PLATFORM RISK • LEGAL TECHNOLOGY • STRATEGIC ANALYSIS

The Playbook for Surviving Platform Risk

. . . Has Been Written Three Times.

MIKE DENKER April 2026

Most legal tech founders haven't read it. On February 3, 2026, a single plugin release erased hundreds of billions in market capitalization and triggered a reckoning that the industry had been told, repeatedly, was coming. The history of technology offers a clear survival guide. The question is whether anyone in legal tech is paying attention.

1

The Claude Crash

On February 3, 2026, a plugin dropped and hundreds of billions of dollars in market capitalization evaporated.

Anthropic, the company behind the AI model Claude, released a legal automation plugin for its agentic desktop application, Claude Cowork. The plugin handles contract review, NDA triage, compliance tracking, and templated legal responses. It is open source, configurable to an organization's own playbook and risk tolerances, and ships on GitHub alongside ten other industry-specific plugins for sales, finance, and HR.[1]

Within hours, Thomson Reuters fell as much as 18 percent intraday. RELX, the parent company of LexisNexis, dropped 14 percent—its steepest single-day decline since 1988. Wolters Kluwer lost 13 percent in Amsterdam trading. DocuSign fell roughly 10 percent.[2] The sell-off spread beyond legal tech into Microsoft, Salesforce, Adobe, and Shopify as investors absorbed the implications: if a foundation model company can ship a domain-specific workflow tool as a free plugin, what exactly is the enterprise software business model?

Jefferies Group coined a term for it: the "SaaSpocalypse."[3]

The legal tech press scrambled to make sense of the carnage. *Artificial Lawyer* called it a "Claude Crash."[4] *LawSites'* Bob Ambrogi observed that the foundation models were "no longer just 'plumbing' underlying legal AI products but rather potential competitors."[5] *Above the Law's* Joe Patrice built on Ambrogi's framing, noting that when an industry is built on what Ambrogi calls a "model + wrapper + workflow" model and the model creator cuts out the middle, "it creates a crisis."[6]

Industry analysts debated whether the sell-off was rational or hysterical. The consensus, such as it was, landed somewhere in the middle: the plugin itself was relatively basic, but it signaled a strategic shift that had existential implications for a certain class of legal technology company.

The interesting thing about this panic is that none of it was new. The script had been performed before. Twice. And both times, the playbook for who survived and who didn't was exactly the same.

This paper tells the story of all three performances.

2

A Brief Theory of Platform Risk

Before we walk through the history, it's worth naming the pattern we're looking at.

In every technology ecosystem, value chains form. A manufacturer makes a thing. Intermediaries add something to it—configuration, customization, integration, a user interface, domain expertise—and deliver it to the end customer. The intermediary exists because the manufacturer either can't or won't serve the last mile. Maybe the manufacturer is a horizontal platform company that serves everyone and therefore can't specialize. Maybe the end customer's needs are too complex or too niche for the manufacturer to address efficiently. Either way, the intermediary fills a gap.

Platform risk is what happens when the manufacturer decides to close that gap.

Sometimes it's a distribution play: the manufacturer goes direct to the customer, cutting the intermediary out of the supply chain. Sometimes it's an absorption play: the manufacturer builds the intermediary's product into its own platform as a native feature. Sometimes it's both at once. But the underlying dynamic is always the same: the upstream provider expands its scope, and the downstream intermediary gets compressed.

The response from the intermediary is predictable too. First, denial: "they'll never do that." Then, indignation: "how dare they compete with us." Then, rationalization: "our customers are loyal." And finally, for the ones who survive, reinvention: "we need to be building something they have no strategic reason to become."

That last phrase is the key. Not something the upstream provider can't build. Large technology companies can build almost anything. The relevant question isn't capability. It's strategic logic. Does the upstream provider have an incentive to become what you are? Or would doing so require them to fundamentally change their identity, their business model, and their relationship with the thousands of other customers who depend on them remaining a horizontal platform?

The survivors in every generation built something that lived in a different category altogether. Not a thicker wrapper. Not a better feature. A different thing—something the upstream provider had no strategic reason to pursue because pursuing it would conflict with their core business model.

This is a pattern that has played out across three distinct technological eras. Each time, the surface details change. The underlying dynamics don't.

