AUGUST 21, 2026 / DEMENTIA CARE / FAMILY CAREGIVING APPS

Daily story-capture companion for dementia families

Care software tracks meds and appointments and stops there; who the person is - their stories, their voice - stays outside the care stack, buried in scattered voice memos and photos that lose context with every phone upgrade.

Daily story-capture companion for dementia families should be tested as a narrow first-win workflow for Adult child caring for a parent living with dementia.

LOW DIFFICULTY

FAMILY SUBSCRIPTION WITH A PREMIUM TIER FOR PRINTED MEMORY BOOKS AND SHARED CAREGIVER ACCESS.

62/100

VALIDATION VERDICT / RESEARCH

Validation is a weighted rubric, not a guarantee. Use the next validation step before building.

Confidence	55%
Lifecycle	Crowding
Timing	37/100
Rubric	INAV-VALIDATION-2026-06-04

CROWDING Window closing

Demand signal	**5.5/10**
Problem severity	**6.3/10**
Willingness to pay	**6/10**
Competitive saturation	**6/10**
Feasibility	**7.8/10**

VERDICT

Research · 62/100

Daily story-capture companion for dementia families should be tested as a narrow first-win workflow for Adult child caring for a parent living with dementia.

THIS WEEK'S TEST

Run a 60-day pilot with twenty families and measure weekly active story capture plus willingness to pay at trial end.

KILL IT IF

Fewer than five qualified buyers agree to discuss the workflow after targeted outreach.

READER DEMAND SIGNAL

Would you build this? Would you pay?

One tap each — anonymous, one vote per question per day. Tallies update as readers weigh in.

Would you build this?

Be the first to weigh in

Would you pay for this?

Be the first to weigh in

VISUAL OPPORTUNITY DASHBOARD

Read the idea like a product signal board.

These visuals are generated from the report's existing scores. They make the decision path scannable without pretending to be live market data.

DEMENTIA CARE / FAMILY CAREGIVING APPS

SIGNAL MODEL

Daily story-capture companion for dementia families

Daily story-capture companion for dementia families should be tested as a narrow first-win workflow for Adult child caring for a parent living with dementia.

