

EXECUTIVE BRIEF

Franchise Performance Intelligence

A governed agent layer that replaces revenue fog with revenue clarity, built as an installable HubSpot app and running in a live portal today.

PRODUCT	STAGE	NEXT MILESTONE	PREPARED BY
Franchise Performance Intelligence (FPI)	Live in a HubSpot portal since 24 July 2026	UNBOUND, Boston, 16 to 18 September 2026	Denamico Franchise Practice

01 The commercial problem

A franchisor sells a system. A franchisee buys the right to run it. What sits between them is a relationship with almost no direct authority, because the franchisee is an independent business owner who has put their own capital at risk. The franchisor cannot instruct; it can only influence. The instrument of that influence is the franchise business consultant, who carries a book of 30 to 150 locations.

Three gaps compound against each other, and each one has a cost that already appears in the P&L:

The gap	What it looks like today	What it costs
Timing	A location is noticed through lagging indicators: royalty reports, quarterly reviews, a location that goes quiet.	By the time a consultant knows, the trouble has been running for a full reporting period or more.
Data	Corporate sees royalty revenue because royalty revenue is what gets reported. Cost of goods, labour, occupancy, and marketing spend are collected as a compliance demand.	Late, partial, unreliable submissions, and no benchmark anybody can act on.
Knowledge	In any hundred-location network, several locations have already solved the problem a struggling location is facing. Nothing moves that solution.	The network quietly re-solves problems it has already solved, at full cost, every time.

02 What FPI is

THE ONE-PARAGRAPH VERSION

A franchisee hands over their monthly profit and loss statement and gets something genuinely useful back in the same interaction. That exchange is the whole product. A deterministic scoring pass grades every location each period and flags the ones deteriorating against their own trend and their cohort. HubSpot Breeze agents then draft a specific intervention, grounded in a library of interventions that have actually worked elsewhere, and route it through two independent human approval gates. The approved asset deploys through a native HubSpot workflow, and a writeback measures whether it worked. The network gets smarter every lap.

It is built as an installable HubSpot public app carrying its own data model, so the second franchise brand gets the entire system the moment they install. Four things it deliberately is **not**: a compliance portal (those extract, and extraction is what produces bad data), a leaderboard (ranking named independent operators is the exact fear that stops submission), a dashboard product (a dashboard tells you a number is bad; FPI tells you what to do and checks whether it worked), and an AI product (the agents are the mechanism, never the pitch).

830k	30–150	1	9	0
Franchise establishments in the United States, all sharing one operating shape	Locations in a single consultant's book, the persona the product serves	Full reporting period earlier than a drifting location is detected	Agent tools published and live, backing seven designed Breeze agents	Per-portal provisioning steps for the data model; it ships inside the app

03 One loop, crossing every surface once

The product is a single operating loop. It touches the CRM, the member portal, the agent layer, the automation engine, and the marketing tools once each per lap, then starts again with better information.

Stage	What happens	Who is in the seat
Sense	A point-of-sale feed and the portal's P&L form write Location and Financial Period records. Two idempotency keys make both paths safe to re-run.	Nobody. This is plumbing.
Recommend	A deterministic scoring action grades every location against its own trend and its cohort and flags deterioration. An agent then drafts an intervention grounded in the play library, with citations.	Nobody yet. Everything is a draft.
Govern	A consultant approves in the agent inbox. Where the play touches customers, the franchisee approves in the portal. Two distinct gates, one record, and they never cross.	Consultant, then franchisee owner.
Deploy	The approved stage fires a native HubSpot workflow that clones and adapts a proven asset into the content approval gate, as a draft.	Whoever holds content approval.
Measure	A writeback records the delta against baseline and rolls the result up to the parent play, which is what makes the next lap sharper.	Nobody. This is the learning.

04 The exchange that makes the data arrive

The hardest input in the system is franchisee-level profit and loss data, and it is also the most valuable, because it is the only thing that turns a revenue number into a margin story. Every attempt to collect it as an obligation produces late, partial, unreliable data. So FPI does not collect it as one. The franchisee submits through the member portal, and what comes back immediately is a benchmark of their own numbers against an anonymised cohort, plus a short list of interventions that have moved that metric for locations like theirs. **The submission is the price of admission to something the franchisee actually wants.**

WHY THIS IS DEFENSIBLE

Two loops compound against each other. The submission loop makes the benchmarks sharper: more submissions produce denser cohorts, and denser cohorts make submitting more worthwhile. The outcomes loop makes the play library smarter: every play tried and measured makes the next recommendation more specific.

A competitor can copy the object model in a fortnight. They cannot copy an accumulated record of what was tried, where, for which cohort, and whether it worked, because that record only exists after a network has run laps. **The moat is time, and it starts accruing on the first install.**

05 What changes for the customer

Change	Today	With FPI
When a consultant learns a location is drifting	When the lagging metric arrives, typically a full reporting period late.	The period it happens. Every location is rescored the moment its financials land, and the drift flag fires a workflow.
What the consultant does about it	Starts from a blank page and their own recollection of what has worked.	Starts from a drafted intervention carrying its own track record across the network.
What the franchisee gets for their financial data	A compliance obligation and no return.	A benchmark against an anonymised cohort and a shortlist of interventions that moved that metric for locations like theirs.

