


Illustrative decomposition of a conversational assistant; not a claim about internal product architecture.

1. Graph view - organize the loops



Edges and shapes

- 01 -> 02 task brief
- 02 -> 03 evidence bundle
- 03 -> 04 draft response
- 04 -> 05 approved response

2. Loop view - zoom into one node

LOOP 03 - Compose Answer

Accepts:	task brief + evidence bundle
Performs:	create one response draft
Returns:	response {text, citations, attachments}
MO:	general reasoning model
TO:	calculator or code tool when needed
M:	conversation + retrieved knowledge
O:	receive from Loops 01/02; pass to Loop 04
CO:	private data handling; allowed tools; time limit

3. Questions to ask

- [] Could any node be split because it performs more than one job?
- [] Can every receiving loop validate the shape it receives?
- [] Where can a human stop, approve, correct, or resume the work?