AM Intelligence to Deploy One of Asia's First NVIDIA Vera Rubin Clusters, Bringing Frontier AI Compute to India

AMI has placed binding order for ~9,000 NVIDIA Rubin GPUs deployed as Vera Rubin NVL72 rack-scale systems in its first AI Factory in Hyderabad, part of its 1 GW compute capacity globally

Hyderabad, India, August 25, 2026: AM Intelligence ("AMI"), the AI infrastructure platform set up by Promoters of Greenko, is striving to become the World's leading, intelligence infrastructure company, powering the AI economy through vertically integrated energy, compute and innovation.

AMI today announced first step towards this vision by placing a firm and binding order for Vera Rubin delivery in Q1 2027 at its first AI factory in Hyderabad. The order covers 9,000 NVIDIA Rubin GPUs deployed as Vera Rubin NVL72 rack-scale systems for the 30 MW of capacity and will make AMI's Hyderabad AI Factory as one of the first frontier AI compute clusters in Asia.

The Hyderabad facility will be the first tranche of AMI's planned 1 GW of Compute-as-a-Service capacity across India, the United States, Finland and Malaysia. AMI plans to bring an initial 200 megawatts (MW) of capacity to market in the near term, involving capital expenditure of more than USD 8 billion.

Anil Chalamalasetty, Chairman, AM Intelligence, said: *"For over two decades, we have focused on transforming electrons into value — addressing critical and emerging societal needs. Today, the electron-to-token opportunity allows us to take this capability further, converting power infrastructure into frontier AI compute at scale. Bringing latest generation NVIDIA Vera Rubin to India marks the next step in this journey and establishes AMI at the forefront of the emerging AI electron-to-token economy."*

By integrating NVIDIA Vera Rubin architecture with high-throughput RDMA over Converged Ethernet (RoCE) networking for fast, efficient data movement, and advanced storage systems, the Hyderabad AI Factory is being designed to run some of the world's largest and most sophisticated AI models, including trillion-parameter models and next-generation agentic AI applications.

NVIDIA Vera Rubin NVL72 introduces NVFP4, a new low-precision computing format designed to run AI workloads more efficiently, alongside next-generation High Bandwidth Memory (HBM4). The architecture is designed to reduce AI inference token-serving costs by up to 10 times compared with the preceding NVIDIA Grace Blackwell generation. AMI's Hyderabad AI Factory is being engineered to deliver approximately 450 exaFLOPS of NVFP4 inference compute, providing massive computing capacity for running advanced AI models at scale.

AMI: Building Toward 5 Gigawatts

Over the last two decades, our focus has been on electron valorisation, creating progressively greater value from energy across energy management, molecules and metals, and now AI tokens, the units of information processed and generated by AI models. This end-to-end approach has been central to our philosophy of addressing critical and emerging societal needs.

Today, intelligence has reached the status of permeating humankind's lives on a daily basis. AMI, as a strategic builder of long-term heavy assets, will leverage its unique power solutions, capex advantage and access to next-generation NVIDIA AI infrastructure to deliver highly cost-competitive electron-to-token economics by integrating energy infrastructure with the computing capacity required to train and run AI models. This will enable wider deployment of AI with a collaborative service offering spanning across hyperscalers, neo cloud providers, sovereign AI initiatives, frontier labs, AI natives and local developer communities.

AMI is developing 5 GW of powered AI data centres across India, the United States and Europe. These facilities are being designed for high-density AI computing, using liquid cooling to manage the significant power and heat requirements of advanced AI chips, while allowing the infrastructure to accommodate successive generations of AI silicon. AMI also intends to bring 1 GW of Compute-as-a-Service capacity to global AI workloads with initial capacity of 200 MW coming online in short-term.

About AM Intelligence:

AM Intelligence (AMI) is building a full-stack AI infrastructure ecosystem spanning energy, data-center infrastructure, hardware, and customized AI models. Backed by Greenko's owned renewable generation and storage, AMI is developing gigawatt-scale, liquid-cooled compute capacity designed to serve global hyperscalers, frontier AI labs, enterprises, and India's sovereign AI initiatives — while widening chipset access for Indian developers building domestic and global AI solutions. AM Intelligence is developing 5 GW of powered AI data centres and 1 GW of Compute-as-a-Service across India, United States and Europe, with 200 MW of Compute-as-a-Service capacity coming to market in the near term.

AMI, promoted by the founders of Greenko, is building globally differentiated platforms across energy, molecules and AI with a focus on enabling next generation of intelligent infrastructure. On electrons, we are building a 50 GW energy infrastructure platform across 20+ states, spanning solar, wind, hydro and energy storage. We are also developing the world's first interconnected energy storage platform, targeting 200 GWh of storage capacity across India.

On molecules, we are developing low-carbon molecule platform, including green ammonia, with a target of 5 Mtpa of green ammonia capacity across multiple locations in India. Our 1 Mtpa green ammonia project currently under construction is expected to be among the world's largest RFNBO-compliant green ammonia facilities, supporting energy diversification for Global markets for multitude of use cases across power generation, mobility amongst others.

++++++

For any media queries, please contact:
AM Intelligence
Media@amintelligencegroup.com

Adfactors PR
Naveen Janapalli
naveen.janapalli@adfactorspr.com | **+91 9396981816**

Misbah Hussain
misbah.hussain@adfactorspr.com | **+91 9820539503**