


**Wall Street, Now Is
the Time to Mitigate
the Next Big Risk:**

Data Center Patents





Wall Street, Now Is the Time to Mitigate the Next Big Risk: Data Center Patents

This is the first in a series of insights and analyses published by Tech+IP Advisory highlighting the growing IP-related issues and risks to investors in the burgeoning AI Infrastructure market. To learn more about Tech+IP or to subscribe to this series, contact us at research@techip.ai

Tech+IP Advisory
October 2026



Introduction

Private capital has committed hundreds of billions of dollars to AI infrastructure (e.g., data centers, compute, memory, networking, power and cooling), underwritten on the expectation of stable, infrastructure-grade cash flows. A material risk, however, is maturing largely outside standard underwriting models: exposure to a massive estate of other companies and licensing entity patents.

In previous innovation cycles, the cost fell on cash-rich device makers and stayed an operating-margin issue; it rarely reached the capital structure. However, as we posit here, the current policy and legal environment, in the US and beyond, and characteristics of the AI infrastructure market itself, suggest this time is different and Wall Street must pay attention to patent risk associated with the AI infrastructure buildout.

Notably, history suggests the AI infrastructure market has reached the stage where where patent risk is beginning to surface: recent litigation activity mirrors the pattern that preceded the Smartphone Patent Wars of the 2010s, but the stakes are dramatically higher with AI.

The upshot is that market participants, including investors and capital providers, now have various strategies and resources available to proactively quantify and manage patent risk within their investments. The time to act is now.



This Time Is Different. It Will Be Worse.

Within the IP and legal communities, the possibility of significant patent disputes surrounding AI and digital infrastructure is no longer theoretical. Patent practitioners have begun openly discussing the prospect of an emerging generation of “AI patent wars,”^{[1][2][3]} while law firms are specifically warning data-center developers and operators about infringement exposure arising from the technologies embedded throughout modern facilities.^{[2][4]} This is different from NVIDIA or AMD chips that may infringe, as those companies are cash rich. Recent commentary has identified cooling, power management, networking, data processing, controls and other **infrastructure technologies** as potential sources of massive patent disputes. Enforcement activity has already begun.

Yet this emerging risk is largely absent from the parallel discussion taking place in the financial markets, which has opined on myriad risks surrounding the AI infrastructure buildout: the risk that expected enterprise AI demand fails to materialize,^[5] supply chain problems,^[6] technological obsolescence,^[7] power availability,^[8] and the enormous capital requirements necessary to build the next generation of computing infrastructure^[9] or the public resistance to data centers.

This is a blind spot for both the Data Centers and their Sponsors and Lenders.

There is a direct precedent. Every major technology wave has followed the same sequence: rapid commercialization and capital deployment first; profit making second, and systemic patent conflict a few years later. Smartphones are the defining example. Fifteen years ago, the smartphone combined wireless, semiconductor, display, software, video and memory technologies into a single product inside a market expanding faster than before. Tens of thousands of patents held by incumbents, former participants, universities, and licensing entities read on those devices. Roughly three to five years after mass commercialization came peak litigation: by 2012, Apple and Samsung alone were directly involved in more than 50 lawsuits worldwide,^[10] and the smartphone patent licensing industry emerged carrying a permanent royalty stack on device revenue still today. Video codecs followed the same arc, ending in patent pools that levy a per-unit charge on anyone shipping video to this day.

AI data centers exhibit the same two ingredients that made smartphones a battleground. They aggregate an even wider set of patented technologies, far beyond GPUs and CPUs: accelerators, memory, optical networking, power conversion, cooling, rack architecture, orchestration software, etc., whose ownership is fragmented across hundreds of companies that often no longer compete in the market - and may not be countersued easily. And they concentrate enormous pools of economic value around an identifiable group of manufacturers, operators and tenants. Furthermore, the cumulative stack of granted AI patents worldwide has grown from roughly 4,000 in 2015 to an estimated 107,000 in 2025,^[11] and if the smartphone analogy holds, the litigation phase should be beginning about now. Three lead indicators, discussed below, confirm the shift is already underway.

But there is one huge difference: the financial consequences of AI data center patent risk may be amplified by the massive lending environment in which it occurs.



