

The Embedded Analytics Playbook for Senior Living SaaS

Turn the data your product already holds into the analytics layer your customers keep asking for. Work through it and you'll leave with an inventory of what your product can already answer, a score for where you sit on the analytics ladder, seven tests to put to any vendor (including us), and a 90-day plan with exit criteria.



A working guide for CEOs, CPOs, and CTOs. About a 20 minute read. Forward it to whoever owns the roadmap.

PLAYBOOK

02

If you only read one page

The argument, end to end. Forward this page to whoever only has two minutes.

1

Your product holds a slice of operational truth about every community it touches that no one else has. Your customers are already assembling that picture by hand, from your exports.

2

They can't build the analytics themselves. Of the 61 largest US senior living operators, **33 have zero data staff**. Whatever analytics they get will come from their software or not at all.

3

Most of your peers can't build it either. Of the 50 mid-market senior living vendors we scored as addressable, **28 have zero data or BI staff and 42 have two or fewer**. Most market "analytics" anyway.

4

The category anchors already built or bought it. For the mid-tier, the choice is to partner for the capability or concede the differentiation in every bake-off.

5

Building it in-house costs **\$150K to \$340K in year one**, 6 to 12 months before customers see anything, and \$470K to \$740K over three years. The visualization layer is the cheap part. Permissions, definitions, and answer accuracy are the expensive part, and they're the part most estimates leave out.

6

The path is a **90-day pilot** with two or three design-partner customers, not a multi-year bet. This guide gives you the inventory, the scorecard, the vendor tests, and the plan.

33 of 61

The largest US senior living operators with zero data staff

42 of 50

Addressable mid-market vendors with two or fewer data or BI staff

6 to 12 months

Before a first in-house build reaches a customer

Where you are, in three numbers

33 of 61

The largest US senior living operators with zero data staff.

When they can't build analytics themselves, the ask lands on their software vendors. That's you.

42 of 50

Mid-market senior living software vendors with two or fewer data or BI staff, out of the 50 we scored as addressable for an embedded analytics layer.

28 of the 50 have none. Most have the word analytics on the homepage today.

6 to 12 months

What a production-grade first build takes before a customer sees anything, per the line-item math in our Buy vs. Build guide.

Add the maintenance tail and the answer-accuracy work and the three-year number is \$470K to \$740K.

If you've read our Buy vs. Build guide, you've seen the shape of this argument. That guide helps you decide whether to build or embed. This one is about what you'd be building or embedding: what data you hold, what it can already answer, how far up the ladder you are, and how to get to the next rung in 90 days without betting the roadmap.

WHAT'S INSIDE

01 Inventory

What your product already holds, and a fill-in table for your own

02 Score yourself on the ladder

Four rungs, sixteen conditions

03 Why this is a product decision

Three forces converging now

04 Three paths up the ladder, priced

Build, partner, white-label

05 The seven tests

To put to any vendor, including us

06 The precedent

White-label isn't exotic in this market

07 The 90-day plan

Three phases, three exit criteria

08 Before it's a build decision

It's a data question

Inventory: what your product already holds

Six categories, the questions operators keep asking each one, and a fill-in table for your own product.

Whatever corner of senior living you serve, your product holds a slice of operational truth no one else has. The table below maps six product categories to the questions operators keep asking that category's software, what you typically already hold, and what other system the answer needs. The questions come from what end users in this category keep asking for after their software already shipped reporting, which is the pattern worth paying attention to.

