Koi Labs × Hukam Capital — Senior-Living AI

The product catalog — what we'd build

Prepared with Gurjote Sethi (Hukam Capital Group) and David Koi (Koi Labs). June 26, 2026.

Every funded player in this space sells the same thing — fall-and-clinical monitoring to the Director of Nursing. We'd build for the *owner's* outcomes instead: faster lease-up, leaner operations, and a stronger exit. The whole stack is camera-free by design — no cameras in resident rooms, no wearables — and we'd prove it in one Hukam building first. What follows is the catalog: seven products we'd actually build and deploy as units. The dozens of individual capabilities are features inside them.

The seven products at a glance

#	Product	What it's for	Our honest read
1	**Sensing & Signal Platform**	The data foundation everything else runs on	Table stakes — but the strategic high ground
2	**Fall Detection & Safety Response**	Camera-free fall coverage, halls to bathrooms	Table stakes; our edge is the camera-free bathroom solve
3	**Resident Wellness & Early-Warning**	Ambient early-warning flags for staff	Under-served, but a commodity baseline + a hard reg ceiling
4	**Family Portal**	The family peace-of-mind app + revenue rail	Strong near-term build; the value is the rail, not the app
5	**Staff & Operations Management**	Lower labor cost, lighter staff workload	Wedge is the care-to-labor link; the paperwork side commoditizes
6	**Operator & Investor Analytics**	The platform in NOI / exit-value terms	**THE wedge — actually differentiated**
7	**Portfolio Intelligence & Data Moat**	The cross-building data flywheel	**The real moat** — but the slowest to mature

Products 1 and 2 are the foundation and the table-stakes safety layer. Products 3, 4, and 5 are genuinely valuable but each has a commoditizing side, so we don't oversell them. Products 6 and 7 are the two wedges no one else is built around — the reason this is a stronger play for an owner than for a software vendor.

The products

1. Sensing & Signal Platform

The camera-free sensing layer and per-resident data store everything else runs on.

What we'd build. The always-on, camera-free sensing layer and per-resident data store every other product reads from. We deploy off-the-shelf motion, door/contact, and optional bedside radar in private rooms — cameras only in common areas — then turn the raw events into a per-resident behavioral baseline with a daily anomaly score. Raw streams are processed on-device at the building and never leave it; only structured, de-identified signals flow upward.

Features

- Per-resident activity baseline + daily anomaly score (deviation from the resident's own 30–45-day norm)
- Sleep & circadian-rhythm tracking (in-bed window, restlessness, night trips, day/night pattern)
- Activity-of-daily-living inference from ambient sensors (eating, dressing, bathing, toileting)
- Bathroom frequency + duration trend (counts and durations only, handled with care)
- Social-engagement / withdrawal metric (common-area dwell, alone vs. group, mealtime attendance)
- Per-resident vitality index (the composite wellness rollup that feeds the family + operator views)
- Sensor ingestion, normalization & the per-building signal store (the integration backbone)

Build difficulty — Medium. The sensing and integration are off-the-shelf; the hard part is that 30–45-day baselines plus high resident turnover leave many residents perpetually in "cold start."

Hardware. Room motion (passive-infrared) + door/contact sensors (incl. fridge/cabinet); common-area cameras (counts and dwell only, no identity stored); optional bedside mmWave radar (a contactless radio-frequency sensor); a per-building edge gateway. No in-room cameras.

Servers — edge / on-prem first. Raw sensor streams are processed on the building's own hardware (event-only retention); the per-building signal store stays on-prem or in a single-tenant private cloud; only de-identified signals sync upward. *Why:* raw sensing is sensitive health data, and "processes on-device, no cloud PII" is itself the privacy selling point.

Serves. Operator, Family, Resident.

Our read. Table stakes — the sensing itself is commodity. But it's the data layer every other product reads from, so owning it is the strategic high ground.

---

2. Fall Detection & Safety Response

Full fall coverage — halls to bathrooms — without a camera in any resident's room.

What we'd build. Full fall coverage without a single camera in a resident's room. Common-area, hallway, and exterior cameras run on-device computer vision, while mmWave radar covers the bedroom and bathroom — where most serious falls happen and cameras are forbidden. Every detection routes to the

nearest staff with escalation and acknowledgement timestamps, and the same layer catches bed/chair pre-fall egress and tagless wrong-hours exits.

