



Reliability Management

AI Content Compliance & Image Manager

User Manual

Auditing published imagery against Article 50 of the EU AI Act

Product	AI Content Compliance & Image Manager
Manual version	1.0
Applies to software version	0.9.0
Issued	12 August 2026
Document owner	Reliability Management
Classification	Internal

This manual documents AI Content Compliance & Image Manager version 0.9.0.

Contents

Before you start

- Who should read what
- What is in this version

What this tool is for

- The one idea
- What the tool will never claim
- What the tool cannot do to your website
- Separation of two questions

Getting started

- What you need
- Install and start
- Ports
- Signing in

Using the tool

- 4.1 Register a website and establish that you may crawl it
- 4.2 Build the inventory: run a crawl
- 4.3 Understand why an image was, or was not, flagged
- 4.4 Work the review queue and record a decision
- 4.5 Correct what the tool got wrong
- 4.6 Read the rule book
- 4.7 Export a report
- 4.8 Decide which language model would answer which question

Reference

- Triage bands
- Origin classes
- Review statuses
- Roles
- Rules in the active rule-set
- The eight model use cases
- Commands

Troubleshooting

Glossary

- Known limits of this version

Before you start

This manual describes the AI Content Compliance & Image Manager — referred to below as "the tool". The tool inventories the images published on websites you own, assesses each one against Article 50 of Regulation (EU) 2024/1689 (the EU AI Act) and adjacent law, and puts the small number that genuinely need a human decision in front of you.

Note Decision support, not legal advice. The rule-set in this build has not been reviewed by legal counsel. Every assessment is engineering's reading of the law, and the tool says so on every screen until someone signs it off.

Who should read what

If you are...	Read
Deciding whether an image needs a label	Chapter 2 (what this is for) and section 4.3 (why an image was flagged)
Running the audit day to day	Chapter 4 in full, then chapter 6 when something looks wrong
Installing or operating the tool	Chapter 3, and the commands table in chapter 5
Reviewing the legal reasoning	Section 4.6 (the rule book) and the glossary in chapter 7
Preparing evidence for an auditor	Section 4.7 (reports) and the note on the manifest at its end
Deciding what a language model may be used for	Section 4.8, and the model reference in chapter 5

What is in this version

Item	Value
Software version	0.9.0
Active rule-set	2026.08.01, 23 rules
Article 50 applies from	2 August 2026
Languages	German and English, switchable at any time
Language models in use	None. See section 4.8.

What this tool is for

The one idea

Article 50 does not say "label every AI image". It splits the duty in two, and puts the halves on different people:

- Article 50(2) puts a machine-readable marking duty on the provider of the system that generated the content — not on the website that publishes the result. If you did not operate the generating system, this duty is not yours.
- Article 50(4) puts a visible disclosure duty on the deployer, and only for deep fakes.

Article 3(60) defines a deep fake as content that resembles existing persons, objects, places, entities or events AND would falsely appear to a person to be authentic or truthful. Both elements, cumulatively. If either is absent, the content is not a deep fake and Article 50(4) does not require a visible label.

So an AI-generated gear icon in a navigation bar needs no label, and a photorealistic image presented as a customer's actual plant needs rather more than one. The tool exists to tell those two apart, with evidence, and to stop the second from drowning in the first.

What the tool will never claim

Reliable detection of AI-generated imagery does not exist. The tool therefore never asserts that an image is AI-generated on the strength of how it looks. It records what it can prove — a C2PA manifest, an XMP or EXIF field, a production record — and grades everything else as a weaker class down to Undetermined, which is the default and an honest answer.

A filename containing the word "midjourney" is a hint. It will move an asset to Likely AI-generated. It will never reach Confirmed AI-generated, because a filename is not proof and the tool is built so that it cannot become one.

What the tool cannot do to your website

Nothing in this application can modify an audited website. The crawler issues GET and HEAD requests only; that restriction lives in the single network chokepoint every outbound request passes through, and is covered by tests. There is no deployment, FTP, CMS-API or git-push dependency anywhere in the product. Remediation is something you carry out yourself, in your own systems, after the tool has told you what and why.

Separation of two questions

The tool keeps two questions apart on purpose, and you will see them as separate panels on every asset:

- Origin — where did this image come from, and what evidence supports that? This is a question about the file.
- Obligation — given that origin, and given how the image is presented on the page, what does the law require? This is a question about the publication.

