


MERLIN CARTER



Staff Technical Writer — Developer Content, Applied AI, Open-Source Documentation

Berlin & Chiang Mai · mcarter@gmail.com · [linkedin.com/in/merlin-carter-bb303ab2](https://www.linkedin.com/in/merlin-carter-bb303ab2)

github.com/merlin-quix · github.com/merlin-lacuna

Summary

I'm a versatile content creator with over 15 years' experience in technical communication. I've worked on a variety of technical documentation projects, as well as technical whitepapers and other sales collateral. I like to approach writing projects like an investigative reporter, probing subject matter experts for useful nuggets of information. Most recently, I've interviewed customers using the "Jobs to be done" methodology to understand what content they might appreciate.

I tend to treat technical documentation as very particular form of product testing. I try out the features I'm documenting, push their limits, and look for the things that trip readers up. This means that I often end up giving product managers proactive feedback on what could be better, alongside writing the docs. Since most of my experience is in documenting developer tools, I'm very used to the docs-as-code approach and can do a bit of Python scripting when I want to automate something tedious.

I've had a personal workshop with two of the authors of IBM's "Developing Quality Technical Information: A Handbook for Writers and Editors" (Michelle Carey and Elizabeth Wilde) enabled me to take a disciplined approach to information architecture, especially according to the DITA guidelines. However, I'm not a DITA purist and don't do much DITA these days but it has given me a good sense for topic-based authoring.

Selected Work

Fine-tuning an SLM for on-device industrial RAG (March 2026)

<https://www.distillabs.ai/blog/how-slms-can-enable-on-device-rag--making-industrial-machinery-more-usable/>

I pitched and wrote an article explaining how Small Language Models (SLM) could be useful for question answering in the security conscious industrial domain. In the process of testing it, I uncovered and reported several product issues.

Klaus Kode — an agentic data integrator built on Claude Code (September 2025)

<https://github.com/quixio/klaus-kode-agentic-integrator>

Designed, built, and documented an open-source agentic tool for building streaming data pipelines using the Quix platform. I was trying to encourage my main customer, Quix, to incorporate AI-coding assistant into their platform (since their users were mechanical engineers rather than software engineers). This was an early prototype that was built on top of the Claude Code SDK. I

built it with the help of Claude Code. Quix has since started incorporating AI agents directly into the web UI of <https://portal.cloud.quix.io/>.

Guide to the Event-Driven, Event Streaming Stack (August 2023)

https://drive.google.com/file/d/1wV-3QgJBODGqWCqczr_dRXv-n32lUutL/view

I wrote a long-form technical guide covering the full EDA component landscape (event buses, pub/sub, message queues, event streaming, stream processing, event stores), with cloud-native and open-source paths for each component. My the idea was to provide objective value to readers without pushing the Quix platform specifically.

What's the best way to manage Helm charts? (May 2020)

- *Full PDF:*
https://drive.google.com/file/d/1mW9q3xNbiLEwU49Ze7UTpv5ZdnZJQ6Eb/view?usp=drive_link
- *Public Version still live on Medium.com (members only):*
<https://medium.com/@Project-A.Tech.Insights/whats-the-best-way-to-manage-helm-charts-1cbf2614ec40>

Best practices article proposed by my team lead. It stemmed from debate at one of the portfolio companies. I interviewed the affected developers and I identified the best practice was just as much of a human' issue as it was a process issue. The article accrued in the first 2 weeks after it was published so it seemed to touch a nerve at the time.

Selected Open-Source Documentation Contributions

The contributions were all made when my main customer, Quix, was trying to forge relationships in the open source ecosystem until this approach was abandoned by the founder. I gently deflected his urges to aggressively self-promote and encouraged him to let me try and provide legitimate value to the affected communities.

LangChain — *langchain-ai/langchain #16771*

A cookbook notebook demonstrating how to manage LangChain chat messages via Apache Kafka, proposing an architecture where the chat frontend and LLM run as separate services communicating over an internal network (a deliberate alternative to the typical REST API pattern).

