


2025

Highlights

 NVIDIA  
GTC















“NVIDIA is hosting the Super Bowl of AI”

*The New York Times*



# “NVIDIA mobilizes IT partners to spread the AI gospel”

CIO Live

GTC used to be a gathering for a few thousand developers. This year, we're filling stadiums. It's been called the "Woodstock of AI"—this year, the "Super Bowl of AI." The only difference? Everyone's a winner. More companies, more industries, more breakthroughs. Every year, GTC just gets bigger because AI solving more interesting problems for more industries and more companies.




# “NVIDIA Showcases AI Future at GTC”

Yahoo Finance

AI started with perception—seeing and hearing the world. Then it learned to generate. Now it can reason, plan, and act. The next wave is physical AI—machines that understand the world and move through it. Each step opens up new possibilities. And we’re just getting started.

2012 ALEXNET



## PERCEPTION AI

SPEECH RECOGNITION  
DEEP RECSYS  
MEDICAL IMAGING

## GENERATIVE AI


DIGITAL MARKETING  
CONTENT CREATION


## AGENTIC AI

CODING ASSISTANT  
CUSTOMER SERVICE  
PATIENT CARE

## PHYSICAL AI

SELF-DRIVING CARS  
GENERAL ROBOTICS





**“NVIDIA’s big show hits its marks”**

THE WALL STREET JOURNAL.

## CUDA-X For Every Industry



“NVIDIA's true strength  
lies in its ecosystem”

ARC Advisory Group

This is what GTC is all about. A long time ago, this was the only slide we had—one library after another. Just as we needed AI frameworks to create AIs, we need frameworks for physics, biology, quantum, and more. We call them CUDA-X libraries—acceleration frameworks for every field of science.


“This is NVIDIA at its best”

Ben Thompson / Stratechery






“Three scaling laws are stacked together and working in tandem”



SemiAnalysis

There are three fundamental questions in AI. How do you solve the data problem? How do you train without humans in the loop? And how do you scale—how do you find an algorithm where the more resources you

provide, the smarter AI becomes? This is where almost the entire world got it wrong. The scaling law of AI is more resilient—and in fact, hyper-accelerated.



**“NVIDIA’s strategy becomes clear: taking on a much larger role than ever before”**

Technalysis

AI is going to go everywhere. The rest of the world has different system configurations, different environments, different needs—enterprise IT, manufacturing, robotics, even new GPU clouds. Each one is building their own stack, and NVIDIA is helping power them all.



**“NVIDIA says every company will become an AI factory”**

Business Insider


Every company that builds something will have two factories: one to make the product, and one to manufacture the intelligence behind it. That second factory—the AI factory—will become essential.

# “NVIDIA launches Dynamo, the operating system of an AI factory”

MarketWatch


We call it NVIDIA Dynamo. It's the operating system of an AI factory. It manages parallelism, batching, memory, context—everything needed to orchestrate massive models across thousands of GPUs. Just like VMware was for enterprise IT, Dynamo is for AI agents.

## Announcing NVIDIA Dynamo Distributed Inference Serving Library





## Inference At-Scale is Extreme Computing



“NVIDIA hasn’t slowed the pace. It is accelerating inference throughput gains”

SemiAnalysis

Inference is the ultimate extreme computing problem. You're trying to generate a whole bunch of tokens—but you're trying to do it as quickly as possible. Smart AIs that are super fast. That's the trade-off. The perfect answer is the upper right of the Pareto frontier—high interactivity, high throughput. That's the best AI factory you can build.