3

Generation 1: When the Manufacturer Went Direct

The Rise of the VAR

To understand the Value-Added Reseller era, you have to understand what buying a computer looked like in the 1980s. It was nothing like walking into an Apple Store. The hardware arrived as components. The software

needed to be installed, configured, and often customized. The network had to be set up. The staff had to be trained. And the manufacturer—Compaq, IBM, HP—had neither the capacity nor the inclination to do any of that for you.

Enter the VAR. Value-Added Resellers purchased hardware from manufacturers, bundled it with software and services, configured it for specific use cases, and delivered a turnkey solution to the end customer.[7] The "value added" was real: a business buying a Compaq server from a VAR wasn't just getting a box. They were getting a configured system, installed software, network integration, training, and ongoing support. By the early 1990s, VARs accounted for roughly 70 percent of the customer base of major distributors like Tech Data.[8] They were the connective tissue of the entire PC industry.

And they thought they were untouchable. As one analyst recalled of the era: "The channel saw themselves as the gods; they were the ones who made a success of the HPs and the IBMs. If a channel partner told those vendors they were not going to carry their products, HP and IBM would go out of their way to court them and keep the channel happy."[9]

Then Michael Dell came along and, as Longbottom put it, told the channel, "We don't need you."

Dell Goes Direct

Dell Computer Corporation, founded in 1984 in a University of Texas dorm room, had a simple insight: if you sell directly to the customer, you eliminate the distributor margin, the reseller margin, and the inventory risk that comes with forecasting demand through a multi-layered supply chain. You build to order. You ship direct. You pocket the savings or pass them to the customer as lower prices.

By the early 1990s, Dell had briefly experimented with retail channels—selling through Best Buy, Costco, and Sam's Club. In 1993, a Bain consultant named Kevin Rollins convinced Michael Dell to pull out of retail entirely. The margins were thin, the channel created complexity, and the direct model was working. Dell left the reseller channel in 1994 and never looked back.[10] Rollins would eventually become the company's CEO.

The impact on VARs was immediate and brutal. Dell didn't just offer lower prices. It offered comparable or better products, configured to order, delivered in days, with no middleman markup. One industry executive described the dynamic: "Customers were getting deals and pricing from Dell and then that reseller would have to take it back to the likes of IBM and say 'you have to try to match this price.' It was difficult to get the likes of IBM to answer that quickly, and in several scenarios Dell was offering products at a price the channel could not match."[11]

Dell had commoditized the box. And the VARs whose entire value proposition was configuring and delivering that box found themselves competing against a company that did it better, faster, and cheaper—without them.

Who Survived and Why

Not every VAR died. The ones who survived had one thing in common: they were building something Dell had no strategic reason to provide.

Dell was a horizontal hardware company. Its business model scaled by selling boxes to everyone. A VAR that specialized in configuring PCs for small businesses was dead—Dell could do that through a website and a call center. But a VAR that had built a proprietary vertical application for, say, managing a dental office's patient records and insurance billing? Dell had no reason to become a dental practice management company. It wasn't a capability gap. Dell had more engineers than most VARs had employees. It was a strategic gap. Becoming a dental practice management vendor would have required Dell to stop being Dell.

The surviving VARs moved up the stack. They became managed service providers, vertical solution vendors, and systems integrators. They stopped adding value to the box and started building value that was independent of the box. The box became a commodity input into something larger that the manufacturer had no incentive to replicate.

The ones who kept configuring boxes? They're gone.

4

Generation 2: When the Platform Absorbed the Feature

Watson, Meet Sherlock

In 2001, a small software company called Karelia Software released an application called Watson. It was a $30 utility for Mac OS X that extended the functionality of Apple's built-in Sherlock search tool. Where Sherlock searched local files, Watson added web services: weather, stock quotes, phone number lookup, package tracking. Watson was clever, well-designed, and popular. It won an Apple Design Award for Most Innovative Product. It built a loyal user base. The name itself was a knowing wink—Dr. Watson as Sherlock Holmes's faithful companion.

In 2002, Apple released Mac OS X 10.2 with Sherlock 3. It included virtually every feature Watson offered. Built in. Free.[12]

Watson's developer, Dan Wood, wrote a blog post about the experience. He recounted that Steve Jobs himself had called to acknowledge that Watson had inspired the new Sherlock. Jobs's framing, as Wood described it, was not exactly sympathetic. "You know those handcars, the little machines that people stand on and pump to move along on the train tracks? That's Karelia. Apple is the steam train that owns the tracks."[13]

Karelia went out of business. The term "Sherlocked" entered the technology lexicon and never left.