YOUR CATEGORY	THREE QUESTIONS OPERATORS KEEP ASKING YOUR SOFTWARE	WHAT YOU ALREADY HOLD	WHAT THE FULL ANSWER JOINS TO
 Workforce and scheduling	- Which communities ran overtime against budget last pay period, and why? - What share of hours went to agency staff, by community and by shift? - What did labor cost per resident day, this month against last?	Schedules, punches, overtime, agency hours, open shifts, fill rates	Census (EHR or billing) for per-resident-day math; the budget (finance) for variance
 EHR and clinical	- How has acuity mix shifted by community over the last two quarters? - What's our fall or incident rate per 1,000 resident days, and where is it moving? - Are med passes on time, by shift and by community?	Assessments, acuity scores, incidents, med pass records, care plans	Staffing (workforce) for acuity-to-staffing ratios; census for the denominator
 Engagement	- Which residents' participation dropped this month? - Does participation predict length of stay or move-outs? - How engaged are families through the app, by community?	Activity attendance, participation logs, family app events, messages	Census and move-out records (EHR or CRM) to connect engagement to retention
 Safety and e-call	- What's our response time by shift, by community, and how does night shift compare? - Which residents are generating repeat alarms? - Did response time move after the staffing change?	Alarms, events, response times, device status	Workforce (who was on shift) and clinical (falls and incidents)
 CRM and marketing	- What's our inquiry-to-tour-to-move-in conversion, by lead source? - How many days from inquiry to move-in, and is it lengthening? - What does the pipeline say census will be in 90 days?	Inquiries, tours, deposits, move-ins, lead sources, notes	Occupancy (EHR or billing) for the census forecast; marketing spend for cost per move-in
 Dining	- What's food cost per resident day by community? - Where is waste concentrated? - Are dietary orders being met, and where are the exceptions?	Menus, orders, purchasing, production, diet orders	Census for the denominator; clinical for diet compliance

Facilities and CMMS vendors fit the same pattern: work order backlog, preventive against reactive maintenance ratios, and cost per unit, all joined to census and to the capital plan.

Notice the fourth column. Almost every question operators actually ask crosses at least one system boundary. Labor against census. Funnel against occupancy. Engagement against move-outs. Each of your categories is one silo of the operator's full picture, and today your customers stitch that picture together the hard way: exports, spreadsheets, and a regional director's Sunday evening.

That stitching work is a product opportunity wearing a support-ticket costume.

Your inventory

Fill this in for your own product. Ten rows is enough to see the shape. The rightmost column is the one to pay attention to.

	QUESTION CUSTOMERS ASK	WE HOLD THE INPUTS Y / partly / N	DEFINITION WRITTEN DOWN AND OWNED Y / N	NEEDS ANOTHER SYSTEM which one	HOW CUSTOMERS GET IT TODAY in-app / export / ticket / they can't
1					
2					
3					
4					
5					
6					
7					
8					
9					
10					


HOW TO READ YOUR RESULTS

Count the rows marked export or ticket. Every one of those is a request for a product you haven't shipped yet, and it's already generating support load. Count the rows where the definition isn't written down. Those are the ones that will produce two different numbers on two different screens the day you ship anything, and they're where a pilot should start. Count the rows that need another system. That's the join list, and it decides how far up the ladder you can go on your own data alone.

Score yourself on the ladder

Four rungs, four conditions each. You're on a rung when all four hold. Sixteen ticks tell you where you sit and what's in front of you.

Most senior living vendors describe their analytics in the same words. The ladder below separates what the word means. Four rungs, four conditions each. You're on a rung when all four hold. Partial rows show you the work in front of you.



RUNG	WHAT IT IS	PRICING POWER
1 Canned reports	A fixed set of reports, per community, usually exported	None. Table stakes
2 Dashboards	Cross-community rollups with drill-down	Modest. Where most analytics marketing points today. High maintenance load
3 Cross-source insight	Your data joined with the operator's other systems	Real. Premium tier or per-community add-on. This is where operators start treating you as the decision layer
4 Conversational analytics, under your brand	Plain-English questions, trustworthy answers, inside your UI	Top of the price card, and the strongest retention story in your deck. This is what customers now mean by analytics, because they've used ChatGPT

The sixteen conditions

Tick every condition that holds today. Sixteen possible.

