

Companion Notes

AI Opportunity Evaluation Pack — Source Tracing

Table of contents

Part 1: Facts Sourced from the Cloudcore Website	1
Executive Positions and Personalities	1
Financial Figures	3
Staffing Numbers	3
The Six AI Opportunities	4
Executive Tensions	4
Cross-References	4
Part 2: Assumptions and Invented Details	4
Revenue Split by Client Sector	4
Proposed AI Budget of \$250,000	5
Tier 1 Support Staff Breakdown (5 of 8)	5
All Quotable Statements	5
Preferred Initiative Assignments	5
Preliminary Opportunity Assessments	6
Specific Details Within Assessments	6

This companion document traces every detail in the AI Opportunity Evaluation Pack to either a specific location on the Cloudcore website or flags it as an assumption invented for the brief.

Part 1: Facts Sourced from the Cloudcore Website

Executive Positions and Personalities

All five executive summaries are grounded in their backstory files. The following details come directly from the site:

Marcell Ziemann (CEO):

- Founded the company; only founder still in executive role — [chatbots/_backstories/marcell_ziemann_c](#)

- Cites 70-85% AI project failure rate — same file
- Will not commit capital without evidence-based plan — same file
- Board wants AI positioning and evidence-based growth strategy — same file
- Lost enterprise deals to competitors with better AI messaging — same file
- Customer retention was 94% before breach — same file

Mark Gonzalez (CTO):

- Joined 3 years ago from larger cloud provider — `chatbots/_backstories/mark_gonzalez_cto.md`
- AI readiness scores: Infrastructure 3-4/5, Data 2/5, Talent 1/5, Governance 1/5 — same file
- “I’d rather do one AI thing well than five things poorly” (paraphrased from backstory) — same file
- 6-12 months to get something meaningful running — same file
- Zero data scientists or ML engineers — same file
- Monitoring stack: Prometheus, Grafana, PagerDuty — same file
- Infrastructure team: 12 engineers — same file

Sarah Thompson (COO):

- ITIL Master certified; joined 2 years ago — `chatbots/_backstories/sarah_thompson_coo.md`
- Service metrics: 4.2hr resolution (target 4hr), 82% satisfaction (target 85%), 68% FCR, 8% churn — same file
- Support team of 8 people — same file
- Concerned about staff anxiety re: AI replacing roles — same file
- Wants AI to augment not replace; gradual rollout — same file

Aisha Rahman (CFO):

- Joined 5 years ago; Big Four tech audit background — `chatbots/_backstories/aisha_rahman_cfo.md`
- Operating margin ~15% — same file
- Board wants profitability in 2 years — same file
- “Competitors are doing AI is not an ROI calculation” (paraphrased) — same file
- Productive tension with CTO on spending — same file
- Post-breach, supports CISO budget requests more readily — same file

Sophia Martines (CISO):

- 4 years at CloudCore; security team of 8 — `chatbots/_backstories/sophia_martines_ciso.md`
- Needs 3+ more people — same file
- 40+ vendor integrations — same file
- No AI governance framework — same file
- Board now requires monthly security updates (was quarterly) — same file
- ISO 27001 took nearly 2 years — same file

Financial Figures

All financial figures are sourced from backstory files:

Figure	Value	Source
Annual revenue	~\$45M	marcell_ziemann_ceo.md, aisha_rahman_cfo.md
YoY growth	~25%	Same files
Operating margin	~15%	aisha_rahman_cfo.md
Series A+B funding	~\$20M	aisha_rahman_cfo.md
Breach total cost (year 1)	~\$3.5M	data_breach_overview.md
Breach direct costs	~\$2.1M	Same file
Insurance coverage	~\$800K	Same file
Lost client ARR	~\$1.5M	Same file
Clients lost to breach	23	sophia_martines_ciso.md
IT/Technology budget share	~40%	aisha_rahman_cfo.md
Security share of IT	~12% (up from 8%)	Same file
Sales/Marketing budget share	~25%	Same file
No dedicated AI budget	Confirmed	Same file
ML engineer market rate	\$180-250K AUD	karen_lee_hr_manager.md
NPS pre-breach	45	lisa_chen_cmo.md
NPS post-breach	28	Same file
Customer acquisition cost	~\$2,400	Same file

Staffing Numbers

Team	Count	Source
Total employees	47	marcell_ziemann_ceo.md, cloudcore_company_overview.md
Perth / Sydney split	35 / 12	karen_lee_hr_manager.md
Infrastructure engineering	12	mark_gonzalez_cto.md
Customer support	8	sarah_thompson_coo.md
Security	8	sophia_martines_ciso.md
Software development	6 + 1 lead	michael_thompson_lead_software_developer.md
Data and analytics	2	jamal_al_sayed_data_analyst.md
Marketing	4	lisa_chen_cmo.md
Compliance	2	emily_chen_head_of_compliance.md

The Six AI Opportunities

The six opportunities and their strategic priority / complexity ratings come directly from `briefs/cloudcore-introduction.qmd`.

