Computer Fraud and Security
ISSN (online): 1873-7056

# Sovereign Cloud Buildout Using Repeatable and Scalable Patterns

**Surya Narayana Lankalapalli**

Microsoft Corporation, USA

**Abstract**

As governments, regulated industries, and critical infrastructure providers accelerate digital transformation, the need for sovereign cloud architectures has become unavoidable. Data residency, legal jurisdiction, operational autonomy, and national security concerns now shape cloud strategy as much as performance or cost. Traditional cloud deployments, often centralized, globally managed, and dependent on foreign operators, struggle to meet these sovereignty requirements. A modern sovereign cloud must guarantee that data, operations, and governance remain under the control of a specific nation or jurisdiction. To achieve these objectives consistently across different agencies, tasks, and areas, it takes more than custom engineering—it needs reliable and scalable designs that make sovereignty a core part of how cloud operations work. This article explores sovereign cloud buildout through the lens of architectural design, research innovations, real-world applications, industry-specific use cases, and the evolving role of AI-driven collaboration.

**Keywords:** Data sovereignty, sovereign cloud, edge computing, regulatory compliance, federated governance

## 1. Introduction

### 1.1 Background and Motivation

The rapid pace of digital transformation in both the public and private sectors has changed how countries think about data, infrastructure, and who controls technology. Governments, healthcare organizations, financial regulators, and defense agencies are moving more and more important workloads to the cloud. However, this move raises important questions about who really controls the data, the infrastructure, and the laws that govern them. The concept of sovereign cloud has emerged in response to these concerns, representing a paradigm shift from convenience-driven cloud adoption to sovereignty-conscious architectural design. Sovereign clouds are designed to keep data and operations under the control of a specific country or regulatory body, unlike traditional cloud deployments that focus on cost and scale.

### 1.2 Problem Statement

Traditional cloud architectures, built for global reach and operational efficiency, did not consider national sovereignty. Centralized control planes are often managed by foreign entities in traditional cloud architectures, rendering them inherently incompatible with the legal and regulatory demands of sovereign governments. Extraterritorial laws, such as the US CLOUD Act, create jurisdictional conflicts that undermine the data residency guarantees required by many nations [9]. Four important laws make this problem worse: the GDPR in the EU, the CCPA in the United States, and the PDPA in Singapore and China each have different rules about how data should be stored, processed, and shared across borders, which no typical cloud system can meet all at once. The lack of standard designs for sovereign systems means that organizations must create custom solutions for every new setup, leading to inconsistencies, higher costs, and regulatory risk [1][2].

### 1.3 Scope and Objectives

This article addresses the architectural, technical, and governance dimensions of sovereign cloud build-out, with a focus on repeatable and scalable patterns that can be applied across diverse jurisdictions and industries. It examines the foundational pillars of sovereignty, explores recent research innovations, reviews real-world deployments, and considers the role of artificial intelligence in automating sovereign operations. The objective is to provide a comprehensive scholarly overview that bridges theoretical frameworks with practical implementation guidance, offering actionable insights for cloud architects, policymakers, and enterprise technology leaders [3][4].

### 1.4 Definition of Key Concepts

Several distinct yet interrelated concepts comprise the sovereign cloud. Data sovereignty is the idea that data is subject to the laws and governance structures of the country where it is stored, such that foreign authorities cannot access or extract it without clear legal permission. Operational sovereignty refers to the restriction of cloud operations—including support, monitoring, and administrative access—to personnel and systems within a given jurisdiction. Technical sovereignty means that only the authorized parties control encryption keys, identity systems, and security policies. Legal sovereignty provides immunity from foreign extraterritorial laws, ensuring that the cloud environment operates exclusively under the legal framework of the host nation. Together, these four pillars define the full scope of what it means to build and operate a truly sovereign cloud [2][3].

