Capgemini
RESEARCH INSTITUTE

Digital Sovereignty

The New Tech Order

Conversations FOR TOMORROW

N°11 — 2026

Executive Conversations

JAI HARIDAS

VP & GM, Regulated & Sovereign Cloud

Google

TUNING THE SOVEREIGNTY DIAL: INNOVATION, CHOICE, AND THE FUTURE OF CLOUD

Jai is a seasoned technology leader with over 28 years in the industry, specializing in cloud computing and the construction of scalable distributed systems for the past 18 years. His expertise spans storage, compute, databases, reliability, compliance, and security. Jai has a well-established track record of successfully growing and leading highly effective teams at renowned organizations such as Microsoft, Meta, and Google.

Since joining Google in 2016, Jai has served as Vice-President and General Manager of Regulated & Sovereign Platforms. In this role, he leads Google Cloud's Data Boundary and Dedicated Cloud Products, a strategic intersection of Cloud, Trusted Infrastructure, and Core. He is also responsible for ensuring Google Cloud Platform (GCP) aligns with global regulations and spearheads Trusted Infrastructure Services, including Confidential Compute along with Key and Secrets Management Systems.

SOVEREIGNTY AND INNOVATION

How would you define sovereignty? What is driving a growing interest?

Jai Haridas: At its core, digital sovereignty is about balancing access to technology innovation within a certain set of constraints and figuring out how to innovate within those constraints. The definition isn't fixed: what the market considers "sovereign" changes over time and varies by geography, industry, and use case.

The one constant is the need to innovate. In the AI era, organizations have to keep innovating to stay competitive, but they also need to meet sovereignty requirements. That is resulting in demand rising for solutions that enable both. The economic stakes make this concrete. Studies show that Europe is expecting roughly €1.2 trillion in economic growth, and most of that is expected to come from AI. But overly restrictive sovereignty-only approaches could reduce that to around two-thirds. That's one reason the topic has moved from cloud teams to the mainstream.

Three specific factors are driving the mainstream shift. The first is every nation's desire for autonomy, particularly over critical infrastructure: energy, finance, health. They want those systems to be autonomous and resilient, resilient against both man-made and natural disasters.

The second factor is that enterprises have shifted to a risk-first approach to

Jai Haridas
VP & GM, Regulated & Sovereign Cloud
Google

> **"Digital sovereignty is about balancing access to technology innovation within a certain set of constraints and figuring out how to innovate within those constraints."**

procurement, rather than the cost-first approach that used to dominate. Under a risk-first framework, they are looking at security, resiliency, and vendor lock-in as distinct risk categories.

The third factor ties back to the innovation imperative. Every business knows it cannot rely solely on internal innovation. They have to take dependencies on outside enterprises, and some of those enterprises are outside their nation. Sovereignty frameworks are how they manage the risk that comes with those external dependencies. But the biggest risk of all is not taking those dependencies, because then they are not innovating.

There is a widespread view that regulation slows innovation. Do you think that framing holds up when it comes to sovereignty?

Jai Haridas: I think the supposed trade-off between sovereignty and innovation is a false choice. Choose only sovereignty and you stop innovating; choose only innovation and you lose security and control. Either extreme can put the business at risk.

Organizations should be able to pick the level of sovereignty their use case requires and still innovate. That's why we make AI available across our sovereign options, from public cloud to dedicated to air-gapped environments, so customers can maintain portability and move at pace. The key is separating requirements from compromises. Encrypting data for a HIPAA workload isn't a compromise; it's a requirement. The compromise would be meeting the requirement but then losing access to AI.

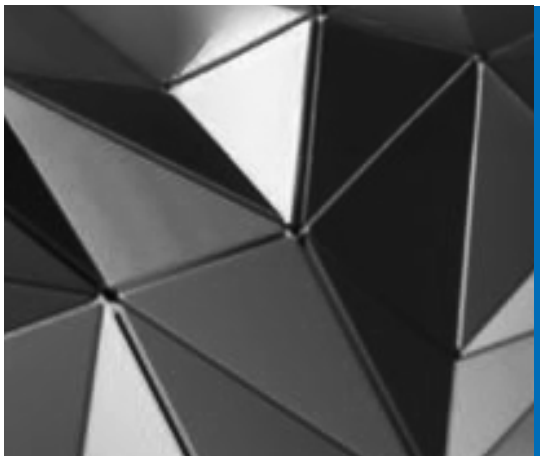

"Organizations should be able to pick the level of sovereignty their use case requires and still innovate."

