











Useful AI Has Arrived




AGENT = LLM + HARNESS



Agents Go to Work

Make a JavaScript Newton's Cradle animation in a single HTML file.



Agents Go to Work

Create a gif of NVIDIA-green dots on black scatter, form Taipei 101 building, morph to "GTC TAIPEI 2026," morph to the NVIDIA eye logo, then scatter and repeat.



CONTEXT

OBSERVE


REASON


ACT




Agents Go to Work

I lost my remote control battery clip, it looks like this. Create a CAD file ready for 3D printing to create a new one.






AGENT = LLM + HARNESS



Full-Stack AI Infrastructure Company



NVIDIA DSX AI Factory Platform

DSX OS


Power Optimization



Infrastructure Software



Platform Software




Libraries and Frameworks


DSX Sim


DSX MaxLPS

DSX Flex



Reference Designs





NVIDIA DSX AI Factory Ecosystem

AI Clouds

firmus

indosat
OREDOO HUTCHISON

GMO
INTERNET

NAVER Cloud

SHARON
- AI -

YTL AI Cloud
YTL GROUP

CoreWeave

Lambda

NEBIUS

NSCALE

GMI

ResetData

Visionbay.ai

YOTTA

I R E N

Mistral AI

AI Factory Software

ddn

emeraldai

MIRANTIS

netris

Open Nebula

phaidra

RAFAY

Red Hat

Spectro Cloud

TrendAI

VAST

WEKA

Design and Construction

AVEVA

BECHTEL

cādence

DASSAULT SYSTEMES

etap

Jacobs

PROCORE

ptc

SIEMENS

Energy and Cooling

ABB

CATERPILLAR

DELTA

EATON
Powering Business Worldwide

ENGIE

GE
GE VERNOVA

HITACHI

mitsubishi electric

nationalgrid

Schneider Electric

SILICON VALLEY POWER
CITY OF SANTA CLARA

TRANE TECHNOLOGIES

VERTIV

Infrastructure and Facilities

AIRTRUNK

BDX
DATA CENTERS

CDC

CSQUARE

DAYONE

oio
Digital Edge PC

DIGITAL REALTY

EQUINIX

NEXTDC

NTT

Princeton Digital Group

sify

STTelemedia
Global Data Centres

Compute Systems

ASUS

DELL Technologies

FOXCONN

GIGABYTE

HPE

Lenovo

PEGATRON

QCT

SUPERMICRO

wistron

wiwynn



 nvidia

 CoreWeave

 **cohere**

 **CURSOR**

 **Jane Street**

ANTHROPIC

OpenAI

 **Mistral AI**

runway


W&B Models


W&B Training


W&B Inference


W&B Weave


CoreWeave Inference


CoreWeave Slurm on Kubernetes


CoreWeave Kubernetes Service


CoreWeave Compute


CoreWeave Storage


CoreWeave Mission Control

NVIDIA AI Platform

cuDNN

Magnum IO

cuOpt

cuVS

cuDF

Omniverse

TRT-LLM

Dynamo

NeMo

PhysicsNeMo

Cosmos

Nemotron



NVIDIA Hopper


NVIDIA RTX PRO Blackwell Server Edition


NVIDIA Grace Blackwell


Confidential Computing


NVIDIA Vera Rubin


NVIDIA ConnectX and BlueField Networking


NVIDIA Quantum IB, Spectrum-X Networking



Multi-Tenant AI Cloud



Nebius Agentic Search

NVIDIA AI Platform

- cuDNN
- Magnum IO
- cuOpt
- cuVS
- cuDF
- Omniverse
- TRT-LLM
- Dynamo
- NeMo
- PhysicsNeMo
- Cosmos
- Nemotron