## 2. Architectural Foundations of Sovereign Cloud

### 2.1 Core Pillars and Design Philosophy

The architecture of a sovereign cloud is not defined by any single technology or product but by a coherent set of design principles that collectively enforce control, compliance, and autonomy. At the foundation lies the recognition that sovereignty must be a built-in property of the cloud environment rather than an overlay applied after deployment. This means that there needs to be a clear design approach that focuses on control based on laws and regulations at every level, from the physical hardware and network setup to how identities are managed and data is secured. The hardware root of trust backs this up: in trusted Telco Cloud environments, the ETSI NFV reference architecture is split into three main parts: the MANO, NFVI, and VNF layers. At each of these layers, trust needs to be built separately using Trusted Platform Module (TPM) chips that store cryptographic keys and check system components. A security problem at any one layer can lead to issues with keeping workloads separate or shutting them down, suggesting that it's time for strong security measures to be built into every level of the system instead of relying on manual checks.

| Architecture Layer | Role in Sovereign Cloud | Trust Enforcement Mechanism |
|---|---|---|
| Management and Orchestration (MANO) | Coordinates provisioning, lifecycle management, and policy enforcement across the cloud environment | The security orchestration component (TSecO) verifies VNF integrity and manages binding policies before workload launch |
| Network Function Virtualization Infrastructure (NFVI) | Provides the underlying hardware, storage, and network resources on which sovereign workloads execute | TPM chips store cryptographic hash measurements; remote attestation servers continuously verify platform integrity state |
| Virtualized Network Functions (VNF) | Delivers virtualized sovereign services such as firewalls, identity functions, and lawful intercept capabilities | VNF image integrity verified via hash digest and signing authority before launch; binding policies restrict VNFs to trusted platforms within sovereign jurisdictions |

Table 1: ETSI NFV Architecture Layers and Trust Components [7]

### 2.2 Sovereign Landing Zones

Within a sovereign cloud environment, a sovereign landing zone acts as a repeatable blueprint for the deployment of all workloads. It encapsulates the governance, security, and compliance controls required by a given jurisdiction and ensures that every new workload automatically inherits these controls upon deployment. Key parts include different network designs that prevent unauthorized access, separate control systems that restrict admin access to approved local staff, local identity providers, and key management systems that keep encryption keys within the sovereign area. The landing zone incorporates policy-as-code frameworks that automatically apply legal requirements, transforming laws into consistent rules across all deployments. Three well-known certification standards are used as benchmarks to make sure that sovereign landing zones meet clear and checkable regulatory requirements. These are ISO 27001 for managing

information security, ISO 27701 for managing privacy information, and Germany's C5 standard for cloud computing security.

### 2.3 Modularity and Scalability

One of the most significant advances in sovereign cloud design is the move toward modularity, which allows sovereign components to be assembled, reused, and expanded for different agencies and workloads without requiring custom engineering for each deployment. Modular sovereign services—like sovereign identity, sovereign key vaults, and sovereign monitoring pipelines—can be created, approved, and used on their own as parts of a bigger sovereign system. This modularity is essential for scalability, as it allows a single sovereign pattern to be instantiated across dozens of agencies or regional deployments while maintaining consistent governance and compliance. The 3-layer ETSI NFV architectural structure supports this modular approach even more by clearly separating the management, infrastructure, and workload tiers. Each tier can be governed separately, but they all work together to enforce sovereignty [7].

## 3. Innovations Enabling Scalable Sovereign Architectures

### 3.1 Edge Computing and Distributed Sovereignty

The rise of edge computing has changed how sovereign cloud architecture is conceptualized, moving the control boundary from just centralized data centers to include distributed edge nodes located at the edges of the network. Edge environments present unique sovereignty challenges, as data is generated, processed, and stored across geographically dispersed locations that may span multiple jurisdictions. To keep data sovereignty at the edge, it's important to set up local rules, such as managing keys on-site, ensuring automatic compliance, and routing data based on location, all without relying on central control systems. A complete checklist for keeping data local in edge deployments covers 7 different areas—legal and compliance research, hosting operations, engineering and administration, operations, maintenance, and related governance—highlighting that managing data sovereignty at the edge involves many aspects that need to be considered throughout the entire deployment process, not just at one technical level.