Patent Risk Translates to *Real* Financial Risk

Three distinctions from the Smartphone Patent Wars make this time around much more concerning:

1. The knock-on effects of the “Royalty Tax”

The AI infrastructure buildout, separate from chip companies and frontier labs, is being financed with leverage: project finance, private credit, asset-backed structures and high-yield debt now sit alongside sponsor equity in most capital stacks. Financial leverage transforms patent royalties from a legal expense into a meaningful credit and valuation variable. A company with substantial margins and little debt can absorb an incremental licensing obligation; a leveraged, capital-intensive company has far less room: an incremental 3%, 5% or 7% claim on revenue does not merely reduce accounting earnings. It reduces EBITDA, raises leverage ratios, constrains further borrowing, and ultimately impairs equity value.

At Tech+IP, we refer to this economic burden as the “Royalty Tax”: the cumulative claim that third-party patent owners may ultimately impose on participants in the AI and digital infrastructure value chain. It need not arrive through a single blockbuster judgment. More likely, it emerges through a combination of litigation, settlements, bilateral licenses, patent-pool payments and portfolio-wide licensing programs, each individually manageable, but collectively capable of becoming a recurring claim on industry margins. This is no longer purely hypothetical: on 10 September 2026, Via Licensing formally launched its data-center DRAM pool ^[28], with license fees ranging from roughly \$0.005 to \$0.020 per server processing unit per hour, charged at the data-center level rather than on the traditional per-device basis.

Across the capital stack, the question is fundamentally the same, whether current assumptions adequately reflect the IP costs that may accompany industry maturation: consensus margins for equity investors; projected EBITDA and free cash flow for credit investors; diligence, underwriting and exit assumptions for private equity and infrastructure sponsors; and the cash flows supporting the debt for lenders.

Exhibit 1: What a Royalty Tax Does to a Leveraged Asset
Illustrative asset: 50% EBITDA margin, 6.0x net debt/EBITDA, 16x EV/EBITDA


Royalty on Revenue	3%	5%	7%
EBITDA reduction	-6%	-10%	-14%
Net leverage change	6.0x → 6.4x	6.0x → 6.7x	6.0x → 7.0x
Equity value	-10%	-16%	-22%

Source: Tech+IP Analysis. Assumes no operating offset and excludes litigation costs.

Cerebras, an early darling of the AI boom, illustrates how even a modest patent “Royalty Tax” could have meaningful financial consequences for AI infrastructure companies. With the company still investing heavily in semiconductor development, manufacturing capacity, and data-center infrastructure, a hypothetical 5% royalty on revenue would divert cash that could otherwise fund growth, roughly \$44 million annually on approximately \$885 million of projected 2026 core revenue,^[12] a burden that grows as revenues scale. The effect is magnified at the EBITDA line: though Cerebras is EBITDA-negative today, if its forecasted margins of roughly 50% materialize, a 5% royalty on revenue would equate to a roughly 10% claim on EBITDA. Cerebras is a new entrant among large, well-patented incumbents with a comparatively thin patent position (compare Cerebras’ patent portfolio of 23 patent families to Broadcom’s roughly 14,000 - and they are quite aggressive in asserting).^[13]

Chart 1: Active Family Patent Count

Selected AI compute peers vs. Cerebras



Source: Tech+IP Research
Note: Not to scale in order to show different orders of magnitude.

2. Injunctions are back in play

Historically, the risk of an injunction prohibiting the sale, use, or manufacture of a product or service has been an ultimate point of leverage for a patent holder. During the Smartphone Patent Wars era that leverage was largely unavailable in the US. After the Supreme Court’s 2006 decision in eBay v. MercExchange, courts required stricter conditions for injunctions. Thus, monetary damages became the primary remedy; an outcome that largely favored the infringers.

Today, the policy environment in the US, and globally, is shifting in favor of patent holders, and injunctive relief is increasingly available. Europe’s Unified Patent Court, which didn’t exist three years ago, can issue injunctions effective across the largest European markets, and has granted preliminary injunctions in roughly 58%^[15] of the cases. The US International Trade Commission provides similar leverage: under

Section 337, its principal remedy is not damages but exclusion of infringing products at the US border. And in China, where cessation of infringement is the default remedy, courts grant injunctions in over 90% of proven cases.^[16] For an industry dependent on globally sourced hardware, the threat of interrupted supply can transform negotiating leverage and stall a buildout.