NVIDIA Hopper



NVIDIA RTX PRO Blackwell Server Edition




NVIDIA Grace Blackwell



Confidential Computing



NVIDIA Vera Rubin



NVIDIA ConnectX and BlueField Networking



NVIDIA Quantum IB, Spectrum-X Networking



Google

THINKING
MACHINES

Training and Inference

NVIDIA AI Platform

cuDNN

Magnum IO

cuOpt

cuVS

cuDF

Omniverse

TRT-LLM

Dynamo

NeMo

PhysicsNeMo

Cosmos

Nemotron



NVIDIA
Hopper



NVIDIA RTX PRO Blackwell
Server Edition




NVIDIA
Grace Blackwell



Confidential
Computing



NVIDIA
Vera Rubin



NVIDIA ConnectX
and BlueField Networking



NVIDIA Quantum IB,
Spectrum-X Networking

 nVIDIA

NAVER Cloud



HyperCLOVA X

NAVER World Model



NAVER Cloud AI Factory



NVIDIA AI Platform

cuDNN

Magnum IO

cuOpt

cuVS

cuDF

Omniverse

TRT-LLM

Dynamo

NeMo

PhysicsNeMo

Cosmos

Nemotron



NVIDIA
Hopper



NVIDIA RTX PRO Blackwell
Server Edition




NVIDIA
Grace Blackwell



Confidential
Computing




NVIDIA
Vera Rubin



NVIDIA ConnectX
and BlueField Networking



NVIDIA Quantum IB,
Spectrum-X Networking



NVIDIA AI Platform

cuDNN

Magnum IO

cuOpt

cuVS

cuDF

Omniverse

TRT-LLM

Dynamo

NeMo

PhysicsNeMo

Cosmos

Nemotron



NVIDIA Hopper



NVIDIA RTX PRO Blackwell Server Edition




NVIDIA Grace Blackwell



Confidential Computing



NVIDIA Vera Rubin



NVIDIA ConnectX and BlueField Networking




NVIDIA Quantum IB, Spectrum-X Networking



Training and Inference

NVIDIA AI Platform

cuDNN	Magnum IO	cuOpt	cuVS	cuDF	Omniverse	TRT-LLM	Dynamo	NeMo	PhysicsNeMo	Cosmos	Nemotron
											
NVIDIA Hopper	NVIDIA RTX PRO Blackwell Server Edition	NVIDIA Grace Blackwell	Confidential Computing	NVIDIA Vera Rubin	NVIDIA ConnectX and BlueField Networking	NVIDIA Quantum IB, Spectrum-X Networking					



Training and Inference

NVIDIA AI Platform

cuDNN

Magnum IO

cuOpt

cuVS

cuDF

Omniverse

TRT-LLM

Dynamo

NeMo

PhysicsNeMo

Cosmos

Nemotron



NVIDIA Hopper



NVIDIA RTX PRO Blackwell Server Edition




NVIDIA Grace Blackwell



Confidential Computing



NVIDIA Vera Rubin








NVIDIA ConnectX and BlueField Networking



NVIDIA Quantum IB, Spectrum-X Networking




TrendAI Fireworks AI Reflection wlstron

- 
GMI Model Garden
- 
GMI Model Fine Tuning
- 
GMI Prime Inference
- 
GMI Dedicated Endpoints
- 
Agent Execution Infra
- 
Agent Dedicated Isolation
- 
Agent Distribution Ecosystem
- 
GMI Compute
- 
GMI Hybrid Clusters
- 
GMI Kubernetes Control

System Builders and OEMs

NVIDIA AI Platform

- cuDNN
- Magnum IO
- cuOpt
- cuVS
- cuDF
- Omniverse
- TRT-LLM
- Dynamo
- NeMo
- PhysicsNeMo
- Cosmos
- Nemotron
- 
NVIDIA Hopper
- 
NVIDIA RTX PRO Blackwell Server Edition
- 
NVIDIA Grace Blackwell
- 
Confidential Computing
- 
NVIDIA Vera Rubin
- 
NVIDIA ConnectX and BlueField Networking
- 
NVIDIA Quantum IB, Spectrum-X Networking


