FREEHOLD ADVISORY

Designed Around the Fish

What restoration would look like if the system were built for what it's trying to save

Ben Parker

Freehold Advisory

April 2026

The diagnostic question

I built Freehold to answer a question no one was asking: what does the full portfolio look like? Not the federal portfolio or the state portfolio — the *full* portfolio, every obligation, every timeline, every match commitment, every reporting deadline, as it all lands on one desk on a Monday morning in September when the fish are running and three federal reports are due. Assembling that picture required twenty-three data connectors and fourteen factors of structural complexity. What it revealed, across fifty organizations in Washington State's restoration sector, is that roughly eighty percent have structural portfolio conflicts, and approximately seventy percent of those conflicts are invisible to any single data source.

That is not a finding about management. It is a finding about architecture. The system was never designed as a system. It accreted. And a diagnostician who can see the whole patient can also ask a harder question: what would the system look like if it were designed for what it's trying to save?

How the system got here

No one planned this. The Endangered Species Act created obligations to recover listed species. The Boldt Decision affirmed tribal co-management of fisheries. The Clean Water Act funded water quality improvements. The Pacific Coastal Salmon Recovery Fund directed money to states. RCO administers state grants. NOAA administers federal ones. USFWS runs a separate set. EPA funds water quality. USDA funds land conservation. Each program was a reasonable response to a real problem. But each brought its own application process, reporting cycle, compliance framework, and fiscal calendar.

Multiply this by fourteen Regional Fisheries Enhancement Groups, thirteen tribes with sovereign co-management authority, dozens of conservation districts, land trusts, and watershed councils. An organization accessing five federal programs and two state ones isn't navigating one system seven times. It's navigating seven different systems, each with its own logic, its own deadlines, its own definition of success. The accretion was rational at every step. The result is not.

Pull the thread

Start with one fish.

A coho salmon returning to spawn in a South Puget Sound creek needs a few things. Not complicated things. Old things. Cold water. Passage, no barriers between the estuary and the gravel. Flow. Cover.

Now pull the thread.

Cold water means riparian canopy. Shade. Trees within fifty feet of the stream bank, and ideally standing long enough to matter — decades, not grant cycles. Riparian restoration is funded through SRFB, NRCS EQIP, EPA Section 319, and sometimes USFWS Partners for Fish and Wildlife. Four programs, four reporting systems, four timelines. One stream bank.

Passage means culverts that aren't barriers. Washington's fish passage inventory identifies thousands. Fixing one requires engineering design (a planning obligation), then construction (a restoration obligation), then permits: an HPA from WDFW, Section 404 from the Army Corps, sometimes SEPA review. The planning money comes from one source on one timeline. The construction money comes from another source on another timeline. The permits run on a third. The in-water work window, the weeks when you can actually put equipment in the stream, runs July through August. Miss it, and you wait a year.

Flow means water in the stream when the fish need it. Which means water rights, minimum instream flow rules, agricultural withdrawals, municipal demand. Ecology sets the rules. Counties enforce land use. Irrigation districts manage allocation. The fish has no standing in the water rights system. Its needs are represented by regulatory minimums set by an agency that is also managing competing demands.

Cover means habitat complexity: large wood, pool depth, floodplain connectivity. Which means land use at the landscape scale. Floodplain development, levee setback, agricultural conversion. FEMA manages flood risk. Counties manage zoning. Land trusts acquire parcels with funding that carries its own match requirements and reporting timelines. Conservation districts work with willing landowners, one fence line at a time.

One fish. One creek. And already the thread touches NOAA, USFWS, EPA, USDA, WDFW, Ecology, the Army Corps, FEMA, a county government, a conservation district, a land trust, an enhancement group, a tribal nation, and a regional recovery board. Each with its own mandate, its own funding, its own reporting system, its own clock.

The fish doesn't know any of this. It needs cold water, passage, flow, and cover. The system that is supposed to provide those things has organized itself around *inputs* (money, mandates, jurisdictions) rather than around *outputs*: living fish in living streams.

This is the core design failure. Not that the agencies are wrong, or the programs are wasteful, or the people aren't committed. They are, on all counts. The failure is architectural. The system is organized around its funding sources rather than around the thing it's trying to fund.

Five design failures

Freehold's diagnostic work reveals five structural problems that persist regardless of which policy pressure intensifies, which party holds power, or which budget cycle the sector is in. These are not failures of execution. They are failures of architecture.

