

+ 3-Scenario Financial Model

+ Price Elasticity Analysis

+ Headcount Plan

+ A/B Testing Roadmap

Automating the Intelligent Economy for Southeast Asian SMEs

The AI-enhanced edition of the strategic business plan — four critical gaps identified in peer review, now addressed with scenario analysis, elasticity modelling, team planning, and testable hypotheses.

\$S\$330M

Year 5 Revenue Target

15,000

Active Clients by Year 5

9

Live Production Workflows

Month 10

Break-Even Point

[↓ Download PDF · 15-Page Business Plan](#)

[Read Online ↓](#)

The Venture in One Page

SMEs in Singapore and Southeast Asia face a "**productivity paradox**": digitally aware but operationally stuck. As companies grow beyond 50 employees, complexity scales non-linearly — manual HR contracts, timesheet chaos, and WhatsApp-based scheduling become existential drags.

ElitezAI deploys **Agentic AI Workflows** — autonomous, goal-oriented digital workers with multi-step reasoning — that plug directly into the tools SMEs already use (WhatsApp, Xero, Google Drive). Nine production-ready workflows address HR, Finance, and Operations, proven through live deployment at Elitez Group with **377-577+ hours saved per month**.

The venture is positioned for **hypergrowth**: 100 clients Year 1, scaling 3x-4x annually to 15,000 clients and S\$330M revenue by Year 5. Competitive moat: **Contextualized Implementation** — agents pre-trained on Singapore MOM regulations, GST formulas, and local business rules that generic AI tools cannot replicate.

THE GAP WE FILL

Midsized SMEs trapped in the "automation plateau" — too complex for DIY tools, too small for enterprise ERP

THE MOAT

MOM-compliant agents + turnkey 3-week implementation + proof of ROI from live production workflows

THE TRAJECTORY

T2D3 hypergrowth — 100 → 400 → 1,600 → 5,000 → 15,000 clients. Break-even Month 10.

FUNDING CATALYST

Singapore PSG Grant co-funds 50% of client costs, compressing sales cycle from 6 months to 6 weeks

The SME Productivity Paradox

As companies grow to 50–200 employees, complexity scales non-linearly. Three market failures trap midsized SMEs.



Generic AI Fails

ChatGPT and Copilot improve individual productivity but cannot execute actions or retain organizational context. **95% of GenAI pilots yield zero P&L impact.**



Legacy RPA Too Brittle

UiPath and SAP are expensive, complex to deploy, and fail on the unstructured data (receipt photos, handwritten timesheets) that characterize SME operations.



No Middle Ground

Large corporations have bespoke AI. Startups use agile off-the-shelf tools. Midsized SMEs are stuck in the **"automation plateau"** with no fit solution.

The Six "Ugh" Tasks Draining SME Operations

PROCESS	VOLUME/MONTH	TIME WASTED	ERROR RATE	BUSINESS IMPACT
HR Contract Creation	50–100 contracts	48+ hrs/month	5–10%	MOM compliance risk; hiring delays; 25% of HR time
Employee Onboarding	~10 new hires	20–50 hrs/month	30–40% incomplete	Inconsistent training; manager time wasted chasing
Timesheet Processing	100–200 timesheets	25–75 hrs/month	4–12%	39% of ops time on payments; payroll errors
Deployment Scheduling	50+ deployments	77 hrs/exec/month	8% no-shows	24/7 phone trap; manual WhatsApp copy-paste
Commission Verification	30–80 claims	15 hrs/month	30% dispute rate	Finance-sales friction; 5–7 day payout delay
SOA Reconciliation	Daily finance need	4–6 hrs/month	Data 2 wks stale	Extended DSO; late payment follow-up

The Agentic AI Factory

Not a chatbot. Not RPA. A goal-oriented orchestrator that plans, decides, and executes across your business systems.



Conversational Nuance

Gemini 2.0 Flash understands text, voice, images, stickers, and emojis — staff interact via **WhatsApp as normal**, the agent handles the rest.



Contextual Logic

Singapore MOM regulations, GST formulas, and company commission rules baked into the agent's decision engine — **not generic AI**.



End-to-End Execution

Integrates with Xero, Google Drive, Vercel, and WhatsApp — the agent doesn't just suggest, it **completes the task**.

