

The Unpriced Liability in AI Infrastructure: Patent Risk

A practical *risk mitigation note* for financial sponsors and lenders financing the AI buildout

The possibility of a 5% permanent “IP royalty tax” on AI Infrastructure companies’ revenues is real and underappreciated by the investment community. The time to act is now.

Private capital has committed hundreds of billions of dollars to AI infrastructure on the expectation of infrastructure-grade cash flows. One cost is missing from most underwriting models: **other people’s patents**.

Just this summer, China’s Supreme People’s Court upheld an injunction barring Infineon (ETR:IFX) from selling accused “GaN power” data center chips in mainland China. The plaintiff was an operating rival, not a licensing entity. Within weeks, Infineon’s shares had fallen more than 20% as analysts stripped China GaN growth from their models.

Every major technology wave has followed the same sequence: capital deployment first, profitability, systemic patent conflict three to five years later. In the early smartphone era, Apple and Samsung alone were in more than 50 direct lawsuits, and smartphones still carry a permanent royalty stack. AI data centers combine an even wider set of patented technologies, owned by parties that no longer compete in the market. Litigation involving AI data center technologies roughly quadrupled in 2024, and first-half 2026 filings are pacing to a record - and this is before anyone is making any real profits.

The litigation phase is beginning now, and AI data center patent wars will dwarf smartphone patent wars.

Why the Financial Consequences Are Larger This Time

1. Leverage turns royalties into a credit variable. In the smartphone era, royalties were an operating margin issue for cash-rich device makers. AI chip makers have cash too. But AI infrastructure is financed with project finance, private credit, asset-backed structures, and high-yield debt. For a leveraged asset, an incremental claim on revenue doesn’t merely reduce earnings: it reduces EBITDA, lifts leverage, tightens covenant headroom and impairs equity.

The Royalty Tax: DefinitionX

The cumulative claim that third-party patent owners may ultimately impose on participants in the AI infrastructure value chain. It rarely arrives as a single judgment. It accumulates through settlements, bilateral licenses and patent-pool payments, each individually manageable, collectively a recurring claim on margins.

What a Royalty Tax Does to a Leveraged Asset

Illustrative asset: \$100 million revenue, 50% EBITDA margin, 6.0x net debt/EBITDA, 16x EV/EBITDA

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

Source: Tech+IP analysis. Assumes no operating offset; excludes litigation costs. At these assumptions each 1% of revenue paid as royalty removes roughly 3.2% of equity value.

The “Royalty Tax” early pricing signals are here. Via Licensing’s “DRAM Memory Patent Pool”, launched in September 2026, charges the server operator \$0.020 per hour for each HBM-type processing unit, about 0.5% to 1% of compute revenue at indicative \$2 to \$3 per GPU-hour rental rates. That is one pool covering one layer of the stack. New NAND, Ultra Ethernet and server-side media pools are in formation. Individual patent owners (both component and system manufacturers and patent licensing entities) will also join the fray.

2. Injunctions and exclusion orders are back in play. For most of the smartphone era, US courts rarely ordered infringers to stop; they awarded only financial damages, an outcome that often favored the infringer. That has changed. Europe’s Unified Patent Court grants preliminary injunctions in roughly 58% of cases, the US International Trade Commission’s principal remedy is exclusion of imports at the border, and Chinese courts grant injunctions in over 90% of proven cases.^[1]

That matters most when the plaintiff doesn’t want money. A licensing entity sues for a royalty, so the dispute is ultimately a negotiation over price. An operating competitor sues to protect market share. For it, a court order halting the rival’s sales or operations is the goal, a long fight that drains a leveraged new entrant’s cash is a strategy, and any settlement comes with cross-licenses or field-of-use limits that cap the target’s value after the case closes. Infineon is the large-cap example. Shoals Technologies, a mid-sized power-infrastructure maker, is the other: in August 2026 it won a \$96 million verdict against a rival together with an injunction and an ITC exclusion order barring the rival’s product from the US market.

