Freeform Dynamics in association with Fsas Technologies a Fujitsu company

Sovereign AI: The Next Chapter

Time to look beyond risk and compliance to the strategic architectural advantage

Freeform Dynamics, 2026

Running AI applications on infrastructure that you own and control can be an effective way to get around the privacy, compliance and lock-in concerns associated with cloud-based AI platforms. But as the technology underpinning this 'Sovereign AI' approach has matured, a new set of benefits has emerged that are more concerned with business value, agility and competitive

KEY TAKEAWAYS

- **The risk and compliance benefits of Sovereign AI are now well known**
  Cloud-based AI services have many advantages, but they can also create data and compliance risks, and may lead to vendor lock-in. Sovereign AI is now broadly acknowledged as an option for mitigating such risks when the situation demands.

- **The strategic value-based benefits are less well understood**
  Sovereign AI enables unconstrained innovation, the ability to shape and right-size models for specific use cases, and the flexibility to adapt as requirements evolve - all of which translate to speed, agility and competitive advantage.

- **Unlocking the full benefits requires a focus on architecture**
  An effective AI strategy will embrace both cloud and sovereign elements. The key is to match platform choices to the needs of each workload rather than defaulting to one model for everything. Get the architecture right and the rest follows.

- **The right architecture becomes more critical as AI complexity grows**
  Greater componentization, blending of models and data sources, and complex workflow automation define the direction of travel. Each step up in sophistication amplifies the importance of sound platform and architectural decisions.

- **Pre-integrated system stacks can shortcut time to value**
  For many organizations, the fastest path to a well-architected Sovereign AI system is a pre-integrated technology stack, combining validated hardware, software and tooling to create a modern, open platform that's supported as a coherent whole.

Copyright 2026 Freeform Dynamics Ltd

Introduction

Generative AI burst into mainstream awareness through early cloud services from the likes of OpenAI, Google and Microsoft. Since then, the progress and antics of hyperscalers have dominated AI related news, keeping these players front-of-mind. Understandably, early enterprise adoption revolved around cloud platforms, and whether consciously or not, most ended up thinking of AI and cloud as going hand-in-hand.

In the early days, generative AI and the cloud naturally went hand in hand, but a range of questions emerged.

But as practical experience was gained, a more nuanced picture emerged. The large, general-purpose models underpinning cloud AI services came with a set of concerns that were difficult to engineer away. Training data of uncertain provenance generated output that raised questions about accuracy, bias and legal exposure. And the nature of the services themselves meant your data, your prompts, and your usage patterns were flowing through infrastructure you did not own, in jurisdictions you might not have chosen, under terms that could change without notice.

Enter Sovereign AI

Running AI on infrastructure you own and control - what the industry now broadly refers to as Sovereign AI - was the logical response. But in the early days, the tooling was immature, deployment experience was thin, and the practical effort involved in assembling a working on-premises AI environment was significant. If you went down this path at all, it was probably because you had to - because a particular workload involved sensitive data, or because regulatory constraints ruled out the cloud option.

The landscape has since shifted

Two key things have changed in the interim period. First, the perception of risk around cloud-hosted AI has, if anything, intensified. Major AI providers continue to resist calls for meaningful regulation and transparency, while geopolitical tensions have raised legitimate concerns about foreign government influence over critical infrastructure. High-profile data handling incidents, along with unpredictable costs, have then eroded trust in contractual assurances.

Thinking of Sovereign AI as a defensive or fallback option doesn't reflect the attributes of modern solutions.

The second big change is that the hardware, software, model ecosystem, and the body of practical experience needed to implement Sovereign AI has advanced considerably, creating a range of new benefits. Let's explore this further.

Copyright 2026 Freeform Dynamics Ltd

Sovereign AI is now a first-tier option

If your only reason for considering on-premises AI is to avoid the risks associated with cloud, you're only seeing half the picture. Modern sovereign AI platforms deliver a set of advantages that exist on their own merits, independent of any comparison with cloud. You can experiment freely with different models without worrying about consumption-based costs escalating. You can select and right-size models for specific use cases - e.g. a focused industry or domain-specific model rather than a sprawling general-purpose one - getting better results with fewer resources. And you can adapt your environment on your own terms, without waiting for a provider's roadmap or navigating contractual constraints.