The return of injunctive relief matters most when the plaintiff doesn't want money. Patent assertion in AI infrastructure comes from two directions. The familiar one is the licensing entity, whose objective is royalty payments; disruptive, but ultimately results in a negotiation over price. The more consequential one, for investors, is the operating competitor, for whom litigation is an instrument of market competition — and for whom these remedies are not bargaining chips but the objective itself. When the plaintiff competes with the defendant, the calculus changes in three ways:

- a. The injunction is the objective, not the bargaining chip.** A court order that halts a rival's sales or operations protects the plaintiff's market share directly; no cash settlement may fully substitute for it.
- b. Capital attrition is a strategy.** A well-funded incumbent can outlast a leveraged new entrant. Prolonged litigation drains cash, distracts management, and creates precisely the uncertainty that pressures covenants, refinancing, and exit timing.
- c. Settlements carry structural terms.** Resolution with a competitor rarely ends at money: cross-licenses, supply arrangements, or field-of-use and geographic limits can cap margins and enterprise value long after the dispute closes.

Wall Street has already seen this play out. In June 2026, China's Supreme People's Court upheld an injunction barring Infineon from selling accused GaN power products in mainland China,^[17] a suit brought not by a licensing entity but by an operating rival. The damages were immaterial; the market impact wasn't: within few weeks, Infineon's shares had fallen more than 20% (as shown in the stock-price chart below) as analysts stripped China GaN growth from their models.^{[18][19]} The ongoing disputes among Navitas, Infineon, Samsung, Micron and YMTC follow the same logic: well-funded campaigns determine which players are permitted to build, operate, and profit from the physical layers of the AI economy.^[20]

Chart 2: Infineon Stock Price Following the China Injunction

Xetra: IFX daily close; 20 May - 10 September 2026



Source: Yahoo Finance, IFX.DE unadjusted Xetra daily closes. Legal dates: Reported Suzhou and SPC rulings.


The dynamic is not limited to mega-caps. In August 2026, Shoals Technologies, a power-infrastructure maker, won a \$96M willfulness verdict^[21] against competitor Voltage LLC with an immediate injunction and a final International Trade Commission (ITC) exclusion order barring imports. A mid-sized operating company, asserting its own patents, secured remedies that bar its rival's product from the US market.

3. Patent risk is spreading up and down the AI technology stack.

Perhaps the most consequential shift is in how infringement is claimed. Assertions increasingly target not individual products, but the way power, cooling, networking, and orchestration software are combined and operated together, at scale. Because the operator defines those configurations, exposure attaches at the integration layer, which no single vendor controls, and no single vendor is likely to indemnify.^[22]

Three practical consequences follow:

- a. **Every participant in the chain carries risk:** Component manufacturers, Original Equipment Manufacturers (OEMs), contractors, developers, operators, and tenants alike. There is no safe seat.
- b. **Vendor indemnities are structurally misaligned with the threat:** They typically cover a product as delivered, not as integrated with other suppliers' systems under the operator's direction, which is exactly the theory now being pleaded. Strict notice provisions and the difficulty of coordinating multiple vendors (often competitors) in a unified defense further erode coverage.
- c. **The facility itself anchors venue:** An asset located in the Eastern or Western District of Texas can be sued there, regardless of where its owner sits.




Three Lead Indicators Show the Shift Is Already Underway

The transition from innovation to patent monetization follows a familiar pattern. Patent filings migrate toward the technologies attracting the largest capital investment; ownership begins to consolidate through secondary transactions and licensing vehicles, and litigation expands from component suppliers to commercial implementers. AI infrastructure now exhibits each of these indicators.

Indicator 1: Patent filings are following capital

Chart 3: AI Patent Grants



Source: IFI Insights: Inventing AI

The filing mix has visibly shifted from AI model architectures toward the physical layer: cooling, power delivery, memory, networking, and interconnect, the technologies every AI facility must deploy.^[23]

Indicator 2: Litigation is accelerating and targeting operators

The volume of patent litigation is rising sharply in technologies fundamental to AI infrastructure, while the profile of defendants is changing:


- **Hardware disputes are intensifying.** US district court cases involving DRAM more than doubled in 2025 (+119% year-over-year), while NAND-related cases grew by 56%.^[24]

- **Operators are increasingly being named.** Plaintiffs are no longer confining suits to chip and equipment manufacturers; they are moving directly to the asset-owner level. Valtrus, asserting a former Hewlett Packard Enterprise portfolio spanning cooling, power and facility management, has sued Equinix, Evoque, Google and, in 2026, a coordinated wave of colocation operators. ThroughPuter has pursued Microsoft, Google and Amazon across the US and the UK.^[4]
- **Weaponization of Courts.** Plaintiffs are exploiting a widening divergence in remedies, steering disputes toward the UPC and the ITC where, as discussed above, injunctions and exclusion orders remain readily available.