Collapsing the two is the mistake the whole product is designed to prevent. An image can be confirmed AI-generated and carry no obligation at all; an image of unclear origin can carry a very large one.

Getting started

What you need

- Node.js 22 or newer, and pnpm
- Docker, for the local Postgres and object storage
- A terminal, and about ten minutes for the first run

Install and start

From the repository root, once:

```
pnpm install
cp .env.example .env
docker compose up -d
pnpm db:generate && pnpm db:migrate
SEED_ADMIN_PASSWORD='choose-something' pnpm db:seed
```

The seed creates the tenant, an administrator, the rule-set and the websites listed in scripts/seed.ts. If you would like five illustrative assets to look at before you crawl anything real, add:

```
pnpm db:seed:demo
```

Then, in two terminals:

```
pnpm dev
pnpm worker
```

Note The worker is not optional. Crawling, downloading and analysis all run as background jobs. With the worker stopped, the Start crawl button will appear to succeed and nothing will happen — the job is queued and waits. This is the single most common cause of "the crawl is stuck".

Ports

The ports are pinned deliberately off the defaults, because the machine this was built on runs several other local backends. If one is taken later, move this service — do not stop whatever holds the port.

Service	Port
Web application	3110
Postgres	55433
Object storage (MinIO): API / console	9110 / 9111

Signing in

1. Open <http://localhost:3110> in a browser.
2. Enter the administrator email from the seed and the password you set in SEED_ADMIN_PASSWORD.
3. Press Sign in.

The sign-in screen. Use DE / EN at the top right to switch language at any time — the choice is remembered against your account.

The dashboard is where you land. It answers one question: where does your published imagery stand today.

The dashboard. The framed sentence under the banner is the most important one in the product, and it is on the first screen on purpose.

The eight tiles are live filters — clicking Critical takes you to the asset list already filtered to that band. Coverage tells you how much of your estate has actually been looked at, which is the number that decides how much the rest of the dashboard is worth.

Using the tool

This chapter is organised around jobs, not screens. Each section says what you are trying to achieve, how to do it, what you get, and what the tool will refuse to do.

4.1 Register a website and establish that you may crawl it

What you are trying to achieve

Bringing a domain into scope, and recording — before any request is made — whether you are entitled to crawl it and whether the AI Act reaches it at all.

How to do it

1. Open Websites.
2. In Add website, enter the Domain and a display name, and press Create.
3. Open the new entry and answer the EU nexus questions.
4. If the site belongs to a partner rather than to you, record the written clearance reference before attempting a crawl.

AI Content Compliance

DE

Sign out

Dashboard

Websites

Assets

Review queue

Rule book

Reports

Models

Websites

The domains in scope for auditing

Domain	Owner	EU nexus	Pages	Assets	Last crawl
containerpark24.com Containerpark24 The EU nexus of this website has not been assessed yet. Answer it once here — it applies to every asset on the site.	Containerpark24	Needs manual review	0	0	No crawl has run yet.
containerpark24.de Containerpark24 (DE) The EU nexus of this website has not been assessed yet. Answer it once here — it applies to every asset on the site.	Containerpark24	Needs manual review	0	0	No crawl has run yet.
demo.aica.local Demo (illustrative data)	Internal	Clear EU nexus	3	5	No crawl has run yet.
rcm-analysis-review.com RCM Analysis Review The EU nexus of this website has not been assessed yet. Answer it once here — it applies to every asset on the site.	Reliability Management	Needs manual review	0	0	No crawl has run yet.
reliabilitymanagement.co.uk Reliability Management The EU nexus of this website has not been assessed yet. Answer it once here — it applies to every asset on the site.	Reliability Management	Needs manual review	0	0	No crawl has run yet.
wrenchworth.com Wrenchworth The EU nexus of this website has not been assessed yet. Answer it once here — it applies to every asset on the site.	Wrenchworth	Needs manual review	0	0	No crawl has run yet.
apps.wrenchworth.com Wrenchworth Apps The EU nexus of this website has not been assessed yet. Answer it once here — it applies to every asset on the site.	Wrenchworth	Needs manual review	0	0	No crawl has run yet.

Add website

Import list

The Websites list. Pages, Assets and Last crawl are live counts, so a row of zeroes means the site has not been crawled — not that it is clean.

What you get

An EU nexus verdict that applies to every asset on the site, so the question is answered once rather than per image. Article 2(1) reaches operators outside the Union where the output produced by the AI system is used in the Union, which is why a domain suffix alone never decides it: the tool asks about EU addresses in the imprint, prices in euro, content in an EU official language, documented EU customers and so on, and weighs them.

