

The Depth Test — Why Most Sovereignty Claims Fail Structurally

SCENARIO PLANNING METHODOLOGY SERIES

Mike Edholm · edholms.com

Sovereignty is being asserted more loudly now than at any point in thirty years, and the word is doing almost no structural work. This essay proposes a test. Sovereignty is not a legal property but a positional one: it depends on which layer of a technology stack a claim actually reaches, how hard that layer is to replace, and whether the holder is free to act on it. Four questions apply the test. Two documented cases show it working, one where a government pulled the lever it held and one where a government has never pulled it, with outcomes that run opposite to intuition. The test explains why most sovereignty claims fail structurally even when the ambition behind them is correct, and it suggests that for most countries the achievable goal is not sovereignty at all.

I.

Sovereignty has returned to the center of political language, and it is being used to describe an unusually wide range of things.

A government negotiating tariffs invokes it. A regulator proposing rules for social media platforms invokes it. A continent announcing an investment program for semiconductors invokes it. A minister defending requirements for product labeling in a national language invokes it. The word carries real weight in each case, and in each case it means something different.

That variety is not a problem of vocabulary. It is a symptom. Sovereignty in its ordinary legal sense — a state has jurisdiction, or it does not — has become a poor guide to what any particular country can actually decide. The legal question still has answers, and the answers are still mostly clear. They simply no longer predict outcomes.

What follows proposes a different way to ask the question, borrowed from how technology systems are built rather than from how states are described. It is a test, in four questions, and it can be run on any sovereignty claim before the arguing starts. Most claims fail it. Some fail in ways that matter and some in ways that do not, and telling those apart is most of the value.

II.

Start with why the legal framing explains so little.

The parties that determine what a country can do are frequently not states. They are frequently not located in the country whose choices they constrain. And they are frequently not doing anything a lawyer would recognize as governing.

A company with a four-decade lead in optical engineering constrains more national choices than most treaties. An insurance mutual disciplines an industry no legislature reaches, by declining to write cover. A currency's settlement system reaches into transactions between two parties who have never set foot in the country that issues it, and neither can route around it at acceptable cost. None of these is an exercise of jurisdiction. All of them determine outcomes more reliably than jurisdiction does.

The useful question is therefore not who has authority. It is who has to pass through you, and what does it cost them to go around. That question can be asked of a state, a company, a standards body, or a payment system, and it produces answers that can be compared across all four. The legal framing cannot do that, because it can only tell you whether an entity is a state.

Asking the positional question changes what counts as evidence. Statutes stop being the primary material and market structure becomes it: how many suppliers exist at a given layer, how long a replacement takes to build, whether an adequate substitute exists at all. Those things are readable from public data before anybody denies anybody anything, which means the test can be run in advance rather than after the fact.

III.

Three terms carry the argument, and each needs defining precisely, because the rest of it uses them precisely.

A **stack** is a set of layers, each depending on the ones below it. Chips sit under hardware.

Hardware sits under operating systems. Operating systems sit under applications. The pattern is not confined to computing: in finance, payment interfaces sit above clearing systems, which sit above settlement in a currency. In shipping, freight brokerage sits above vessels, which sit above the classification societies that certify them. Every industry with infrastructure has a stack, and the layers are ordered by dependence rather than by importance.

Depth is how far down a layer sits and how hard it is to replace. Three readable properties make a layer deep: few suppliers, long replacement times, and no adequate substitute. A layer with dozens of interchangeable suppliers is shallow no matter how much revenue passes through it. A layer with one supplier and a fifteen-year replacement horizon is deep even if hardly anyone outside the industry can name the company. Depth determines consequence — the deeper the layer at which a party is denied, the more severe and persistent the damage, and the less any workaround higher up compensates.

The money stack

One dollar payment between two non-American banks, and what each layer would take to relocate.

