

The Coetzee Liquidity Protocol

An Operational Model for Niche Commerce Authority
in High-Velocity Search Markets

ABSTRACT

The Coetzee Liquidity Protocol (CLP) is an operational methodology for SME commerce operators competing in high-velocity, attribute-driven niche markets against generalist marketplaces. It formalises the conditions under which a niche merchant can achieve superior machine-readable product authority — not through logistical scale, but through semantic depth. The protocol introduces a five-variable Search Confidence formula governed by a binary Resilience Floor precondition, a Specificity Threshold model for niche qualification, an Economic Viability gate, and a Harvest/Pivot rule triggered by a defined Entropy Signal. The CLP is the second pillar of the SEO Gurus Methodology Library, complementing the Coetzee Convergence Framework (CCF). Together they constitute a two-pillar architecture separating the Trust Economy from the Liquidity Economy.

v0.4 · Post-Adversarial Draft · March 2026

Erwee Coetzee · SEO Gurus · seo-gurus.co.za

Session: WBT-2026-03-08-ZA-PILLAR-2-STRESS-TEST

Published under Creative Commons Attribution 4.0 International (CC BY 4.0)

If $R_f = 1 \rightarrow SC = (L_q \times U_i \times T_f) \times B_g \times E_c$

CLP v0.4 Master Formula

Table of Contents

1.	The Liquidity Hypothesis
	1.1 Browsing, Auditing, and Routing
	1.2 The Structural Implication
2.	Market Classification
	2.1 Defining High-Velocity Commodity Markets
	2.2 Qualifying and Disqualifying Criteria
3.	The Search Confidence Formula
	3.1 The Resilience Floor Precondition (Rf)
	3.2 Formula Variables Defined
4.	The Five Operational Pillars
	4.1–4.5 Pillars I through V
5.	The Specificity Threshold
	5.1 Threshold Assessment Criteria
	5.2 The Economic Viability Calculation
	5.3 SA Data Tax Reference Ranges (2026)
6.	The Entropy Signal and Harvest/Pivot Rule
7.	Relationship to the Coetzee Convergence Framework
8.	Adversarial Stress-Test Summary
9.	Falsifiability Conditions
10.	References
App. A	CLP Resilience Audit Checklist
App. B	Specificity Threshold Self-Assessment

SECTION 01

The Liquidity Hypothesis

The Coetzee Liquidity Protocol is predicated on a specific and testable behavioural distinction. In high-friction search environments, users audit. In high-velocity commodity environments, users — and in 2026, increasingly their AI agents — route. The routing decision happens before the human is involved.

1.1 Browsing, Auditing, and Routing

BROWSING	AUDITING	ROUTING
Low-friction, exploratory	High-friction, verification-driven	High-velocity, attribute-driven
User selects on surface relevance	User verifies credibility before commitment	Agent matches SKU attributes to intent
Standard relevance optimisation	Coetzee Convergence Framework (CCF)	Coetzee Liquidity Protocol (CLP)
Example: pizza recipe	Example: bespoke diamond ring Cape Town	Example: 3-person ultralight tent 20D nylon, in stock, CPT delivery

1.2 The Structural Implication

A standard SEO approach focused on keyword relevance does not pass an agentic routing check. The routing agent matches a constrained intent to a structured product record. If the merchant's product record is incomplete, ambiguous, or unavailable at query time, the agent routes to the next result. The CLP is the engineering protocol for a commerce presence that passes that routing check — and that survives when the infrastructure that performs it fails.

"In high-velocity markets, search engine confidence and agentic routing confidence are not separate problems. They are the same problem expressed at different speeds."

— Erwee Coetzee, CLP v0.4

SECTION 02

Market Classification

The CLP does not apply to all commodity markets. The Specificity Threshold defines the minimum conditions a category must meet. Misclassification — applying the CLP to a below-threshold market — produces negative returns: data maintenance costs without viable margin capture.

