1. What is the investor's role?

The investor's role is primarily as a capital partner enabling growth. Investors are, not expected to operate the business, manage staff, or make day-to-day decisions. The company remains founder-led to preserve speed, clarity, and execution quality. Investors may contribute optional value through introductions, strategic advice, and feedback—especially around enterprise opportunities, but operational authority remains with management to avoid delays and conflicting decision-making.

2. ROI on investment and how much equity percentage of the company?

This round is, structured as a Brainpower AI economic participation pool, not a broad company-wide equity offer. A maximum of 20% of Brainpower AI economic participation is allocated to investors in exchange for a total of OMR 75,000. Each investor's percentage is proportional to their share of the raise. Returns are driven by Brainpower AI's performance, profitability, licensing income, and/or an acquisition or buyout event. These projections are illustrative only; actual results may differ materially, and there can be no certainty of outcomes. This is a performance-based participation structure designed to align investor outcomes with Brainpower AI's growth specifically.

3. What is their involvement in the company?

Investors receive information and economic rights tied to Brainpower AI's performance, but they do not receive operational control. Their involvement is typically limited to quarterly updates, KPI visibility, and optional advisory discussions. Investors do not approve hires, marketing experiments, product changes, or spending line-by-line. This structure is intentional because early-stage execution requires rapid iteration, and the company must be able to act quickly based on product and customer feedback.

4. How long until investors get the ROI?

ROI timing depends on the route: profit distributions or liquidity events. For early-stage software, a realistic expectation is that meaningful liquidity typically takes 2–5 years, with a base target of 24–36 months to reach strong revenue milestones that enable either profit distributions or credible acquisition interest. If enterprise conversion is slower or churn is higher than expected, the timeline extends and ROI may not occur. The plan is to reach break-even within the available runway and then grow into enterprise licensing for accelerated profitability and valuation.

5. What would be the distribution of ROI and the method?

Distributions occur via (a) profit distributions and/or (b) exit proceeds. Profit distributions are based on net profit attributable to Brainpower AI, calculated after operational costs, AI infrastructure costs, marketing spend, staffing, taxes (if applicable), and agreed reserves for reinvestment. Exit proceeds are, distributed proportionally if Brainpower AI is acquired, licensed, or bought out. The agreement will define the distribution schedule (for example, quarterly or annual) and the calculation basis to avoid ambiguity.

6. What is the exact legal structure of the investment?

The investment is, made into Smarthinkerz LLC under a written agreement that defines each investor's economic participation tied specifically to Brainpower AI. This avoids legal ambiguity because a product brand alone cannot typically hold equity unless it is a separate legal entity. The agreement defines the "Brainpower AI participation pool," the investor's share within that pool, and the events that trigger distributions (profits or exit events). It also defines investor reporting and transfer restrictions.

7. What traction do you currently have (paying users, revenue, active users, contracts)?

As of 28 Feb 2026, Brainpower AI is pre-launch with zero current revenue and users. There are 0 paying users and OMR 0 in MRR, with a prototype in internal testing and early feedback from 5 founder-network testers on workflow usability. The enterprise pipeline currently consists of informal discussions with 2 potential GCC partners for post-launch pilots. Full details are provided in the "Current Traction" section of the Topline Investor Brief.

8. What are the exact financial projections for the next 12, 24, and 36 months?

Projections start from zero traction, assuming post-raise launch in Months 1-3. They are Projections start from zero traction, assuming launch in Months 1–3 after this round. They are assumption-driven: subscriber growth of 15–35% month-on-month post-launch, ARPU of OMR 15–25, churn of 10–18%, and 1–4 new enterprise contracts per quarter from Month 6 onward with ACV of OMR 3,000–8,000. Marketing scales only if early metrics, such as 50%+ beta retention, are achieved.

The detailed conservative, base, and upside projections for 12, 24, and 36 months—covering MRR, enterprise contracts, total revenue, key assumptions, and sensitivities—are summarized in the projections table in this section. Sensitivity factors include: no product-market fit in the first 6 months (delaying milestones by 9–12 months), a 15% increase in CAC (extending runway needs by roughly 20%), and inflation scenarios with 5–10% annual cost increases (potentially delaying break-even by 3–6 months). An edge case 2026 recession variant assumes a 10% revenue haircut and a 12-month delay to exit timing. These projections are illustrative; actual results may differ materially and there can be no certainty of outcomes.

