Capgemini
RESEARCH INSTITUTE

Digital Sovereignty

Balancing Control, Dependencies, and Innovation

Conversations
FOR TOMORROW

N°11 — 2026

CALISTA REDMOND

VP Global AI Initiatives, Nvidia

BUILDING NATIONAL AI CAPABILITY: INVEST LOCALLY, ENGAGE GLOBALLY

Calista Redmond is the Vice President of Global AI Initiatives at NVIDIA, where she leads strategic collaborations with regions, ecosystem partners, and product groups to advance national AI programs worldwide. Before joining NVIDIA, Calista served as CEO of RISC V International, driving global adoption of the open RISC V instruction set architecture across industry, academia, and government. Previously, she spent more than a decade at IBM in leadership roles including VP of the IBM Z Ecosystem as well as open source initiatives such as the Open Mainframe Project, OpenPOWER, and OpenDaylight. Earlier in her career, Calista was an entrepreneur in four successful technology start ups. She holds degrees from the University of Michigan and Northwestern University and is based in Washington, D.C.

THE RISE OF NATIONAL AI

The conversation around national AI has intensified recently. What is driving that urgency?

Calista Redmond: Generative AI has really changed the strategic value here. National and enterprise data becomes the raw material for productivity and the raw material for competitiveness for industry. The urgency comes from dependency risk. If a country is challenged in accessing, tuning, auditing, or operating AI that subscribes to their local regulations, and is instead inheriting someone else's language coverage, policy assumptions, and service levels, then both generative AI and inference AI become affected.

As we move from proof of concepts and experimentation into bringing AI into critical workflows, this raises the bar. Now, it includes where the data lives, the provenance of the model, cybersecurity, audit, and other operating controls. It really weaves together a full strategic view of that value. At its core, it comes down to cultivating a local ecosystem and ensuring compute access.

Calista Redmond
VP Global AI Initiatives | Nvidia

How does NVIDIA think about its role in this landscape?

Calista Redmond: Every nation needs to build their own national AI capability. This means countries should be able to produce intelligence using their own data, their own infrastructure, their own workforce, and their business networks. They are no longer just consumers of intelligence produced somewhere else. They should be able to harness the power of AI to produce intelligence that is meaningful for moving their countries forward.

> **"Every nation needs to build their own national AI capability"**

NVIDIA is there to support governments and their national champions: local model builders, startups, and existing industry champions. It really takes all parts of the stack, from different levels of infrastructure and energy through to LLMs and applications that serve the purpose they are trying to accomplish. A common first steppingstone is developing a local model that brings local data and local expertise, reflects that geography and domain, and sits on top of global-grade infrastructure. That is where we team together and take a collaborative approach.

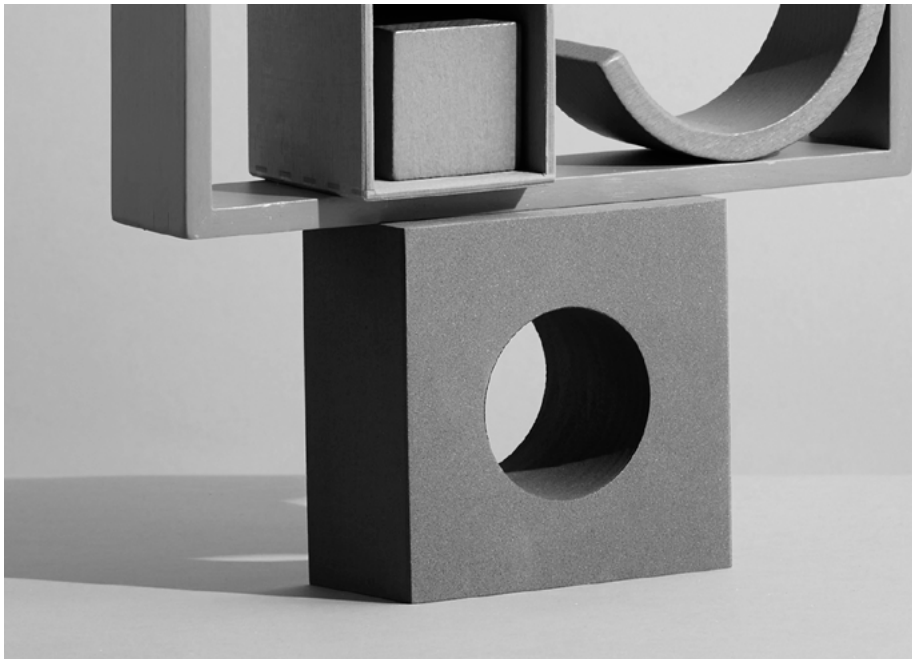

BUILDING NATIONAL AI

If you were advising a government on building national AI capability from the ground up, what are the foundational building blocks?

