Acryl Takes GPUBASE to North America as First-Half Revenue Rises 150%

KOSDAQ-listed AI infrastructure firm presents 1,272-GPU verification at UKC 2026, meets NVIDIA co-founder Malachowsky and reports □21.2 billion order backlog



Acryl management with NVIDIA co-founder Chris Malachowsky (center) at UKC 2026, Orlando, Aug. 7.

SEOUL, SEOUL, SOUTH KOREA, September 5, 2026 /EINPresswire.com/ -- SEOUL, South Korea — Acryl Inc. (KOSDAQ: 0007C0), an AI infrastructure software company, has begun its entry into the North American market with GPUBASE, its GPU operations platform, on the back of first-half revenue that rose 150% year on year to KRW 12 billion (US$8.6 million) and an order backlog of KRW 21.2 billion (US$15.3 million) at the end of June.

> Competition in AI infrastructure is moving from how many GPUs a company has secured to how efficiently and reliably it operates the compute it holds
>
> *Jin Park, CEO of Acryl Inc*

Acryl took part in the 39th US-Korea Conference on Science and Technology (UKC 2026), held August 5–8 in Orlando, Florida, as the only Korean company on the panel of the official Korea-US AI cooperation symposium and delivered the anchor presentation at the official sponsor forum. During the conference the company visited HiPerGator, the University of Florida's supercomputing facility, to discuss cooperation with its operations team, held a roundtable with SelectFlorida, the state's economic development organization, and on August 7 met NVIDIA co-founder Chris A. Malachowsky to introduce the company and GPUBASE.

A full-stack platform for heterogeneous GPU clusters

GPUBASE pools GPUs and NPUs from different manufacturers into a single resource. It partitions a single GPU so that multiple jobs can share it and reduces network bottlenecks, raising the

realized performance of equipment already purchased. The commercial product has three parts: Booster, which reduces data-transfer bottlenecks between GPUs; Control Tower, which manages GPU, server, storage and network resources as one pool; and Monitor Tower, which shows infrastructure and job status in real time. Its network technology is based on the research of Acryl CTO Yeom Ik-joon, a professor of software at Sungkyunkwan University, including PeRF, which distributes data-transfer paths, and UL-MPRDMA, which uses multiple paths at the same time.

Acryl's entry point in the North American market is heterogeneous accelerator environments rather than single-vendor ones, and Ethernet-based clusters rather than InfiniBand. Universities and public research institutions are the first target; references secured there will be used to extend into data center and cloud operators.



Acryl executives at the University of Florida East Campus Data Center, HiPerGator racks, Aug. 5



Acryl executives in conversation with co-founder Malachowsky

1,272 GPUs verified: utilization above 90%, training time cut by up to 96%

The evidence Acryl presented in North America is the K-Scale Evaluation, a performance verification the company began in June in Amazon Web Services and other global cloud environments and whose first stage, the Horizontal Phase, was completed in July. In tests run independently on 1,272 GPUs of seven heterogeneous types across the three major global cloud environments, each with a different network architecture, GPU utilization rose above 90%, training time under high load fell by up to 96% — a speed-up of up to 24 times — and job queue wait time fell by up to 93%. NCCL all-reduce communication bandwidth improved by between 141% and 2,375% depending on load level, and fault detection and recovery completed within five minutes. The next step is a performance test at single-cluster scale of 1,000 GPUs, with a KOLAS-accredited test report to follow.

Cost effects are presented as the company's own estimates. On a basis compiled from published

papers and commercial cases, network total cost of ownership on a 512-GPU NVIDIA H100 cluster comes to 48% of the InfiniBand figure, a saving of about US$2.2 million. Raising average GPU utilization from 55% to around 85% would cut the GPU purchase cost needed for the same throughput by about 35%.

US$4 million cumulative supply to Hancom's modular data center, KRW 6.7 billion Ethernet fabric project

Commercial supply continues in Korea. In July Acryl signed a contract to supply GPUBASE to the modular data center (MDC) construction project led by Hancom. Following the basic tier, an advanced-tier contract worth KRW 3.3 billion (US$2.4 million) was added on July 29, bringing cumulative supply for the project to KRW 5.5 billion (US$4 million). The advanced tier runs on NVIDIA H200 NVL GPU servers and direct-liquid-cooled racks, with delivery scheduled to complete in the third quarter. A July 24 memorandum of understanding with Megazone Cloud extends coverage to large AI data centers and distributed AI infrastructure, and earlier agreements with MangoBoost and FuriosaAI widened the verification scope to Korean-made DPUs and NPUs.
In June Acryl was selected as lead organization for a project commissioned by Korea's Institute of Information & Communications Technology Planning & Evaluation (IITP) to develop an Ethernet-based GPU cluster network fabric. With a total budget of about KRW 6.7 billion (US$4.8 million), the project runs through December 2028 jointly with Yonsei University, Sungkyunkwan University and Ajou University. Its goal is to replace InfiniBand, NVIDIA's proprietary networking technology, with open Ethernet.

First-half results

Acryl reported on August 18 consolidated first-half revenue of KRW 12 billion, up 150% and equal to 90% of its full-year 2025 revenue. The operating loss ratio narrowed to 19% from 44% a year earlier as the company continued up-front investment in technology and computing infrastructure. The order backlog at the end of June was 49% higher than at the end of March. Medical AI contributed KRW 8.5 billion (US$6.1 million) in new orders this year on the combination of NADIA-ANE, an AI-native electronic medical record platform, and ALLM.H, a healthcare foundation model deployed on hospitals' own GPU servers. On September 1 the company reported that its 31-billion-parameter successor model, ALLM.H2 FRANK, answered all 435 physician-examination questions on the KorMedMCQA benchmark.

Jin Park, CEO of Acryl, said: "Competition in AI infrastructure is moving from how many GPUs a company has secured to how efficiently and reliably it operates the compute it holds. We will make joint verification and business cooperation concrete on the basis of the operating problems North American universities, research institutes and data centers are actually facing."

About Acryl Inc.

Acryl is an AX (AI transformation) infrastructure software company founded in March 2011 and listed on the KOSDAQ in December 2025. Its main products are the GPU operations platform GPUBASE; the integrated AI platform Jonathan; the industry-specific foundation model family ALLM; and the healthcare AI solution NADIA. Headquarters: 704 Seolleung-ro, Gangnam-gu, Seoul, Korea. www.acryl.ai
Won-dollar conversions at US$1 = KRW 1,390. Cost-saving figures are Acryl's own estimates.

Jerry (Jaehyuck) Yang
Acryl Inc
+82 10-9949-0296
jerry2@acryl.ai

This press release can be viewed online at: https://www.einpresswire.com/article/939925221

EIN Presswire's priority is source transparency. We do not allow opaque clients, and our editors try to be careful about weeding out false and misleading content. As a user, if you see something we have missed, please do bring it to our attention. Your help is welcome. EIN Presswire, Everyone's Internet News Presswire™, tries to define some of the boundaries that are reasonable in today's world. Please see our Editorial Guidelines for more information.
© 1995-2026 Newsmatics Inc. All Right Reserved.