

INTERVIEW: Maase Pursues AI Infrastructure in New Business Initiative

The company plans to invest up to 5 billion yuan over the next five years on its Stars Distributed Intelligent Computing Center AI infrastructure project

By Doug Young

As a listed company, the company's history dates back to 2019 when it was engaged in financial services. But for **Maase Inc.** (Nasdaq: MAAS) ("Maase" or the "company"), history really begins in May 2025, when it took on its current name and embarked on a brief acquisition spree that forms its current core assets. Among its three acquisitions, Huazhi Future has quickly become Maase's primary focus. The other two came in the new energy and health and wellness sectors.

The Huazhi acquisition takes Maase directly into the hot area of AI infrastructure, which revolves around the data centers filled with servers, GPUs and other hardware needed to run complex AI applications. Following the acquisition, the company moved quickly to secure an additional \$50 million in fresh funding, and laid out a RMB 5 billion (\$747 million) roadmap for developing data centers running its own self-developed large model.

Spearheading the new AI infrastructure push is Li Zhifeng, Huazhi Future's previous chairman who was named as Maase's chief technology officer after the Huazhi acquisition. Li sat down with Bamboo Works to discuss Maase's AI infrastructure roadmap in more detail, including its development timeline and how it differs from other developers piling into the market.

Bamboo Works: Let's start with the new \$50 million investment. Can you tell us who was the investor, and what drew them to Maase?

Li Zhifeng: The investor in the PIPE financing is a financial investment institution with a focus on AI infrastructure. From their perspective, we sit at the intersection of two things institutional investors are watching closely right now — the global computing power shortage, and the growing premium being placed on AI safety and data sovereignty. Our pitch wasn't "We buy and resell GPU time." It was that we're building a closed loop — infrastructure, our own AI models and enterprise delivery — with signed, revenue-generating contracts already in place, not just a roadmap slide.

Q: Tell us a little bit about the Stars Distributed Intelligent Computing Center launched in April, which is the backbone of your AI infrastructure initiative. Where is it and what kind of services does it provide?

A: The Stars Project was announced by our subsidiary Huazhi Future, together with two partners — China Power Computing Technology Application (Beijing) Co., Ltd. and Sino-International Yuzhi (Chongqing) Co., Ltd. It's designed to support China's national "East-to-West" and "Xinjiang-to-Chongqing" computing-resource-transfer initiatives.

The architecture is "dual-core + multi-level edge + a unified dispatch platform" consisting of two centralized training-and-inference hubs —one in Yinchuan, Ningxia (with an initial phase already operational, targeting roughly 512 high-performance servers); and one in Xinjiang (partially operational, targeting roughly 256 high-performance servers) — plus a planned network of 50 to 100 containerized data center, each carrying 10 to 16 containerized compute modules, for low-latency inference closer to end users.

Q: What's the capacity for the center's first edge node, and what plans are there to scale it up over time? What are the challenges involved in the replication and expansion of modular edge nodes?

A: The first node was deployed in April in the town of Lengshui of Shizhu county in Chongqing, with a designed compute capacity of 4,000 petaflops. On scaling: the modular, containerized design is meant to make replication faster than traditional data-center construction — no long civil-works cycle, it's largely plug-and-play.

But real challenges remain: securing power-supply arrangements and local-government incentive agreements site-by-site, each of which currently requires its own negotiation; financing each additional node or cluster, since we haven't yet finalized definitive financing arrangements for the broader network; and the normal execution risks of any distributed infrastructure build — construction timing, equipment delivery, and achieving utilization once nodes go live.

Q: At the time of the Stars Center announcement in April, the company said the project would have eventual investment of RMB 5 billion. What's the time horizon for that investment, and where does the company expect to get that money?

A: The planned construction period is 60 months — five years — rolled out in five phases from pilot validation to large-scale implementation. So RMB 5 billion is the ceiling for the full build-out. As for funding: we expect financing for individual sites — starting with the Shizhu site — to be arranged at the project-company level, through a mix of equity investment and debt financing from external investors, potentially alongside local incentive arrangements like preferential land-use and tax policies that are still under discussion with the local government.

Q: Tell us about some of the customers for the Stars Center so far, including who they are, the size of their contracts, and what services are they using?

A: I want to draw a distinction here that matters: the signed commercial contracts we've announced to date sit within Huazhi Future's broader computing-power and algorithm-solutions business rather than being revenue specifically generated by the Stars Center's own infrastructure.

With that distinction clear, our recent disclosed wins include: a computing-services delivery to Guangzhou Benyun Artificial Intelligence Technology, valued at RMB 1.65 million, covering five AI computing nodes plus environment deployment, fully paid and accepted; and a 12-month, RMB 76.8 million agreement with Beijing VirtAI Technology for over 450-petaflop-class high-performance computing infrastructure.

These and several other deals collectively represent well over RMB 100 million in signed, disclosed contract value across the past few months, which we think demonstrates real, paying demand for our computing services — even as the flagship Stars infrastructure itself is still being built out.

Q: How does Maase plan to differentiate its AI services from all the many other companies out there that are setting up similar operations?

A: Most companies in this sector are competing on raw scale — more GPUs, bigger clusters, a general-purpose model with no real safeguards behind it. We think the differentiator that actually matters going forward is trust: can an enterprise or government customer put sensitive data and mission-critical workloads on your platform with confidence? Our approach is “safety by architecture,” not safety as an afterthought.

That starts with our own computing infrastructure layer, so we're not just renting capacity from a third party we can't control. It continues with Lingyanmiaoyu (referred to as 'Lingyan'), our large language model, which is designed from the ground up around secure, on-premises private deployment, data localization, and content risk-control, rather than being a public API we bolt safety filters onto later; and it extends to how we deliver — data governance, compliance support, and ongoing operations as part of the contract, not just a one-time model handoff.

Q: Tell us more about the Lingyan mixture-of-experts (MoE) large model that you're using. What's the future international roadmap for Lingyan?

A: Lingyan is our proprietary large language model, built on a mixture-of-experts architecture with a total parameter scale of RMB 9 billion. We chose MoE first for efficiency — dynamically activating selected expert networks on demand instead of the full parameter set on every request, which lowers inference cost and improves concurrency. But the reason that efficiency matters strategically is that it's what makes safety-hardened, privately-deployed AI commercially viable rather than a research showcase — enterprise and government customers that require on-premises, data-

sovereign deployment need a model that's both controllable and cost-effective enough to run that way at scale.

The Bamboo Works offers a wide-ranging mix of coverage on U.S.- and Hong Kong-listed Chinese companies, including some sponsored content. For additional queries, including questions on individual articles, please contact us by [clicking here](#).

To subscribe to Bamboo Works' free weekly newsletter, [click here](#).