Balancing digital sovereignty with security and reliability

Overcoming common tensions that complicate sovereignty-driven IT and security procurement

The landscape of IT and security procurement is undergoing a significant shift. For years, effectiveness, operational efficiency, and cost were the most common considerations. Today, a fourth pillar is becoming equally important: sovereignty.

This shift has a variety of overlapping motivations. Some organizations want to ensure security and IT service continuity in the face of hypothetical future geopolitical disruption. Others are concerned about protecting data from foreign government requests. Others still want to support their local technological ecosystem, viewing sovereignty through the lens of economic growth and resilience.

As understandable as these motivations are, 'building sovereignty into our digital environment' is a more difficult task than it might seem. Sovereignty is not something you can buy. It's a broad and constantly evolving concept, and little settled regulation exists to define what it should look like in different technical contexts. In addition, certain interpretations of digital and data sovereignty can cause tradeoffs between other important goals like security, resiliency, and operational efficiency.

This article examines those tradeoffs and offers strategies technology leaders can use to find the right sovereignty strategy for their organization.

Key takeaways

- Sovereignty initiatives can create difficult tradeoffs with security and resilience
- The right architecture can give organizations avoid these hard choices

Tension 1: Scale and threat intelligence

Companies often prioritize sovereignty to protect service availability. Unfortunately, using localized sovereign services in order to prevent disruption by foreign governments can increase the risk of disruption by cyber attackers.

The interconnected nature of the public Internet means companies in any part of the world face cyber attacks from all over. This is why it's important for cybersecurity services to draw on a diverse, global body of threat intelligence. If, for example, a new zero-day exploit or malicious IP address is identified in one region, that intelligence can be instantly applied to protect companies everywhere.

A large mitigation capacity is also important for DDoS attacks in particular. Throughout 2025, Cloudflare network data revealed an average 33% increase in hyper-volumetric DDoS attacks every three months.

Unfortunately, those two qualities may be harder for locally isolated security infrastructure to achieve. Localized infrastructure, by definition, has a narrower field of vision, and is less likely to be able to scale up to match the largest global DDoS botnets. How do organizations balance the risk

of attack-related disruption against that of disruption by foreign governments?

Tension 2: Customer and user experiences

A similar tradeoff can emerge outside of a cyber defense context. When security and IT services are too regionally isolated, latency and limited network capacity can slow digital experiences: for customers, and for internal users trying to access the network.

Even companies operating in a single country may have issues when customers travel, or during sudden traffic surges. A large French ehealth service faced the latter challenge: when they were mentioned in a series of government health communications, traffic to their web application spiked as high as 10x. As a European company, they faced strict limitations about where traffic could be inspected. But no local app delivery providers were able to handle that much traffic, leaving them with a seemingly impossible choice.

Tension 3: Visibility and efficiency

Limiting security and IT purchasing decisions to in-region vendors can hurt operational efficiency and visibility for larger organizations in particular. A company headquartered in one country may have operations across many. If they maintain a separate set of security and IT services for each of those jurisdictions for sovereignty reasons, the complexity of managing and interlinking all of them can quickly spiral out of control. Each local provider likely has its own API, its own logging format, and its own management interface — a significant burden for the average IT department, which is only projected to grow by 1.3% in 2026.

Like the others, this tension has real implications for data security and resilience. When an incident occurs, the time required to correlate logs from three or four different regional providers can be the difference between a minor localized issue and a major breach.

Strategies for navigating sovereignty tensions

The complexity and fluidity of the sovereignty landscape means no organization will ever fully avoid these tensions. Still, strategies exist to help minimize them:

1. Define and prioritize sovereignty risk across your security and IT stack

Anyone with security experience knows that you can never completely eliminate every risk — you need to prioritize them, and reduce each type using the resources available. Sovereignty is no different. Rather than applying the same broad sovereignty rules for any new IT or security service purchase, start by listing specific sovereignty risks across different services, and then prioritize how they affect your to overall availability.

For example: in a web application context, is it most important where traffic decryption and inspection happens? Where customer logs are stored, or how those logs are backed up? Or, in a network access context: does it matter more where internal applications live, or where access policies are enforced? By defining these priorities clearly, you can make more informed decisions — and more clearly articulate to other stakeholders why you've made them.