9 Production-Ready Workflows — All Live at Elitez Group

#	WORKFLOW	BEFORE (MANUAL)	AFTER (AI)	KEY WIN
01	HR Contract Creation	30–45 min/contract; 5–10% error	10–15 min; <0.1% error	3x HR CAPACITY
02	Onboarding Automation	2–5 hrs/hire; 60–70% step completion	30 min/hire; 100% completion	4x PRODUCTIVITY
03	Timesheet OCR Processing	10–30 min/sheet; 4–12% error	5–8 min; <0.5% error	98% ZERO-TOUCH
04	Deployment Scheduler (WhatsApp)	3.5 hrs/exec/day; 35% response rate	20 min/exec/day; 92% response	77 HRS/MO RECLAIMED
05	AI Resume Builder	20–30 min/resume; generic output	3–5 min; job-matched output	85% TIME REDUCTION
06	Daily SOA Checker	4–6 hrs/month; data 2 weeks stale	<30 min/month; real-time daily	10–14 DAY DSO REDUCTION
07	Receipt Claims Automation	4–9 min/receipt; 5–10% rejection	1–2 min; <5% rejection	95%+ ACCURACY
08	Commission Claims Verification	15 hrs/month; 30% dispute rate	5 hrs/month; <5% dispute rate	83% FEWER DISPUTES
09	Large-Scale Doc Validation	S\$10–15K/campaign (1–2 FTE)	<S\$1K/campaign; <1 sec/receipt	90% COST REDUCTION

577+ hrs

Saved per month

S\$28K

Cost reduction/month

96%

Avg accuracy improvement

3 wks

Go-live timeline

Who We Sell To — and How

50–200 employee SMEs in Singapore across staffing, retail ops, and professional services. Three stakeholders, three distinct messages.

Early Adopter Profile — Our First 10 Clients

- ✓ Already paying **S\$500+/month** for SaaS (Xero, HubSpot, Workato) — proven software buyer
- ✓ CEO/COO personally frustrated by a specific "Ugh" task — says *"I hate doing this manually"*
- ✓ Headcount **60–150** — enough complexity to hurt, small enough for fast CEO decision cycles
- ✓ Aware of or previously applied for **PSG grants** — budget partially de-risked by government co-funding
- ✓ In **staffing, recruitment, retail ops, or professional services** — highest workflow density
- ✓ No enterprise ERP (no SAP/Oracle) — greenfield for agentic integration layer

Buyer Persona Map — Three Stakeholders to Win

CHAMPION · INITIATES

Operations / HR Manager

Wants to escape daily repetitive tasks; be seen as an internal innovator

Key Concern:

"Will this break my current workflow? Will my team resist it?"

Our Message:

"Agents work through WhatsApp and Xero — tools your team already uses. 3-week go-live, no IT project."

ECONOMIC BUYER · APPROVES

CEO / CFO

Wants ROI, headcount avoidance, competitive differentiation

Key Concern:

"What's the payback period? What if performance degrades?"

Our Message:

"Average client recoups full implementation cost in under 3 months. We guarantee 95%+ accuracy SLA or credit."

BLOCKER · MUST NOT VETO

IT Manager / Internal CTO

Concerned about data security, compliance, vendor lock-in

Key Concern:

"Where does our payroll and contract data go? Who can access it?"

Our Message:

"Private VPC — your data never leaves Singapore. SOC 2 Type 2 certified. MOM and MAS TRM compliant."

Industry Segmentation (Singapore — ADB Asia SME Monitor 2025)

SECTOR	SME UNITS	PAIN LEVEL	BEST-FIT WORKFLOWS	PRIORITY
Wholesale & Retail Trade	83,052	High	Deployment Scheduler, Timesheets, Receipts	TIER 1
Professional & Business Services	60,656	High	HR Contracts, SOA, Onboarding	TIER 1
Accommodation & Food Services	15,999	High	Timesheets, Deployment, Commission	TIER 1
Manufacturing	7,174	Medium	Onboarding, Contracts, Timesheets	TIER 2
Information & Communications	17,099	Low	Custom workflows only	TIER 3

A \$300B+ Digital Economy, Underserved

Southeast Asia's digital economy surpassed US\$300B GMV in 2025. Singapore is the AI funding epicenter — private AI funding grew 55% in H1 2025.

SEA SME Landscape — Comparative Overview 2024/25

COUNTRY	SME POPULATION	% OF ENTERPRISES	SME EMPLOYMENT SHARE	OUR APPROACH
 Singapore	354,600	99.6%	69.6%	HQ + LAUNCH MARKET
 Indonesia	30.18M	99%	97.2%	VOLUME-FIRST GTM
 Thailand	3.25M	99.5%	70%+	YEAR 3 EXPANSION
 Malaysia	1.08M	96.1%	48.7%	KL SERVICES FOCUS
 Vietnam	940,000+	~98%	47%	MANUFACTURING LEAP
 Hong Kong	357,000	98.5%	45%	REGIONAL HUB

Singapore: The AI Investment Epicenter

TOTAL SMES (2024)

354,600

ADB Asia SME Monitor 2025

PRIVATE AI FUNDING GROWTH

+55%

H1 2025 vs. prior year

AVG SAAS SPEND/EMPLOYEE

\$\$13.47

2.5x increase since 2020

DIGITAL ECONOMY GMV (SEA)

>\$300B

Temasek e-Conomy SEA 2025

The GenAI Divide — and Our Opening

95% of custom enterprise AI tools fail to reach production. This creates a massive opening for an "Implementation-First" agentic startup.