Timeline	Scenario	MRR (OMR)	Enterprise Contracts	Total Revenue (OMR)	Key Assumptions	Risks/Sensitivity
12 Months	Conservative	1,000	1	12,000	15% MoM growth from launch, 18% churn, CAC OMR 70	Validation delays; +10% churn adds 4 months to break-even. Inflation variant: +5% costs delays by 2 months extra.
	Base	2,000	2	24,000	20% MoM, 15% churn, CAC OMR 60	Tied to app ASO and authority content; aligns with Brief's Year 1 target.
	Upside	4,000	3	48,000	25% MoM, 12% churn,	Faster beta feedback; recession edge: If

Timeline	Scenario	MRR (OMR)	Enterprise Contracts	Total Revenue (OMR)	Key Assumptions	Risks/Sensitivity
					CAC OMR 50	enterprise slows, cap at 2 contracts.
24 Months	Conservative	5,000	3	60,000	Steady post-validation, 16% churn.	Enterprise cycles >9 months; cost inflation (+10%) reduces margins by 15%.
	Base	10,000	5	120,000	Break-even at ~OMR 8K MRR; matches Brief's Year 2.	Assumes 70% pilot conversion; downturn mitigation: Pivot to B2B for stability.
	Upside	20,000	8	240,000	Referrals accelerate.	If CAC drops to OMR 40; edge: Global recession caps growth at 15% MoM.
36 Months	Conservative	20,000	8	240,000	Renewal focus, 14% churn.	Market competition; +10% inflation delays profitability by 6 months.
	Base	50,000	15	600,000	5-8x ARR multiples potential; aligns with Brief's Year 3.	Global rollout success; recession variant: Adjust ACV down 20% for conservatism.
	Upside	100,000	25	1,200,000	Viral enterprise adoption.	If ACV hits OMR 10K; edge: Economic downturn extends timeline to 42 months.

9. What is the total amount you are raising in this round?

The total raise is OMR 75,000. This is, sized to provide enough runway to complete product maturation, run global marketing experiments, and convert early enterprise pilots—without over-raising before the growth engine is validated.

10. How much runway will this funding provide?

Runway depends on the monthly burn rate and how aggressively marketing is scaled. A disciplined burn model targets 12–18 months of runway with room for controlled marketing tests. The company will commit to milestone-based spending: marketing spend increases only when conversion, retention, and CAC meet predefined thresholds. This approach protects investor capital and improves the odds of reaching break-even before runway ends.

11. How exactly will the invested funds be allocated ?

Funds are allocated across specific categories: development, marketing, infrastructure, operations, legal, and reserve. The goal is balanced execution—building a strong product while also creating global demand and enterprise traction. A category-based allocation summary is investor-friendly and avoids micromanagement. The allocation will be reviewed quarterly against progress and adjusted based on measurable performance.

12. What is the competitive landscape (ChatGPT, Copilot, Gemini, Claude) and why will Brainpower AI win?

ChatGPT, Copilot, Gemini, and Claude are horizontal AI platforms designed for broad tasks. They provide general capability but are not purpose-built decision systems. Brainpower AI is positioned as a vertical Decision Intelligence application layer that embeds structured frameworks, guided workflows, and repeatable decision outputs. We do not compete at the model layer; we leverage foundational models as infrastructure and compete at the workflow and outcome layer. This mirrors how vertical SaaS wins alongside platform giants by becoming embedded in specialized workflows where users pay for structure, repeatability, and results

13. What is the exit strategy?

There are multiple exit paths: (1) acquisition by a productivity SaaS company, AI workflow company, or enterprise software vendor; (2) licensing buyout of Brainpower AI's workflows and platform; (3) a growth round with partial liquidity for early investors; or (4) long-term profitability with ongoing distributions. The strategy focuses on achieving enterprise contracts and strong retention, as these drive acquisition interest and valuation multiples.

