

The AI employee a person **stands behind.**

Every company wants AI agents. Almost none can operationalise them — because nobody owns the output. Donna sells the missing layer: agent-speed work with a named human's signature on every outcome, built by a 15-year executive assistant and operated by a Canadian carrier. The wedge is an EA seat. The endgame is the trust layer for the agentic workforce.

PREPARED FOR FIBERNETICS / NEWT LEADERSHIP · SEPTEMBER 2026 · CONFIDENTIAL · MARKET & FINANCIAL FIGURES ARE ILLUSTRATIVE PLANNING ASSUMPTIONS

01 · EXECUTIVE SUMMARY

Supervision is the product.

Gartner forecasts that 40% of agentic-AI projects will be scrapped by 2027 — not because models fail, but because organisations can't govern them. Meanwhile a senior executive assistant costs \$120–162K CAD loaded and there aren't enough of them. Donna closes both gaps at once: an AI agent that drafts at machine speed, inside an operating model where a trained person reads everything before it ships under a customer's name.

We are not building a better model. We are productising the thing Fibernetics has actually done for fifteen years: run business-critical operations that other companies bet their day on, with a person accountable when it matters. Two engines drive the business:

- **The Desk (revenue now):** supervised executive seats at **\$1,500 CAD/month** — Donna drafts, a NEWT-staffed reviewer signs, the customer approves only where they set the line. Live in beta today with a published price and a 5-seat capacity cap.
- **The Platform (scale):** the reviewer dashboard itself, sold as software to the world's ~2M executive assistants — each EA becomes the accountable human for their own fleet of agents and carries 3× the executives. Free for EAs during beta; converts to a SaaS seat.

The billion-dollar shape (illustrative): ~\$100M ARR at a software multiple — roughly 4,000 supervised seats plus 25,000 EA platform seats — reached by expanding from the EA agent into a roster (Finance, Support, Voice) on the same supervised-desk substrate, distributed through Fibernetics' existing carrier and VAR channels.

2	\$1,500	15 yrs	\$100M
PRODUCTS LIVE IN BETA	CAD / SEAT / MO · PUBLISHED	EA CRAFT ENCODED	ARR TARGET · ILLUSTRATIVE

02 · COMPANY & PRODUCT OVERVIEW

NEWT's first AI employee.

Mission: make agentic AI usable by real businesses by attaching a named, accountable human to every outcome.

What it is: Donna is an AI executive assistant personified on Dawn Gamble — a real senior EA of fifteen years at Fibernetics — and operated under her review. She runs calendars, inboxes, travel, meeting prep and follow-ups. Nothing she produces reaches another human until a person has read it. Donna ships under **NEWT**, the business division of Fibernetics, a Canadian carrier running business voice and network since 2011.

Stage: live early beta. Two marketing sites are deployed with a working application pipeline; pricing is published; the first five supervised seats are open; the reviewer dashboard — **Donna Desk** — is live and built out well past MVP, encoding the operating desk Dawn already runs.

ORIGIN	POSITIONING	PARENT
Born from a working desk The blueprint is Abby — the agent Zoe built that automates a real controller's finance reporting, document pipelines and website upkeep, cutting hours of manual work to minutes. Donna generalises that pattern to the executive desk.	AI-enhanced human Clients trust a person, not a model. The visible face is a trusted EA; the AI turbocharges her output. Marketing never says "replace your EA" — it says one assistant now covers three desks.	Carrier, not startup Phone systems, business internet and virtual receptionists already run on Fibernetics infrastructure daily. Donna inherits the operating discipline, the brand and the customer base.

03 · PROBLEM & SOLUTION

An assistant they can't hire. An agent they can't trust.

THE PROBLEM

Both options fail

Leaders lose ~18 hrs/week to admin. A senior EA costs \$10–13.5K CAD/month loaded and takes months to find. Raw AI agents are instant and cheap — and one confidently wrong email under an executive's name ends the pilot. 40% of agentic projects will die on exactly this (Gartner).

THE SOLUTION

Draft → review → release

Donna drafts at agent speed. A named person at the NEWT desk reads everything outbound within 30 business-minutes. The customer approves only money, commitments and their own name. Every action is logged, attributable, reversible. Silence is a no.