2.1 Defining a High-Velocity Commodity Market

Three concurrent conditions must be present for the CLP to apply:

CONDITION	DEFINITION	CLP IMPLICATION
1 — Attribute-Driven Selection	Decision determined by specification, compatibility, availability, and price — not brand trust.	Ui (Use-Case Interpretability) is the primary moat variable.
2 — Machine-Routable Intent	Intent specific enough to be resolved by structured data. An agent can match the query to a SKU.	Lq (Liquidity) enables the routing. Without it, the merchant is invisible.
3 — Margin-Viable Niche	Gross margin sufficient to fund the data maintenance cost required for Attribute Certainty.	Ec (Economic Viability) gates the entire protocol.

2.2 Qualifying and Disqualifying Criteria

QUALIFYING MARKETS	DISQUALIFYING CONDITIONS
Specialist outdoor and hiking equipment (technical SKUs)	True commodity markets — AA batteries, generic consumables. Amazon wins on volume economics.
Compatibility-dependent automotive and mechanical components	High-friction trust markets — luxury goods, professional services. Apply the CCF instead.
Industrial replacement parts and specialised tooling	Category gross margin sub-15% — data maintenance cost cannot be amortised.
Specialised electronics accessories with fitment constraints	Pure service businesses unless the service is productised into defined deliverable SKUs.
Pet nutrition, hunting optics, precision sport equipment	Fewer than 5,000 monthly searches in the targeted specification query cluster.

SECTION 03

The Search Confidence Formula

The CLP formalises commodity Search Confidence (SC) as a governed equation. It is a multiplicative model: weakness in any variable suppresses the output of all others. The Resilience Floor (Rf) is a binary precondition whose failure renders the formula inoperable regardless of the values of all other variables.

CLP v0.4 MASTER FORMULA

$$\text{If } Rf = 1 \rightarrow SC = (Lq \times Ui \times Tf) \times Bg \times Ec$$

IF RESILIENCE FLOOR FAILS

$$\text{If } Rf = 0 \rightarrow SC = 0 \text{ (regardless of all other variables)}$$

3.1 The Resilience Floor Precondition (Rf)

The Resilience Floor is the lesson of the February 2026 Merchant Centre Content API outage — a 17-day disruption during which merchants operating on a feed-only architecture reported revenue losses of up to 40%. Minimum threshold for $Rf = 1$: at least 30% of revenue must be sustainable without Merchant Centre feed distribution (the Revenue Independence Ratio).

REVENUE INDEPENDENCE RATIO

$$RIR = (\text{Revenue_direct} + \text{Revenue_organic} + \text{Revenue_email}) \div \text{Revenue_total} \geq 0.30$$

3.2 Formula Variables Defined

VAR	NAME	DEFINITION	ROLE
Lq	Liquidity Baseline	Accuracy and freshness of stock, pricing, and availability data across all feed surfaces.	Enables discovery. Does not differentiate from large marketplaces.
Ui	Use-Case Interpretability	Degree to which taxonomy, attributes, and content make a product unambiguously matchable to a narrow buyer intent. The Boolean True return.	Primary moat variable. Where SMEs structurally outperform generalist marketplaces.
Tf	Trust Floor	Merchant reliability signals using 2026-available infrastructure: Merchant Centre history, verified GBP, structured return/shipping policies, review integrity.	Binary pass/fail in practice. Must meet minimum threshold.

VAR	NAME	DEFINITION	ROLE
Bg	Buyer-Goal Fit	Merchant ability to satisfy a specific constrained intent better than a generalist marketplace. Alignment between niche taxonomy and target query cluster.	Amplifier. Answers: why does an agent choose you when Amazon exists?
Ec	Economic Viability	Whether margin captured by the niche exceeds the data maintenance cost (Dc). Profitability condition: $(Sv \times Cr \times Cm + Cv + Ac) > Dc$.	Business model gate. High Ui is worthless if the niche cannot fund its data layer.