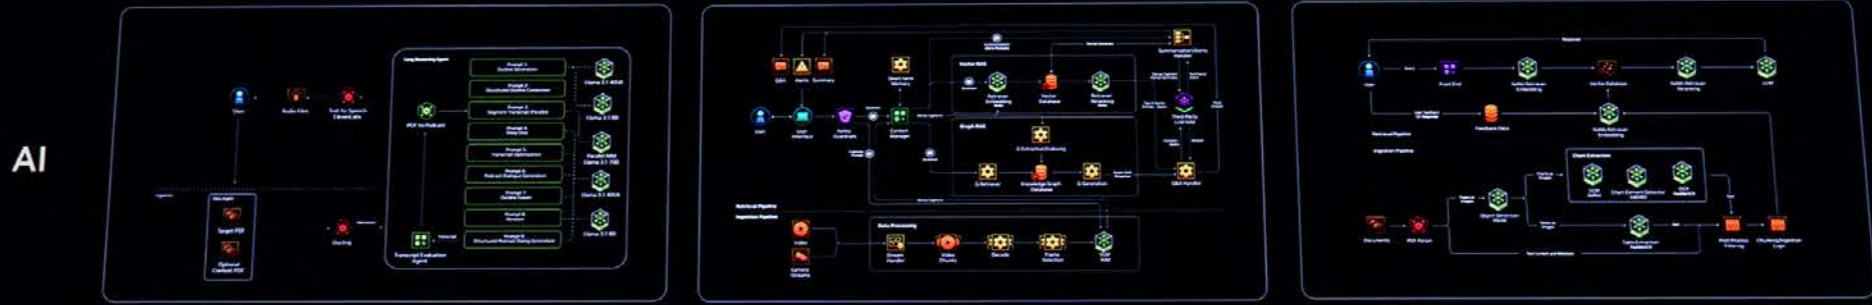


# Reinventing \$500B Enterprise IT For the Age of AI

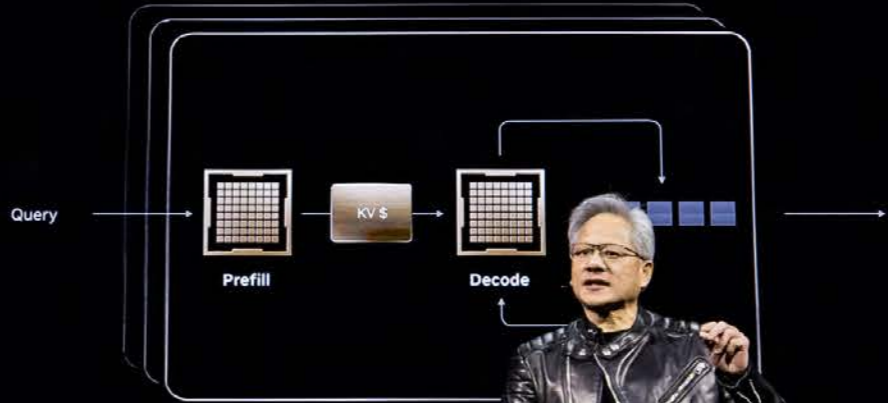
“Tech giant details massive AI opportunity ahead”

Investors Business Daily


We're building AI infrastructure for the cloud, for enterprise IT, and for robots. And now, sovereign AI is emerging as the fourth. Just as we've modernized the cloud, we're now reinventing the world's \$500 billion enterprise IT industry for the age of AI.



OS



INFRASTRUCTURE



Compute




Networking



Storage

## Announcing Storage Leaders Build AI Data Platforms for Enterprise AI



“Storage players ride the NVIDIA bus at GTC 2025”


Computer Weekly

Storage has to be completely reinvented for the age of AI. It won't just retrieve data—it will embed and interpret it. We're working with the entire storage industry to build systems that are GPU-accelerated and semantics-driven. This is the enterprise AI infrastructure of the future.

# “NVIDIA unveils powerful new Blackwell Ultra chips at Super Bowl of AI”

New York Post

Blackwell is in full production, and the ramp has been incredible. Customer demand is strong—and later this year, we'll introduce Blackwell Ultra: more memory, more bandwidth, and support for longer context and deeper reasoning. It's the next step in building AI infrastructure.



“We saw NVIDIA up close and personal last week, and were blown away”

Jim Cramer / CNBC

Lenovo

CoreWeave  
DELL Technologies

Microsoft  
Azure

aws

Meta

Google Cloud

ORACLE  
CLOUD  
Infrastructure

Hewlett Packard  
Enterprise

AIVRES

ASUS

GIGABYTE


Ingrasys

Inventec

PEGATRON

SUPERMICRO

QCT






“NVIDIA’s silicon photonics-based switch enable clusters with millions of GPUs”

Tom’s Hardware


To scale to hundreds of thousands of GPUs, we need to move beyond copper. That’s where silicon photonics comes in. We’ve built the world’s first 1.6 terabit-per-second co-packaged optics system—saving megawatts of power and unlocking the next generation of AI factories.