The chart below illustrates this broader view, and you can see, both deployment models have genuine strengths, and neither is inherently superior. The right choice depends on the characteristics of the workload and what you are trying to achieve.

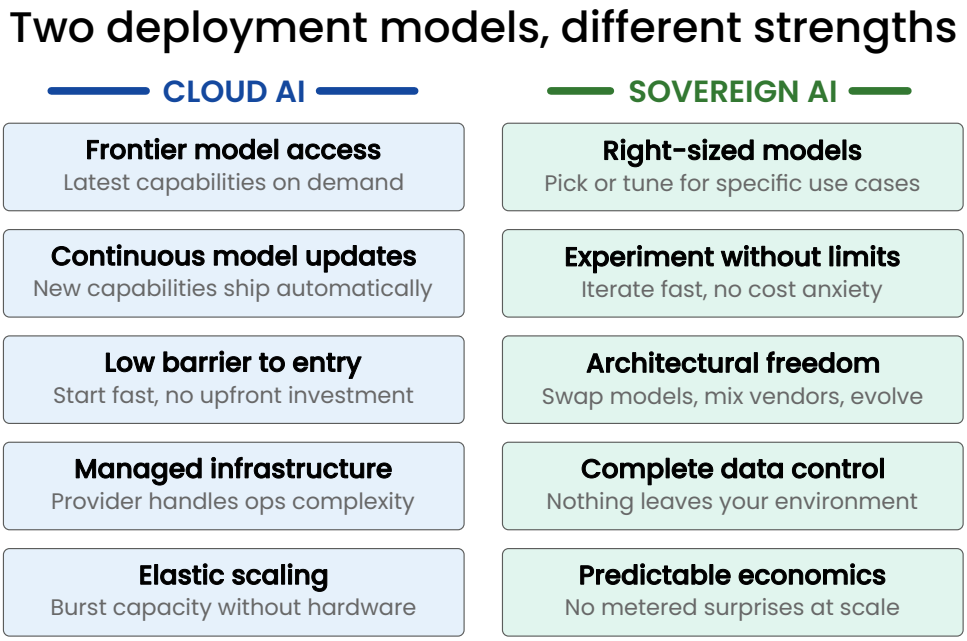

A case of 'and' not 'or'

For most organizations, the practical reality is that cloud AI and sovereign AI will coexist. Some workloads - rapid prototyping against frontier models, or burst capacity for occasional large-scale processing - are well suited to the cloud. Others - production applications involving sensitive data, workloads that demand predictable costs at scale, or use cases requiring specialized models tuned to your domain - are better served by a sovereign environment. The important thing is to make that choice deliberately for each workload, based on its specific requirements, rather than defaulting to one model for everything. This is fundamentally a question of architecture, and it is where the real strategic value lies.

Copyright 2026 Freeform Dynamics Ltd

Think architecture first

If there is one principle that experienced IT professionals recognize instinctively, it is that architecture is the foundation on which everything else depends. Get it right, and you can adapt, scale, and evolve as your requirements change. Get it wrong, and you accumulate technical debt that becomes progressively harder and more expensive to unwind.

The same principle applies to AI, and it is one that deserves more attention than it currently receives. A well-designed sovereign AI environment treats the LLM not as the centerpiece, but as one component in a broader stack - selected for the job at hand, much as you would choose the right database or orchestration engine for a given application. The architecture shown below illustrates the key layers and cross-cutting elements that a mature sovereign AI deployment needs to address.