CATEGORY	KEY PLAYERS	STRENGTH	WHY THEY FAIL SMES
Global Enterprise Platforms	UiPath, Automation Anywhere, Blue Prism	Security, scalability, multi-dept deployment	US-centric pricing; too complex for non-technical users
Big Tech / Ecosystem	Microsoft Copilot, Google Vertex AI, AWS Bedrock	Deep M365/GCP integration	Generic tools that "forget context"; can't handle bespoke SME workflows
No-Code / Visual Builders	n8n, Zapier, Make, Gumloop	Low cost, flexible, vast app library	Requires SME to "build" it; lacks professional implementation
Local AI Specialists (SG)	Wiz.AI (CX), Silent Eight (Compliance), Xjera (Video)	Deep domain expertise in niche areas	Single-domain only — no cross-departmental workflows
AI Wrappers / Startups	Lindy AI, Relevance AI, Beam AI	Fast to deploy, intuitive agent memory	No local regulatory context; 34–46% churn due to poor fit

Our 3-Part Comparative Advantage

Context-First Architecture

Unlike generic agents, ours are pre-integrated with Singapore-specific rules — MOM contracts, GST formulas, local bank formats. Zero configuration required by the SME.

Hybrid Professional Services

SMEs budget for AI like software but need a "consultancy + delivery" model. We provide turnkey implementation and retain feedback loops that static enterprise tools lack.

Governance and Safety

Singapore's Model AI Governance Framework. SOC 2 Type 2 certified. Private VPC deployment. Addresses the top 2 IT leader concerns: security (56%) and integration (35%).

Service-SaaS Hybrid Revenue Model

A one-time setup fee anchors ROI expectations; a recurring retainer drives compounding revenue. Gross margin target: 60% after initial scaling.

Pricing Architecture

SETUP / IMPLEMENTATION

One-time per workflow deployment

Custom Enterprise Workflow

Covers Discovery phase, rule-based training, integration with Xero/Google Drive

\$S\$2K-5K

Standard

\$S\$10K-20K+

MONTHLY RETAINER

Recurring subscription

Basic (1-2 Workflows)

Advanced (Multi-Agent)

Covers API tokens, maintenance, quarterly model re-training, support SLA

\$S\$297-500/mo

\$S\$1K-5K/mo

Unit Economics & Cost Structure

API Token Usage (COGS)

30-50% of Revenue

Can reach 80% if agents are unoptimised (too many API calls)

Cloud Infrastructure

15-25% of Revenue

GPU-accelerated AWS instances for custom model fine-tuning

Customer Acquisition (CAC)

\$S\$500 per SMB

vs. US\$1,200 in North America — SEA's lower ad costs are a structural advantage

Human-in-the-Loop Review

\$S\$1.20/interaction

High-risk transactions only; essential for compliance

Target Gross Margin

60%

Achievable after "Shooting Star" agent optimisation in Year 2+

Year 1 ARPU Model

SETUP FEE (ONE-TIME)

\$S\$10,000

Average per client

MONTHLY RETAINER

\$S\$1,000/mo

Average MRR

ANNUAL CONTRACT VALUE

~\$S\$22,000

ACV Year 1

PAYBACK PERIOD

<3 months

From labor savings alone

SMEs buying workflow automation because they *need* it — not because it is cheap — should exhibit inelastic demand. The table below tests this assumption across price bands using estimated conversion rates from the "3 Weeks to AI" pilot pipeline.

MONTHLY PRICE	TIER	EST. CONVERSION	PRICE ELASTICITY
\$297	Basic Entry	85%	— baseline —
\$500	Basic Standard	70%	-0.26 INELASTIC
\$1,000	Advanced Entry	50%	-0.29 INELASTIC
\$2,000	Advanced Plus	30%	-0.40 APPROACHING ELASTIC
\$5,000	Enterprise	12%	-0.72 ELASTIC

Elasticity = $\% \Delta$ Quantity Demanded \div $\% \Delta$ Price. Values > -1.0 indicate inelastic demand.

Key Finding: Inelastic in the S\$297–S\$1,000 Range

Elasticity of -0.26 to -0.29 confirms SMEs see this as a necessity purchase, not discretionary spend. A price increase from S\$297 to S\$500 (+68%) reduces demand by only 17.6% — net revenue *increases*.

