

The background features a dark blue field with abstract, glowing patterns. On the left, there are vertical red and orange bars of varying heights. To the right, there are blue, pixelated shapes resembling continents or data clusters. Scattered throughout are various colored circles (red, orange, yellow, white) and thin vertical lines with dots at their ends. A solid red horizontal banner spans the middle of the image, containing the title text.

NVIDIA GTC 2026 Recap



Arató Bence

A BI Consulting igazgatója, két évtizedes szakmai múlttal az adattárházak, BI és analitika területén

- **Adatstratégiai tanácsadás**

Adat- és BI stratégia kialakítása, architektúratervezés, eszközválasztási tanácsadás, pályáztatás, minőségbiztosítás, mentoring.

- **Oktatás**

A BI Akadémia vezető oktatója (DW/BI architektúrák, adatmodellezés, adatvizualizáció). A CEU vendégtanára.

- **Rendezvények**

A Budapest BI Fórum és a Budapest Data+AI Fórum konferenciák főszervezője.

- **Kutatások**

A magyar Data Leaders Survey és a Data Salary Survey szakmai vezetője.



BI Consulting

Ajánlók

Budapest Data+AI Fórum (május 18-20.)

BUDAPEST **DATA + AI** FORUM

Data Platforms • Data Engineering
+
AI • Machine Learning • Data Science

May 18

DATA Day

In-person conference day

May 19

AI & ML Day

In-person conference day

May 20

DATA + AI Day

Online-only conference day

budapestdata26.hu

Előzetes a programból

Data előadók

AI & ML előadók

Kiemelt AI & ML előadók ezen a fülön tekinthetők meg. További előadók pedig ezen az oldalon érhetők el.



Balázs Lóránd
AUMOVIO (HU)
Establishing and scaling AI development @ AUMOVIO
ENTERPRISE AI CASE STUDY AI ADOPTION



Márton Hárs
MOL Group (HU)
AI Playbook supporting Explosive Retail Growth
ENTERPRISE AI CASE STUDY



Lotár Schin
Nokia (HU)
The Future of AI at Nokia
ENTERPRISE AI CASE STUDY



Zoltán Prekopcsák
Prezi (HU)
From SaaS to Agent-Native at Prezi
AI INNOVATION CASE STUDY



Eszter Windhager
Onekey (UK)
AI-Powered Firmware Health-Check
AI INNOVATION AGENTIC AI



Kornélia Papp
Zurich Insurance (CH)
From Clay to Fire: Building AI Products
ENTERPRISE AI APPLIED AI CASE STUDY



Wesley Boelrijk
KLM Royal Dutch Airlines (NL)
Bayesian Uncertainty Quantification for F1 Racing Strategies
AI TECHNOLOGY PYDATA







Peter Gostev
Arena (UK)
Battle of the labs - who is winning?
AI TECHNOLOGY GENERATIVE AI



Markus Nissl
Independent (AT)
AI-Driven Interfaces: Designing the New Grammar of Interaction
AI TECHNOLOGY AI ADOPTION

Bejelentett előadókink

 Matthias Rochtus Matthias Rochtus - Datashift - Senior Consultant	 Eugene Fidelin Engineering Manager & Frontend SME @ eBay	 Gábor Ertelt One Magyarország Zrt., Data Engineering Senior Manager	 Peter Vincze Head of Data & Analytics Platform at Generali Insurance Ltd.	 Jochen Christ Co-founder and CTO, Entropy Data	 Aris Tzoumas Staff Engineer @ RudderStack	 Evgeniya Alekseeva Customer Master Data Governance Lead @ Daikin Europe
 Guillermo Sanchez Dionis Developer Relations & Product at DuckDB Labs	 Balazs Sziladi Head of AI - Magyar Telekom	 Kornelia Papp Head of Conversational AI & Automation, Zurich Insurance	 Péter Szabolcs Lead Data Scientist at Mol Group's Digital Factory	 Mihaly Kavasi Fastrack Recognized Solution Architect, MCT, Group Manager - Avanade	 Timothy Spann Senior Solutions Engineer	 Lotar Schin AI Domain Architect @ Nokia
 Dennis Nerush Director of AI Engineering @Elementor	 Markus Nissl CEO @ data-insights.ai & markusnissl.com	 Naci Simsek Ververica, Manager - Customer Success Technical Engineering	 Balazs Lorand AUMOVIO Autonomous Mobility, Managing Director	 Catalin Gheorghiu Epam Systems	 Ferenc Györ CEO @ DAI Group	 Eszter Windhager ONEKEY GmbH, senior data scientist
 Wesley Boelrijk Lead Data Science @ KLM Royal Dutch Airlines	 Gergely Szücs 4iG Nyrt, Group Head of Data Domain	 Marton Hars Head of Data Analytics and AI, at MOL Group's Digital Factory	 Zsolt Csako Sr. BI Manager, ESAB corporation	 Egor Tarasenko Data & AI Engineer at Ponder	 Miky Schreiber Data Engineering platform TL, Next Insurance	 Peter Gostev AI Capability Lead @ Arena
 Zoltán Prekopcsák SVP Strategy & Analytics @ Prezi	 Radovan Bacovic Staff Data Engineer at GitLab	 Bence Arato Senior data architect & Data strategy advisor	 Gábor Gollnhofer Managing Director @ Meta Consulting	 Marin Niehues Founder & Principal Consultant	 Gergely Várhelyi-Tóth Advanced Technology Scientist @ TÜVRheinland, Founder @ roboGaze	

