CASE STUDY · INDIGENOUS MEDIA GROUP

# What We Built for IndigeFish Fishing Co.

A commercial forecasting platform covering 250 miles of Southern California coast. Six public data feeds, a prediction engine that explains its own reasoning, and a working subscription business — designed, built, and operated in-house.

- Forecast platform
- Field data capture
- Subscription business

**CLIENT**
IndigeFish Fishing Co.
Southern California

**SCOPE**
Product design, architecture, engineering, data integration, billing, hosting and operation

**BUILT BY**
Indigenous Media Group

**CONTACT**
Robert Jacobo III, President
rjacobo@4bmg.org · indigenousmediagroup.org

SECTION 01

# Where it started

Federal agencies publish ocean data continuously, accurately and for free. Almost nobody can use it. The distance between a public dataset and a decision somebody can act on is the whole problem — and it is the problem we were asked to close.

**6**

**Public data feeds**

Tides, buoys, weather and more — all free, all public domain, none usable as published.

**250**

**Miles of coast**

Santa Barbara to the Mexican border. Five counties, one weather system.

**81**

**Species**

Each with its own habitat, temperature band and feeding behaviour.

**$0**

**Cost of the data**

Every input is public. The value was never the data. It was the translation.

### FOR THE USER

A wasted morning costs fuel, daylight and the one free day somebody had. The information that would have prevented it was published and sitting there, in a format built for researchers.

### FOR PUBLIC DATA

Ocean monitoring is paid for by everyone. Without translation, a public asset stays theoretical for the public that funded it.

### FOR THE BUSINESS

IndigeFish had an audience and a voice, and no product underneath it. An audience with nothing to sell is a cost, not a business.

**The part worth remembering**

This was a translation problem, not a data problem. Nothing needed to be discovered — it needed to be made usable by a person standing outdoors with one free hand. That distinction decided every technical choice that followed.

**What we built first**

The prediction engine, before a single screen was drawn. A polished interface sitting on a model that does not work is worse than nothing, because it earns trust it cannot keep. The model was built, measured against real conditions and corrected before anybody saw a pixel.

SECTION 02

# What we delivered

Three connected systems on one platform — a forecasting product, a field data recorder, and the commercial machinery to run it as a business. Live, taking payments, and maintained by us.

### THE FORECAST PLATFORM

Six live data sources normalised into hourly predictions across the full Southern California catalogue — ranked, filterable by species, and explained factor by factor. Seven days ahead, rebuilt daily.

### FIELD DATA CAPTURE

An offline-first logbook that records conditions at the moment of the entry, a shared coastal map and a community feed — the dataset that makes every later forecast better.

### THE BUSINESS SYSTEM

Subscriptions with trials and entitlements, a compliant email engine, a self-serve competition platform, a referral programme and internal analytics.

**The principle underneath all of it**

Every third-party dependency is optional. Remove the payment provider and billing turns off. Remove the email service and notifications turn off. Remove any single data source and one reading goes quiet. The core product keeps serving in every one of those cases. A platform that collapses because one vendor had a bad afternoon is not finished.

### BUILT FOR THE CONDITIONS

This product is used outdoors, one-handed, in direct sunlight, often on one bar of signal and a phone the user does not want to drain. Data is written to the device first and synced when the network returns. Charts are driven by a slider rather than a fingertip, because outdoors a finger covers the number being read. These are not polish items — they decided the architecture.

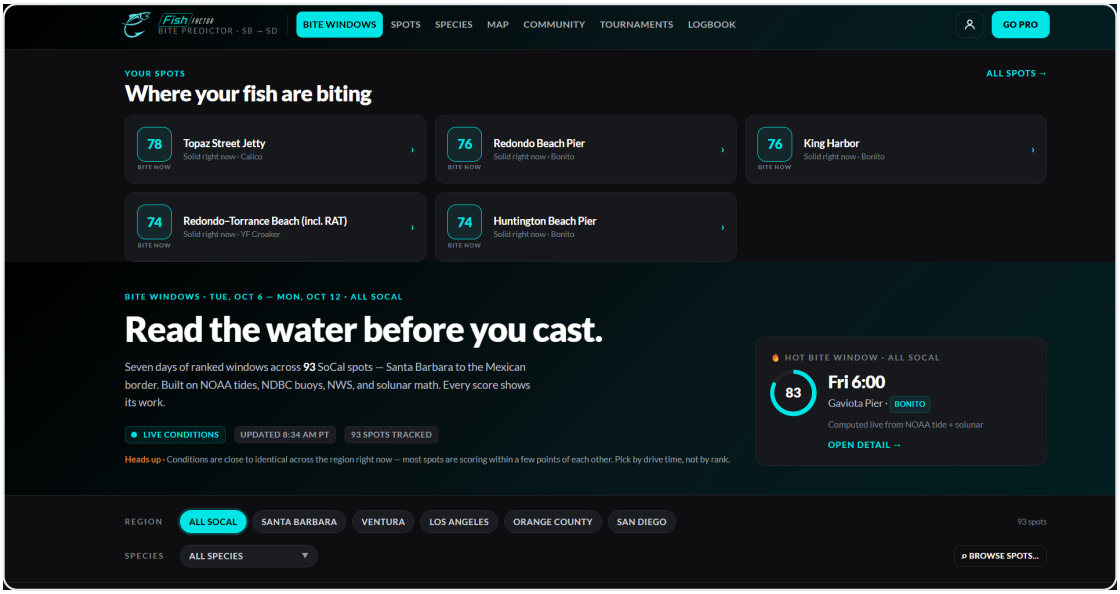
### BUILT TO BE MAINTAINED

Every database change is a numbered, reviewable, reversible migration. The prediction model carries a version stamp, so any forecast is attributable to the logic that produced it and a change can be rolled back. Documentation is checked against the running system rather than trusted, because a description that has drifted is worse than none.