InfluxDB — *influxdata/docs-v2 #5342*

A downsampling tutorial for InfluxDB Cloud Serverless using Python tumbling windows, positioned as a pure-Python alternative to the existing SQL/InfluxQL approach and as a Pythonic replacement for the Kapacitor+Flux/TICKscript toolchain. Merged into InfluxDB's own docs site.

Upstash — *upstash/upstash-web #229*

A full technical blog post on building a CDC pipeline from PostgreSQL to Upstash Vector via Upstash Kafka and Python. It covered change data capture, real-time streaming, and vector database upserts for AI applications.

Aiven — *aiven/aiven-docs #173*

Step-by-step integration guide for connecting Aiven Kafka as a broker to Quix, with screenshots and sidebar navigation updates.

Upstash — *upstash/docs #151*

Integration guide for using Upstash Kafka as a broker with Quix, with screenshots and navigation config.

Talks, Video, and Multimedia

Knowledgebase Ninjas Podcast

Participated in an early episode of this vendor run podcast (Document360) giving my views on how technical documentation should not be a siloed function.

<https://podcasts.apple.com/us/podcast/interconnectivity-great-documentation-with-merlin/id1496613504?i=1000483484937>

Powering Vector Search with Real-Time Data (Quix) — *~April 2025*

<https://www.youtube.com/watch?v=9pG1pspgELU>

Co-presented with Zain Hasan (Developer Advocate, Weaviate), hosted by Kris Jenkins (Developer Voices). Live RAG and vector search using real-time transcription into Weaviate, with a companion CDC-to-Weaviate code template repo built alongside the talk.

Build an industrial data pipeline using AI and MCP servers (Quix) — *~June 2025*

<https://www.youtube.com/watch?v=Y6QBakE9zB8>

Early public experiment with Model Context Protocol for agentic construction of industrial data pipelines — published during the first wave of MCP adoption.

Watch an AI build an industrial data pipeline from a sketch (Quix) — *~June 2025*

https://www.youtube.com/watch?v=P9vVQPJ_bRU

Demo of agentic pipeline construction from a sketch as input. The experimental work that led directly to the development of Klaus Kode.

Using AI to map OPC UA data by making it read a signal manual (Quix) — *~June 2025*

<https://www.youtube.com/watch?v=uwGAgv-KWwo>

Demo of LLM-assisted parsing of an industrial signal reference manual to auto-map OPC UA data streams. I explored a similar problem pattern later in the Distil Labs article on fine-tuning SLMs for industrial documentation RAG.

Real-Time Python Aggregations with WarpStream (Quix) — *August 2024*

<https://www.youtube.com/watch?v=RuiweYsnkt0>

A walkthrough of the Quix integration with WarpStream. This was part of an ongoing campaign to ally with technical partners (WarpStream has since been purchased by Confluent).

Quix repositioning — industrial launch video (Quix) — April 2025

<https://www.youtube.com/watch?v=kux4Q-56d8M>

Cut and edited the animated launch video for Quix's industrial repositioning. It was designed to be the flagship video that ran on autoplay on the Quix YouTube channel.

Experience

Freelance Technical Content Creator

Mar 2022 – Present

Anchor client: Quix (Python stream processing, event-driven architecture, and — from late 2024 — industrial data processing). Additional clients including Distil Labs and Chainlink Labs.

- **Quix.io:** Primary technical content creator for Quix's developer-facing output: long-form technical guides, end-to-end tutorials, interview-based deep-dives, and explanatory content for data engineers, data scientists, stream-processing developers, and (from 2024 onward) industrial automation engineers.
- Worked directly in the docs-as-code pipeline: Git-based workflow, staging and production environments, Markdown/HTML source, Quix platform documentation site.
- **DistilLabs:** (Mar 2025): ran a controlled fine-tuning experiment for on-device industrial RAG and wrote the accompanying technical article (see above).
- **Chainlink Labs:** Six-month contract with (2022) writing thought-leadership content on building and scaling Web3 startups, for the Startup with Chainlink program.
- **Rows.com** (now part of Superhuman) 3 month contract in 2022 writing user-facing technical documentation.

