How to write an AI Content Policy

And what your DAM needs to support it

Version 1.3

AI involvement: This guide is AI-assisted. Research and structure were developed in collaboration with Claude. Drafting was supported by ChatGPT. Final editorial direction and review were human.

www.qbankdam.com

AI content governance starts in the DAM, not the AI tool

Most conversations about AI content focus on the tools themselves. Which tools support C2PA. Which tools add watermarks. Which providers will comply with the EU AI Act.

Those questions matter. But they are only one part of the picture. The bigger challenge starts after the content leaves the AI tool.

Once an AI-generated asset is downloaded, uploaded to a DAM, edited, approved, distributed across channels, and published, the operational questions begin.

→ Who tracks the AI involvement?
→ Who decides whether disclosure is required?
→ Who confirms the rules were followed?
→ Who can prove how the asset was handled six months later?

That work does not happen inside the AI tool. It happens in the content infrastructure around it.

This is why AI governance starts in the DAM.
The AI tool creates the content. The DAM governs it.

When an asset contains Content Credentials, the provenance stays with the file. The DAM can map relevant information into metadata so it can be used in approvals, publishing rules and disclosure workflows. When no credential exists, the same governance information can be added manually.

This guide is built around that idea. An AI content policy only works when the systems behind it can operationalize the rules in practice.

For most organizations, the DAM is where that happens.

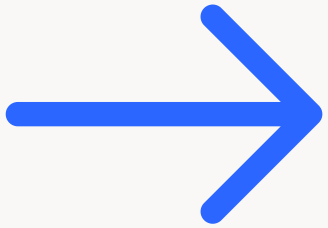

Why this matters?

AI changes content operations in ways traditional governance models were never designed for.

With traditional content, the chain is usually clear. You know who created the asset, where it came from, when it was approved, and how it was distributed.

AI complicates that process.

An image may be generated from a text prompt. A voiceover may be synthetic. A video may depict something that never happened. A document may be partially written by AI and then edited by a person.

That creates new governance questions.

Is this fully AI-generated or AI-assisted?
Was it reviewed by a human?
Can it be used in this channel?
Does it require disclosure?
Can we prove where it came from?

Since 2 August 2026, Article 50 of the EU AI Act includes transparency obligations for certain providers and deployers of AI systems. The exact requirement depends on the organisation's role, the type of content and how it is used. Many organisations choose to apply broader transparency rules through their own policy.

AI content introduces operational risk that most content governance models were never built to manage.

An AI content policy gives organizations a way to:

- identify AI involvement
- apply rules consistently
- control distribution
- document approvals
- maintain traceability

Without that structure, AI content spreads faster than governance can keep up.

What an AI content policy actually does

At its core, an AI content policy does three things.

It defines what counts as AI involvement

Not all AI content is the same.

There is a major difference between:

- a fully AI-generated campaign image
- a product photo enhanced with generative fill
- a document checked with AI writing assistance

The policy needs shared definitions so teams apply the right rules consistently.

It defines what content can go where

Different channels carry different levels of risk.

AI-generated social media content is very different from AI-generated regulatory documentation. Internal training material is different from customer-facing healthcare communication.

The policy maps content types to approved channels and use cases.

It connects rules to operational systems

A policy document alone changes nothing.

Governance only works when the rules connect to workflows, metadata, approvals, and distribution controls inside the DAM.

That is what turns policy into practice.

Where this policy fits

The EU AI Act includes transparency requirements for certain AI-generated and AI-manipulated content.

Providers of generative AI systems may need to ensure their outputs can be identified through machine-readable marking.

Organizations using AI professionally may need to disclose deepfakes and certain AI-generated text published on matters of public interest.

Not every AI-generated asset requires a visible label. However, many organizations choose broader transparency rules based on their brand, audience and risk.

This guide focuses on the operational side:
How organizations classify, approve, track and distribute AI-generated or AI-assisted content, while managing provenance and disclosure requirements.

It is not a broader AI ethics framework. It is a practical governance guide for content operations.

AI involvement: This image is fully AI-generated. Created with Gemini for illustrative purposes.

Two questions to ask before you start:

Before you start building governance around AI-generated content, settle two things first: what AI you are actually working with, and what kind of governance model your organization needs. These decisions shape everything that comes after.

AI involvement: This image is fully AI-generated. Created with Gemini for illustrative purposes.

Question 1:

What AI are you actually working with?

Before you write rules, take stock. The policy has to fit the AI use you actually have, not the AI use you imagine.

