EN

GTC 2026 — Post B1

10 AI Announcements for the Enterprise

GTC 2026 Series · BOTUM Analysis · March 2026

March 2026

Table of Contents

Introduction

On March 16, 2026, Jensen Huang took the stage at the SAP Center in San Jose to deliver one of the densest keynotes in NVIDIA history. For over two hours, he unveiled a complete vision of enterprise agentic AI — from chips to agents, from Earth to orbit. BOTUM was there, taking notes.

Here are the 10 announcements that truly matter for solutions architects.

---

1. OpenClaw + NemoClaw — The World's Most Popular Open Source Agent

Jensen Huang was unequivocal: "OpenClaw is the most popular open source project in the history of humanity."

This open source project enables deploying autonomous AI agents with a single command and has become the de facto standard for agentic development.

With NemoClaw, NVIDIA goes further — a secure agent runtime built on OpenClaw, designed for enterprise:

- Native sandboxing: each agent runs in an isolated environment
- Least-privilege model: agents receive only strictly necessary permissions
- Integrated privacy router: automatic filtering of sensitive data before LLMs
- Complete audit trail: every agent action is tracked

Security partners: Cisco, CrowdStrike, Google, Microsoft Security, TrendAI.

i BOTUM angle: At BOTUM, we've been using OpenClaw for months, well before GTC. NVIDIA's validation confirms our technological vision and head start.

---

2. NVIDIA Agent Toolkit — The Complete Agentic AI Stack

NVIDIA presents a unified ecosystem for deploying AI agents at enterprise scale. 4 pillars:

- Nemotron: open source models optimized for agentic AI (multi-step reasoning, planning, tool use)
- AI-Q Blueprint: hybrid architecture reducing costs by 50% via intelligent routing between frontier and Nemotron models
- OpenShell: secure runtime for agent execution with network policy management
- cuOpt: GPU-native mathematical optimization library

17 major partners: Adobe, Salesforce, SAP, ServiceNow, Siemens, CrowdStrike, Atlassian, Cadence, Synopsys, IQVIA, Palantir, Box, Red Hat, Cisco, Dassault Systemes, Cohesity and Amdocs.

---

3. Vera Rubin Platform — The Agentic Era Platform

Vera Rubin — named after the astronomer who discovered dark matter — is the next-generation computing platform. A complete system: 7 chips, 5 rack-scale systems.

- NVIDIA Vera CPU: optimized for orchestrating AI workloads at rack scale
- Rubin GPU: new graphics architecture with integrated HBM3e
- BlueField-4 STX: next-gen SmartNIC for storage and networking

Orders already placed: over $1 trillion for the 2025-2027 period.

4. Feynman (2028) — The Next Generation Takes Shape

- Rosa CPU: named after Rosalind Franklin, optimized for agentic pipelines
- LP40 LPU: Low-Power Latency Processing Unit for ultra-low latency inference
- TSMC A16 process (1.6nm) with custom HBM and 3D stacking
- Kyber interconnect: optical scale-up network for 100,000+ GPU clusters

5. Dynamo 1.0 — The OS for AI Factories

If Vera Rubin is the hardware, Dynamo 1.0 is the operating system. Orchestrates thousands of GPUs, manages request queues, guarantees latency SLAs.

Already in production: AWS, Azure, GCP and Oracle Cloud. It's the Kubernetes of agentic AI.

6. DSX Air — Digital Twin for AI Factories

DSX Air (Omniverse DSX Blueprint) lets you create a complete digital twin of an AI Factory — rack layout, thermal flows, network consumption — before placing the first server. Strategically major for teams planning large-scale deployments.

7. NVIDIA Space-1 — AI Takes Orbit

NVIDIA Space-1: AI data centers in Earth orbit. The Vera Rubin platform adapted for space deployments — remote sensing data processing, real-time Earth observation, ultra-low latency satellite communications.

8. CUDA Turns 20 — And DLSS 5 Arrives

Jensen pays tribute to CUDA, the GPU language launched in 2006 — the foundation of modern AI. "CUDA is the flywheel of accelerated computing"

Bonus: DLSS 5, real-time 4K neural rendering.

9. Nemotron Coalition — The Open Frontier Model Alliance

NVIDIA launches the Nemotron Coalition: an alliance to develop open frontier models — auditable LLMs deployable on-premise.

A sovereign alternative to proprietary models for enterprises that cannot send their data to external clouds.

10. Oracle AI Database 26ai — Agent Meets Database

Oracle and NVIDIA announce Oracle AI Database 26ai with GPU-accelerated vector indexes. A critical building block for production RAG architectures.

What This Means for Solutions Architects

Key Announcement	Architect Impact
OpenClaw + NemoClaw	Core competency to acquire immediately
Agent Toolkit	Reference stack for enterprise deployments
Vera Rubin	Hardware platform of the agentic era
Dynamo 1.0	Kubernetes for AI — SREs must master it
AI-Q Blueprint	50% reduction in inference costs
Nemotron Coalition	Sovereign alternative to proprietary models
Space-1	Orbital compute — emerging use cases

At BOTUM, we've been preparing for months — and this blog is the proof. 10 technical guides on our OpenClaw deployment are available today.

Need help with your agentic stack?
BOTUM teams help organizations evaluate and implement their AI agent infrastructure.
Contact: www.botum.ca/contact

-> Online version: blog.botum.ca/gtc2026-b1-10-announcements-enterprise-ai/