

MoneyHero

SEO Recovery & Growth Plan

Singapore & Philippines Markets

Senior Associate SEO — Case Study Submission

Markets: Singapore (EN) | Philippines (EN)

CONFIDENTIAL

1. Executive Summary

A simultaneous ~15% month-over-month decline in organic traffic and credit-card lead conversions across Singapore and the Philippines points to shared-driver problems rather than isolated editorial variance. Three converging factors are most likely responsible: an algorithmic reassessment applying stricter trust standards to financial content (YMYL core update), a shared template performance regression degrading mobile experience metrics (Core Web Vitals), and multi-market targeting failures causing ranking signal dilution through hreflang, canonical, and schema misalignment.

These issues are compounded by an accelerating shift in Google's search results. AI Overviews now answer many financial comparison queries directly, measurably reducing click-through rates even when rankings remain stable. The strategic response must therefore address not only traditional SEO recovery but also position the platform to be cited within AI-generated answers—a first-class growth lever in the 2026 search landscape.

This plan operationalises recovery across three parallel streams: diagnosis-first technical audit with clear thresholds and fixes, a “Search Universe” strategy focused on AI Overview and LLM citation visibility (AIO & GEO), and stakeholder-ready prioritisation with confidence-ranged timelines. The twelve-week roadmap moves from stabilisation through optimisation to growth, with measurable checkpoints at each phase boundary.

2. Abbreviations

Abbreviation	Definition
AIO	AI Overview — Google's AI-generated summary at the top of search results
CLS	Cumulative Layout Shift — visual stability metric (threshold: ≤ 0.1)
CTR	Click-Through Rate — percentage of impressions resulting in a click
CWV	Core Web Vitals — Google's UX metrics (LCP, INP, CLS) used as ranking signals
E-E-A-T	Experience, Expertise, Authoritativeness, Trustworthiness — Google's quality framework
GA4	Google Analytics 4 — current web analytics platform
GEO	Generative Engine Optimisation — improving visibility in AI-generated answers
GSC	Google Search Console — monitoring search performance and indexing
INP	Interaction to Next Paint — responsiveness metric (threshold: $\leq 200\text{ms}$)
LCP	Largest Contentful Paint — loading performance metric (threshold: $\leq 2.5\text{s}$)
LLM	Large Language Model — AI systems (ChatGPT, Perplexity) that cite web sources
SEO	Search Engine Optimisation — improving organic search visibility
SERP	Search Engine Results Page
YMYL	Your Money or Your Life — Google's classification for financial/health content

Part 1: Diagnostic & Technical Audit

This section addresses the “why”: identifying the root causes behind a simultaneous traffic decline across both markets.

3. Hypothesis Generation

Three distinct patterns can produce similar traffic declines across SG and PH simultaneously but require different remediation strategies. All three can co-occur given the shared infrastructure across SingSaver, Seedly, and market-specific templates.

Hypothesis	Why It Hits Both Markets	Risk	Primary Tool
A. Core Update / YMYL Reassessment	Both markets share templates, editorial format, and scoring methodology. A core update can reduce rankings across the vertical when Google applies stricter trust standards to financial comparison content.	High	GSC + Ahrefs
B. Core Web Vitals Regression	A single front-end deployment (JS framework update, tag manager bloat, ad widget) can degrade LCP/INP/CLS across both locales and brands if they share components. Affects conversions even when rankings hold.	Medium-High	GSC + GA4
C. Hreflang / Canonical / Schema Failures	Shared hreflang/canonical patterns can cause wrong-country URLs to rank, duplicates to compete, and cross-brand cannibalisation. Simultaneously, AI Overviews suppress CTR even when positions are stable.	High	SF + GSC

4. Audit Actions

4.1 GSC Performance Analysis

Report: Performance → Search results (Web). Compare last 28 days vs previous 28 days. For core update correlation, compare 7 days after rollout completion against 7 days before rollout started.

Filters: Country = Singapore, then Philippines. Page regex = `.*\/credit-card\/.*`. Device = Mobile first.

Thresholds: Avg position drop ≥ 2.0 across top 20 queries for ≥ 14 days; impressions drop $\geq 10\%$; clicks drop $\geq 15\%$. If impressions and position remain stable but CTR drops $\geq 15\%$, the driver is SERP feature displacement (AI Overviews), not ranking loss.