The screenshot shows the 'AI Content Compliance' dashboard. The top navigation bar includes 'Dashboard', 'Websites', 'Assets', 'Review queue', 'Rule book', 'Reports', and 'Models'. On the right, there are buttons for 'DE' and 'Sign out'.

The main content area is titled 'demo.aica.local' with a subtitle 'Demo (illustrative data)'. It is divided into three sections:

- Details:** A form with fields for Owner (Internal), Brand (Demo), Platform (static), Mode (ANALYSIS_ONLY), Automated changes allowed (No), Operated by us (Yes), and Notes (Demo data for evaluating the tool. Safe to delete).
- Start crawl:** A section with a warning: 'The crawler only ever issues GET and HEAD requests. Nothing in this application can modify an audited website.' It includes input fields for Max depth (3), Max pages (200), and Request delay (ms) (1000). There is a checkbox for 'Respect robots.txt: Yes' and a blue 'Start crawl' button.
- EU nexus:** A section with a note: 'Article 2(1) reaches operators outside the Union where the output produced by the AI system is used in the Union. A domain suffix alone never decides this.' It contains two columns of checkboxes:
 - Left column: EU address in the imprint (3) [checked], Delivers products or services into the EU (3) [unchecked], Prices shown in euro (2) [checked], Explicitly addresses EU customers (2) [unchecked].
 - Right column: Registered EU establishment (3) [unchecked], Documented use by EU customers (3) [unchecked], Content in an EU official language (2) [checked], EU customer references (2) [unchecked].

A website record. The crawl settings sit next to the sentence describing what the crawler is allowed to do, and the EU nexus indicators are weighted, not counted.

What the tool will not do

- It will not crawl a partner site until a written clearance reference has been recorded. This gate is enforced in the server action and again in the command-line script, because a convenience script that skips the check is how the check stops meaning anything.
- It will not treat "needs manual review" as a nexus answer. Until you answer, the site's assets are assessed but the nexus question stays visibly open.

4.2 Build the inventory: run a crawl

What you are trying to achieve

A complete list of the images actually published on a site — including the ones nobody remembers, such as CSS backgrounds, Open Graph cards and favicons.

How to do it

1. Open the website record.
2. Set Max depth, Max pages and Request delay (ms). The defaults are 3, 200 and 1000.
3. Press Start crawl.
4. Watch the crawl run appear on the same page. The worker does the work; the page reflects it.

For many domains at once, use the shell instead:

```
pnpm crawl <domain> [depth] [pages] [delayMs]
```

What you get

Every asset is stored once per tenant and identified by the hash of its bytes, with a separate record for each place it appears. One image used on three pages is one asset with three

placements — so a decision you make about it is made once, and the pages that would be affected are listed for you.

Originals are archived byte-identical and are never re-encoded, because re-encoding destroys exactly the provenance metadata the assessment depends on.

What the tool will not do

- It issues GET and HEAD only. It cannot change, upload to or delete anything on the site it is reading.
- It respects robots.txt by default, including the crawl delay, and skips disallowed paths. The override exists for sites you own, is off by default, and writes an audit entry naming the person who used it.
- It will not follow a URL to a private or internal address. Every outbound request passes one guarded chokepoint, and the check is repeated on each redirect rather than only on the first URL.

4.3 Understand why an image was, or was not, flagged

What you are trying to achieve

Being able to explain a verdict to someone who disagrees with it — including a regulator — without saying "the system decided".

How to do it

1. Open Assets, or click a tile on the dashboard to arrive pre-filtered.
2. Narrow the list with Triage band, Status, Origin and Website, then press Filter.
3. Open an asset to see the full assessment.

AI Content Compliance

DE Sign out

Dashboard Websites Assets Review queue Rule book Reports Models

Assets

Every image, icon, video and document found

Filter

Triage band: All

Status: All

Origin: All

Website: All

Filter

Showing 5

All

Preview	Origin	Type of image	Triage	Recommended action	Placements	Status
 1d518be111	Undetermined Inferred from evidence	Team photo	Review immediately	Legal review	1	Legal review
 f6fcacd184	Confirmed AI-generated Inferred from evidence	Industrial facility	Review immediately	Remove the asset	1	Legal review
 26f0859b6a	Declared AI-generated Declared by a person	Technical illustration	Improve documentation	Add machine-readable metadata	1	Needs review
 f54c243c85	Undetermined Inferred from evidence	Social / Open Graph graphic	Improve documentation	Verify the licence	1	Automatically analysed
 8c458f3bd6	Confirmed AI-generated Inferred from evidence	Icon	Improve documentation	Verify the licence	2	Automatically analysed

Decision support, not legal advice. [User manual](#)

The asset inventory. Origin and triage are separate columns because they answer separate questions.