# “NVIDIA networking steals the show for the second GTC in a row”

VentureBeat

To build AI factories at scale, we had to reinvent networking. Spectrum-X brings AI performance to Ethernet, and with our silicon photonics technology, we can now connect hundreds of thousands of GPUs while saving megawatts of power. This is the network fabric of the AI era.




Spectrum-X  
Integrated Silicon Photonics  
2nd Half 2025

Quantum-X  
Integrated Silicon Photonics  
2nd Half 2026

## NVIDIA Photonics

CPO Co-Invention With Ecosystem Partners

- 1st 1.6T Silicon Photonics CPO Chip - New Micro Ring Modulators (MRM)
- 1st 3D-Stacked Silicon Photonics Engine with TSMC Process
- High-Power, High-Efficiency Lasers
- Detachable Fiber Connectors
- 100's of Patents, Licensed to Partners



“NVIDIA's Spark desktop AI supercomputer arrives this summer”

Engadget


This is what a PC should look like in the age of AI. DGX Spark delivers 20 petaFLOPS of performance, 72 CPU cores, and high-bandwidth memory—all in a workstation. It's built for researchers and developers powering the next wave of AI.

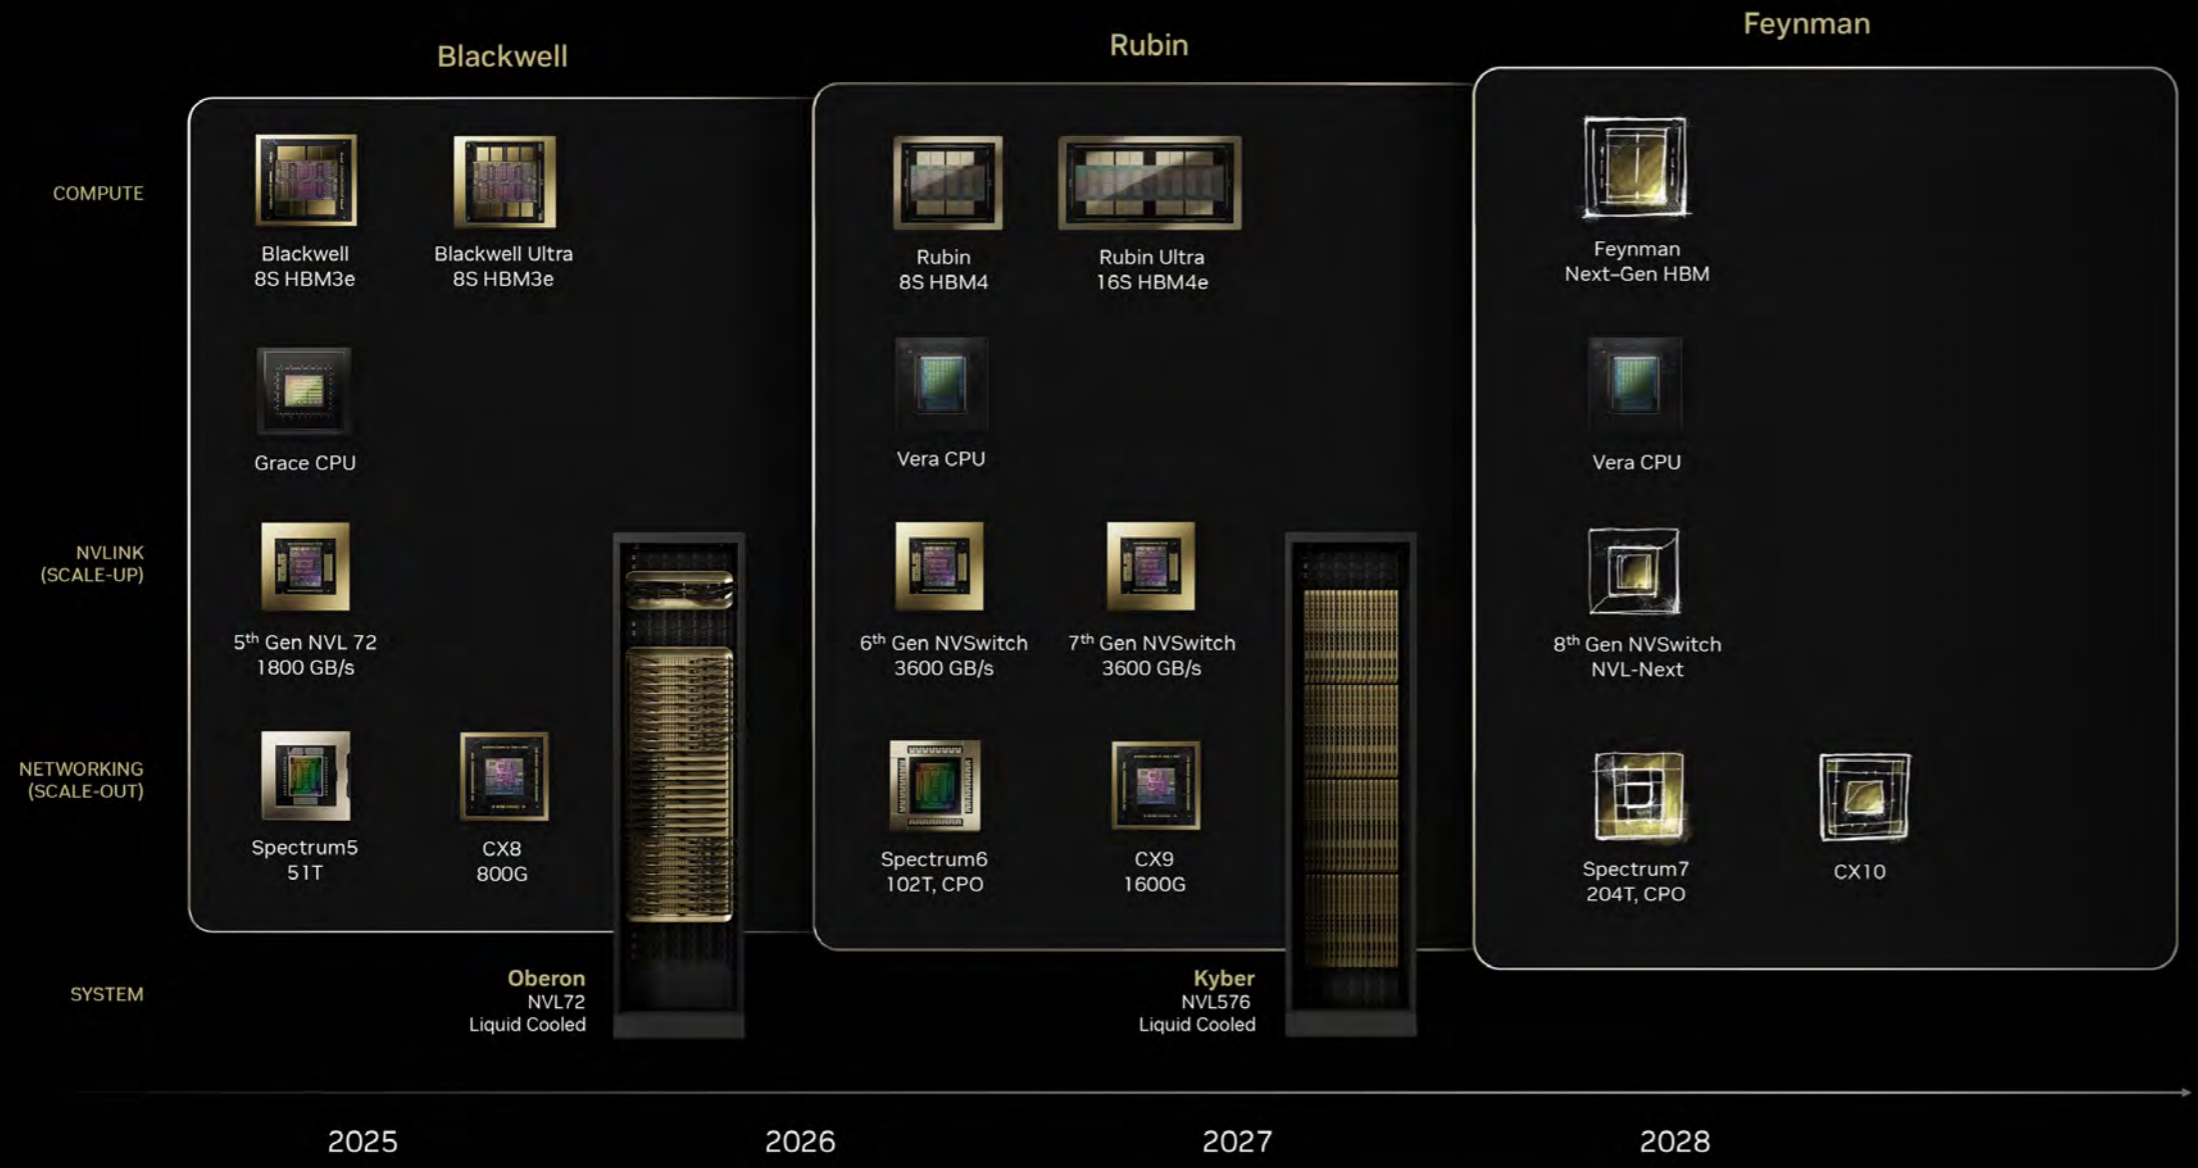


# “What’s NVIDIA’s next big thing? Superchip, named after astronomer who discovered dark matter”

The Wall Street Journal

Rubin is our next-generation platform—brand new CPU, GPU, memory, networking, everything. It’s designed for the future of AI: trillion-parameter models, sovereign infrastructure, and massive-scale AI factories.






“NVIDIA delivers multiyear roadmap at GTC”

Bloomberg

No company in history has ever laid out a roadmap like this—multiple years, multiple architectures, every detail. But we do it because we’re building infrastructure. The world is counting on us, and our partners need trust, no surprises, and confidence in execution.



NVIDIA AI Infrastructure for Enterprise Computing



“NVIDIA GTC 2025:  
AI Matures into Enterprise Infrastructure”

Bain & Company

We’re not a chip company—we’re an infrastructure company.

From RTX Pro to DGX Station, from B200 to GB300, we build the full stack for AI: systems, networking, software, everything.

This is NVIDIA.

