FOR DOCUMENT360

Beyond the prompt: Architecting trust in AI documentation with Life Sciences precision

Eliza Marin · 8 April 2026

What’s the plan?

1. Strategic governance
2. Ownership
3. Accessibility

Eliza Marin · 8 April 2026

1. Strategic content governance: From static documentation to behavioral logic

Eliza Marin · 8 April 2026

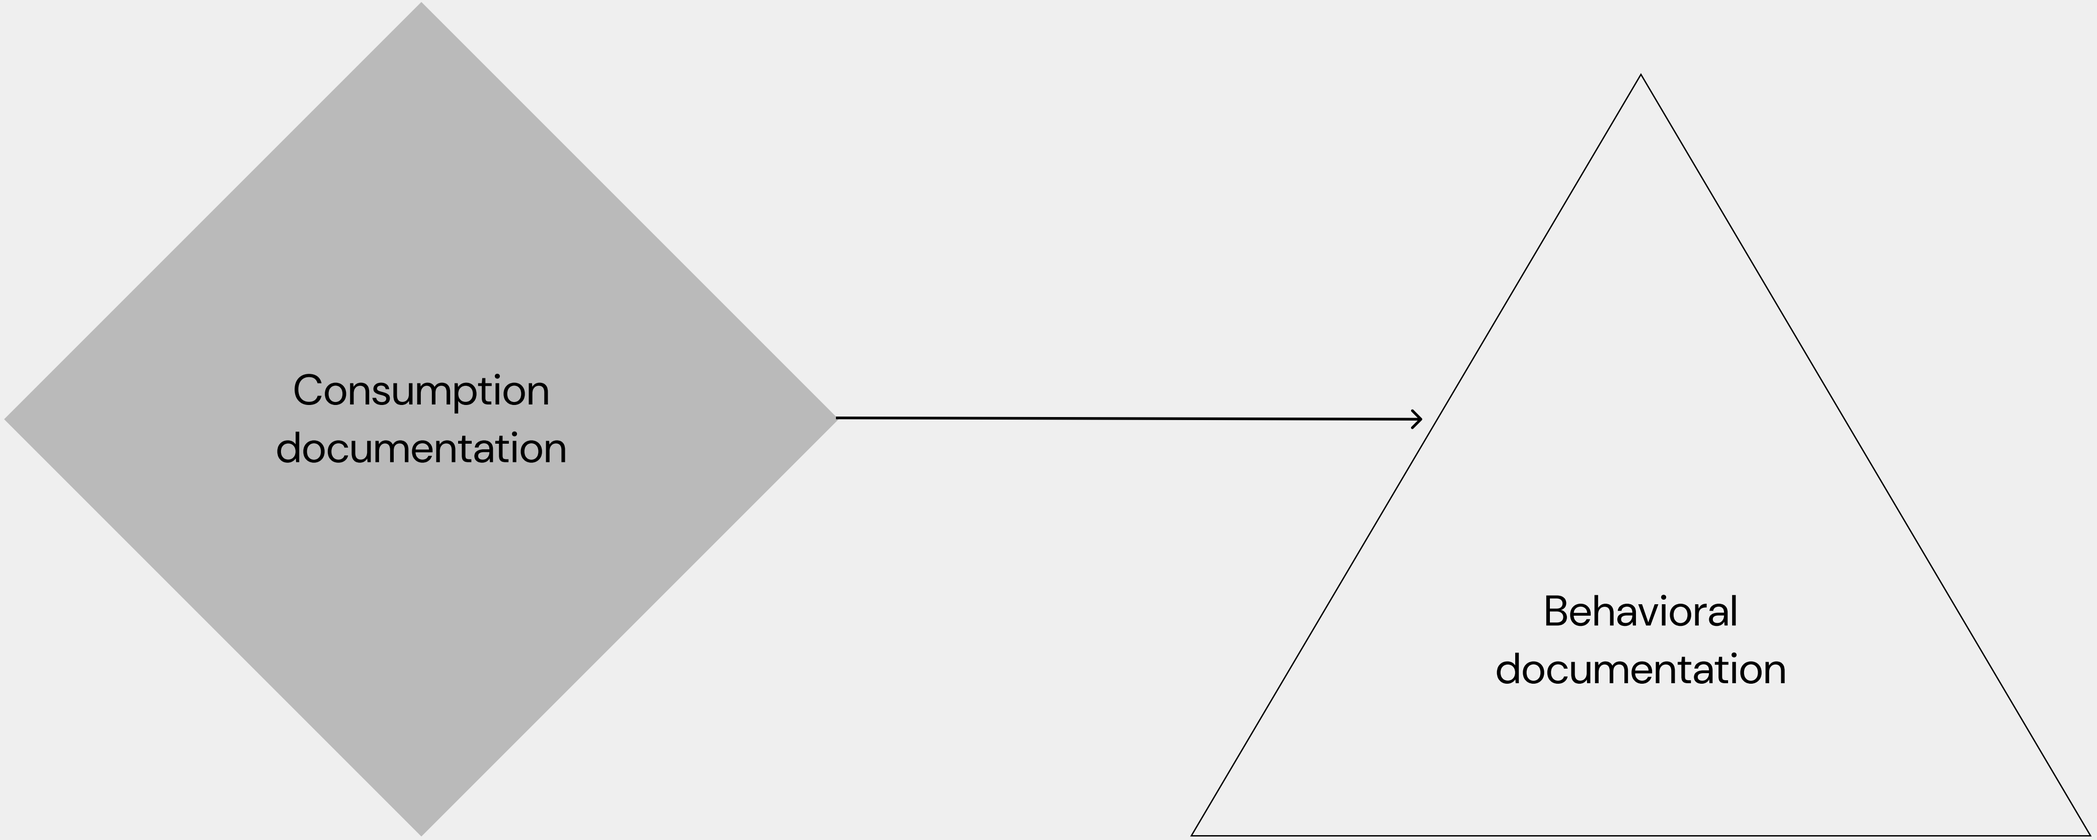

Human-read

you document the link interaction

AI-read

you **govern** the rules of its derivation.

Eliza Marin · 8 April 2026

Classic

Create a Y account and Log In

1. Navigate to the Y Cloud sign-up page to get started.
2. Do one of the following to sign up:
   - Click **Continue** with Google.
   - Enter the details for an existing business email account. This will be your username on Y Cloud; temporary email addresses are not allowed.
3. Follow the on-screen instructions. (**Note**: The steps vary depending on your account type.)
   - **Google account**: you might be asked to provide 2-factor authentication.
   - **Business email account**: you'll receive an email from Y Cloud with a link to set your password, then you'll need to sign in.
4. Follow the rules for setting up your initial password.
   - At least 1 uppercase character.
   - At least 1 lowercase character.
   - At least 1 number character.
   - At least 1 of * $ - + ? _ & = ! % @.
   - Between 12 and 128 characters.
5. Click **Get Started**.

Markdown

```
---
topic_id: AUTH-SIGNUP-001
service_name: Y Cloud
feature_area: Identity & Access Management (IAM)
user_intent: Account Creation / Onboarding
security_level: Tier 1 (Public Sign-up)
last_verified: 2026-04-06
---