The measurement writeback turns this from an assertion into a number. Every deployed play records its baseline, its result, and the delta against the metric it targeted, so after a few laps a franchisor can answer a question they currently cannot: did the programme cause anything.

THE WHOLE ADOPTION ARGUMENT, IN A FRANCHISEE'S WORDS

Your numbers stay yours. What you get back is about your business, not your ranking. No peer is ever named. **Every design decision behind this loop exists to make that true rather than merely stated,** because a franchisee who feels graded stops submitting and the whole loop starves.

06 Why a franchise brand can actually adopt this

For a product that reads independent operators' financial data and can cause marketing to be sent to their customers, governance is not a compliance appendix. It is the part of the design most likely to determine whether a brand can adopt it at all.

The obvious version of this product autonomously optimises franchisees. That version is not merely worse, it is commercially fatal. Franchisees are independent owners, and a system that acts on their business without consent is a liability, not a growth tool.

The governing principle: the system recommends and equips; a human approves. The autonomy setting exists in the codebase and always resolves to gated.

Four gates, and two of them never cross

Gate	Who holds it	What it means
Consultant review	A consultant, in the agent inbox	The agent's draft is good enough to show a franchisee. It does not advance the record toward deployment.
Franchisee consent	The owner or their staff, in the portal	This may go to my customers. Terms name what is sent, to whom, from which brand, and over what period.
Content approval	The brand's existing approvers	Native HubSpot content approval plus subscription and consent enforcement on every send.
The autonomy dial	The franchisor	Configuration, not approval. The branch exists and always resolves to gated.

An unrecognised role is refused rather than defaulted, because guessing wrong either fabricates a consent or discards a real one. **No surface may present a review stamp as an approval.**

Three protections a franchisee will ask about by name

- **No peer is ever named, and no other location's figures ever reach a franchisee's screen.** Reach is re-derived on every request rather than inherited from a login, and a denial is byte-identical to a record that does not exist.
- **A percentile never appears without its cohort size.** Sitting in the 40th percentile means nothing until you know the comparison was against fifty locations rather than five.
- **The anonymity floor cannot be configured away.** The absolute minimum cohort is three and the default is five, and the effective floor is the larger of the two. Below it, the result carries no numbers at all.

Two positions taken deliberately, in advance

Consent. At least one brand in this vertical carries a marketing opt-out class action, which is why consent is load-bearing rather than boilerplate. Consent copy sits in the same block as the control, declining is never blocked by the checkbox, cloned assets are drafts that cannot be published by accident, and the installer halts rather than substituting a subscription type the customer never agreed to.

Regulatory posture. Under the EU AI Act, social scoring and employment decisions are unacceptable or high-risk uses. FPI grades business locations, never people: every tool operates on a location record, none takes a person as a subject, and nothing returns an individual assessment. That was decided in advance rather than found in review.

HOW A RESULT IS ALLOWED TO BE DESCRIBED

Two reporting rules are enforced in the data itself rather than in a prompt, because a rule that lives only in a prompt survives until someone edits it. **Confidence is history, never a score:** a consultant reads "used with 14 locations, moved the metric for 11", never a percentage. **A percentile never appears without its cohort size.** Nobody in the chain, human or model, can quote a precision the sample does not support.

INDEPENDENTLY EXERCISED, AGAINST A LIVE PORTAL

24 of 24 isolation tests passed with none inconclusive, including four consultants' books resolving live and proving pairwise disjoint. **18 of 18 governance-gate tests passed**, including the hard case of a consultant who also owns a location. One guard test was deliberately run against a broken build and failed alone, as it should, because a guard test that has never been watched to fail proves nothing.

07 The economics: built once, installed many times

Franchising is large and structurally repetitive, which is the commercial point. What varies between two franchise brands is the chart of accounts, the fiscal calendar, the word they use for a location, the primary metric, and the brand identity. Every one of those is a configuration row rather than a code change.

SHIPS WITH THE APP Once per developer, forever • Four app objects, 84 properties, eight labelled associations • Two CRM record cards and a settings page • Nine agent tools and three workflow actions • A 30-module member portal theme • Nine native workflow definitions • 46 backend functions across 42 routes • The isolation and refusal contracts	CONFIGURATION One row per brand, never code • The word for a location and for a period • Currency, locale, timezone • Seven chart-of-accounts labels • Fiscal calendar, one of three shapes • Primary metric and variance tolerance • Benchmark floor, which cannot go below three • Logo, three colours, typeface	MANUAL, PER PORTAL The real per-install cost • Creating each agent and its instructions • Binding each tool to each agent • Roughly 10 to 15 interface selections • Enabling three ownership workflows • Creating the subscription type HubSpot exposes no API for any of this. It becomes install-guide work, not engineering.
---	--	--

The install path is twelve checklist items and the runner is idempotent. It never deletes, never retypes a property, and never overwrites a value an administrator edited. It reports per-item status rather than one pass or fail, because a portal without Content Hub Enterprise fails one step while everything else succeeds.