The insight competitors keep missing: **supervision is not overhead to eliminate — it's the feature customers are actually buying.** The EA-facing version inverts the same product: the customer's own assistant becomes the reviewer, and Donna makes her the most capable person in the building instead of the most threatened.

04 · PRODUCT & HOW IT WORKS

One agent, one desk, three authority tiers.

The customer draws the line once on an onboarding call; it never moves without them.

TIER 1 · AUTOMATIC

Reversible and internal: calendar moves, filing, triage, drafts into the drafts folder. Donna acts; the desk audits after the fact.

TIER 2 · PERSON-CHECKED

Leaves the building but recoverable: external replies, scheduling with outsiders, chasing commitments. A named reviewer reads in full and sends, fixes, or escalates.

TIER 3 · YOUR SIGNATURE

Money, promises, the customer's name: waits indefinitely for an explicit yes. No timeout, no default-approve.

The Desk dashboard — the real product

What Dawn uses today becomes what we sell tomorrow: a single pane of glass where one reviewer runs many principals. Per-person buckets (email, calendar, tasks, notes, documents), a review queue ordered by risk tier, an insights page that surfaces patterns (missed sales calls, recurring requests) as proactive suggestions, and one-tap escalation to the principal. Channels meet people where they are: email and calendar natively, iMessage/SMS domestically, WhatsApp for international scale, Teams and call transcription next.

Architecture

- **Claude (Anthropic) + Mosaic, then the box** — frontier reasoning plus an instant-deploy beta scaffold; the same stack shipped both sites and the live product in days, not months. Production lands on a Fibernetics-operated box in

Canada: the migration is two documented seams (database and password auth — everything else is already self-contained; see DEPLOYMENT-BOX.md in the repo).

- **Model failover** — automatic in-family fallback is live today (Opus → Sonnet on transient failure; a weaker model degrades polish, never safety — everything still passes review). Second-provider and local fallback land in the funded plan.
- **Data residency** — production runs on Fibernetics' own hardware in Canada (the final host is a box we operate, not a cloud); the beta scaffold runs on vetted cloud infrastructure under the DPA, stated plainly to customers. Model inference is Anthropic's API until the funded second-provider/local failover lands. Content scoped to connected mailboxes; not used to train general models; exported and removed on exit.
- **No new portal for customers** — Donna works inside the email, calendar, phone and finance tools the customer already runs. The dashboard is for the reviewer.

WORKING TODAY

Live: both product sites, the six-question qualification quiz, published pricing, a verified application pipeline (form → database → notification email, tagged by seat) — and the **Desk dashboard**, well past MVP: risk-ordered review queue (high-stakes replies surface first, replies hard-locked to human review), AI triage of pasted or forwarded mail with voice-matched reply drafts and per-executive voice profiles, meeting-notes and contact extraction with source-quote receipts, follow-up chasing, a travel desk that assembles itineraries and preps re-books, expenses to claim-ready CSV, real calendar sync with protected focus blocks, the “Donna noticed” anticipation ledger, one-tap escalation to the principal, morning briefing emails, full activity record with compliance export, one-file data export, and universal ingest (Gmail/Outlook connectors, drag-anything drop, a watched desktop folder). Applications double as a workable CRM with Donna drafting the promised two-business-day reply for approval. Not yet live (and said so publicly): Finance, Support and Voice agents; deep OAuth write-back to Google/Microsoft calendars is in the funded plan.

05 · MARKET ANALYSIS

A labour market priced for scarcity, meeting a software market priced for scale.

All figures below are illustrative planning assumptions from published 2026 ranges, stated so they can be checked — not performance guarantees.

~\$150B NA EXEC/ADMIN LABOUR SPEND	~2M EAS & ADMIN ASSISTANTS, NA	\$47B AGENTIC-AI MARKET BY 2030 (EST.)	40% AGENTIC PROJECTS SCRAPPED BY '27
---	---	---	---

- **TAM — the work itself:** executive/administrative support labour in North America alone runs on the order of \$150B/yr (≈2M roles × ~\$75K average loaded cost). Donna monetises a slice of that work at a fraction of its labour price.