The Pattern Repeats

Watson was the first and most famous case, but it was hardly the last. In 2005, Apple released Mac OS X Tiger with Dashboard, a widget system that replicated the functionality of Konfabulator, a popular roughly $25 JavaScript widget application. Yahoo later acquired Konfabulator, but the market was gone. In 2012, Apple shipped a native notification system in Mountain Lion, and Growl—an open-source notification framework that hundreds of Mac apps relied on—became unnecessary overnight. In 2022, Apple launched Continuity Camera, allowing iPhones to serve as Mac webcams—eliminating the need for Camo, an app that had built an entire business around that exact functionality.[14]

The list keeps growing. In 2024 alone, Apple announced features that replicated the functionality of TapeACall (call recording), Grammarly (writing assistance), 1Password (password management), and Otter (voice transcription).[15] By 2025, the annual WWDC "Sherlocked" roundup had become its own genre of tech journalism.

The mechanism was different from Dell's. Dell eliminated the middleman through distribution efficiency. Apple eliminated the middleman through platform absorption—building the intermediary's product directly into the operating system. The intermediary didn't lose on price or speed. It lost on relevance. Why would a user pay $30 for a feature that ships free with the next OS update?

The Browser War: Platform Risk at Scale

The same pattern played out at much larger scale in the browser wars of the mid-1990s. Netscape Navigator held over 80 percent of the web browser market by 1995.[16] It was the gateway to the internet for millions of users. It was, briefly, the most widely used software in the history of computing.

Microsoft saw Netscape as an existential threat—not because the browser itself was dangerous, but because if Netscape controlled the browser, it could potentially become a platform that made Windows less relevant. Microsoft's response was total war. It built Internet Explorer, gave it away for free, and bundled it with every copy of Windows 95. It struck exclusive deals with PC manufacturers to pre-install IE. When Compaq attempted to ship Netscape Navigator as the default browser, Microsoft threatened to revoke Compaq's Windows license.[17]

Intel's vice president testified that a senior Microsoft executive had outlined a strategy to "cut off Netscape's air supply."[18] Microsoft spent over $100 million per year on IE development, employing over 1,000 people on a product it was giving away. The math was simple: Windows revenue dwarfed anything Netscape could ever earn from browser sales.

By 2001, Internet Explorer held over 90 percent of the market. Netscape was dead. The Department of Justice found that Microsoft had violated antitrust laws, but the ruling came years after the damage was done.[19] By the time the legal system caught up, the market had already moved on.

The Survivors of Generation 2

The companies that survived platform absorption shared a characteristic: they built something the platform had no structural way to replicate by simply adding a feature.

Cross-platform compatibility was one escape route. Apple could Sherlock any Mac-only utility, but it couldn't Sherlock a tool that worked across Mac, Windows, and Linux. The platform's absorption power ended at the boundary of its own ecosystem. This is part of why 1Password has survived despite Apple building a native password manager—1Password works everywhere, and Apple's Passwords app only works on Apple devices.

Deep vertical workflows were another. Apple could build a notification system, but it wasn't going to build an enterprise IT management platform. It could build a basic writing tool, but it wasn't going to build a legal document automation system with clause libraries, approval workflows, and compliance tracking.

Network effects were the strongest defense of all. A feature that gets better because multiple parties use it is inherently resistant to platform absorption, because the platform would need to replicate not just the software but the network of participants. Apple could build a messaging app, but it couldn't absorb Slack, because Slack's value lived in the network of teams and integrations that had formed on it.

The browser war and the Sherlocking pattern are two expressions of the same underlying dynamic: the platform owner leveraging its position to eliminate an intermediary. In Sherlocking, the weapon is feature absorption. In the browser war, the weapon was distribution control. Both are platform risk. And both will reappear, combined, in Generation 3.

The core question from Generation 1 applied again in Generation 2, with one refinement. In the VAR era, the diagnostic was: "Does the manufacturer have a strategic reason to serve your customer directly?" In the platform era, the diagnostic became: "Is what you've built a feature of the platform or a product that exists independent of it?"

Venture capitalists have been asking that question for two decades. They just didn't realize it was about to become the defining question in legal technology.

"The platform would need to replicate not just the software, but the network of participants. That's not a feature addition. That's a pivot."