Executive Tensions

The following tensions are established in backstory files:

- CEO vs CISO (speed vs governance) — `marcell_ziemann_ceo.md`, `sophia_martines_ciso.md`
- CTO vs CFO (build vs buy) — `mark_gonzalez_cto.md`, `aisha_rahman_cfo.md`
- CEO/CTO vs COO (efficiency vs staff) — `sarah_thompson_coo.md`

Cross-References

All website URLs reference real pages on the Cloudcore site.

Part 2: Assumptions and Invented Details

The following details were invented for the brief. They are plausible and do not contradict existing content, but they are not sourced from the website.

Revenue Split by Client Sector

Sector	Share	Basis for Invention
Professional Services	30%	Backstories mention SMEs as primary market but no percentage is given anywhere
Healthcare	25%	Backstories emphasise healthcare as important and compliance-heavy, but no revenue figure exists
Finance	20%	Same; mentioned as key sector but no percentage
Education	15%	Mentioned as a sector; described as “price-sensitive” in marketing backstory
Other	10%	Catch-all for retail, manufacturing, government mentioned in passing

Note: The sales data CSV (`data/cloudcore-sales-data.csv`) contains industry breakdowns at the transaction level (Education, Manufacturing, Healthcare, Retail, Finance, Technology) but these are for individual support tickets and customer records, not company-level revenue attribution. The percentages above are invented.

Proposed AI Budget of \$250,000

The handoff document suggested a \$200-300K range. The \$250K figure was chosen as a midpoint. No specific budget figure appears anywhere in the repo. The backstories confirm there is no dedicated AI budget, and the CFO’s position is that spending requires a business case.

Tier 1 Support Staff Breakdown (5 of 8)

The repo confirms 8 support staff total and a tiered support model, but does not specify the split across tiers. The 5/2/1 breakdown (Tier 1/2/3) was invented as plausible for a team of that size.

All Quotable Statements

Every direct quote attributed to an executive was written for the brief. The backstories contain paraphrased positions and opinions but no direct quotes in this format. Each quote was crafted to match the established voice:

- Marcell: cautious, breach-conscious, strategic
- Mark: technically honest, pragmatic, anti-hype
- Sarah: customer-first, protective of team
- Aisha: numbers-focused, ROI-driven, anti-waste
- Sophia: governance-first, breach-scarred, risk-aware

Preferred Initiative Assignments

The backstories establish each executive’s general AI interests but do not explicitly name a single “preferred initiative” from the six in the introduction brief. The assignments were inferred:

Executive	Assigned Initiative	Inference Basis
Marcell	Customer support chatbot	Backstory emphasises customer experience as top strategic priority

Executive	Assigned Initiative	Inference Basis
Mark	Predictive maintenance	Backstory says “maybe predictive maintenance” as starting point; infrastructure data is his strongest asset
Sarah	Churn prediction	Backstory lists “customer health scoring and churn prediction” as where she sees potential
Aisha	Lead scoring	Backstory says she’d see value “if it helped us win enterprise deals”; lead scoring is lowest complexity
Sophia	Security threat detection	Backstory says AI could help with “threat detection, anomaly detection, automated response”

Preliminary Opportunity Assessments

The data readiness, stakeholder support, and ethical risk assessments for each of the six opportunities were composed for the brief. They draw on backstory content (e.g., Mark’s data readiness scores, Sophia’s governance concerns) but the specific assessments per opportunity are invented.

Specific Details Within Assessments

- “Support ticket history exists (3 to 4 years)” — Jamal’s backstory says “3-4 years” of historical data, but does not specify this applies to support tickets specifically
- “500 to 800 daily alerts” — sourced from docs/policies/ (incident response policy implementation notes)
- HubSpot lead tracking described as “basic” — inferred from Tom Bradley’s backstory describing limited marketing automation
- “\$2,400 customer acquisition cost” as baseline — sourced from Lisa Chen backstory

This companion document is for instructor reference. It is not intended for student distribution unless adapted.