- **SAM — supervised seats + EA software:** SMB and mid-market executives without an EA (the NEWT customer profile) plus the ~2M EA population as software seats. Illustratively: 500K reachable executive seats × \$18K/yr + 2M EA seats × ~\$1.2K/yr ≈ a \$11–12B serviceable pool.
- **SOM — the plan's horizon:** ~4,000 supervised seats and ~25,000 platform seats ≈ **\$100M ARR** — under 1% of SAM.
- **Why now:** three-quarters of enterprises say they've adopted agentic AI; almost none have anything genuinely live. The distance between the two is supervision — a job that has to be staffed, trained and productised, which is an operations problem before it is a software problem. Carriers do operations.

06 • TARGET CUSTOMERS

Who buys, who uses, who spreads it.

BUYS FIRST

The exec who can't hire

SMB founders and executives (the Jordan-at-Delta-Air profile): drowning in admin, can't justify or find a \$130K EA, burned or scared by raw AI. Buys the supervised seat.

USES & CHAMPIONS

The overloaded EA

Supports more principals than the role was scoped for. Gets Donna free in beta, becomes the reviewer, covers 3+ execs — and becomes the product's evangelist instead of its victim. Converts to a paid software seat.

DISTRIBUTES

Channel partners

VARs already selling NEWT phone systems experience Donna internally, then resell her to their book. International reach follows WhatsApp, where numbers don't gate by country.

07 • COMPETITIVE ANALYSIS

Everyone else sells the model. We sell the signature.

	HUMAN REVIEW BUILT IN	NAMED ACCOUNTABLE PERSON	RUNS IN YOUR TOOLS	PRICE SHAPE	WEAKNESS DONNA EXPLOITS
Donna (NEWT)	Yes — core design	Yes, you meet her	Yes	\$1,500 CAD/seat/mo · EA seat SaaS	—
Autonomous AI EAs (Lindy, Martin, Ohai)	No — autonomous by design	No	Partial	\$30–500/mo	One wrong email under your name ends the pilot; no one owns the output
AI email tools (Fyxr, Superhuman AI)	You are the reviewer	No	Yes	\$30–90/mo	Moves the reading to your evening instead of removing it
Suite copilots (Microsoft Copilot, Gemini)	You are the reviewer	No	Native	\$20–40/user/mo	Generic; no service, no accountability, no desk
Human VA services (Belay, Time etc, Athena)	Human does the work	Rota / offshore	Yes	\$1,600–3,400+/mo part-time	Human speed at human cost; one set of hands; vacations
Hiring a senior EA	N/A	Yes	Yes	\$10–13.5K/mo loaded	Can't hire them fast enough; single point of failure

The moat, honestly

- **The reviewer network + correction loop.** Corrections don't go into a ticket queue — the reviewer is the designer, so a correction changes the product that week. Every supervised seat compounds a proprietary dataset of "what good judgment looks like" that model vendors don't have.
- **The EA flywheel.** Every other vendor threatens EAs; we arm them. The free assistant seat recruits the exact population that (a) buys software seats, (b) staffs our desk as we scale, and (c) sits between every vendor and every executive inbox as gatekeeper.

- **Carrier trust + channel.** Fifteen years of business-critical uptime, data in Canada, one accountable throat — plus an installed base and VAR network no agent startup can rent.
- **Switching cost.** Connected mailboxes, calendars, tone models and authority tables make seat N+1 cheaper for us and leaving costlier for them.

08 • GO-TO-MARKET

Concentric circles, each one funding the next.

- **Circle 0 — the 5-seat beta (now).** Local, friendly businesses within ~100 km for in-person support; warm leads first. Free during beta, published \$1,500 price, capacity cap stated plainly. The intake pipeline is live and tags applicants by seat.
- **Circle 1 — VARs as customers (Q1'27).** Resellers of NEWT phone systems get seats before they sell them — the demo is their own desk. Their book of SMBs becomes the first paid cohort.
- **Circle 2 — the EA community (rolling).** donna-for-eas is the acquisition engine: free assistant seats, a written "pitch your exec" note in the product, and the only vendor message in the category that EAs can forward without fearing for their job. EAs pull Donna into companies; execs don't have to be sold.
- **Circle 3 — carrier bundle + international.** Donna joins the NEWT bundle (voice + internet + receptionist + AI employee, one bill, one throat). WhatsApp unlocks non-NA markets where iMessage/SMS numbers fragment.