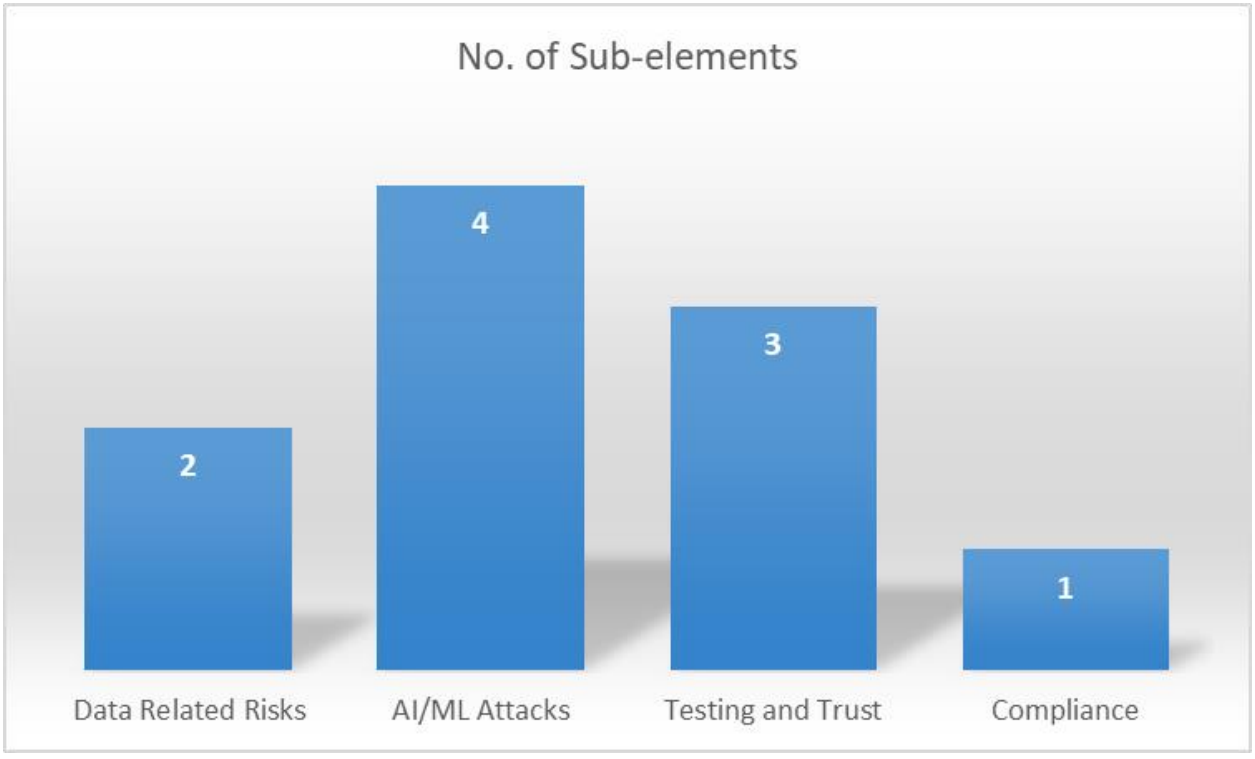


Fig. 1: AI Risk Categories [8][9]

### 3.2 Data Localization and Federated Governance

Data localization—the requirement to store and process certain categories of data within national borders—represents one of the most operationally complex aspects of sovereign cloud implementation. Organizations operating across multiple jurisdictions must navigate a patchwork of data localization regulations that differ in scope, enforcement, and

technical interpretation [5][6]. Three major national sovereign cloud projects show how different countries are putting this requirement into practice: the EU's Gaia-X project aims to create a federated, interoperable cloud infrastructure based on European data protection values; China's Cybersecurity Law says that sensitive data must stay in China with strict access controls; and India's Personal Data Protection Bill sets strict localization rules for certain types of sensitive data while allowing limited cross-border flows under certain conditions [9]. Federated governance models have emerged as an effective solution, enabling national authorities to set fundamental sovereignty rules while allowing different agencies or cloud regions to operate independently, sharing governance responsibilities and reducing the likelihood of failure.