FORECAST PLATFORM · 03

# The home screen

The first screen anybody sees. It has one job: answer whether it is worth going right now, and say where.




### ON THE PAGE

Saved locations across the top with a live score each, then the day's strongest window anywhere on the coast — a score, a place, a species and a time.

### WHY IT WORKS

Three badges state when conditions were read and how many locations are being tracked. The product shows the age of its own answer before it is asked.

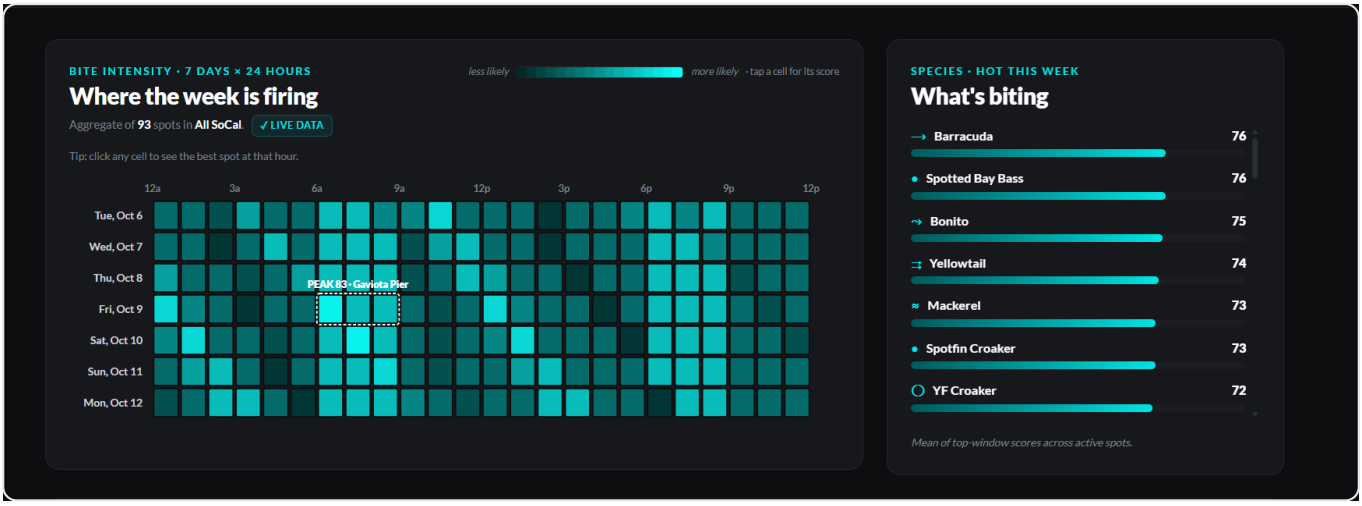
### WHAT IT CHANGES

A decision that used to mean cross-referencing three websites takes seconds — and a note warns when scores are clustered, so nobody over-reads a one-point gap.

FORECAST PLATFORM · 04

# Seven days at a glance

A different question from the home screen. Not “should I go now” but “which morning this week is worth taking off work”.




### ON THE PAGE

Seven days across, twenty-four hours down, every cell one hour at one place. Beside it, the species running hottest this week, ranked.

### WHY IT WORKS

It turns a planning problem into a visual one. Nobody reads this grid — they see the peak marked on Friday dawn and plan around it.

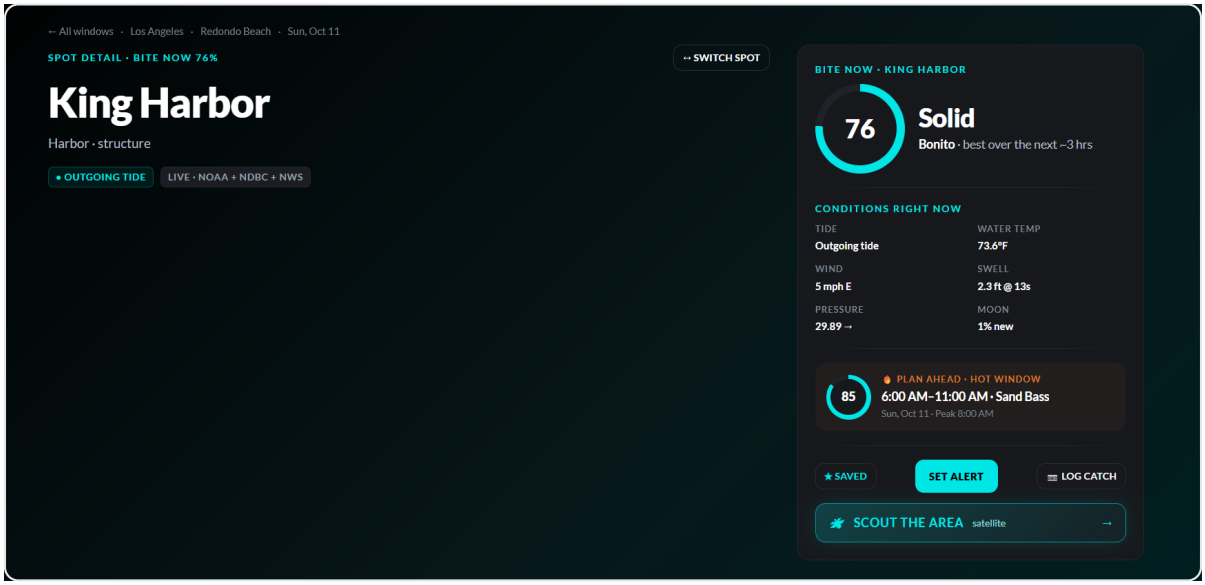
### WHAT IT CHANGES

Clicking any cell names the best location at that hour. The grid is not a chart; it is a route into the decision.