We've seen this play out in practice. A US government agency wanted to adopt AI early but couldn't obtain GPUs from a traditional government cloud provider. With GCP [Google Cloud Platform], they were able to run IL5 workloads under FedRAMP in accredited public cloud regions with sufficient capacity, keeping compliance controls without giving up AI.

In France, an insurance company was already running in public cloud with local controls. When we built Trusted Cloud by S3NS, a SecNumCloud-compliant dedicated environment, it took roughly three years to design and implement. But when the customer finally migrated, the experience was so seamless that they mentioned,"Why did it take you three years just to change the DNS name?" For us, that was genuinely a compliment, because the work involved lifting an entire operational and technical machinery and making it run correctly inside a constrained, SecNumCloud-compliant environment. But from the customer's perspective, the migration was so seamless that nothing seemed to have changed. That portability is the power of choice.

Ultimately, this is an "and" proposition, not an "or" proposition. Customers often use multiple deployment models based on data classification and can port the same application across environments – preserving choice, control, and innovation.

There is a widespread view that regulation slows innovation. Do you think that framing holds up when it comes to sovereignty?

Jai Haridas: It is an important shift in perspective. Sovereignty often starts as a compliance conversation, but it can quickly become a competitive differentiator, because it's really about operational control and survivability.

Take Trusted Cloud by S3NS. We originally built it for France to offer a SecNumCloud-compliant solution, but now we're seeing European financial institutions ask for it even when no regulation requires them to. What they're buying is resilience. Some banks are designing their backup strategy

"Sovereignty often starts as a compliance conversation, but it can quickly become a competitive differentiator, because it's really about operational control and survivability."

around a sovereign setup: they keep running on public cloud, but if something happens, they fail over to a SecNumCloud instance. In that environment, S3NS is operating the cloud and is in control, and Google has no access to "push a button" and disable an account. That gives organizations complete operational control, and that is something serious financial institutions have started to value independent of regulatory mandates.

What we are also seeing is enterprises outside regulated industries looking at sovereign cloud as a privacy and competitive advantage. They are not approaching it as a compliance burden. They are recognizing that they can innovate much faster using a sovereign cloud capability without having to build a parallel platform themselves. And US companies are as interested in this as any European organization.

At the end of the day, the questions sovereign cloud answers are the same ones every CXO faces with their board: What's your disaster recovery and business continuity plan? Who can access data? Can it be exfiltrated? What are the controls, and how are they enforced? "Sovereignty" is just a label; the underlying need is universal and strategic.

""Sovereignty" is just a label; the underlying need is universal and strategic."

OPERATIONALIZING DIGITAL SOVEREIGNTY

When you have conversations with clients about sovereign cloud architecture, is the discussion primarily about technical controls or operational governance?

Jai Haridas: The honest answer is both, but the direction of the conversation has shifted in a meaningful way towards auditable technical controls.

> **"Security and sovereignty are two sides of the same coin."**

The most sophisticated customers – those who deeply understand the legal frameworks and have also experienced the shortcomings of solutions that rested purely on legal controls – are now asking explicitly for auditable technical controls. They have seen what happens when a sovereign solution is built entirely on contractual and legal mechanisms without the underlying technical enforcement, and they do not want to be in that position again.

There is also a broader principle at work: security and sovereignty are two sides of the same coin. You cannot have one without the other. And because security is no longer simply a matter of legality, the questions customers ask about security naturally lead them into the territory of technical controls. They want auditable controls. It is not enough for us to say a control exists. We need to demonstrate that it operates continuously, is actively monitored, and remains auditable.

As customers move beyond data sovereignty into operational sovereignty, survivability, and resiliency, they also need to address the legal dimension. SecNumCloud, for example, has over 300 technical controls. I have not seen a more stringent regulatory framework anywhere in the world. That comprehensiveness, both technical and legal, is what makes it a genuine benchmark.

In almost all cases, customers have arrived at technical controls either through a security-focused conversation or through direct experience of what solutions based purely on legal controls cannot protect them from.

How critical are partnerships to delivering sovereign cloud at scale?

Jai Haridas: There is a fundamental reason why organizations are seeking sovereignty in the first place, and a parallel reason why they trust local entities more than they trust anyone else. We can establish local operations, local legal entities, and local commitments, but the fact of who the parent company is remains. Partnerships with trusted local entities are how we address that trust deficit in a genuine and structural way.

