uptime INTELLIGENCE

INTELLIGENCE UPDATE

What cloud sovereignty really means

Dr. Owen Rogers 15 Jan 2026

In 2025, the close bonds between the US and Europe faltered, leading some European enterprises to question their continued use of cloud services provided by US hyperscalers (see *Tariff tensions undermine trust in cloud hyperscalers*). To address concerns about legal authority over workloads and data, several US cloud providers rolled out "sovereign public cloud" products aimed at ensuring data and operations remain within European borders. Concurrently, European cloud providers have been marketing themselves as geopolitically neutral alternatives.

However, the features of sovereign cloud offerings differ significantly among providers. Some seek to secure sovereignty through commitments to data residency, while others adopt organizational frameworks to minimize the risk of unauthorized access. These discrepancies pose challenges for businesses trying to discern which products truly fulfill their needs and which amount to little more than the clever marketing of standard cloud services.

This report defines six levels of sovereignty that clarify the differences between these products and their associated risks. While Europe has emerged as a focal point for cloud sovereignty, similar debates are taking place globally. Accordingly, this report focuses on Europe, but the framework it presents is applicable to other national and regional contexts.

European context

It is important to situate Uptime Intelligence's model within the broader European regulatory discussion on cloud sovereignty. The EU is developing a Cloud Sovereignty Framework that defines sovereignty through a set of eight policy objectives and evaluates compliance using minimum assurance criteria. The framework is intended to support harmonized regulation, certification and procurement across EU member states. It focuses on the presence and maturity of specific controls rather than on a single hierarchy of outcomes.

This report provides a complementary, risk-based view to explain how different combinations of legal authority, operational control and technology ownership translate into materially different sovereignty outcomes. Together, the two approaches address different needs: the EU framework defines how sovereignty is assessed, while Uptime Intelligence's model explains why these distinctions matter to enterprises and policymakers.

Overview

Here, we define six levels of cloud sovereignty for public cloud. Level 0 represents the highest level of protection against unauthorized access and the lowest reliance on foreign powers, while Level 4 is on the opposite end of the spectrum. The six levels are summarized in **Figure 1**.