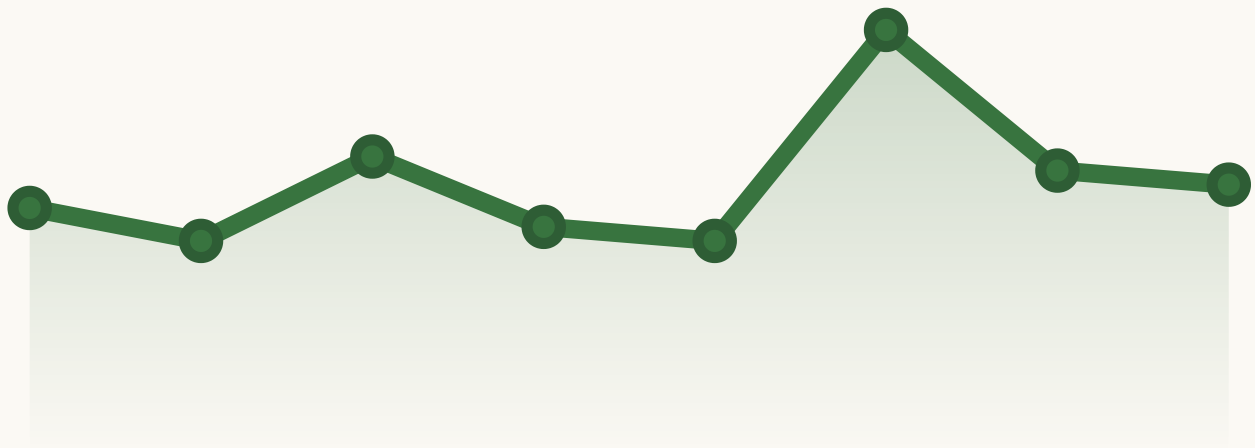

VALIDATION
62/100
Research

CONFIDENCE
55%
Editorial confidence

SCORE AVG
7.3/10
Scorecard average

PROOF
5.5/10

Proof signal average

SCORE RADAR

Decision balance

VALUE EQUATION

Offer strength

DREAM OUTCOME	8/10
PERCEIVED LIKELIHOOD	6/10
TIME DELAY	8/10
EFFORT AND SACRIFICE	7/10

MARKET MAP

Demand-led wedge

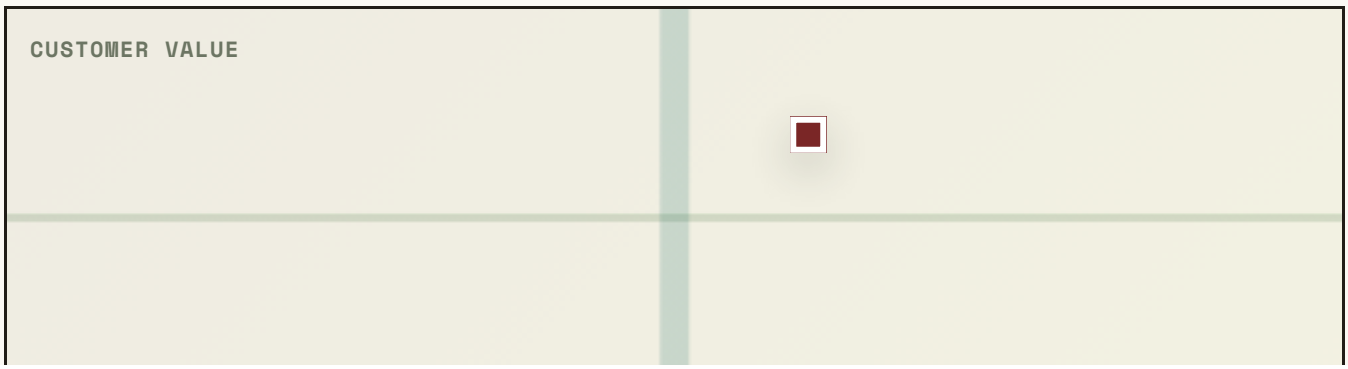

UNIQUENESS

High value plus high uniqueness deserves deeper research; lower uniqueness requires a clear distribution advantage.

VALIDATION FUNNEL

From pain to product.

1	**Buyer pain** Adult child caring for a parent living with dementia	5.3/10
2	**Concierge proof** Run a 60-day pilot with twenty families and measure weekly active story capture p...	5.5/10
3	**Paid wedge** Concierge review or paid template	8/10
4	**Repeatable product** Family subscription with a premium tier for printed memory books and shared careg...	6.4/10

EVIDENCE HEATMAP

Signal intensity.

WHY NOW
5/10
Demand visibility

WHY NOW
8/10
Tooling readiness

WHY NOW
4/10
Budget clarity

WHY NOW
6/10
Competitive window

PAIN
5/10
Repeated workflow friction

MONEY
4/10
Budget hypothesis

URGENCY

6/10

Switching pressure

DISTRIBUTION

7/10

Reachable buyer language

LIFECYCLE TIMING

Crowding (37/100): demand exists, but funded or visible competitors are compressing the window.

Deterministic stage assignment from re-check status, demand signals, complaint echo, and competitive saturation.

37/100

CROWDING

Adoption substrate is up 536.6% across matched packages.

2 matched company signals raise saturation.

Demand

70/100

Not old enough for a 30-day re-check yet.

Saturation

60/100

2 funded signals across 2 matched competitor signals.

Complaint echo

22/100

Matched adoption substrate is up 536.6%.

VALIDATION ENGINE

Evidence-backed idea-validation score.

The score uses a versioned 2026 rubric across demand, problem severity, willingness to pay, competitive saturation, and feasibility.

62/100

Research

Research is the current validation verdict: feasibility is the strongest signal, while demand signal is the main evidence gap to close before scaling the build.

Rubric version: INAV-VALIDATION-2026-06-04 / generated August 21, 2026

Demand signal

5.5/10

24% WEIGHT

Demand looks thin because the report has 2 source-backed signal(s), an editorial confidence of 55/100, and a defined buyer in Dementia care / family caregiving apps.

- Reminiscence therapy shows measurable benefits for mood and cognition in people living with dementia.
- Target buyer: Adult child caring for a parent living with dementia

Problem severity

6.3/10

22% WEIGHT

Problem severity is thin when the buyer pain, customer value, and dream-outcome scores are combined.

