







George
Kurtz

Arthur
Kaindl

Frédéric
Godemel

Arkady
Volozh

Pascal
Daloz

Deirdre
Bosa

David
Silver

Jensen
Huang

Robin
Rombach

Gavin
Baker


Artem
Sokolov

Ola
Källenius

Jeannette
zu Fürstenberg

Zhaopeng
Chen

What Do L'Oréal, Revolut, Spotify, Merck, Vodafone, Zara,
Porsche, Nestlé, and HSBC Have in Common?





\$10B



Over the Last Year

14 New Investments

In Start-Ups This Week

12K+

Companies in Our Ecosystem





AI Powers a New Industrial Revolution

APPLICATIONS


MODELS

AI FACTORIES

COMPUTE

ENERGY


Building Europe's AI Economy



The NVIDIA DSX Full-Stack AI Factory Platform









CUDA-X Libraries



Platform Libraries




Open Models




NVIDIA DSX

NVIDIA Vera Rubin Powers the Whole Agent

Vera Rubin NVL72
Compute Node
Model Processing




Vera MGX
Compute Node
Secure Agent Processing




Vera CMX
Storage Node
Knowledge Serving




Vera RTX
Compute Node
CUDA-X Tools Processing





NVIDIA Vera Rubin
NVL72

“Vera Rubin will generate significantly cheaper tokens than the next-leading accelerator and thus make you significantly more money.”





Announcing NVIDIA Vera Rubin DSX POD Full-Stack Modular AI Factory Reference Architecture

Factory-Built, Integrated, Pre-Tested Clusters
Delivery-to-Deployment Cut From Months to Days
The Container is the AI Factory

ANDURIL

ARMADA

Bull

Crusoe

DAYONE

DELL Technologies

DELTA

EATON

GIGABYTE™

HPE

Ingrasys®

Lenovo


PEGATRON

QCT


Schneider
Electric

VERTIV


wlstron



One CUDA Architecture. Many Markets. Productive, Fungible, and Durable.




Inference Inflection Arrives



Useful AI Creates Demand for More Intelligence


Agents used **34x** more tokens today than at the start of Feb

Weekly tokens used by agents across OpenRouter | Sep 2, 2025 - Oct 5, 2026





Subscribe for more token insights: openrouter.ai/data

AI Factories Produce Intelligence. Agents Put Intelligence to Work.



Announcing NVIDIA OpenShell Reference Architecture for Safe AI Agents





Over 100 Partners Advance Agent Safety with NVIDIA OpenShell

ANTHROPIC



Google Cloud



JPMorganChase



OpenAI




servicenow

SIEMENS



A new approach to agent security with Claude Managed Agents and NVIDIA

On September 28, NVIDIA announced the Open Agent Safety Platform, an open software platform and reference system design for strengthening AI security. Anthropic and NVIDIA collaborated to bring additional layers of security and control to the agent stack.




Open Models Turn Expertise Into Advantage




Ising Quantum Computing



Earth-2 AI Physics



BioNeMo
Biology AI



Nemotron
Agentic AI



Cosmos
Physical AI



GR00T
Robotics





Alpamayo Autonomous Vehicles

Turn Proprietary Data Into Sovereign Intelligence

NVIDIA Nemotron is “Incredibly Customizable”



ALICE

AI21



Turn Proprietary Data Into Sovereign Intelligence

NVIDIA Nemotron is “Incredibly Customizable”





Turn Proprietary Data Into Sovereign Intelligence

NVIDIA Nemotron is “Incredibly Customizable”

Harvey

Trajectory




Turn Proprietary Data Into Sovereign Intelligence

NVIDIA Nemotron is “Incredibly Customizable”


SKYDANCE

SCALE



Turn Proprietary Data Into Sovereign Intelligence


NVIDIA Nemotron is “Incredibly Customizable”




An Ecosystem to Help Build Your Super Agents

 adaption	 AfterQuery	 AI21	 Applied Compute	 baseten
 cohere	 Deep Cogito	 distil labs	 Fireworks	<i>Handshake</i>
 inferact	 LangChain	 Letta	 MERCOR	 Mistral
 MULTIVERSE COMPUTING	 Palantir	 PRIME Intellect	 RadixArk	 Reflection
 river	 SCALE	THINKING MACHINES	 Trajectory	 Vercel

NVIDIA Agent Knowledge Foundation. Turn Company Data Into Company Knowledge.