5

Generation 3: When the Model Became the Product

The Legal Tech Wrapper Landscape

To understand why the Claude Crash hit so hard, you have to understand what the legal technology market looks like in 2026.

Over the past three years, a wave of legal tech companies has emerged offering some variation of the same core proposition: give us your contract playbook—your positions on indemnification caps, confidentiality terms, liability limitations, governing law—and our AI will review incoming contracts against that playbook, flag departures, and generate redlines to conform the agreement to your standards.

It sounds like a significant innovation. And in fairness, the user experience these companies have built around the underlying AI capability is often polished and well-designed. But the core function—read a contract, compare it to a set of rules, suggest changes—is a prompt chain. It's a workflow wrapper around a general-purpose large language model. The AI doing the actual analysis is Claude, or GPT, or Gemini. The legal tech company's contribution is the UI, the playbook structure, the Microsoft Word integration, and (in some cases) pre-built playbook content drafted by attorneys.

The market has stratified into what might be called a wrapper spectrum.

At the thin end are companies whose entire product is essentially "upload a contract, get it checked against your playbook, receive redlines." The AI does the work. The wrapper provides the interface. This is the Dell-box-configuration of legal tech: useful, but replicable by anyone who controls the underlying model.

In the middle are companies that have layered on meaningful additions: pre-built attorney-drafted playbooks, integration with contract lifecycle management systems, clause libraries, and benchmarking against market standards. These additions create switching costs and immediate out-of-the-box value that a raw AI model doesn't provide. LegalOn, for instance, markets "Day 1 readiness" through pre-built playbooks that require no custom setup.[20] That's a real differentiator—until the foundation model ships with its own pre-built playbooks, which is a matter of when, not whether.

At the thick end are the legacy information providers—Thomson Reuters, LexisNexis, Wolters Kluwer—who sit on decades of proprietary legal data: curated case law, contract databases, searchable regulatory archives. Their moat isn't the AI. It's the data the AI needs to be genuinely useful for complex legal work—and courts are actively deciding how high that wall stands. In February 2025, a federal district court ruled that using Westlaw's proprietary headnotes to train a competing AI was copyright infringement, not fair use, rejecting the first major AI training fair-use defense to come before a U.S. court.[31] The case is now fully briefed before the Third Circuit on interlocutory appeal, with ten amicus briefs filed on the challenger's side and a decision pending—making it the most-watched AI copyright case in the country. These companies are also not waiting for the courts to settle the question. CoCounsel is now integrated into Westlaw. LexisNexis struck a strategic partnership with Harvey to bring Harvey's AI platform inside its

proprietary data environment.[32] The strategy is deliberate: make the data and the AI inseparable, and raise the switching cost for anyone who wants both.

The threat to the thick end is real but different in kind, and operates on a longer timeline. A foundation model plugin can't replicate 150 years of editorial curation, Key Number taxonomy, or Shepard's treatment flags. But as AI reasoning improves and general-purpose models get better at navigating public-domain sources, the perceived value of the editorial layer will erode incrementally for routine work. The fortress is not impregnable—it is simply not vulnerable to the same weapon that destroyed the thin-wrapper companies. A foundation model can review a contract against a playbook, but it can't tell you what's market for a particular clause type in a particular industry without access to proprietary data that doesn't exist on the open internet. That distinction is still the difference between a feature and a fortress. For now.

The Plugin That Ate the Wrapper

On February 3, 2026, Anthropic released a legal plugin for Claude Cowork that was, by most accounts, a fairly basic implementation. It could review documents, flag risks, triage NDAs, track compliance, generate legal briefings, and produce templated responses. It was configurable to an organization's playbook and risk tolerances. It shipped on GitHub as open source, alongside plugins for ten other business functions.[21]

Anthropic explicitly noted that the tool assists with legal workflows but does not provide legal advice, and that outputs should be reviewed by licensed attorneys.[22] It was not, by any measure, a sophisticated legal technology platform.

But that was precisely the point.

The plugin didn't need to be sophisticated to be devastating. It just needed to be good enough to make a certain class of legal tech product look unnecessary. If the foundation model company can ship a "good enough" version of your core feature as a free, open-source plugin —one of eleven plugins it released on the same day, almost as an afterthought—then the value of your wrapper just collapsed.