Features

- Common-area fall detection (halls, lounges, dining, exterior) via on-device computer vision on existing cameras
- In-room & bathroom fall detection, camera-free via mmWave radar (the central hard problem)
- Bed / chair-exit pre-fall alert (the get-there-before-the-fall egress signal)
- “Long-lie” / time-on-ground minimization + an automated fall-response workflow (alert nearest staff, escalate, log)
- Tagless elopement / wrong-hours exit detection (no wearable tag to lose or remove)
- In-room wandering/pacing + hallway near-miss/instability detection (stumbles, furniture-grabbing)

Build difficulty — High. Camera-free in-room and bathroom fall detection via radar is the central unsolved problem and the headline liability; common-area computer vision is the easy, table-stakes part.

Hardware. Common/hall/exterior cameras + on-device computer vision; mmWave radar in bedroom + bathroom; bed-presence/bedside radar; door sensors (elopement); motion sensors (wandering); nurse-call integration + a staff app.

Servers — edge / on-prem mandatory. Raw video and radar are processed on-device — event clip only, never a livestream to the cloud — and alert routing stays local for life-safety latency; only event metadata + response-time data reach the cloud. *Why:* shared-space video privacy, sub-second alert latency, and the fact that an alert-routing failure is itself a liability event.

Serves. Operator, Family, Resident.

Our read. Table stakes, and the most crowded, best-funded fight in the category. Our only real edge here is the camera-free bathroom/bedroom solve — we deliberately don’t try to out-monitor the incumbents.

---

3. Resident Wellness & Early-Warning

Early-warning health signals from ambient behavior — as flags for staff, never diagnoses.

What we’d build. The predictive layer that runs on top of the sensing platform. It turns weeks of ambient behavior into early-warning flags — rising fall risk from gait-speed decline, a urinary-tract-infection screen from bathroom patterns, cognitive-decline and dehydration trends, and a fused “something’s off” signal — each delivered as a prompt for a human to assess, never a diagnosis. It also tracks recovery after a fall or hospitalization against the resident’s own baseline.

Features

- Fall-risk forecast from gait-speed decline + post-fall / post-hospitalization recovery tracking
- UTI (urinary-tract-infection) early-warning screen (bathroom-visit stats + nighttime usage)
- Cognitive-decline trend from long-run mobility frequency + stability
- Dehydration / reduced-intake proxy (drop in kitchen/fridge access or dining attendance)

- Sleep-disturbance & nighttime respiratory / heart-rate signals via bedside radar
- “Something’s off” composite (multiple signal deviations fused into one early-illness flag)
- Medication-routine adherence proxy (missed med-cabinet access — a proxy, not proof)

Build difficulty — High. The decline / UTI / cognitive models are research-grade with real accuracy ceilings (UTI screening tops out around 75% sensitivity), and any “detect or predict disease” claim risks pulling the software into FDA medical-device regulation — so it must stay screening-only.

Hardware. Reuses the platform’s room motion + door/contact sensors; adds bedside mmWave radar where nighttime vitals or sleep signals are needed. No new in-room cameras.

Servers — edge for the raw radar physiology; model inference on-prem or in a HIPAA business-associate-covered private cloud; only de-identified trends reach the cloud. *Why:* vital-sign and decline inference is sensitive health data and the screening claims must be auditable — keep the raw signal at the edge.

Serves. Operator, Family, Resident.

Our read. Genuinely under-served — few incumbents productize the predictive-from-ambient angle — but it runs on a commodity baseline and under a hard regulatory ceiling (screening flags, never diagnosis). Real, not a slam dunk.

---

4. Family Portal

The family peace-of-mind app — and the venture’s recurring-revenue rail.

What we’d build. The family-facing app and the venture’s recurring-revenue rail. It turns a resident’s signals into a plain-language daily/weekly wellness narrative, proactive “good day” pings, and operator-gated anomaly alerts, plus natural-language Q&A (“Has Dad been eating?”) answered strictly from structured signals. It doubles as a privacy-first tour asset — “we watch over your parent without a camera in their room” — and gives residents consent controls over what family can see.