# Create a Y Account and Log In

## 1. Access Gateway
**Action:** Navigate to the [Y Cloud Sign-up Page](https://cloud.y.com/signup).

## 2. Authentication Methods
| Method | Process Logic | Requirements |
| :--- | :--- | :--- |
| **Google SSO** | OAuth 2.0 Flow | May require 2FA (MFA) |
| **Business Email** | Manual Registration | **Strictly no temporary addresses** |

## 3. Account Verification (Conditional Logic)
- **IF Google Account:** Follow on-screen 2FA instructions.
- **IF Business Email:** 1. Await system-generated "Set Password" email.
    2. Execute "Initial Sign-in" via the provided secure link.

## 4. Password Governance (Policy: IAM-PW-01)
To ensure account integrity, the AI must validate that the user's password meets these **Mandatory Constraints**:

* **Length:** 12 – 128 characters.
* **Complexity (Uppercase):** Minimum 1 character.
* **Complexity (Lowercase):** Minimum 1 character.
* **Complexity (Numeric):** Minimum 1 character.
* **Complexity (Special):** Minimum 1 from this set: * $ - + ? _ & = ! % @.

**Final Action:** Click **Get Started** to initialize the workspace environment.
```

Critical = governed

- **Intent mapping**
- **Guardrails**
- **Action Verbs**

The Invisible = automated

- **Precise URL syntax**
- **Repetitive screenshots**

2. Ownership

Deconstructing the logic

If AI has direct access to the API schema and the code, what do tech writer do?

1

Writing behavior rules for AI.

Instead of writing a topic

2

Checking if the rules match the code.

Instead of checking if docs match UI, API, and general technical accuracy.

3

Proving the AI is following the rules.

Instead of thinking the AI is smart enough.

Who does what?