This is the Dell-goes-direct moment and the Sherlock moment happening simultaneously. The foundation model provider is both going direct to the end user (bypassing the legal tech intermediary) and absorbing the core feature into its own platform (making the intermediary's product redundant). It's the worst of both Generation 1 and Generation 2 at once.

The Immediate Fallout

The market's reaction was severe but not indiscriminate. Thomson Reuters fell as much as 18 percent intraday in a single session. RELX experienced its worst single-day decline since 1988, falling 14 percent. Wolters Kluwer dropped 13 percent. The London Stock Exchange Group fell more than 8 percent amid broader fears about AI's impact on data companies.[23]

Industry analysts were divided on whether this was rational. *Artificial Lawyer* argued it was irrational, pointing out that Thomson Reuters, LexisNexis, and Wolters Kluwer are "legal data fortresses" with decades of curated, proprietary content that no plugin can replicate.[24] *Above the Law's* Joe Patrice offered a more colorful assessment, noting that from his experience with Claude Cowork, "it screws up most workflows that involve more than a couple tasks."[25]

But the debate over whether the stock reaction was proportionate is almost beside the point. The signal was clear: the foundation model provider had crossed a line. Instead of quietly powering legal tech companies through its API, it had shipped its own competing product directly to end users.[26] And this is where the historical pattern becomes useful. Because the question of who's actually at risk and who isn't has been answered before. Twice.

The Double Game

There's an additional dimension to the current generation that didn't exist in the same way during the VAR or platform eras, and it's worth naming explicitly.

Anthropic is simultaneously selling its AI capabilities to legal tech companies and competing with those same companies through its own plugins. Harvey, one of the most prominent legal AI startups (valued at $11 billion as of March 2026, with over $1 billion in total funding), is built on Claude.[27] Harvey is also a launch partner on Anthropic's new Marketplace.[28] Anthropic wants Harvey to succeed as a customer and partner while also shipping a free legal plugin that competes with parts of Harvey's offering.

Artificial Lawyer described this dynamic with a football (soccer) analogy: "It's a bit like hiring a great goal scorer for the football team who insists on scoring at least one own-goal in every match, but who everyone still agrees is a great player."[29]

This tension was sharpened further on March 17, 2026, when Anthropic launched its Claude Partner Network, a program designed to help enterprises adopt Claude directly, backed by $100 million in funding for implementation partners including Deloitte and Accenture.[30] *Artificial Lawyer* noted the inherent contradiction: Anthropic is encouraging legal tech vendors to build on its models while simultaneously recruiting consulting firms to sell Claude directly to those vendors' customers.

This tension isn't new. It's a direct echo of the channel conflicts Dell faced when it began selling directly, or Apple's complicated relationship with the developers whose features it absorbed. But in the AI era, the dependency is even more acute: these legal tech companies don't just sell through the platform; they're built on it. Their product literally cannot function without the foundation model. That's a level of platform dependency that makes the VAR's relationship with Compaq look casual.

6

The Survivor's Playbook

Three generations. Three different technologies. The same pattern.

In every generation, the upstream provider expanded into territory previously occupied by intermediaries. In every generation, the intermediaries who survived shared the same characteristics. And in every generation, the ones who died shared the same blind spot: they believed the upstream provider "wouldn't" come for them, when the real question was whether the upstream provider had a strategic reason to.

Here is what the survivors had in common.

They Operated in a Different Category

The surviving VARs didn't build a better-configured box. They built vertical solutions that required domain expertise the manufacturer had no incentive to develop. The surviving ISVs didn't build a better Mac feature. They built cross-platform tools and enterprise workflows the platform owner had no incentive to create. The distinction is subtle but critical: it's not about being "better" than the upstream provider. It's about being different in a way the upstream provider has no strategic reason to replicate.

Applying this to legal tech: a company that reviews a contract against a playbook is operating in the same category as a foundation model with a legal plugin. It's doing the same thing, just with a nicer interface. A company that changes the fundamental interaction model of how contracts get negotiated—that's a different category. The foundation model provider has no strategic reason to become a multi-party negotiation platform. It's not in their lane. Not because they lack the engineering

talent (they obviously don't), but because becoming that would require them to stop being a horizontal platform and start being a vertical workflow company. That's not a feature addition. That's a pivot.

Their Value Lived Between Parties, Not Within One

This is perhaps the most underappreciated distinction. Every product that has been Sherlocked, Dell'd, or wrapper-collapsed shares a common trait: it optimized a single party's experience. Watson made one user's Mac search better. Netscape gave one user a browser. A contract review tool helps one party assess one contract.