14. What happens if additional funding is required later?

If additional funding is needed, it will be milestone-triggered and transparent. Investors may be offered pro-rata rights to maintain their participation share. If investors choose not to participate, dilution of their relative share may occur depending on the structure. The agreement will define future-round principles clearly: when funding may be sought, what milestones trigger it, and how existing investors are treated.

15. Who owns the intellectual property (IP) of Brainpower AI?

All IP will be owned by Smarthinkerz LLC. This includes source code, branding, design assets, decision frameworks, and proprietary workflow logic. All contractors and developers will sign IP assignment and confidentiality agreements to prevent future disputes. Clean IP ownership is essential for enterprise trust and any future acquisition.

16. What are the biggest risks right now, and how are you mitigating them?

Key risks include customer acquisition cost inflation, churn, enterprise sales cycle length, and AI infrastructure cost volatility. Mitigation plans include authority-led acquisition to reduce CAC, onboarding and habit-building loops to reduce churn, a pilot-to-contract enterprise approach to manage sales cycle risk, and multi-provider routing plus usage governance to manage AI cost volatility. We explicitly track these risks and communicate them to investors along with mitigation progress.

17. What governance rights do investors have?

Investors typically receive information rights and defined economic rights. Governance will remain streamlined so execution stays fast. If needed, a small set of "reserved matters" can be defined (for example, issuing more than the capped pool, selling major assets, or changing distribution rules). Day-to-day approvals will remain with management, since heavy governance at this stage slows execution and harms growth.

18. What happens if the founder leaves or becomes unable to operate?

A continuity plan will be in place: documented processes, secured admin credentials, infrastructure access controls, and a clear method for appointing an interim operator. For added investor confidence, the agreement can require the company to designate a backup manager and keep key operations documented. This reduces "key person" risk and supports long-term resilience.

19. Are there any existing liabilities, debts, or legal exposures?

As of the date of this memorandum, the company has no material financial liabilities, debts, or legal disputes beyond normal operating obligations; any changes will be disclosed to investors. The company will also disclose any future debts, obligations, pending disputes, or major contractual liabilities in the investment agreement, and describe them clearly with associated terms. Investor trust depends on full transparency in this area.

20. What is the realistic downside scenario?

In the downside scenario, the product fails to achieve scalable acquisition or retention, enterprise deals take longer than expected, and revenue does not cover costs. In that case, the company may reduce burn, pivot focus toward higher-value B2B licensing, or ultimately wind down. Investors must understand that capital can be lost in early-stage investments; this is a normal risk profile and is stated clearly and professionally in the agreements.

21. How is the 20% Brainpower AI allocation legally structured and documented?

The agreement defines a fixed "Brainpower AI participation pool" capped at 20%. Each investor's share in this pool is recorded and cannot be diluted beyond defined rules without formal approvals. The agreement also defines what counts as Brainpower AI revenue and profit, and how distributions are calculated and scheduled.

In simple terms, investors own a fixed percentage of a capped 20% pool that tracks Brainpower AI's distributable profit and exit value.

22. How will Brainpower AI revenue be separated from other Smarthinkerz LLC activities?

Brainpower AI will be tracked as a separate internal cost center with distinct revenue lines, AI/API costs, marketing spend, and staffing allocations. This enables investors to see Brainpower AI unit economics clearly, without confusion from other company operations. It also supports enterprise reporting and any potential future spin-out or acquisition of the Brainpower AI business.

23. Are there audited or independently reviewed financial statements?

At the early stage, the company will rely on management reporting. As enterprise revenue grows and crosses agreed thresholds, an external accountant review can be introduced annually. Investors will receive quarterly management statements plus an annual external review once it becomes cost-effective, which improves credibility without adding heavy overhead too early.

24. What is the current monthly burn rate?

At this pre-launch stage, monthly burn is kept lean and focused on product development and essential operations. As the round is deployed, burn will be reported transparently and broken down by category (development, infrastructure, marketing, operations). Burn will be managed through milestone-based spend rules, with increases in marketing and hiring only when performance metrics justify it. Investors will see burn evolution in the quarterly reports and KPI dashboard.