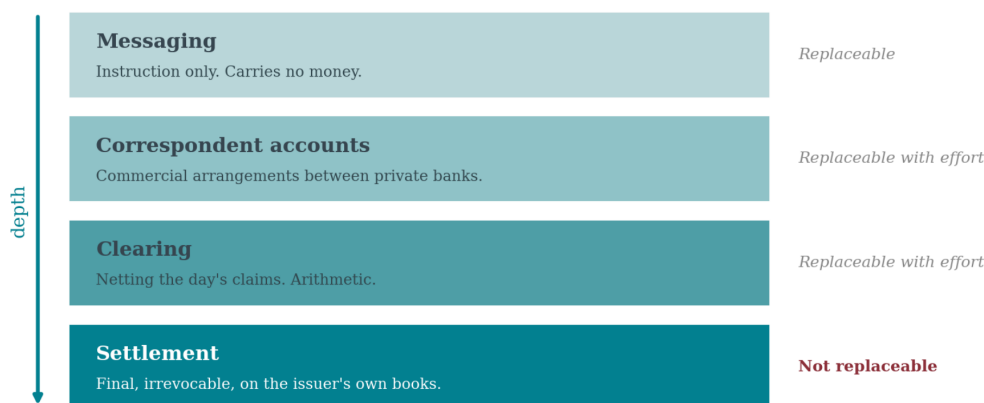


Figure 1. Depth is readable before anything is denied. One payment, four layers, and only the bottom one has no substitute while the currency in use is the dollar.

A **channel** is what makes a rule bind when it is not law. Somebody must be able to make non-compliance expensive, and whatever that somebody controls is what participants have become dependent on. Channels rank the same way layers do: by how many parties must pass through them, and what it costs to route around.

None of these terms is new. Treating a stack's layers as the organizing structure for sovereignty comes from Benjamin Bratton. The conditions under which a network position becomes leverage come from Henry Farrell and Abraham Newman, whose work on weaponized interdependence supplies the third question below. What is proposed here is only that the terms be used as a test, applied before a claim is evaluated rather than after it has been argued about.

IV.

The test is four questions. They run in order, and a claim that fails an early one does not need the later ones.

First: which layer does the claim actually reach? Not which layer it invokes rhetorically — which one it touches. A rule governing what platforms may show reaches the application layer. A rule governing how a product is labeled reaches the application layer. A subsidy for domestic fabrication plants reaches a considerably deeper one. The distinction is not about ambition or seriousness. It is about position in the stack, and the answer is usually obvious once the question is asked.

Second: how deep is that layer? Count suppliers. Estimate replacement time. Ask whether an adequate substitute exists. A claim reaching a layer with many suppliers and short replacement times is a claim about something that can be routed around, and it will be routed around if routing around becomes worth the cost.

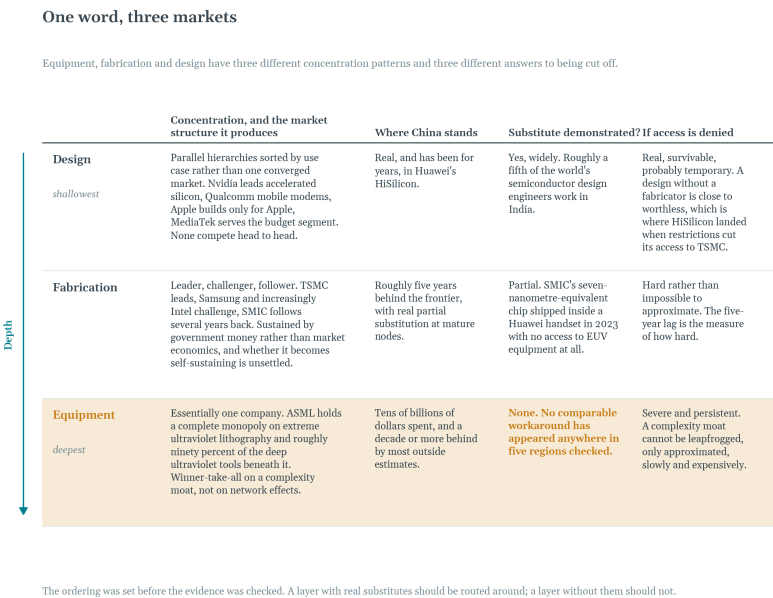


Figure 2. The second question, applied. One word covers three layers with three different concentrations, three different substitution records, and three different answers to being cut off.

Third: do you hold both conditions, or only one? Farrell and Newman name two requirements before a position becomes leverage. Jurisdiction means the state can compel the entity holding the node — a legal question, usually with a clear answer. Institutional freedom

means the state can decide, on its own judgment, whether to use that compulsion. The second is political and has almost nothing to do with law. A government can hold unambiguous authority over a firm and remain unable to exercise it in the direction it would choose, because doing so would cost something it values more elsewhere. The two fail differently and are repaired differently: missing jurisdiction is remedied by legislation, reasonably quickly, while missing institutional freedom is not remedied in any short period, because what constrains it is the country's entire external position.