FORECAST PLATFORM · 05

# The location page

Every location gets its own page, its own data sources, and its own answer.




### ON THE PAGE

A live score with a plain-language label, the species currently leading, six current readings, and the next strong window with its peak hour.

### WHY IT WORKS

Each location is bound to its own tide station and nearest buoy. The user does not need to know that 73.6°F sits above a species' preferred band — the model already weighed it.

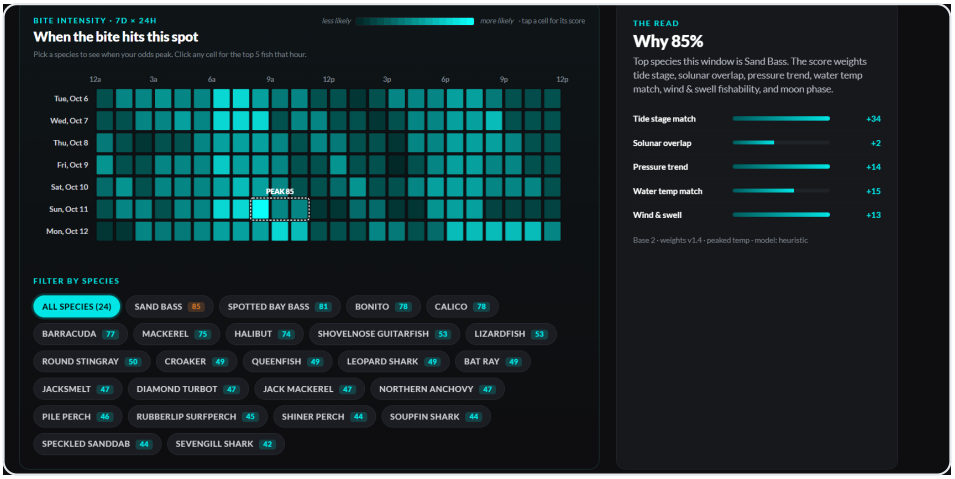
### WHAT IT CHANGES

One tap sets an alert for this place, one logs a catch against it, one opens satellite imagery to scout access before driving out.

FORECAST PLATFORM · 06

# The score shows its work

The most important screen in the product, and the one we would put in front of any sceptic.




### ON THE PAGE

The score broken into the factors that produced it — tide stage +34, pressure trend +14, water temp +15, wind and swell +13, solunar +2 — with the model version printed underneath.

### WHY IT WORKS

The contributions add up to the number displayed. That sounds trivial and is not: guaranteeing it took real engineering, and most competing products never attempt it.

### WHAT IT CHANGES

A black box that says 85 earns nothing from a sceptical user. A number that shows its reasoning can be argued with — and when it is wrong, the user can see which input misled it.

**Why we hold this standard**

The people using this are sceptical by trade and they talk to each other. A tool that overstates is found out inside a season. Explainability is not a feature here; it is the only defensible position for a product making predictions a user can personally check against reality.

FORECAST PLATFORM · 07

# The conditions underneath

For users who want to check the working rather than take the answer.




### ON THE PAGE

A full tide curve with now and the next hot window both marked, high and low either side, and the whole week's bite intensity for this one location.

### WHY IT WORKS

The charts were built rather than imported, so they could be driven by a slider instead of a fingertip. Outdoors, a finger covers the value being read.

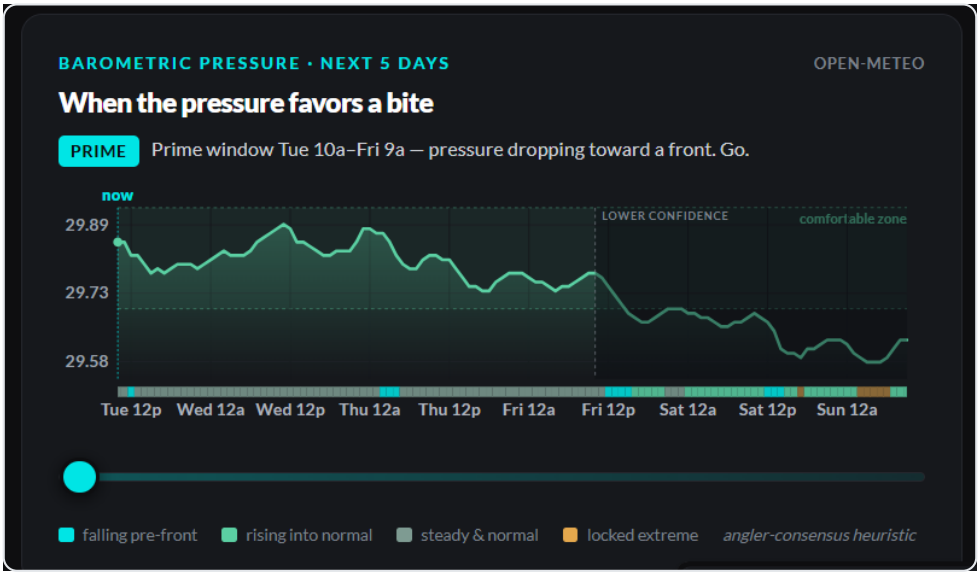
### WHAT IT CHANGES

Experienced anglers can audit the model against their own judgement. That is how a forecasting product earns trust from people who already know the water.