In France, we partnered with S3NS - a subsidiary of Thales, for the dedicated cloud and the SecNumCloud offering, and have extended that partnership to the air-gapped cloud as well. Beyond France, we have partnered with Proximus and others. Our approach is partnership first, always.

These partnerships are win-win. Our partners bring perspectives on the local market, local risk, and local customer expectations that we simply do not have from our position as a US-headquartered company. Combining those two perspectives produces a much stronger solution than either party could offer alone, and one that local customers can trust more deeply.

The Thales partnership also has a very specific operational dimension. For our dedicated cloud, updates do not lag behind our commercial cloud. We run a five-day rollout cycle. When updates are released, we first send them to a quarantine environment where Thales runs their own tests to validate the software and verify there are no backdoor entries or vulnerabilities. Thales has complete autonomy in that environment and can stop rollouts if required. That means we are patching the dedicated cloud at the same pace as our commercial cloud, but with an independent third-party validation step that our sovereign customers require.

SIX MONTHS IN AI IS LIKE SEVEN DOG YEARS.

Customers want access to the latest AI models, but also need to remain sovereign. How do you see that tension playing out?

Jai Haridas: The pace of change in AI makes this urgent. In the past six months alone, the model leaderboard changed multiple times. Six months in AI is like seven dog years. If you are in a sovereign environment that is six months behind on AI innovation, you have missed multiple generational changes in capability. I have seen companies that are not using AI get left behind so rapidly that their backlogs are not even being reviewed while their competitors are already shipping features.

For customers choosing between sovereign environments, the first question should always be security: which environment or solution provides the strongest security for your specific workload? From there, they can work through their sovereignty requirements in detail. Where is data being processed? What data is leaving the jurisdiction? Who has access to it? Can it be exfiltrated? Is the most sensitive data encrypted? Can external key management systems be used to maintain client-side encryption while still benefiting from cloud services? And can air-gapped cloud environments be used to process highly classified data while still accessing AI capabilities?

That last point matters. We have customers running AI workloads in air-gapped environments. They are not compromising on AI capability by choosing air gapped. The choice of deployment model is driven by data classification and use case, not by a forced trade-off between security and innovation.

On the question of model lock-in specifically, if avoiding vendor lock-in is a requirement, there are open-source models available. Customers can use these models – from Gemma to a growing range of other open-source options

"For customers choosing between sovereign environments, the first question should always be security."

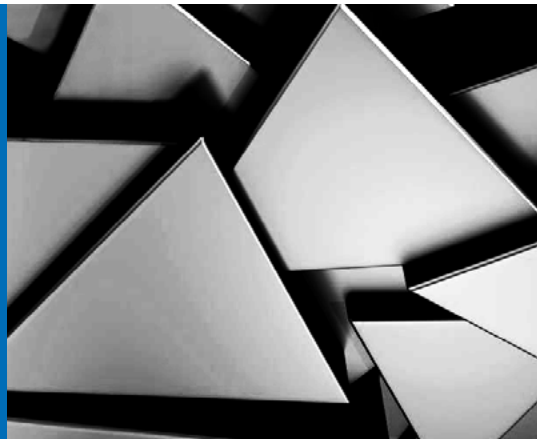

– without locking themselves into a single vendor's proprietary stack. Many customers today are running multiple models simultaneously behind the scenes, mixing open-source and proprietary depending on the workload.

The question customers should be asking is: which vendor gives me the flexibility to access AI capabilities without compromising on my sovereignty requirements, and which vendor has built sovereignty into the foundations rather than layered it on top?

THE CXO VIEW ON SOVEREIGNTY

What are the key questions a CXO should be asking when evaluating a sovereign cloud solution?

Jai Haridas: Many CXOs have resigned themselves to the belief that sovereignty necessarily limits their ability to innovate. I think that is a false, damaging choice. Sovereignty is becoming a baseline requirement. Geopolitical conditions are shifting faster than ever, and organizations need the most flexible sovereign cloud solutions available. The question is not whether to pursue sovereignty, but how to do so without sacrificing competitive position and innovation capacity.

With that framing in mind, here are the four things I would tell a CXO to focus on.

> "Sovereignty is becoming a baseline requirement."

