

# What happens to internal comms when AI starts writing the news

Tom McIlroy and Sean Zintl, Sage

---

Tom McIlroy and Sean Zintl, take it away. Brilliant. Thanks, Cai. Thanks, Cai. Yeah, absolutely. We're going to go through this probably at a slightly faster pace than we anticipate, just because we're a FTSE 100 company and we've got results tomorrow, so naturally meeting rooms are a bit in demand. But thank you, Cai, for inviting us along, and it's a real pleasure to talk to all of you, not just the SWOOP team, everyone who's joined today. A bit of context before we do introductions. Sage is a FTSE 100 company in software development for accountancy, payroll, platforms and things like that.

And we have about eleven and a half thousand colleagues across many countries. Twenty three is the current count. So we are dealing with a lot of the interesting headwinds and opportunities that AI and the world itself are presenting right now. But in terms of us, I'll hand over to Sean to introduce himself very quickly. Yeah, thanks. So I manage our digital channels, the intranet being one of them. Great that Tom laid out our organisation. And one of the challenges we have at Sage is a very decentralised approach to communications and product development and marketing across those regions.

So it's been really inspiring to hear what a lot of you have been talking about. Governance. It's the drum I keep beating. As you all know, it's very difficult to implement, possibly even more so at Sage. And that's why we are very interested in AI and what it can do for us. Brilliant. Sean's right at the heart of not just our channel strategy, but making sure people use the tools and channels that we've got. Effectively, not just in our comms team, but more broadly across the business in every country.

For my part, I'm involved in more of the data and analytics side, telling the story about what these platforms, SWOOP being one of them, tell us about how people are using our various channels and reading our various comms. And then also driving more tech transformation across the business. We just rolled out the 365 Copilot to every one of our colleagues after initial 4000 person pilots. So we have a fully AI driven business now. So naturally, we are right in the thick of what happens when you give colleagues powerful AI tools. But the infrastructure, as Jarrod was mentioning earlier, may not be up to snuff.

And we're currently trying to tackle those questions. So what we're going to cover over the next probably 15 minutes is a demo of a content agent we built, automating a very similar process in our comms team, talking about what feeds into that, talking about some of the obstacles that we see personally coming over the horizon, and then leading over to Sean to talk a little bit about how we're going to prepare for those challenges by really setting up the

governance and thinking about how we want Sage colleagues to use SharePoint and the various channels we've got in place.

Like any smart AI tool, I'm going to cite Jarrod as offering a huge amount of context probably for what we're going to talk about. So if we do miss anything, Jarrod probably covered it quite effectively. So I will point you to his presentation, because it's a really well thought through and detailed one. So thank you for that. And then a final sort of takeaway for you to kind of think about how you roll out governance in your Teams, because we appreciate not everyone is or everyone is at different stages of their own journey.

So, as I said, we're in a significant rollout of AI across our business. Every Sage colleague now has access to the most advanced Microsoft Copilot tools, which allows them to do everything you can do on ChatGPT and to some extent, Claude with Anthropic. So we are very much have 2026 colleagues working on what is probably like a 2010 framework. That's not a sign of how outdated our channels are, but SharePoint is very much the context model for a lot of these AI things. And I think every company is rushing now to figure out how you power these AI tools effectively.

And governance is the big part of that. But for us, where we're encouraged to think about AI across Sage, our internal comms team was thinking about tasks we could automate and get AI to power that took up, that used time typically in our Teams. And one of the big ones that came up was content formatting. If you've got background in terms of content, you know that often you have to turn a SharePoint article into a newsletter, into a Viva Engage post, and often that requires manual rewriting. But probably one of the more basic features I can call now with AI is content production.

It seems so commonplace now, but really, probably two years ago, it was magic sauce and everyone was amazed by it. So we thought about creating an agent around that. But to make that reliable, if you've ever used ChatGPT or Claude or any of the other models, just kind of raw to produce content. You know, they go off on big tangents when they don't have that context. And so one of the big things we wanted to do was make sure that the agent had the context, which basically was us pouring decades of experience into its rule set.

So how should they approach different content structures, the length, the clarity, the tone of voice, using exemplars that we've already got in the business that humans have produced. And we think that's the new opportunity for a lot of internal content professionals, is owning the structure and the platform there as much as producing content that tells a story. So I'm going to do a really swift transition to the agent in question. You can tell we're an eight billion pound market cap business with that one. But this is the agent creator.

