



Bags, Docs, and Forecasts

GenAI as a force multiplier for a three-person team

Nathan Tallman · Academic Preservation Trust

DPC Workshop · Fantastic Futures 2026

3.2

people on the whole APTrust
team

GenAI has not let us do things we could not do.

It has let us do things that were always justified and never urgent enough to reach the top of the list.

01

A tool

s3-to-aptrust

02

A migration

aptrust-docs

03

A method

storage projections

01

s3-to-aptrust

A member's plain S3 bucket needs to become BagIt bags in APTrust

The naive path

- Pull it all down, bag it locally, push it back up
- Egress costs, weeks of transfer, local storage nobody has

The move

- Do the whole thing in the cloud - data never touches a local disk
- The hard part is bit integrity, not transfer

Why it is an AI story: a single-member migration tool with no roadmap that may never run again. A three-person team does not write that tool - until drafting gets cheap.



**The code was the cheap part. Knowing
what correct looked like was the
expensive part.**

gitlab.com/aptrust/s3-to-aptrust

02

aptrust-docs

Not a technical problem - a nobody-has-three-uninterrupted-weeks problem

4 → 1

MkDocs instances, on
different hosts,
consolidated into one site

What AI did

- Reconciled four config setups
- Mapped redirects so nothing 404s
- Caught broken cross-refs at scale
- Accessibility fixes → WCAG 2.2 AA

What it did not do

- Decide the information architecture
- Judge member needs vs. staff needs
- Judge what is public, what is deprecated

Docs are a good early target because the failure mode is visible. **A broken link announces itself. A bad storage forecast does not.**




AI didn't tell us how to organize the docs. It removed the excuse for not doing it.

docs.aptrust.org · shipped mid-August 2026

03

Storage projections

Deposits are lumpy and human-driven - the data has no clean signal



Chat vs. skill

- A codified procedure, not a conversation
- Anyone on the team can run it

The real risk

- Confidence is the default output
- The discipline: refusing false precision

I do not trust the forecast. I trust the range - and I trust that the method is consistent year over year, so our errors are comparable and we can see when reality drifts outside the envelope.



**The most important design decision
was making it produce a range instead
of a number.**

What a "skill" is: platform.claude.com/docs/en/agents-and-tools/agent-skills/overview

What stayed human

The pattern across all three: lowered activation energy for work that was always justified but never urgent.



Domain knowledge



**Information
architecture**



**Knowing when the
output is too
confident to trust**

Genuine open question: who maintains an AI-drafted migration tool in three years, when the person who prompted it has moved on? It was built with disposable intentions - so does that matter? Or is the cost of disposable infrastructure worth it?

Links

Happy to drop these in the chat or share afterwards

01 **s3-to-aptrust** gitlab.com/aptrust/s3-to-aptrust

01 **Baglt spec (RFC 8493)** rfc-editor.org/rfc/rfc8493.html

02 **aptrust-docs repository** github.com/APTrust/aptrust-docs

02 **The site, and the announcement post** docs.aptrust.org · aptrust.org/2026/08/19/

03 **What an agent "skill" is** platform.claude.com/docs/en/agents-and-tools/agent-skills/overview