Calista Redmond: I think about it across six areas.

First, data. Make sure they have trusted, permissioned, very high quality national and sector datasets, including the IP rights and data sharing rules. High quality input matters.

Second, compute and infrastructure. The ability to combine accelerated compute, networking, and storage, with clear energy resources, security, and the orchestration that nations require for training, fine-tuning, and simulating workloads.

Third, models. I would expand that beyond a single local model to a portfolio of open, commercial, and locally trained and fine-tuned models. This enables model choice across workloads and makes token utilization more efficient.

Fourth, software. This means the right tools to customize, evaluate, guardrail, deploy, and manage everything around those models. This is the orchestration piece.

Fifth, applications and agents. Applications can be citizen services, fraud detection, public health, customs, and tax copilots. Many of these in government are about efficiency. Agents go further and can boost the capacity of a nation, through digital twins, logistics, weather, and agriculture. These are things governments support not just for operational efficiency but for strengthening a nation.

Sixth, people and ecosystems. Across the countries I interact with, a common theme I come across is about workforce development and cultivating talent. That talent comes from AI literacy, so that every discipline embraces AI as part of their craft, and from AI engineers and developers building capabilities into startups and filling out the ecosystem. Part of what governments do here is transform university and research output into commercialization.

If you look across all six, the framing is: "invest local, engage global." You need local strength and capacity, but you need global engagement for developer communities, for supply chains and resilience, and for markets, so you can tap into the economic value of that local investment.

What is the right governance posture for countries that want control over AI models and workflows, particularly those relying on frontier models they do not own?

Calista Redmond: The governance challenge is shifting from content quality to action quality. Agents can now plan, call tools, access systems, and trigger downstream actions. Governance needs to cover the entire workflow, not just the content being generated.

That requires a tiered risk approach. Low-risk agents can handle high automation. In a document review workflow, for example, they can process at speed and flag only the anomalies. High-impact agents require human approval, separation of duties, policy checks, and clear accountability measures. The goal is automation at the right levels and human intervention at the right levels.

Some of the core controls include privileged tool access, role-based permissions, sandboxed execution environments, human-in-the-loop gates at critical points, and the more traditional disciplines: audit logs, red teams, and rollback capability. These are not new concepts in software development. What is new is applying them to AI workflows and then continuously checking for compliance.

"The governance challenge is shifting from content quality to action quality"

Does regulation slow innovation?

Calista Redmond: It depends entirely on the quality of the regulation. Good regulation accelerates adoption and innovation. Murky or bad regulation slows it down.

Positive regulation creates trust, procurement clarity, and a common language for public and private sector deployment. If everyone knows the rules and the framework, it is easier to get going. The key is making requirements risk-based and outcome-focused rather than prescriptive. Focus on safety, auditability, and data residency as the critical elements, rather than a single frozen technical approach. Think about regulatory sandboxes, validated reference architectures, and certification pathways. What are the methods and paths that can be used as a standard to accelerate time to production, while safeguarding citizens and enterprises?

The reverse is also true. Murky or bad requirements can force every workload into the same local-only architecture, prohibit model choice, or make compliance impossible, especially for startups and SMEs. The principle should be to make it easier to build the right way rather than to operate without governance or figure it out later.

GOOD REGULATION ACCELERATES ADOPTION AND INNOVATION. MURKY OR BAD REGULATION SLOWS IT DOWN.

How strategically important is compute infrastructure for a country pursuing AI ambitions?

Calista Redmond: Compute is foundational. Data and talent cannot become a national strength or capability without accessible compute. You need compute for training, tuning, simulation, and inference. If local talent does not have access, or you cannot instrument local data on that compute, that is going to slow innovation capability. Compute is really driving who gets to experiment, who can build local language models, and who can deploy low-latency services and operate AI at scale.