COMPUTE IS REVENUE

EXTREME CO-DESIGN — TTFT, TOKEN/WATT

SCALE & EXPERIENCE — TTFT, MTBI

CUDA ECOSYSTEM & INSTALLED BASE — USEFUL LIFE

AI FACTORY
REVENUE




COMPUTE IS REVENUE

EXTREME CO-DESIGN — TTFT, TOKEN/WATT

SCALE & EXPERIENCE — TTFT, MTBI

CUDA ECOSYSTEM & INSTALLED BASE — USEFUL LIFE

AI FACTORY
REVENUE




COMPUTE IS REVENUE

EXTREME CO-DESIGN — TTFT, TOKEN/WATT

SCALE & EXPERIENCE — TTFT, MTBI

CUDA ECOSYSTEM & INSTALLED BASE — USEFUL LIFE

AI FACTORY
REVENUE




COMPUTE IS REVENUE

EXTREME CO-DESIGN — TTFT, TOKEN/WATT

SCALE & EXPERIENCE — TTFT, MTBI

CUDA ECOSYSTEM & INSTALLED BASE — USEFUL LIFE

AI FACTORY
REVENUE




COMPUTE IS REVENUE

EXTREME CO-DESIGN — TTFT, TOKEN/WATT

SCALE & EXPERIENCE — TTFT, MTBI

CUDA ECOSYSTEM & INSTALLED BASE — USEFUL LIFE

AI FACTORY
REVENUE




COMPUTE IS REVENUE

EXTREME CO-DESIGN — TTFT, TOKEN/WATT

SCALE & EXPERIENCE — TTFT, MTBI

CUDA ECOSYSTEM & INSTALLED BASE — USEFUL LIFE

AI FACTORY
REVENUE





Vera Rubin in Full Production

ODMs and OEMs

DELL Technologies

HPE

Lenovo



CLOUDIAN

ddn

Everpure

Hitachi Vantara

IBM

NetApp

NUTANIX

vast

WEKA

AIVRES

ASUS

FOXCONN

GIGABYTE

Inventec

PEGATRON


QCT

wlstron

wiwynn




NVIDIA Vera – CPU for the Age of AI




Vera Rubin
GPU


Grace
BlueField



Vera
Compute



Vera BlueField
Storage



NVIDIA Vera

CPU for Agents

NVIDIA-Custom Olympus Core

88-Core / 176-Threads
Spatial Multithreading

10-Wide Instruction Fetch/Decode

2MB L2 per Core / 164MB L3

250W - 450W TDP

1.2 TB/s LPDDR5X ECC

40% Lower Loaded Latency
(vs x86)

3.4 TB/s Core-to-Core BW


1.4 TB/s PCIe Gen6

1.8 TB/s NVLink C2C



NVIDIA Vera


3X Bandwidth per Core vs. x86






NVIDIA Vera

1.8X Agentic Sandbox Performance vs. x86







“NVIDIA Vera is the most formidable competition to Intel and AMD x86_64 processors ever realized.”



Michael Larabel

Author, Phoronix CPU Benchmark Suite

Vera CPU 3X Faster SQL Data Processing

“With Starburst on NVIDIA Vera CPUs with Dell, agents can instantly query governed enterprise data at scale, unlocking the ground truth to power decisions at the speed the agentic era requires.”



— Justin Borgman, CEO and Founder of Starburst



Vera CPU 6X Faster Real-Time Streaming Data Processing

“In collaboration with Redpanda and HPE, using NVIDIA Vera CPUs, we will be scaling our capacity, while further optimizing latency to power a high-performance, resilient and AI-ready market infrastructure.”