- Care software tracks meds and appointments and stops there; who the person is - their stories, their voice - stays outside the care stack, buried in scattered voice memos and photos that lose context with every phone upgrade.
- Reminiscence therapy shows measurable benefits for mood and cognition in people living with dementia.

Willingness to pay

6/10

20% WEIGHT

Willingness to pay is weak; the model has a monetization hypothesis, but it must still be proven through paid pilots or explicit pricing objections.

- Family subscription with a premium tier for printed memory books and shared caregiver access.
- Run a 60-day pilot with twenty families and measure weekly active story capture plus willingness to pay at trial end.

Competitive saturation

6/10

18% WEIGHT

No source-backed direct match is recorded yet, so saturation risk is treated as unknown rather than proof of novelty.

- Existing-product check has no named direct match.
- Competitive score rewards a narrow wedge, not absence of research.

Feasibility

7.8/10

16% WEIGHT

Feasibility is strong for a low build if the MVP is limited to the first measurable workflow.

- Run a 60-day pilot with twenty families and measure weekly active story capture plus willingness to pay at trial end.
- Consumer willingness to pay in caregiving apps is historically weak despite strong emotional resonance.

Next validation step

Run a 60-day pilot with twenty families and measure weekly active story capture plus willingness to pay at trial end.

VALIDATION SPRINT

Seven days to a build / kill decision.

Derived from this report's own validation test, channels, offers, and kill criteria.
Each day has a threshold, so the week ends in a decision instead of a feeling.

DAY 1

Build the buyer list

List 50-100 named adult child caring for a parent living with dementia prospects from Community pain posts and Direct outreach — names, not categories.

Threshold: 50+ named, reachable buyers on the list.

DAY 2

Join the watering holes

Join and observe Reddit / forums, Launch communities, Review and alternative pages. Collect the exact words buyers use for this pain.

Threshold: 10+ verbatim pain quotes captured.

DAY 3

Send first outreach

Send the cold outreach template (below) to 15 buyers from the day-1 list, personalized with one detail each.

Threshold: 15 sent; 3+ replies of any kind.

DAY 4

Run buyer interviews

Hold 15-minute calls using the interview script (below). Listen for current workarounds and what they cost.

Threshold: 3+ completed interviews.

DAY 5

Run the report's validation test

Run a 60-day pilot with twenty families and measure weekly active story capture plus willingness to pay at trial end.

Threshold: Problem resonance: 5+ calls or 10+ detailed replies.

DAY 6

Make the smoke offer

Offer "Concierge review or paid template" at $19-$99 to every interviewed buyer. Manual delivery is fine — payment is the signal.

Threshold: 1+ pre-commitment (payment, signed LOI, or scheduled paid pilot).

DAY 7

Decide against the kill criteria

Score the week against this report's kill criteria, then take the stated next validation step: Run a 60-day pilot with twenty families and measure weekly active story capture plus willingness to pay at trial end.

Threshold: A written build / keep-testing / kill decision.

Pass signal

Pass: thresholds on days 3, 4, and 6 are met — proceed to the next validation step with real buyer language in hand.

Fail signal

Kill or rethink if the week confirms: Fewer than five qualified buyers agree to discuss the workflow after targeted outreach.

RESEARCH WORKFLOW

Decision scorecard.

The report is structured to force a yes, no, or test decision instead of leaving the reader with a loose brainstorm.

Opportunity

6/10

PROMISING

Daily story-capture companion for dementia families has an editorial confidence score of 55/100 before live buyer validation.

Problem

5/10

PROMISING

Care software tracks meds and appointments and stops there; who the person is - their stories, their voice - stays outside the care stack, buried in scattered voice memos and photos that lose context with every phone upgrade.

Feasibility

8/10

STRONG

A low build can work if the MVP stays limited to the first repeated workflow.

Why now

10/10

EXCEPTIONAL

Dementia prevalence is climbing with the aging population, speech-to-text and photo organization are finally effortless, and reminiscence therapy has mainstream clinical acceptance as a non-drug intervention.

MARKET AND MONEY

Business fit and offer ladder.

Revenue potential

$250K-$2M ARR potential if the wedge proves budget urgency and becomes a recurring workflow.

Execution difficulty

Execution is low; the main constraint is staying narrow enough for a first proof loop.

Go-to-market

Start with manual concierge output, direct outreach, and community proof before paid acquisition.