Figure 1 The six levels of cloud sovereignty for public cloud

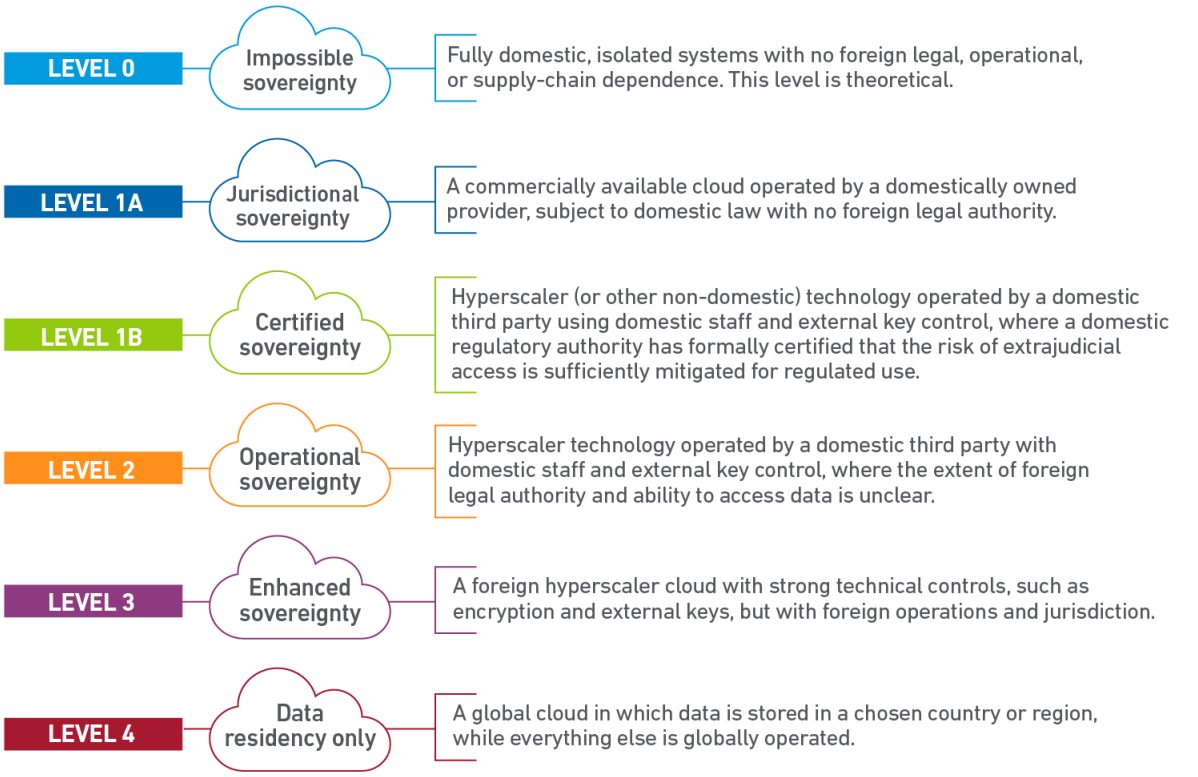

Table 1 defines these levels of sovereignty in terms of the following characteristics:

- **Provider ownership:** Which courts can claim authority over the cloud provider?
- **Legal jurisdiction:** Which laws govern access to the data stored in the cloud? (See *In cloud and colo, whose laws rule the data?*)
- **Operations and staffing:** Where are staff located and under which laws can they be compelled to act?
- **Control plane:** Where is the control plane located, and which governments could potentially compel access to or control over it? The control plane provides management and orchestration across a cloud environment. While the operator of the control plane cannot necessarily see data, it can control where it resides and how it is accessed. The operator of the control plane may also have visibility into metadata on confidential data that may reveal patterns or relationships between entities. (See *Cloud resiliency: plan to lose control of your planes*)
- **Key management:** Where are access keys located, and which governments could potentially

compel access to those keys? (See *Key players: cloud control and the colo advantage*)

- **Supply chain dependency:** Which nation states provide the technology supporting the platform, and therefore have some control over IP and delivery?

Table 1 Sovereignty level breakdown

	← Increasing sovereignty / Decreasing flexibility					
Level	**0**	**1A**	**1B**	**2**	**3**	**4**
Description	**Impossible sovereignty**	**Jurisdictional sovereignty**	**Certified sovereignty**	**Operational sovereignty**	**Enhanced sovereignty**	**Data residency only**
Provider ownership	Domestic	Domestic	Foreign platform, domestic operator	Foreign platform, domestic operator	Foreign platform, operated domestically by a subsidiary	Foreign
Legal jurisdiction	Domestic	Domestic	Domestic, as formally assessed by regulator	Unclear — foreign law may apply	Foreign law may apply	Foreign law may apply
Operations and staffing	Domestic	Domestic	Domestic	Domestic	Domestic	Foreign
Control plane	Domestic	Domestic	Foreign IP, domestically operated	Foreign IP, domestically operated	Foreign platform, domestically operated	Foreign platform
Key management	Domestic	Domestic	External keys held by the domestic operator	External keys held by the domestic operator	External keys as an option	External keys as an option
Data and metadata location	Domestic	Domestic	Domestic	Domestic	Domestic	Data domestic, metadata may be foreign
Supply chain dependence	None — domestic	Some — unavoidable	Some — unavoidable	**Medium** (foreign tech, domestic operation)	**High** (same stack, more controls)	**High** (global stack)

UPTIME INSTITUTE 2026

For simplicity, the term "domestic" is used to describe situations in which a given characteristic is being controlled within the same legal jurisdiction as the data itself. This usually means the EU, the UK and/or Switzerland. Conversely, "foreign" is used to describe scenarios involving extrajudicial access or control. Note: this analysis does not include private clouds or on-premises appliances offered by public cloud providers.

In general, an increase in sovereignty is usually associated with an increase in management overhead and a reduction in flexibility. At Level 4, enterprises have access to the widest choice of services and locations for their applications, delivered using a single supplier and management interface. At Level 0, enterprises have fewer capabilities available to them — to operate in multiple jurisdictions, they must manage multiple cloud environments.

Level 0: Impossible sovereignty

Level 0 represents a theoretical level of sovereignty, used here as a device to demonstrate that IT

systems are, in practice, always reliant on many nations for manufacturing and operations, and may therefore be subject to access by foreign governments and law enforcement agencies. At this level, the public cloud is implemented, maintained and operated with no foreign legal, technical or supply-chain influence. Achieving this would require fully national ownership of hardware, firmware, software, cryptography, operations and intellectual property. While some governments and defense organizations — notably in the US, China and Russia — may achieve this level of sovereignty, it is impossible for enterprises.