MESSAGE DISCIPLINE

To executives: "hire your first AI employee — a real person signs her work." To EAs: "you already are the assistant; this is how you cover three of them." Never "replace your EA." The honesty register — published pricing, stated limits, no SOC 2 yet, seats genuinely capped — is itself the differentiator in a category drowning in overclaim.

09 • BUSINESS MODEL

Two engines, one substrate.

STREAM	PRICE (CAD)	REVIEWER	ILLUSTRATIVE GROSS MARGIN	ROLE IN THE PLAN
Supervised executive seat	\$1,500 / mo	NEWT desk	~50–60% at 8–10 seats/reviewer	Revenue now, trust proof, correction data
EA platform seat	Free in beta → ~\$99–299 / mo	The customer's EA	~85%	Scale engine; no desk to staff behind it
Agent roster add-ons (FIN / SUP / VOX)	Per-agent seat pricing	Mixed	Blended	Expansion revenue on the same substrate
Carrier bundle uplift	Bundled MRR	—	—	Retention + distribution, defends core NEWT base

Unit economics — supervised seat (illustrative)

revenue / seat = **\$1,500 / mo**
reviewer (loaded ~\$7.5K/mo) across **8–10 seats** ≈ **\$750–940 / seat**
AI compute + infra ≈ **\$75–150 / seat**
→ gross margin ≈ **\$400–700 / seat / mo** (~50–60% at desk maturity)
platform seat @ \$149: cost to serve ≈ **\$20** → ≈ **85% margin**

The strategic point: the supervised seat is priced at roughly **a seventh of a loaded EA and below a part-time offshore VA** — while the platform seat costs less than an EA's monthly parking. Neither engine asks the customer to believe anything about AI; both ask them to believe a person they've met.

The desk is the factory.

- **Review SLA:** outbound checked in under 30 business-minutes (2 hrs off-hours); when we'll miss it, Donna hands the customer the finished draft to send themselves — the customer is never blocked on us.
- **Desk scaling:** Dawn is reviewer #1 and the training standard. Each new reviewer is an experienced EA, trained on the authority-tier playbook, shadowing live seats before holding their own. Target ratio: 8–10 supervised seats per reviewer at maturity; customers are told by name when anyone new touches their account.

- **Onboarding:** Day 0 thirty-minute call and authority table; days 1–3 shadow mode (Donna reads, touches nothing); days 4–10 Tier 1 live; days 11–30 the customer moves the line.
- **Reliability:** operated on carrier discipline — monitored, multi-model failover, incident paths that already exist for voice and network customers.
- **Security roadmap:** data-processing terms before anything connects; scoped mailbox access, revoked on disconnect; SOC 2 Type I in the funded plan. Until certified, we say so plainly — the honesty is load-bearing.

11 • TEAM

The person, the builder, the carrier.

PRODUCT • OPERATIONS

Dawn Gamble

Senior Executive Assistant, Fibernetics — fifteen years running the founders' desk. Designed Donna from her own job, reviews every beta seat, and trains the desk as it grows. The face customers meet on day zero.

ENGINEERING • AI

Zoe Cross

Built and shipped the stack: the agent architecture (proven on Abby's finance desk), both product sites, and the live application pipeline. Owns the Claude + Mosaic platform, failover design and the Desk dashboard build.

PARENT • CHANNEL

Fibernetics / NEWT

Canadian carrier since 2011: infrastructure, uptime discipline, installed SMB base, VAR network, brand, and the founders' direct sponsorship of the project.

12 • FINANCIAL PLAN

From five seats to \$100M ARR — the arithmetic, not the poetry.

Every figure is an illustrative planning assumption for decision-making, not a forecast.

PHASE (ILLUSTRATIVE)	SUPERVISED SEATS	PLATFORM SEATS	ARR (CAD)	DESK STAFF	STANCE
Beta → mo 6	5 → 25	50 free	\$0 → \$450K	Dawn + 2	Investment
Year 1	100	500 @ \$149	≈ \$2.7M	~12	Near break-even
Year 2	400	3,000	≈ \$12.5M	~45	Profitable desk
Year 3–4	1,500	10,000	≈ \$45M	Regional desks	Roster live (FIN/SUP/VOX)
Year 5 target	4,000	25,000+	≈ \$100M	Franchise-model desks	≈ \$1B at ~10× ARR

