

Entity Clarity Score



Measures how precisely and consistently a content unit names, references, and contextualises its entities — the people, organisations, frameworks, products, tools, and concepts the content is about.

Weight: 0.20 in RASA composite

WHAT IT MEASURES

AI systems do not read content as humans do — they resolve entities, matching named references against knowledge graphs and embedding indexes. ECS measures whether every entity in a content unit can be unambiguously resolved, correctly attributed, and reliably linked.

AFFECTS

- Knowledge Graph Attribution
- Entity-Specific Retrieval
- RAG Synthesis Confidence
- GEO Knowledge Positioning

4 PRIMARY FACTORS THAT DRIVE ECS SCORES

F-01

Named Entity Precision

Entities should be named exactly — full official names, registered product names, canonical framework names. Description substitutes ("a Bangalore-based AI company") are not resolvable entity references.

F-02

Terminological Consistency

The same entity must be referenced by the same name throughout. Varying between "the RASA framework," "this system," and "the framework" forces AI co-reference inference — a direct ECS penalty.

F-03

Absence of Pronoun Ambiguity

Pronouns ("it," "they," "this") become unresolvable when a chunk is extracted and processed in isolation. High-ECS content eliminates pronouns in favour of consistent named references throughout.

F-04

Contextual Grounding of New Entities

Every new entity introduced must be grounded on first mention with a categorical descriptor: "RASA-Analyst, the official evaluation engine for the RASA framework." Bare names alone create unresolvable nodes.

COMMON FAILURE MODES

- F1 Weak Entity Clarity**
Referencing organisations, tools, and frameworks through pronouns, generic descriptions, or category labels. AI cannot ground these and misattributes or excludes the content.
- F2 Inconsistent Terminology**
Using multiple names for the same entity creates separate candidate nodes, fragmenting retrieval signal. Especially damaging in RAG, where chunks from one document are reconciled at synthesis time.

ECS SCORE REFERENCE SCALE



ECS & KNOWLEDGE GRAPH ALIGNMENT

High ECS actively contributes to AI systems' understanding of an entity — its name, relationships, domain, and authority. For organisations building semantic authority via

9-10

EXCEPTIONAL

"Nebula Personalization Tech Solutions Pvt. Ltd. developed the RASA framework, a scoring system for AI retrieval compatibility."

7-8

STRONG

"The RASA framework evaluates content across five dimensions. It was developed by Nebula to improve AI retrieval outcomes."

5-6

MODERATE

"The company's framework helps with AI search. Their tool gives scores from 1-10 and the system improves content for retrieval."

3-4

WEAK

"They launched it last year to help with this kind of problem. The team said it works well with the new approach they use."

1-2

NO ENTITIES

"The company launched a new tool last year. It is very useful for businesses in this space."

Generative Engine Optimization (GEO), ECS is the dimension most directly linked to long-term knowledge graph positioning.

ECS IMPROVEMENT CHECKLIST

- ✓ Use full official names on first reference — never introduce via pronoun
- ✓ Choose one canonical name per entity and use it throughout
- ✓ Replace pronouns with named references for standalone chunks
- ✓ Ground every new entity with a brief categorical descriptor
- ✓ Audit for synonyms and variants — pick one and delete the rest
- ✓ Match entity names exactly across pages, schema, and external references

RASA COMPOSITE FORMULA

$$\begin{aligned}
 &RP \times 0.25 + SCC \times 0.20 \\
 &+ \mathbf{ECS \times 0.20} + SCI \times 0.20 \\
 &+ CGP \times 0.15
 \end{aligned}$$

PUBLISH ≥ 6.0 composite
 REJECT ≤ 6.0 on ECS alone

DD1
Retrieval Probability

DD2
Semantic Chunk Coherence

DD3 · ACTIVE
Entity Clarity Score

DD4
Synthesis Compatibility Index

DD5
Citation & Grounding Potential