4.2 GA4 Conversion Validation

Report: Acquisition → Traffic acquisition. Organic Search; Landing page = `/credit-card/`; Country = SG vs PH.

A disproportionate drop in key event rate ($\geq 10\%$ relative) on mobile vs desktop can indicate CWV-driven conversion loss rather than pure ranking decline.

4.3 Core Web Vitals (GSC)

Report: CWV report, Mobile first. Compare against last front-end release date.

Metric	Good	Needs Improvement	Poor
LCP	≤ 2.5 seconds	2.5–4.0 seconds	> 4.0 seconds
INP	≤ 200 ms	200–500ms	> 500 ms
CLS	≤ 0.1	0.1–0.25	> 0.25

Action trigger: $\geq 20\%$ increase in “Poor” URLs for LCP or INP within the credit-card group, coinciding with the decline window.

4.4 Screaming Frog Validation

Config: Hreflang extraction, canonical extraction, schema validation, JavaScript rendering enabled. Crawl both brand domains.

- **Hreflang:** Missing return links, non-200 hreflang URLs, references not aligned with canonicals.
- **Canonicals:** Self-referencing enforced, no canonical pointing to 404, Google-selected canonical matches intended (URL Inspection).
- **Structured data:** Product/Review validation errors, policy compliance (markup = visible content), manual action checks.
- **Indexing:** Spikes in “Duplicate without user-selected canonical” in the credit-card section.

5. Cannibalisation Resolution: SingSaver vs Seedly

SingSaver and Seedly are targeting identical credit card queries with overlapping content. The objective is distinct brand ownership by intent—not merely a ranking swap—executed at URL level with canonicals, hreflang, and crawl behaviour aligned.

5.1 Brand Intent Separation

Dimension	SingSaver	Seedly
Primary Intent	Transactional comparison: “apply now,” eligibility, best-for lists, application CTAs.	Informational depth: finance education, community reviews, “how to choose” guides.
AI Citation Role	Authoritative product comparison source for transactional queries.	Independent, earned-authority voice—generative engines favour third-party sources.

5.2 Technical Directives

1. Self-referencing canonicals on all indexable pages. Cross-domain canonicals only for explicit content retirement.
2. Hreflang within each brand domain independently with full return-link sets (en-SG, en-PH, tl-PH, x-default) per ISO 639-1 + ISO 3166-1 requirements.
3. 301 redirects for deprecated duplicates: assign ownership to the intent-aligned brand, redirect the other.
4. Internal linking reinforced per brand’s topical cluster. Differentiated titles, headers, and structured data.

5.3 Rollout Sequence

- **Days 1–3:** Export top 100 credit-card pages from GSC. Map keyword ownership via Ahrefs (filtered by AI Overview presence).
- **Days 4–7:** Canonical + hreflang corrections in staging. Screaming Frog validation: missing return links = 0, non-200 hreflang URLs = 0.
- **Week 2:** Controlled release by template: infrastructure first, then redirects, then content rewrites.
- **Weeks 2–6:** Monitor GSC indexing reports, CTR/position on affected queries, GA4 key event rate recovery.

Part 2: Search Universe Strategy (AIO & GEO)

This section addresses the pivot from “ranking #1” to “winning the answer” across AI Overviews, Perplexity, and ChatGPT Search.

6. AIO-Optimised Content Framework

Using the “Best Travel Credit Cards” page as a blueprint, the content template is designed to satisfy three retrieval systems simultaneously: classic SEO (rankings + rich results), AI Overviews (chunkable, corroboratable fact blocks), and LLM answer engines (citation-ready content with explicit sourcing). This differs from a traditional SEO page in three key ways:

- **Structure for extraction, not just reading:** Content is segmented into self-contained fact blocks that can be cited independently by AI systems, rather than flowing as continuous editorial prose.
- **Corroborability over persuasion:** Every claim is anchored to verifiable data (issuer terms, rates, dates) with inline citations—LLMs and AI Overviews preferentially cite sources they can cross-reference.
- **Decision architecture over listicle format:** Interactive decision trees and “best for” segmentation replace flat ranked lists, matching how AI systems decompose comparison queries into conditional recommendations.