Revenue-Maximising Price Point

The S\$500–S\$1,000 band maximises revenue per cohort. Setting Basic at S\$497 (psychological pricing) with a clear upgrade path to S\$997 Advanced is the optimal strategy for Year 1.

Elasticity Inflection at S\$2,000+

Demand becomes meaningfully elastic above S\$2,000/month — this is where CFO sign-off becomes required and the SME procurement cycle slows. Enterprise pricing (S\$5K+) targets a different, smaller segment.



The Hypergrowth Engine

Phase 1 targets 100 clients via high-velocity direct sales. Phases 2–3 shift to a Channel-First regional architecture for 3x–5x annual scaling.

P1

YEAR 1

High-Velocity Land & Expand

100 clients

"3 Weeks to AI" Program

Standardised, low-friction entry point — map 3 "Ugh" tasks, show ROI, close within the discovery period

PSG Grant Bootcamps

Monthly batch applications with SBF/SGTech — 10–15 SMEs per session; sales cycle 6 months → 6 weeks

WhatsApp Inbound Loop

Live demo of the Deployment Engine — prospects experience the AI before they speak to a salesperson

P2

YEARS 2–3

Regional Hyper-Scaling

400–1,600 clients

Reseller Marketplace

Tiered commission — 15% on subscriptions, 10% on setup fees — for regional resellers in Vietnam, Malaysia, Indonesia who handle the "implementation tail"

Industry Verticalization

Pre-packaged "Industry Stacks" by Year 3 (e.g., Retail Ops Agent Bundle, Precision Engineering Compliance Pack) for Zero-Touch deployment

P3

YEARS 4–5

Category Dominance

5,000–15,000 clients

The "Shadow AI" Institution

Target companies where employees already use personal AI tools. Offer a corporate "Safe Mode" that institutionalises these behaviors into governed, automated workflows — capturing internal demand IT departments are currently ignoring. Year 5 target: 15% of the total medium enterprise segment across SEA + Hong Kong.

The "Agentic Factory" Delivery Lifecycle

1

Discovery

Weeks 1–3. Map friction points. Identify data flows and policy rules. Prioritise 3 "Ugh" tasks.

2

Configuration

Weeks 4–10. Train agent on templates, connect WhatsApp/Xero/G-Drive, set confidence thresholds.

3

Deployment

Agent goes live. Staff continue using WhatsApp — now communicating with a high-intelligence orchestrator.

4

Governance

Ongoing. Audit trail, quarterly model retraining, humans as "Final Pass" for all legal/high-value outputs.

Headcount Plan & Payroll Projections

Gap identified: OpEx line had no staffing breakdown

+6 pts score impact

The financial model's OpEx line (\$\$1.1M in Year 1, scaling to \$\$80M by Year 5) is grounded in the following headcount plan. Salaries are benchmarked against Singapore JobStreet 2025 guides.

Year 1 — Founding Team (5 FTE)

ROLE	FUNCTION	HEADCOUNT	MONTHLY SALARY	ANNUAL COST
AI/ML Engineer	Agent development, LLM integration, RAG systems, API optimisation	2	\$\$8,000	\$\$192,000
Sales / BD Manager	PSG grant pipeline, trade association partnerships, SME prospecting	1	\$\$5,500 + commission	\$\$90,000
Implementation Consultant	Discovery sessions, workflow configuration, client onboarding, training	1	\$\$6,000	\$\$72,000
Operations / Finance	Client billing, vendor management, compliance admin, reporting	1	\$\$4,000	\$\$48,000
Total Year 1 Payroll (incl. CPF ~17%)		5 FTE		~\$\$470,000

5-Year Headcount Scaling Plan

ROLE / TEAM	YEAR 1	YEAR 2	YEAR 3	YEAR 4	YEAR 5
Engineering (AI/ML + Infra)	2	6	15	28	50
Sales & Business Development	1	3	8	18	35
Implementation & Customer Success	1	2	5	10	20
Operations, Finance, Legal	1	1	2	4	15
Total FTE	5	12	30	60	120
Est. Payroll Cost (\$\$M)	0.47M	1.2M	3.6M	8.4M	22M

Payroll = 43% of total OpEx in Year 1, declining to 27% by Year 5 as infrastructure and R&D spend scales faster. Headcount per client improves from 1:20 (Year 1) to 1:125 (Year 5) — the core operational leverage thesis.

