

A practical guide:

How you can manage AI-generated content in QBank

For our customers building governance around
AI-generated assets



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.

Version 1.4 - Updated August 25th 2026

Copyright © QBNK Company 2026
www.qbankdam.com

Build AI governance that actually works

AI-generated and AI-modified content is already part of how organizations create and distribute content. It also raises new questions.

How was the content created? Has it been reviewed? Does it need to be disclosed? Where can it be used?

The EU AI Act introduces transparency requirements for certain AI-generated and modified content. For content teams, this means understanding what an asset is and applying the right rules before distribution.

QBank gives you the structure to stay in control. Content Credentials are preserved with the original file, while relevant information can be mapped to

QBank metadata. If no credentials exist, the information can be added manually. With property sets, dynamic folders, approval information and channel filtering, you can manage how AI-involved content is reviewed, disclosed and distributed.

Strong governance does not need to be complex. It starts with clear decisions, simple metadata and workflows that people understand.

This guide shows how you can build that structure in QBank.



Glossary

The discussions around AI governance and the EU AI Act introduce a lot of new terminology. To make the guides easier to navigate, we've included a short glossary of some of the most relevant concepts connected 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.

Step 1: Map where AI shows up in your content

Before you build anything, understand the scope. AI content rarely sits in one place. It is usually spread across teams, tools, channels and use cases.

Start by asking:

- Which teams use or plan to use AI?
- What is being created or modified? Images, product visuals, social posts, video, voiceovers or text?
- Which AI tools are being used?
- What AI-generated or AI-assisted content already exists in QBank?
- Does any content already include Content Credentials or other information about how it was created?
- Who should the policy cover? Employees, agencies, suppliers or partners?

The goal is a clear overview of who uses AI, what they create, which tools they use and where the content ends up.

You cannot govern what you have not found.

You do not need to update every older asset at once. Start by applying clear rules to new content, then review existing assets based on risk and whether they may be reused, updated or published again.

If you find content with unclear AI involvement, mark it as Unknown until it can be reviewed.

Step 2: Define your AI content policy

Before configuring QBank, decide how your organization wants to manage AI-generated and AI-modified content.

Your policy can apply across the company or focus specifically on how content is managed in QBank. Either way, the rules should be clear and consistent across teams.

Start by defining your levels of AI involvement:

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

Deepfakes and other sensitive synthetic content should be treated as a separate risk, not another level of AI involvement.

Then decide:

- What types of AI content are allowed?
- Which AI tools are approved?
- Who reviews and approves AI content?

- Where can the content be used?
- When is visible disclosure required?
- How should content with unknown AI involvement be handled?
- What additional rules apply to deepfakes or other sensitive content?

The policy does not need to be long. One page can be enough.

What matters is that the rules are agreed, documented and understood before they are turned into metadata, review steps and distribution rules in QBank.



[Download the AI Content Policy guide here](#)



Step 3: Design an AI involvement property set

Once the policy is clear, build the structure in QBank.

Create a property set called AI information and attach it to categories where AI content may appear, such as images, video, audio and documents.

Start with four core properties:

- **AI involvement:** Fully AI-generated, AI-assisted or modified, human-created or unknown.
- **Source of metadata:** Content Credentials, manual entry, other metadata or unknown.
- **Disclosure requirement:** Not required, required or review required.
- **Approved channels:** Where the asset can be used, such as internal channels, websites, social media or partner portals.

Add other properties when they are relevant:

- **AI tool:** The tool used to create or modify the content.
- **Created by:** The person or team responsible.
- **Creation date:** When the content was created.

- **Approval status:** Pending, approved, restricted or rejected.
- **Approver:** The person responsible for approving the asset.
- **Source asset references:** Related assets linked through a media picker.

Deepfakes can be recorded separately when they require additional review. Add a property called **Sensitive content**, if this is needed.

Information can be mapped from Content Credentials when available or added manually. The original credential remains with the file, while QBank properties make the information useful for search, review and distribution.