*"A niche is only viable if it wins both the search algorithm and the profit equation.
Visibility without margin is not a strategy — it is a subsidy to the marketplace."*

— Signal Viability Principle, CLP v0.4

SECTION 04

The Five Operational Pillars

The CLP is applied as a five-pillar sequential architecture. Pillar I establishes the structural precondition. Pillars II-IV constitute the discovery and dominance layer. Pillar V is the competitive defence discipline.

The Resilience Floor (Rf)

Objective (Rf): Ensure the business can sustain a minimum revenue floor during a 14-day platform outage. Required: crawlable product pages with independent organic ranking; Product schema on every PDP mirroring all feed attributes (the Crawlable Shadow); multi-platform feed distribution (Google, Bing, Meta); email and returning customer base. During the February 2026 outage, long-tail organic traffic for technical SKUs spiked 300%. Merchants with a functioning Crawlable Shadow captured it. Feed-only merchants did not.

The Liquid Tip (Lq)

Objective (Lq): Achieve and maintain feed accuracy sufficient to qualify for agentic routing consideration. Liquidity is table stakes — not a moat. Real-time or near-real-time stock and pricing pushed via the Content API. Out-of-stock items are an authority killer: an agent that routes to an unavailable product does not route to that merchant again. Every feed attribute must also appear in on-page Product schema (JSON-LD Parity).

Attribute Certainty (Ui)

Objective (Ui): Return a Boolean True on compatibility and specification queries where the generalist marketplace returns a Maybe. Agentic routing models are calibrated to avoid ambiguity because the cost of returns is embedded in confidence scoring. For most SME niches, Attribute Certainty does not require a proprietary database — leasing specialist fitment data via API is sufficient. Economic gate: four sales at 35% margin on an R8,000 AOV cover the annual cost of an R8,000/month data subscription.

Verified Trust (Tf)

Objective (Tf): Meet the Trust Floor required for agents and users to consider the merchant credible — using 2026-available infrastructure only. Components: Merchant Centre account history and policy compliance; verified Google Business Profile; structured return and shipping policy markup; review integrity from a single consolidating source; domain age and brand search consistency. An Expert-Verified Returns layer in technical categories creates a trust signal that agentic systems prioritise.

The Moving Taxonomy

Objective (Bg + moat defence): Maintain a moving specificity frontier that extends taxonomy depth faster than generalist marketplaces can standardise the discovered attributes. Amazon's feed team is reactive: when a niche generates visible transaction volume, they begin enriching their own attribute data. The Moving Taxonomy ensures the SME's Ui advantage survives. In practice: extend category depth beyond the visible query cluster; add compatibility attributes before they appear in marketplace schema standards; use customer query logs to discover edge-case attributes. This proprietary edge-case dataset is the one layer that cannot be purchased by a marketplace.

SECTION 05

The Specificity Threshold

The Specificity Threshold defines the minimum conditions a product category must meet before the CLP generates a positive return. Below the threshold, generalist marketplaces win on volume economics. Above it, the SME can dominate on semantic depth — for as long as the Moving Taxonomy discipline is maintained.

5.1 Threshold Assessment Criteria

CRITERION	BELOW THRESHOLD	ABOVE THRESHOLD	ASSESSMENT METHOD
Attribute Complexity	3-5 generic attributes	10+ attributes incl. compatibility, fitment, certification	Count disambiguating attributes for confident purchase
Marketplace Taxonomy Depth	Amazon / Takealot equivalent granularity	Marketplace listing returns ambiguous or incomplete attributes	Manual audit of top 3 marketplace listings
Category Gross Margin	Sub-15% — data layer unviable	≥25% — data layer fundable	Margin calculation vs estimated Dc
Long-Tail Query Specificity	1-2 word head terms, marketplace PPC	5-10 word specification queries, technical vocabulary	Keyword research for attribute-specific long-tail
Decision Risk	Wrong product easily returned	Wrong product causes cost, delay, or operational failure	Customer journey audit: does buyer verify before transacting?

