AI MARGIN

Board Report

Feb 15-Feb 21, 2026 | Peter Zsobrák's Organization

YOUR AI FEATURES ARE 32pp BELOW HEALTHY MARGINS

Current: 33.4% | Benchmark: 65% | Gap: 31.6pp

Based on public SaaS gross margin data (typical AI-heavy SaaS range: 60-70%)

Unrealized margin: ~$6.1K/mo

At 65% margin on current $19.3K revenue, profit would be ~$12.5K/mo vs current $6.5K/mo.

AI SPEND	SAVINGS FOUND	NET MARGIN	ANNUAL ROI
$12.8K /mo	**~$7.2K** /mo	**33.4%**	**0-14x**
$154,200/yr	~$85,953/yr	Target: 65%	on $499/mo

SAVINGS BREAKDOWN

Guaranteed (no tradeoff):	**~$300/mo**	*AI Image Generation*
Requires validation:	**~$7.2K/mo**	*Code Assistant*

Conservative total: ~$2.4K/mo | Full potential: ~$7.2K/mo

(Guaranteed ~$300 + 30% of validation savings ~$2.1K) Full potential assumes successful validation and no quality degradation.

KEY FINDING

Code Assistant + Document Summarizer generate just $701/mo profit on $7.1K revenue (9.9% margin). Improving these features to 50% margin would add ~$2.8K/mo to your bottom line.

EXECUTIVE SUMMARY

AI Margin identified ~$2.4K-~$7.2K/mo (~$29,386-~$85,953/yr) in optimization potential. 4 of 5 features are profitable. 2 are healthy (>40% margin), 2 are critical (<15%), 1 is losing money.

If no action taken: Annual AI gross profit shortfall vs benchmark: ~$73,140.

AI Profitability Score: **63/100** (Moderate)

TOP RISK	AI Image Generation (-20.0% margin, ~$300/mo loss)
TOP OPP	Code Assistant margin is thin at 10%

Feature P&L

Feature	Revenue	AI Cost	Profit/mo	Margin	Status
AI Image Generation	$1.5K	$1.8K	-$300	-20.0%	Losing
Document Summarizer	$3.1K	$2.8K	$298	9.6%	Critical
Code Assistant	$4.0K	$3.6K	$403	10.1%	Critical
AI Chatbot	$8.5K	$4.2K	$4.3K	50.6%	Healthy
Semantic Search	$2.2K	$450	$1.8K	79.5%	Healthy

Of the table content, 5 unique features tracked. All spend is allocated to features.

PROVIDER BREAKDOWN

Provider	Cost/mo	Share	Requests	Models
OpenAI	$9.1K	70.6%	110.2K	5
Anthropic	$3.8K	29.4%	25.5K	1

MODEL DETAIL

Provider	Model	Cost/mo	Share	Requests	Tier
OpenAI	gpt-4o	$5.8K	44.7%	35.0K	Frontier
Anthropic	claude-3-5-sonnet-20241022	$3.8K	29.4%	25.5K	Frontier
OpenAI	dall-e-3	$1.8K	14.0%	1.2K	Frontier
OpenAI	gpt-4o-mini	$1.1K	8.3%	24.0K	Standard
OpenAI	text-embedding-3-large	$405	3.2%	45.0K	Standard
OpenAI	text-embedding-3-small	$45	0.4%	5.0K	Standard

Margin Spread & Risk Exposure

MARGIN DISTRIBUTION

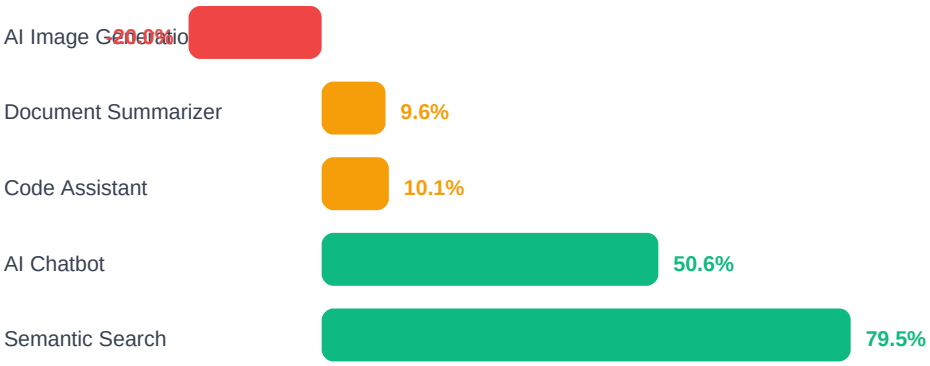

RISK EXPOSURE

Frontier Model Dependency — 88.2%

88.2% of spend (by cost) on frontier models (GPT-4o, Claude Sonnet, etc.). Price cuts help but lock-in risk remains.

Revenue Coverage — 1.5x

At 1.5x coverage, a 20% cost increase would reduce margin from 33.4% to 20.1%. Target: >2x coverage for resilience.

Single Provider Concentration — 70.6% OpenAI

70.6% of spend with OpenAI. Consider multi-provider strategy to reduce outage and pricing risk.

Negative Margin Features — 1 of 5

1 feature currently losing money. Immediate action required to stop cash burn.