FORECAST PLATFORM · 08

# Turning data into a recommendation

The difference between displaying a number and telling somebody what it means.




### ON THE PAGE

Five days of barometric pressure, shaded by what each stretch means — falling pre-front, rising into normal, steady, locked extreme — with the confidence horizon marked.

### WHY IT WORKS

The banner names the prime window in plain language and tells the angler to go. The chart supports the recommendation rather than replacing it.

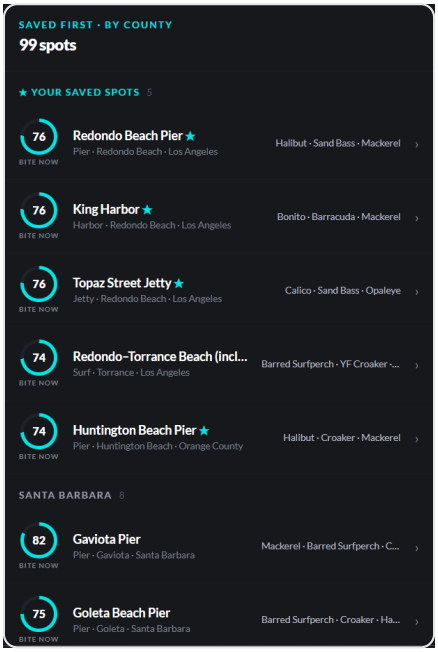
### WHAT IT CHANGES

Pressure is the input anglers argue about most and understand least. Labelling the pattern, and saying where the forecast gets less certain, is more useful than a line on a grid.

FORECAST PLATFORM · 09

# The full catalogue

A reference asset that cannot be scraped from anywhere, because it did not exist until it was assembled.




### ON THE PAGE

Every usable location across five counties, grouped by region with a live score and leading species. Saved places rise to the top.

### WHY IT WORKS

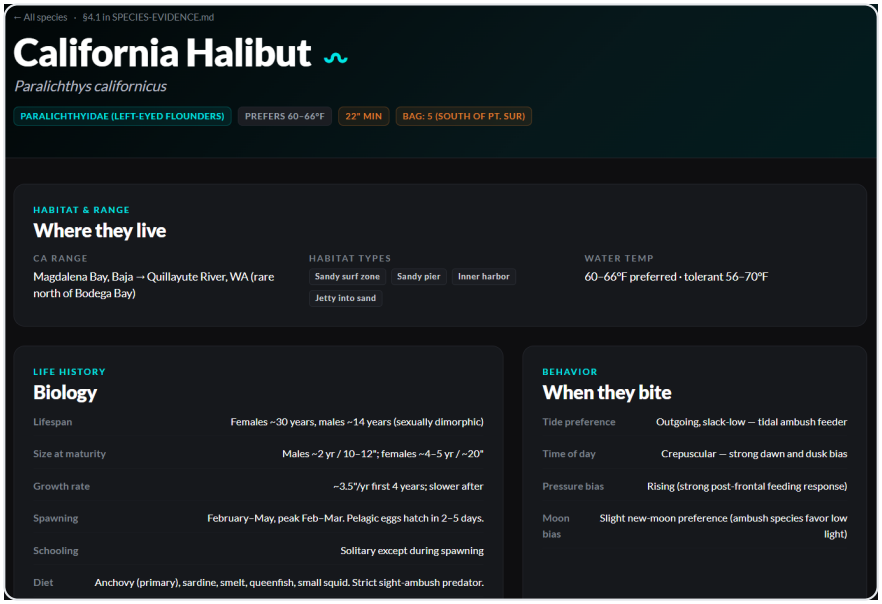
Each record holds a tide station, nearest buoy, depth range, facing direction, structure type and species list — local knowledge, structured.

### WHAT IT CHANGES

A national competitor can add the region to a map. They cannot add which end of a jetty holds fish on an outgoing tide.

# Species knowledge

The domain layer that makes the forecast species-aware rather than generic.




### ON THE PAGE

Range and habitat, preferred and tolerated water temperature, life history, diet and predators, and the tide, time, pressure and moon conditions each species feeds on.

### WHY IT WORKS

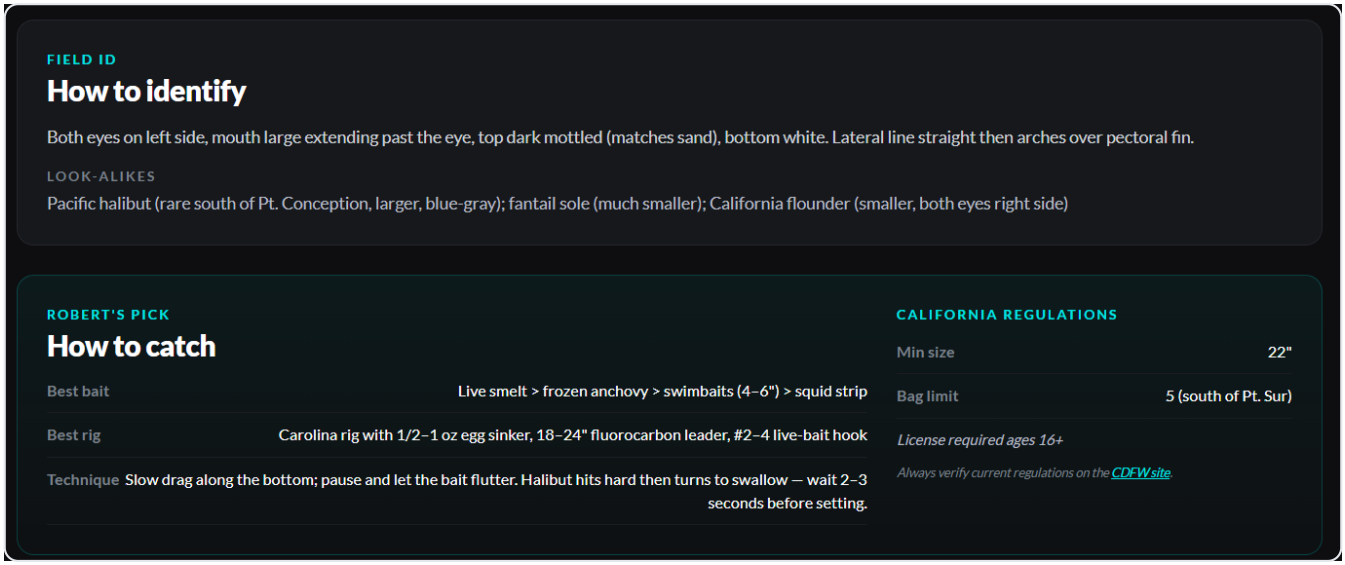
Conditions that are excellent for one species are unremarkable for another, and the model knows the difference.