I will give a coffee to anyone who comes up with a really snazzy name for it. But this is designed on the three six five Copilot platform and basically takes whatever you put into it and then rolls a series of prompts. You can see here at the bottom of the screen to help format a lot of the content. And it can take things from very basic information, very succinct information or opaque presentations in PowerPoint. So to give you a live demonstration.

And as much as I've rehearsed this, you never know quite how it's going to react.

But this is a very basic update that we shared actually publicly. We've got a new chief products officer in the business and a new chief strategy officer. This is all public. So don't worry, I'm not revealing any secret confidential information. But this might be the sort of example of a comms that you have to produce your leadership team and then have to format into multiple channels. So I'm just going to take that very succinct prompts.

I'm going to select the first one, which is create first drafts for all of our major content channels, which is Viva Engage, a managing newsletter, a colleague newsletter and a SharePoint article. We call our SharePoint Your Sage. I'm just going to post it in there and let it do its thing. I should say what it's thinking, I would love if Copilot showed you an interesting fact while this was all kind of processing just to kind of use up the time. Exactly. Or something to kind of or an animation like a little person running. You never quite know how long it's going to take.

This is also a slightly dumbed down version of something we have, because one of the powerful parts of the agent is that it can access resources on SharePoint. But I've kind of stripped that back a little bit so that nothing too confidential gets shared. And any second now. It should produce all of the content. I am using a chatty thing, longer model just to give it a bit of extra context. Hopefully any second. If it doesn't, I did have one I prepared earlier, so I might shift to that. It doesn't do it in the next five seconds. I think it's being a bit shy.

This is the great thing with demos. It works perfectly every time you rehearse. And then she asked, did you try to get things rolling? It could equally be trolling. Right. Here we go. Here's one I produced earlier. Right. So you can see here I use exactly the same prompts earlier and then it produced the managed newsletter. Off the top, key things here are tone of voice and structure. So with managed newsletters, we want to make sure the managers are thinking about what to share with their Teams, which they it talks about here.

We have the colleague newsletter, which kind of says the same thing without that manager nuance. We have a Your Sage article which talks about the information it can glean. It kind of infers things about, you know, what the role of the chief product officer is. Very basic stuff. You will see here as well. It prompts us to confirm things, the remit of certain people in this instance, what products they might be working on. And at the bottom, we also get a summary of any references that have been used, which aren't in this case. And then Viva Engage posts.

The idea is to get a kind of almost ready bit of content ready to go that users can then shape. You'll see here. I was going to ask it to make the newsletter tone a little bit warmer and to make the Viva Engage post come from our chief people officer, Amanda Custin, and put it in the personal tone of voice. So here you can see very quickly that it does a little bit more of that or brings a little bit more of that warmth into the post. It's not as cold and objective about the news. And then the Viva Engage post starts with I'm pleased.

We could send that to Amanda right now. She could post it. It saved our team maybe 30 minutes of writing this. So just to quickly show you what's behind this, because there is a lot of context behind the scenes, this all ties into that conversation about governance. You know, at the end of the day, AI is a really powerful tool, but it's only really powerful if you've got really strong data powering it, rules powering it. You only get strong data rules when you have strong governance. So I do think the gold rush for AI is soon going to move to people who can structure this.

But here you've got the default rules and everything from the length that it should be to the tone of voices. There is a separate rule about no em dashes. Unfortunately, em dash I think is going to die out probably in the next six months if things continue this way. Two to three short paragraphs for something like a managed newsletter. So you get the idea there. Then it's got the context examples here. So, for example, we've engaged posts in your post that humans created that we used examples. And then the very strict prompts that we use designed around the workflows we have.

So that is the agent. I'm just going to buzz through what I had to say after that so we can get to Sean. It's producing almost channel ready content or channel ready content that can be adapted, saving time there. What's powering it? The instructions, the context, the examples, the defined rules, everything that people in internal comms learn over the course of their career. So it's about sharing knowledge with the AI, not having us replace our knowledge with the AI. But what it doesn't also replace is editorial judgment. Teams can totally adapt this. Chat AI is supposed to be conversed with. So it's a starting point.

It doesn't get you from zero to 10. It gets you from zero to six. And the team takes it the rest of the way. Very briefly here. The reason governance is so important, if I can say a TLDR, is that the era of subsidised AI models will end at some point. Companies that are funded by venture capital will need to not only generate revenue, they'll need to generate profits. And we're already seeing that in stories. Companies are tracking AI tokens. We are seeing companies running out of budget with their AI in months rather than a year.