Content Layer	Implementation
Above-the-Fold Answer	80–120 word summary with top 3 picks. Eligibility baseline, methodology anchor, timestamped freshness signal.
Card Fact Blocks	Per card: “best for,” key earn rate, costs/tradeoffs, “who should avoid.” Standalone citation-ready snippets.
Structured Data	Product schema (ItemList wrapper), Review snippets where visible review exists, consistent fields across SG/PH.
Decision-Tree Section	“Find your best card in 60 seconds”—maps spending patterns to recommendations. LLM-friendly conditional logic.
Transparency Signals	Methodology page (scoring rubric, affiliate disclosure, changelog). Author credential pages with per-market expertise.

6.1 Localisation Framework

Singapore: Miles-centric positioning, MAS regulatory alignment, travel insurance and lounge access as primary comparison dimensions.

Philippines: Cashback prioritisation, instalment culture integration, Tagalog hybrid support (tl-PH) for educational content. Same data fields and table structures across all variants.

7. E-E-A-T & Authority Strategy

AI engines and modern SERPs reward sources that are easy to corroborate. For YMYL queries, the bar is higher. The strategy operates across three pillars:

Brand Citation Campaigns: Publish original datasets per region (SG miles valuations, PH cashback economics). Pitch to tier-1 finance publishers for citations and authoritative backlinks.

Editorial Transparency: Crawlable “How we rank credit cards” methodology page with scoring rubric, affiliate disclosure, update cadence, and changelog. Author credential pages with per-market expertise.

LLM Retrieval Alignment: ChatGPT Search and Perplexity include source citations—“being cited” is a measurable KPI. Optimise through consistent brand naming, structured knowledge graph references, and semantic consistency across markets.

8. AI-Leveraged Scaling Workflow

This workflow uses AI tools for speed while preserving compliance, financial accuracy, and market nuance. The goal is rapid keyword clustering and content localisation across SG and PH without sacrificing quality.

8.1 Tooling

- **Ahrefs:** Keyword discovery + AI Overview opportunity mapping (SERP feature filters, AI citations module).
- **ChatGPT:** Intent-based keyword clustering and initial content drafting.
- **NeuronWriter:** Semantic optimisation scoring and competitor content-gap analysis.
- **Screaming Frog:** Post-production validation of hreflang, canonical, and structured data compliance.

8.2 Keyword Clustering Workflow

Export the credit-card keyword universe from Ahrefs for both SG and PH locations. Feed the keyword CSV into ChatGPT with a structured clustering prompt:

```
Cluster these keywords into intent-based groups for a credit card comparison site.
Output JSON: cluster_name, intent (transactional/informational),
primary_keyword, secondary_keywords, recommended_page_type
(hub/list/review/guide), and market (SG, PH-EN, PH-TL).
Rules:
- Keep clusters mutually exclusive.
- Map each cluster to a single page.
- Flag clusters likely to trigger AI Overviews (comparison/explanatory queries).
```

Validate clustering output against Ahrefs SERP overlap data to confirm intent alignment. Feed validated clusters into NeuronWriter to score semantic gaps against top-ranking competitors and identify content angles for differentiation.

8.3 Localisation Workflow

For each content piece, generate locale-specific variants using a structured localisation prompt:

Localise this section for:

- (A) Singapore English (en-SG)
- (B) Philippines English (en-PH)
- (C) Philippines Tagalog (tl-PH)

Constraints:

- Preserve identical data fields and table structure across locales.
- Adjust examples, terminology, and user concerns per market.
- Do not change financial facts. Flag missing data as [NEEDS SOURCE].

Market nuances: SG skews miles-centric with dense premium card competition; PH emphasises cashback and instalment flexibility with broader eligibility. These assumptions should be validated against GA4 query-to-conversion paths.

8.4 Quality Assurance (Non-Negotiable)

All AI-assisted content passes through a four-gate QA process before publication:

5. Human financial accuracy check: Content Ops verifies rates, fees, eligibility, and regulatory references against issuer source material.
6. Compliance validation: affiliate disclosures present, no misleading claims, structured data matches visible content.
7. Technical validation: Screaming Frog crawl confirms self-canonical, hreflang return links, and schema passing validation.
8. Tone consistency: editorial review ensures consistent brand voice across AI-drafted variants.