### WHAT IT CHANGES

Every claim cites its section in an internal evidence file. The rule is that a claim must be falsifiable and defended, and where the evidence and the catalogue disagree, the catalogue is wrong.

# All the way down to the rig

Knowing when to go is half the decision. The other half is what to do when you get there.




### ON THE PAGE

Field identification with look-alikes, then bait, rig and technique in the owner's own words — next to the legal minimum size and bag limit.

### WHY IT WORKS

Regulations sit beside the advice rather than buried in a policy page, with a standing instruction to verify against the official state source.

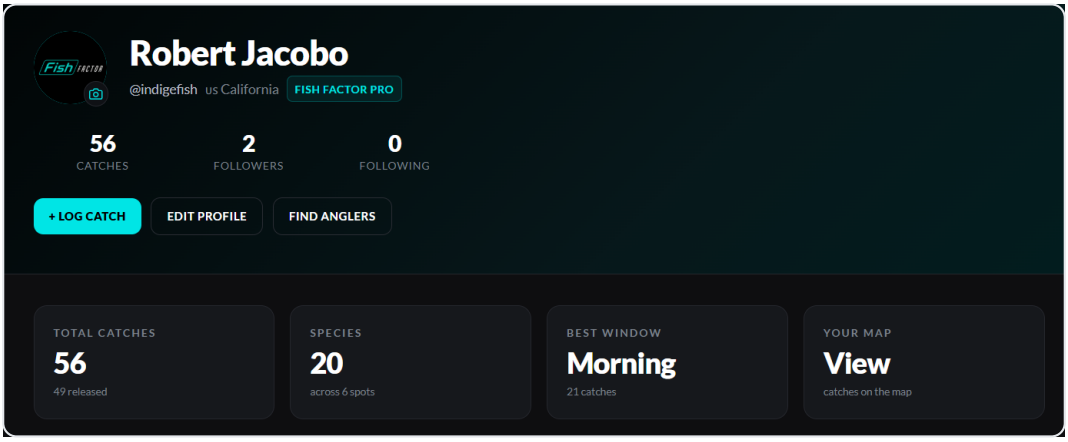
### WHAT IT CHANGES

The product answers the whole question — where, when, what, how, and whether it is legal to keep — instead of only the easy part.

FIELD DATA CAPTURE · 12

# The logbook

Strategically the most important system in the product, and the one users notice least.




### ON THE PAGE

A personal record — species, size, place, photograph, position — and, captured automatically, the exact conditions at the moment of the entry.

### WHY IT WORKS

Built offline-first. An entry made far from signal saves to the device instantly and uploads itself when the network returns, retrying on reconnect and on app focus.

### WHAT IT CHANGES

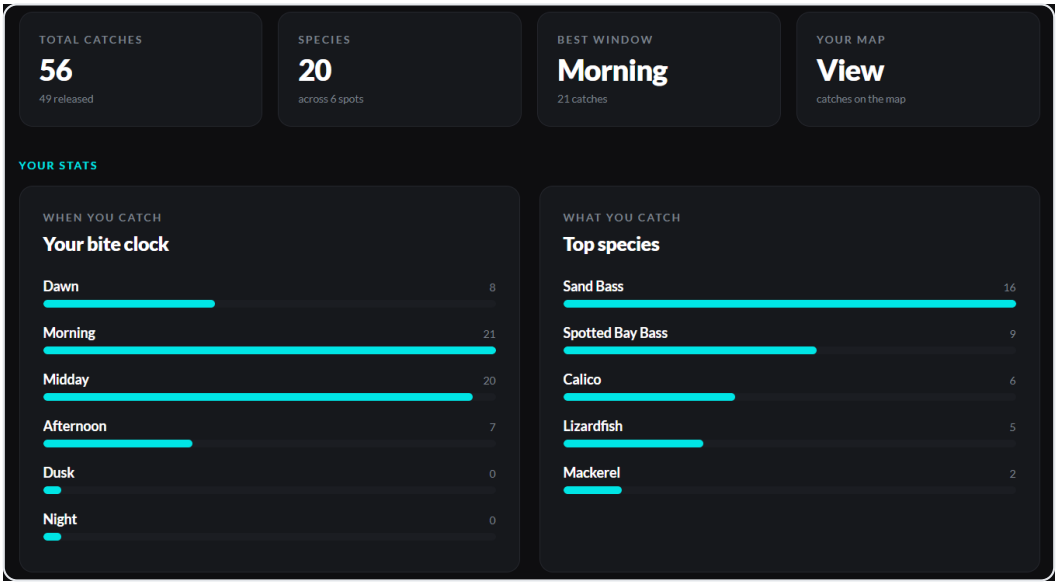
Every entry is a labelled example: these conditions, this place, this time, this result. That dataset compounds, and a competitor entering later cannot backfill it.