SCENARIO SENSITIVITY

Scenario	Current	Projected	Impact
Provider raises prices 20%	33.4%	20.1%	-13.3pp
Implement low-risk swaps (none identified)	33.4%	33.4%	+0.0pp
Implement all recommendations*	33.4%	70.5%	+37.1pp
Reprice to 65% target margin	33.4%	65.0%	+31.6pp

** Assumes successful execution and no revenue impact.*

Smart Recommendations

IMMEDIATE ACTIONS (0-30 days)

CRITICAL **AI Image Generation is losing money!** Save ~$300/mo

'AI Image Generation' costs $1,800/mo but generates only $1,500/mo. You're losing $300/mo on this feature.

Action: Either reduce costs (switch models, reduce usage) or increase pricing to cover the $1,800 cost.

HIGH **Code Assistant margin is thin at 10%** Save ~$1.6K/mo

'Code Assistant' has a 10% margin ($4,000 revenue - $3,596 cost). A healthy SaaS AI margin is typically >50%.

Action: Optimize model costs or adjust pricing. Even a 40% improvement would add $1,597/mo.

HIGH **Document Summarizer margin is thin at 10%** Save ~$1.3K/mo

'Document Summarizer' has a 10% margin ($3,100 revenue - $2,802 cost). A healthy SaaS AI margin is typically >50%.

Action: Optimize model costs or adjust pricing. Even a 40% improvement would add $1,252/mo.

OPTIMIZATION OPPORTUNITIES (30-90 days)

MEDIUM **Document Summarizer uses gpt-4o for summarization** Save ~$2.0K/mo

'Document Summarizer' (summarization) uses the frontier model gpt-4o, costing $2,802/mo. Simpler tasks like summarization often work well with cheaper models.

Action: Switch to gpt-4.1. Estimated savings: $1,997/mo (-89%). (Test in staging before switching.)

MEDIUM **High frontier model concentration** Save ~$2.0K/mo

74% of your AI spend ($9,530/mo) is on frontier models. If even 30% of these tasks could use mid-tier models, you'd save ~$2,001/mo.

Action: Review features using frontier models and consider GPT-4o-mini or Claude Haiku for simpler tasks.

LEAKAGE ALERTS

No cost leakage detected. All feature costs are within expected ranges.

Action Plan

Prioritized actions with estimated impact. Each card shows current state, recommended change, and projected savings.

#1 AI Image Generation is losing money!

Save ~$300/mo

This week | Effort: Low | Payback: Immediate

'AI Image Generation' costs $1,800/mo but generates only $1,500/mo. You're losing $300/mo on this feature.

Action: Either reduce costs (switch models, reduce usage) or increase pricing to cover the $1,800 cost.

#2 Code Assistant margin is thin at 10%

Save ~$1.6K/mo

0-30 days | Effort: Medium | Payback: 1-2 weeks

'Code Assistant' has a 10% margin ($4,000 revenue - $3,596 cost). A healthy SaaS AI margin is typically >50%.

Action: Optimize model costs or adjust pricing. Even a 40% improvement would add $1,597/mo.

#3 Document Summarizer margin is thin at 10%

Save ~$1.3K/mo

0-30 days | Effort: Medium | Payback: 1-2 weeks

'Document Summarizer' has a 10% margin ($3,100 revenue - $2,802 cost). A healthy SaaS AI margin is typically >50%.

Action: Optimize model costs or adjust pricing. Even a 40% improvement would add $1,252/mo.

#4 Document Summarizer uses gpt-4o for summarization

Save ~$2.0K/mo

30-60 days | Effort: High | Payback: 30 days

'Document Summarizer' (summarization) uses the frontier model gpt-4o, costing $2,802/mo. Simpler tasks like summarization often work well with cheaper models.

Action: Switch to gpt-4.1. Estimated savings: $1,997/mo (-89%).

#5 High frontier model concentration

Save ~$2.0K/mo

30-60 days | Effort: High | Payback: 30 days

74% of your AI spend ($9,530/mo) is on frontier models. If even 30% of these tasks could use mid-tier models, you'd save ~$2,001/mo.

Action: Review features using frontier models and consider GPT-4o-mini or Claude Haiku for simpler tasks.

TOTAL POTENTIAL IMPACT **~$7.2K/mo (~$85,953/yr)**

Score Breakdown

63

/100

AI Profitability Score: Moderate

Some features need attention. Review margins and cost trends.

COMPONENT SCORES

Margin Health

Weight: 40%

45

Weighted average margin health across 5 features

Loss Ratio

Weight: 25%

80

1 of 5 features losing money

Leakage Trend

Weight: 20%

50

No trend data yet

Model Efficiency

Weight: 15%

100

100% of cost allocated to features

METHODOLOGY

Margin Health (40%)

Weighted average of feature margins, penalizing negative-margin features more heavily.

Loss Ratio (25%)

Proportion of features with negative or near-zero margins relative to total feature count.

Leakage Trend (20%)

Month-over-month cost trend analysis. Rising costs without revenue growth reduce the score.

Model Efficiency (15%)

Ratio of frontier vs. mid-tier model usage. Over-reliance on expensive models reduces efficiency.

This report is generated by AI Margin based on your connected provider data. Savings estimates are projections based on current usage patterns and model pricing. Actual savings may vary. Model quality assessments are heuristic-based. Always test model changes in staging before production deployment.