— Lynn Martin, President of NYSE Group





NVIDIA Vera Opens Large Market

CPU for Agents

Early Adopters

OpenAI

ANTHROPIC

SPACEX

Cloud

NEBIUS

ORACLE
CLOUD
Infrastructure

CoreWeave

NSCALE

Crusoe

firmus

Lambda

together.ai

Ecosystem Partners

DELL Technologies

HPE

Lenovo

SUPERMICRO

AIVRES

ASRock
Rock

ASUS

COMPAL

FOXCONN

GIGABYTE

hyve
solutions

Inventec

MITAC
MITAC Computing

msi


PEGATRON

QCT


wlstron

wiwynn

AGENT = LLM + HARNESS



Announcing NVIDIA Agent Toolkit for Enterprise AI



Announcing Cadence Chip Design Super Agent



u_dut.core0.pipe.decode.mux.cycle[15:0]	00EA	00EA	
top.core2.l1d.data.pll_locked	0		
tb.dut.gpc[2].mon.imm[7:0]	41	40	
top.core0.irq.parity[63:0]	B7E99ACAA9706800	B7E99ACAA9706800	
tb.dut.way.clk_div2	1		
u_dut.warp[26].mux.user[63:0]	CF8D446ABC2CB800	CF8D446ABC2CB800	
u_dut.clk_gen.gpc[0].tpc[1].sm.warp[14].ch[14].way.operand[15:0]	E9C5	E9B5	
tb.dut.core3.l2.addr.resp[31:0]	EB67146A		
tb.dut.core0.gpc[4].tpc[1].sm.warp[19].beat[31:0]	B70B341C	B70B3420	
u_dut.nvlink.fsm.valid	0		
top.core3.clk_gen.tpc[2].tag[15:0]	6548	6548	
tb.dut.core2.nvenc.gpc[3].tpc[11].sm.warp[19].state.clk_sys	1		
u_dut.nvlink.tpc[15].data.beat[7:0]	61		
u_dut.core0.nvlink.cmd.slot[31:0]	0000006D	00000055 // 0000005D // 00000065 // 0000006D // 00000075	
top.core0.gpc.tpc[5].state.row[31:0]	DEEF580F		
tb.dut.core0.slot[31:0]	000000C4	000000B4 // 000000C4	
top.mem.grant	0		
top.cuda.addr[31:0]	A5CC8BF7		A5CC8BF7
top.ch[14].state.mask[7:0]	5B	4B	5B
top.core0.tensor.gpc[6].tpc[7].sm.warp[10].key[63:0]	27D415B60D77C600	27D415B60D77C600	
u_dut.tex.mux.result[63:0]	117B355B70944C00	117B355B70944C00	11
top.core3.mem.gpc[4].tpc[1].sm.warp[21].busy	1		
u_dut.dma.irq.ce	0		
u_dut.tensor.ctrl.beat[7:0]	22	F0	22
u_dut.core2.pcie.ch[5].err.parity[7:0]	FB	3C	FB
top.mem.addr.row[63:0]	94A53FDEF10D2800		94A53FDEF10D2800
top.cycle[31:0]	000000F7	000000EF // 000000F7	

Announcing NVIDIA Nemotron 3 Ultra


Frontier Smart

	Nemotron 3 Ultra (550B)	GLM 5.1 (744B)	Kimi K2.6 (1T)	Qwen3.5 (397B)
Agent Productivity PinchBench	91%	84%	91%	89%
Long-Horizon Planning EnterpriseOps-Gym	33%	40%	29%	30%
Coding Terminal-Bench 2.0	54%	64%	67%	53%
Instruction Following IFBench	82%	77%	74%	78%
Knowledge Work GDPVal-AA	1,448	1,594	1,508	1,192
Professional Work Tasks ProfBench (Search)	56%	46%	56%	53%
Long Context Ruler @1M	95%	N/A (max 256K)	N/A (max 256K)	90%