25. What is the projected monthly break-even point?

Break-even is defined as the point where recurring revenue covers both fixed operating costs and variable AI inference costs. In our base case, this occurs at an MRR level in the OMR 8,000–10,000 range, assuming target gross margins and a lean fixed cost structure. A professional answer will remain: "We reach break-even at

X MRR given Y gross margin assumptions and Z fixed cost structure," and these inputs will be updated as real data comes in.

26. What are the current customer acquisition costs (CAC)?

CAC will be measured per channel (organic, paid search, paid social, partnerships) once acquisition campaigns begin. Early-stage CAC can be noisy; what matters is the system to drive it down. Our plan focuses on authority marketing (content, LinkedIn, YouTube), referrals, and product-led growth loops to reduce CAC over time, with a target CAC payback period of under 9 months. CAC by channel and payback period will be tracked on the KPI dashboard.

27. What is the lifetime value (LTV) of a customer?

LTV is driven by ARPU and retention. For subscriptions, LTV improves with strong habit loops, reusable templates, and "workflow embedding" that makes Brainpower AI part of the user's regular decision process. For enterprise customers, LTV improves through renewals and seat expansion across teams. LTV assumptions will be tied to actual retention cohorts and renewal behavior rather than optimistic guesses, and reported alongside CAC to monitor unit economics.

28. What is the churn rate?

Churn will be measured monthly and tracked by cohort once users are live on the platform. The product roadmap includes churn-reduction mechanisms such as better onboarding, more frequent "aha" moments, decision vault usage, weekly review prompts, and enterprise features that lock in workflows. Investors will see churn data and cohort curves on the KPI dashboard, and we will treat churn as a core product metric—not just a finance metric.

29. What is the product roadmap for the next 12 months?

The 12-month roadmap focuses on three themes: launch readiness, retention, and enterprise readiness. Key items include: app store launch polish (iOS/Android listings, demo video, ASO), expansion of decision frameworks and templates, improved onboarding and in-app guidance, team workspaces and shared decision vaults for enterprises, reporting dashboards, and multi-provider AI routing. Each roadmap item is tied to a specific objective: improving retention, improving acquisition, or unlocking enterprise revenue.

30. What technical risks exist in the current architecture?

Technical risks include scaling reliability, AI latency and cost, data security, and consistency across mobile and web platforms. Mitigation measures include load testing and monitoring, caching and usage governance to manage AI calls, secure coding practices, and an architecture designed for incremental scale-out rather than big-bang rewrites. As the platform grows, we will harden the stack further based on enterprise requirements and security reviews.

31. Is the AI proprietary or built on third-party APIs?

Brainpower AI uses third-party foundational model APIs for language intelligence, while our proprietary value lies in workflows, decision frameworks, user experience, orchestration logic, and enterprise packaging. This reduces R&D risk and accelerates time-to-market, while still allowing defensibility at the workflow and customer relationship layer rather than the model layer.

32. If third-party APIs are used, what is the dependency risk?

Dependency risks include pricing changes, rate limits, performance fluctuations, and policy changes from model providers. We mitigate these by using multi-provider routing, designing model-agnostic workflows, implementing fallback models, and applying usage controls to keep costs within defined thresholds. The goal is to avoid being "single-vendor trapped" and to keep flexibility as the AI infrastructure landscape evolves.

33. Do you own all source code and IP assignments from developers?

All developers and contractors will sign IP assignment, confidentiality, and delivery clauses so that all work product is owned by Smarthinkerz LLC. This prevents future disputes and is essential for enterprise adoption and acquisition readiness. We treat IP hygiene as non-negotiable and will maintain appropriate documentation to demonstrate clean ownership.

34. What is the cybersecurity posture of the platform?

Security measures include secure authentication, encrypted storage where applicable, access control and least-privilege policies, and regular monitoring of the infrastructure. As enterprise deals grow, we will increase security hardening with admin logging, policy controls, and contract-driven compliance requirements. Security posture will evolve alongside customer needs and will be reflected in technical documentation for enterprise buyers.