Features

- Plain-language daily/weekly wellness narrative + proactive reassurance pings (“walked more this week”)
- Natural-language Q&A about a resident (retrieval-only — the AI answers solely from stored signals, with “ask staff” routing)
- Family alert on a meaningful anomaly (operator-controlled thresholds)
- Privacy-first tour / lease-up demo experience (“no cameras in the room”)
- Family peace-of-mind subscription (the recurring-revenue rail the venture owns end-to-end)
- Resident-governed sharing & consent controls (dignity + consent)
- *Out of scope (listed, not built):* a resident-facing in-room voice companion — blocked by all-party-consent wiretap law

Build difficulty — Medium. The narrative and pings are easy; the risk is the natural-language Q&A — it needs strict retrieval-only scoping and refuse-and-route guardrails so it never hallucinates a health answer.

Hardware. None of its own — a family mobile/web app that reads from the sensing platform and the wellness layer.

Servers — cloud app for delivery; the AI runs in the cloud but retrieval-only over the resident's structured, consented signals (never raw sensor data), behind a HIPAA business-associate-covered endpoint if any health data is involved. *Why:* family delivery is internet-facing, but health-question hallucination and consent demand strict scoping — and the model never sees raw resident data.

Serves. Family, Resident, Investor, Operator.

Our read. A strong near-term build and the recurring-revenue rail, and family conversational AI is comparatively unbuilt by the sensor incumbents — but family apps are a crowding space. The durable value is the revenue rail and the consent/tour story, not the app by itself.

---

5. Staff & Operations Management

Turns sensor signals into lower labor cost and lighter staff workload.

What we'd build. The operations layer that turns sensor signals into less work and lower labor cost — the largest line in the building. It ranks alerts to fight alarm fatigue, auto-writes shift handoffs, maps staffing to when and where activity clusters, verifies rounds and response times, and tracks staff and maintenance labor utilization. AI assistants draft care notes and audit-prep, and the layer writes back into the building's electronic health record (EHR) system of record rather than replacing it.

Features

- Priority alert feed (rank anomalies by urgency, suppress alarm fatigue)
- Shift-handoff auto-summary (what changed for each resident overnight)
- Staffing-to-need heatmap + agency-labor / overtime reduction tracking
- Response-time analytics + round/check verification (alert → acknowledge → resolve; was the 2-hour check done)
- Staff & maintenance labor-utilization tracking (time-on-task vs. idle, by zone and shift)
- Care-note / incident-report drafting + compliance / audit-prep assistant (AI drafts, staff edit)
- EHR write-back / interoperability (push signals into PointClickCare / MatrixCare / August Health)
- Staff natural-language query + agentic anomaly triage ("show declining residents"; human-gated)

Build difficulty — Medium core. The hard pieces are EHR write-back into the incumbent systems (the #1 interoperability barrier operators cite) and human-gated agentic triage; labor-utilization tracking also needs a worker-monitoring policy.

Hardware. Reads the platform's signals; adds common-area cameras / zone sensors for labor-utilization and optional staff badges; a staff mobile app; nurse-call + EHR integration.

Servers — hybrid. The priority alert feed runs at the edge for latency; AI drafting + analytics run in the cloud (HIPAA business-associate-covered where health data is touched); EHR write-back goes through the healthcare data-exchange standard (FHIR); worker-monitoring video stays on-prem. *Why:* alerts need local latency, note-drafting touches health data, and worker-monitoring video must stay on-prem for labor-law and labor-relations reasons.

Serves. Staff, Operator, Investor.

Our read. The wedge is connecting care signals to labor cost — workforce-scheduling tools don't touch care data. The AI note-drafting and audit-prep sub-features are commoditizing fast, so we lead with the labor / utilization link, not the paperwork.

---

6. Operator & Investor Analytics

The whole platform expressed in the investor's vocabulary: occupancy, NOI, and exit value.

What we'd build. The wedge no incumbent is built around: the whole platform expressed in the investor's vocabulary — occupancy, net operating income, and exit value. It ties fall-reduction, ER-avoidance, labor savings, acuity-capture, and length-of-stay to dollars and produces the board-ready "what this did to NOI and value" report. Around that sit lease-up / unit-turn intelligence, a vibrancy dashboard for tours, move-out-risk prediction, and an owner-facing copilot that answers "which buildings drove the falls drop?"