A valid credential does not automatically show that an asset was created using AI.



Consider storing generation prompts when useful, but avoid including sensitive information.

Example: Metadata tagging for AI governance

Keep the structure simple and easy to act on.

AI involvement:	Fully AI generated	Dropdown
Source of metadata:	Content Credentials.	Dropdown
Disclosure requirement:	Required	Dropdown
Approved channels:	Web, Portal, E-com	List

Additional properties, when relevant

AI tool:	Google Gemini	Dropdown
Created by:	Marketing team	Dropdown
Approval status:	Approved for distribution	Dropdown
Source asset references:	Product image 2026	Media Picker

This structure makes it easier to:

- Find and review AI-involved content.
- Control where assets can be used.
- Apply the right disclosure rules.
- Keep governance consistent across teams.

Information can be mapped from Content Credentials when available or entered manually. Only include an AI tool if it can be confirmed. Start with what you need. Add more information when it supports a real decision.



Ribbed Beanie Cobalt Blue Studio

AI-information

AI involvement: Fully AI Generated

AI tool: Google Gemini

Created by: [Marketing team](#)

Creation date: 2026.06.18

Approved channels: Web, Portal, E-com

Disclosure requirement: Required

Step 4: Decide how originals and derivatives are handled

The original AI generated file should be the **source of truth in QBank**.

That original stays intact and keeps its provenance information or manifest. When the asset is distributed through templates, QMO, or download variations, that chain may not follow the new file. But the original is still stored in QBank and can always be traced.

If the manifest is important for a specific use case, make sure the original is also available as a download option in the portal or channel where the asset is shared. The key is to communicate clearly which file carries the manifest.

Approved versions can be used when you need a specific version with its own preserved manifest, for example for legal, brand, or regulated use. In that case, upload the approved version as a separate asset and connect it back to the original with a media picker field.

For disclosure in templates or other formats, use metadata to define what should happen during distribution. For example, whether an AI label or watermark is required, or whether the original file needs to be available to preserve its Content Credential. This creates a practical foundation for automated disclosure rules across templates, portals, and channels.

Today, AI labels or watermarks can already be added through templates in QBank. The limitation is that the logic is not yet automatic. This means customers may need separate templates, for example one version with a watermark and one without.

QBank is working on rule based watermarking, where templates can react to metadata automatically. For example, if the disclosure requirement is set to “Required”, the AI label or watermark could be applied automatically during distribution.

This is the direction: less manual template management, fewer duplicate templates, and stronger governance at the point where the asset is actually used.

Today, QBank preserves original assets and their provenance information. QBank can already read available provenance information and map relevant details to metadata. The next step is making validation and trust information easier to use and carrying provenance through more of the distribution workflow.

The direction is clear: move from simply storing provenance information to actively using it as part of the governance layer.

The original AI-generated file should be the source of truth in QBank.

Today, QBank stores the original file and keeps it intact. If the file includes provenance information or a manifest, that information stays with the original. When the asset is transformed, downloaded as a variation, or used through templates or QMO, that information may not always follow the new file. But the original is still stored in QBank and can be traced.

If the manifest is important for a specific use case, make sure the original file is available as a download option in the portal or channel where the asset is shared. The important part is to make it clear which file carries the provenance information.

When an approved version already has its own manifest, for example for legal, brand or regulated use, upload that version as a separate asset and connect it back to the original with a media picker field.

What you can do in QBank today	What is coming / evolving
Keep the original file and its Content Credentials intact.	Explore becoming a trusted C2PA signer so changes made in QBank can be added to the asset's history.
View embedded C2PA information directly in QBank.	Make it easier to see whether a Content Credential is valid and trusted.
Map C2PA information to your own metadata fields.	Use that information in more automated workflows and distribution rules.
Record AI involvement, disclosure requirements and approved channels.	Show AI-involved assets with a small icon in QBank and QPortal.
Link original files, approved versions and source assets.	Improve how provenance follows assets across versions and formats.
Add labels or watermarks through templates.	Automatically add a watermark when an asset's metadata says disclosure is required.