Walk through the questions:

- Which teams are using AI today, and for what?
- Which AI tools are in play? (Midjourney, Adobe Firefly, ChatGPT, Synthesia, ElevenLabs, internal models, others?)
- What kind of content is being produced? Imagery, video, voiceover, written copy, technical drawings?
- Is the AI generating content end-to-end, or assisting with edits to human-created work?
- Are any of the AI systems or use cases classified as high-risk under the AI Act? This is a separate assessment from AI content transparency.
- What's the volume? Occasional experimental use, or hundreds of assets a week?
- Do the tools add C2PA Content Credentials or other machine-readable markings?
- What should happen when an asset has no embedded provenance information?

You're looking for the shape of the problem. A retail company using Firefly for lifestyle imagery has a different policy need than a medtech company evaluating AI for clinical documentation.

This step also surfaces shadow AI: content created with AI tools but never recorded as such. Assets made available before 2 August 2026 do not require retroactive labelling, but reviewing older content is still worthwhile if it remains in use or may be published again.

Question 2:

Is this a company-wide policy or a DAM-specific policy?

These are two different documents, and it's worth being clear which one you're writing.

A company-wide AI content policy governs how AI is used across all content production in the organization. It applies to anything anyone creates, whether or not it ends up in the DAM. It's the right level for setting overall direction, and it usually sits with marketing leadership, legal, or a cross-functional governance group.

The DAM-specific policy also defines how embedded provenance is mapped into operational metadata and what users must provide when provenance is missing. What metadata gets recorded, what approval flow it goes through, what rules apply at distribution. It's narrower, more operational, and usually owned by whoever runs the DAM.

The two can be the same document, or they can be separate documents with the DAM-specific one drawing from the company-wide one. Either works.

What matters most is using the same definitions across the organization. Marketing, legal and DAM teams should agree on what counts as human-created, AI-assisted or fully AI-generated content. Otherwise, the same asset may be classified differently depending on who reviews it, making approval, disclosure and publishing rules harder to apply consistently.

For the rest of this guide, we use "AI content policy" to mean either case. The structure is the same. The scope is the choice you're making.

AI involvement: This image is fully AI-generated. Created with Gemini for illustrative purposes.

Understanding the levels of AI involvement

Most policies need to distinguish at least 4 levels. The level affects what rules apply.

Fully AI-generated. Content created by AI without a human-created original. Examples include prompt-generated images, synthetic voiceovers and AI-written text. These assets may require additional review or disclosure, depending on how they are used.

AI-assisted or modified. Human-created content that has been edited or enhanced with AI, such as generative fill, background removal, audio cleanup or writing assistance. The key question is whether AI meaningfully changed the content.

Human-created. Content created without meaningful AI involvement. Basic editing, such as cropping, color correction, format conversion or spell-checking, will usually remain in this category.

Unknown. Content where the level of AI involvement cannot be confirmed. Missing Content Credentials or other provenance information should not automatically be interpreted as proof that an asset is human-created.

These distinctions help organizations apply rules that match the content and the actual risk.

What about deepfakes?
Deepfakes are not a separate level of AI involvement. They are a higher-risk type of content that may be fully AI-generated or created by modifying existing material. They may require additional review, disclosure or distribution restrictions.

AI involvement: This image is fully AI-generated. Created with Gemini for illustrative purposes.

The structure of an AI content policy

Now the policy itself. Each section that follows covers what goes in it, why it matters, how organizations land differently, and what your DAM needs to be able to do to support that section in practice.

1 Purpose and scope

What goes in this section.
Say why the policy exists, what content it covers, and who must follow it.

Keep it short. This section should make the boundaries clear from the start.

Why it matters.
Unclear scope creates friction. Every unclear case becomes a new discussion: does this apply to agencies? To internal content? To partner portals? To archived assets?

A good scope removes guesswork.

What organizations decide.
Decide whether the policy applies to:

- content created by employees
- content created by agencies and suppliers
- internal content
- external content
- partner and contractor-facing content
- all asset types, or only selected ones

Most organizations should avoid assuming that “internal” means low risk. Internal content often reaches employees, consultants, partners, and other groups outside a narrow technical setting.

What your DAM needs.
Your DAM needs to separate assets by context, channel, audience, and use case.
The same asset may be safe in one channel and restricted in another. That distinction needs to live in metadata and distribution rules, not in someone’s memory.