**The detail we care most about**

The confirmation never claims more than it can prove. It says “saved to this device” and upgrades to “saved to your account” only once the server has confirmed. The naive version — write locally, fire an upload, show a tick — is what most teams ship, and it loses work while telling the user everything is fine.

# What the data gives back

Collecting outcomes is only worth it if the angler gets something for it.




### ON THE PAGE

The angler's own bite clock by time of day, their top species, their most productive places, and the conditions their catches actually happened in.

### WHY IT WORKS

It is the same analysis the forecast runs, pointed at one person's history instead of the coast. The data they contributed comes back as insight.

### WHAT IT CHANGES

It turns a logbook from an admin chore into a reason to keep logging — which is what keeps the dataset growing.

# The coastal map

Shared results across 250 miles, with precision governed by the user rather than the platform.




**ON THE PAGE**

Every shared result placed on the coastline, filterable by county, over satellite imagery.

**WHY IT WORKS**

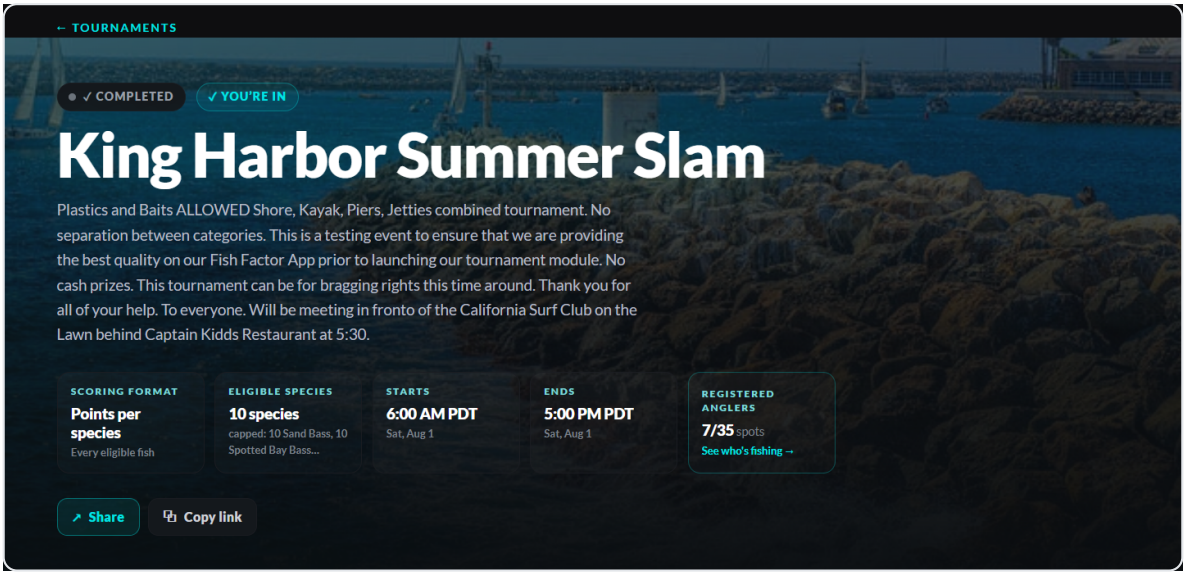
Location precision is tiered. Free accounts see a general area; paid accounts see exact positions — and the person who recorded it decides what is shared at all.

**WHAT IT CHANGES**

Fishing locations are sensitive information. Treating them that way is both the right default and, in this market, a competitive position.

# Self-serve competitions

The deepest subsystem in the platform, and entirely self-serve — users create and run their own events without us.




### ON THE PAGE

Scoring format, eligible species with per-species caps, start and end times, registered anglers against a capacity limit, and a shareable public link.

### WHY IT WORKS

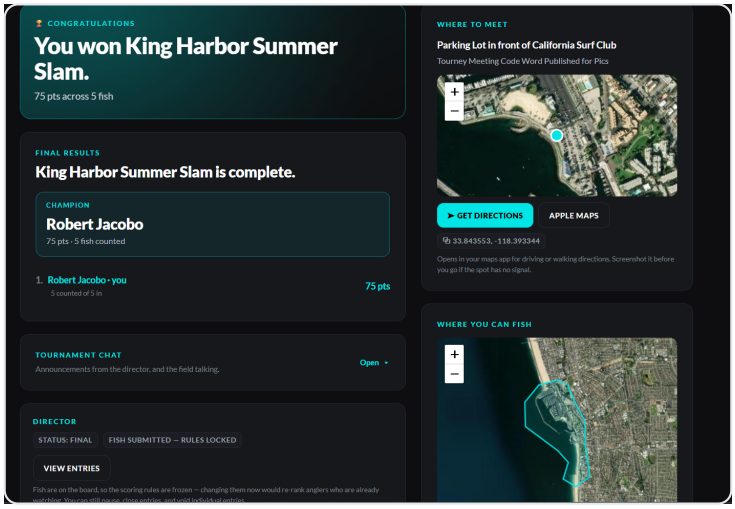
Competitions award things there is only one of, so correctness under simultaneous entry is not optional. Scoring and capacity logic runs in the database, not the application.

### WHAT IT CHANGES

A local business or community group can run a weekend event with no software of their own — a partnership instrument, not just a feature.

