E2E Networks Limited

CIN- L72900DL2009PLC341980

Uppal's Genesis A-32, Block B, Mohan Cooperative Industrial Estate,
Badarpur, New Delhi 110044, Phone No. +91-11-4084-4964

Email: cs@e2enetworks.com, **Website** https://www.e2enetworks.com/

Date: November 02, 2026

To, National Stock Exchange of India Ltd. ('NSE') Exchange Plaza, 5th Floor, Plot No. C/1, G Block, Bandra - Kurla Complex, Bandra (E), Mumbai – 400 051, Maharashtra, India **(Script Symbol: E2E)**	To, The Manager- Listing BSE Limited, Phiroze Jeejeebhoy Towers, Dalal Street, Mumbai- 400 001 Maharashtra, India **(Scrip Code: 544783)**

Sub: Intimation of Product Launch under Regulation 30 of SEBI (Listing Obligations and Disclosure Requirements) Regulations, 2015 ('SEBI Listing Regulations')

Dear Madam/Sir,

Pursuant to Regulation 30 read with Schedule III, Part A, Para B of the SEBI (Listing Obligations and Disclosure Requirements) Regulations, 2015, we wish to inform you that the Company has launched a new product, viz. "Blaze"

The details of the said Product as required in terms of SEBI Listing Regulations read with SEBI Master Circular No. HO/49/14/14 (7) 2025-CFD-POD2/I/3762/2026 dated January 30, 2026 are as follows:

S. No	Particulars	Details
1.	Name of the product	Blaze
2.	Date of Launch	02-11-2026
3.	Category of the product	A pay-per-token inference service for open AI models
4.	Whether caters to domestic/ international market;	Both
5.	Name of the countries in which the product is launched (in case of international).	Globally

A copy of the Press Release captioned as **"E2E Networks launches Blaze, bringing pay-per-token AI inference to India's Sovereign AI Factory"** is enclosed herewith.

This intimation is also being uploaded on Company's website and can be accessed at https://www.e2enetworks.com/

Kindly take this on record.

Thanking You,

Yours faithfully,

For E2E Networks Limited

RONIT Digitally signed by RONIT Date: 2026.11.02 16:39:16 +05'30'

Ronit
Company secretary & Compliance Officer
ICSI M. No.: A59215

E2E Networks launches Blaze, bringing pay-per-token AI inference to India's Sovereign AI Factory

Built on NVIDIA DSX OS, Blaze gives India's AI builders access to leading open models through a single API powered by NVIDIA GPU infrastructure in India

- A single OpenAI-compatible API for open models including NVIDIA Nemotron, Sarvam, DeepSeek, Qwen, GLM and Gemma
- Engineered for coding agents and long-context workloads, with custom NVIDIA CUDA kernels and NVIDIA Dynamo-based serving
- With TIR, gives labs and enterprises one India-hosted path from training and fine-tuning to production inference

New Delhi, India - November 2, 2026 - E2E Networks Ltd. (NSE: E2E; BSE: 544783) today officially launched Blaze, a pay-per-token inference service for open AI models. Blaze gives developers and enterprises a single OpenAI-compatible endpoint to leading open models. Powered by NVIDIA GPU infrastructure and open-source components from NVIDIA DSX OS, Blaze turns E2E's India-hosted AI factory capacity into production inference services, while E2E manages model hosting and scaling. Customers pay only for the tokens they use.

Indian AI companies are moving from pilots to production, making inference cost and latency critical to product viability as AI agents drive up token consumption. Data sovereignty and keeping compute within India are key priorities, additionally the rising AI-assisted development costs are adding pressure on CIOs to improve token efficiency.

"Blaze is where you can run your production inference," ***said Tarun Dua, Founder and Managing Director, E2E Networks****. "There is customer demand for frontier-class open weight models running on an NVIDIA Cloud Partner (NCP) with NVIDIA reference design clusters as Sovereign Compute. Blaze provides them this on E2E Cloud in India. Blaze provides easy integration with all E2E Cloud platform services"*

Built for how AI is used now

Blaze is designed for inference-heavy products where latency and cost per token matter most, including voice, coding agents, document AI and OCR, image and video generation, and language applications.

For latency-sensitive voice and chat applications, Blaze offers fast time to the first sentence, while open weight models benefit from high kv cache reuse to reduce inference cost.

Bringing AI factory capacity into production with NVIDIA DSX OS

Blaze integrates open-source components from NVIDIA DSX OS, the open, modular software layer within the NVIDIA DSX AI Factory Platform. NVIDIA Dynamo provides distributed model serving, with KV cache-aware routing, disaggregated prefill and decode, and fast data transfer between GPUs through NVIDIA NIXL. NVIDIA Grove and KAI Scheduler coordinate inference workloads across E2E's Kubernetes fleet. NVIDIA BlueField DPUs, NVIDIA DOCA and NVIDIA Infra Controller enforce hardware-level isolation between customers, which matters for banks, healthcare providers and public-sector bodies with strict data rules.

"Access to AI computing and inference services is essential to advancing India's AI ecosystem," said ***Vishal Dhupar, Managing Director for South Asia at NVIDIA****. "NVIDIA's accelerated computing platform provides a foundation for cloud providers like E2E Networks to make AI more accessible to India's developers and enterprises."*

One platform from training to tokens

Blaze and TIR, E2E's AI platform, provide a flexible path for customers to scale AI workloads. **Blaze** delivers API-based, token-efficient inference for open-weight models, with optimized performance and seamless scaling. **TIR** enables self-hosted inference and model pre-training and post-training on high-performance NVIDIA Blackwell GPU clusters, through managed Kubernetes, offering greater control over infrastructure and model customization. Customers can serve their own or fine-tuned models on TIR's managed Kubernetes with Integrated NVIDIA Dynamo. E2E's AI factory capacity is expanding to about 14.6 MW, including a 1,024-GPU NVIDIA Blackwell GPU cluster live since May 2026 and additional 5.4 MW of Blackwell capacity including HGX B300 systems and other Blackwell GPUs.

About E2E Networks Ltd.

E2E Networks Ltd. (NSE: E2E; BSE: 544783) is building India's sovereign AI Factory, providing AI infrastructure and cloud services to AI-native companies, startups, enterprises, higher education and research and public-sector programs.

Media contact

Gouri B | gouri.b@2020msl.com | +91 79028 04053

K Sai Sharanya | ksai.sharanya@2020msl.com | +91 80882 65334

Anika Khanna | anika.khanna@2020msl.com | +91 99728 35554

E2E Marketing: Ashish Dixit | marketing@e2enetworks.com | +91-11-40844965