Sovereign AI architecture

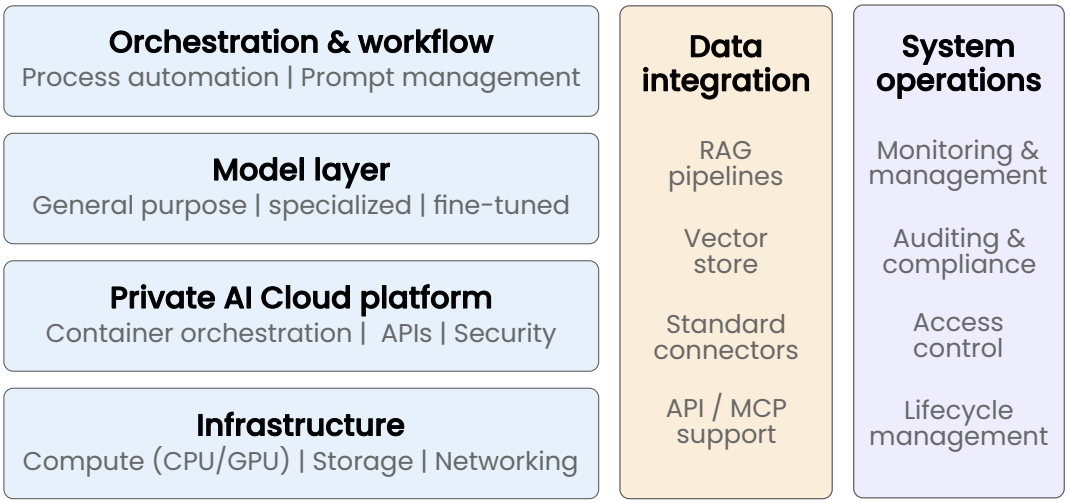

The critical design principle is openness. Each layer should support a range of options - different models, different data integration patterns, different orchestration approaches - connected through open APIs and standard interfaces. This is what gives you the freedom to select, tune and right-size your solution for each use case, swap components as better options emerge or requirements change, and avoid the kind of lock-in that comes from building on proprietary, tightly-coupled foundations.

This matters more with each step up in complexity. As your AI environment evolves - incorporating more models, integrating with more business data sources, and orchestrating increasingly complex workflows, every architectural decision you made early on gets amplified. An environment built on a solid, open foundation absorbs this evolution naturally. One built on expedient choices or proprietary dependencies will struggle, and the cost of retrofitting grows with every layer of sophistication you add - particularly important for anyone considering Agentic AI, for example.

Copyright 2026 Freeform Dynamics Ltd

From principle to practice

The architectural principles outlined so far are sound in theory, but translating them into a working Sovereign AI environment is not trivial. You need hardware capable of handling AI workloads efficiently, a platform layer that provides container orchestration and standard APIs, a curated selection of models, data integration tooling, and the operational instrumentation to run it all in production. Assembling these components from scratch, ensuring they work together reliably, and building the operational expertise to support them is a significant undertaking - particularly if your team's AI experience to date has been largely cloud-based.

This is where pre-integrated technology stacks earn their place. A well-designed stack brings together validated hardware, software and tooling into a coherent foundation that you can deploy with confidence and build on over time. The key is knowing what to look for as not all stacks are created equal, and the wrong choice can simply replace one form of lock-in with another.

Here are some key questions to ask when evaluating different options.

Platform evaluation checklist

1. **Complete, integrated stack?**
   Not a collection of parts. Components designed to work together.
2. **Model choice and flexibility?**
   Range of sizes and specializations. Open architectures. Freedom to swap.
3. **Sovereignty by design?**
   Architectural guarantees, not policy promises. Data physically cannot leave.
4. **Enterprise-grade operations?**
   Monitoring, auditing, access controls, traceability. POC to production.
5. **Realistic on-ramp?**
   Practical entry point without massive upfront commitment. Path to scale.
6. **Partner, not just supplier?**
   Use case discovery, architecture design, implementation support.

With these criteria in mind, it is useful to look at a real-world example to get a feel for what is currently available. We will do this with a tour of the sovereign AI portfolio from Fsas Technologies, a Fujitsu company and the sponsor of this paper. Please note that this is for illustrative purposes only and should not be taken as a supplier or solution recommendation by Freeform Dynamics.

Copyright 2026 Freeform Dynamics Ltd

Example of a real-world solution

Fsas Technologies, a Fujitsu company, has built its AI portfolio around the principles outlined in this paper. Rather than offering a single product, Fsas provides a range of pre-integrated solutions spanning the full stack - from validated server hardware and storage through to ready-to-deploy AI environments - designed to match different organizational needs and levels of ambition.