3. Exposure attaches to the operator, and vendor indemnities typically do not cover it. Assertions increasingly target not individual products but the way power, cooling, networking and orchestration software are combined and operated at scale. The operator defines that configuration, so exposure lands at the integration layer, which no single vendor controls or is likely to indemnify. Valtrus, asserting a former Hewlett Packard Enterprise portfolio, has sued Equinix, Google and, in 2026, a coordinated wave of colocation operators.

Via's DRAM pool likewise licenses the server operator, not the memory vendor, per hour of operation, a cost no product-as-delivered indemnity is drafted to absorb. The facility itself anchors venue: an asset in the Eastern or Western District of Texas can be sued there regardless of where its owner sits. Component makers, OEMs, developers, operators and tenants all carry risk: there is no safe seat.

4. Today's IP licensing market is far more sophisticated. In previous times, companies could deploy a tried-and-true playbook to delay or avoid the onslaught of attacks from dedicated licensing groups – sometimes referred to as Patent Assertion Entities (PAEs). Today's PAEs are different: they're seasoned, sophisticated and well-funded. Couple that with a regulatory environment less favorable to the playbook, and the potential for consequential exposure is significant.

Five questions to ask a target or portfolio Data Center or borrower this quarter

1. What is our own patent position, and could we counterclaim if a competitor sued?
2. Which competitors and licensing entities hold patents in our technology stack, and do we have a landscape of their portfolios?
3. Do our vendor indemnities cover the system as integrated under our direction, or only the product as delivered? What are the caps, exclusions and notice terms?
4. Have we been approached by Via, Sisvel or any other licensing entity, and who is tracking those requests?
5. Are any of our suppliers, or their competitors, currently in patent litigation?

What to Do Before Someone Asserts

Patent risk is far cheaper to manage before a suit, and the risk of injunctions can be much better addressed. Before assertion, a Data Center can identify the patents that matter, buy defensive positions, fix its contracts and take licenses on its own terms. After assertion, it is forced to negotiate on the plaintiff's timetable, pays litigation costs whatever the outcome and, where an injunction is available, does so under the threat of interrupted operations.

Tech+IP Advisory helps sponsors, lenders and their portfolio companies identify, quantify and mitigate patent exposure. Our **Patent Risk Assessment Tool** maps products and facility designs against the relevant third-party patent estates, identifies areas of heightened litigation and licensing activity, and estimates the potential economic burden, in view of estimated real-world risks. Based on the findings, Tech+IP executes the mitigation: acquiring highly targeted patent portfolios that read on competitors' products to provide counterclaim leverage and negotiating currency; strengthening indemnification terms across supplier and customer agreements; evaluating patent pools, defensive aggregators and cross-licenses; resolving specific exposures through targeted licenses; and monitoring patent owners, transactions and assertion activity on an ongoing basis. We've worked for the world's largest companies on defensive strategies, from Google, Amazon and Meta to Intel, IBM RedHat, and Roku.

Tech+IP Advisory, a carveout from Houlihan Lokey, is a boutique investment bank focused exclusively on intellectual property and deep-tech, with more than \$2 billion of patent transactions advised and a client base of private equity, private credit and infrastructure investors. This note summarizes the first article in Tech+IP's Insights series on IP risk in AI infrastructure; the full article and references can be found at techip.ai. Contact: research@techip.ai

^[1] Mark Cohen, Reforms to China's Patent Injunction Practice in a Domestic Comparative Perspective, China IPR, 29 January 2026.

Sources: Tech+IP AI Data Center Patent Litigation Tracker (2026); MMLC Group; Semiconductor Today; Yahoo Finance; USITC 337-TA-1438; Troutman Pepper Locke; JUVE Patent; China IPR; Bloomberg Law; Via Licensing Alliance; ip fray; CNET. GPU-hour rental rates are indicative market figures.