No shared ledger. No single system sees the full portfolio of obligations landing on one organization. USASpending tracks federal awards. RCO PRISM tracks state projects. ProPublica reports annual financials. WDFW tracks the fish. No one sees what converges on one desk. I had to build twenty-three data connectors to assemble what should already be infrastructure. In any comparable sector with this level of public investment, a shared ledger would exist. In restoration, it doesn't.

Blind loading. Because there is no shared ledger, every funder adds obligations without seeing what's already there. NOAA awards a cooperative agreement in March. RCO funds a salmon recovery project in June. EPA adds a water quality grant in September. Each award is meritorious. Each adds reporting requirements, match commitments, and timeline constraints. None of them knows about the others. Freehold's data shows that financial complexity and institutional complexity load work onto the same organizations independently; the correlation between them is effectively zero. Two weather systems converging on the same coast.

Two incompatible clocks. The biological clock runs on seasons: in-water work windows open in July, planting season starts in October, spawner surveys run August through October. These are set by water temperature and species physiology. The fiscal clock runs on politics: the federal fiscal year ends September 30, the state biennium follows its own rhythm, and quarterly reports land on calendar quarters regardless of whether work was possible that quarter. End-date clustering, the most common conflict pattern in Freehold's data, is almost entirely an artifact of these two clocks colliding on the same desk every autumn.

Reporting without knowledge. A typical restoration organization reports to six or more systems on different cycles. Each system asks for overlapping but different information. Each has its own format, its own portal, its own deadlines. The aggregate cost is significant. The aggregate knowledge is fragmented. No system synthesizes these reports into a portfolio-level picture. The reporting burden produces compliance artifacts that satisfy individual funders without informing the sector about what is actually happening at the organizational level.

Merit without capacity. Salmon Recovery Funding Board scores projects on technical merit, cost-effectiveness, and ecological benefit. These are the right criteria for evaluating projects. But the system has no mechanism to ask: does this organization have the capacity to execute this project alongside everything else it's already doing? A technically excellent project loaded onto a desk already carrying twenty active obligations and match commitments exceeding thirty percent of annual revenue is a worse investment than a good project loaded onto a desk with bandwidth. The system can see projects. It cannot see desks.

What a designed system would look like

These are not policy recommendations. They are design principles: the architectural choices that would follow from taking the fish seriously as the organizing unit. If the sector wanted to solve these problems structurally, these are the constraints any solution would have to satisfy.

The watershed as the unit of account. Funding would flow to watersheds, not to programs. A coordination mechanism at the watershed scale would receive federal, state, and private funding and allocate it based on ecological priority and organizational capacity. The more than thirty-five federal assistance listings currently funding restoration work in Washington would consolidate into one interface at the watershed level.

The structure for this partly exists. Lead Entities (NOPLE, WRIA 8, WRIA 9, and others) already coordinate project prioritization at the watershed scale. But they coordinate between fragmented funding systems. They recommend; they don't allocate. A designed system would close that gap.

The hardest constraint is also the most important one. Tribal nations hold co-management authority that predates every agency, every program, and every coordination structure in the system. The Boldt Decision did not create a seat at someone else's table. It established a separate, co-equal authority over fisheries management. Any watershed-scale redesign that treats tribal co-management as one voice among many, rather than as a co-governing authority, will fail, and should. This is not a design constraint to work around. It is a design constraint to build around.

The biological calendar as the master clock. Reporting and compliance cycles would follow ecological cycles, not fiscal ones. Construction projects would report after the work window closes. Planting projects would report after the growing season demonstrates survival. Habitat projects would report when monitoring data is available — not when the federal quarter ends.

This means asking federal agencies to accept that their grantees' work runs on a clock they don't control. It means trusting that money spent in August on in-water work is better reported in November, when results are observable, than in September, when the arbitrary deadline falls.

This single change would eliminate more portfolio conflicts than any amount of process streamlining. When the biological clock is the master clock, the fiscal clock becomes a back-office function. Not eliminated — reconciled.

Portfolio visibility as infrastructure. What Freehold currently assembles from twenty-three data sources would be built into the system's architecture. Before a new obligation loads onto an organization, the full portfolio picture (active obligations, match commitments, reporting deadlines, capacity utilization) would be visible to the awarding entity. Not as a diagnostic after the fact, but as a design input before the award.

This is the principle that would make Freehold's core diagnostic unnecessary. That is the point. A system that needs an outside practice to assemble its own picture of itself has failed at self-awareness. The information already exists. It is scattered across databases that don't talk to each other. Connecting them is an infrastructure problem, not a knowledge problem.

Capacity alongside merit. Project evaluation would include organizational carrying capacity as a factor. Not replacing ecological merit, which matters, and SRFB's technical review process is one of the strongest elements of the current system. But recognizing that a project's probability of successful completion depends on the organizational context it lands in.