What you get

The asset page is laid out as the argument itself. Preview and file facts on the left; Origin with the evidence that supports it in the middle; the Assessment, including both cumulative elements of Article 3(60) and the rule-set version that produced it, on the right.

AI Content Compliance

DESign out

DashboardWebsitesAssetsReview queueRule bookReportsModels

Assets

ImageReview immediatelyLegal review

Preview

PLACEHOLDER — not a real photograph

Stands for a photorealistic image presented as a real customer facility.

File type

image/png

Dimensions

1200 × 675

File size

36.5 kB

Content hash

f6fcacd1849e98e5fe4bc6bd168a1113349c2ebd8109fa58f34e5754bf4bab

Content Credentials (C2PA)

Absent

Externally hosted

No

Origin

Re-run assessment

Inferred from evidence

Confirmed AI-generated (Confidence 0.90)

Declared by a person

Not set

Licence type

Not set

Rights holder

Not set

Consent

Unknown

Evidence

XMP metadata field (0.90)

xmp:CreatorTool: Midjourney v6

Assessment

Deep-fake risk

Legal review required

Resembles an existing person, object, place or event

Yes

Could falsely appear authentic

Yes

Disclosure form

Visible disclosure required

Prominent, adjacent to the image

Rule-set version

2026.08.01 / 1.0.0

Article 3(60) requires BOTH elements. If either is absent, the content is not a deep fake and Article 50(4) does not require a visible label.

Assessment

Origin documentation

Partially documented through file metadata.

2 / 5

Deception risk

High risk of being misinterpreted as real.

5 / 5

Person and event nexus

A real organisation or industrial facility.

3 / 5

Licence and rights clarity

Plausible use, but not evidenced.

2 / 5

Transparency status

Required transparency is probably missing.

0 / 5

Recommended action

PRIMARY

Remove the asset

Urgency: Immediate

A synthetic image is presented as a real customer installation, reference or testimonial. Our policy does not allow this to be fixed with a label: replace it with genuine material or remove it.

A fabricated "customer refinery". Confirmed AI-generated at confidence 0.90 — on an XMP field, not on appearance — and both elements of Article 3(60) answered Yes.

Below that, the recommended action and its alternatives. Each carries its legal or policy basis, who is responsible, whose approval is needed, how to carry it out and how to verify it afterwards. Nothing is offered without saying which rule it comes from.

AI Content Compliance

DESign out

DashboardWebsitesAssetsReview queueRule bookReportsModels

Legal or policy basis: Internal company policy - POL-NO-FABRICATED-REFERENCES@1.0.0

Effort: Small

How to do it: Remove the image and adjust the layout so the page still works without it.

Verify afterwards: Re-crawl and confirm the asset no longer appears and the page renders correctly.

Alternatives

Replace with a real photograph

Urgency: Immediate

A synthetic image is presented as a real customer installation, reference or testimonial. Our policy does not allow this to be fixed with a label: replace it with genuine material or remove it.

Legal or policy basis: Internal company policy - POL-NO-FABRICATED-REFERENCES@1.0.0

Effort: Large

How to do it: Commission or select a genuine photograph of the actual subject, using the generated image briefing as the shot list.

Verify afterwards: The new asset is in place, its origin is documented, and the old asset is no longer referenced anywhere.

Fallback: Replace with an abstract illustration

Legal review

Urgency: Short term

This asset is AI-generated but was not generated by or for us. Whether the Article 50(4) disclosure duty reaches a publisher in that position is unsettled, so this needs a legal decision rather than an automatic one.

Legal or policy basis: Legal interpretation (unsettled) - AIA-DEPLOYER-ROLE-UNSETTLED@1.0.0

Effort: Medium

How to do it: Route the asset to a legal reviewer with the assessment, the evidence and the specific question that needs answering.

Verify afterwards: A legal reviewer has recorded a decision with reasoning.

Verify the licence

Urgency: Short term

No licence and no rights holder are recorded for this asset. Establish and record where it came from and on what terms.