08 Where it stands today

Every figure below was read from the repository, the cloud account, or the HubSpot build API on 4 August 2026, rather than copied from an older document.

Component	State	Evidence
The HubSpot app	Live	Build #15 successful, 28 components, installed in a live portal since 24 July 2026. A real authorisation round trip has completed and real records have been written across all four objects.
The backend service	Live	46 functions and 42 routes, all present and matching the manifest. The rate limiter survived a burst of roughly 530 writes including a back-off.
The agent tools	Live	Nine published, endpoints routed, and reachable through two independent invocation paths so that beta access lapsing is not a single point of failure.
The Breeze agents	2 of 7	Two built and verified end to end against live tool output, one in build, four designed with their instruction blocks written. This is configuration work, not platform work.
The member portal	Built	Eight pages, 30 modules, three serverless endpoints. 52 of 52 structural checks passing.
Quality gates	Green	759 backend tests passing, strict typecheck clean, theme validation blocking on failures, and the two live test plans above.

WHAT THE LIVE AGENTS ACTUALLY DID One agent returned 13 coached locations for a consultant, correctly ranked, and quoted a two-day-old score <i>as of its date</i> rather than as a current figure. A second ran three tools against one location, reported a suppressed benchmark as insufficient reporting rather than as poor performance, and reasoned from weak lead flow with steady conversion to a demand problem rather than a conversion problem, then picked the intervention whose target metric matched. That last step is real diagnostic work, and it is the thing the consultant-as-hero narrative is actually selling.

09 What is open, stated plainly

Ten defects and gaps are documented by name and severity in the full dossier. Four of them are worth an executive's attention, because each one carries either a date or a decision.

Open item	Severity	Why it matters at this level
No batch benchmarking pass	Last blocker	Cohort comparison currently returns below-floor everywhere, so briefings say there are too few comparable locations. That is the anonymity floor working correctly, and it is also what an audience reads as failure. Scoped, and the per-location writer already exists.
One isolation path fails open	Highest	In the one module where a bug would be a data breach rather than a defect. Latent rather than universal, the fix is on a branch, and the verification takes two minutes.
Scoring writer not yet deployed	Closed in code	Written, merged, and tested. Until it deploys, six properties stay empty in every portal and the drift-alerting workflow has nothing able to fire it.
Per-install agent configuration	Structural	Not fixable by us. It is roughly 10 to 15 manual selections per customer portal and it is the one cost the built-once rule cannot remove. It should be priced into every implementation.

10 The road to September, and the offer past it

Six items sit between here and the pitch: batch benchmarking, the fail-open verification, the third agent, deploying the scoring action, the custom domain, and configuring agents four through seven. Only the first two are engineering risk. The rest is sequencing.

THE STRATEGIC READ

The contest is a deadline, not the goal. What survives the event is a productised offer: an installable HubSpot app that a franchise brand buys, that Denamico implements, and that gets better for every brand as each one runs laps. The marketplace listing is a form, the tool ids, and a demo video. The second franchise client installs the same app.

HubSpot has told Denamico that its top 2026 partners lean into agentic frameworks, apps, and case studies. Entering advances partner standing whether or not it wins, which makes the effort defensible on its own terms.

11 Decisions that belong to leadership

- **Named-customer permission.** A public pitch that names a live customer, and the data access behind it, is a relationship decision rather than a build decision. The fallback is anonymised but real, and it is materially weaker on the story-and-stage axis.
- **Marketplace posture.** Distribution is marketplace and unlisted, which permits 25 installs before a public listing is required. Listing is a form when it is wanted, and it is not on the critical path.
- **Implementation pricing.** The manual per-portal configuration above is real labour on every install. It should appear in the statement of work rather than being absorbed.
- **Autonomy thresholds.** The dial ships gated. Whether any play type ever graduates toward automatic execution, and on what evidence, is a franchisor-facing policy decision that should be made deliberately rather than by accretion.
- **Ownership of the cohort roll-up.** Whatever writes the best-fit cohort property steers which locations every intervention is offered to. The recommendation is that it be computed from measured outcomes rather than written by an agent.

28	84	759	42	12
HubSpot app components shipping in the current build	Properties across four objects, shipped as code inside the app	Backend tests passing, with a strict typecheck clean	Live isolation and governance tests passed, none failed	Checklist items in an idempotent, per-item install path

Verified 4 August 2026 against the repository, the cloud account, and the HubSpot build API.

Prepared by Denamico's Franchise Practice from the Franchise Performance Intelligence submission dossier and the system design of 18 July 2026. Every claim in this brief is carried in the full dossier with its supporting evidence. Where a figure here disagrees with an earlier document, the dossier of 5 August 2026 is correct.