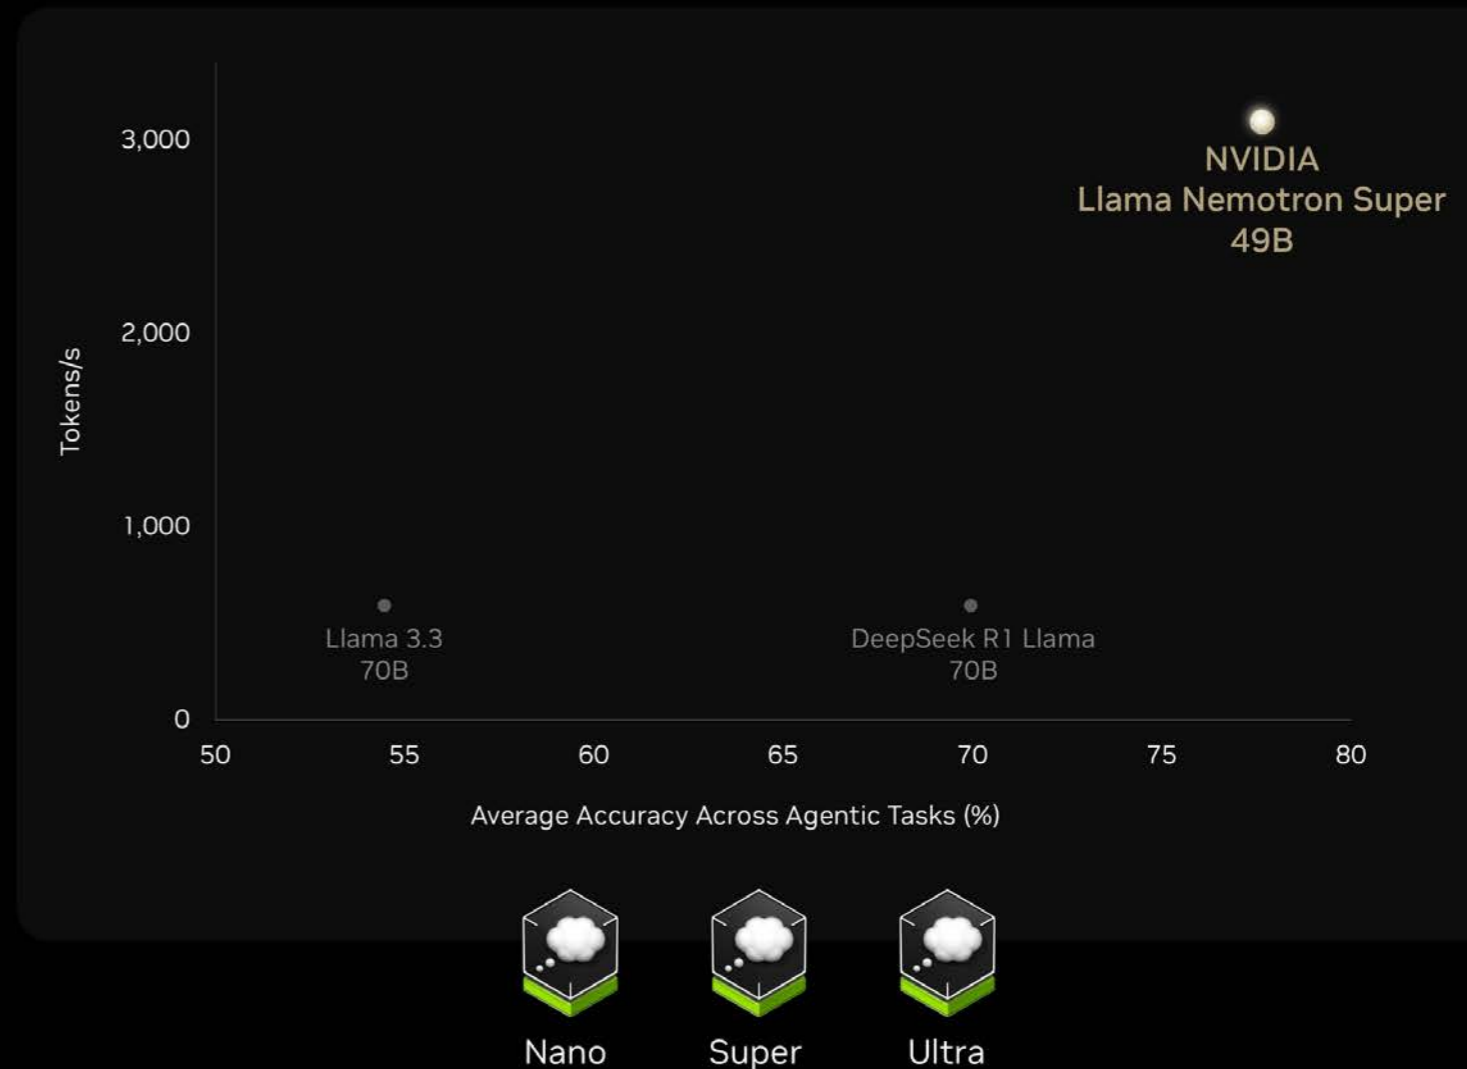
# “NVIDIA’s new reasoning models pave way for advanced AI agents”

SiliconANGLE

Our Llama Nemotron model is built for reasoning—step-by-step thinking, planning, and problem solving. It’s not just about answering questions. It’s about understanding context, breaking down complex tasks, and taking action.

## Announcing NVIDIA Llama Nemotron Reasoning

Distilled, Quantized, Aligned, and Optimized by NVIDIA





## “NVIDIA launches blueprint to revolutionize weather forecasts”

IT Brief

We’re seeing more extreme weather events than ever—threatening lives and property. The NVIDIA Omniverse Blueprint for Earth-2 will help industries prepare for and mitigate climate disasters with faster, more accurate, energy-efficient forecasts.



## ROBOTS – THE NEXT MULTI-TRILLION DOLLAR INDUSTRY




## “NVIDIA Dreams of a Robot Future”

The New York Times

The time for robots has come. Everything that moves—cars, forklifts, surgical arms, warehouse bots—is becoming autonomous. To build them, we need three computers: one to

train the AI, one to simulate it, and one to run it. We call that physical AI—AI that understands the laws of physics, that can reason, plan, and act in the real world.