Features

- NOI & valuation-impact model + board-ready "what this did to NOI and value" report
- Occupancy, lease-up & unit-turn / move-in-readiness intelligence
- "Vibrancy" / community-health dashboard (aggregate activity + social engagement as a lease-up sales asset)
- Length-of-stay / move-out-risk prediction (catching decline-driven churn before it hits the rent roll)
- Care-acuity capture / level-of-care justification (objective data backs the billing tier)
- Liability / risk-reduction quantification (fewer falls → lower insurance & litigation exposure)
- Comparable-building benchmarking (this building vs. portfolio norms)
- Operator/owner analytics copilot + an investor-memo / diligence drafting agent

Build difficulty — Medium. Most of it is defensible analytics over de-identified + financial data; the Hard, higher-liability sub-model is move-out-risk / length-of-stay prediction, and every financial claim must be auditable.

Hardware. None of its own — it ingests occupancy / door-sensor signals (incl. vacant-unit sensors for unit-turn) plus leasing, property-management, and building-financial data.

Servers — cloud. It operates on aggregated / de-identified building signals + financials, so no raw resident health data is required. *Why:* the data here is de-identified and financial, so cloud scale is appropriate; the one hard constraint is a defensible audit trail behind every NOI / valuation claim.

Serves. Investor, Operator.

Our read. THE wedge. Every incumbent's ROI story is clinical; no one packages monitoring + operations into an occupancy / NOI / exit narrative for a buy-improve-sell owner. This is the product that is actually differentiated.

---

7. Portfolio Intelligence & Data Moat

The cross-building data flywheel that compounds as Hukam buys, improves, and sells.

What we'd build. The compounding, cross-building asset — the reason this is stronger for an owner than for any software vendor. As Hukam instruments, stabilizes, and resells buildings, de-identified outcome data pools into a flywheel that makes every model better and every next building lease up faster and exit higher. It benchmarks buildings against each other, packages a repeatable instrument → baseline → optimize → resell playbook, and conveys an "AI-instrumented, NOI-optimized" diligence asset at sale.

Features

- Portfolio benchmarking (rank buildings on safety, response time, efficiency, vibrancy)
- Model-improvement flywheel (more buildings → labeled outcomes → better fall / decline / anomaly models)
- Transferable acquisition playbook ("instrument → baseline → optimize → resell", repeatable per building)
- "AI-instrumented building" exit premium (the model + data history conveys at sale as a diligence asset)
- Underwriting intelligence (use operating data to underwrite the next acquisition better)
- Design feedback to acquisitions (which layouts/units drive falls or isolation → what Hukam buys / renovates)

Build difficulty — High. The flywheel needs many buildings × confirmed-outcome labels × years, plus heavy cross-building de-identification and data governance — the slowest capability to mature.

Hardware. None of its own — a central aggregation + machine-learning layer over every building's de-identified data.

Servers — cloud, multi-tenant, but fed only by de-identified cross-building aggregates with data-use agreements and strong security, built on portable standards so it conveys cleanly at building sale. *Why:* pooling raw resident health data across buildings isn't allowed without de-identification and consent; the moat lives on de-identified, labeled outcomes — and state health-data law (with private lawsuits in the strictest states) makes governance non-optional.

Serves. Investor, Operator.

Our read. The real moat — structurally stronger for an owner who keeps the data across a cycling portfolio than for any pure-software vendor. The catch: it's the slowest to mature, so it's a later-phase prize, not a day-one build.

How we'd build it

What gets installed in a building

It all runs on one off-the-shelf sensing layer, installed once per building. Each resident unit gets a motion (PIR) sensor, a door/contact sensor, and bedside-and-bathroom mmWave radar that detects falls without a camera — and never a camera or wearable in the room itself. Cameras stay in hallways, common areas, and the exterior; a per-building edge gateway processes the raw streams on-site, so only de-identified signals leave the building. That one layer feeds all seven products — and four of them (the family portal, owner & investor analytics, the portfolio data moat, and most of staff & operations) are software with little or no hardware of their own.