1 Canned reports

- ☐ A fixed report library ships with the product
- ☐ Reports are per community, not cross-community
- ☐ Users export to CSV or PDF to do anything further
- ☐ A new report is an engineering ticket

2 Dashboards

- ☐ Rollups across communities and regions
- ☐ Drill-down from a number to its rows
- ☐ Filters by date, region, and community
- ☐ Refreshed on a schedule the customer didn't set

3 Cross-source insight

- ☐ At least one join to a system you don't own (census, labor, finance) is live in production
- ☐ A non-engineer on your team can add a view without a ticket
- ☐ Each metric has one written definition, applied in one place
- ☐ Role-based permissions hold across the joined data, not just your own

4 Conversational analytics, under your brand

- ☐ Users ask in plain English inside your product
- ☐ The same question returns the same number as the dashboard
- ☐ Permissions are enforced when a model, not your UI, writes the query
- ☐ Answer accuracy is tested after every model or schema change, by a named owner

Scoring

0 to 4

RUNG 1

Your exports are your analytics strategy. Your inventory's rightmost column is probably mostly export and ticket.

5 to 8

RUNG 2

You have dashboards and a maintenance bill. Customers still export to answer the next question.

9 to 12

RUNG 3, OR CLOSE

Pricing power starts here. The conditions you're missing are usually the definition and the permissions rows, and they're the expensive ones.

13 to 16

RUNG 4

You're ahead of the category. The question is whether the accuracy-testing row is real or aspirational; see test B in section 05.

WHERE THE FIELD SITS

Of the 50 mid-market senior living vendors we scored as addressable in August 2026, roughly 40 ship basic reporting or dashboards (rungs 1 and 2) and about 8 market a named analytics product. Almost nobody in the mid-market has shipped rung 4 under their own brand. That's the gap prospects notice in a bake-off, and it's why the first mover in each category gets the reference customers.

WHERE REVENUE ATTACHES

Rungs 3 and 4 are where revenue attaches. How to price them is Playbook 3, Pricing and Packaging the Analytics Layer. This guide gets you there.

Why this is a product decision, not a feature request

Three forces are converging on senior living software right now.



Your customers can't build it themselves

33 of the 61 largest US operators have zero data staff. The analytics they want will come from their vendors or it won't come at all. Every export your product generates is a request for a product you haven't shipped yet.



The anchors are already moving

PointClickCare, Yardi, and Aline have built or bought analytics capability. In a consolidating market, the mid-tier's choice is to partner for that capability or concede the differentiation to the platforms.



Your owners are asking

If you're PE-backed, your board has asked for the AI story this year. "We're evaluating" holds for one meeting. A shipped, revenue-attached analytics layer holds much longer. And the cost of waiting is now measurable: every quarter on rung 1 or 2 is a quarter of renewals negotiated without the module the competition is demoing.

Analytics is a rare kind of module. It raises price and retention at the same time. It's demoable in every sales cycle, it gets stickier with every month of accumulated insight, and it moves your product up the stack, from the system that records what happened to the system that decides what happens next.

Systems of record get replaced. Systems of decision get renewed.

Three paths up the ladder, priced

Three ways to get from where you scored to rung 3 or 4, with very different cost, speed, and ownership tradeoffs.