> "Data and talent cannot become a national strength or capability without accessible compute"

Building local champions, whether model builders, application developers, or companies applying AI to existing industries, is critical for competitiveness and for a country moving forward. Compute is the foundation of all that.

Not every workload needs frontier-scale systems, but every country does need a tiered plan. I work with countries that have national-scale AI factories alongside enterprise data center capacity, cloud access, and edge infrastructure. Think about the right infrastructure and models to overcome the challenge at hand. By making those careful choices, a nation can build AI infrastructure very similar to a national utility.

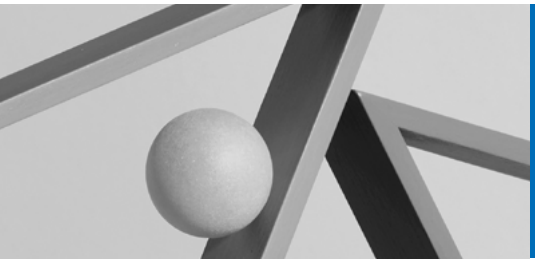

> "Not every workload needs frontier-scale systems"

AI FACTORIES AND NATIONAL INFRASTRUCTURE

You use the term "AI factory" frequently. What distinguishes an AI factory from a conventional data center?

Calista Redmond: A traditional data center is more of a house where you churn information. An AI factory manufactures intelligence. It generates intelligence you could not get to before. Data comes in, models and agents run at scale, and that generates useful digital output.

An AI factory is purpose-built and unifies five layers, what we call the five-layer cake: energy, chips, infrastructure, models, and applications. Governments often think about this as a public AI factory or an on-premises implementation covering government workloads, or as a local cloud that serves incubators, industry, SMEs, and startups across the jurisdiction.

It becomes critical infrastructure when it starts to support cross-sector workloads, for example in public services, cybersecurity, healthcare, manufacturing, agriculture, or transportation. If your country has big economic drivers in manufacturing, agriculture, or textiles, you want that to be part of what the government is supporting at this level. And it is optimized differently from a data center. You are fine-tuning for a specific workflow or model, focused on energy efficiency, and operating in a secure and continuous manner.

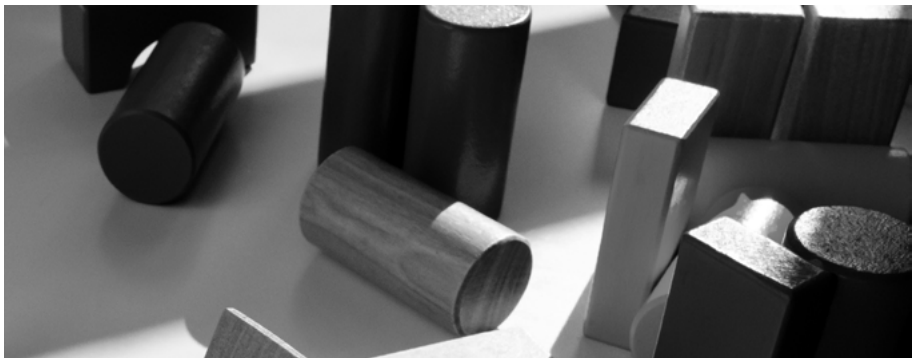

What role do open ecosystems play in building national AI capability?

Calista Redmond: Whenever I speak with a developer, a company, or a country or city, I always encourage collaboration. At this stage, no one should work alone, and no one should start with a blank sheet of paper. In just about every industry there is a running start with AI, an application to learn from and build on, and much of this is open.

Open ecosystems fundamentally reduce single-provider dependence. Local teams can inspect, adapt, benchmark, and improve their own work. They can take open models and implement their own language, laws, culture, and industries into them, which reflects national AI capabilities. These open models and tools support local language fine-tuning, RAG, academic research, and startup experimentation.