An organization showing high concurrent load (many active obligations, compressed reporting timelines, elevated match burden) would either receive capacity support alongside a new award, or the award would be structured for phased execution rather than immediate full loading. The system would see the desk, not just the project.

This requires something the current system doesn't have: a working definition of organizational carrying capacity. Freehold has built one. It's not the only possible approach, but it demonstrates that the measurement is feasible. What gets measured can be managed.

Outcomes over activities. Report on fish, not on process. Spawner returns. Barriers removed. Riparian miles restored. Floodplain acres reconnected. Water temperature trends at established monitoring stations.

This is the hardest principle because it requires the deepest trust. Activity-based reporting exists because funders need assurance that public money is spent responsibly. That need is legitimate. But the current approach, documenting every expenditure across six different portals, produces accountability for process without accountability for outcomes. The monitoring infrastructure to track ecological results largely exists: WDFW tracks escapement, NOAA tracks habitat, USGS tracks hydrology. What's missing is the willingness to let outcomes serve as the primary measure of whether the investment is working.

What the sector would give up

A designed system would cost something. Being honest about those costs is more useful than being optimistic about the benefits.

Organizational autonomy. Many restoration organizations define themselves by their independence: their relationships with their watershed, their community, their board. A system organized around watersheds rather than organizations changes those identities. An enhancement group that currently manages its own portfolio of grants would become part of a watershed-level allocation system. Even if the ecological outcomes improve, the loss of independence is real.

Agency authority. NOAA, EPA, USFWS, and WDFW have distinct mandates and distinct institutional cultures. A coordination layer that channels their funding through a common interface reduces their individual discretion. Agencies exist to exercise judgment about how their mandates are fulfilled. Asking them to coordinate allocation is asking them to share that judgment. Some will see this as overdue efficiency. Others will experience it as a loss of mission.

Institutional knowledge of the current system. Organizations have adapted. They've hired grant writers who know the SRFB process. They've built reporting workflows for six portals. They know which program officers to call and when. A redesigned system makes that institutional knowledge less valuable. The people who are best at navigating the current system have the most to lose from changing it.

Control through compliance. Activity-based reporting gives funders traceability. They can verify that each dollar was spent as intended, that deliverables were produced, that match was documented. Outcome-based accountability trades that granular control for something more meaningful — did the fish come back? — but the trade requires a level of trust that the current system was not built to support.

The window

The Infrastructure Investment and Jobs Act expires after federal fiscal year 2026. The State of the Sound 2025 documents an approximate twenty-five percent functional reduction to key salmon recovery programs. Organizations have already begun to feel it: paused projects, deferred maintenance, staff uncertainty.

This is not a future scenario. It is the current budget trajectory.

When there is less money in the system, the proportion spent on administration becomes harder to defend. When organizations are cutting ecological programs to maintain compliance staffing, the question of whether the system's architecture is worth preserving gets asked out loud, by the people carrying the cost, if not yet by the people who designed it.

The IIJA was the largest federal investment in restoration infrastructure in a generation. If its expiration results only in less funding moving through the same fragmented architecture, the same blind loading, the same incompatible clocks, the same reporting without knowledge, that would be a failure of imagination on the far side of a crisis.

The alternative is to use the disruption as a design opportunity. Not just how much money flows to restoration, but how it flows. Not just whether to fund the next round, but whether the round should work the way it has always worked.

The X-ray and the blueprint

I built Freehold to provide a structural picture that no single data source can see alone. Freehold assembles data, measures complexity, maps conflicts, and describes what emerges without judgment. The complexity Freehold observes is not a failure of the organizations carrying it. It is a feature of a system that loads work onto desks without seeing what's already there.

That diagnostic work is necessary today. It will remain necessary for as long as the current architecture persists. But I would be incomplete in my diagnosis if I described the system's structural problems without noting that structural problems can have structural solutions. The five design principles outlined here (watershed-scale allocation, biological calendars, portfolio visibility, capacity-based funding, and outcome accountability) are not aspirational ideals. They are the minimum ar-

chitectural changes required to address the conflicts Freehold observes in the majority of portfolios it has analyzed.

I wrote previously about the five pressures that converge on a single restoration director's desk. This essay asks what a system would look like if those pressures didn't have to converge. If the architecture distributed load intentionally rather than accidentally, and measured success by what happened in the stream rather than what was reported in the portal.

Freehold doesn't prescribe. It diagnoses. But a diagnostician who can see the whole system can also describe what a healthy one would look like.

The fish don't know which agency funded the culvert. They just need to get through.