And soon we're going to see that agents that don't have restrictions and don't kind of just wander around for context are not just going to be an editorial issue. They are going to be a budget issue at the end of the day. Tokenisation is going to be the currency of these AI models. It's never been cheaper to experiment. I would recommend anyone experiment while. Claude is like 20 quid a month because it's not going to stay 20 quid a month for the basic plan. And companies are going to learn very quickly that there are much higher costs when these companies want to be profitable.

So don't build powerful agents that can start with a prompt and build you a conspiring with the comms. Structure them around very specific workplace and join those together. And these are examples of what we're already seeing in the last few weeks. NVIDIA's vice president of applied deep learning is already talked about the fact compute power outweighs the cost of employees doing the same tasks. Microsoft has already started to subtly move some of its developer tools, its token based pricing. So this is coming. Governance is the

solution. So think about that. It all comes together.

This is a budget issue as much as a kind of IT issue at the end of the day. It will be a big budget issue. Sure. Yeah. So, Tom, thanks immensely for that. So you can see what. For internal communicators, this is quite scary because, you know, the team that I work in. Outcomes communicators, I have a background in journalism, spent many years analyzing how to craft the perfect intro. And suddenly with one fell SWOOP or a few clicks of the keyboard, those skills seem to be almost being made redundant. So, you know, what does this look like for communicators going into the future?

And what I'm trying to get our team to think about is it's less about presenting information and helping surface information. So what I mean by that is, you know, to Tom's point, they would craft an article, an email, an entry. They would spend a lot of time formatting it on the various channels and then they would present it. What we're really seeing now with Copilot coming in is that presentation level is disappearing and you don't know where colleagues are picking up information. It's simply being surfaced.

So, you know, many of the presenters on this call have spoken at length about governance and absolutely inspiring the level of governance that you have implemented. This is certainly the journey that we are on and trying to enforce this across. So I mentioned earlier, we have a very decentralized system at Sage. We have regional Teams. We have global Teams. We work in multiple different languages. You know, each one of those Teams with their own pinch points, if you like. In my team, we have about six, eight, I think, at the moment, business partners that partner with different parts of the business.

And what I'm urging them to think about is instead of spending more time writing content for those business areas, that they partner with them to make sure that their content is there so that we can use these agents to very quickly extract it. Sorry, I'm not going to go through all of this. I know it's very wordy, but I hope you get the gist of that. So, yeah. I think it's a non-starter to say everybody recognizes this. If we can have this information in a place that is machine readable and very quickly accessible, we can pick it up. The other part of this is translation.

Now, we've experimented with various translation tools. We have embedded into SharePoint. We use Microsoft Translator. Our comms Teams hate it. They ask us, please don't use it. We've used DeepL, much better. But for instance, in German, for any German speakers out there, they prefer the informal version, but plural form of address. Now, to try and explain that to an English speaker who doesn't speak German is like, what? So we just need to get all of these prompts and build them in and talk to IT to see if we can then input copilots and make the default translator.

But with all of this, again, it's the metadata that we put around a lot of this information. Fortunately, we do have quite a robust policy section that is managed by two Teams. All the global policies are managed by our legal team, so very strict version control. All of our people policies are managed by the people team. But at the moment, as a colleague, if you

want to find out what is applicable to you, you need to go to the global policies, look for something there. Then you need to go to the HR team to see what's applicable for you there.

Wouldn't it be great if we could just draw that and contextualize that for the individual? So that is another area that we're really looking to explore with AI. Is that hyper-personalization or contextualization or what we call segmentation? AI can do this at a level that we can't even begin to think about doing manually. Voila. With all of this, trust is the word that keeps coming up. How do we build trust in the outputs? How do we build trust in the accuracy of the information? Fortunately, Sage as an organization, as Tom explained earlier, we support small businesses. We operate in multiple countries.

In each one of those countries, we need to be 100% compliant with any legal obligations. Tax, whatever the legal framework in that country is. We've been embedding AI into our products for the longest time, but they have to be 100% accurate. It's taking that mindset and reapplying that to our colleague communications Teams to say, listen, all of those skills that you've learned, in my case as a journalist, verify things to sources. Is it really genuinely accurate? You need to be looking at this content. And then, sorry, just going back to that one. What Tom is particularly driving from a metrics perspective is what we're recognizing.