### 3.3 Policy-as-Code and Automated Compliance

One of the most important changes in sovereign cloud architecture is the use of policy-as-code frameworks to automate compliance. Policy-as-code integrates regulatory requirements directly into deployment pipelines, ensuring constant and consistent enforcement of compliance across all workloads. Such an approach is better than relying on periodic manual audits or static configuration checklists [5][6]. Three main issues in data protection indicate the importance of this automation: increasing cybersecurity threats that change quicker than manual defenses can keep up with, the difficulty of managing compliance with different regulations at the same time, and the ongoing risk of insider threats that need constant monitoring instead of just occasional checks. When these three challenges are handled using policy as code instead of manual methods, compliance is transformed from a one-time task into an ongoing requirement that adapts as regulations change.

## 4. Industry-Specific Applications and Real-World Deployments

### 4.1 Government and Public Sector

Government agencies represent the most demanding sovereign cloud consumers, requiring strict operational sovereignty, immutable audit trails, and guarantees that citizen data never leaves national borders. National digital services—including digital identity platforms, tax administration systems, and public health registries—rely on sovereign landing zones to ensure that sensitive citizen data is processed and stored exclusively within the jurisdiction [7][8]. A key requirement for government systems is to sort tasks into three different trust levels: tasks that don't need trusted resources, tasks that should use trusted resources with some flexibility, and tasks that must use trusted resources without exception. If a trusted resource is unavailable in the hard-trust classification, the system halts the work rather than shifting it to an untrusted environment. This is essential in government situations where the security of citizen data has legal and constitutional importance.

| Sector | Primary Sovereignty Requirement | Key Compliance Frameworks | Trust Model | Certification Standards |
|---|---|---|---|---|
| Government | Operational sovereignty, immutable audit trails | GDPR, national data laws | Hard-trust / Soft-trust / No-trust | ISO 27001, C5 |
| Healthcare | Data residency, cross-border research compliance | HIPAA, GDPR | Soft-trust | ISO 27001, ISO 27701 |
| Financial Services | Transaction auditability, local key management | GDPR, CCPA, Basel III | Hard-trust | ISO 27001, ISO 27701, C5 |
| Defense | Air-gapped deployments, cryptographic sovereignty | National security law | Hard-trust only | ISO 27001, C5 |

Table 2: Industry Sector Sovereign Cloud Requirements and Applicable Standards [7][8][9]

### 4.2 Healthcare and Critical Infrastructure

Healthcare organizations face a dual mandate that makes sovereign cloud architecture particularly critical: they must protect patient privacy and comply with stringent data residency regulations while simultaneously enabling secure cross-border collaboration for research and clinical purposes. Sovereign cloud patterns address this by establishing clear

boundaries around where patient data can be processed and stored while permitting controlled data sharing across borders under defined compliance rules [8][9]. Critical infrastructure sectors—including energy, transportation, and telecommunications—similarly require sovereign cloud environments that maintain operational autonomy while benefiting from the elasticity and resilience of cloud computing. In healthcare, sovereign cloud services support genomics research and global medical teamwork, requiring that sensitive genomic data be stored according to strict laws while still allowing for international research.

Rich partnerships ensure compliance with the three standards of ISO 27001, ISO 27701, and the C5 framework.