The direction is to move from storing provenance information **to making it easier to use that information in everyday governance.**

Step 5: Configure dynamic folders as governance surfaces

QBank does not need heavy workflow machinery to make governance visible. Dynamic folders can surface the right assets automatically, based on metadata.

When metadata changes, the folder updates with it. Use them as live governance views. Each folder has a clear purpose, a clear owner, and a clear action. This makes it easier to find what needs attention without adding noise, manual lists, or extra process.

Dynamic governance in practice

Build these dynamic folders in QBank

Assets with unknown AI involvement.

Assets where AI involvement is empty or set to Unknown.

AI-involved assets pending approval.

Assets with AI involvement set, but approval status is still pending.

Assets created with unapproved tools.

Assets created with tools that are deprecated, prohibited, or no longer approved.

Assets missing a disclosure decision.

AI generated assets where disclosure requirement has not been set.

AI content needing review

Approved AI assets that need to be reviewed again after a defined period.


Assets with manually entered AI information.

Assets where the source of metadata is manual entry.

Make ownership clear

Decide who reviews each folder, how often, and what action they should take. The folders become the workflow.

Step 6: Set up the upload and review flow



Metadata only works if it is completed and kept up to date. The upload and review flow makes that happen.

QBank gives assets two key states: uploaded and approved. Metadata can be added in both states. Fields can be required when someone edits an asset, even if every field cannot be enforced at the approval transition.

A practical flow:

- **At upload:** Embedded C2PA information can be reviewed directly in QBank. If mappings are configured, relevant information can also be added to QBank properties.
- **Complete the metadata:** The uploader reviews the information and adds anything that is missing. If AI involvement cannot be confirmed, select Unknown.

- **Before approval:** The approver checks AI involvement, source of metadata, disclosure requirement and approved channels before approving the asset.
- **For older or batch content:** A responsible person works through the assets awaiting classification and adds missing information based on risk and usage.

Remember that a valid Content Credential does not automatically prove that AI was involved.

Give approvers a simple checklist: what to check, what to record and what should block distribution. This turns the policy into consistent everyday practice.

Step 7: Configure distribution rules, templates, and disclosure logic

This is where governance becomes real.

Each publishing channel in QBank can have filtering rules that decide which assets are allowed through.

For each channel, define:

- What approval status is required?
- Which AI involvement values are allowed?
- Which approved channel value is required?
- Are deepfakes or other sensitive assets blocked, if separately classified?
- Which disclosure requirements can the channel support?
- Does the approver field need to be filled?
- Are source asset references required?

The disclosure requirement should stay as clear metadata. Values such as Not required, Required or Review required are easier to govern than rules hidden in individual templates or channels.

The logic should sit in the distribution setup. That means portal rules, templates, watermarking workflows and channel-specific configurations.

For example, an asset marked as Required should only be distributed through channels that can support the necessary label or watermark. If the original Content Credential also needs to be available, make sure the channel can provide access to the original file.

Conditional watermarking will make it possible to apply labels or watermarks automatically based on metadata. For example, when an asset is marked as Fully AI-generated and disclosure is Required, the right watermark can be added during distribution.

QBank is also exploring ways to show AI involvement with a small icon in QBank and QPortal, making relevant assets easier to identify.

The metadata records the requirement. Channel rules and templates help apply the right action.

Put your channel rules into practice

Different channels can follow different rules, depending on how the content is used and who will see it.

Public website: Only show approved assets that are allowed for public use. If disclosure is required, make sure the right label or watermark is applied.

Social media: Allow AI-involved content only when the channel is approved and any required disclosure can be included.

Partner portal: Show assets approved for partner use. Exclude content that is restricted, awaiting review or not approved for that audience.

Internal channels: Allow a wider range of content when appropriate, while keeping assets with unknown AI involvement under review before external use.