Even if such a level of sovereignty were feasible, unauthorized access could still occur. Many governments have reciprocal agreements with other governments and may cooperate with foreign law enforcement bodies in accessing data subject to these agreements. In addition, foreign intelligence services may attempt extrajudicial access even in the absence of formal agreements.

Level 0 shows that enterprises must accept that, in a globalized world, they inevitably rely upon non-domestic products or services at some point in the supply chain. A nation state may seek to disrupt this supply chain or embed backdoors in technology to enable covert access to data. State-sponsored actors have been implicated in major supply chain compromises, including suspected backdoors in Juniper firewalls discovered in 2015 and the SolarWinds software compromise in 2020, which has been widely attributed to Russian intelligence services. Even with rigorous planning and robust supply-chain controls, such risks cannot be entirely eliminated.

Level 1A: Jurisdictional sovereignty

Level 1 represents the highest level of sovereignty achievable for most enterprises and is divided into two types: A and B.

In Level 1A, the cloud provider is domestically owned and governed, operates exclusively under domestic law, and runs a domestically controlled control plane with EU-based staff and operations. Data and metadata remain within the jurisdiction, and external courts cannot compel a foreign parent company to intervene. While some reliance on global hardware and software supply chains remains unavoidable, Level 1A removes foreign legal authority from those supply chains and provides strong independence for critical, highly regulated workloads. A Level 1A provider may further bolster its credentials through appropriate regulatory certification.

However, the use of a European provider does not automatically protect customers from foreign access. In 2024, Canadian authorities sought to compel OVHcloud to provide data relating to a European citizen, arguing that OVHcloud's corporate presence in Canada, which includes a subsidiary and data center, brought the company within Canadian jurisdiction, despite the data being hosted in Europe. OVHcloud appealed the request, and a legal challenge remains ongoing. This demonstrates that even the use of a Level 1A sovereign cloud does not guarantee that data cannot be accessed by every jurisdiction, for now and in the future.

Public clouds operating at Level 1A include:

- OVHcloud (France).
- Scaleway (France).
- Exoscale (Switzerland).
- IONOS Cloud (Germany).
- T-Systems Sovereign Cloud (Germany).
- Aruba Cloud (Italy).
- Elastx (Sweden).
- Civo (UK).

Some nations have created their own public clouds primarily for governmental use, owned and operated entirely by government bodies. These also fall into category 1A:

- GI Cloud MeghRaj Initiative (India).
- Qatar Cloud.

Fully government-owned and operated public cloud platforms are rare, as they require the state to assume the cost, operational complexity, and innovation burden typically borne by commercial hyperscalers.

Level 1B: Certified sovereignty

Level 1B describes cloud services where sovereignty is established through formal regulatory certification, rather than through domestic ownership of the underlying technology stack (as in Level 1A). In this model, hyperscaler technology is used, but the cloud service is operated by a domestic provider, with domestic staff, domestic operational control and exclusive control of encryption keys. The service is certified by a national authority as sufficiently protected against unauthorized foreign access for regulated use.

This model provides sovereignty in a regulatory sense, even if it is not structurally sovereign.

In the event of access by an extraterritorial government body, the enterprise may be able to argue that it acted in accordance with regulator-approved controls and obligations, and therefore should not be held responsible for the breach. Certification can potentially shift some responsibility away from the enterprise and onto the cloud provider, regulator and national government.

Level 1B public cloud offerings today include:

- S3NS PREMI3NS (Thales with Google Cloud, ANSSI SecNumCloud certification, France).

Governments are likely to favor certification-based approaches because they help protect sensitive data, avoid the cost and complexity of building and operating national cloud platforms, and allow responsibility for security and operations to remain with hyperscalers or other third-

party providers.

Level 2: Operational sovereignty

Level 2 describes cloud environments that use hyperscaler technology and are operated by a domestic partner but without formal regulatory certification establishing legal sovereignty. These environments typically implement strong operational and technical controls — including domestic operations, restricted access and external or customer-managed encryption keys — to reduce the likelihood of foreign access. Data and metadata are hosted domestically, and access may be limited to locally authorized personnel.