2 Definitions

What goes in this section.
Define the terms the rest of the policy depends on.
At minimum, define the four levels of AI involvement:

- Fully AI-generated.
- AI-assisted or modified.
- Human-created.
- Unknown.

Also define related terms that affect policy decisions, such as deepfakes, synthetic content and disclosure.

Why it matters.
Most AI content debates start with unclear language.
If marketing, legal and DAM teams use different definitions, the same asset may be classified differently and the rules applied inconsistently.
Clear definitions make governance easier to understand and follow.

What organizations decide.
Decide where your lines are.

For example:

- Does limited use of generative fill count as AI-assisted?
- Does AI spell-checking count as meaningful AI involvement?
- How should content be handled when AI involvement is unknown?
- Are deepfakes prohibited, restricted or allowed with review?

The important thing is that the same definitions are used across legal, marketing, operations and the DAM.

What your DAM needs.
Your DAM needs controlled metadata fields that match the definitions in your policy. Use fixed values for AI involvement and manage related considerations, such as disclosure or deepfake risk, separately when relevant.

3 Approved and prohibited uses

What goes in this section.
Define what AI content is allowed, what is not allowed, and what needs extra review.

This is where the policy becomes useful.

Why it matters.
“Use AI responsibly” is not a rule. It is a hope.

Teams need clear guidance they can act on. They need to know what is safe, what is blocked, and what needs approval before it moves forward.

What organizations decide.
Examples:

- a retail brand may allow AI lifestyle imagery but block AI product photography
- a medtech company may allow AI illustrations but block AI clinical device imagery
- a B2B company may allow AI campaign visuals but block AI depictions of real customers
- a financial services company may allow AI for internal training but require legal sign-off for customer-facing use

The right answer depends on brand, audience, regulation, and risk appetite.

What your DAM needs.
Your DAM needs to connect use cases to distribution rules.
If AI-generated product imagery is prohibited, the DAM should help prevent it from reaching product pages, sales portals, or partner channels.
A rule that cannot be enforced will eventually be bypassed.

4 Approval requirements

What goes in this section.
Define who approves AI content, what they check, and when legal or compliance needs to be involved.

Why it matters.
Approval is where governance either works or becomes theatre. Someone needs to verify that the asset is correctly classified, allowed for the intended use, approved for the channel, and disclosed where required.

What organizations decide.
Some organizations route all AI content through one central approver. Others assign approval by content type, market, channel, or risk level.

A practical model is to keep the existing content approver, but add AI-specific checks to the approval process.

Even if the file includes information about human involvement, the final approval should still be recorded in the DAM.

What your DAM needs.
Your DAM needs approval states and distribution gates.
Unapproved AI content should not be able to flow into live channels by accident. The DAM should also make review queues visible, so the right people see what needs action.

5 Disclosure requirements

What goes in this section.
Define when, where and how AI involvement should be disclosed.

Distinguish between embedded provenance, such as Content Credentials, and visible disclosure, such as labels, icons or watermarks.

Why it matters.
Embedded provenance informs systems, not audiences.

Visible disclosure may be required by law or company policy, but not all AI-generated content needs a label.

What your DAM needs.
Your DAM should preserve embedded provenance, record disclosure requirements as metadata and ensure that content is only distributed to channels that can support the required disclosure.

What organizations decide.
Some organizations disclose AI involvement broadly. Others disclose only when required.

Decide:

- Which content requires disclosure.
- Which channels can support it.
- What the label, icon or wording should say.
- Who checks that disclosure is applied correctly.

For video and audio, disclosure should be clear enough for the audience to notice and understand.

EU icons for AI-generated content
The EU has created a set of icons to help organizations clearly label AI-generated or modified content. The icons are free to use but optional.

›› View and download the EU icons

6 Channel rules

What goes in this section.
Map which channels can use which kinds of AI content, and under what conditions.

This is one of the most important sections in the policy.

Why it matters.
AI governance breaks when every publishing decision depends on personal judgment.

Channel rules remove that uncertainty. They turn the policy into clear operational decisions.

What organizations decide.
Examples:

- social media may allow AI-generated illustrations
- product pages may block AI-generated product visuals
- regulated documentation may require legal review
- partner portals may allow only approved and disclosed AI content
- deepfake content may be blocked from all external channels

Your channel rules should connect directly to your disclosure rules. If a channel cannot support disclosure, it should not receive content that requires it.

What your DAM needs.
Your DAM needs metadata-driven distribution control.
Different portals, integrations, and publishing destinations should follow different rules. Without that control, the policy depends on people remembering every restriction every time.