Announcing NVIDIA Nemotron 3 Ultra

5X Faster




Announcing NVIDIA Nemotron 3 Ultra


30% Lower Cost




Enterprise Software Platforms Build Agents on NVIDIA




RTL Engineer for EDA




Detection Triage, Malware Analysis, Threat Hunting




Engineering and Scientific Workflow Automation




AI FDE Supply Chain Agent




Expense Accounting, Supply Chain



IT Services and Workflow Automation



Autonomous Chip Design




Autonomous Debug and Verification Engineer for EDA

40 Years of Personal Computing



NVIDIA and Microsoft Reinvent PC



Local LLM


DeepSeek


Gemma


GLM



GPT-OSS



Kimi



MiniMax


Nemotron


Qwen


WINDOWS





Announcing NVIDIA and Microsoft Reinvent PC


Powered by RTX Spark

Blackwell RTX GPU
1 Petaflop FP4 AI Performance


20 Core Grace CPU
Custom Built with MediaTek

128 GB Unified Memory
600 GB/s NVLink C2C

Full NVIDIA Stack
CUDA | TensorRT | NVFP4
RTX Ray Tracing | DLSS | Reflex | G-SYNC



Accelerate Every AI. Every App. Every Game.





Announcing New Adobe Premiere and Photoshop for RTX Spark

Up to 2X Faster
Creative Agent-Ready

Announcing NVIDIA RTX Spark Laptops

Smart. Powerful. Beautiful.



Available Fall

acer

ASUS

DELL

GIGABYTE™



Lenovo

Microsoft

msi




Three Revolutionary Windows Machines

One Architecture. Agent-Ready.




A New Line. A New Beginning.






Compute is Data

WORLD FOUNDATION MODEL








Announcing NVIDIA Cosmos 3 — The Frontier of Physical AI

Artificial Analysis Text to Image

 NVIDIA	Cosmos3-Super-Text2Image	 Open Weights	1,243
 HiDream	HiDream-O1-Image-Dev-2604	 Open Weights ⓘ	1,190
 Alibaba	Qwen Image Max 2512	 Open Weights	1,161
 Black Forest Labs	FLUX.2 [dev]	 Open Weights	1,158
 Fal	FLUX.2 [dev] Turbo	 Open Weights ⓘ	1,156
 Fal	FLUX.2 [dev] Flash	 Open Weights ⓘ	1,143
 Black Forest Labs	FLUX.2 [klein] 9B	 Open Weights	1,126

Artificial Analysis Image to Video

 NVIDIA	Cosmos3-Super-Image2Video	 Open Weights	1,246
 Lightricks	LTX-2 Pro	 Open Weights ⓘ	1,192
 Lightricks	LTX-2 Fast	 Open Weights ⓘ	1,178
 Lightricks	LTX-2.3 Fast	 Open Weights	1,162
 Lightricks	LTX-2.3 Pro	 Open Weights	1,152
 Tencent	HunyuanVideo-1.5 (Fal)	 Open Weights	1,118
 Alibaba	Wan 2.2 A14B	 Open Weights	1,110

PAI-Bench World Generation Accuracy

Cosmos3-Super	83.9	89.5
Cosmos3-Nano	83.7	89.4
Source	82.6	87.1
Veo-3	82.1	86.7
Cosmos-Predict2.5-14B	81.0	83.8
Cosmos-Predict2.5-2B	81.0	84.0
Wan2.2-I2V-A14B	80.6	84.1

Physics-IQ Physics Accuracy

Cosmos3-Super + WMReward (BoN) (multiframe (v2v)) 🏆	63.4%
Magi-1 + WMReward (BoN) (multiframe (v2v)) 🏆	62.6%
Cosmos3-Super (multiframe (v2v)) 🏆	59.7%
Cosmos3-Nano + WMReward (BoN) (multiframe (v2v))	57.7%
Magi-1 (multiframe (v2v))	56.0%
Cosmos3-Nano (multiframe (v2v))	50.2%
Cosmos3-Super + WMReward (BoN) (i2v) 🏆	48.9%
Sora2 + WMReward (BoN) (i2v) 🏆	46.4%
Wan2.2 + WMReward (BoN) (i2v) 🏆	44.4%