We are moving away from unique views, visitors, because as information is serviced through Copilot, you can't track that as effectively as previously. We're having to go back to source, maybe look at the citations. How often was something referenced through an AI query? And then really shifting the question away from how many views did this have to questions like, did colleagues understand the strategy or has this led to a different outcome? And so, yeah, fundamental head shifts. Mentioned this, citations. Thank you. I think this probably came from some of the SWOOP reference material. News is becoming the thin layer, knowledge is becoming the thick layer.

So again, it's that shift in mindset from communicators presenting information where colleagues are just actively, in many cases, ignoring it or tuning out. And this is not just internally at Sage. It's a worldwide phenomenon, right? I don't know how many of you guys have actively stopped reading news. I certainly used to be a doomscroller first thing in the morning. I've stopped doing that just for my own mental health. And many colleagues, that is something, there's just too much information. So it's really being selective about what you engage with, how you engage with it, the way you engage with it.

So, again, for communicators, this is really a tough one to wrap their heads around. Because with an intranet, what you do have is the opportunity to engage colleagues with something that they might not have known about. So as they're coming in to do something, they'll see the news articles. Oh, look at this. But if they're not looking at it, then that is going to become more of a challenge. And, yeah, I think a lot of this has been covered by many of the previous presentations.

But it's just re-emphasizing that idea that AI is not necessarily going to supplant human work, much as we are dancing around this issue. But it is certainly going to amplify it. And all of those judgment calls that humans are capable of need to be brought in ever more. Thank you very much. Yeah, just basically, Sean and Tom, it's been a really thought-provoking presentation. And I think this whole, like, what's going to happen with AI and the internal comms profession is obviously something I think many of us are thinking about. There are some questions I want to make sure that we got to.

Are you OK if I ask you some of the questions? Go for it. Because there's something like, obviously, people are very interested in what you're doing. And there are some at a very practical level. So, would it be possible to get a copy of those prompts? I think people are interested in that. So, Tom, I'm not sure that's a question for you, but I'm sure for you as well. But, you know, is it possible to share the instructions with the community? Or is it something I think people understand if it's confidential. But I think there's an eagerness to say, oh, can we try that out?

Oh, absolutely. Look, I think prompting is half the art of using AI. In a separate work stream, I've actually just taken on reaching out across the entire corporate affairs function to find out who's using AI and where. Because I think everybody is reinventing the wheel in dark little pockets. So, prompting library is huge. Just to say, well, if you're doing content creation, these are the type of prompts you would use. I mentioned the one about translations. So, it really depends on the tone of voice, what the output is.

For any colleagues in the UK, I don't know if you've come across the GCSE model for prompting AI. I don't know. It stands for Goal Contest. GCSE. God, yeah, yeah, yeah. But essentially what it comes down to is you need to prompt the AI to tell it what it is that you're trying to achieve. So, you need to give it a certain amount of context. You need to outline the goal. And then also qualify that. Well, that's why they've used the goals for context expectations. Thank you, Emily. And that's exactly why they've used that acronym in the UK.

I think, though, on the point, if you're wanting that prompt that powers 360, I'll have one. I'm happy to share it. I'll strip it out. Just bear in mind there are some nuances. I'll make sure we share the prompt with you so you can just drop that into the back end of whatever you use. And just before we go to the panel session, just also a very quick question. I think it was Rachel who was saying that she found that sometimes asking Copilot directly get better answers than when you're coding it from within the agent. Is that something you've noticed that there's a difference?

Are you happy with the answers you're getting from the agent? I think we're comfortable with the answers. I mean, there's always kind of like what I think is good codes might slightly be different from someone else's. There's always a personal element to it. But I would say if you find there's a pattern, then obviously bring that into the workflow. But generally speaking, as long as the AI has very rigid rules, you probably didn't glean it from the very succinct part I showed you. But I'm very strict with it. Like I'm a very nasty parent to the

agent.

So I would say just be harsh with it behind the scenes. And it should generally behave itself. But also, there's new language models all the time. And I have to admit, just moving from like Copilot, which is the broad one to ChatGPT's ones, there are variances. So the tech's always moving as well, which doesn't help. But I have noticed it's certainly gotten better from the time that we've been using it, particularly with the agent. And so if you've not built an agent before, you can actually give it a set of instructions.

So Tom was pointing behind the scenes, you actually prime it to say, I want you to always do this. Once you've got those instructions, you then go to the front end. And that's where your prompt and your source will always be different, depending on the output that you require. But any other questions? I mean, naturally, I appreciate you've got to reach your panel, but you've got our emails and the presentation. Find us on LinkedIn. Happy to answer your questions and share the props with you. [swoopanalytics.com](http://swoopanalytics.com)