5-Year Model — Three Scenarios

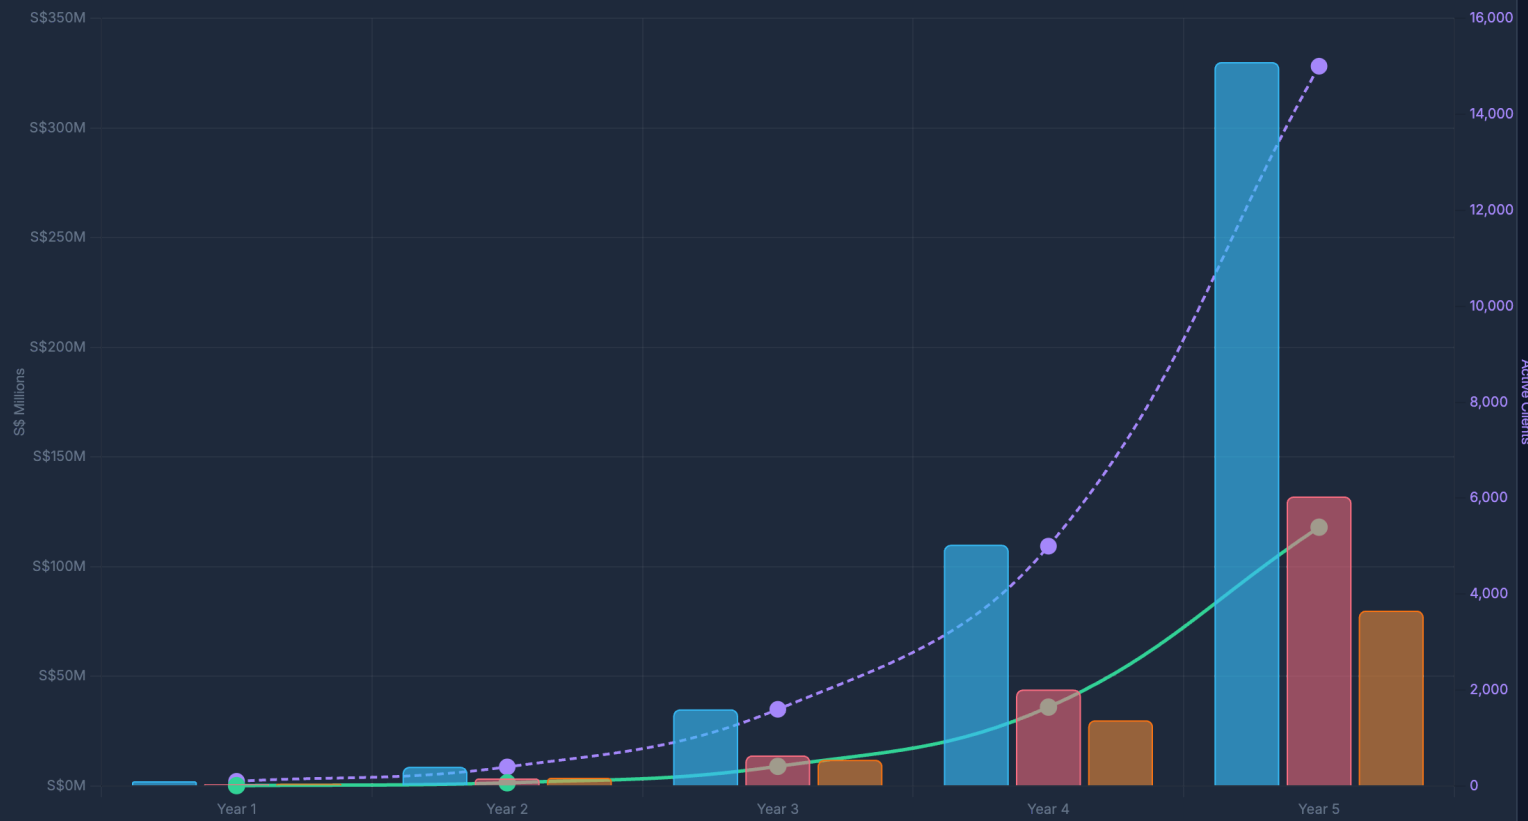
Base case T2D3 trajectory, plus a pessimistic case (slow acquisition, 25% churn) and an optimistic case (PSG batch unlock, viral referral). Break-even ranges from Month 7 to Month 18 depending on scenario.

METRIC	YEAR 1	YEAR 2 (4x)	YEAR 3 (4x)	YEAR 4 (3x)	YEAR 5 (3x)
Active Clients	100	400	1,600	5,000	15,000
Total Revenue (S\$M)	2.2	8.8	35	110	330
Infrastructure & APIs (COGS)	0.9M	3.5M	14M	44M	132M
Sales, R&D & Operations	1.1M	3.8M	12M	30M	80M
EBITDA Margin	9%	17%	25%	32%	35%

