

Positioning and why

- What are the Content Agent capabilities we can talk about and show publicly?
 - Ans: Citations, guided workflow, multiple content types, can pull in context from external integrations, etc
- Why does Content Agent write a brief before it writes anything, even when nobody asked for one? What is the reasoning behind making the brief a required artifact rather than an optional one?
 - Ans: We think of brief as a way to gather user intent before the guided workflow begins. To ensure everything is captured.
- Why can't Content Agent run on a schedule the way Research Agent can? Is that a deliberate position or a current limitation?
 - Ans: It was a deliberate choice to ensure content was generated with human guidance.. so it captures original thinking, POVs etc. Some content types are candidates for full autonomy, say we delegate the content agent to update a product name from X to Y or update logos etc. But content that captures original thinking needs to be guided by humans. We are deliberate with our autonomy levers - users can choose Guided vs. Autopilot mode. We can publish about this autonomy slider as team is currently working on this and it should be on prod soon enough.
- Why does the agent verify competitor claims during the brief run rather than checking the draft afterward? (the "Exclude Competitor content" toggle button)
 - Ans: Because it clearly lets the agent know whether this content can have competitor claims or not. Otherwise there is no way for the agent to know that. Hence, it needs to be set prior.
- If a competitor wanted to build this, what would be hard to copy? What in our architecture makes it possible?
 - Ans: The editor is complex to build, and tuning an agent to it is hard. Its doable, but very tricky, sensitive and time consuming.
- What is Content Agent deliberately not built to do? What guardrails did we put in place that help it produce accurate output?
 - Ans: Not go too out of template, both a boon and bane. It doesn't try to recreate the template completely, just what it needs to modify, like a patch.

How it works

- What does the agent actually do in each of the three plan steps? "Find relevant product context and supporting angles," "verify current competitor updates," "set the product update brief."
 - Ans: Steps are fixed always, its more like a default if nothing else is asked. It usually searches amongst the ideas and knowledge bases, templates, etc and chooses the relevant ones.
- Are those steps fixed, or does the plan change by content type or template?

- Ans: Like I said above, its a default if you ask it do anything else, it will follow that order.
- What does the agent retrieve from the knowledge layer on its own, versus what the user attaches?
 - Ans: If users have attached knowledge, it doesnt search. Only when not provided its searches.
- Why three files and three knowledge sources at the brief stage? What breaks above that?
 - Ans: Too much context can cause weird behaviours not fully understood, so best to not have too much. Dumping context is worse than having no context.
- Are the brief's keyword selection limits the same as Copilot's one primary and five secondary?
 - Yes
- How does the brief carry through to the outline and the draft? Is it re-read at each stage or compiled once?
 - Ans: A mix, as in we have already searched for knowledge, ideas, etc. The outline works with the brand ontology mostly, the draft just implements after re reading any visuals, or referenced assets. Since compilation isn't something we can have when working with LLMs, so we have our workarounds as mentioned above.
- Can a user regenerate from the outline without redoing the brief, and what is preserved?
 - Ans: Yes, you can go back to any stage and regenerate the next stage. Currently, it wont be preserved as you are regenerating it.
- Which model runs each stage, is it selectable per piece, and what are the current selector contents? Release notes show Claude Opus 5 and GPT-5.4 moving in three weeks.
 - Ans: To move between stages, we can choose models from the selectors. The copilot is fixed, we have gpt-5.6-terra at the moment. We can choose Opus also, but its pricier.

Citations and grounding

- How does a citation get attached to a claim? At generation time, or matched afterwards?
 - Ans: As of now its generation time, but we can occasionally as the copilot to modify an existing citation.
- What can be cited: public URLs, uploaded documents, Fathom recordings, Research Agent signal, all of it?
 - Ans: All of it, but not sure what the research agent signal means here. We can point to external urls, or internal knowledge base items, so that should cover all the cases.
- Do citations survive into WordPress, HubSpot, LinkedIn, and every export format?
 - Ans: Case dependent, not all platforms support it like we do.
- What happens when the agent wants to make a claim it has no source for? Does it drop the claim, flag it, or write it uncited?

- Ans: It is instructed to drop it. Even within claims, it tried to find a public url, if it cant, it falls back to the knowledge item it came from.
- Does Exclude Competitor Content filter the sources the agent searches, the output, or both?
 - Ans: It filters the sources primarily, but it trickles down into the output as well.