35. Are there regulatory risks related to AI compliance or data protection laws?

Yes, especially when operating across multiple regions with different privacy and AI regulations. We will mitigate these risks by minimizing sensitive data storage, using clear and transparent user consent terms, implementing safe data-handling policies, and adjusting compliance posture (for example, data residency and retention controls) for enterprise customers as needed. Regulatory changes will be monitored and incorporated into product and contract updates over time.

36. What geographic markets are you targeting first?

The first targets are high-income, English-speaking professional markets and regions with strong AI adoption. Initial focus will be on the GCC, US/UK, and selected hubs such as Singapore and Australia, followed by broader expansion through app distribution and partnerships as retention and unit economics support it.

37. What is your pricing strategy and why?

Pricing will be tiered to support both global acquisition and enterprise upsell. Individual and small-team subscription tiers drive user volume and habit formation, while enterprise licensing packages drive higher-value recurring revenue. Pricing will be tested and optimized through experiments and will be anchored to clear, measurable value: time saved, decision clarity, and improved workflow outcomes.

38. What proof do you have of product-market fit?

Product-market fit is ultimately demonstrated through retention, repeat usage, willingness to pay, referrals, and enterprise pilot conversions. In the early stages, proof will be partial: conversion from trials to paid plans, weekly and monthly active usage, and successful pilots with clear outcome metrics. Over time, enterprise renewals and expansion across teams will become the strongest indicators of product-market fit.

39. What milestones will trigger the next funding round?

Milestones for a potential next round include reaching target MRR levels, stabilizing churn, achieving predictable CAC and payback periods, and converting a defined number of enterprise pilots into annual contracts. Funding will be tied to hitting these measurable outcomes rather than to arbitrary dates, so that valuation and terms are grounded in demonstrated performance.

40. At what valuation do you expect the next round to occur?

Rather than fixing a valuation too early, the next round's valuation will be tied to measurable metrics such as ARR, growth rate, retention, and enterprise contract count. We aim to reach a level of traction where standard early-stage SaaS multiples (for example, 5–10× ARR depending on quality of revenue) can be justified by data, not just narrative.

41. What dilution should current investors expect in future rounds?

Future rounds may dilute earlier investors depending on size and structure. Offering pro-rata rights allows current investors to maintain their participation share if they choose. Our principle is to be transparent about potential dilution, tie new capital raises to clear milestones, and ensure that any dilution is offset by value creation at the Brainpower AI level.

42. What is your personal financial commitment to Brainpower AI?

The founder has significant "skin in the game." To date, I have committed over 18 months of focused work to Brainpower AI and invested approximately OMR 10,000 of personal capital into product development, infrastructure, and initial market validation activities. Going forward, I am fully dedicated to Brainpower AI as my primary focus, and my personal reputation and future upside are directly tied to the platform's long-term success.

43. What is the worst-case scenario, and what does downside protection look like?

In the worst-case scenario, the product does not scale, enterprise adoption is slower than expected, and investor capital may be lost. Downside protection comes from disciplined burn management, the ability to pivot toward higher-value enterprise licensing, and transparent reporting so investors can see risks early. Optional buyback mechanisms could be considered later only if the business generates sufficient cash flow; no buybacks will be promised that cannot be realistically financed.

44. Why invest now versus waiting 6–12 months?

Investing now secures early participation in the capped 20% pool before scaling milestones increase pricing and terms. This round funds the next execution phase: global app distribution, retention improvements, and conversion of enterprise pilots. Waiting may reduce risk but also reduces upside and access, as later rounds are likely to be raised at higher valuations if targets are met.

45. Why are you confident you are the right founder to execute this at scale?

The strategy is clear and realistic: we are not trying to "beat ChatGPT" at the model layer. We are building a vertical Decision Intelligence system that leverages foundational models, focuses on workflow embedding, and scales via global app distribution plus enterprise licensing. The plan is milestone-driven and measured (CAC, churn, retention, enterprise conversion), which is how scalable companies are built, and the founding team's background aligns with executing this plan.