### 4.3 Financial Services and Defense

Financial institutions and defense organizations represent the two most stringent sovereign cloud use cases, each demanding the highest levels of assurance, auditability, and technical control. Banks and payment networks must simultaneously contend with 3 primary categories of AI-related risk in sovereign multi-cloud environments—data-related risks, AI and ML attacks including data poisoning and model We need to reduce risks from inversion, testing, and trust compliance by using sovereign cloud systems that follow Zero Trust security, fully encrypt data for storage, transfer, and processing, and implement role-based access control. Defense and national security settings need even stricter controls, such as systems that are completely isolated or partially separated, cryptographic systems managed locally, and identity frameworks that are specific to the country. In these settings, the 3-layer ETSI NFV architecture is important: trust needs to be maintained separately in the MANO, NFVI, and VNF layers, and if trust is broken in any layer, the system should stop the workload in a controlled way instead of moving it to keep the sovereign boundary safe.

## 5. AI and Human-AI Collaboration in Sovereign Cloud Operations

### 5.1 AI-Driven Compliance and Anomaly Detection

Artificial intelligence is rapidly becoming an indispensable component of sovereign cloud operations, enabling organizations to enforce compliance and detect anomalies at a scale and speed that exceeds human capability. Machine learning models that analyze data from sovereign cloud systems can quickly spot when policies are not being followed, unusual access behaviors, and incorrectly set up resources much faster than older monitoring methods. Three categories of AI-related risk—data-related risks, AI and ML attacks, and testing and trust compliance issues—define the threat landscape that AI-driven compliance tools must address. A three-part strategy that uses Zero Trust security, complete encryption, and federated learning forms the technical basis for achieving these objectives while keeping sensitive data within the sovereign area. This capability is especially critical in sovereign cloud environments, where a single compliance failure can have significant legal, regulatory, and national security consequences [9].

| AI Risk Category | Risk Description | Mitigation Strategy | Sovereignty Mechanism |
|---|---|---|---|
| Data-Related Risks | Unauthorized data access, cross-border exposure | Data localization, access controls | Geofencing, local KMS |
| AI and ML Attacks | Data poisoning, model inversion, adversarial inputs | Zero Trust security, encryption | End-to-end encryption across 3 stages |
| Testing and Trust Compliance | Policy non-compliance, lack of transparency | Federated learning, explainable AI | Sovereign boundary enforcement |

Table 3: AI Risk Categories and Mitigation Strategies in Sovereign Multi-Cloud Environments [8]

### 5.2 Automated Remediation and Workload Isolation

In addition to finding problems, AI-driven automation allows sovereign cloud environments to fix compliance issues automatically, without needing people to step in, which greatly speeds up the time it It takes time to resolve difficulties and reduces the risk of a sovereignty breach. Automated remediation workflows can quickly respond to issues by fixing misconfigured resources, changing compromised encryption keys, isolating non-compliant workloads, and creating regulatory incident reports—all within seconds of spotting a problem. A structured AI compliance automation pipeline includes 6 steps: checking current compliance processes, looking into AI technologies, setting goals for using AI, creating a data plan, testing AI solutions, and training staff on AI tools before expanding successful projects. This 6-step

process makes sure that automated fixes are put in place in an organized and reliable way, instead of just reacting randomly, and that the technology behind them works completely within the set limits.

### 5.3 Human-AI Collaboration and Societal Impact

While AI automation significantly improves the operational efficiency of sovereign cloud environments, the human element still plays a crucial role in maintaining sovereignty and accountability. This involves providing ethical oversight and exercising contextual judgment. Human-AI collaboration models see AI as a helpful tool that lightens the workload for human operators, allowing them to concentrate on important strategic and investigative tasks while AI takes care of regular monitoring and problem-solving. The five key components of data rights—the location of data storage, its mobility, the explanation of its use, who can access it, and its quantity—establish the limits that human operators must monitor, while AI manages the enforcement tasks beneath them. This 5-aspect framework makes sure that automation doesn't take over completely and that people are still responsible for their actions. This concept is important in government situations where decisions about data access, managing workloads, and responding to incidents have legal and national security consequences that need human approval. At a societal level, the development of sovereign cloud systems supported by AI automation and human supervision significantly affects a country's control over its digital space, how much citizens trust government online services, and the nation's overall ability to defend against cyber threats.