5.2 The Economic Viability Calculation

Before committing to a niche, calculate the Net Moat against the Data Tax:

NET MOAT — ECONOMIC VIABILITY GATE

$$\text{Net Moat} = (\text{Margin} \times \text{Conversion_Agent}) - (\text{Data Tax} + \text{Friction Cost})$$

PROFITABILITY CONDITION (ANNUAL)

$$(\text{Sv} \times \text{Cr} \times \text{Cm}) + \text{Cv} + \text{Ac} > \text{Dc}$$

Where: Sv = monthly search demand; Cr = conversion rate improvement from Attribute Certainty; Cm = category gross margin per unit; Cv = customer lifetime value; Ac = acquisition cost avoided via organic discovery; Dc = annual data maintenance cost.

5.3 South African Data Tax Reference Ranges (2026)

LAYER	DESCRIPTION	ANNUAL COST (ZAR)	SUITABLE FOR
Lightweight	Manual compatibility matrices maintained internally	R20,000 – R80,000	Small niche categories, limited SKU range
Licensed Database	Subscription to industry fitment databases via API	R100,000 – R400,000	Automotive, technical equipment — established fitment ecosystems
Full Data Engineering	Custom compatibility dataset built and maintained internally	R500,000+	High-margin verticals where the proprietary dataset is itself a moat

Data provider concentration risk note: If the API provider from whom compatibility data is leased is acquired by a marketplace competitor, the leased moat becomes a shared moat. The CLP recommends supplementing leased datasets with a proprietary edge-case layer built from customer query logs.

SECTION 06

The Entropy Signal and Harvest/Pivot Rule

Every niche has a discovery half-life — estimated at 18-36 months in typical above-threshold categories. The Entropy Signal is the observable indicator set that the Ui advantage is being eroded. The Harvest/Pivot Rule defines the required response.

ENTROPY SIGNAL	INDICATOR	REQUIRED ACTION	WINDOW
Marketplace Attribute Parity	Amazon/Takealot adds your niche attributes to their core schema. Listings return comparable attribute completeness.	PIVOT: Deepen taxonomy. Extend to sub-niche attributes and edge-case compatibility data not yet standardised.	180 days from first parity signal
Margin Compression	CPC for niche keywords exceeds 30% of category gross margin. Organic click share declining.	HARVEST: Cease aggressive acquisition. Redirect to LTV — email, repeat purchase, bundle architecture.	Immediate on breach of 30% threshold
Data Provider Concentration Risk	Your leased compatibility database provider acquired by or enters preferred partnership with a marketplace competitor.	EXIT / BUILD: Migrate to alternative provider. Accelerate proprietary edge-case dataset from customer query logs.	60 days from acquisition announcement
Query Cluster Commoditisation	Long-tail query cluster attracts high-volume generic PPC. Adjacent category CPCs increase 40%+.	BRAND / COMMUNITY SHIFT: The cluster is becoming paid-media commodity. Shift to brand-search and direct channel.	90 days from CPC escalation

"If a marketplace achieves 90% attribute parity with your niche taxonomy, the moat is breached. You have 180 days to extend the specificity frontier by 20%, or pivot the revenue floor to brand and community loyalty."

— CLP Harvest Rule, v0.4

SECTION 07

Relationship to the Coetzee Convergence Framework

The CLP and the CCF are architecturally distinct responses to structurally different search economies. The central discovery of the stress-test session was the formal separation of two search economies that most existing SEO theory conflates.

DIMENSION	CCF	CLP
Market Type	High-friction, high-trust	High-velocity, high-specificity
Buyer Behaviour	Auditing — verification before commitment	Routing — attribute matching before human involvement
Primary Signal	Authority and credibility convergence	Inventory interpretability and attribute specificity
Moat Variable	Entity signal convergence ($S_d + C_c + P_h$)	Attribute Certainty (U_i) — Boolean True returns
Platform Risk	Low — not feed-dependent	High — requires R_f binary gate
Competitor	Trust-deficit operators in the same niche	Generalist marketplaces with superior logistics
Scoring Tool	Entity Confidence Score (ECS) 0.00–1.00	Net Moat = (Margin × Conversion) – Data Tax

The correct diagnostic question before selecting a framework: does the buyer in this market audit or route? If they audit, apply the CCF. If they route, apply the CLP. If the market contains both buyer types at different intent stages, a hybrid architecture is indicated.