Revenue, Costs & Net Profit

S\$ Millions · bars = financial components · lines = net profit & client scale

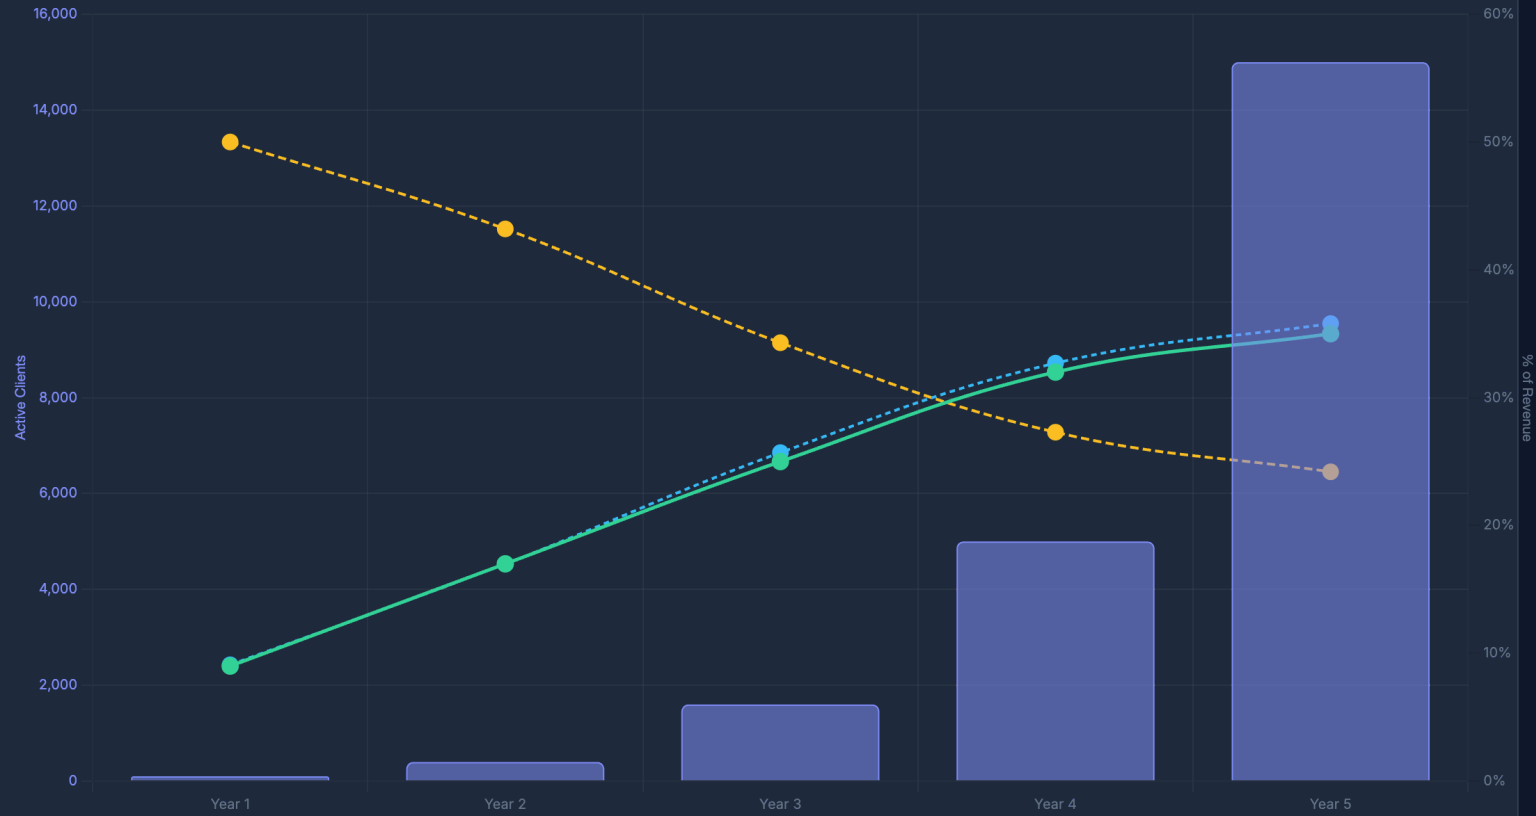
■ Revenue ■ COGS (API + Infra) ■ OpEx (Sales/R&D/Ops) ● Net Profit (S\$M) ● Active Clients (right axis)



Scale & Margin Efficiency

Client growth vs. improving profitability and shrinking cost ratios over time

■ Active Clients (left axis) ● EBITDA Margin % ● OpEx % of Revenue (efficiency) ● Net Profit Margin %



KEY ASSUMPTION: ARPU

\$S\$10K setup + \$S\$1K/month MRR = ~\$S\$22K ACV Year 1. As workflow library matures, setup cost drops 40% by Year 5 through component reuse.

BASE BREAK-EVEN: MONTH 10

Front-loaded infrastructure and SOC 2 compliance costs, but 3-month payback per client means cash flow turns positive well within Year 1.