Part 3: Stakeholder Management & Prioritisation

9. Stakeholder Script

When the Head of Business Development (a non-SEO audience) asks “Why is traffic down, and when will it come back?”—the following response manages expectations while demonstrating control:

Meeting-Ready Script:

"The traffic decline is driven by changes in how Google displays search results. For financial comparison queries, Google increasingly answers questions directly with AI-generated summaries—which reduces clicks to all websites, even when our rankings haven't changed."

"At the same time, Google appears to be applying stricter quality standards to financial content. We're also seeing technical issues where our Singapore and Philippines pages are competing against each other instead of ranking independently."

"We're already running a structured diagnostic and have identified specific technical fixes we can ship within 2–3 weeks. Conversion recovery from performance improvements typically follows shortly after. The broader ranking recovery is less predictable—we should see progressive improvement over 6–8 weeks rather than a single bounce-back."

"Our growth pivot is to ensure MoneyHero is the source that AI answers cite. This is now measurable and trackable—we're treating 'being the cited answer' as a core KPI alongside rankings and conversions."

10. Prioritisation Matrix

Tasks ranked by impact-to-effort ratio with limited dev resources. Schema takes highest priority; long-tail production is deferred until foundational issues are resolved.

Rank	Task	Effort	Impact	Rationale
P1	Implement Product + Review Schema	30–60 hrs	High	Improves rich result eligibility and AI Overview citation readiness simultaneously. Must comply with structured data policies and match visible content.
P2	Fix broken 404 links	8–20 hrs	Medium	Reduces crawl waste, preserves internal link equity. Low risk, fast execution—clears technical debt before larger changes.
P3	Update “Last Updated” dates	4–12 hrs	Low–Med	Freshness signal improvement. Only effective when coupled with substantive content updates—cosmetic-only changes risk quality issues per Google's core update guidance.

P4	Write 50 new long-tail articles	50–140 hrs	Medium	Deprioritised: if the decline is driven by CTR suppression or technical dilution, volume won't fix it and risks diluting topical focus. Best executed after templates and authority improve.
-----------	--	------------	--------	--

11. Recovery Roadmap

The twelve-week plan moves through three phases with confidence intervals reflecting the reality that root causes may overlap.

Phase	Duration	Key Actions	Confidence
Stabilise	Weeks 1–2	Technical audit (GSC/GA4/Ahrefs/SF). Cannibalisation fixes. Schema deployment across credit-card pages.	70% in 2–4 weeks; 90% in 4–8 weeks
Optimise	Weeks 3–6	AIO content rebuild. CWV template fixes. Internal linking restructuring. Authority reinforcement launch.	60% GA4 lift in 2–6 weeks post-release
Grow	Weeks 6–12	AI-assisted content scaling with editorial QA. Data PR campaigns. AI citation monitoring. Monthly reporting.	50% partial recovery in 6–12 weeks

12. Measurement Framework

KPI	Tool	Configuration
Organic clicks & impressions	GSC	Performance (Web). Country = SG vs PH. Page regex = credit-card. Device split.
CTR & position diagnostics	GSC	Compare date ranges. Annotate CTR vs position divergence.
Indexing & canonical health	GSC	Page indexing report + URL Inspection. Monitor Google-selected canonical.
Core Web Vitals	GSC	CWV report (LCP/INP/CLS). Mobile first. Compare against last release.
Sessions & conversions	GA4	Traffic acquisition. Organic Search. Landing page = credit-card. Country + device.
AI Overview presence	Ahrefs	SERP features filter (AI Overview). AI citations module. SG + PH locations.
Technical errors	SF	Hreflang + structured data validation. Both brand domains. Return links + canonical consistency.

Conclusion

The cross-market traffic decline reflects a convergence of algorithmic reassessment, AI-driven SERP displacement, and technical targeting failures rather than any single cause. Recovery requires parallel action across all three vectors: resolving the technical infrastructure that splits ranking signals between brands and markets, rebuilding content architecture to be citation-ready for AI search features, and reinforcing authority signals that meet the elevated trust standards applied to financial content.

The phased roadmap delivers measurable stabilisation within the first four weeks through technical fixes within direct control, followed by content and authority investments that address slower-moving but higher-impact competitive dynamics. The strategic pivot to “win the answer”—optimising for citation visibility in AI Overviews and LLM search—positions the platform not just for recovery but for growth in a search landscape where traditional click-through is structurally declining.

End of Report