- **Cost structure:** reviewers (the dominant variable cost, ~50% of supervised-seat revenue at maturity), AI compute (~5–10%), platform engineering (fixed, small — the stack is deliberately lean), sales via existing channel (low CAC by design).
- **Break-even (illustrative):** ~35–40 supervised seats carries the early team; the platform seat's 85% margin pulls blended margin up as mix shifts toward software — which is the whole strategic design.
- **Funding need: \$1.5M CAD for 18 months** (internal venture allocation or seed): Desk dashboard deepening (OAuth write-back, channels), 6 reviewers hired and trained, SOC 2 Type I, WhatsApp + voice channels, VAR partner program, and marketing that stays in the honesty register.
- **Use of funds:** ~45% engineering, ~30% desk staffing & training, ~15% compliance/security, ~10% GTM.

Three horizons.

Horizon 1 · now → 6 mo

Prove the desk

Fill 5 beta seats from live intake. Deepen the shipped Desk dashboard (Google/Microsoft OAuth write-back, iMessage/SMS channel). Convert beta → paid with an explicit yes. Hire reviewers 2–3. SOC 2 program start. VAR pilot with 2 resellers.

Horizon 2 · 6–18 mo

Scale the two engines

EA platform seat GA with self-serve pricing. FIN agent live (invoice chasing, expense coding — reviewed before anything posts). WhatsApp channel for international. 100+ supervised seats, 500+ platform seats.

Horizon 3 · 18 mo +

Own the category

VOX voice agent on NEWT's own telephony. Donna-certified reviewer program → a marketplace where any certified EA sells supervised capacity on our rails, and we take the platform cut. The trust layer for agentic work.

14 · RISKS & MITIGATIONS

What kills this, and what we've done about each.

Reviewer scaling is the bottleneck

→ By design: the platform seat needs no desk (customer's EA reviews), the certification program turns the labour constraint into a marketplace, and seats-per-reviewer rises as Tier 1 automation earns trust.

Key-person risk: Dawn

→ The authority-tier playbook and training program exist precisely to encode her judgment; reviewer #2–3 are hired in Horizon 1; the correction loop writes her standards into the product weekly.

Suite copilots commoditise drafting

→ Let them — drafting was never the product. Copilots make our COGS cheaper and still leave the customer as their own reviewer at 11 p.m. Accountability plus service doesn't ship in a suite SKU.

Services margin drags the multiple

→ Mix-shift is the plan: supervised seats are the wedge and the data engine; platform seats at ~85% margin become the majority of ARR by Year 3 (illustrative).

No security certification yet

→ Stated plainly on the public site today; DPA before anything connects; scoped access, revoked on disconnect; SOC 2 Type I funded in the ask. Honesty now beats implication later.

Model/vendor dependency

→ In-family failover is already live (Opus → Sonnet on transient failure); a second provider is in the funded plan — and the review layer means a weaker model degrades speed, not safety.

A public error under a customer's name

→ Tier design makes it structurally hard (Tier 3 waits indefinitely); full action logs make any incident auditable; the correction becomes a product change the same week — and we say so.

Beta converts poorly at \$1,500

→ Price is published now specifically so conversion is tested honestly from day one; fallback positioning against \$1,600–3,400 part-time VAs holds margin; platform seat catches the price-sensitive.

15 • THE ASK & NEXT STEPS

Green-light the venture. The beta is already live.

- **Decision:** constitute Donna as a NEWT product line with **\$1.5M CAD / 18-month** allocation and Dawn formally allocated to it.
- **Next 30 days:** fill the 5 beta seats from live intake (first target: the warm Delta Air lead); put the shipped Desk dashboard in Dawn's daily hands; begin reviewer #2 recruitment from the EA community funnel; kick off SOC 2 scoping.
- **Success gate for the next tranche:** 5 seats live with retention through day 30, review SLA held, and a signed VAR pilot — evidence, not projections.

THE ONE-SENTENCE CASE

Everyone is racing to make agents more autonomous; the customers with money are begging for someone to make them **accountable** — and accountability is the only part of this Fibernetics has already been selling for fifteen years.

Market sizes, financial projections, margins and seat counts in this document are illustrative planning assumptions drawn from published 2026 ranges and internal estimates; they are not forecasts or performance guarantees. Product claims describe the live beta as deployed. Prepared with the founders' knowledge for internal decision-making.