The portfolio includes turnkey private AI environments for organizations that want to get started quickly, with options ranging from single-site deployments to multi-location architectures where satellite servers communicate with a central system via queries rather than data transfer - maintaining sovereignty across distributed operations. For organizations wanting maximum architectural flexibility, the Private AI Cloud platform is built on Kubernetes with OpenAI-compatible APIs, meaning your existing development skills transfer directly. At the model layer, the portfolio supports a range of open source models alongside Fujitsu's own Takane LLM, with the freedom to fine-tune, substitute, and right-size for specific use cases. All of this runs on validated hardware - Fujitsu PRIMERGY servers for mainstream deployments and Supermicro GPU servers for more demanding workloads, plus Eternus and NetApp storage.

Fsas sovereign AI portfolio mapped to reference architecture

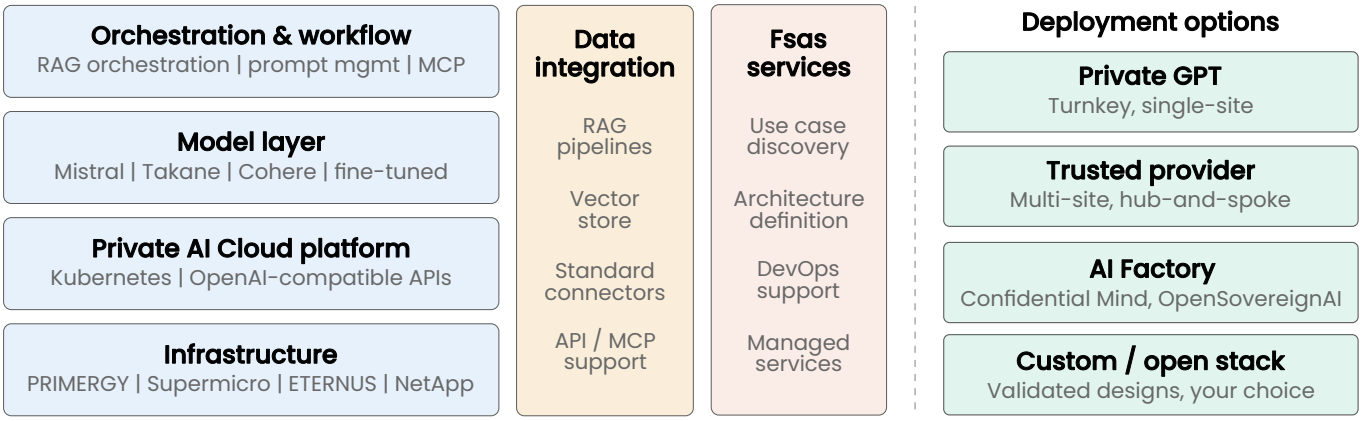

Importantly, Fsas approaches engagements as a partner in your architectural decisions, not just as a technology supplier. From initial use case discovery workshops through to implementation and ongoing support, the engagement model is designed to help you reach a working sovereign AI capability without the false starts that come from assembling components independently. A "fail fast" philosophy means that if your requirements are better served by a different approach, Fsas will tell you. The company says its emphasis is always on finding the right fit. If you have in-house AI expertise, the platform is open enough to customize and extend. If you need more hands-on guidance, the consulting and engineering capability is there to draw on.

Copyright 2026 Freeform Dynamics Ltd

Making it work for you

If there is one message we hope you take from this paper, it is that the deployment choices you make now for AI will have consequences that extend well beyond the immediate project. They will shape your cost structure, your flexibility, your ability to innovate, and your exposure to risks that are still evolving.

The good news is that the thinking required is not unfamiliar. If you have experience with cloud strategy, hybrid infrastructure, or any form of platform architecture, you already have most of the mental models you need. With this in mind, sustainable success with Sovereign AI platforms is largely about the application of sound architectural principles to a new and fast-moving technology area.

As a practical starting point, consider where your organization, or a particular use case, sits on each of the dimensions shown below. There are no right or wrong answers here - the point is to make the trade-offs visible and ensure your deployment choices reflect a deliberate strategy rather than an unexamined assumption.