Founder fit

Best for an AI-assisted solo founder who can interview the buyer and ship a focused first version quickly.

1. Lead magnet

Daily Story-capture Companion For Dementia Families checklist

Free

Helps Adult child caring for a parent living with dementia audit the painful workflow before buying software.

Capture qualified leads and learn the buyer's exact language.

2. Frontend offer

Concierge review or paid template

$19-$99

Delivers the first useful output manually before automation is trusted.

Validate urgency, workflow fit, and willingness to pay.

3. Core offer

Daily story-capture companion for dementia families focused SaaS

$49-$499/month

Turns the recurring manual workflow into a repeatable product loop.

Create the recurring revenue product after the narrow wedge survives tests.

4. Continuity

Monitoring, benchmarks, and monthly reporting

$99-$1,000/year add-on

Keeps the buyer engaged with ongoing proof, saved time, or reduced risk.

Increase retention and make the product part of a routine.

5. Backend offer

Done-with-you setup, agency, or team rollout

Custom

Adds implementation help, integrations, and workflow migration.

Capture higher-value accounts once the productized wedge is proven.

ECONOMICS

Price-anchored revenue scenarios.

Derived from this report's "Core offer" offer-ladder stage ($49-$499/month). These are price-anchored scenarios, not market-size claims.

Proof

$490-$4,990 MRR

10 CUSTOMERS

Ten paying customers proves willingness to pay and funds continued validation.

Wedge

$2,450-$24,950 MRR

50 CUSTOMERS

Fifty customers in one niche makes the workflow the default in that circle and feeds referrals.

Vertical leader

$12,250-$124,750 MRR

250 CUSTOMERS

A few hundred accounts in one vertical is a real business before any horizontal expansion.

Break-even

At $49-$499/month, 1 customers cover the stated Local-first MVP budget: $0-$10K before paid acquisition. budget within a month; fewer if they land at the top of the range.

Sizing the buyer universe

Size the buyer universe in one day: count adult child caring for a parent living with dementia reachable through the report's channels (directories, associations, communities) until the list stops growing — the test only needs the first 100 names, not a TAM estimate.

Pricing benchmark

No public look-alike products were recorded in this report, so price against the manual workaround's time cost, not against software.

REFLEXIVITY

Mixed reflexivity — execution over secrecy.

Does publishing this opportunity strengthen it or crowd it? low confidence.

Publish or protect

No dominant reflexivity signal. Publish the analysis, but note that execution speed and distribution matter more than secrecy here; revisit if the space shows saturation.

Signals

No dominant network or scarcity marker — a mixed case where execution speed and distribution matter more than secrecy.

EVIDENCE

Why now and proof signals.

Why now

5/10

Demand visibility

Reminiscence therapy shows measurable benefits for mood and cognition in people living with dementia.

Build only if the complaint repeats across interviews, posts, or existing workflow artifacts. alz.org

8/10

Tooling readiness

AI-assisted product work and managed infrastructure reduce the first-version cost.

The first release should automate one high-friction step rather than become a broad platform. en.wikipedia.org

4/10

Budget clarity

Family subscription with a premium tier for printed memory books and shared caregiver access.

Ask for money during validation before building the full workflow. alz.org

6/10

Competitive window

The wedge is specific enough to test without claiming the whole market.

Position around one buyer and one measurable first-win outcome. alz.org

Proof signals

5/10

Pain: Repeated workflow friction

Reminiscence therapy shows measurable benefits for mood and cognition in people living with dementia. alz.org

4/10

Money: Budget hypothesis

Adult child caring for a parent living with dementia is the first group to test because the monetization path is: Family subscription with a premium tier for printed memory books and shared caregiver access. alz.org

6/10

Urgency: Switching pressure

Urgency becomes real only if the current workaround costs time, risk, money, or reputation every week. en.wikipedia.org

7/10

Distribution: Reachable buyer language

The first channel should be whichever source lane already contains the buyer's vocabulary. alz.org

POSITIONING

Market gaps and execution plan.

Underserved segments

- Adult child caring for a parent living with dementia who still run the workflow in spreadsheets, generic docs, email, or chat threads.
- Small teams in Dementia care / family caregiving apps that feel the pain weekly but are too narrow for broad incumbents.
- New adopters who need guided proof before committing to a larger platform.