"Most SEO theory collapses both economies into one model. This two-framework architecture separates them. The question 'audit or route?' precedes every other strategic decision."

— CLP v0.4 — Two-Economy Framework

SECTION 08

Adversarial Stress-Test Summary

CLP v0.4 is the result of six rounds of structured adversarial pressure-testing in Session WBT-2026-03-08-ZA-PILLAR-2-STRESS-TEST. The table below documents each material kill attempt, its validity assessment, and the protocol response.

KILL ATTEMPT	VALIDITY	PROTOCOL RESPONSE	VER
Liquidity is a race to the bottom — Amazon always wins on speed and scale	Partially valid for generic markets. Invalid above the Specificity Threshold.	CLP scope restricted to above-threshold niches only. Threshold criteria formalised.	v0.2
Cold-start problem defeats new entrants — agents favour historical transaction volume	Valid in generic categories. Diminished where Data Granularity outweighs transaction history.	Tf (Trust Floor) introduced as solvable via 2026-available signals. Cold-start disadvantage acknowledged at launch.	v0.2
Data Tax exceeds SME capacity — building compatibility databases is a losing investment	Valid for proprietary builds. Invalid for leased API-based data layers in qualifying niches.	Ec (Economic Viability) gate and Data Tax Reference Ranges introduced. Lease-don't-build as default.	v0.3
Platform concentration risk — 100% feed dependency is a single point of failure	Valid. February 2026 MC outage validated empirically (17 days, 40% revenue loss).	Rf (Resilience Floor) introduced as binary precondition. If Rf = 0, SC = 0.	v0.4
Moving Taxonomy is a process answer, not economic — Amazon will acquire the data provider	Valid. Data provider concentration risk is real.	Provider concentration risk added to Entropy Signal table. Proprietary edge-case layer from query logs as non-acquirable hedge.	v0.4
The framework is a compliance checklist for second place, not a business model	Valid at v0.1. Protocol has since acquired Rf, Ec, and the Harvest/Pivot Rule.	Formal acknowledgement: CLP is a justification framework for properly architected niche commerce.	v0.4

SECTION 09

Falsifiability Conditions

The CLP predicts that implementing the full five-pillar architecture for a verified above-threshold commodity niche will produce measurable improvement in at least two of the following outcomes within 90 days of full implementation:

- Improvement in Google Search Console impressions and click share for the targeted long-tail specification query cluster, measured against the pre-intervention 90-day average.
- Increase in Merchant Centre product approval rate and feed attribute completeness score.
- Measurable organic traffic to product and category pages independent of feed-driven Shopping surfaces — verifiable in GSC by filtering for non-Shopping placements.
- Improvement in qualified conversion rate from specification-query organic traffic versus head-term traffic baseline over a comparable session volume period.

The protocol is weakened — not invalidated — if a niche merchant with materially better taxonomy depth, compatibility metadata, and feed accuracy fails to outperform a generalist marketplace over a 90-day observation period. Such cases should be reviewed for Specificity Threshold misclassification, Economic Viability miscalculation, or implementation failure before the protocol is considered invalid.

CLP v0.4 is published as a post-adversarial draft. It has not yet been validated against a live commerce category baseline. Empirical validation in a defined niche, with a documented pre-intervention query cluster and organic traffic baseline, is the identified next phase. Practitioners who apply the protocol and observe results are invited to submit case observations for inclusion in the evolving validation dataset.