	 Build	 Partner	 White-label
		Integrate a BI product	An analytics and AI layer under your brand
YEAR-ONE COST	\$150K to \$340K: visualization layer about \$60K (1 to 2 months), multi-tenancy and row-level security \$90K to \$160K (4 to 6 months), AI query layer \$80K to \$120K (4 to 6 months)	Vendor fee plus integration engineering. The fee usually scales with communities or end users; get the unit in writing (test H below explains why it's not in this table)	Vendor fee. Same advice on the unit. Integration is lighter because the layer runs on your existing database and permission model rather than beside them
TIME TO FIRST CUSTOMER	6 to 12 months	2 to 4 months, typically	60 to 90 days, which is why the plan in section 07 is 90 days
ONGOING COST	At least 20% of an engineer, permanently, for schema changes, broken charts, tenant edge cases, and performance. Add answer evaluation if there's an AI layer. Three-year total \$470K to \$740K once the enterprise-tenancy rework arrives	The fee, plus the reconciliation work: two metric definitions, two permission models, two places a number can come from	The fee. Accuracy, evals, and model churn are the white-label provider's problem by contract; make sure the contract says so (test B)
WHO OWNS ACCURACY AFTER LAUNCH	You. A named engineer, or nobody	Split. Their engine, your definitions, and the seams show when they disagree	The white-label provider, if you've tested it. This is the whole reason to pick this path and the whole reason to run the seven tests first
WHO OWNS PERMISSIONS	You, on every query path including the AI one	Two systems to keep in sync. The gap between them is where the regional director sees the wrong region	Inherited from your product's model, if the layer runs on your database. Verify with test C
WHAT YOUR CUSTOMERS SEE	Your brand	A second product's look, in a frame. Customers can usually tell	Your brand, end to end, if the license tier says so (test G)
RIGHT ANSWER WHEN	Analytics is your core product, not a module wrapped around one, and you'll fund a real data team with a leader who has shipped production analytics before	You need dashboards fast, definitions are simple, and AI isn't in scope yet	Analytics is a feature your customers demand, your data team is small or nonexistent, and there are bake-offs this quarter where this is the gap. This is where most of the mid-market lands

The build column uses the line-item math from our Buy vs. Build guide. The full build-versus-embed argument, including where in-house builds stall, is in our Buy vs. Build guide. Short version: **the demo is a weekend, and the last 20% is the whole game.**

The seven tests to put to any vendor, including us

The question to ask, what a good answer looks like, what the vague one sounds like, and how to check it live.

Rungs 3 and 4 live or die on three things: consistent, defensible answers; permissions that survive AI writing the query; and answers that stay right after the model underneath changes. Those three requirements break into seven tests. Each one below gives you the question to ask, what a good answer looks like, what the vague answer sounds like, and how to check it live.

You're not looking for a yes. You're looking for a mechanism. If the answers are confident but general, the expensive half of the work hasn't been done.

Two things about this list. First, it works for any vendor, including ones that aren't us, and we'd rather you scored every vendor against it than took our word on any of it. Second, it was built from two sources: what end users in senior living keep asking for after their software already shipped reporting, and a sweep of what six embedded analytics vendors actually publish in their own documentation. Where the list says "few vendors document this," it means few, and we're not exempting ourselves.

A

Whose definition of occupancy wins, and is it applied or merely suggested

 ASK

Where does the definition of occupancy live, who can change it, and is the AI capable of computing its own version instead?

 GOOD ANSWER

One place, owned by you, applied at the moment the question is asked, so the same question returns the same number on every surface and for every user. The vendor can show you, for one answer, which stored definition it used.

 VAGUE ANSWER

"The AI uses your metrics." Every vendor claims consistent, governed definitions. In the documentation sweep, none of the six published a precedence rule for competing definitions or an audit trail showing which definition a given answer used, and two conceded in their own docs that the AI can author its own formula and that predefined metrics merely "may help." In senior living, an occupancy number can feed a financing covenant, a state report, or a family conversation. A plausible answer that disagrees with the last one is worse than no answer.

 TEST IT LIVE

Ask the same metric three different ways in three different phrasings and confirm the number is identical every time. Then ask the vendor to show you which stored definition one of those answers used. A screenshot of a dashboard is not the answer. A lineage or inspector view is.

B

Who tests that the answers are still right after the model changes

 ASK

What happens to our answers the week after you upgrade your model, and who is accountable for finding out before our customers do?



GOOD ANSWER

A named owner, a standing set of known question-and-answer pairs run on a stated cadence and after every model or schema change, a pass rate you can see, a notice period before the model underneath changes, and ideally the ability to run your own question set against the system and get a score.