# Run through to a result

Anyone can start a competition. Finishing one correctly is the hard part.




### ON THE PAGE

Final standings with the winner's points and fish counted, a meeting point with directions, the fishable boundary drawn on satellite imagery, and a moderated chat.

### WHY IT WORKS

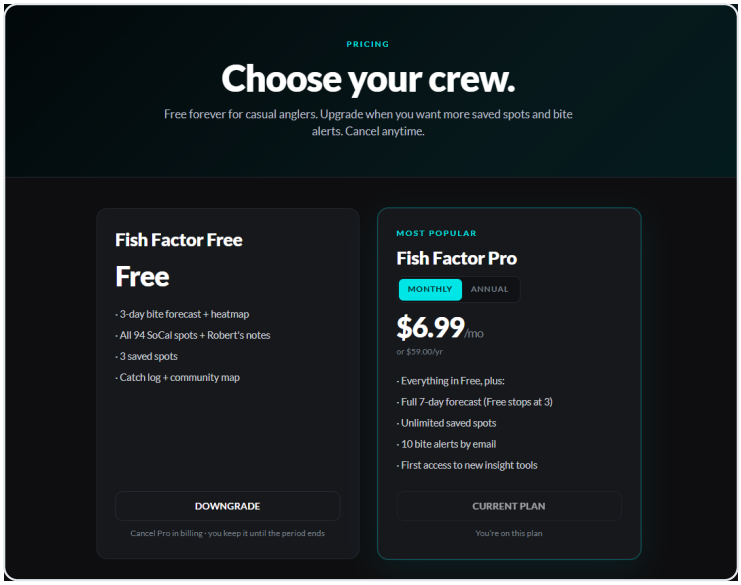
Once fish are on the board the scoring rules freeze. Changing them afterwards would re-rank anglers who are already watching, so the system refuses.

### WHAT IT CHANGES

The organiser gets a defensible result. That is what makes an event repeatable rather than a one-off argument.

# Subscriptions

A real commercial system — trials, billing cycles, entitlements and the edge cases that decide whether a customer trusts it.




### ON THE PAGE

Two plans, a free trial, monthly or annual billing, and a plain statement of what each tier includes. Cancelling is one click and the plan runs to the end of the period.

### WHY IT WORKS

The free tier is genuinely useful rather than a crippled demo — the whole catalogue, a real three-day forecast, the full logbook. The people who most need the answer are often least able to gamble on it.

### WHAT IT CHANGES

Granted and complimentary accounts are recorded separately from paid ones, so a billing event can never revoke access that billing did not grant.

ENGINEERING · 18

# When the federal data source went down

In July 2026 the national service providing tide predictions failed across the country. Every product built on that single endpoint lost its most important input.

### WHAT HAPPENED

The endpoint returned no data for every station, nationwide. Observations and reference information stayed available. Most tools built on it showed stale numbers or fell back to a placeholder.

### WHAT WE UNDERSTOOD

Tide predictions are not measurements. They are computed from each station's published harmonic constituents — and those constituents are served by a different interface that never went down.

### WHAT WE BUILT

A local harmonic solver. The platform now retrieves each station's constituents, caches them for thirty days and computes the predictions itself, with the correct astronomical arguments, nodal corrections and datum offset.

### Why this is the story we tell first

The easy response was to display stale data and wait for a federal agency to fix its own service. The response delivered was to understand the underlying science well enough to replace the dependency entirely. The platform has had no single point of failure on its most critical input since.

### The general principle

Every external dependency is treated as something that will eventually fail. Each has a defined degraded state, and none of them can take the core product down. That posture is cheap to design in at the start and expensive to retrofit after an outage.

# How we keep a system honest

The failures that matter most in a live system are not the loud ones. They are the operations that fail quietly and report success. Our standards are built around finding those before a user does.

### NOTHING FAILS SILENTLY

Every database operation, external call and scheduled job is checked and reports its own outcome. A job that does nothing must say so, with the reason, rather than returning a healthy result and an empty count.

### THE INTERFACE STATES THE TRUTH

A confirmation is shown only once the system has confirmation. Where an action has not completed, the product says which state it is actually in rather than the optimistic one.

### VERIFIED AGAINST THE SYSTEM

Before any claim is made — in the product, in a report, or in a document like this — it is checked against the running system. Comments and prior documentation are hints, never authority.

**Removing a feature because the measurement said to**

The prediction model included one factor added on reasonable theoretical grounds. We measured it across 462 location-by-species combinations and it returned exactly the same value every time — a constant wearing a variable's clothing, adding a flat amount to every score while occupying a line in every user-facing explanation. It was switched off and folded into the base value. The capability remains behind a flag, ready if the evidence ever changes.

**Why we work this way**

Systems that fail loudly get fixed. Systems that fail quietly get believed — and a number that is quietly wrong is more dangerous than a page that is visibly broken, because somebody will make a decision on it.

# Data stewardship

How user information is held, who can reach it, and what we will not do with it. For an organisation handing its data to a vendor, this is the section that matters most.

### Least privilege, enforced by the database

Access rules live in the database itself rather than only in application code, so a request can only ever reach the records it is entitled to — regardless of which part of the system is asking.

### Privileged operations stay server-side

Anything that grants access, changes entitlements or touches another user's record runs only on the server under controlled credentials. Nothing privileged is reachable from a browser.

### Automated jobs fail closed

Scheduled and administrative endpoints refuse to run without valid authentication — including when the credential is missing entirely. The safe failure is to do nothing.

### Sensitive data is separated by design

Precise location is held apart from anything displayed publicly, and the person who recorded it controls what is shared. Defaults favour privacy over visibility.