The products that resist platform absorption tend to create value in the interaction between parties. Slack wasn't a better notification tool; it was a network. Salesforce wasn't a better database; it was a system of relationships between sales teams, customers, and workflows. These products are resistant to absorption not because they're technically complex, but because their value is emergent from participation. The platform would need to replicate not just the software, but the network of participants using it.

In legal technology, the implications are direct. A tool that helps one lawyer review one contract is doing something a large language model can do natively. A platform that creates a shared space where two parties negotiate against transparent standards—where the value comes from both parties being present on the same platform—is doing something structurally different. The upstream provider would have to build and manage a multi-party marketplace, which is about as far from "ship a plugin on GitHub" as you can get.

They Owned the Workflow, Not Just the Intelligence

Foundation models are extraordinarily good at analysis, generation, and reasoning. What they're not good at—what they're not designed to be—is a system of record. They don't manage state across time. They don't coordinate handoffs between human actors. They don't enforce approval chains or audit trails. They don't manage the messy, multi-step, multi-party workflows that characterize real business processes.

The legal tech companies that are most exposed are the ones that use AI for intelligence but haven't built the surrounding workflow infrastructure that makes that intelligence operationally useful. Review a contract? The model can do that. But manage the lifecycle of a thousand contracts across fifty counterparties, with version control, approval workflows, compliance tracking, and institutional memory of how each counterparty negotiates? That's a system, not a feature.

The surviving VARs understood this instinctively. They stopped selling boxes (intelligence) and started selling managed services (workflow). The surviving ISVs did the same: they stopped building features (intelligence) and started building platforms (workflow). The legal tech companies that will survive the wrapper reckoning will be the ones that own the workflow, not just the AI that powers part of it.

They Were the Steam Train, Not the Handcar

Steve Jobs's metaphor to Dan Wood was cruel but instructive. Karelia was a handcar. Apple was the steam train that owned the tracks.

The lesson isn't "don't be a handcar." Sometimes handcars are the right business for the moment—they're fast, capital-efficient, and can cover a lot of ground while the steam train isn't looking. The lesson is: don't be a handcar on somebody else's tracks. If your business runs on rails the upstream provider owns, you're one product announcement away from irrelevance.

The surviving companies in every generation built their own tracks. They created value that existed independent of the platform, that couldn't be absorbed

by a feature update, and that required a strategic identity change—not just an engineering effort—for the upstream provider to replicate.

7

Five Questions for the Current Moment

Whether you're a legal tech founder evaluating your own position, a VC assessing a portfolio company, or a GC deciding which tools to invest in, the historical pattern suggests five questions that separate the survivors from the casualties.

1. *Is what you're building a natural extension of the foundation model's trajectory?*

   Contract review against a playbook is a prompt chain. NDA triage is classification. Compliance flagging is pattern matching. These are things large language models do natively and will only get better at. If your core feature is something the foundation model can do with a well-crafted prompt and a bit of orchestration, you're in the kill zone. The Cowork legal plugin should be Exhibit A. It's not a sophisticated product. It doesn't need to be. It's good enough to make the thin-wrapper products look like a surcharge on capability that's becoming free.

2. *Does your product require the foundation model provider to change what it is in order to replicate you?*

   This is the strategic identity test. Not "could they build it?" but "would building it require them to stop being a horizontal platform and become a vertical product company?" Anthropic's business model scales by being the foundation for thousands of applications across every industry. The moment it starts building deep vertical products for legal, it's competing with its own customers in finance, healthcare, and everywhere else. The incentive structure doesn't support it. If your product would require the upstream provider to undergo a strategic identity change to replicate, you're probably safe. If it would only require them to ship another plugin, you're probably not.

3. *Does your product create value between parties or for a single party?*

   Single-party optimization—review my contract, draft my brief, summarize my document—is squarely in the foundation model's wheelhouse. Multi-party value creation—platforms where the interaction between participants generates the value—is structurally resistant to absorption. A foundation model can't ship a multi-party marketplace as a plugin. It can ship a contract reviewer. Know which one you are.

4. *Do you own the workflow, or just the intelligence?*

   If the AI component of your product were available for free tomorrow—which, after February 3, is not a hypothetical—would your customers still need you? If the answer is yes, because you manage the state, the handoffs, the approvals, the audit trail, and the institutional memory of their contracting process, then you're selling a system. If the answer is no, because the AI analysis was the whole point, then you're selling a feature. Features get absorbed. Systems get adopted.