SECTION 10

References

- [1] Spence, M. (1973). Job Market Signaling. *The Quarterly Journal of Economics*, 87(3), 355–374. <https://doi.org/10.2307/1882010>
- [2] Barnard, J. (2020). Fundamentals of Entity SEO and the Entity Home. Kalicube. <https://kalicube.com/resources/>
- [3] Google LLC. (2024). Search Quality Rater Guidelines. <https://guidelines.raterhub.com/searchqualityevaluatorguidelines.pdf>
- [4] Google LLC. (2026). Merchant Centre Product Data Specification. <https://support.google.com/merchants/answer/7052112>
- [5] Google LLC. (2026). Content API for Shopping. <https://developers.google.com/shopping-content>
- [6] eFulfillment Service. (January 2026). 2026 Ecommerce Attribute Completion Benchmark Report. [Industry publication.]
- [7] Coetzee, E. (2026). The Coetzee Convergence Framework v1.3. SEO Gurus. <https://seo-gurus.co.za/methodology/coetzee-convergence-framework>
- [8] Session Log: WBT-2026-03-08-ZA-PILLAR-2-STRESS-TEST. (2026). Adversarial Stress-Test of the Coetzee Liquidity Protocol. Internal session record, SEO Gurus, Cape Town.

APPENDIX A

CLP Resilience Audit Checklist

Use this checklist to assess whether a client's commerce architecture meets the Resilience Floor (Rf) precondition. All Gate 1 items must pass before proceeding to Gates 2 and 3.

REF	GATE	CHECK ITEM	PASS CONDITION	✓
R1.1	GATE 1 / Resilient Base	Product pages crawlable and indexable without feed dependency	GSC shows organic impressions for product URLs independent of Shopping tab	■
R1.2	GATE 1	Product schema deployed on all PDPs with minimum required attributes	Zero critical errors in Google Rich Results Test for Product schema	■
R1.3	GATE 1	Category taxonomy pages rank independently in organic search	At least 3 category pages in GSC top-10 for relevant specification queries	■
R1.4	GATE 1	Revenue Independence Ratio (RIR) $\geq 30\%$	Direct + organic + email revenue = $\geq 30\%$ of total monthly revenue (prior 90 days)	■
R2.1	GATE 2 / Feed Layer	Merchant Centre feed has $\geq 95\%$ product approval rate	MC dashboard shows $\leq 5\%$ disapproved products	■
R2.2	GATE 2	Stock and price signals update within 2 hours of inventory change	Content API push confirmed; or automated crawl $\leq 24h$ with manual process for critical changes	■
R2.3	GATE 2	Feed distributed to at least 2 platforms beyond Google	Active product feeds on Bing Shopping and/or Meta Catalogue with $\geq 90\%$ approval rate	■
R3.1	GATE 3 / Audience	Email subscriber list exists and generates measurable revenue	At least one revenue-generating email campaign per 90 days with attributable transactions	■
R3.2	GATE 3	Returning customer rate $\geq 15\%$ of monthly transactions	GA4/CRM data confirms $\geq 15\%$ of monthly orders from returning customers	■

APPENDIX B

Specificity Threshold Self-Assessment

Score each criterion 0-2: 0 = clearly below threshold, 1 = borderline, 2 = clearly above threshold. Total score $\geq 8/10$ indicates the category likely qualifies and warrants a full Economic Viability calculation before commitment.

CRITERION	SCORE (0-2)	NOTES
1. Attribute Complexity — product requires 10+ disambiguating attributes		
2. Marketplace Gap — top 3 marketplace listings return ambiguous or incomplete attributes for equivalent product		
3. Category Margin — gross margin $\geq$ 25% per unit		
4. Long-Tail Query Volume — 5-10 word specification queries with technical vocabulary, $\geq$ 500 monthly searches		
5. Decision Risk — wrong product creates significant cost, delay, or operational failure		
TOTAL SCORE (≥ 8 = proceed to Economic Viability calculation)	/ 10	

This self-assessment is an indicative screen, not a definitive qualification. A full CLP Threshold Assessment includes category margin analysis, competitor feed audit, and keyword cluster verification. Request a full assessment: seo-gurus.co.za/contact/