VAGUE ANSWER

"We always use the latest and best model." That's a statement about the model, not about your answers. In the sweep, one vendor in six shipped a customer-runnable evaluation feature and one held an audited AI-management certification. None of the six published a testing cadence, an accountable owner, or an accuracy number, and two market the inability to pin a model version as a benefit. This is the single most underestimated line item in every build plan and every vendor evaluation, and it's why "AI insights" sits on roadmaps for three straight quarters. Anthropic's own analytics agent went from **21%** to **95%+** answer accuracy on the same model; the model was never the hard part.



TEST IT LIVE

Ask for the evaluation set, the pass rate, who runs it, on what cadence, and the notice period for a model change. Ask whether you can run your own set of ten questions and see a score. Write the answer into the contract.

C

Fail-closed tenant isolation on the AI query path, provable by you

 ASK

What happens when the AI can't work out which tenant or which role is asking? Does the query fail, or does it return something?



GOOD ANSWER

Enforcement sits on every path a model can take to the data, not on the interface in front of it. If a filter can't be applied, the query fails with an error rather than silently returning unfiltered rows. The vendor can show you the same question asked by two roles returning two different result sets, and can show you a preview of what a given user would see before you roll a policy out.



VAGUE ANSWER

"The AI respects our permissions." Dashboards are a solved problem; AI-written queries are not. In the sweep, three of six vendors documented the mechanism on AI paths well, and two documented fail-open edge cases in their own docs, one where omitting a parameter silently drops a role and one where an empty access claim erases all row security. One warned that hiding data with dashboard filters is not a secure practice, which is worth asking every vendor about because it's a common shortcut. Get this wrong in a market that handles resident health information and you haven't shipped an analytics feature.



TEST IT LIVE

Log in as a regional director and as an executive director. Ask the same question. Then ask a question designed to cross the boundary ("show me every community's overtime") from the ED login and watch what comes back. Then ask what happens when the tenant identifier is missing from the request.

D

What leaves your environment when the AI is switched on, feature by feature

 ASK

Which AI features send data outside our environment and which don't, and how long are prompts and conversations kept?



GOOD ANSWER

A per-feature list of what is transmitted when AI is enabled, the named model provider and region, the retention period for prompts and conversation history, and a no-training statement, all in writing.



VAGUE ANSWER

"Your data never leaves your cloud" on the marketing page, with an admin guide that requires you to supply a third-party model key. In the sweep, two of six vendors published a per-feature disclosure. Three carried a live contradiction between the marketing claim and the technical documentation. Retention for prompts was unpublished or unfindable at four of six. The no-training statement is the easy one; five of six publish it. The retention period is the one nobody volunteers, and in a HIPAA-adjacent market it's the one your customers' compliance teams will ask you for.



TEST IT LIVE

Ask for the four items in writing before the second call. A vendor who has them sends them the same day.

E

Whether a non-engineer can add a view, and whether your customer can

 ASK

Can our customer, not our analyst, answer a new question without a ticket, and which plan tier does that require?



GOOD ANSWER

Someone adds a new view live, as your end customer would, without touching code, and the vendor states plainly what license that user needs. Bonus: when the AI answers, your customer can export that answer as data, not just as an image.



VAGUE ANSWER

"Fully self-service." All six vendors claim it. For the vendor's own staff most deliver it. For your end customers, which is the case that matters when you're embedding, the picture splits hard: one caps embedded end users at view-only permanently, one gates end-customer authoring to a plan tier roughly two and a half times its entry price, and one documents that the mode exists but never says what a user can actually create. This is the loudest signal in the senior living complaint research: end users describe report-building as confusing, tedious, or not flexible enough, and keep asking for easier custom reports after the vendor already shipped an upgrade.



TEST IT LIVE

Ask the vendor to add a new view as your end customer would, live, start to finish. Ask what license that user needs. Then ask the AI a question and try to export the answer.

F

Cross-source joins, and what they quietly require



ASK

What does it take to answer one question across two of our systems? Does it need a warehouse, a license upgrade, or a copy of our data on your infrastructure?