Senior Technical Content Writer — Project A Ventures

Oct 2019 – Mar 2022

Berlin-based VC with in-house technical experts deployed to portfolio companies.

- Covered developer innovation and new technologies at Project A, a venture capital firm with in-house operational support for early-stage startups. Writing spanned DevOps best practices, front-end development, and data science.
- Held talks on technical writing
- Seconded to Quix to help build their technical content function from the ground up. They became my freelance anchor client after I decided to become self-employed.

Senior Technical Writer — HERE Technologies

Jul 2017 – Sep 2019

Berlin.

- Worked across a large suite of developer-facing products as part of a central documentation practice team supporting writers and contributors throughout the company. Acquired deep understanding of developer audiences at scale.
- Trained product owners on how to create effective technical content.
- Developed terminology and writing guidance for use across teams.

- Edited technical content directly in source code and in technical authoring tools; managed contributions from multiple authors via Gerrit and Git.
- Checked code examples for completeness; streamlined the process of continuous documentation and integration.

Acrolinx

Oct 2006 – Jan 2017

Berlin. Eleven years across three roles in a content-quality enterprise software company (the "Grammarly for enterprises" category).

Senior Content Writer

Oct 2009 – Jan 2017

- Developed user documentation for Acrolinx's enterprise server product and related plug-in software. Worked with developers and product management to synthesise complex technical specifications into clear user-facing content.
- Migrated the entire publishing infrastructure from Author-it to the DITA Open Toolkit. Customised the DITA Open Toolkit to match the Acrolinx design style.
- Worked with DevOps to integrate publishing automation into continuous integration systems.
- Developed guidelines for terminology, style, and tone of voice across the product line.

Product Manager

Oct 2008 – Oct 2009

- Analysed requirements and developed specifications for new product features. Assessed product design and usability. Researched competitor products and academic papers on comparable systems.

Technical Support Analyst

Oct 2006 – Oct 2008

- Provided technical support for enterprise content quality software. Built out the support knowledge base and systems to serve help content more effectively.

[Ask.com](https://www.ask.com)

May 2000 – Dec 2005

Auckland, then London.

Commercial Technologies Manager (2003–2005) · Commercial Operations Coordinator (2002–2003) · Editorial Assistant (2000–2001)

- Worked with development teams to build internal systems for sales forecasting and ad revenue management. Collected requirements from stakeholders, wrote detailed specifications, worked with external partners on bespoke ad-serving software. Earlier, as Editorial Assistant, served as New Zealand category expert for the Ask Jeeves International editorial team.

Tools, Methodologies, and Ecosystems

Documentation DITA, DITA Open Toolkit, reStructuredText, Markdown, MyST, the Good Docs Project ecosystem, Developing Quality Technical Information (Hargis et al.) as a reference for quality criteria and content-type separation.

Workflow Docs-as-code, Git, GitHub/GitLab PR review cycles, staging/production splits, continuous documentation, terminology management, content strategy across teams.

Developer Linux, command line, Python (working knowledge, including for data-cleaning, synthetic data generation, and tutorial code), Apache Kafka and the broader stream processing ecosystem, cloud-native developer tooling.

Applied AI Claude Code SDK (Klaus Kode built on it), Model Context Protocol, fine-tuning small language models for on-device inference (distil labs platform, HuggingFace publication), LLM-as-a-Judge evaluation, synthetic data generation via open-weight models (gpt-oss-120b on Groq), retrieval-augmented generation for domain-specific technical documentation.

Education

Meisterschüler, Fine Art (majored in sculpture) — Universität der Künste Berlin, 2009–2015

Bachelor of Communication Studies (majored in multimedia publishing) — AUT University, Auckland, 1995–1998

Languages

English (native), German (professional working), Thai (B1 level)

Interests

Contemporary art and media art.

- I still help my former art school friends with technical consulting for their art projects
- I regularly attend the Ars Electronica festival

Learning languages, Swing dancing, Buddhism