Feature gaps

- A narrow workflow that reaches value without configuration-heavy onboarding.
- A buyer-facing proof artifact that shows time saved, risk reduced, or communication improved.
- A handoff path from manual concierge service to repeatable software.

Differentiation levers

- Use specificity as the wedge: one buyer, one workflow, one measurable result.
- Show proof earlier than broad competitors with before-and-after examples and small pilot data.
- Keep implementation lighter than incumbent suites or generic AI assistants.

Execution snapshot

Type	**Consumer app product**
Timeline	**2-4 weeks**
Budget	**Local-first MVP budget: $0-$10K before paid acquisition.**
Initial offer	**Concierge review or paid template**

Build only the first-win workflow for "Daily story-capture companion for dementia families" and keep research, setup, and exceptions manual until the wedge is proven.

Weekly

Community pain posts

Use communities and forums where Adult child caring for a parent living with dementia already describe the painful workflow.

Problem teardown, interview ask, and short demo clip / 5 qualified calls or 10 detailed replies in 7 days

Daily during validation

Direct outreach

Direct conversations are the fastest way to verify budget ownership and switching cost.

Concierge pilot offer with a manually prepared sample / 3 paid pilots, LOIs, or budget-owner follow-ups

Bi-weekly

Searchable comparison content

Alternative and comparison pages reveal objections, pricing language, and buying intent.

Before-and-after page or alternatives memo for the exact workflow / Organic clicks, booked demos, or waitlist joins from comparison intent

Once MVP is clickable

Launch directory

Launches test whether the promise is legible to people outside the first interview set.

Single-purpose demo and first-win story / 25% demo completion or 10 waitlist joins

COMPETITIVE LANDSCAPE

Alternatives, incumbents, and whitespace.

This section names likely workarounds and public players so the report can argue where the wedge is still open.

Daily story-capture companion for dementia families should be positioned against generic AI assistants, no-code workarounds, and any vertical incumbent that already owns Dementia care / family caregiving apps. The opening is a narrower first-win workflow for Adult child caring for a parent living with dementia.

WORKAROUND

Airtable

No-code database

Competes when the first version can be modeled as a lightweight database and workflow view.

WORKAROUND

Notion

Workspace and documentation

Competes when buyers can solve the pain with templates, checklists, and shared pages.

WORKAROUND

Asana

Project management

Competes where the buyer can express the workflow as tasks, owners, and due dates.

ADJACENT

HubSpot

CRM and marketing platform

Competes for sales, marketing, client follow-up, webinar, and service pipeline workflows.

DIRECT

ServiceTitan

Field service platform

Relevant to field service, HVAC, appliance repair, contractor, and service dispatch ideas.

Whitespace

- A narrow workflow that reaches value without configuration-heavy onboarding.
- A buyer-facing proof artifact that shows time saved, risk reduced, or communication improved.
- A handoff path from manual concierge service to repeatable software.
- Use specificity as the wedge: one buyer, one workflow, one measurable result.
- Show proof earlier than broad competitors with before-and-after examples and small pilot data.
- Keep implementation lighter than incumbent suites or generic AI assistants.
- Own the specific buyer workflow instead of selling a broad AI assistant.

Positioning moves

- Lead with the exact buyer: Adult child caring for a parent living with dementia.
- Show a proof artifact for: Run a 60-day pilot with twenty families and measure weekly active story capture plus willingness to pay at trial end.
- Name the generic-assistant workaround directly and explain what it misses.
- Offer concierge setup before promising a full platform.

Public source

Airtable

https://www.airtable.com/

Public source

Notion

https://www.notion.com/

Public source

Asana

https://asana.com/

Public source

HubSpot

https://www.hubspot.com/

Public source

ServiceTitan

https://www.servicetitan.com/

Public source

Report source

https://www.alz.org/

Public source

Report source

https://en.wikipedia.org/wiki/Reminiscence_therapy

Who's already moving in Software & AI

Public companies and funding signals the intelligence graph links to this vertical (related by keyword overlap — sized players, not direct competitors). Source: /graph.json .

FIELD SERVICE MANAGEMENT

$625M

ServiceTitan

Operations software for contractors and field-service trades: scheduling, dispatch, quotes, jobs, and crew management.

IPO · 2024-12-12

AUDIENCE COMPANION

Segments, channels, and intent language.

The companion is also published as a standalone HTML page and Markdown file for research handoff.