# “Robots Invade NVIDIA’s GPU Technology Conference”

CNET

In robotics, training is just the beginning. Post-training, AI models need to reason, adapt, and simulate in real time. Test-time scaling becomes critical—robots don’t just recall, they think and act in dynamic environments. That’s why we’re building physical AI infrastructure.



“NVIDIA’s Cosmos makes robot training freakishly realistic—and that changes everything”

VentureBeat

Omniverse is our operating system for physical AIs. We use it to condition Cosmos, our generative model that understands the physical world. Together, they create infinite training environments that are grounded in physics and fully under our control.



NVIDIA Omniverse With Cosmos  
Physical AI Digital Twin Operating System



NVIDIA Omniverse



Cosmos




# “NVIDIA teams up with Disney and Google to build smarter robots”

The Globe and Mail

Today, we're announcing something really, really special—a partnership of three companies: Google DeepMind, Disney Research, and NVIDIA. We call it


Newton. It's a physics engine designed for fine-grained control, soft bodies, tactile feedback—everything you need to train robots with precision.



“Cutest robots I’ve ever seen”

Monica White / Digital Trends






## “Disney’s robotic droids are the toast of Silicon Valley”

The Wall Street Journal

This is Besh, one of the BDX droids built by Disney Research, powered by two NVIDIA computers and trained in simulation. Tactile feedback, soft-body physics, fine motor skills—Disney’s BDX droids have got it all. This is what happens when physical AI meets a little bit of magic.

Announcing NVIDIA Isaac GROOT N1  
Humanoid Foundation Model



“NVIDIA’s latest GROOT AI model just took another step closer to fully humanoid robots”

TechRadar

We’re open-sourcing GROOT N1—our foundation model for humanoid robots. It learns by watching humans, understanding language and video to generate actions. With GROOT,

developers can teach robots skills through demonstration. Companies like Boston Dynamics, Figure, and Unitree are already building on it. This is how we’ll bring physical AI to life.

# “GM’s partnership with NVIDIA could change driving forever”


Autoblog

We’re partnering with GM to build their future self-driving fleet. AI in the car, AI in the factory, AI in the enterprise—every part of the company will be powered by NVIDIA. The time for autonomous vehicles has arrived.






NVIDIA GTC 2025



BLACKWELL IN FULL PRODUCTION

\$1T Computing Inflation Point



BLACKWELL NVL72 AND DYNAMO 40X FOR INFERENCE

Reasoning 100X One-Shot Blackwell 40X Hopper



VERA RUBIN NVIDIA PHOTONICS AI INFRASTRUCTURE FOR AI CLOUDS

Annual Rhythm for the World to Build-Out AI Infrastructure



AI INFRASTRUCTURE FOR \$500B ENTERPRISE IT

New Compute, Networking, Storage, Software



AI INFRASTRUCTURE FOR ROBOTS

Physical AI For \$50T Industrial and Robotics

“Jensen Huang’s keynote was fantastic”

Jim Cramer / CNBC



What the press said...

**“NVIDIA is the undisputed tone-setter of the entire technology industry.”**

Creative Strategies

**“Analysts roundly agree that NVIDIA remains far ahead”**


The Wall Street Journal

**“GTC felt more bullish than ever”**

TechCrunch

**“NVIDIA delivered at GTC again”**

Quartz



What the financial analysts said...

**“NVIDIA remains the king of the hill”**

Citi

**“NVIDIA continues to dominate the AI value chain”**

Bank of America

**“From the Woodstock of AI to the Superbowl”**

Daiwa


**“It was the best of times, it was the best of times”**

Cantor

**“Blackwell ramp was the man of the match”**

Itau




























245K

Conference  
Registrations

27K

In-Person  
Attendees

1,200

Talks and  
Sessions

43M

Keynote  
Views

58K

Press  
Articles

400

Partner  
Sponsorships

We talked about several things.

Blackwell is in full production. We're building AI infrastructure for the cloud, for enterprise, and for robots. We're laying the blueprints not just for our company—but for the industries that are here, the ones that aren't here, and the entire supply chain.

I want to thank all of you for coming to GTC.  
**This is an extraordinary moment in time.**

Jensen Huang