2. Understand vendors' programmability and customizability

In some circumstances, a blanket "only in my home region" approach to IT and security procurement may be painting with too broad of a brush. Looking for providers who let you customize many of the specific sovereignty qualities mentioned in strategy #1. Doing so may let you control where data lives, and who can access it, without sacrificing security, performance, or operational efficiency.

For example: does it matter which network locations your encrypted traffic passes through, or which ones your encrypted data lives on, if you can customize where decryption takes place and where those keys are stored? Or, in an application security context, is it possible to ensure that traffic decryption and processing takes place only within your jurisdiction while still using a larger network to mitigate DDoS attacks? Often — though admittedly not always — such flexibility will meet your specific sovereignty needs.

3. Explore interoperability and failover options

As mentioned, localized 'sovereign' security and IT providers can have unintended resiliency problems. Rather than make a difficult 'global vs local' decision, consider a multi-provider approach for key services, in which global and local providers share the load in some way. That way, you get the benefit of a larger local provider while still being prepared for the hypothetical future possibility of it being compelled to turn off services.

This approach looks different, and works better or worse, in different use cases — another reason to do the prioritization exercise described in step one. For example, a company using a global ZTNA service may want to be sure that provider allows for easy transfer onto

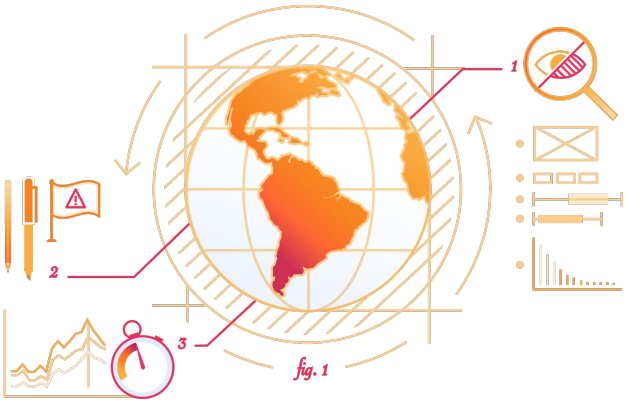

fig. 1

a local 'shadow VPN.' Another company may use a primary/secondary DNS model for a similar reason, or split a percentage of traffic across global and local DDoS mitigation providers. Such an approach has benefits outside of a sovereignty context, since many regulations (e.g. the European Union's Network and Information Security 2 Directive and Digital Operational Resilience Act) have strong business continuity requirements.

4. Understand vendors' approaches to legal transparency

Technical controls are of course only half of the sovereignty equation. The other half is trust: in a hypothetical geopolitical conflict, will your security and IT vendors prioritize their customers' data rights and service continuity?

No one can predict the future. But it's worth investigating any vendor's prior track record about similar situations. How transparent are they about prior government requests? What is their policy for handling those requests? Are there other, non-geopolitical circumstances in which they have been asked to shut off services for certain customers? How strongly did they maintain their stance, even when public opinion was against them?

A thorough review of these issues can make the trust question feel more concrete and comparable — an important step in an area as fraught as sovereignty.

Cloudflare can help resolve sovereignty tensions

As mentioned, sovereignty is a principle and a set of related practices, not something you can buy. Still, Cloudflare can help organizations exercise their own inherent sovereignty while avoiding the aforementioned tensions. The technical controls supporting our security, connectivity, and developer services — along with our governance posture — help organizations move from a “sovereignty vs. security/resilience" mindset to a “sovereignty and security/resilience” mindset.

- **Technical controls:** Customers can constrain where decryption and L7 inspection occur, how TLS key custody is handled where required, and where certain customer logs are processed and exported.
- **Governance posture:** We publish transparency reporting about government requests and have a track record of challenging unreasonable legal demands.
- **Resiliency and interoperability:** We support multi-provider and multi-environment architectures so customers can design continuity and exit options aligned to their risk posture.

Learn more about our sovereignty approach here, or talk to our team to get answers to your specific questions.