R-Bench Embodiment Video Generation

Cosmos3-Nano	Open-source Top-1	2	0.584	0.550
Cosmos3-Super	Open-source Top-2	4	0.581	0.487
Wan2.2-A14B		10	0.507	0.381
Cosmos 2.5		11	0.464	0.358
HunyuanVideo 1.5		12	0.460	0.442
LongCat-Video		13	0.437	0.372
Wan2.1-14B		16	0.399	0.344
LTX-2		17	0.381	0.284
Wan2.2-5B		18	0.380	0.331




RoboArena Action Policy Score

Cosmos3-Nano-Policy	1870	122.1
Spirit v1.6	1785	43
DreamZero ✕	1732	37.5
WALL-OSS	1657	66.4
pi05_droid ✕	1555	24.3
balsomsoc	1550	33.9
pi0_fast_droid ✕	1549	23.3
floppy_bagel_195k_temp1p0	1522	151.0

RoboLab Manipulation Policy Success

1	Cosmos3-Nano-Policy
2	μ2-vla
3	DreamZero ➤
4	WALL-OSS
5	pi05_droid ➤
6	balsomsoc
7	pi0_fast_droid ➤
8	floppy_bagel_195k_temp1p0

VANTAGE-Bench World Understanding

Cosmos3-Super 🏆	65.07	86.97	76.23
  			
Gemini 3.1 Pro 🏆	61.92	76.15	67.17
 			
Cosmos3-Nano 🏆	60.67	74.11	75.57
  			
Qwen3-VL-32B-Instruct 🏆	55.05	64.33	72.39
			

TAR Anomaly Detection Reasoning

1	N/A	Cosmos3-Super	0.5484	0.8250	0.8500	0.3107
2	38	Qwen3.6-27B	0.5146	0.8000	0.8500	0.2428
3	N/A	Cosmos3-Nano	0.4348	0.7375	0.8250	0.2728
4	N/A	Qwen3.5-27B	0.3875	0.7750	0.8000	0.2434
5	N/A	Cosmos-Reason2-32B	0.3752	0.6625	0.7750	0.2465
6	N/A	Gemini-3.1-Pro-Preview	0.3538	0.6875	0.7750	0.2251
7	N/A	Gemma-4-31B-It	0.3389	0.6250	0.7500	0.2302
8	N/A	Cosmos-Reason2-8B	0.3381	0.4375	0.7250	0.1779
9	N/A	Qwen3-VL-8B-Instruct	0.3096	0.6000	0.6250	0.2346

Pick up the screwdriver and put it on the top rack.

Plan:

1. Lower the gripper to grasp the screwdriver.
2. Lift the screwdriver from the table.
3. Move the screwdriver horizontally toward the white rack.
4. Place the screwdriver on the designated shelf of the white rack.



4X Speed


Announcing AI Leaders Join NVIDIA Cosmos Coalition




NVIDIA DRIVE Hyperion Robotaxi Platform

Announcing Alpamayo 2 Super Robotaxi Reasoning Model

Alpa Simulators



Alpamayo and OmniDreams
Open Models




Training Data Generation



NVIDIA Halos OS



DRIVE AGX Thor



Hyperion

Global Automakers Build Robotaxis on NVIDIA DRIVE Hyperion







NVIDIA Isaac GR00T Humanoid Robot Platform



Isaac Simulators




GR00T Open Models



Training Data Generation



NVIDIA Isaac OS



Jetson Thor



Unitree H2 Plus

Announcing NVIDIA Isaac GR00T Reference Humanoid Robot

Jetson Thor Onboard Compute

2070 FP4 TFLOPS | 14-Core Arm CPU | 128 GB memory


31-DOF Unitree H2 Plus Robot



22-DOF 5-Finger Sharpa Hands

GR00T Software Stack

Data Generation and Policy Validation
Unitree Whole-Body Control | GR00T 1.7 Model

Available October 2026







Vera Rubin in Full Production



NVIDIA Vera
CPU for Agents




NVIDIA Agent Toolkit
for Enterprise AI



NVIDIA and Microsoft
New Line of PCs



The Next Wave
is Physical AI



NVIDIA DSX AI Factory Platform