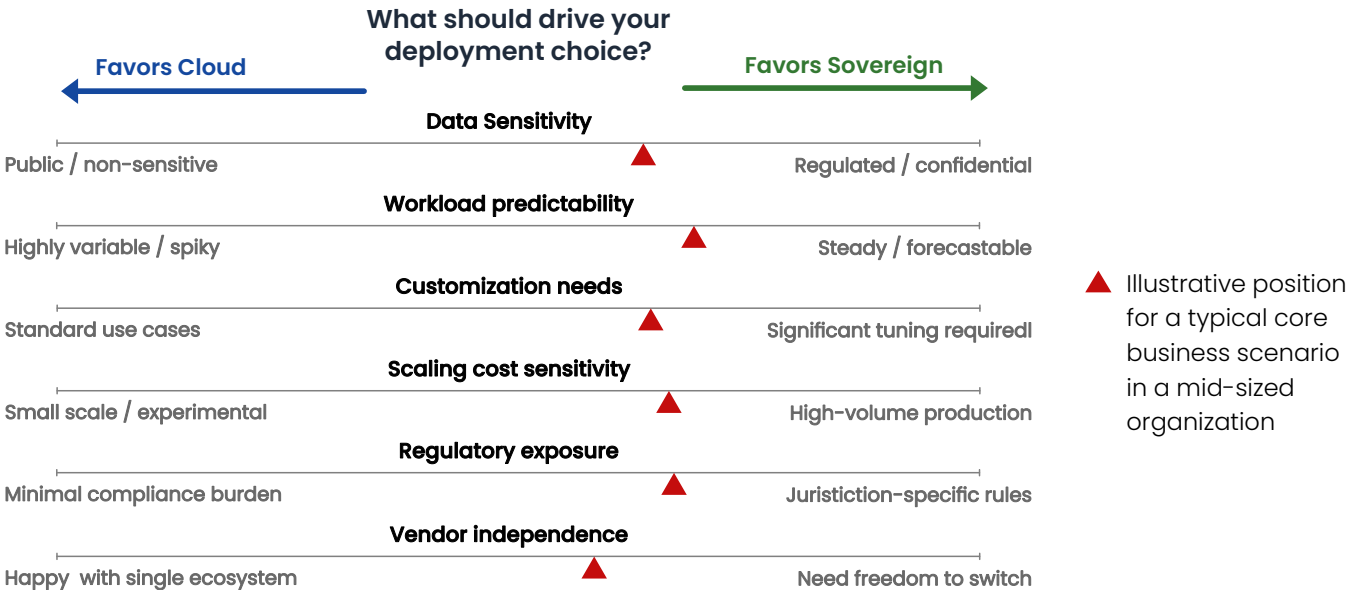

For most situations, the picture that emerges will be mixed. Some workloads will clearly be suited to cloud, some to Sovereign AI, and for others might depend on factors specific to your business. The goal is not to pick a side, but to build a foundation that lets you place each workload where its requirements are most naturally fulfilled. If you need some informed guidance, a discovery workshop with an experienced partner can help you map your priority use cases against the above dimensions.

From there, now the friction associated with Sovereign AI has been dramatically reduced, you can make deployment decisions on a deliberate and genuinely objective basis rather than passively adopting the historically assumed default. The bottom line is that the right decisions made now will pay dividends well into the future.

Copyright 2026 Freeform Dynamics Ltd

About Freeform Dynamics

Freeform Dynamics is an IT industry analyst firm. Through our research and insights, we help busy IT and business professionals get up to speed on the latest technology developments and make better-informed investment decisions.

For more information, please visit: www.freeformdynamics.com.

About Fsas Technologies

Fsas Technologies is a leading global provider of data centre infrastructure and advanced computing solutions, operating as a wholly owned subsidiary of the Fujitsu Group. Formed on 1 April 2024 by consolidating Fujitsu's server and storage business units, the company focus is on "sovereign" and secure AI, specifically designed for European regulatory and data privacy needs.

For more information, please visit: https://eu.fsastech.com/eu/.

Terms of use

This document is Copyright 2026 Freeform Dynamics Ltd. It may be freely duplicated and distributed in its entirety on an individual one to one basis, either electronically or in hard copy form. It may not, however, be disassembled or modified in any way as part of the duplication process. Hosting of the entire paper for download and/or mass distribution by any means is prohibited unless express permission is obtained from Freeform Dynamics Ltd or Fsas Technologies. The contents contained herein are provided for your general nformationand use only, and neither Freeform Dynamics Ltd nor any third party provide any warranty or guarantee as to its suitability for any particular purpose.