Legal or policy basis: Binding law - COPYRIGHT-LICENCE-UNDOCUMENTED@1.0.0

Effort: Small

How to do it: Locate the licence or purchase record, confirm it covers the current use, and attach the document to the asset.

Verify afterwards: The asset shows a licence type, a rights holder and a stored document.

Fallback: Replace with licensed stock

Document the origin

Urgency: Short term

No licence and no rights holder are recorded for this asset. Establish and record where it came from and on what terms.

Legal or policy basis: Binding law - COPYRIGHT-LICENCE-UNDOCUMENTED@1.0.0

Effort: Very small

How to do it: Find out where the asset came from — commissioning records, the designer, the stock account, the generator history — and record it against the asset with any supporting document.

Verify afterwards: The asset shows a declared origin and at least one evidence item.

Alternatives to the primary recommendation. "Legal review" cites AIA-DEPLOYER-ROLE-UNSETTLED — the tool presents an open question as open rather than deciding it.

Then the rules that applied, each tagged with its category, and the deep-fake screening as a plain list of questions and answers.

AI Content Compliance

DESign out

DashboardWebsitesAssetsReview queueRule bookReportsModels

This asset is AI-generated but was not generated by or for us. Whether the Article 50(4) disclosure duty reaches a publisher in that position is unsettled, so this needs a legal decision rather than an automatic one.

This rule rests on our interpretation of an unsettled question, not on settled law. (Confidence 0.50)

BINDING LAW

How the disclosure must be made

AIA-50-5-DISCLOSURE-MANNER@1.0.0

A disclosure is required but none is present. It must be visible at first exposure, in the page language, and readable by assistive technology.

BINDING LAW

Potentially misleading imagery

UCPD-MISLEADING-IMAGERY@1.0.0

A synthetic image is being presented as documentation of something real. That is capable of misleading a viewer regardless of the AI Act.

BINDING LAW

Licence not documented

COPYRIGHT-LICENCE-UNDOCUMENTED@1.0.0

No licence and no rights holder are recorded for this asset. Establish and record where it came from and on what terms.

INTERNAL COMPANY POLICY

Company policy: no fabricated references

POL-NO-FABRICATED-REFERENCES@1.0.0

A synthetic image is presented as a real customer installation, reference or testimonial. Our policy does not allow this to be fixed with a label: replace it with genuine material or remove it.

INTERNAL COMPANY POLICY

Company policy: voluntary disclosure for photorealistic imagery

POL-VOLUNTARY-DISCLOSURE@1.0.0

A voluntary disclosure is suggested for this photorealistic AI image. This is company policy, not a legal obligation, and a reviewer may decline it.

INTERNAL COMPANY POLICY

Company policy: industrial imagery must be plausible

POL-INDUSTRIAL-ACCURACY@1.0.0

Work through the industrial plausibility checklist before approving this image: PPE, safe working situations, coherent machinery, plausible piping and instrumentation.

VOLUNTARY CODE OF PRACTICE

Content Credentials (voluntary)

COP-AIA-50-7-MARKING@1.0.0

Consider adding Content Credentials to this asset. This is a voluntary good practice, not a legal obligation.

Deep-fake screening

Was the content artificially created or manipulated?

Does it resemble an existing person, object, place, entity or event?

Could it falsely appear authentic or truthful?

Is the artificial nature obvious to an average person?

Is there an artistic, creative, satirical or fictional use?

Is a real person depicted in a way they have not authorised?

Is it used in marketing, education, reporting or another publicly relevant context?

Does a clear visible disclosure already exist?

Does machine-readable provenance information exist?

Is a manual or legal review required?

Yes

Yes

Yes

No

Unknown

Unknown

Yes

No

No

Yes

• The image is presented as documentation of a real customer, reference or testimonial. That is the highest-risk case we recognise and it goes to legal review.

The deep-fake screen. Unknown is a permitted answer and is shown as Unknown, not silently treated as No.

Software v0.9.0 · Page 11 of 25

The contrast is the point. The gear icon below is also Confirmed AI-generated, and attracts no disclosure recommendation at all: an icon resembles no existing person, object, place, entity or event, so the first element of Article 3(60) is not met and Article 50(4) is not engaged. What it does attract is a licence question, which is a different area of law.

AI Content Compliance

DE Sign out

Dashboard Websites Assets Review queue Rule book Reports Models

0b9ad8e891166ec7c

Content Credentials (C2PA)
Absent

Externally hosted
No

Transparency status

Required transparency is probably missing.

