Cloud Sovereignty Framework

This three-level framework with eight criteria is an analytical model synthesized from the core principles and definitions used by European regulatory bodies, industry analysts, and the cloud providers themselves. It is designed to create a consistent and measurable "scorecard" to compare different sovereignty strategies.

Layer A [1,2]: Foundational Layers (Where is the data?)

This layer addresses the most fundamental questions of sovereignty: the physical and legal location of data.

[1] Data Residency

Data residency refers to the physical, geographic location where an organization's data is stored at rest. This is the most basic level of sovereignty control. An organization achieves data residency by choosing to deploy its applications and store its data in cloud data centers located within a specific country or region, such as Germany or the EU. This decision is often driven by the need to reduce latency for local users, but it is also the first step toward meeting regulatory requirements. All major cloud providers allow customers to select a specific region (e.g., Frankfurt, Paris, Dublin) for their primary data storage.

[2] Data Localization

Data localization is a stricter, legally mandated version of data residency. It refers to laws that require data generated within a country's borders to be initially collected, processed, and stored within that same country. While data residency is often a choice, data localization is a legal obligation. This concept is critical because it often extends beyond just customer content to include metadata (such as system logs, user permissions, and resource tags), which can be challenging for globally architected cloud platforms to keep entirely within one country's borders.

Layer B [3,4,5]: Control Layers (Who can access and manage the data and systems?)

This layer moves beyond the physical location of data to address the human and technical controls governing access to that data and the systems that run it.

[3] Data Access Sovereignty

This level involves the technical and organizational controls that limit access to data from outside a specified jurisdiction, particularly by the cloud provider's own personnel. The goal is to ensure that no one, including a cloud administrator, can access customer data without explicit, auditable authorization from the customer. This is often achieved through technical means, such as Google's Key Access Justifications (KAJ), which requires a valid reason for every access request, or the hardware-level isolation of the AWS Nitro System, which is designed to prevent any operator access to data while it's being processed.

[4] Operational Sovereignty

Operational sovereignty focuses on who manages and operates the cloud infrastructure. It requires that all personnel involved—including technical support, system administrators, and on-call engineers—are located within a specific jurisdiction and are subject to its laws. This control is designed to mitigate the risk of foreign legal compulsion or unauthorized access through global "follow-the-sun" support models, where staff in different countries might have privileged access to systems. Examples include Google's "EU Data Boundary and Support" package, which guarantees support from EU-based personnel, and the announced AWS European Sovereign Cloud, which promises that all operations and support will be handled exclusively by EU residents.

[5] Control Sovereignty

This is the highest level of technical control, where the customer uses cryptographic technologies to gain ultimate authority over their data, making it verifiably inaccessible to the cloud provider. This is achieved through two primary mechanisms:

- **Customer-Controlled Keys:** Using services like an External Key Manager (EKM or XKS), customers can store and manage their encryption keys in a system completely outside of the cloud provider's infrastructure. The cloud provider never has access to the key material itself, only a reference to it, giving the customer a "kill switch" to revoke access at any time.
- **Confidential Computing:** This technology protects data *while it is in use* (i.e., being processed in memory). It uses a hardware-based Trusted Execution Environment (TEE), or "enclave," to create an isolated, encrypted space where code and data are protected from being viewed or modified by anyone, including the cloud provider's hypervisor or operators.

Layer C [6,7,8]: Assurance Layers (What are the surrounding guarantees?)

This final layer assesses the broader architectural, legal, and compliance frameworks that support a provider's sovereignty claims.

[6] Legal/Contractual Posture

This refers to the specific legal agreements, data processing addenda, and public commitments a provider offers. It is the legal foundation of their sovereignty promise. This includes contractual terms that define data boundaries, commitments to challenge government data requests, and addenda that help customers comply with regulations like GDPR and address the requirements of legal rulings such as Schrems II. Examples include Microsoft's broad EU Data Boundary commitment and the specific service terms for Google's Assured Workloads.

[7] Isolation Model

This describes the fundamental architectural approach a provider takes to deliver sovereignty. There are two main models:

- **Configured Sovereignty on a Shared Cloud:** This model uses the provider's standard global infrastructure but applies software-defined controls, policies, and specific operational boundaries to create a sovereign environment. This is the approach taken by Google's Assured Workloads and Microsoft's EU Data Boundary. It offers broad service availability but relies on logical separation.
- **Dedicated/Isolated Sovereign Cloud:** This model involves building a physically and operationally separate, independent cloud for a specific jurisdiction. This is the approach of the announced AWS European Sovereign Cloud. It promises the highest level of isolation but may initially launch with a more limited set of services.

[8] Ecosystem/Compliance

This level evaluates a provider's alignment with recognized European and national compliance standards, as well as the strength of its partner ecosystem. Third-party certifications, such as Germany's BSI C5 (Cloud Computing Compliance Criteria Catalogue), provide independent validation of a provider's security claims. Furthermore, a mature ecosystem of local partners, like T-Systems in Germany for Google Cloud, is often essential for implementing and managing high-assurance sovereign solutions. This also includes a provider's engagement with pan-European initiatives like GAIA-X.