Primary audience

Adult child caring for a parent living with dementia is the first audience because the report already names a repeated pain, reachable channels, and a validation test that can be run before software is complete.

DAILY WORKFLOW | STORY VALIDATION | DAILY AI | STORY AUTOMATION

ELDERCARE | CAREGIVING | DEMENTIA CARE / FAMILY CAREGIVING APPS

First validation channels

- **Reddit / forums:** Post a problem teardown for Dementia care / family caregiving apps and ask how people solve it today.
- **Launch communities:** Ship a narrow demo and watch which promise gets clicks.
- **Review and alternative pages:** Write an alternatives page that owns one narrow use case.
- **Community pain posts:** Problem teardown, interview ask, and short demo clip

TIME TO EXECUTE

Execution-readiness scorecard.

The score turns the report into bottlenecks, accelerators, and a dated first-month launch plan.

83/100

Ready to test

Daily story-capture companion for dementia families scores 83/100 for execution readiness. The recommended next step is Run a 60-day pilot with twenty families and measure weekly active story capture plus willingness to pay at trial end.

Execution scorecard is generated from report validation, confidence, feasibility, founder fit, and difficulty.

Bottlenecks

- Consumer willingness to pay in caregiving apps is historically weak despite strong emotional resonance.
- Sustained daily engagement from the person with dementia depends on caregiver facilitation.
- A broad AI assistant can flatten differentiation unless the wedge is painfully specific.
- The first release can become a generic dashboard if the job is not named tightly.
- Needs real buyer access, not only desk research.

First milestones

- 2026-08-21: Frame the wedge
- 2026-08-24: Interview 10 people who match the buyer persona.
- 2026-08-28: Ship a clickable demo or concierge workflow that produces the first useful artifact.
- 2026-09-04: Run one paid pilot or collect explicit pricing objections before automating the rest.

FRAMEWORKS

Value equation, matrix, and ACP.

FOUNDER LENS

Fit, roast, and kill criteria.

10/10

Founder fit

A solo or AI-assisted founder with direct access to Adult child caring for a parent living with dementia.

ADVANTAGES

- Can talk to the buyer before writing much code.
- Can ship a narrow first-win demo quickly.
- Can use local-first research artifacts to keep validation moving without a large team.

GAPS

- Needs real buyer access, not only desk research.
- Needs proof of budget or repeated urgency.
- Needs a crisp wedge before broad product work starts.

Roast

Promising enough to test, not strong enough to build broadly.

BLIND SPOTS

- Consumer willingness to pay in caregiving apps is historically weak despite strong emotional resonance.
- A broad AI assistant can flatten differentiation unless the wedge is painfully specific.
- The first release can become a generic dashboard if the job is not named tightly.

HARD QUESTIONS

- Who wakes up already trying to solve this?
- What do they stop paying for or stop doing when this works?
- What proof would make a skeptical buyer trust it in one screen?
- What is the smallest paid version of this idea?

Kill criteria

- Fewer than five qualified buyers agree to discuss the workflow after targeted outreach.
- No buyer can name a current cost in time, money, risk, or reputation.
- The first demo does not produce a clear next step, paid pilot, or specific objection.

Next actions

- Write the one-sentence promise and test it in the strongest channel.
- Create the lead magnet and use it to recruit interviews.
- Build the smallest demo that proves the first win.

ACTION SUITE

Move from reading to testing.

Local-first handoff cards copy prompts or structured data without requiring an account.

BUILD THIS IDEA

Copy the focused build brief for a coding agent.

COPY

ROAST

Copy the critique lens and blind spots before committing time.

COPY

LANDING PAGE

Copy a landing-page brief based on buyer, pain, and validation.

COPY

BRAND PACKAGE

Copy positioning inputs for naming, messaging, and design direction.

COPY

AD CREATIVES

Copy campaign angles for buyer-problem validation.

COPY

EXPORT DATA

Copy structured JSON for a research engine, roadmap tracker, or another agent.

COPY

FOUNDER FIT

Copy the founder-fit self-check before entering build mode.

COPY

FIRST CONTACT KIT

Outreach template and interview script.

Built from this report's buyer, pain language, and channels. Personalize one detail per message — these are starting points, not spam ammunition.