First, evaluate which vendor best balances sovereignty, security, and innovation together. Not one over the others. All three simultaneously. That is the question they should be asking of every vendor in an evaluation. Do not accept a trade-off as inevitable.

Second, assess vendor lock-in carefully, and do not limit that assessment to software. Enterprises that have been through complex migrations understand that lock-in is not just about whether you are using SQL Server versus Spanner. It also operates through licensing mechanisms, and licensing lock-in can be just as constraining as technology lock-in.

Third, understand your own business and your own customer base deeply before you choose a vendor. Who are your customers? What data classifications apply to your various workloads? If the same application needs to serve both commercial sector customers and government agencies, you should not need to build three separate stacks. You should be asking which vendor enables a single stack with sufficient portability to operate across public cloud, dedicated cloud, and sovereign environments while meeting the requirements of each.

Fourth, distinguish between sovereignty requirements and what I call "emotional sovereignty." Decisions driven by fear or emotion rather than a careful, use-case-by-use-case risk analysis tend to over-constrain organizations, push them into overly restrictive solutions, and unnecessarily compromise innovation. The discipline of evaluating use cases, data classifications, and real risk before choosing a solution is what separates the organizations getting this right from those that are not.

"The question is not whether to pursue sovereignty, but how to do so without sacrificing competitive position and innovation capacity."

What are the biggest misconceptions you encounter among CXOs when it comes to sovereign cloud?

Jai Haridas: Two misconceptions come up consistently.

The first is about migration complexity. When organizations have built applications on a public cloud and need to move to a dedicated cloud environment, the assumption is almost always that this will require a massive uplift: rewriting the stack, rethinking the architecture, and significant cost and time.

The second and more pervasive misconception is that sovereignty is only a European concern. That has fundamentally changed. US companies are as interested in sovereign cloud as any European organization. At its core, sovereign cloud is about security controls: controlling data exfiltration, managing who has access to data, and ensuring business continuity and disaster recovery. These are the standard questions every CXO already answers to their board, regardless of geography. Sovereign cloud is a rigorous, structured approach to answering those questions. The framing as a European regulatory topic obscures the fact that the underlying concerns are universal.

THE FUTURE OF SOVEREIGN CLOUD

Looking five years out, how do you see sovereign cloud evolving?

Jai Haridas: Sovereignty will become a baseline requirement, embedded into every data center and every deployment. The geopolitical environment is changing faster than ever. Organizations need the most flexible sovereign cloud architecture available, and they need it as a foundation, not as an add-on.

More specifically, I think sovereignty will evolve from a discrete product choice into something more like a dial or a knob. Today, if a customer wants SecNumCloud-level compliance, they are moving to a dedicated cloud instance. That is a big architectural shift. What I would want to see, and what I think the industry is moving toward, is a model where sovereignty is a continuously adjustable parameter. A customer dials it to the minimum level that meets their current requirements. With geopolitical shifts, as regulation changes, or a risk event occurs, they dial it up. It should be as simple as that. The power of that kind of dynamic sovereign control would be genuinely transformative. We're not there yet. Today, organizations essentially choose a band: public cloud, dedicated cloud, or air gapped. But the direction of travel is toward something far more fluid.

If you had a magic wand, what would you change to the current approach to sovereignty?

Jai Haridas: If I had a magic wand, I would build a system where sovereignty and security levels are adjustable dynamically based on data classification and risk conditions, without requiring customers to rebuild or migrate their applications. That is very difficult to achieve – both for the customer in how they build applications, and for the cloud provider in how the underlying infrastructure is engineered. But that is the destination I would aim for.

The second thing I would change is regulatory clarity. Many customers today are not just paralyzed by having too many regulations; they are paralyzed by undefined or ever shifting regulations. They do not know what to build towards. They do not know what the next change is going to require. That uncertainty is genuinely costly for long-term architecture decisions. If we could give every jurisdiction a clear, well-defined, stable regulatory framework for digital sovereignty, organizations could make confident architectural choices and build towards a known target. That would accelerate the entire market.

MANY CUSTOMERS TODAY ARE NOT JUST PARALYZED BY HAVING TOO MANY REGULATIONS; THEY ARE PARALYZED BY UNDEFINED OR EVER SHIFTING REGULATIONS.

Jai Haridas
VP & GM, Regulated & Sovereign Cloud
Google

"The framing of sovereignty as a European regulatory topic obscures the fact that the underlying concerns are universal."

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/