That is not governance. That is risk.

7 Tools and vendors

What goes in this section.
List which AI tools are approved, which are not, and how new tools are reviewed.

Why it matters.
Without tool governance, shadow AI spreads fast.

People use what is easy. That often means consumer tools, free trials, or agency workflows with unclear data rules.

Once that content enters the DAM, tracking it later becomes much harder.

What organizations decide.
Some organizations keep a tight approved list. Others allow more tools through enterprise licenses and central procurement.
Some also treat enterprise and consumer versions differently.

When reviewing tools, consider:

- Do they support Content Credentials or other provenance information?
- Can information about AI involvement and the tool used be mapped into the DAM?

Review the approved list regularly. AI tools change quickly.

What your DAM needs.
Your DAM should record which tool was used to create or modify each AI-involved asset.
Use a controlled field for “generation tool” or similar. Keep it aligned with the approved tool list.
That way, you can find assets created with deprecated, risky, or unapproved tools later.

8 Record-keeping and traceability

What goes in this section.
Define what information should be recorded for AI-involved content, including how it was created, approved, disclosed and distributed.

Why it matters.
Sooner or later, someone will ask:

- Which assets involve AI?
- Where did the information come from?
- Who approved them?
- Was disclosure required?
- Was it applied?

Without records, the answer is guesswork.

What your DAM needs.
Your DAM should preserve embedded provenance, make relevant information usable as metadata and maintain records of approvals, changes and distribution. Information can be mapped from Content Credentials or entered manually, provided the source is clearly identified.

What organizations decide.
Start with the essentials:

- AI involvement.
- Source of metadata.
- Disclosure requirement.
- Approved channels.

Add other information when relevant:

- AI tool.
- Creator and creation date.
- Source asset references.

Use existing approval status and approver information where available.

Consider storing prompts or generation instructions when they support traceability, review or recreating an asset. As prompts may contain sensitive information, record them only when needed and limit access appropriately.

9 Roles and responsibilities

What goes in this section.
Name who owns the policy, approves content, manages the DAM setup and reviews the rules over time.

Why it matters.
Policies without owners drift.

If no one is responsible, no one updates the policy, checks adoption or fixes gaps when workflows change.

What organizations decide.
Larger organizations may use an AI governance group with legal, marketing, compliance, IT and DAM owners.

Smaller organizations may assign one clear owner and a small approver group.

Also decide who records AI involvement when provenance is missing and who keeps metadata and distribution rules up to date.

Both can work. The problem is when ownership is vague.

What your DAM needs.
Your DAM needs roles and permissions that match the responsibility model.
Approvers need approval rights. Uploaders need clear input requirements. Admins need control over metadata, workflows and distribution settings.

The platform should reflect the accountability structure.

10 Review and updates

What goes in this section.
Define how often the policy is reviewed, what should trigger an update and who is responsible for keeping it current.

Why it matters.
AI does not stand still.

Tools change. Regulations change. Internal use cases change. A policy that is not reviewed becomes outdated fast.

What organizations decide.
Annual review may be enough for some organizations.

Others need quarterly or twice-yearly review, especially if AI content is used at scale or in regulated contexts.

Also define trigger events, such as:

- New or discontinued AI tools.
- New regulations or disclosure requirements.
- Changes to Content Credentials or provenance standards.
- New content channels.
- Incidents or misuse.
- Major workflow changes.

What your DAM needs.
Your DAM should make policy changes enforceable on existing assets, not just new ones.
If a tool becomes prohibited, you need to find all assets created with it. If a channel rule changes, the DAM should help update distribution access.

A policy update should not become a manual cleanup hunt.

The minimum viable AI content policy

If you are starting from scratch, do not try to solve everything at once.

Most organizations need a policy that is practical enough to use before it becomes comprehensive enough to impress legal.

Start with the basics first.

Most organizations can mature the framework over time. The important thing is **to start before unmanaged AI content spreads** across teams, channels, and asset libraries.

A minimum viable AI content policy should cover five things:

1. Definitions for AI involvement
Define fully AI-generated, AI-assisted or modified, human-created and unknown. Define deepfake separately as a content characteristic.

2. Approved and prohibited use cases
Clarify what is allowed, restricted, prohibited, or requires extra review.

3. Approval rules for external distribution
Name who reviews AI content before it reaches public channels.

4. Disclosure requirements
Define when disclosure is required by law or company policy, where it appears and which channels can support it.