Templates and BrandOS

- What is a template technically, and what does it control: structure, length, tone, all three?
 - Ans: It can control all three, but mostly used for 1,2.
- What is the difference between Global and Workspace templates, and who can author them?
 - Ans: A workspace template is only visible for selection for the users of the workspace. Global templates are available across all workspaces. As of now only Magi team can do so, but once this is ready to expose to customers, we shall do that.
- How does BrandOS reach the model: injected wholesale, retrieved selectively, or compiled into rules?
 - Ans: Answered previously.
- Are the 12 default legal and compliance rules enforced during generation, or checked afterwards?
 - Ans: During generation as well, and also flagged during brand audit.
- How do BrandOS images get generated, and what constrains them to brand?
 - Ans: We use models like gpt-image-2, our brand ontology has instructions and guidelines for all cases.

Versioning

- What creates a new version of a draft (what qualifies as V2 exactly?), and what is preserved between versions?
 - Ans: A change by a new user, withing a window say 10 min, where other users are not editing. If a different user in the same window, its always a new version. Cannot have multiple users in the same version.
- Is the outline versioned separately from the brief? What about the draft?
 - Ans: Outlines dont have versioning yet, is on the roadmap. Only draft for now.
- Can someone open an earlier version, compare two, or roll back? What does the user actually see?
 - Ans: Yes. The see the preview of the old version. When we do restore, it goes to the top as a new version, so no versions are ever lost.
- When a brief is edited after an outline or draft exists, what happens downstream? Is the draft flagged as stale, regenerated, or left alone?
 - Ans: Changes dont trickle downstream yet, its left alone for now.
- Does the version history persist after publishing, and does it travel with the piece into exports?

- Yes it persists, it is decoupled from that. Any export happens on a particular version only.
- Who can see version history: the author, anyone with access, admins?
 - Ans: Author, and anyone with access at this point. Admins obviously.

Live collaboration

- What does live collaboration cover on a Content Agent piece: the draft only, or the brief and outline too?
 - Ans: Only draft for now
- Can two people work on the same brief at once, and what happens if both edit the same field?
 - Ans: Usually wont happen, but whoever is slightly earlier gets the first version, the other the next.
- Does the document know who wrote or changed what, and is that attribution visible to other people or only in history?
 - Ans: Its visible in the version history only.
- What happens if someone edits the draft while the agent is generating? Is the document locked, queued, or merged?
 - Ans: Its decoupled, unless its a major change, where that element is itself deleted, it will work. We shall discover bugs soon, its only been 1 day or release.
- Is the Review gate on the brief a single-person action or can it be assigned? Does anyone get notified?
 - Ans: We dont have this gate yet. No notifications yet.
- Does the “Share” button on the brief panel give view or edit access, and does it work for people outside the workspace?
 - Ans: Not right now, content can only be shared via email to the user who is already a member of the workspace

Learning

- Does AIQ learn from anything in the Content Agent flow, and is that learning workspace-wide or per-user?
 - Ans: It learns from explicit user feedback at the moment, as well as implicit patterns we try to find periodically. It starts per user, but common feedback / patterns are moved into workspace AIQ occasionally.
- Where is the draft-edit learning beta now, and when does it ship?
 - Ans: Not developed yet.
- Does the Content Agent learn anything from an edited or rejected brief?
 - Ans: Not yet.

Integrations

- Is the GitHub integration available to customers, and what would a marketing team use it for? Integrations still carries a Beta badge and every repo listed is ours.

- Ans: Yes available. Things like product release notes, or sometimes repos are for meeting recordings, we can bring that context into magi easily. Yes integrations are beta. Only in our workspace the repos would be ours, because that is what we connected it to. Customers would connect it to their own github, and will see their own workspaces.
- What is the full current list of knowledge-side integrations that are live? Search Console appeared in a product prompt and is not live.
 - Ans: The integrations page on the app is the source of truth here. We have slack, github, notion, google drive, fathom meeting recordings
- What is Slack's actual status and nature as an integration? Product truth says partial.
 - Ans: Yes, its only for knowledge yet, we have can make it more like github / notion soon. These are read only integrations, later we can have write to slack as well, once headless magi orchestrator agent is developed.

Boundaries

- What is Content Agent's failure mode? When it is wrong, how is it wrong, and how would a user notice?
 - Ans: Anything that is out of template it struggles with. We need to have a template for it now. Making it not needing templates always will be the biggest unlock.
- What is the longest or most complex piece it handles well, and where does quality drop?
 - Ans: Anything where context is missing it does not do well, the usual ai slop starts showing up.
- Anything shipping in the next four weeks that would date this post?
 - Ans: We didn't have cancelation of a request ability, we have that on staging now. Not really content agent, but its important UX. More work on content quality is wip. Auto progress to draft stage instead of always navigating brief->outline->draft is also WIP.

Origin

- What was the version of Content Agent you built first, and what changed after people used it?
 - Ans: We had an older version of the editor built using [Craft.js](#) which we scrapped to move to the new tiptap based editor. The content agent had to be rebuilt in both cases. So yes, the older version was built for this old editor.