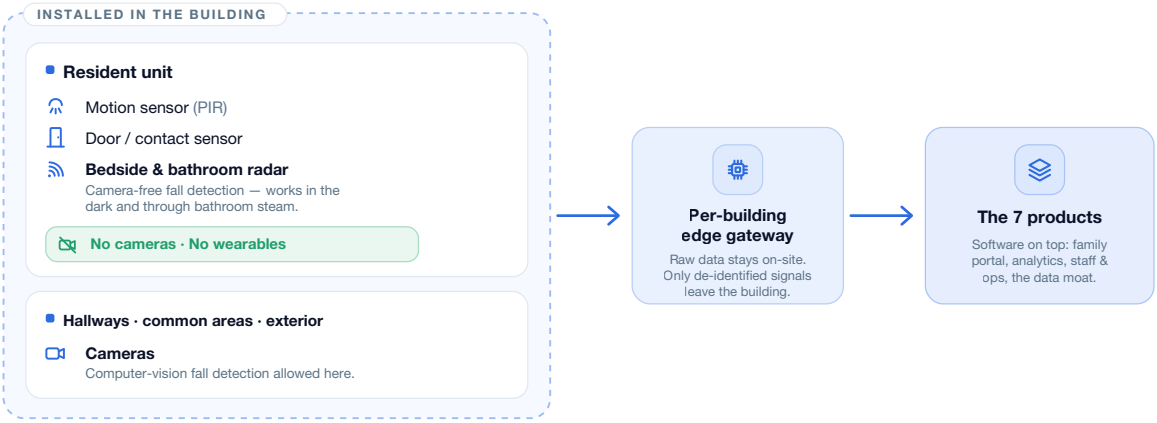

What gets installed in a building: one sensing layer per building — a motion (PIR) sensor, a door/contact sensor, and bedside/bathroom radar in each resident unit, with cameras only in hallways and common areas — feeds an on-site edge gateway, which in turn feeds the seven software products.

Three principles run through every product:

- **Camera-free by architecture.** No cameras in resident rooms, no wearables. Off-the-shelf mmWave radar covers the bathroom and bedside (where most serious falls happen and cameras are forbidden); motion and door sensors cover the rest of the room; cameras stay in halls, common areas, and exterior only. It's the strongest privacy posture available and a real tour and lease-up differentiator — not just a constraint.
- **Off-the-shelf hardware first, our own only later (if ever).** The defensible intellectual property is the software and the owner/family interfaces, not the sensor. We buy the sensing and build the intelligence. Proprietary hardware comes only from revenue, and only if the margin math demands it.
- **On top of what the building already runs, never a rip-and-replace.** The intelligence layer sits above the existing EHR and nurse-call systems and pushes alerts into the workflow staff already use.

And the path is phased: instrument one Hukam assisted-living building end-to-end as the design-partner proof → package that into a repeatable install → license it to other value-add owners. The first building is the proof, the first revenue, and the seed of the data moat.

What we'd navigate

This is the part that's genuinely ours to handle — the AI, privacy, and liability discipline that makes a system like this safe to deploy. None of it is a blocker; each is a known, navigable design choice:

- **Camera-free privacy.** No in-room cameras, and audio off by default, sidesteps the state "granny-cam" and wiretap-consent rules inside resident rooms — and turns privacy into a selling point rather than a liability.
- **Consent and data residency.** Raw sensing stays on-device at the building; only structured, de-identified signals move upward. Resident consent and de-identification are designed in from day one, because even buildings outside HIPAA are still bound by federal and state health-data law.
- **Liability discipline.** Everything is framed as wellness and alerts that prompt a human — never diagnosis, never a guaranteed outcome. That keeps the system clear of medical-device regulation and the "we'll catch every fall" liability trap, and positions it as a supplement to staff, not a replacement.

These are ours to own, and we know how to handle them.

Next steps

Three things to do together, in order:

1. **Pick the design-partner building.** Choose one Hukam assisted-living building to instrument end-to-end.
2. **Scope the first pilot.** Decide what ships in building one versus later — likely camera-free fall coverage, a behavioral baseline feeding the existing nurse-call workflow, the family app, and an owner-facing readout.
3. **Ideate the product priorities together.** Walk the seven products, rank them against Hukam's actual operating pain, and lock the first build order.

From there: prove it in one building, then decide what to productize and license. That's the whole plan — one building first, built for the owner.