

The Real Story Behind AI in Content Editing

CMS Summit
Frankfurt 2026

*The last 12 months taught me one thing: **the feature question is usually the wrong question** to start with.*




Ondřej Chrastina

Developer Advocate



ondrej.chrastina.dev



The Copy-Paste Ritual

The hidden productivity tax on every AI-generated paragraph.

01

The Export

Strip data from your CMS or business app into a raw text block.

02

The Context

Open a new tab and paste into ChatGPT, Claude, or Gemini.

03

The Prompt

Brainstorm and iterate until the output looks "mostly right."

04

The Return

Copy the AI response and paste it back into your original app.


05

The Cleanup

Spend 30 minutes fixing broken links, tables, and formatting.

"We aren't automating workflows, we're just shifting the manual labor to browser tabs."

GenAI workflows need improvement



Standalone tools: ChatGPT, Claude, Gemini, etc.

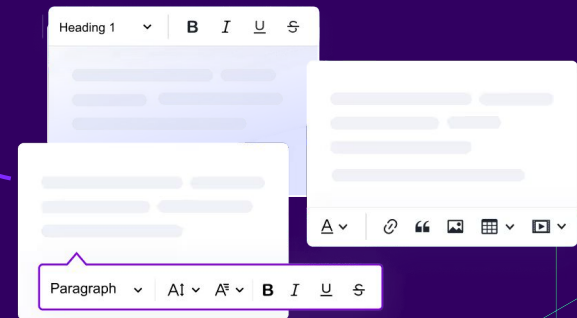
Infrastructure

- Data access
- Security
- Visibility
- Shadow AI

Interoperability
issues that kill
momentum

Content workflows

- Formatting
- Collaboration
- Versioning
- Track changes



Business applications:
CRM, CMS, LMS, Internal App, etc.



At the heart of the transformation

10%

**of the world's adult population
had adopted ChatGPT by
September 2025¹**

GenAI is the fastest-spreading technology in human history. And when we include Gemini, Claude, and open-source models, adoption becomes even more striking.

3.36x

**increase of work-related usage
of ChatGPT from June 2024 to June
2025¹**

The transformation is already reshaping knowledge work. Nearly half of all employees (45%) and more than two-thirds of leaders (69%) were using generative AI for work².

~40%

**of all work-related AI messages
are writing tasks¹, the most
common use case**

This isn't coincidental. Large Language Models are fundamentally writing engines, designed for exactly the work that knowledge workers in modern organisations do most.

¹ How People use ChatGPT, page 1–2, OpenAI & National Bureau of Economic Research, 2025

² Work Trend Index Annual Report: The year the Frontier Firm is born, page 15, Microsoft, 2025

GenAI workflows need improvement

Fragmented UX and friction come at a cost

- The back-and-forth between your business app and a standalone GenAI tool **disrupts productivity, creates room for mistakes, and kills momentum.**
- Knowledge workers don't want an AI sidekick that lives in a separate browser tab—they **want AI inside their primary workspace.**
- People recognize that standalone AI lacks context and convenience; when AI is embedded within their regular tools and data flows, it becomes far more effective and trustworthy.

Embed [AI] to eliminate friction. Don't make customers do the heavy lifting... Embed into existing workflows and add value right where your customers already are.

Saanya Ojha, Partner at Bain Capital Ventures

45%

**would use gen AI tools more frequently
if they were integrated into their daily
workflows***

* Reconfiguring work: Change management in the age of gen AI,
McKinsey & Company, 2025

About This Talk

*What we found surprised us, **let's find if you have similar findings!***

Our insights are drawn from direct field research, not marketing decks:

1. **Enterprise discovery calls**
 - What we build?
2. **The Real Ask Was Underneath**
 - What enterprise want first?
3. **Are going agentic?**
 - What future holds?



Discovery calls


What we found out first
from enterprise clients &
see in our projects

Today's reality
vs.
production expectations.

What stories do we hear?

The urgency to “do AI now” is real
—but without a clear workflow problem,
speed turns into scattered pilots.

- Fear of being late again drives **top-down pressure** to move fast.
- “**AI everywhere**” creates activity, not necessarily outcomes.
- Start with a **specific workflow + measurable goal**.



What stories do we hear?

Enterprises need a platform they can trust.

- **Trust gates adoption:** security, compliance, and control.
- **AI raises the stakes** for residency, audits, IP, and access.
- **Trust needs proof:** clear partnerships and certifications that stand up to security reviews.



Breadth vs. Depth

General AI is broad, but your workflows need depth.

- ChatGPT-style tools do “**a bit of everything,**” but won’t go deep into your rules and edge cases.
- Your products are **intentionally deep**: templates, approvals, permissions, terminology, compliance.
- The win is **AI inside the workflow**, so creation, review, and governance stay in one place.




From the model call to user value

A model call is easy, but the user value comes from the surrounding workflow.