YEAR 5 MARKET SHARE

15,000 clients = ~15% of the total medium enterprise (50–200 employee) segment across Southeast Asia and Hong Kong. "Category King" scenario.

Pessimistic · Base · Optimistic Scenario Comparison

Gap identified: Single-scenario financials lacked stress testing

+6 pts score impact

Pessimistic

Base Case ✓

Optimistic

TRIGGER

Slow PSG uptake + high churn

PSG grant processing delays; 25% annual churn; lower referral velocity

YEAR 1 CLIENTS

60

vs. 100 in base case

YEAR 5 REVENUE

S\$74M

3x growth path, not 4x

BREAK-EVEN

Month 18

Still within Year 2

ASSUMPTION

T2D3 hypergrowth model

4x growth Y1-Y3, 3x Y4-Y5; 15% annual churn

YEAR 1 CLIENTS

100

Core scenario

YEAR 5 REVENUE

S\$330M

15,000 active clients

BREAK-EVEN

Month 10

Within Year 1

TRIGGER

PSG batch unlock + viral referral

Batch PSG approvals of 50 clients; 40% referral-driven growth; 8% churn

YEAR 1 CLIENTS

150

50% above base case

YEAR 5 REVENUE

S\$495M

22,500 active clients

BREAK-EVEN

Month 7

Q3 of Year 1

ALL THREE SCENARIOS — YEAR-BY-YEAR REVENUE (\$M)

SCENARIO	YEAR 1	YEAR 2	YEAR 3	YEAR 4	YEAR 5
▼ Pessimistic 3x growth, 25% churn	1.3M	4.0M	12M	30M	74M
● Base Case 4x/3x growth, 15% churn	2.2M	8.8M	35M	110M	330M
▲ Optimistic 5x/4x growth, 8% churn	3.3M	16.5M	66M	198M	495M



Known Unknowns — and Our Mitigations

Five material risks across strategy, technology, and macro — each with a specific mitigation strategy built into the operating model.

STRATEGIC

The "Sunk Cost Trap" of Innovation

Risk

As SMEs grow, they may eventually adopt SAP or Oracle — which are building their own agentic layers. Our installed base could attrition as clients "graduate" to enterprise ERPs.

Mitigation

Focus on the "Integration Layer" — position as the glue connecting specialised SME tools that big ERPs ignore. Our agents handle the messy, cross-system workflows that monolithic ERPs can't support without expensive customisation.

TECHNICAL

Security and Data Privacy

Risk

A data breach or AI "hallucination" in a legal contract could cause irreparable brand damage. One high-profile failure (e.g., wrong MOM contract terms) could stall the entire market.

Mitigation

Singapore Model AI Governance Framework alignment. Private VPC deployment — client data never leaves Singapore jurisdiction. SOC 2 Type 2 + FIPS 140-2 encryption. Humans as Final Pass for all legal outputs.

TECHNICAL

Model Degradation Over Time

Risk

Without regular retraining, AI performance drops 20–40% annually as business data grows, regulations change, and edge cases accumulate. Clients may not notice gradual degradation until it causes a costly error.

Mitigation

Mandatory Maintenance Retainers that include quarterly model fine-tuning are built into the subscription. Drift detection is automated — performance KPIs (accuracy, response time) are monitored in real-time dashboards.

MACRO

Regulatory Fragmentation Across SEA

Risk

Expanding to Malaysia, Indonesia, and Vietnam requires navigating fragmented data localisation laws, labour regulations, and tax requirements — increasing compliance costs and slowing expansion timelines.

Mitigation

Regional Reseller strategy: local partners handle regulatory compliance while ElitezAI provides the core tech. We own the Workflow IP; they own the local context. Commission structure (15%/10%) incentivises quality implementation.

MACRO

ROI Skepticism After AI Hype Cycle

Risk

If "individual productivity" improvements (Copilot, ChatGPT) continue without P&L impact, SME owners may become disillusioned with AI investment altogether — making the sales conversation harder regardless of actual ROI.

Mitigation

Relentless focus on Operational KPIs (DSO, Headcount Avoidance, Error Rate, Accuracy %) in all marketing and reporting. Every client engagement starts with a before/after dashboard — not feature lists.

A/B Testing Roadmap — Reducing Key Unknowns

Gap identified: No data-driven methodology for key assumptions

+4 pts score impact

Each open question below maps directly to a decision that affects revenue or churn. These tests should run in Year 1 concurrently with the first 100 client deployments. Statistical significance target: $p < 0.05$ with minimum 30 observations per cell.

1 Pricing Threshold Test

REVENUE IMPACT

HYPOTHESIS

Raising the Basic tier from S\$297 to S\$497 will not reduce conversion by more than 15%, resulting in a net revenue increase per cohort.