5. Metadata fields that record AI involvement
Start with AI involvement, source of metadata, disclosure requirement and approved channels or use. Add the AI tool, creator, creation date, approver or source asset references when relevant.

These five areas create a workable operational foundation.

Start with a few core properties and add further information only when the asset or intended use makes it relevant.

Glossary

AI governance introduces a lot of new terminology. Here are some of the concepts most relevant to AI-generated content, transparency and provenance.

AI Act and governance terms

EU AI Act
The European Union's regulation for artificial intelligence. For content teams, it introduces rules around transparency and how AI-generated content is identified and disclosed.

Article 50
The part of the EU AI Act that focuses on transparency. It covers how AI-generated content should be marked and when people need to be told that AI was involved.

Fully AI-generated content
Content created entirely by AI. Examples include generated images, synthetic voiceovers and text written by an AI tool.

AI-assisted or modified content
Content created by a person and then edited or improved using AI. Examples include generative fill, background removal and AI writing assistance.

Human-created content
Content created without meaningful use of AI. Basic editing, cropping, color adjustments and spell-checking usually fall into this category.

Unknown AI involvement
Content where it is unclear whether AI was used. Missing provenance information does not automatically mean the content was created without AI.

Synthetic content
Content that has been fully or partly created or changed using AI. This can include images, videos, audio or text.

Deepfake
AI-generated or altered content that looks or sounds like a real person, place or event. It may need additional review, disclosure or restrictions.

Disclosure
Telling people that content was created or changed using AI. This can be done through labels, icons, watermarks or explanatory text.

AI involvement: This image is fully AI-generated. Created with Gemini for illustrative purposes.

Provenance and transparency terms

Provenance

Information about where an asset came from, how it was created and how it has changed over time.

C2PA

An open standard for attaching signed provenance information to digital content. It can help show how an asset was created or edited and whether AI was involved.

Content Credentials

A user-friendly term for C2PA-based provenance information. Content Credentials can provide details about an asset's origin, creation and editing history.

Manifest

A record containing provenance information, such as the tools used, edits made and AI involvement. It can be embedded in the file or linked to it.

Machine-readable marking

Information added to content so software can identify whether AI was involved. It helps systems recognize AI-generated content but is not the same as visible disclosure.

Watermarking

Adding a visible or invisible mark to content. Visible watermarks can help disclose AI involvement, while invisible marks can help identify it.

If you’re already a QBanker

How you can configure the policy in QBank.

AI involvement property sets

Create a dedicated property set for AI governance.

Start with:

- AI involvement.
- Source of metadata.
- Disclosure requirement.
- Approved channels.

Add the AI tool, creator, creation date, approver or source asset references when relevant.

Information can be mapped from Content Credentials or entered manually. The original credential stays with the file, while the property set makes the information usable in QBank.

Attach the property set to the asset categories where AI content may appear.

Dynamic folders as governance workflows

Dynamic folders can automatically surface assets that need attention.

Examples:

- Assets where AI involvement is unknown.
- AI-involved assets pending approval.
- Assets missing disclosure information.
- Assets with manually entered AI information.
- AI-involved assets restricted from certain channels.

This helps turn your policy into an operational workflow.

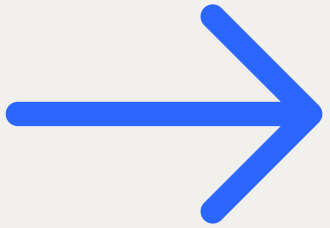

Portal filtering to enforce channel rules

QBank portals and publishing environments can filter assets based on metadata values. This allows different channels to follow different governance rules.

For example:

- A portal may allow AI-assisted or modified content.
- Another may block fully AI-generated assets.
- A regulated channel may only allow approved assets that meet its disclosure requirements.

This is where DAM governance becomes real operational control.

Preserved provenance on originals

When an original asset includes Content Credentials, QBank preserves that information with the file.

Relevant details can also be mapped into QBank properties, making them usable in approval, publishing and disclosure workflows.

Transformed versions may not retain the original credentials. When provenance needs to accompany the distributed asset, make the original file available.

Audit trail and reporting

QBank audit logs help track:

- Uploads.
- Approvals.
- Metadata changes.
- Workflow actions.
- Distribution history.

The audit trail can also be downloaded, making it easier to document how content was handled and share information for reviews, audits or compliance checks.

AI involvement: This image is fully AI-generated. Created with Gemini for illustrative purposes.

Take your content operations to the next level.

Visit www.qbankdam.com