- **Demos work fast:** wire an API, get output, celebrate.
- **Real value** = how users trigger, guide, review, refine, approve, and publish.
- **A big leap ahead:** making this repeatable and consistent without building the whole UX yourself



Gaps and bridges




Speed

Shorten time-to-setup by months to solve real problems faster.




Trust

Use an enterprise-grade and compliance-ready solution.




Breadth vs. Depth

You know the depth. Don't waste time on general purpose tools.



API vs. Flows



Use partners for key parts that take years to build.



What we built


We Thought Features
Were the Answer

Today's reality
vs.
production expectations.



What We Built

The feature layer: Useful, context aware tools built for enterprise teams



Context Injection (Foundation): Always on brand guides, templates, and glossaries across all features.

Quality Bonus: Chunk streaming vs word by word changes perceived quality.

Enterprises liked these. And then came the real questions.

The Real Ask Was Underneath

We Thought Features
Were the Answer

Today's reality
vs.
production expectations.

The Infrastructure Spectrum

Deployment is a question of compliance posture, not just technical preference.

SAAS



Managed

Fastest start, fully handled scaling and updates.

BYOK



On-premises

Your own model provider keys and config.

HOSTED



Self-hosted

Your servers, custom telemetry and data residency.

GAPED

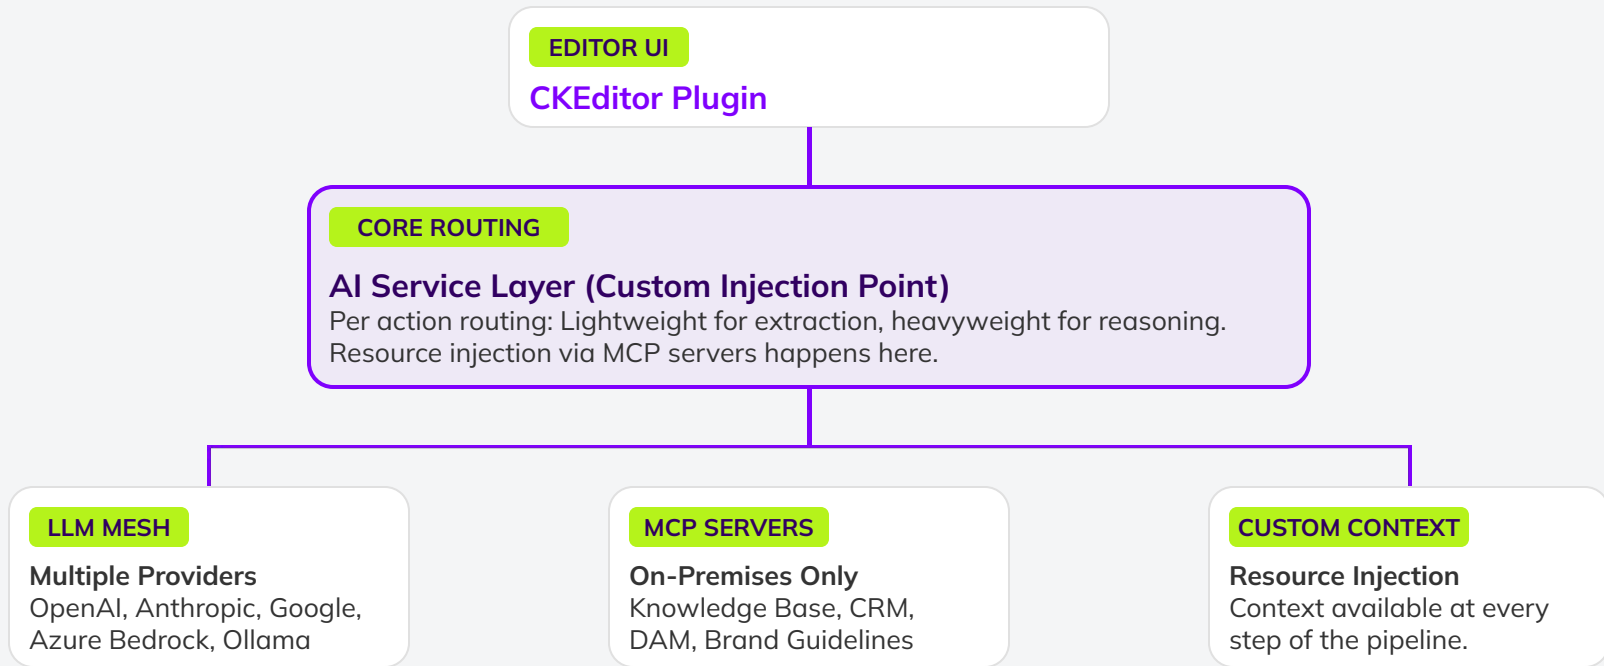


Air-gapped

Custom local model, zero external calls.

"On-prem isn't paranoia: it's a procurement and compliance requirement."

Architecture: From Plugin to Pipeline



"Enterprises don't want a chat widget: they want an AI layer they can wire into pipelines."



Your AI Infrastructure Today

Where does your organization sit on that spectrum?