NVIDIA Agent Knowledge Foundation



ParseBench



nvidia/Nemotron-3.5-Parse




lightonai/LightOnOCR-3-0.8B



PaddlePaddle/PaddleOCR-VL-1.6

4




PaddlePaddle/PaddleOCR-VL-1.5

5



Datalab-to/surya-ocr-2


NVIDIA Agent Knowledge Foundation



LMEB

- #1 [nvidia/Nemotron-3.5-Embed-8B-BF16](#) PARETO
- #2 [BAAI/bge-multilingual-gemma2](#)
- #3 [nvidia/Nemotron-3-Embed-1B-BF16](#) PARETO
- #4 [DataScience-UIBK/uibk-embedding-v1](#)
- #5 [tencent/KaLM-Embedding-Gemma3-12B-2511](#) PARETO


NVIDIA Agent Knowledge Foundation



ViDoRe v3

- #1 [nvidia/Nemotron-3.5-Rerank-8B-BF16](#) PARETO
- #2 [Qwen/Qwen3-VL-Reranker-8B](#)
- #3 [Qwen/Qwen3-VL-Reranker-2B](#) PARETO
- #4 [nvidia/llama-Nemotron-rerank-vl-1b-v2](#)
- #5 [Qwen/Qwen3-VL-Embedding-2B](#) PARETO

NVIDIA Agent Knowledge Foundation




BeyondArena		
0	Kumo-Tabular	1426
1	TabPFN-3.5	1376
2	Kumo-Tabular-Medium	1376
3	Causilo	1327
4	Kumo-Tabular-Small	1320



NVIDIA CUDA-X Tools Give Agents the Tools of the Trade

Simulators | Solvers | Databases | Digital Twin Engines
Optimization Engines | Search Engines | Recommendation Engines









∞ Meta




Aerial	ALCHEMI	cuEquivariance	cuTILE	ModelExpress	nvTIFF	Automodel	cuFFT	cuSOLVER	DAQIRI	MONAI	Omniverse	BioNeMo	cuGraph	CUDA CCCL
Dynamo	NeMo Gym	NIXL	OptiX	Cosmos	cuML	FlashDreams	Newton	nvCOMP	Retriever	cuBQL	cuProj	cuVS	Foundation Stereo	cuDART
NVENC/NVDEC	Sionna	CUDA Graphs	CUDA-Q	cuRAND	Holoscan	nvJPEG	Thrust	cuDNN	CuTe DSL	Earth-2	Kumo	nvmath-python	Triton	cuPotential
cuEmbed	cuTensor	Megatron	NCCL	NVSHMEM	Warp	cuDF	cuMotion	cuSPARSE	Isaac Lab	nvJPEG2000	Transformer Engine	CUDA Math API	cuLitho	CuPy
cuQuantum	GSF Ontology	nvImageCodec	TensorRT-LLM	cuBLAS	cuOpt	FlashInfer	JaxPP	NPP	RMM	CompileIQ	cuISS	FastGen	Nemotron	cuSOLVERMp
Parabricks	PhysicsNeMo	BioLR	cuFolio	deepeeRT	GR00T	NeMo	OpenShell	AmgX	cuDSS	cuEST	CUTLASS	ModelOpt	nvtrust	NeMo RL