5. *Are you on your own tracks or someone else's?*

   This is the hardest question and the most important one. If your business exists entirely within the ecosystem of a single foundation model provider—if your product is built on their API, distributed through their platform, and differentiated primarily by a UI layer on top of their intelligence—you are a handcar on someone else's tracks. That doesn't mean you can't build a valuable business. Handcars covered a lot of ground in the VAR era, and plenty of Sherlocked apps had profitable runs before the platform caught up. But you should be clear-eyed about the structural risk. The steam train is coming. The only question is when.

8

The Script Doesn't Change. Only the Stage Does.

In the 1990s, a generation of VARs watched Dell go direct and wondered what happened to the business they thought they owned. Most of them are gone. The survivors built vertical solutions that Dell had no reason to become.

In the 2000s, a generation of software developers watched Apple absorb their features into the operating system and Microsoft bundle a free browser into Windows, and wondered what happened to the products they thought they'd built. Most of them are gone. The survivors built cross-platform tools and enterprise workflows that Apple had no reason to replicate.

In 2026, a generation of legal tech companies is watching foundation model providers ship legal workflow plugins and wondering what happens next. Some of them will be gone within two years. Some of them will thrive for decades. The difference will not be the quality of their AI. It will be whether they built a feature or a category.

The playbook for surviving platform risk has been written three times. The pattern is clear. The framework is tested. The diagnostic questions are the same ones that separated survivors from casualties in 1994 and 2002 and every year since.

The only question left is whether anyone in legal tech is reading it.

NOTES

[1] Anthropic, *Claude Cowork Plugins*, GitHub (February 3, 2026). The legal plugin was released alongside plugins for sales, finance, HR, marketing, customer support, data analysis, product management, biology research, productivity, and enterprise search.

[2] "Market Reaction or Overreaction? Anthropic's Legal Plugin and the Facts So Far," *ComplexDiscovery* (February 4, 2026). See also "Thomson Reuters, RELX, and Wolters Stocks Crushed After Anthropic Debuts Claude Legal Plug-In," *Morningstar* (February 4, 2026).

[3] "Anthropic Enters Legal Tech, Legal Tech Enters Freefall," *Above the Law* (February 4, 2026), citing the Jefferies Group's characterization.

[4] "Claude Crash Impact on Thomson Reuters + LexisNexis is Irrational," *Artificial Lawyer* (February 4, 2026).

[5] Robert Ambrogi, "Anthropic's Legal Plugin for Claude Cowork May Be the Opening Salvo in a Competition Between Foundation Models and Legal Tech Incumbents," *LawSites* (February 3, 2026).

[6] Joe Patrice, "Anthropic Enters Legal Tech, Legal Tech Enters Freefall," *Above the Law* (February 4, 2026).

[7] A value-added reseller (VAR) adds features or services to an existing product, then resells it as an integrated or complete "turnkey" solution. See Wikipedia, "Value-added reseller." See also "VAR Business Model: Challenges, Risks & How to Solve Them," *VARStreet* (January 2026).

[8] "Tech Data Corporation Company Profile," *Reference for Business*. By 1992, approximately 70 percent of Tech Data's 25,000 customers were value-added resellers.

[9] Clive Longbottom, founder of analyst firm Quocirca, quoted in "CRN Flashback: When Dell's Direct Model Rattled the Channel," *CRN/ChannelWeb*.

[10] Dell sold through Best Buy, Costco, and Sam's Club in the early 1990s but exited the reseller channel in 1994 after Bain consultant Kevin Rollins demonstrated the long-run margin disadvantage. See "Dell," Wikipedia; "A New Channel Strategy for Dell," Kellogg School of Management, Northwestern University.

[11] Steve Ellis, managing director of APSU, quoted in "CRN Flashback: When Dell's Direct Model Rattled the Channel," *CRN/ChannelWeb*.

[12] "What Does It Mean When an App Gets 'Sherlocked' by Apple?" *MakeUseOf* (September 2023). See also "Sherlocked: The Controversial Practice of Apple Copying Apps," *App Developer Magazine* (December 2024).

[13] Dan Wood of Karelia Software, quoted in "Everything Apple Sherlocked in macOS Ventura, iOS 16, and iPadOS 16," *AppleInsider* (June 2022).