Fourth: what is the position for? Either it is a weapon, valued for the ability to deny, or it is a guarantee of a seat, valued because nobody can proceed without you. Those are not two descriptions of the same asset. They make opposite predictions about what happens when the position is used, which means the record can settle which one is right.

V.

Two governments held comparable positions and made opposite choices.

In July 2019 Japan imposed licensing requirements on exports to South Korea of three materials essential to making chips and displays. The trigger was not commercial: a Korean court had ruled the previous year that Japanese companies owed compensation for wartime forced labor.

The position was real. Japan produced something like ninety percent of the world's fluorinated polyimide and the large majority of photoresist, and South Korea took roughly nine tenths of its photoresist imports from Japanese suppliers who fed Samsung, SK Hynix and LG. Japan held jurisdiction over the node and complete freedom to act on its own judgment. Both conditions were satisfied. Japan used them.

Korea did not capitulate. It reallocated. Buyers moved sourcing toward Belgium, the United States and Taiwan, Korean domestic production rose, and Japanese suppliers themselves increased production inside South Korea — the outcome least often noticed and the most revealing. Faced with losing the customer, the holder of the chokepoint relocated part of the chokepoint into the customer's territory. By 2023, with the political relationship repaired, Japan lifted the restrictions and recovered its trade.

What Japan spent was not revenue. It was the assumption of reliability that made the position valuable in the first place. A supplier who has demonstrated that it will withhold has told every customer in the world something they will act on for a decade.

The Netherlands holds a position at least as deep and has never used it. ASML has a complete

monopoly on extreme ultraviolet lithography and roughly ninety percent of the deep ultraviolet immersion tools beneath it; tens of billions of dollars of Chinese spending has dented neither. In June 2026 it became the most valuable company in European history.

The Dutch government has jurisdiction over that node completely. What it does not fully have is the second condition. It has repeatedly had to balance pressure from Washington against its own commercial relationship with China — tightening export terms in 2023, extending restrictions to servicing installed machines in 2024, managing an unresolved dispute with the United States over an allegation that a restricted system reached China anyway. Those constraints arrived from an alliance position, not from any limitation in the technology.

Read as a weapon, that is failure: a deep node its holder cannot fire at will. Read as a seat, it is close to complete success. American export policy toward China requires Dutch cooperation and has to be negotiated rather than imposed. No party in the system, however powerful, has behaved as though the node were expendable.

The record settles it. Japan used its position, recovered its trade, and holds a less secure node than before. Europe has not used its position and holds one that no amount of motivated spending has dented. If the distinction were merely rhetorical, those outcomes would run the other way.



Figure 3. Both governments held a deep node. One fired it and one has not, and the outcomes ran opposite to what the coercive reading predicts.

Which gives the sharpest version of the argument. **Sovereignty is the capacity to decide alone. Indispensability is the guarantee that others cannot decide without you.**

They overlap in nothing except that both make a party hard to ignore. A sovereign actor can act against the wishes of everyone else and accept the consequences. An indispensable actor cannot act alone at all, and does not need to, because no arrangement can be built that excludes it. Europe is indispensable and not sovereign, and it chose that. Reading the second half of that sentence as a failure to achieve the first mistakes two different axes for one.

VI.

Most sovereignty claims fail the test at the first question. They reach a shallow layer while the layers beneath it are held elsewhere.

The second most common failure is subtler and more expensive: the claim reaches a genuinely deep layer, the ambition is correct, and nothing in the plan obliges anyone to do anything.

Europe supplies thirty years of evidence on that point, and the pattern repeats with unusual precision. In December 1993 the European Council commissioned a report on the continent's position in information technology. A high-level group chaired by Martin Bangemann delivered it to the Corfu summit in June 1994. The membership was industrial rather than academic, including the heads of Volvo, Olivetti and what would become BT. Its diagnosis was that entrenched positions had put Europe at a competitive disadvantage against the United States and Japan, that the continent had the technological and entrepreneurial capability, and that it was not converting it. It observed, memorably, that political attention was too intermittent and that the private sector was waiting for a signal.