Role	Regulated method	Trust result
Engineering Implements the state layer (what the system can do)	GAMP 5 config baseline Feature flags, tier locks, region gates exposed as machine-readable metadata. **Every param: owner, version, tested state.**	**Zero-gap accuracy** AI queries verified config before responding. Cannot describe a feature as available if the baseline marks it disabled.
Product Defines the policy layer (what the system is allowed to do)	21 CFR 11 + EU Annex 11 **Access controls**: who can read or trigger a record. Data integrity: **what counts as a valid electronic record.** GDPR at data arch layer.	**Structural compliance** AI enforces policy by architecture, not prompt. Cannot suggest Pro to Basic or breach regional rules — boundaries are encoded.
Tech writer Architects the logic layer (how intent maps to governed rules)	ALCOA+ + controlled vocab Human-read tags become machine-executable rules. **One term, one meaning.** **Every chunk: Req. ID, owner, validation status.**	**Verified authority** Every AI response traces to an **approved**, **versioned content chunk. Auditable** under OQ/PQ. **Writer's authority is enforceable.**

3. Accessibility

1

Safety-critical doc standards

IEC 82079 & ISO 15223. If an action can cause data loss, permission escalation, or irreversible state change, the constraint must appear in the content chunk before the instruction, not after, so the AI surfaces it first in any retrieval.

2

Controlled vocabulary

One term, one meaning, no synonyms. Terminology must be attributable, legible, and unambiguous across the entire record lifecycle.

3

Logic mapping

Describe the **function** (e.g., "Trigger the user update" and not “Click the top-right blue button.").

4

One topic, one intent

DITA principle., widely adopted in regulated industries precisely because atomic topics map cleanly to Requirement IDs, can be individually versioned and approved, and allow AI retrieval to return a precise answer without pulling in adjacent context that may carry different validation status.

5

Test the logic

Test if the AI guides a user through a voice-only interface without **zero direction language**.

Handover — Documentation method comparison

	Classic tech writing	Regulated method	Trust result
SME review	**ACCURACY CHECK** Steps match the app Ensures documented steps reflect current UI and behaviour.	ATTESTATION · 21 CFR Every answer has a parent rule SME sign-off tied to Requirement ID, legally binding.	LEGAL DEFENSIBILITY Traceable, signed-off source Every AI response has provable parent rule, audit-ready.
Validation base	**OBSERVATIONAL** I tested this and it works Validation based on writer's direct experience.	REQUIREMENT-BASED · GAMP 5 Tied to a specific Requirement ID Content validated against defined requirement, not testing.	AUDIT-READY TRACEABILITY Provable chain from response System demonstrates which requirement authorised response.
Language	**ADAPTIVE** Synonyms for natural reading Writers vary terminology to match reader register.	CONTROLLED VOCAB · ASD-STE100 One term, one meaning Every action maps to exactly one approved term.	REDUCED HALLUCINATION AI cannot drift between synonyms One term anchors one action, preventing semantic drift.
Change trigger	**MANUAL** Writer updates after change Gap between product state and docs is structural.	EVENT-DRIVEN · GAMP 5 Code changes flag docs for review Config or code change automatically triggers review flag.	PERPETUAL SYNC AI and codebase stay aligned No lag between product state and documented behaviour.
Audit trail	**IMPLICIT** Depends on writer's notes No formal record of which requirement drove which update.	FORMAL RECORD Change control log per requirement Every update links to a requirement, timestamped and signed.	ALGORITHMIC PROOF Source data structures Audit system queries requirement graph to prove lineage.
Risk surface	**HUMAN MEMORY** Writer recalls, maybe records testing Documentation accuracy depends on individual recollection.	REQUIREMENT DRIFT Does content match intent? Risk is semantic mismatch between rule and implementation.	MODEL HALLUCINATION Did the AI invent? Risk is synthesis beyond sourced requirements and rules.

ASD-STE100 (Aerospace & Defense standard)

ASD-STE100 is an Aerospace & Defense standard maintained by ASD/ISPE, adopted across regulated industries including medical devices.
In Life Sciences, equivalent terminology discipline is enforced through ALCOA+ data integrity principles: Attributable, Legible, Contemporaneous, Original, Accurate.
The three methods reflect increasing formality and defensibility: Traditional writing prioritises readability; regulated approaches add traceability; trust-result architectures reduce hallucination risk through source control and semantic precision.

The glass box

Question

A user asks a high-stakes question. Instead of a helpful but unverified hallucination, the system triggers a **governed response protocol.**

Match

The AI scans your **structured writing** to find the specific **Requirement ID** (e.g., REQ-101) authorized or prohibited for that intent.

Verification

The AI performs a real-time check against **code base**. Verifies the person's permissions and if the system is in the correct state to perform **REQ-101.**

Trust

AI delivers the answer and cites its **requirement ID**. Every word is now a Glass Box, 100% traceable and audit-ready.

Let’s talk.

Thank you.

Mulțumesc.