- **SaaS:** Fast and fully managed (OpenAI, Anthropic Cloud)
- **BYOK:** Using your own enterprise API keys and config
- **Hosted:** Managed instances within your own cloud VPC
- **On-Premises:** Local servers running internal open models
- **Air-Gapped:** Zero external calls for total data residency

Custom LLMs: The Harder Question

"The hard part isn't running a model, it's knowing whether it's good enough for your specific editorial tasks."

THE ENTERPRISE ASSUMPTION


Bringing our own LLM is the hard part.

Focus on hosting infrastructure, latency, and tokens.

THE EDITORIAL REALITY

Quality tracking is the true bottleneck.

Frontier models carry years of accumulated quality history. Custom models without that history are a bet, not a solution. CKEditor maintains per-model history specifically for editorial task performance.



The 3B Problem

Small, local models are fast and free, but often fail at complex editorial reasoning.

Model Agility Beats Model Loyalty

Model Loyalty **VS.** Model Agility

THE STRATEGY

Don't commit to one model

Avoid vendor lock-in by designing for interchangeability rather than static integration.

Route per action

Use lightweight models for extraction and heavyweight frontier models for complex reasoning.

Leverage quality history

Frontier model support provides a quality benchmark that acts as a valuable organizational asset.

Agility = Intelligence

The power to switch per task, per provider, or per context without rewriting your core pipeline.

INTELLIGENT ROUTING

USER REQUEST / EDITORIAL TASK

Extraction Task
Lightweight (e.g. 8B)

Reasoning Task
Heavyweight (Frontier)

Unified Content Pipeline

Lessons Learned

The Individual to Enterprise Ceiling

WORKS FOR INDIVIDUALS

Drop-in AI Features

Quick actions like "make longer", "summarize", or "fix tone" provide immediate, standalone value for single users.

HITS A CEILING FOR ENTERPRISE

- No way to inject business context
- Cannot customize underlying logic
- Cannot use infra layer without feature layer

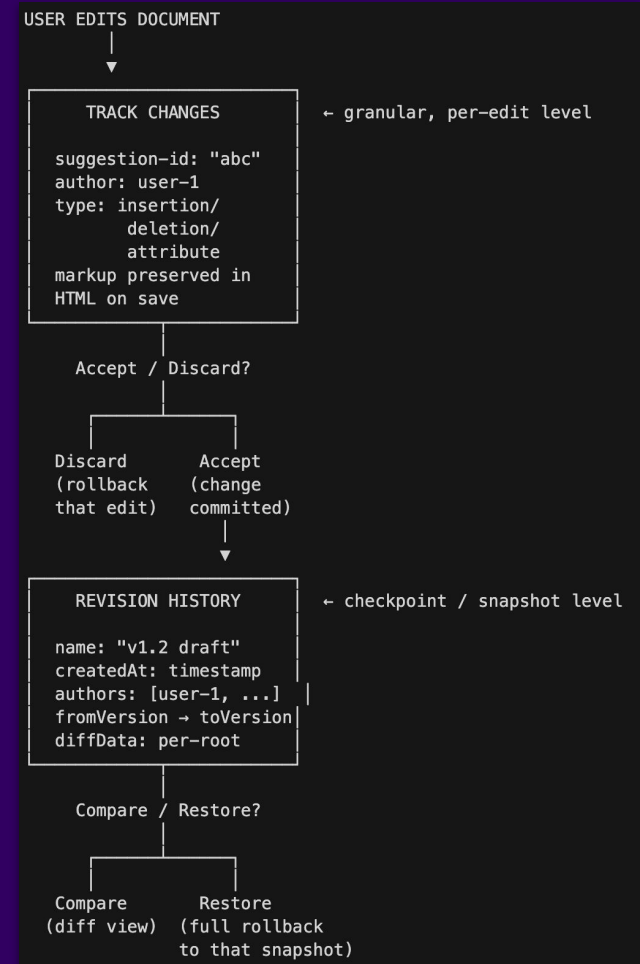

The real enterprise need is controllable infrastructure, not smarter quick actions.

Are going agentic?

What feature holds?

Tomorrow's
outlook

AI Interoperability with multi (user/agent) systems



AI as User vs. AI as Process

AI generates a change

Path A: Frontend (AI as user)

Human-in-the-loop Governance

- AI acts as a named participant
- Change enters Track Changes with agent ID
- Human reviews each suggestion individually

AI is Author / Human is Approver

Path B: Backend (AI as process)

Automated Pipeline Governance

- AI runs a pipeline or automation
- Result is delivered to a reviewer
- No individual author identity tracked

Reviewer is Sole Accountable Party

Five Things to Take Back

- 01** Start with **one workflow** , not a platform vision.
- 02** **Compliance is a feature** : own your deployment layer.
- 03** Frontend customization matters only once **backend trust** is solved.
- 04** Generic AI is breadth; **Enterprise AI** is depth + domain context.
- 05** **Model agility** beats loyalty: route per action, track quality.

Let's have lunch!



Ondřej Chrastina

Developer Advocate



ondrej.chrastina.dev