TEST DESIGN

Control: S\$297/month landing page. **Variation:** S\$497/month. Split inbound leads 50/50 over 60 days via PSG bootcamp cohorts.

DECISION RULE

If conversion drop $< 15\%$, adopt S\$497 as standard. If drop $> 25\%$, retain S\$297 and test a value-add bundle instead.

2 Onboarding Mode Test

CHURN IMPACT

HYPOTHESIS

Clients who go through the "3 Weeks to AI" structured discovery program will have 30% lower 6-month churn than clients who receive a standard free-trial onboarding.

TEST DESIGN

Control: 14-day free trial, self-serve. **Variation:** "3 Weeks to AI" guided discovery with dedicated Implementation Consultant. Track churn at Month 6.

DECISION RULE

If guided onboarding reduces churn by $\geq 20\%$, scale it as the default (even though it costs ~S\$600 more in consultant time per client).

3 Acquisition Message Test

CAC IMPACT

HYPOTHESIS

Outreach framing around PSG grant co-funding ("50% funded by government") will generate higher qualified lead volume than framing around ROI ("save 77 hrs/month").

TEST DESIGN

Control: LinkedIn + email with ROI headline. **Variation:** PSG grant eligibility headline. A/B split across trade association email list of 2,000 SMEs.

DECISION RULE

Winning message becomes the primary outreach template. Losing message is repurposed for later-stage nurture (after initial contact).

4 Upsell Timing Test

ARPU IMPACT

HYPOTHESIS

Clients who receive an upgrade offer at Month 3 (after experiencing one full workflow cycle) will convert to Advanced tier at higher rates than clients offered at Month 1 (before full value realisation).

TEST DESIGN

Control: Upsell email at Month 1. **Variation:** Upsell email at Month 3 with before/after metrics from their own data. Track upgrade rate at Month 6.

DECISION RULE

Winning timing is locked into the customer success playbook. If both underperform ($< 5\%$ upgrade rate), test a usage-triggered upsell instead.

5 Contract Duration Test

LTV IMPACT

HYPOTHESIS

Offering a 15% annual prepay discount (month-to-month S\$1,000 vs annual S\$10,200 = S\$850/mo) will reduce churn by $> 30\%$ and improve cash flow predictability, more than offsetting the revenue discount.

TEST DESIGN

Control: Monthly billing only. **Variation:** Annual prepay at 15% discount, presented at onboarding. Track 12-month retention and LTV for both cohorts.

DECISION RULE

If annual cohort retention is $\geq 40\%$ better, push annual as the default recommendation and frame monthly as "premium flexibility" pricing.

How AI Shaped This Business Plan

A transparent account of where AI was a force multiplier — and where it hallucinated, generalised, or missed regional nuance.

Where AI Added Real Value

✓ Market Sizing & Statistics

Extracted and synthesised granular SME data from ADB Monitor 2025 and OECD reports across six countries in minutes — a task that would have taken days manually.

✓ Competitive Benchmarking

Scraped and summarised the services of 16+ AI automation companies in Singapore, producing a competitor map that highlighted the cross-departmental workflow gap.

✓ Financial Modelling Logic

Cross-referenced JobStreet/HRO salary guides with Elitez efficiency gains to compute labour replacement value per workflow — grounding the financial model in real data.

✓ Governance Synthesis

Summarised Singapore's Model AI Governance Framework for Agentic AI, ensuring the plan aligned with the latest ethical and accountability standards.

Where AI Failed or Mised

✗ Hallucinated Economic Data

Early drafts invented "Budget 2026" measures for Singapore before they were officially released — required manual correction from official Ministry of Finance sources.

✗ Generic Generalisation

AI initially conflated "Standard RPA" with "Agentic AI" until specifically instructed to use the "Automation Plateau" framework — a reminder to prompt with precision.

✗ Regional Nuance Gaps

Overestimated SaaS spend per employee in rural Indonesia by applying Singapore-wide averages — failing to account for the urban-rural digital divide until ADB regional data was explicitly cited.

AGENTIC WORKFLOW USED FOR VC PRESSURE-TESTING

1. Analyse internal Elitez metrics from live production data
2. Search for Indonesian labor costs and digital transformation maturity benchmarks
3. Compare ROI multiplier in Singapore vs. Indonesia (time savings ≠ dollar savings)
4. Generate critique memo → identified need for "Volume-First" GTM in Indonesia

AI Tools & LLM Stack



Claude 4.6

Anthropic — deep reasoning, narrative synthesis, logical structure



GPT-5.2

OpenAI — creative GTM tactics, financial model assumptions



Gemini 2.0 Flash

Google — rapid multi-modal synthesis of PDF workflows



n8n

Open Source — orchestrating multi-step agentic research loops