Lila Sciences Lab Assistant Agents Use NVIDIA BioNemo and CUDA-X Tools



BioNeMo MONAI Parabricks CodonFM Boltz-2 Proteina GenMol Omniverse Cosmos Isaac

Siemens Factory Agents Use NVIDIA Cosmos and CUDA-X Tools



NeMo Helix

Nemotron

Omniverse

Isaac



cuOpt

Kumo

Cosmos

Metropolis

Dassault Systèmes Virtual Companions Use NVIDIA Nemotron and CUDA-X Tools



- Nemotron Parse
- cuBLAS
- cuDSS
- OptiX
- cuFFT
- cuSOLVER
- cuSPARSE
- NeMo Eval Author
- NeMo Compass
- NeMo AutoModel

3D

VR

3DEXPERIENCE | 3DSwym VW-Tiguan_Gen4_MY2028 program milestone

Search

Jessica DUPUIS

VW-Tiguan_Gen4_MY2028 program milestone


AURA, LEO, Marc FISCHER, Sophie LAURENT

mechanical Engineering

SAG simulation

SAG simulated on the detailed assembly - 100 kg static load applied at the latch point.

SAG Simulation



Result: 9.6 mm deflection — within the predicted range (9.4 mm ± 0.6), under the 10 mm limit.

Verdict: OK

LEO 7:52 PM

Mechanical Engineering

Side-crash simulation

Side Crash solved under Euro NCAP MDB and FMVSS 214 pole-test boundary conditions.

Side Crash Simulation



Result: 120 mm intrusion (limit 130 mm), energy absorption target met.

Verdict: OK

+

Start writing











AURA may display inaccurate information. Please check the answers.

Side Crash Simulation

Open with





Enterprise Platforms Build Agents on NVIDIA Agent Toolkit





Announcing Telecom Networks
Become Distributed AI Infrastructure
With NVIDIA Aerial AI-RAN





Announcing Siemens Healthineers Next-Generation Adaptive Radiotherapy



NVIDIA AI for Healthcare

Nemotron

Omniverse


Holoscan

Cosmos-H

Isaac-H

IGX Thor

NVIDIA Hyperion: Building a Global Robotaxi Fleet and Ride-Hailing Ecosystem



NVIDIA AI for Robotaxis

Alpamayo

Halos SAFE

Halos OS

Cosmos

Hyperion

AlpaSim

Bolt

Grab

lyft

Uber

Mobility
Platforms

AUTOBRAINS

DEEPROUTE.AI

momenta

Motional

AV
Software

nuro

pony.ai

TIER IV

waabi

wayve

WeRide

ZYT

BYD

DiDi

GEELY

HYUNDAI
MOTOR GROUP

Auto
Manufacturers

LUCID



NISSAN
MOTOR CORPORATION

STELLANTIS



VOLVO

ZEEKR

Announcing Europe's Automakers Build on NVIDIA Hyperion



NVIDIA AI for Robotaxis

Alpamayo

Halos SAFE

Halos OS

Cosmos

Hyperion

AlpaSim

NVIDIA Isaac: Building the Platform for the Next Great Industry

Isaac OS
NVIDIA Jetson Thor
Onboard Compute
Halos Safety OS



NVIDIA AI for Robotics

GR00T



Omniverse


Halos

Cosmos

Isaac Lab

Isaac OS

				Model Builders
Physical Intelligence				
				Robot Platforms
				
				
				
				



Isaac OS

NVIDIA Jetson Thor
Onboard Compute

Halos Safety OS

Agile ONE SN
31-DOF Body


WUJI Hand 2 20-DOF
5-Finger Hands

Announcing Agile ONE SN Humanoid Robot Built on Isaac GR00T



Available Q1 2027

NVIDIA AI for Robotics





*“NVIDIA RTX Spark changes what
a Windows laptop can be”*

— CGMagazine

*“NVIDIA’s RTX Spark is an
earthquake for the PC industry”*

— PCMag UK

acer

ASUS

DELL Technologies

GIGABYTE™



Lenovo

Microsoft



 Choose project  This computer

Download, install, and setup Hermes Agent with OpenShell on my RTX Spark.

Always use CUDA, RTX and OpenGL Accelerated tools.

Use the best open model that will fit on my RTX Spark PC.

I prefer Nemotron whenever similar performance.

+  Approve for me

GPT-6 Astra High ▾