The acceleration is starkly visible in the filing curve: litigation activity involving AI data center technologies roughly quadrupled in 2024, and filings during the first half of 2026 are already pacing to exceed every prior full-year record.^[20]

Chart 4: AI Data Center Patent and IP Disputes

Cases across cooling and power, memory, compute, networking, and virtualization




Source: Tech+IP 2026

Indicator 3: The patent transaction market is heating up

Patents are often bought before they are asserted. Listings of AI and machine-learning patents on the brokered market have risen every year but one since 2015, as operating companies and licensing vehicles position portfolios for leverage and defense. Brokered-market activity has historically led litigation by twelve to twenty-four months.

Chart 5: AI and ML Patents on the Brokered Market



Source: 2025 Richardson Oliver Patent Market Report

At the same time, collective structures are emerging. Via Licensing has now formally launched its data-center DRAM pool, signing founding licensors and disclosing pricing—and Sisvel launched an Ultra Ethernet pool, both targeting hyperscalers, cloud providers and data-center operators, echoing the licensing models that later became standard in smartphones and video codecs.^{[26][27]}



Managing Patent Risk

Patent risk is generally less costly to manage before anyone asserts it. Before a suit, a company can quietly identify which patents matter, buy defensive positions, fix its contracts, and take licenses on its own terms. After a suit, it's negotiating on the plaintiff's timetable, paying litigation costs whatever the outcome and, where an injunction is available, doing so under the threat of interrupted operations.

Knowing where that exposure sits, and what it's likely to cost, requires patent-market, technology and financial expertise in one place. **That's what Tech+IP Advisory provides to sponsors, lenders and their portfolio companies. The action plan below shows how we turn that analysis into decisions.**

■ A Patent Risk Management Action Plan

Almost always, an appropriate starting point is a Patent Risk Assessment designed to identify and quantify a company's principal areas of existing and prospective patent exposure. Such an assessment should evaluate the company's products and technologies against relevant third-party patent estates, identify areas of heightened litigation or licensing activity, and assess the company's existing defensive patent position. The objective is not to eliminate patent risk, but to identify where it's concentrated, estimate its potential economic significance, and determine which exposures warrant mitigation.

Based on the findings of that assessment, companies, sponsors and other investors can consider a targeted set of risk-mitigation measures, including:

Defensive Portfolio Development. Strengthen patent positions in areas of identified exposure, particularly through targeted acquisitions of third-party patents that can provide defensive leverage, counterclaim capability, or negotiating currency.

Contractual Risk Allocation. Review and, where possible, strengthen patent indemnification provisions across supplier, customer and strategic-partner agreements, with particular attention to liability caps, exclusions, defense obligations and the creditworthiness of the indemnifying party.

Third-Party Risk Management Solutions. Evaluate participation in defensive aggregators, patent pools, cross-licensing arrangements, and other collective risk-management structures.

Targeted Licensing/Exposure Resolution. Where material exposure is identified, consider proactively securing licenses or otherwise resolving specific patent risks before they mature into litigation.

Ongoing Patent-Risk Monitoring. Establish monitoring around relevant patent owners, litigation, patent transactions and emerging assertion activity so that changes in the risk environment can be incorporated into investment and operating decisions.

For the investment community, the broader implication is straightforward: patent risk management should begin before litigation. A relatively modest investment in identifying, quantifying and mitigating patent exposure can be substantially more efficient than addressing that same exposure after an assertion has occurred, when strategic options have narrowed and the economic leverage has shifted to the patent holder.

