

Brett Fletcher

Enterprise AI Transformation and Delivery Leader

Oxford, Oxfordshire | brett@bfletcher.co.uk | bfletcher.co.uk | linkedin.com/in/bfletcher-pm

SUMMARY

Most of my decisions come down to the same trade: how much faster we can work, and what that puts at risk. I cancelled a part-built tooling and stack migration in my first year because it was heavier than the company could carry. I set company-wide policy requiring other departments to declare changes before they reached customers. I stopped buying the platform our product knowledge runs on and started building one AI can work on, with the controls to make that safe.

SELECTED ACHIEVEMENTS

- **Current.** Started from the adoption problem rather than the tooling one: how to make VS Code familiar enough to writers that moving to Markdown stopped being a change-management exercise. Built the prototype myself with Claude Code and GitHub Copilot, and the demo did the selling: three people from other teams joined to deliver it. Pilot Q3 2026, production Q4, against a committed \$250,000 annual reduction in company-wide tooling licensing.
- **Current.** Asked to set the knowledge and content approach for a 700-person organisation adopting AI agents in its day-to-day work. Proposed three things and got all three: rename the function to Knowledge Operations because its remit had outgrown technical content, move to Markdown so models can operate on the content directly, and build a governance pipeline for anything a model presents. When a request came to feed the whole Help corpus to a model, I refused it, because nothing existed to control what happened to the content afterwards.
- **2014.** Inherited an incomplete DITA XML migration with content converted, build targets running and oXygen licensed, while documents were still maintained in Word. Cancelled it inside my first year, on my authority as the function owner, and moved to Confluence: already licensed, faster to publish, open to contribution from anyone in the company, and hireable without XML skills. The company was not large enough to carry the complexity. The function later grew to eight without needing a single XML specialist.
- **2014 to 2016.** Fourteen development teams reported customer-facing change by email, direct message and verbal request, and the omissions landed on my team. Made flagging mandatory in Jira, sponsored by the CTO and agreed in the development lead forum: unflagged work was excluded from release notes and the omission stayed with the originating team. Misses fell from two or three per release to near zero.
- **2014 to 2016.** Development teams were changing the interface without telling QA, Documentation or anyone else downstream, so tests were not adjusted, published guidance went stale, and customers met changes nobody had prepared for. Proposed a written company-wide policy requiring planned interface changes to be recorded and reviewed before implementation, with anything that bypassed it liable to be backed out. Put enforcement with the QA lead rather than my own function, so the check did not depend on me.
- Video production depended on an external voiceover supplier, so changing one line meant re-recording the whole thing. Weighed building it internally against buying, and bought, on cost and time to value. Turnaround went from about a week to about an hour, roughly \$1,000 per video of external cost became a flat subscription, and updates moved from full re-recordings to regenerating a single line.

PROFESSIONAL EXPERIENCE

One Network Enterprises / Blue Yonder

Dallas, TX, then Oxford, UK following relocation in 2024

Senior Product Development Project Manager | 2024 to present

Six-person function across 14 product lines and 23 products, with direct reports in India, the United States and Europe. Same reporting line and portfolio through the 2024 acquisition. Nominated to represent the function on the company-wide AI council, and founded a cross-department working group on AI in documentation and communication.

Technical Content Delivery Manager | January 2016 to 2024

Built the function from two people to eight, then held full portfolio coverage after a reduction to six. Selected and implemented a replacement authoring platform.

Principal Technical Writer | January 2014 to January 2016

Sole member of the function for the first eighteen months, reporting directly to the co-founder and CTO.

Earlier Experience

AT&T | Senior Technical Writer | 2013 to 2014

Innodata Isogen | Technical Writer, promoted to Senior Technical Writer | 2008 to 2013, placed on client engagements at Cisco Systems, Blackboard and the American Psychological Association.

EDUCATION

University of Essex. MSc, Artificial Intelligence. In progress, part-time, started January 2026.

Applied Python machine-learning coursework covering predictive-maintenance datasets, model comparison, class imbalance, precision, recall, and business trade-offs.

University of North Texas. BA, English: Technical Writing.

Brookhaven Community College. Certificate, Computer Science and Programming.

SELECTED PROJECTS

Hensington | local-first AI assistant. Built a working voice assistant running entirely on owned hardware: a custom ESP32-S3 voice node, local speech transcription, inference against a locally hosted Gemma model, and persistent memory that retains household context between sessions. Normal conversation never reaches a cloud model. Direct experience of the deployment, latency and data-residency trade-offs that enterprise AI decisions turn on.

TECHNICAL

Python | Git and GitHub | Azure DevOps | CI/CD | Local LLM deployment | Model evaluation and selection | Structured content pipelines | Systems integration and prototyping