Use the same metadata across channels:

- AI involvement shows how the asset was created or modified.
- Disclosure requirement shows whether a label is needed.
- Approved channels defines where the asset can be used.
- Approval status confirms whether the asset is ready for distribution.

If a channel cannot support the required disclosure, the asset should not be distributed there.

The policy defines the rules. QBank helps apply them where content is used.



Step 8: Build audit views with dynamic folders

You need to answer questions about AI content quickly, whether they come from legal, compliance, leadership or customers.

Dynamic folders make this easier. Create one main folder for AI governance, with subfolders for the views you need.

Examples:

- Distributed AI assets: AI content used within a defined period.
- Pending approval: AI content waiting for review.
- Unknown AI involvement: Assets that still need classification.
- Assets by AI tool: Content created with a specific tool.
- Disclosure required: Assets needing a label or watermark.
- Sensitive content: Deepfakes or other sensitive assets, if classified.
- Manual AI information: Assets with manually entered information.

Give relevant teams access to the views they need.

QBank's audit trail can also be downloaded to document how assets were uploaded, reviewed and distributed.

The setup does not need to be complex. It just needs to be ready when someone asks.



Note: This is an illustrative example, not an exact representation of how the view looks in QBank.

Step 9: Train the team and keep governance aligned

AI governance only works when people understand the rules and QBank stays aligned with them.

Train uploaders, approvers, reviewers and agencies on:

- The AI policy.
- Metadata fields and classifications.
- Content Credentials and manual information.
- Approval requirements.
- Distribution restrictions.
- Where to ask questions.

Use consistent classifications: fully AI-generated, AI-assisted or modified, human-created and unknown.

Keep the governance structure aligned as your policy evolves. This includes:

- Updating approved AI tools.
- Reviewing metadata and terminology.
- Adjusting disclosure and channel rules.
- Updating approval workflows.
- Removing outdated values.

Give the governance structure a clear owner responsible for maintaining the setup and keeping it aligned with the organization's AI policy.

For many organizations, an annual review is a good starting point. Organizations using AI at scale or in regulated industries may need more frequent reviews.

Also review the setup when:

- AI tools change.
- Disclosure requirements change.
- New channels are added.
- Approval workflows change.
- The AI policy is updated.
- C2PA standards or AI regulations evolve.

Governance only works if the metadata, workflows and policies stay aligned.

Step 10: Plan for what comes next

AI governance standards are evolving quickly. You do not need to solve everything today, but you should understand where the industry is moving.

QBank already makes it possible to review embedded C2PA information and map relevant fields to your metadata structure.

Areas to watch:

- Clearer visibility of credential validation and trust.
- Conditional disclosure and watermarking based on metadata.
- AI involvement indicators in QBank and QPortal.
- Stronger provenance support across versions and derivatives.
- The possibility of adding signed QBank manifests.
- Evolving transparency requirements under the EU AI Act.

The structure in this guide is designed to support that direction. Property sets, dynamic folders, portal rules, templates and governance workflows create a foundation that can evolve over time.

The direction is clear: move from manually managing AI governance to making provenance, metadata and automation part of the content workflow.





A starting point, not a finished system

AI governance is not a one-time project. It becomes part of how the organization works with content.

The good news is that most of the foundation already exists in QBank. Property sets, dynamic folders, templates, portal rules and governance workflows support a structured approach to AI-generated and AI-modified content.

Content Credentials add another layer. When available, they provide provenance information that can be reviewed and mapped to QBank metadata. The important part is not building the perfect system from day one. It is building a structure that people understand, trust and use consistently.

Start simple. Define the policy, structure the metadata, decide how assets should move through the organization and make governance visible through workflows and dynamic folders. From there, the setup can evolve with new standards, regulations and ways of working with AI.

The direction is clear: make governance part of the content flow itself.



**Take your content
operations to the
next level.**

Visit www.qbankdam.com