

THE
AMERICAN LAWYERLatham Is Buying Its Own AI Servers:
'Flexibility is Critical'

By Andrew Maloney and Christine Simmons

September 10, 2026

Big Law firms are partnering with the biggest names in AI to get early access to high-powered models and collaborate with them to develop and roll out tools for legal work. But Latham & Watkins, one of the highest-grossing law firms in the world, is going a step further, purchasing its own data servers and tailoring its own AI models, the firm confirmed on Thursday.

Latham has purchased several Nvidia GPU servers over the past few years, including recently Nemotron 3 open-weight models, a Latham spokesperson said Thursday, allowing the firm to further customize its AI infrastructure.

The firm said in a statement that the idea is to diversify its approach to the rapidly-ascendant technology.

"The future of AI will involve many different models, platforms, and deployment approaches, not a single solution," said Rene Mendoza, chief information officer at Latham, in a statement to Law.com. "That's why flexibility is critical. Our infrastructure strategy gives us greater flexibility in how we develop, test, and deploy AI capabilities, while ensuring we are not dependent on any single vendor or technology provider."



Photo: Diego M. Radzinski/Law.com

The servers, with multiple GPUs, are based in a data center that the firm leases space from, and only its employees have access to it, the Latham spokesperson said.

Like other firms, Latham is still working with other AI companies, including OpenAI, Harvey and Microsoft, and has licensed all the major frontier models: ChatGPT, Claude and Gemini.

But in buying several of its own servers over the past few years, each with multiple GPUs, Latham may very well be the first Big Law firm to get into the hardware game to this extent.

Latham, which has more than 3,700 lawyers, also has more than 900 technologists across

engineering, cybersecurity, information, governance and innovation, the firm spokesperson noted.

A Latham spokesperson declined to comment further, including how much it has invested and spent on the servers and related needs.

Big Law has announced a variety of approaches so far for AI needs. For instance, Kirkland & Ellis has partnered with Palantir and Sylo. Freshfields announced a partnership with Anthropic, and A&O Shearman announced a revenue-sharing agreement with Harvey.

Rudy DeFelice, a former Big Law lawyer and executive vice president in charge of AI strategy at Harbor Global, said he's not heard of other Big Law firms going the route Latham has, but noted other firms may eventually follow.

The administrative burdens and costs may be the largest disadvantage for law firms maintaining their own AI servers. A large firm may need to spend tens of millions of dollars to buy 40 or 50 servers for the models and data it needs for thousands of lawyers, DeFelice said, not to mention the cost to maintain the servers in a secure and climate-controlled facility.

"It's not a trivial amount of money," he said. "For a major firm, it could easily be an eight-figure investment," he said, speaking generally. He said he wasn't aware of Latham's needs and costs.

Still, DeFelice pointed to the several benefits for law firms in the long run for maintaining their own AI infrastructure.

Building and maintaining a firm's own architecture, he said, gives firms control of client data, especially when some clients have much higher confidentiality and security requirements. Some AI platforms hold onto customer data for training or to assess security risks. "When it's in-house, you have control over all of it," he said.

Secondly, by having data servers and architecture in house, firms have predictability over costs and control of their own economics, he said, noting the unpredictability of AI data expenses now.

And having its own servers allows a law firm to further customize their own data needs through open-weight models, he said. While Claude or ChatGPT offer their own algorithms in off-the-shelf AI platforms, a firm may not be able to adjust their own weighting for evaluating data, including giving weight to different factors in data.

"If you have your own infrastructure, you can adjust your own weighting," he said. "It allows you to affect how the algorithm will perform."

"It takes a lot of expertise to make those customizations; it could make the model perform a lot more like you teach your own lawyers," he added.