

At Armory we have a number of experts that are answering questions every day in a number of private formats, including:

Jira Tickets

Slack Threads

Salesforce Tickets

Dropbox Paper Documents

Intercom Chats

I really wish there was an automated system that we could give access to our raw data sources that could consume all of these private answers and automatically turn them into sanitized public content. Specifically, the goal would be to bolster our credibility in Google's eyes via SEO that's generated from our internal data exhaust from efforts we're already undertaking internally. A huge secondary benefit would be that we can give better answers to people who are asking questions — or keep them from having to ask questions in the first place.

Here are some specific examples of use cases:

We are members of several Slack communities that have between 6k to 50k members each. We spend a fair amount of time answering technical questions about how Spinnaker (an OSS project) works in the general channel of these Slack communities. We are often answering the same questions over and over again — but with small variations each time. Ideally, some (AI-based?) system would categorize all our answers and generate some kind of public content (say in a knowledge base) such that we could just reference that answer the next time someone asks. A huge bonus here would be some Slack plugin that helps our engineers query the system before answering. So the next time we get a question, about, say "RBAC controls in Spinnaker" in the Slack channel, an engineer could type /bot RBAC or some similar keyword and get presented with a list of existing URLs that have detailed answers, and choose to just post one of those URLs as the answer, and/or write some more detail in (and if s/he wrote in more detail, the system would auto-update the public answer w/ an updated version of that detail.)

We also have private Slack channels for our prospects & customers where we are answering the same kinds of questions over and over again. We do this over and over again in a couple hundred private Slack channels. It just drives me crazy that we aren't able to get this out into the public realm so Google could index it, and to help people to be more self-service about their interactions w/ Armory. I would LOVE to have a bot that ingests all these conversations, and then intelligently sanitizes them to remove anything that could be customer-specific, and then posts them to our KB at <http://kb.armory.io> ... right now we have a channel where we intend to publish them to the KB someday but it's not automatic.

One idea would be to have a "tinder for content" type experience, where an AI bot based system could vacuum up all private content (from Jira tickets... Intercom... Slack... etc. — be happy to give it full access to all our data exhaust) and then write up a KB post and post to Slack I've created a new KB post for you based on this Slack conversation you had: [URL]. If it looks ok to publish, just respond with "Ship It!". If it needs to be edited in any way, just respond with "Edit" and I'll take you to a page to make changes.

- Your friendly AI ContentBot Same workflow w/updating existing content based on new information. (our KB is powered by Jekyll Github pages, so it could just issue a PR when someone says "Ship It!") As a Version 2.0 feature, I'd love to have this bot integrate w/ Intercom, so when people ask us questions, it could respond w/ the best answers from our KB. (That shouldn't be too hard to

integrate; we could just use Slack as the mechanism if required, since humans already respond to Intercom messages via their Slack integration).

The main headline here is this: Right now, content is living in private venues and not being promoted to public consumption, and it's incredibly expensive for us as a business from an opportunity cost perspective not to have the amazing knowledge living in the heads of our engineers / customer success department / sales engineers ... getting out into the world in frictionless, sanitized ways. As a CS Leader, this just drive me up the wall. That we are already doing 95% of the work, and don't have an easy way to apply a tool that would 10x to 100x the benefit we could get from that "last content mile" being locked away in private conversations.

I'm 100% sure that we are missing out on 6 and even 7-figure contracts by having this information locked away currently, because we're in a fast-emerging market (cloud-native software delivery) where companies need expertise that few people have (and that we have).

I would happily pay \$200/lead for leads that come in from a system like this (that's our average today via SDRs) from our ICPs (Global 2K Enterprises w/ a hybrid cloud strategy). But I'd probably pay much more. I could absolutely see us paying \$500/month+ for something like this if it worked (inbound leads are more qualified). Honestly, probably more because there's no good way to do this today w/o requiring the time of very expensive people (\$260k annual engineers, for example).