GOOD ANSWER

A plain statement of the mechanism and its cost. Joins across sources without a required warehouse, without a tier gate, and with a stated performance profile, or an honest account of which of those it does require.

VAGUE
ANSWER

"Yes, we support multi-source." In the sweep, four of six could join across sources and each charged a different architectural bill: one requires the top license tier, one stops pushing joined queries down to the database and sells its own acceleration store as the remedy (which moves where your data lives), one requires materializing both sides into vendor-managed storage, and one doesn't document cross-connection joins anywhere, as a capability or a limitation. This matters in senior living specifically because so many platforms were assembled by acquisition, with engagement, communications, clinical, and financial data in separate systems. It's the question that gets discovered late, after the contract.



TEST IT LIVE

Bring two real exports from two real systems to the demo and ask for one number that needs both. Then ask where the data went to produce it.

G

The license line that removes the vendor's name



ASK

Which license tier removes your name from our footer, and is the answer the same in the embed mode we'll actually ship?



GOOD ANSWER

The clause, in the contract, not on the pricing page. And a straight answer on whether it applies to the embed mode you're buying.

VAGUE
ANSWER

"True white-label" in the marketing copy. All six vendors claim it. In the sweep, one documents it as included on every plan. Two gate it behind a premium tier. Two assert it only in marketing with no supporting documentation, in one case with the vendor's own element names still visible in the buyer's page. One vendor's white-label documentation applies only to a legacy mode that new customers aren't provisioned into. Branding removal is license-dependent at most vendors in this category, and that includes Querri, where full attribution removal sits at a premium tier or a custom agreement. Ask for the tier and the clause.



TEST IT LIVE

Ask for the contract language. Then ask to see the footer in the embed mode you'll ship.

H

The one we left off the list, and why

NOT A TEST

The obvious eighth test is how the vendor counts your end customers, since rolling analytics out to every community is priced very differently per seat, per community, per question, or flat. It's not on the list because most vendors publish the unit and only two publish a real number, so a public checklist item can't discriminate. It belongs in your negotiation instead, with two things worth knowing: where a vendor requires you to bring your own model key, the model bill is an uncapped variable line the pricing page doesn't mention, and at least one vendor publishes AI overage per credit without publishing what a credit buys. **Get the unit, the cap, and what's excluded, in writing.**

SCORING THE SEVEN

- ☒ Mechanism shown
- ☐ Claimed only
- ☐ No

Score each vendor A through G as mechanism shown, claimed only, or no. **Two or more "claimed only" on A, B, and C means you're being sold the demo, not the product.** Take the same scorecard to your own engineering lead if you're considering the build path; the questions are identical, and so is the standard.

06

SECTION 06

The precedent


The white-label motion isn't exotic in this market. At least one senior living platform has shipped its product under a partner's brand for years, in this exact vertical, successfully. Embedding an analytics layer is the same motion pointed at a new capability.

The precedent question is settled. The only open question is who moves first in your category, because the first mover gets the reference customers and every later entrant gets compared to them.

The 90-day plan

Three phases from first inventory to a launch-ready package, each with an exit criterion.

Three phases from first inventory to a launch-ready package. Each has deliverables, an owner, and an exit criterion. If a phase doesn't clear its exit criterion, don't start the next one; that's the whole point of running it as a pilot instead of a roadmap item.



Days 1 to 30 Inventory and definitions

Owner: Product, with one engineer who knows the schema.

DELIVERABLES

- 1 Your completed inventory table from section 01, at least ten rows, with the rightmost column filled in honestly.
- 2 A written definition for every metric in the inventory that customers ask about by name: occupancy, census, move-in, labor cost per resident day, whatever your category's five are. One sentence each, one owner each. Where two teams define it differently today, pick one and write down why.
- 3 A data-readiness read: which sources are clean enough to answer the inventory questions, which joins break, and what's missing. This is exactly what a data readiness assessment produces, and it's worth getting before any vendor conversation, including one with us.
- 4 Your ladder score from section 02, with the missing conditions listed.