References

- ^[1]Michael Fainberg, AI Data Center Patent Wars: Energy, Cooling, Compute, LinkedIn, 2026.
- ^[2]Lauren Castle, Data Center Growth Brings Wave of Patent, Trade Secret Disputes, Bloomberg Law, 20 August 2026.
- ^[3]Sanjana Mishra, 'Reading the Tea Leaves': Practitioners Predict Data Centre Patent Litigation Spike, Managing Intellectual Property, 12 June 2026.
- ^[4]Stephenson Harwood, Should Data Centres Prepare for Patent Litigation in the UK?, 15 April 2026.
- ^[5]Goldman Sachs, The AI Investment Boom: When Will It Pay Off?, 2 June 2026.
- ^[6]Pepijn Bourgonje, Data Center Construction Supply Chain Risks, Caliber, 9 June 2026.
- ^[7]Scope Ratings, Research Announcement Addressing Technological Obsolescence in Data-Center Investment, 5 September 2025.
- ^[8]J.P. Morgan, Financing AI Infrastructure and Data Centers, 10 August 2026.
- ^[9]Amelia Michael and Ben Cottier, AI Data Center Cost Breakdown, Epoch AI, 14 May 2026.
- ^[10]Charles Cooper, Apple v. Samsung: 50 Suits, 10 Countries and Counting, CNET, 28 April 2012.
- ^[11]IFI CLAIMS Patent Services, IFI Insights: Inventing AI, 28 July 2026.
- ^[12]Cerebras Systems Inc., Cerebras Systems Fast Inference Cloud Business Nearly Quadruples in Second Quarter 2026, Exhibit 99.1 filed with the U.S. Securities and Exchange Commission, 12 August 2026.
- ^[13]Lumenci, NVIDIA Patent Portfolio Analysis, updated 29 August 2025.
- ^[14]eBay Inc. v. MercExchange, L.L.C., 547 U.S. 388, 15 May 2006.
- ^[15]Simmons & Simmons, UPC Statistics: 60% Success Rate for Patentees, Rapid Proceedings and Further Key Insights, JUVE Patent, 10 February 2025.
- ^[16]Mark Cohen, Reforms to China's Patent Injunction Practice in a Domestic Comparative Perspective, China IPR, 29 January 2026.
- ^[17]Xia Yu, China's Top Court Affirms Injunction Barring Infineon from Selling GaN Products, MMLC Group, 15 June 2026; Semiconductor Today, China's Supreme Court Upholds Injunction Against Infineon, Says Innoscience, 15 June 2026.
- ^[18]Simply Wall St, Is Infineon (XTRA:IFX) Rethinking Its Power Semiconductor Strategy After China's GaN Sales Ban?, 23 June 2026.
- ^[19]Yahoo Finance, Infineon Technologies AG (IFX.DE): Historical Data, accessed 8 September 2026.
- ^[20]Tech+IP Advisory, AI Data Center Patent Litigation Tracker, internal database, n = 138, 22 July 2026.
- ^[21]U.S. International Trade Commission, Certain Photovoltaic Trunk Bus Cable Assemblies and Components Thereof, Investigation No. 337-TA-1438, Notice of Final Determination and Limited Exclusion Order, 25 June 2026; Troutman Pepper Locke, Troutman Pepper Locke Secures \$96.4 Million Jury Verdict and Injunction for Shoals Technologies Group in Landmark Patent Infringement Trial, 27 August 2026.
- ^[22]Michael Chibib, Conor M. Civins and Christopher L. Crumbley, Patent Risk in Data Centers: A Legal Q&A for Operators and Developers, National Law Review (Bracewell LLP), 27 May 2026; C. Erik Hawes, Manita Rawat and Jason E. Gettleman, A Practical Guide to IP Risk and Litigation in Data Center Development, Morgan Lewis, 8 June 2026.
- ^[23]Shannon Williams, AI Data Centre Cooling Patents Rise 50% amid Energy Concerns, Data Center News, 23 October 2025; reporting data from Appleyard Lees, Inside Green Innovation: Progress Report, 5th ed., October 2025.
- ^[24]Nisha Shetty, DRAM, NAND Patent Disputes Surge in US District Courts, IAM, 23 June 2026.
- ^[25]Richardson Oliver Insights, The 2025 Richardson Oliver Patent Market Report: The Brokered Patent Market in 2024, 2025.
- ^[26]Via Licensing Alliance, Via Licensing Alliance Announces NAND Memory Licensing Program Expansion and Issues Formal Call for Patents, 22 May 2026.
- ^[27]Sisvel, Sisvel Announces Call for Patents for New Ultra Ethernet Pool, 29 April 2026.
- ^[28]ip fray, Via Signs Founding Licensors to Data Center DRAM Program, Announces New Cloud Media Pool for genAI, 10 September 2026.