### What we will not do with a client's data

User information is not sold, is not shared with any advertising network or data broker, and is not used to train anything outside the client's own product. Accounts and infrastructure are registered in the client's name. If a relationship ends, everything transfers back in full — ownership is the deliverable, management is the service.

### On describing security publicly

This document describes our practices and deliberately not the internal structure of any live system. Specific configurations, schemas and defensive details belong in a technical review under agreement, not in a portfolio. We are glad to walk a client's technical staff through the detail directly.

SECTION 21

# How a build like this runs

The same sequence we apply to any platform engagement, whatever the subject.

| PHASE | WHAT HAPPENS |
|---|---|
| **Discovery** | The subject is mapped and every factual claim is traced to a source and marked verified or not. Where sources disagree, we bring the conflict to the client rather than choosing the number we prefer. |
| **Model** | The logic that decides what the product says is built and measured before any interface. A polished screen over a model that does not work earns trust it cannot keep. |
| **Data** | Every external source is integrated with its own cache, its own failure mode and a defined degraded state. No source is permitted to take the product down. |
| **Build** | The public surface first, so there is something real to react to within weeks. Anything holding user data second, because that is never rushed. |
| **Prove** | Security review, a backup restored onto a clean machine, a failure pass, and verification against the running system rather than the documentation. Gates that must pass before user data goes in. |
| **Operate** | Hosting, monitoring, backups, scheduled jobs, updates, content changes and support under one agreement, with one number to call. |

### What we will not do

We do not present an estimate as a certainty, build a feature because a competitor has one, or describe planned work as though it were finished. We do not weaken security or privacy controls to ship faster, and we do not put a number on a screen we cannot explain the origin of. Where we are uncertain, the product says so — because a tool that overstates is found out, and everything else depends on not being.

SECTION 22

# What this demonstrates

Fish Factor is included as evidence of delivery capability. The subject matter is incidental — what it proves is the ability to take a problem from nothing to a live, maintained system and make the judgement calls that decide whether it survives real users.

### Full-stack delivery, one accountable party

Concept to live operation with no gaps and no subcontracted modules. Product design, architecture, data integration, interface, billing, communications, deployment and ongoing operation.

### Public and scientific data pipelines

Six external feeds normalised, cached and made fault-tolerant — including replacing a federal computation locally when the agency's own service failed nationwide.

### Systems that work in the field

Offline-first capture, reliable sync, honest state reporting, and interfaces designed for sunlight, one hand and a poor connection.

### Commercial and operational maturity

Subscription billing with real edge cases, compliant email, versioned releases, scheduled jobs, internal reporting and a rollback path.

**PLATFORMS ON PUBLIC DATA**

Government, scientific and environmental feeds turned into something a non-specialist can act on.

**FIELD AND LOW-CONNECTIVITY SYSTEMS**

Anywhere users are outdoors, away from reliable signal, and cannot afford to lose what they entered.

**EXPLAINABLE DECISION TOOLS**

Any subject where a system's output has to be defensible to someone who will check it against reality.

SECTION 23

# Ownership, stated plainly

The common pattern for a Native-owned brand is licensing somebody else's platform — paying rent on technology it does not control, and generating data that accrues to a stranger. This is the other thing.

### Native-built, not commissioned

The architecture, the prediction model, the database, the billing integration and the deployment pipeline were produced in-house by a Native-owned company. Nothing here was bought in and rebranded.

### The asset stays with the owner

Domain, hosting, database, repository and customer relationships are held in the client's name. The intellectual property and the revenue stay with the business whose name is on it.

### Verified rather than asserted

This is live software serving real users and taking payments. It can be opened, used and inspected — a different kind of evidence from a capability statement on a form.

### Built by someone in the field

The local knowledge in the location catalogue came from a Native angler who fishes that coast himself. That is why the detail is right, and it is not something a national competitor can hire for.

### Why we say it this way

We make no claim that traditional knowledge is encoded in the model, and we do not dress the work in that language. The true version is simpler and stronger: the person who built this has a relationship to this place, the technology is owned by the people whose name is on it, and the whole thing can be checked by anyone who wants to.

SECTION 24

# Working with Indigenous Media Group

INDIGENOUS MEDIA GROUP

A Native-owned multimedia and digital strategy company built to serve Tribal Nations, Indigenous communities and the organisations that work with them.

### WHAT WE DO

Tribal government websites, secure member and employee portals, commercial software platforms, digital strategy, branding, photography and video, podcast production and digital storytelling.

### HOW WE WORK

In-house, end to end. Design, build, content, hosting and ongoing management from one team, so there is one point of accountability rather than four vendors pointing at each other.

### HOW WE CHARGE

One monthly fee covering hosting, security, backups, monitoring, content changes and support. No hours to count and no overage invoices.

### WE VERIFY BEFORE WE CLAIM

Every factual statement in this document was checked against the live system before it was written down. What is built is described as built; planned work is labelled as planned. We would rather hand over a thinner document than one that does not survive the first click.

### OWNERSHIP IS THE DELIVERABLE

Accounts in the client's name, credentials held in a managed vault so no staff member has to administer anything, and a full transfer within thirty days at no charge if the relationship ends. Management is the service. Ownership is never the leverage.

## Let's talk about what you need built

Whether that is a Tribal government platform, a member portal, or commercial software like the one in this document — the first conversation is free and carries no obligation. We will tell you honestly what we think you need, even when that is less than we would like to sell you.

**Robert Jacobo III**
President
Indigenous Media Group
rjacobo@4bmg.org
indigenousmediagroup.org