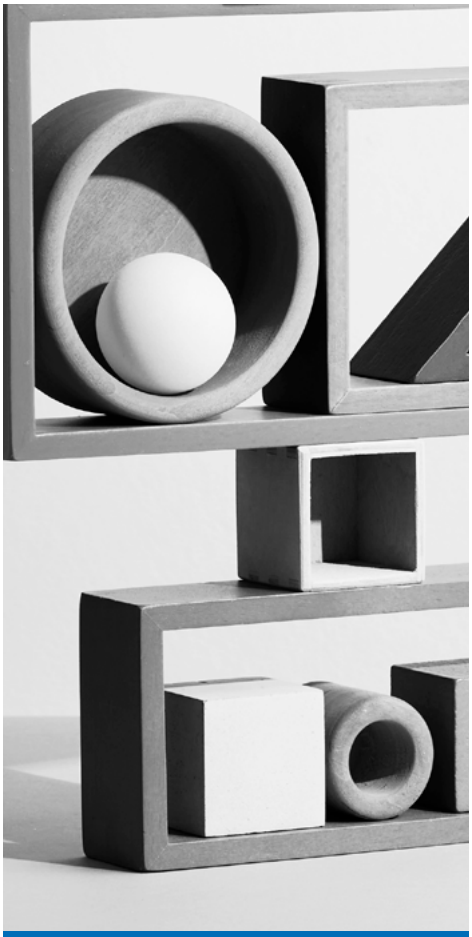

"Open ecosystems fundamentally reduce single-provider dependence"

Countries become technically stronger when they get to collaborate with peers, learn together, and leapfrog their own technical capacity building. I do want to be clear that open does not mean free-for-all or unmanaged. Countries still need to manage licensing, provenance, security scanning, and guardrails to ensure deployment reflects their intentions. At NVIDIA, we are big proponents of broad open ecosystems. Part of that is the principle of using the right model for the right workflow. You do not need a frontier model for everything, and tuning is critical for every workload.

COUNTRIES BECOME TECHNICALLY STRONGER WHEN THEY GET TO COLLABORATE WITH PEERS, LEARN TOGETHER, AND LEAPFROG THEIR OWN TECHNICAL CAPACITY BUILDING.

What are the biggest misconceptions you encounter when leaders think about national AI strategy?

Calista Redmond: The biggest misconception is that national AI capability means building everything alone or buying a single national model. The better framework is capability and control. Data, compute, and models work within a larger ecosystem that has local components and the strength of a global community.

"The biggest misconception is that national AI capability means building everything alone or buying a single national model"

The second misconception is that AI is only software. AI at national scale is infrastructure: energy, workforce development, procurement, and operating discipline. That is a fundamental thing to consider.

The third is the claim that governance slows innovation. In fact, the right governance accelerates adoption. It gives citizens, regulators, and others the confidence to use AI in high-value workflows. We all started using AI to plan a vacation, and now we are using it to instrument business decisions and outcomes. AI is not an application you are rolling out. It is a new production system, a new platform on which you can accelerate innovation and grow economies. That is more than an app.

If you had to distill your advice to a national AI policy leader into three priorities, what would they be?

Calista Redmond: First, choose three to five high-value missions. Not applications, missions. Unpack the datasets and identify use cases with measurable outcomes such as citizen services response times, fraud detection, or supply chain resilience. Build a collaboration roadmap that identifies the major stakeholders and maps the milestones you can celebrate and the outcomes you can anticipate.

Second, develop a secure AI factory plan. Understand the tiers, the hybrid approach, and the data controls, and start implementing. Without accessible compute, you are limiting innovation, utility, and workflows.

Third, build responsible AI into the operating model from the start. This includes model evaluation, guardrails, human oversight gates, and procurement rules that demand accountability. Think about what responsible AI means for talent, workforce, and the accountability measures that accompany rollout.

And if there is one overarching message, it is this: invest locally, engage globally.

Calista Redmond
VP Global AI Initiatives | Nvidia

"Invest locally, engage globally"

About Capgemini

Capgemini is an AI-powered global business and technology transformation partner, delivering tangible business value. We imagine the future of organizations and make it real with AI, technology and people. With our strong heritage of nearly 60 years, we are a responsible and diverse group of over 420,000 team members in more than 50 countries. We deliver end-to-end services and solutions with our deep industry expertise and strong partner ecosystem, leveraging our capabilities across strategy, technology, design, engineering and business operations. The Group reported 2025 global revenues of €22.5 billion.

Make it real | www.capgemini.com

About Capgemini Research Institute

The Capgemini Research Institute is Capgemini's in-house think tank on all things digital. The Institute publishes research on theimpact of digital technologies on large traditional businesses. The team draws on the worldwide network of Capgemini experts and works closely with academic and technology partners. The Institute has dedicated research centers in India, Singapore, the United Kingdom, and the United States. It was recently ranked number one in the world for the quality of its research by independent analysts.

https://www.capgemini.com/insights/research-institute/