keeps users on the platform is the voice they cloned, the Voice Library where creators list and monetize their work, and the hosted agents built on top. That positioning carried it past \$600 million in ARR by July 2026.¹²

[Jasper](#) learned the lesson too late. It wrapped GPT-3 in tuned prompts and marketing templates and rode that head start to a \$1.5 billion valuation, until ChatGPT handed the same capability to everyone for free and its revenue slid from roughly \$120 million toward \$55 million inside a year.¹³ It has since pivoted from consumer product to enterprise, building its moat in [Jasper IQ](#), proprietary models for brand voice and style.

A playbook emerges from all three. The orchestration lead is borrowed time, and the point of holding it is to convert it: **first into data, then into a model, then into the**

user-made assets that make leaving costly. Higgsfield's recent moves place it squarely in the workflow stage, pushing people to build because that usage is the data advantage it spends on models of its own: a proprietary reasoning engine that plans narrative, camera logic, and visual consistency, fed by an in-house lineup that includes [Soul 2.0](#), [Higgsfield DoP](#), [Higgsfield Popcorn](#) and [Cinema Studio](#).

But raw model quality commoditizes as rivals catch up, so the pressure is on orchestrators to build what the models cannot reach. What lasts is a character cloned in [Soul ID](#), a color signature in [Soul HEX](#) and months of consistent frames. Owning a model was never about having the best model, it was about pulling users deep enough into the product that they build something worth staying for.

The Path Forward

In an age of open protocols, platform lock-in is far harder than it used to be. When MCP lets a customer swap one model for another without leaving the window, the old playbook of trapping users behind a proprietary integration stops working. That has made lock-in more expensive, which is why the leading labs are racing to embed themselves in enterprise clients, sending forward-deployed teams onsite for months to wire custom code into core workflows, train models on proprietary data, and rebuild the org around the system.¹⁴

We expect Higgsfield and others to keep pushing users to **build more and build more often**. Every generation, every accepted frame, every edit becomes training fuel for the in-house models that turn a workflow into something a rival cannot clone. Higgsfield is already walking this path, orchestrating partner models through its own reasoning engine while shipping its own models off the back of what its users taught it.

Ballmer was right that platforms live and die on what gets built on top of them but the developer was only ever a proxy for the real prize: the accumulated, non-portable work that made leaving painful. The free trial is cheap if what users build becomes both the reason they stay and the fuel that makes staying worthwhile.

-
- ¹ [Higgsfield AI X account](#), July 21st, 2026
 - ² Higgsfield AI, [Supercomputer Apps](#), 2026
 - ³ Activant Capital, [Is Video AI Having Its GPT-3 Moment?](#), 2026
 - ⁴ YouTube, [Steve Ballmer: Developers](#), 2008
 - ⁵ Higgsfield, [Pricing](#), 2026
 - ⁶ Higgsfield, [Higgsfield MCP for any AI](#), 2026
 - ⁷ Anthropic, [Donating the Model Context Protocol and establishing the Agentic AI Foundation](#), 2025
 - ⁸ MindStudio, [What Is Cursor Composer 2? The AI Coding Model Built for Cost-Efficient Sub-Agent Work](#), 2026
 - ⁹ Forbes, [Cursor Hits \\$4 Billion Annualized Revenue Ahead Of SpaceX IPO](#), 2026
 - ¹⁰ TechCrunch, [SpaceX to acquire Cursor for \\$60B in stock, days after blockbuster IPO](#), 2026
 - ¹¹ ElevenLabs, [What is ElevenLabs?](#), 2026
 - ¹² LinkedIn, [Mati Staniszewski LinkedIn post](#), 2026
 - ¹³ Medium, [The Wrapper Economy Is Collapsing](#), 2026
 - ¹⁴ Activant Research, [Forward Deployed Lock-In](#), 2026

Disclaimer: The information contained herein is provided for informational purposes only and should not be construed as investment advice. The opinions, views, forecasts, performance, estimates, etc. expressed herein are subject to change without notice. Certain statements contained herein reflect the subjective views and opinions of Activant. Past performance is not indicative of future results. No representation is made that any investment will or is likely to achieve its objectives. All investments involve risk and may result in loss. This newsletter does not constitute an offer to sell or a solicitation of an offer to buy any security. Activant does not provide tax or legal advice and you are encouraged to seek the advice of a tax or legal professional regarding your individual circumstances.

This content may not under any circumstances be relied upon when making a decision to invest in any fund or investment, including those managed by Activant. Certain information contained in here has been obtained from third-party sources, including from portfolio companies of funds managed by Activant. While taken from sources believed to be reliable, Activant has not independently verified such information and makes no representations about the current or enduring accuracy of the information or its appropriateness for a given situation.

Activant does not solicit or make its services available to the public. The content provided herein may include information regarding past and/or present portfolio companies or investments managed by Activant, its affiliates and/or personnel. References to specific companies are for illustrative purposes only and do not necessarily reflect Activant investments. It should not be assumed that investments made in the future will have similar characteristics. Please see "full list of investments" at activantcapital.com/companies/ for a full list of investments. Any portfolio companies discussed herein should not be assumed to have been profitable. Certain information herein constitutes "forward-looking statements." All forward-looking statements represent only the intent and belief of Activant as of the date such statements were made. None of Activant or any of its affiliates (i) assumes any responsibility for the accuracy and completeness of any forward-looking statements or (ii) undertakes any obligation to disseminate any updates or revisions to any forward-looking statement contained herein to reflect any change in their expectation with regard thereto or any change in events, conditions or circumstances on which any such statement is based. Due to various risks and uncertainties, actual events or results may differ materially from those reflected or contemplated in such forward-looking statements.