Jegyvásárlási akciók

Jegyvásárlás

Jelenleg a **Standard árak** vannak érvényben április végéig, valamint csoportos részvétel esetén további kedvezményeket kínálunk. Májusban a konferencia előtt a drágább, Late Bird árak lépnek érvénybe.

→ **További részletek és regisztráció** ←

All-Inclusive jegyek

Hibrid

Standard:
320 000 Ft

Late Bird 05.01-től:
400 000 Ft

Online

Standard:
190 500 Ft

Late Bird 05.01-től:
237 500 Ft

2 napos Data jegyek

Hibrid

Standard:
200 000 Ft

Late Bird 05.01-től:
250 000 Ft

Online

Standard:
130 500 Ft

Late Bird 05.01-től:
162 500 Ft

2 napos AI&ML jegyek

Hibrid

Standard:
200 000 Ft

Late Bird 05.01-től:
250 000 Ft

Online

Standard:
130 500 Ft

Late Bird 05.01-től:
162 500 Ft

DAI26-DATADUCK20 (20% kedvezmény)

Jegyvásárlási akciók



DAI26-DATADUCK20 (20% kedvezmény)



GTC 2026

NVIDIA GTC 2026



AI Conference & Expo March 16–19, 2026 | Keynote March 16 | Workshops March 15 | San Jose, CA & Virtual

EN


Keynote ▾ Explore ▾ Sponsors & Exhibitors More ▾

Watch On Demand

Explore the Best of GTC 2026

Get a recap of GTC news and learn about the breakthroughs shaping the next generation of AI with session replays.

Watch On Demand



Overview

 Highlights Featured Episodes Session Recommender Speakers Key Topics

NVIDIA GTC wrapped up a fantastic week in San Jose, bringing together thousands of developers, researchers, and business leaders to explore AI breakthroughs that are transforming every industry.


<https://www.nvidia.com/gtc/>

Data Foundations

Structured Data


Structured Data is the Ground Truth of AI

\$120B Structured Data Ecosystem – The Ground Truth of Enterprises



Unstructured Data

Unstructured Data is the Context of AI
100's of Zettabytes Per Year of Unstructured Data – Growing Exponentially



A billió dolláros kérdés

Inference workloads

Inference Inflection Arrives
10,000X ChatGPT Compute



1T Revenue in 2026-27

NVIDIA GTC Keynote 2026



I see through 2027. At least \$1 trillion.




Short vs. Long scale

Value	Zeros	us US English	short scale	hu Hungarian	long scale
10^6	6	million		millió	
10^9	9	billion		milliárd	
10^{12}	12	trillion	1 trillion USD	billió	1 billió USD
10^{15}	15	quadrillion		billiárd	
10^{18}	18	quintillion		trillió	
10^{21}	21	sextillion		trilliárd	
10^{24}	24	septillion		kvadrillió	

A magyar éves GDP nagyjából 240 milliárd dollár

Inference workloads

Inference Inflection Drives Strong Growth NVIDIA Full-Stack Expanding AI to All Regions and Industries




Extreme Co-Design


Extreme Co-Design

NVIDIA Extreme Co-Design Revolutionized Token Cost “GB NVL72 Inference King”

Tokens per Watt Drives Factory Revenue



Performance Drives Token Cost



Extreme Co-Design

AI Factories are the Industrial Infrastructure of the AI Era
Inference is the Workload
Tokens are the New Commodity
Compute is Revenue



Extreme Co-Design




Vera Rubin POD

AI FLOPS	60,000 PF
Memory Bandwidth	92 TB/s
All-to-All Bandwidth	260 TB/s
Scale-Out Radix	131,072

Extreme Co-Design

NVIDIA Vera Rubin 7 Chips – 5 Rack Systems AI Factory for the Agentic AI Frontier

1 GW AI Factory	X86 + Hopper	Vera Rubin
# of GPUs	600K	300K
AI FLOPS	1.2 ZFLOPS	16 ZFLOPS
All-to-All Scale-up	7.2 TB/s	260 TB/s
Memory BW-per-Domain (GROQ SRAM)	2 EB/s	100 EB/s
Tokens per Second	2M	700M