Cold outreach message

QUESTION ABOUT DAILY WORKFLOW

HOW ARE YOU HANDLING CARE SOFTWARE TRACKS MEDS AND APPOINTMENTS AND STOPS THERE;...

15 MINUTES ON A DEMENTIA CARE / FAMILY CAREGIVING APPS WORKFLOW?

```
Hi {{firstName}},

I'm researching how adult child caring for a parent living with dementia handle
this today: Care software tracks meds and appointments and stops there; who the
person is - their stories, their voice - stays outside the care stack,...

I'm not selling anything yet — I'm testing whether "Daily story-capture companion
for dementia families" is worth building, and I'd rather learn from people living
the workflow than guess.

Would you trade 15 minutes for first access (and a say in what gets built) if it
goes ahead?

{{yourName}}
```

COPY MESSAGE

Buyer interview script

1. Walk me through the last time this happened: Care software tracks meds and appointments and stops there; who the person is - their stories, their voice - stays outs... What did you actually do?
2. What does that workaround cost you — in hours, money, or risk — in a normal month?
3. What have you already tried or bought to fix it, and why didn't it stick?
4. If "A daily companion: morning check-in with date, weather, and relatives' photos with names and connec..." existed, what would have to be true for you to switch in the first week?
5. Who else feels this worse than you do — and would you introduce me?

WHERE TO SEND IT

- Community pain posts — Problem teardown, interview ask, and short demo clip
- Direct outreach — Concierge pilot offer with a manually prepared sample
- Searchable comparison content — Before-and-after page or alternatives memo for the exact workflow
- Reddit / forums — Post a problem teardown for Dementia care / family caregiving apps and ask how people solve it today.
- Launch communities — Ship a narrow demo and watch which promise gets clicks.

HANDOFF

Build and review prompts.

Build prompt

Build a narrow MVP for "Daily story-capture companion for dementia families" for Adult child caring for a parent living with dementia. Preserve the evidence, build only the first-win workflow, include source links, and treat Run a 60-day pilot with twenty families and measure weekly active story capture plus willingness to pay at trial end. as the first acceptance gate.

Review prompt

Review the "Daily story-capture companion for dementia families" MVP for over-breadth, unsupported claims, weak buyer proof, privacy risk, and missing validation instrumentation. Do not approve expansion until the kill criteria and success metrics are measurable.

organization / alz.org

Alzheimer's Association

Prevalence data and caregiver-burden research documenting the population this companion serves.

encyclopedia / en.wikipedia.org

Reminiscence therapy - Wikipedia

The evidence-backed intervention the app operationalizes through daily prompts and recordings.

PIVOT MAP

If this exact wedge isn’t yours, these are adjacent.

Derived deterministically from this report's buyers, vertical language, and business model.

Same problem, different buyer: Budget owner who feels the operational cost of the broken workflow.

The workflow pain in this report is not exclusive to adult child caring for a parent living with dementia. Budget owner who feels the operational cost of the broken workflow. faces the same friction with their own budget and urgency.

First test: Re-run day 3 of the sprint (15 outreach messages) against this buyer only, and compare reply rates before changing anything else.

Same workflow, adjacent vertical: pick the nearest regulated niche

No second vertical matched this report’s language strongly, which usually means the wedge is horizontal. Horizontal wedges win by going vertical first.

First test: Pick the vertical where the pain costs the most per incident and rewrite the promise in its vocabulary.

Same wedge, alternate model: a productized service (fixed-price, done-for-you delivery)

This report monetizes via ”Family subscription with a premium tier for printed memory books and shared caregiver access.”. Concierge delivery validates willingness to pay before any software exists and earns the workflow knowledge the product needs.

First test: Offer both versions on day 6 of the sprint and let the first pre-commitment choose the model.

CONNECTIONS

Where this report sits in the intelligence graph.

Links from the ontology layer. Declared links are explicit in the research record; inferred links are keyword overlap and labeled as such. Full graph at /graph.json.

EVIDENCE INDEPENDENCE 87/100

5 source domains, 20 evidence edges. Dominant family: local:data/regulatory-clock.json. Audit all provenance .

Related reports (shared keywords)

- Daily AI check-in calls for seniors living alone — daily ai, daily workflow
- Daily gum-line photo scoring for cleaning-gap prevention — daily ai, daily workflow
- Daily postpartum check-ins for the first two weeks home — daily ai, daily workflow

FULL NARRATIVE