Blue Sheen

AI ROI report

Personalized ROI analysis for Acme Mid-Market Co.

May 23, 2026

Prepared for Acme Mid-Market Co.

Executive summary

Acme Mid-Market Co. submitted the following inputs to the on-page AI ROI calculator:

- Team size affected: 312
- Average fully-loaded hourly cost: $85
- Estimated hours saved per person per week: 4
- AI tool cost per person per month: $40
- One-time implementation budget: $150,000
- Risk discount applied: 25%
- Time horizon: 24 months

Conservative case (75% of estimated hours, 24 month horizon): $4.13M total benefit; $299K total cost; **1,279% ROI**; payback in 1.2 months.

Expected case (calculator inputs as given): $5.51M total benefit; $299K total cost; 1,743% ROI.

Aggressive case (1.25x inputs): $6.89M total benefit; same cost; 2,205% ROI.

These numbers look implausibly high — that is normal for a knowledge-work AI deployment with these inputs. The likely-real ROI is lower because:

1. Adoption is rarely 100% in year one (we model ramp below)
2. Hours-saved estimates are typically overstated; we recommend tracking measured hours per Layer 1 of the McKinsey AI ROI methodology
3. There are change management costs not captured in the $150K implementation budget

The realistic 24-month picture, adjusted for ramp + measurement reality, is **312-415% ROI**. Still excellent. Still defensible to a CFO.

Methodology

We use the three-layer McKinsey + BCG framework adapted for mid-market healthcare SaaS:

- **Layer 1 — Utilization**: weekly tracking of active users + prompts per day. Direct measure.
- **Layer 2 — Productivity**: hours saved per person per week. Self-reported initially; calibrated against measured outputs over time.
- **Layer 3 — Business outcomes**: revenue impact, customer satisfaction, churn — tracked monthly, attributable with multi-month lag.

Your calculator inputs primarily address Layer 2. This report extends to Layers 1 and 3.

Sensitivity analysis

The single most-sensitive variable is hours-saved-per-person-per-week. A ±30% variance changes ROI dramatically:

Hours saved/wk	Annual benefit	First-year net	Payback (months)
3 (conservative)	$1.45M	$0.95M	3.0
4 (your estimate)	$1.93M	$1.43M	2.3
5 (aggressive)	$2.42M	$1.92M	1.9

The second-most-sensitive variable is adoption ramp. If full adoption takes 12 months instead of 3:

Months to full adoption	Year-1 net	Year-2 net	24-mo ROI
3 (your estimate)	$1.43M	$1.93M	1,022%
6 (realistic)	$0.96M	$1.93M	868%
12 (conservative)	$0.48M	$1.93M	705%

We recommend planning to the 6-month-adoption case.

Industry benchmarks

Per BCG 2026 enterprise AI report, healthcare SaaS companies measuring rigorously see:

Metric	Sector median	Top quartile
Productivity gain on knowledge work	18%	41%
Time to measurable ROI	9 months	4 months
% of pilots reaching production	21%	67%
Year-1 cost overrun	35%	9%

Your team's profile (no formal AI strategy, ChatGPT Team + Copilot already deployed, internal champion present, mature data platform) maps to the second quartile. With Blue Sheen's roadmap, top-quartile is realistic in 18 months.

Vendor recommendations

Based on your stack (Python/Django + React + Postgres + Snowflake + AWS) and your use case profile:

Layer	Build / Buy / Partner	Recommended vendor
LLM foundation	Buy + use existing	Anthropic Claude Team (procurement Q1) + retain GitHub Copilot Enterprise
CS draft assist	Buy	Forethought or Ada (RFP both)
SDR research	Buy	Clay (or Apollo + Claude API combo)
Data layer	Build on existing	dbt + Snowflake Cortex
Vector DB (RAG)	Buy managed	Pinecone or Weaviate Cloud
Observability	Buy	Langfuse for LLM observability — integrate with existing Datadog

Total tooling cost year 1: estimated $145K (in addition to your existing ChatGPT + Copilot spend).

Implementation roadmap (90 days)

Days 1-30: AI Governance Committee formed; AI Acceptable Use Policy adopted (see separate Blue Sheen deliverable). Anthropic procurement begins. One ML/AI engineer hired or reassigned.

Days 31-60: Use case #1 (CS draft assistant) piloted with 8 CS agents. Use case #3 (SDR research) piloted with 4 SDRs. Use case #4 (PR review with Copilot) measured baseline established.

Days 61-90: Pilot retro + decision to scale. KPI dashboard live. First quarterly business review of AI portfolio.

KPI dashboard template

Weekly auto-populated from your data warehouse:

Layer	Metric	Source	Target Q1
1 Utilization	Weekly active users	LLM platform logs	60% of seats
1	Prompts per active user / day	LLM platform logs	12
2 Productivity	Hours saved (self-reported)	Weekly pulse survey	4
2	Tickets resolved per CS rep	Zendesk	+20% vs baseline
2	PR review cycle time	GitHub	-25% vs baseline
3 Business	Customer first-response time	Zendesk	-30%
3	New qualified meetings / SDR / week	Salesforce	+50%
3	NPS (CS team)	Quarterly survey	+5 points

Risks and mitigations

Risk	Likelihood	Mitigation
Adoption ramp slower than 3 months	High	Build the 6-month case; assign internal champion per team
PHI leakage incident	Medium	AI policy + training (already in motion)
Vendor lock-in	Medium	Multi-model architecture; prompt portability audit quarterly
Cost overrun (>20%)	Medium	Monthly spend review tied to budget

Risk	Likelihood	Mitigation
AI tool sprawl	High	AI Governance Committee owns the approved list
Tool changes hands or pricing	Low	Quarterly vendor review

Conclusion

The math says this works. The execution risk is on adoption, not on technology or pricing. The single highest-leverage move is establishing the AI Governance Committee and the policy in week 1 — that compresses every downstream timeline.

This sample ROI report was hand-crafted by Blue Sheen for a fictional client. Your real report will be tuned to your industry, your data maturity, your existing AI investments, and the specific assumptions you want stress-tested. Request your custom report at bluesheen.com/tools/ai-roi-calculator/.