

Eight ways to score the same evidence

Worked example · 17 September 2026

The agent covers **A and C** and misses **B**. First select the evidence; then assign weights. Every row scores the same answer trace.

Node	Distance	S1 base weight	Covered?
A	Seed · 0 hops	0.8	Yes
B	1 hop	0.6	No
C	2 hops	0.4	Yes

Score = covered selected-node weight ÷ total selected-node weight

Policy	Weights in denominator	Covered / total weight	Score
Seeds only · equal	1	1 / 1	100%
One hop · equal	1 + 1	1 / 2	50%
Two hops · equal	1 + 1 + 1	2 / 3	66.7%
Two hops · half per hop	1 + 0.5 + 0.25	1.25 / 1.75	71.4%
Seeds only · S1	0.8	0.8 / 0.8	100%
One hop · S1	0.8 + 0.6	0.8 / 1.4	57.1%
Two hops · S1	0.8 + 0.6 + 0.4	1.2 / 1.8	66.7%
Two hops · S1 × half per hop	0.8 + 0.3 + 0.1	0.9 / 1.2	75%

The denominator uses **selected node weight**, not tags or edges. At two hops the node fraction is always **2 / 3** here, even when the weighted percentage differs.

What half per hop changes

Multiply by **1** at a seed, **0.5** at one hop, and **0.25** at two hops. The final row applies these multipliers to S1: 0.8 stays 0.8; 0.6 becomes 0.3; 0.4 becomes 0.1.

S1's retained base weight combines selection frequency across question phrasings with the fraction of non-hub tags unique to the node. The weights above are illustrative.

What the result tells us

Discovery includes returned snippets, graph results and source-page reads. **Returned read/view coverage** only credits confirmed physical source-page exposure. Compare models under the same definition.

Seeds-only reaches 100% because B and C are outside its denominator. That does not establish a complete answer. Input defects and incomplete receipts must stay visible; one run does not establish stability or answer quality.