Despite these safeguards, Level 2 services remain subject to foreign legal authority because the service has not yet been formally recognized by a domestic regulator as sovereign. As a result, residual extraterritorial legal exposure persists, even where access is practically constrained. While hyperscalers and their partners may make contractual and public commitments limiting access to customer data, the enforceability of such commitments under conflicting legal obligations has not been conclusively tested across all jurisdictions.

The fundamental difference for enterprises between Level 1B and Level 2 lies in legal protection. In a Level 1B public cloud, if access by foreign governments or agencies were to occur, the enterprise can demonstrate that it selected a cloud service formally approved by a national authority for regulated use. This provides a strong regulatory defense, as the enterprise has complied with applicable legal and supervisory requirements, shifting scrutiny away from its own due diligence and toward the certified service and its governance framework.

Level 2 public cloud offerings include:

- Bleu (Capgemini and Orange with Microsoft technology, France) — pending certification.

Level 3: Enhanced sovereignty

Level 3 introduces operational and cryptographic controls designed to reduce, but not eliminate, the risk of foreign interference. Infrastructure is typically operated by EU-based staff or trusted domestic partners, and customers often retain substantial control over encryption keys through external key mechanisms. These measures significantly limit the cloud provider's routine access to customer data and improve auditability and governance. However, the underlying platform remains foreign-owned, and foreign laws may still apply to the provider.

Level 3 is therefore about risk reduction and regulatory alignment, not full legal sovereignty.

Public cloud offerings at Level 3 include:

- AWS European Sovereign Cloud (planned).

- Microsoft Sovereign Public Cloud.
- Oracle EU Sovereign Cloud.

Level 4: Data-only sovereignty

Level 4 sovereignty is limited to control over where data is stored, typically by selecting a specific country or regional data center. While this supports basic data residency requirements, the cloud provider remains foreign-owned, operates a global control plane and is subject to foreign laws. Operational access, metadata processing and support functions may occur outside the chosen region, and foreign authorities can compel the provider to disclose data or metadata.

Level 4 can help meet EU data residency requirements, which can support GDPR compliance, but it offers little protection against extraterritorial legal or operational access.

Cloud provider offerings operating at Level 4 include:

- AWS (standard EU regions).
- Microsoft Azure (standard EU regions).
- Google Cloud Platform (standard EU regions).
- Oracle Cloud Infrastructure (standard EU regions).
- IBM Cloud (standard EU regions).
- Alibaba Cloud (standard EU regions).

> ## The Uptime Intelligence View
>
> Sovereignty is increasingly being used in cloud providers' marketing materials to reassure enterprises that data is secure and compliant with regulations. Enterprises should be cautious about taking such claims at face value, as ultimate responsibility remains with them. Understanding the differences between public cloud offerings is the first step in assessing which venue is best for each workload. Enterprises should also be realistic in setting their goals: not all risks can be mitigated against, and not all workloads will require the highest level of sovereignty. Most organizations will choose a combination of private facilities, colocation providers, hyperscalers and smaller cloud providers that best suits their needs. Further guidance on venue selection will be provided in a future report.

ABOUT THE AUTHOR

Dr. Owen Rogers

15 Jan 2026

Dr. Owen Rogers is Uptime Institute's Senior Research Director of Cloud Computing. Dr. Rogers has been analyzing the economics of cloud for over a decade as a chartered engineer, product manager and industry analyst. Rogers covers all areas of cloud, including AI, FinOps, sustainability, hybrid infrastructure and quantum computing.

orogers@uptimeinstitute.com

About Uptime Institute

Uptime Institute is the Global Digital Infrastructure Authority. Its Tier Standard is the IT industry's most trusted and adopted global standard for the proper design, construction, and operation of data centers – the backbone of the digital economy. For over 25 years, the company has served as the standard for data center reliability, sustainability, and efficiency, providing customers assurance that their digital infrastructure can perform at a level that is consistent with their business needs across a wide array of operating conditions.

With its data center Tier Standard & Certifications, Management & Operations reviews, broad range of related risk and performance assessments, and accredited educational curriculum completed by over 10,000 data center professionals, Uptime Institute has helped thousands of companies, in over 100 countries to optimize critical IT assets while managing costs, resources, and efficiency.