Rubin GPU



Vera CPU




CX9



BF4



NVLink Switch



Spectrum CPO



Groq 3 LPU



Vera Rubin Compute Tray



NVLink Switch Tray



Vera Compute Tray



BF4 STX Server



Spectrum Switch




Groq 3 Compute Tray




Inference at scale


Inference at scale

Inference Performance and Efficiency Drive Company Results






Inference at scale




Inference at scale



Inference Performance
and Efficiency Drive
Company Results

Inference at scale



Inference at scale



Groq, a \$20B-os nem-akvizíció


Platform ▾ Solutions ▾ Learn ▾ Pricing



Developers ▾ Enterprises

Start Building

Unmute



3M DEVELOPERS AND TEAMS

ramp

Dropbox

Vercel

Champion

VW

Canva

Rob

INFERENCE IS FUEL FOR AI

Groq delivers fast, low cost inference that doesn't flake when things get real.

Get Started



OFFICIAL PARTNER



SPEED AT A WINNING COST

THE GROQ LPU

TRY THE SPEED OF GROQ...



inference globally.

affordability at scale.

[Platform](#)  [Solutions](#)  [Learn](#)  [Pricing](#)



[Developers](#)  [Blog](#)

DEC 24, 2025

Groq and Nvidia Enter Non-Exclusive Inference Technology Licensing Agreement to Accelerate AI Inference at Global Scale

Today, Groq announced that it has entered into a non-exclusive licensing agreement with Nvidia for Groq's inference technology. The agreement reflects a shared focus on expanding access to high-performance, low cost inference.


As part of this agreement, Jonathan Ross, Groq's Founder, Sunny Madra, Groq's President, and other members of the Groq team will join Nvidia to help advance and scale the licensed technology.

Groq will continue to operate as an independent company with Simon Edwards stepping into the role of Chief Executive Officer.


GroqCloud will continue to operate without interruption.

Az NVIDIA 20 milliárd dollárt „felvásárolta” a Groq-ot 2025. decemberében
(technilailag licenszelési megállapodás + kulcsemberek átvétele)


Low-latency premium inference



Low-latency premium inference



Low-latency premium inference with Groq 3 LPX




Low-latency premium inference with Groq 3 LPX

NVIDIA Vera Rubin + LPX
Unlocks \$300B Revenue
Opportunity



Rubin + Groq 3 together

Rubin GPU




288 GB HBM4

22 TB/s

50 PFLOPs (NVFP4)

336B Transistors
+ 2.5T (HBM4)

Groq 3 LPU



500 MB SRAM

150 TB/s SRAM Bandwidth


1.2 PFLOPs (FP8)

98B Transistors

Uniting Processors of Extreme Performances

Rubin + Groq 3 together

Rubin GPU




288 GB HBM4

22 TB/s

50 PFLOPs (NVFP4)

336B Transistors
+ 2.5T (HBM4)

Groq 3 LPU



4 GB SRAM

1,200 TB/s SRAM Bandwidth

9.6 PFLOPs (FP8)

784B Transistors

55X


Uniting Processors of Extreme Performances



Valami szokatlan

Personal Agentic AI

Inference Inflection Arrives



Personal Agentic AI

Announcing NVIDIA NemoClaw for OpenClaw


```
ubuntu@brev-ae6ah8vzm: ~  
→ ~ curl -fsSL https://nvidia.com/nemoclawn.sh | bash  
→ ~ nemoclawn onboard
```



Personal Agentic AI


Announcing NVIDIA NemoClaw Reference OpenClaw

NVIDIA Agent Toolkit for Building Specialized Agents




Personal Agentic AI

Enterprise IT Renaissance from SaaS to Agent-as-a-Service



És végül

ChatGPT reviews GTC 2025...




(Blue (28:17): This is how intelligence is made, a new kind of factory generator of tokens, a new kind of factory-generator of tokens, the building blocks of AI. Tokens have opened a new frontier in to aim to get us there faster. so next time we'll be ready. So let's get started. Tokens connect the dots so we can protect our most noble creatures. They turn potential into plenty. And help us harvest amidst hope to bravely go where no one has gone before. And here it is where it all begins.

Speaker 2 (31:30). Welcome to GTC.


Welcome to GTC. What an amazing year. We wanted to do this at NVIDIA founder and CEO, Jensen Huang.

Welcome to GTC. What an amazing year. We wanted to do this at GTC. Most 3130: What an amazing year. I wanted to do this at Nvidia so the magic of artificial intelligence, we're going to bring you you AMD GTC. What an amazing year. I've been here. We're starting. We're on GeForce 5090, the Blackwell generation. I'm comparing it to the GTC 5090, the Blackwell generation. AI, Tokens have opened in to more great better at dissipating energy, and incredible performance. Amanda A;


ChatGPT reviews GTC 2025...



GTC 2026 Hype Words



GTC 2026 Hype Words



Thank you!