EXIT CRITERION

You can name the two or three questions you'll pilot, the definitions behind them are written and owned, and you know which other system each one needs. If you can't, the pilot will spend its 30 days on definitions instead of customers.

Days 31 to 60 Design-partner pilot

Owner: Product, with customer success in the room.

Pick the partners. Two or three customers who keep filing the export tickets. Not your biggest logos; your most vocal reporting complainers, ideally in two different sizes so you learn where the permissions model bends.

DELIVERABLES

- 1 Rung 3 or rung 4 live for those partners, under NDA, on the two or three questions from phase one. Real data, real users, not a sandbox.
- 2 The design-partner agreement, one page, covering:
 - what data access you're granting and to whom
 - NDA terms and the demo rights you keep
 - evaluation rights (you'll be logging their questions and using them, see below)
 - pricing at graduation, so "free during the pilot" has a written end date
 - reference rights, so the partner agrees up front to be a reference call if the pilot works
- 8 **The eval suite, seeded from the pilot.** Every question a partner asks in the pilot gets logged with the answer they got and whether they trusted it. By day 60 you'll have 30 to 50 real questions with known-good answers. That's your regression set. Rerun it after every model change and every schema change, forever. Here's what ten rows look like:

QUESTION AS THE USER ASKED IT	EXPECTED ANSWER	SOURCE OF TRUTH	LAST RUN	PASS
What was occupancy at Maple Grove last month?	91.4%	Census view, occupancy definition v2		
Which communities ran agency hours over 15% in August?	Three named communities	Labor view joined to census		
Show me falls per 1,000 resident days by community, Q2	Table, 12 rows	Incidents view, resident-day definition		

Whether you build, partner, or white-label, this table is yours. It's the thing that turns "the AI is accurate" from a claim into a number you can put in a renewal deck.

EXIT CRITERION

Partners are asking questions in the product without being prompted, the eval suite passes on rerun, and at least one partner has said some version of "can we have this for the whole portfolio." If they're still exporting, you haven't cleared rung 3 and shouldn't price it yet.

Days 61 to 90 Package and price

Owner: CEO or CPO, with sales.

DELIVERABLES

- 1 The tier decision: premium tier, per-community add-on, or enterprise package, and where rungs 3 and 4 sit on the price card. Playbook 3 works through this in full.
- 2 The pricing unit. Per community is the natural unit in this market; multi-site operators already think in per-community economics.
- 3 The launch story for your next renewal cycle: a named module, its own page, its own demo track in every bake-off.
- 4 Grandfathering terms for the design partners, written and loud, so they become the reference calls that sell everyone else.
- 5 The metrics you'll track from day one: attach rate on new deals, upgrade rate at renewal, question volume per account per week, and the retention delta between accounts with and without the module. Set baselines from your first two quarters.

EXIT CRITERION

A named module with a price, on the price card, timed to a renewal wave, with two reference customers and a passing eval suite behind it.

THE FOUR METRICS, FROM DAY ONE



Attach rate

on new deals



Upgrade rate

at renewal



Question volume

per account per week



Retention delta

accounts with and without the module

Before it's a build decision, it's a data question

Every path in section 04 and every phase in section 07 assumes the data you're sitting on can actually power an analytics product. That's worth checking before you spend a quarter or sign a contract, and it's the first deliverable of the 90-day plan.

Start with one export, or challenge us with as many sources and connectors as you'd like. More sources means the report can show how your tables connect across systems and where the joins break. In **3 to 5** business days Querri sends back a branded **4 to 5** page report on what's clean, what's broken, and the questions your customers could already be asking your product. It's free and it's yours to keep either way. Your file comes back or gets deleted on request.

✓ SOC 2 Type II

✓ ISO 27001

✓ HIPAA-ready

Get your free assessment

See an embedded analytics layer running in a senior living product. Book a demo.

querri.com/senior-living/software

The Embedded Analytics Playbook for Senior Living SaaS