From Centralized to Federated: Extending Governance Continuity for a Multi-Sovereign World

A public introduction and an invitation to expert review

Governance Continuity, as originally introduced, defined an architectural property: governance-relevant context — who acted, under what authority, under what policy, with what evidence — should remain intact as activity crosses between independently governed systems. The original framework's reference design achieved this through a single coherence point: one component, one operator, positioned to observe both sides of every boundary crossing.

That design choice, examined closely, is also its principal limitation. A single point of observation is also a single point of control. Within one country, concentrating this capability in one organization, public or private, creates exactly the kind of unchecked power the framework exists to hold accountable: no other party can verify the record, contest it, or catch it when it is wrong. Across borders, the problem compounds. Whichever government has legal reach over that one operator gains structural leverage over every other country relying on it. Sovereignty would, in effect, be surrendered to the operator — or to whichever jurisdiction that operator answers to.

Any credible answer needed to address both forms of centralization at once, not treat them as separate problems.

We explored ways to distribute this authority rather than concentrate it: federating the architecture across many independent operators that do not need to trust one another, rather than centralizing it in one. We use the word *federated* deliberately, in place of *decentralized*. This is not an architecture in which authority is absent. Each country, and within a country each authorized organization, retains real authority over its own infrastructure. What changes is that no single piece ever holds complete authority over an exchange between two parties. Authority is distributed among the actual parties involved, not eliminated.

The result works like this. Instead of one party watching both sides of a cross-border exchange, each side runs its own infrastructure. The two sides jointly produce a record of what happened between them, split in half, with one half signed by each side. Neither can fake the record alone. Neither can erase its half without the other side still being able to prove what happened. Neither has to see inside the other's systems to do it. The same split-authorship principle solves the domestic concentration problem too: even within one country, no single organization is positioned to unilaterally author, alter, or suppress the record of what it did.

The addressing itself carries the same logic. Instead of a system having one fixed identity, each connection carries its own identifier tied to the specific rule that authorizes it: a legal agreement, an approved transfer, a specific permission. If that rule is revoked, every connection relying on it can be shut off instantly and provably, without touching anything else. If a category of data is legally required to stay inside a country, that rule is enforced simply by never issuing an identifier that could leave, so the absence of any crossing becomes proof the rule was followed.

This matters because data protection laws around the world genuinely disagree with each other. Some countries whitelist approved partners. Some restrict specific categories outright. Some require government approval case by case. Some regulate by industry sector. A federated architecture does not need to pick a side among these approaches. Each rule simply becomes a policy the system enforces and records, so the same infrastructure works under every regime at once and is controlled by no single party.

One limit deserves to be stated as plainly here as anywhere in this work. This architecture produces strong evidence: proof of what crossed, when, and under whose authority. It is not, by itself, legal compliance, and it cannot prevent a government from lawfully compelling a company within its own territory to disclose something in secret. No architecture can promise that. What federation does provide is proof that no single party, including us, controls the whole picture.

An invitation to challenge this work

The technical detail behind this concept — the full architecture, its formal properties, and an honest accounting of what remains unproven — is significant enough to warrant real scrutiny before it is offered as a finished contribution. Rather than publish that material openly at this stage, we are making it available under controlled access to a small number of invited reviewers with relevant expertise in distributed systems and network architecture, cryptography and security, and international data protection and technology law.

We are not asking for encouragement. We are asking these reviewers to find what is wrong with it: to test the claims against their own expertise, identify where the architecture overreaches, and press on the places we have already flagged as unresolved. A paper that survives that kind of challenge is worth more than one that has only been read by people inclined to agree with it.

Contributions incorporated into the final published work will be credited by name, at the contributor's discretion, in the paper itself.

If you have relevant expertise and would like to be considered as a reviewer, contact Donald W. Jacobs at Don@ENATIP.com.