## Conclusion

Sovereign cloud development has changed from being a specific compliance need to a common practice, based on clear methods that can be repeated, expanded, and adjusted to meet the needs of different industries and regions. The main ideas of data, operations, technology, and legal control create a clear guide for building cloud systems that stay completely under national authority, while new advancements in edge computing, shared governance, and policy-as-code allow for maintaining control on a large scale across various types of infrastructure. The basic structure includes a 3-layer ETSI NFV architecture, follows 4 main global regulatory frameworks, manages workloads in 3 trust categories, is validated by 3 internationally recognized certification standards, and has a checklist for data localization with 7 components, a framework for data rights with 5 aspects, and a 6-stage pipeline for AI compliance automation, all showing that creating a sovereign cloud is now a repeatable process using real-world examples from government, healthcare, finance, and defense. These demonstrate that sovereignty and innovation are not mutually exclusive—when built on scalable patterns, they reinforce each other. The use of AI for compliance automation and teamwork between humans and AI makes the sovereign cloud better at keeping up with rules, even as regulations change and cloud systems become more complicated. In the future, creating flexible sovereign systems, adjustable compliance rules, and operations supported by AI will be crucial to making sure that the sovereign cloud stays a strong and reliable part of the country's digital plans.

## References

[1] Vasileios Karagiannis, et al., "Data Sovereignty at the Edge of the Network," in 2023 IEEE 7th International Conference on Fog and Edge Computing (ICFEC), 29 January 2024. Available: https://ieeexplore.ieee.org/document/10403098

[2] Christian Esposito, et al., "On Data Sovereignty in Cloud-Based Computation Offloading for Smart Cities Applications," IEEE Internet of Things Journal, Volume: 6, Issue: 3, 12 December 2018. Available: https://ieeexplore.ieee.org/document/8573800

[3] Christian Esposito, et al., "Encryption-Based Solution for Data Sovereignty in Federated Clouds," IEEE Cloud Computing, Volume: 3, Issue: 1, 26 February 2016. Available: https://ieeexplore.ieee.org/document/7420528

[4] Vasileios Karagiannis, "Data Sovereignty and Compliance in the Computing Continuum," in International Conference on Future Internet of Things and Cloud (FiCloud), IEEE Xplore, 08 November 2024. Available: https://ieeexplore.ieee.org/document/10743062

[5] Jaganmohan Reddy Kancharla, et al., "Advancing Data Sovereignty in Distributed Environments: An In-Depth Exploration of Data Localization Challenges," in 2024 International Conference on Computer, Electronics, Electrical Engineering & their Applications (IC2E3), 10 January 2025. Available: https://ieeexplore.ieee.org/document/10827688

[6] Eric B. Blancaflor, et al., "Comparative Analysis of Current Infrastructure of Cloud Computing Unveiling the Trends of Industry 5.0," in 2024 International Conference on Computer, Electronics, Electrical Engineering & their Applications (IC2E3), 16 January 2025. Available: https://ieeexplore.ieee.org/document/10838120

[7] Ian Oliver, et al., "Experiences in Trusted Cloud Computing," Journal of ICT Standardization, IEEE Xplore, 2018. Available: https://ieeexplore.ieee.org/document/10258094

[8] Various Academic Authors, "AI and Multi-Cloud Compliance: Safeguarding Data Sovereignty," IRE Journals, 2024. Available: https://www.irejournals.com/formatedpaper/1705421.pdf

[9] Adedamola Abiodun Solanke, "Sovereign Cloud Implementation: Technical Architectures for Data Residency and Regulatory Compliance," International Journal of Science and Research Archive, April 2024. Available: https://ijsra.net/sites/default/files/IJSRA-2024-0502.pdf