[14] Sherlocking history compiled from *MakeUseOf*; *AppleInsider*; "Apple 'Sherlocking' Highlighted in Antitrust Probe," *iTechPost* (April 2021); and "Apple 'Sherlocked' These Apps at WWDC 2025," *MacRumors* (June 2025).

[15] "'Sherlocked': Apple Accused of Copying Apps' Services for New iPhone," *NPR* (June 2024). Affected apps included TapeACall, Grammarly, 1Password, Otter, Newji, and AllTrails.

[16] "Browser Wars," Wikipedia. Netscape Navigator reached approximately 90 percent browser market share with the release of version 3 in 1996.

[17] *United States v. Microsoft Corp.*, 253 F.3d 34 (D.C. Cir. 2001). See also Wikipedia, "United States v. Microsoft Corp."

[18] Intel Vice President Steven McGeady testified that Microsoft Senior VP Paul Maritz outlined a strategy to "cut off Netscape's air supply." See *United States v. Microsoft Corp.*, trial testimony.

[19] Judge Thomas Penfield Jackson issued findings of fact on November 5, 1999, and conclusions of law on April 3, 2000, holding that Microsoft had engaged in monopolization. The two parties reached a settlement in 2001.

[20] "Best AI Contract Review Tools 2025," *LegalOnTech.com*. See also "AI Contract Review Software: Complete 2025 Buyer's Guide," *LegalOn*.

[21] "Anthropic Unveils Claude Legal Plugin and Causes Market Meltdown," *Legal IT Insider* (February 3, 2026).

[22] "Thomson Reuters Shares Sink After Anthropic Unveils AI-Enhanced Legal Tool," *Investing.com* (February 3, 2026), quoting Anthropic's GitHub documentation.

[23] "Anthropic's Claude Legal Plugin: One Month On, the Market Fallout and What It Means for Legal Teams," *Legal.io* (March 2026). See also "Panic Rises in Legal Industry Due to Anthropic's AI Plugins," *AI Business* (February 6, 2026).

[24] "Claude Crash Impact on Thomson Reuters + LexisNexis is Irrational," *Artificial Lawyer* (February 4, 2026).

[25] Joe Patrice, "Anthropic Enters Legal Tech, Legal Tech Enters Freefall," *Above the Law* (February 4, 2026).

[26] "Anthropic's Legal Plugin for Claude Cowork May Be the Opening Salvo," *LawSites* (February 3, 2026).

[27] Harvey closed a $200 million round at an $11 billion valuation in March 2026, co-led by GIC and Sequoia, bringing total funding to over $1 billion. See "Legal AI Startup Harvey Raises $200 Million at $11 Billion Valuation," *CNBC* (March 25, 2026); "Harvey Confirms $11B Valuation: Sequoia Triples Down," *TechCrunch* (March 25, 2026).

[28] "Anthropic Launches Marketplace for Claude-Powered Software," *The Next Web* (March 2026). Harvey was named as a launch partner alongside Snowflake and Replit.

[29] "Will Anthropic's Claude Partner Network Impact Legal Tech?" *Artificial Lawyer* (March 17, 2026).

[30] *Ibid*. The Claude Partner Network includes implementation partners such as Deloitte and Accenture, with $100 million in funding. Anthropic's Head of Global Business Development stated: "Our partners are instrumental in getting enterprises from proof of concept to production with Claude."

[31] *Thomson Reuters Enterprise Centre GmbH v. Ross Intelligence, Inc.*, No. 1:20-CV-613-SB (D. Del. Feb. 11, 2025). The district court found infringement of over 2,000 Westlaw headnotes and rejected the fair-use defense as a matter of law—the first U.S. court to do so in an AI training context. The case is on interlocutory appeal to the Third Circuit, No. 25-2153, fully briefed as of January 2026, with oral argument not yet scheduled.

[32] "Harvey and LexisNexis Announce Strategic Alliance," *LexisNexis* (June 2025). The partnership integrates Harvey's AI platform with LexisNexis's Protégé, U.S. primary law, and Shepard's Citations.

ABOUT THE AUTHOR

Mike Denker is a legal technology executive and the founder and CEO of Transactency, the first-of-its-kind contract negotiation platform based upon shared playbook methodology. He has over 30 years of experience in legal practice, corporate law, and technology-enabled legal services.