I was in corporate strategy at Ericsson at the time, and part of that work involved carrying the company's view into the European information society discussions that followed, presenting in Stockholm and Barcelona alongside Swedish politicians and other industry representatives. What I watched, without having a framework for it then, was a set of recommendations splitting into two kinds.

Some were instruments. Accelerate telecommunications liberalization. Open the remaining monopoly infrastructure to competition. Set clear timetables and deadlines. Others were aspirations. Make interconnection and interoperability primary Union objectives. Review the standardization process to make it faster and more market-responsive.

Thirty years later, look at which half happened. Telecommunications liberalization proceeded, because it arrived as directives with dates attached and national regulators obliged to act. Interoperability as a primary Union objective produced conferences. The standardization process was not made materially faster. And GSM, running alongside all of this on a spectrum allocation nobody could route around, became the world standard.

That last clause is the whole lesson. GSM did not win because Europe wanted it to win. It won because spectrum is a channel. A frequency allocation obliges every participant without anyone being prosecuted, and a manufacturer who ignores it has no product. The allocation did not need enforcement. It needed only to exist, and everything downstream arranged itself around it.

The pattern is not that Europe misdiagnoses itself. It diagnoses itself accurately, repeatedly, at thirty-year intervals, with distinguished people and considerable evidence. The Draghi Report of September 2024 restated the 1994 verdict with harder numbers: European research and development spending at 2.11 percent of GDP against 2.56 in China and 3.59 in the United States, with the gap widening; four of the world's top fifty technology companies European. Its central recommendation is close to eight hundred billion euros of additional annual investment, roughly five percent of European GDP, at levels the continent has not sustained since the 1970s. The diagnosis is thirty years old. Only the arithmetic is new.

Eight hundred billion euros of recommended annual investment is a number, not an instrument. Spectrum was an instrument. Four decades of accumulated optical engineering is an instrument. The difference between them is the whole of Europe's technology record, and the question any sovereignty program should have to answer before it is funded is not whether the ambition is correct but whether anything in it obliges anyone to do anything.

VII.

Run the test on a live case.

Canada spent much of 2026 negotiating with the United States, with sovereignty named explicitly among its objectives. When the talks collapsed in August, the reasons given were structural rather than commercial: Ottawa said the United States had sought to restrict Canada's ability to make trade agreements with other countries, and to constrain protections for language and culture.

The specific matters in play are instructive. Requirements for French-language markings on products. The Online Streaming Act, which obliges platforms to contribute to Canadian cultural industries. The Online News Act, which requires payment for news. A proposed restriction on social media for those under sixteen. The American trade representative raised digital trade alignment as a component of any agreement.

Every one of those measures sits at the application layer. They govern what may be shown, what must be paid, what must be printed on a box, and who may hold an account. They are real exercises of authority and they are squarely within Canada's jurisdiction.

The difficulty is one layer down. The platforms being regulated run on operating systems Canada does not hold, on cloud infrastructure Canada does not hold, on silicon Canada does not hold. The layers beneath the regulated layer belong, substantially, to the party across the table.

The general point outlives the negotiation, whatever it settles into. A country asserting sovereignty at the application layer, over infrastructure it does not hold, is asserting something real and something fragile at the same time. Both halves of that sentence matter. The assertion is not empty. It is simply not load-bearing in the way the word suggests.

VIII.

The test has limits, and they should be said out loud.

It measures position, not merit. A claim that fails the depth test may still be entirely right, and a country may be correct to make it. Canada's cultural protections are not made wrong by sitting at a shallow layer. They are made vulnerable, which is a different finding and calls for a different response.

And the test does not say what to do. Knowing that a claim reaches a shallow layer tells you the shape of the exposure. It does not tell you whether to accept it, insure against it, or spend a decade building a deeper position. Those are political choices and the test has nothing to say about them.

But it does narrow the field, and it points somewhere counterintuitive. For most countries, full sovereignty over a technology stack is not achievable at any price they are willing to pay, and the programs promising it will keep producing accurate diagnoses and disappointing outcomes. Indispensability is achievable. It requires choosing one node deliberately, building unmatched depth there, and accepting that the instrument works by existing rather than by being fired.

That is a smaller ambition than sovereignty. It has the advantage of being available.

Mike Edholm has spent three decades advising corporations and institutions on scenario planning and strategic foresight, with assignments spanning telecommunications, computing, industrial automation, and international development. This essay is part of a series on scenario planning methodology.