0 / 5

Recommended action

PRIMARY Verify the licence

Urgency: Short term

No licence and no rights holder are recorded for this asset. Establish and record where it came from and on what terms.

Legal or policy basis: Binding law - COPYRIGHT-LICENCE-UNDOCUMENTED@1.0.0 Responsible: Content reviewer
Effort: Small Approval required from: Website owner

How to do it: Locate the licence or purchase record, confirm it covers the current use, and attach the document to the asset.
Verify afterwards: The asset shows a licence type, a rights holder and a stored document.
Fallback: Replace with licensed stock

Alternatives

Add machine-readable metadata

Urgency: Short term

We operated the generating system, so the machine-readable marking duty under Article 50(2) applies to us directly.

Legal or policy basis: Binding law - AIA-50-2-PROVIDER-MARKING@1.0.0 Responsible: Developer
Effort: Small Approval required from: Website owner

How to do it: Write the provenance assertion into the file's XMP/PTC fields before it enters the build pipeline, and confirm the pipeline preserves it.
Verify afterwards: Download the published file and confirm the metadata survived.

Document the origin

Urgency: Short term

We operated the generating system, so the machine-readable marking duty under Article 50(2) applies to us directly.

Legal or policy basis: Binding law - AIA-50-2-PROVIDER-MARKING@1.0.0 Responsible: Content reviewer
Effort: Very small Approval required from: Content reviewer

How to do it: Find out where the asset came from — commissioning records, the designer, the stock account, the generator history — and record it against the asset with any supporting document.
Verify afterwards: The asset shows a declared origin and at least one evidence item.

Rules that applied

BINDING LAW Machine-readable marking (provider obligation)

AIA-50-2-PROVIDER-MARKING@1.0.0

We operated the generating system, so the machine-readable marking duty under Article 50(2) applies to us directly.

Confirmed AI-generated, and correctly not asked to carry a label. The recommendation is about the licence.

What the tool will not do

- It will not recommend a visible disclosure for functional or abstract imagery — logos, icons, diagrams, infographics, abstract illustrations, backgrounds, decoration. This is the rule POL-NO-DISCLOSURE-NOISE, and it exists because a queue full of icons is a queue nobody reads.
- It will not resolve the unsettled question of whether Article 50(4) reaches a publisher who merely licensed an AI image. It states both readings and escalates.
- It will not advance an asset past human review on its own. No code path exists that does.

4.4 Work the review queue and record a decision

What you are trying to achieve

Getting through the assets that actually need a person, in the order that risk — not discovery — dictates.

How to do it

1. Open Review queue. It is sorted by risk, not by when things were found.
2. Open the first asset and read the assessment.
3. Scroll to Record a decision.
4. Choose Acknowledge, Propose an action or Reject.

5. Write a comment. Acknowledge and Reject require one.
6. Press Record a decision.

AI Content Compliance DE Sign out

Dashboard Websites Assets Review queue Rule book Reports Models

Review queue

Sorted by risk, not by discovery order

Decision support, not legal advice. This rule-set has not been reviewed by legal counsel. Treat every assessment as engineering's reading of the law.

Showing 3

Preview	Origin	Type of image	Triage	Recommended action	Placements	Status
	1d518be111 Undetermined Inferred from evidence	Team photo	Review immediately	Legal review	1	Legal review
	f6fcacd184 Confirmed AI-generated Inferred from evidence	Industrial facility	Review immediately	Remove the asset	1	Legal review
	26f8859b6a Declared AI-generated Declared by a person	Technical illustration	Improve documentation	Add machine-readable metadata	1	Needs review

Decision support, not legal advice. [User manual](#)

The review queue. Three assets, not five — the two that need nothing are not here.

AI Content Compliance DE Sign out

Dashboard Websites Assets Review queue Rule book Reports Models

Yes

☐ Generated by or commissioned by us

Comment

Save

Record a decision

Decision

☐ Acknowledge → **Needs review** (Comment required)

☐ Propose an action → **Action proposed**

☐ Reject → **Rejected** (Comment required)

Comment

Record a decision

History

Nothing here yet.

Decision support, not legal advice. [User manual](#)

Recording a decision. Each option names the status it moves the asset to, so the consequence is visible before you commit to it.

What you get

The asset moves to the named status, and the decision is appended to its history with your name, your role and the time. The decision history is append-only at the database grant level: the application role does not hold UPDATE on it, so a decision cannot quietly be rewritten later — not by a user, and not by a future version of this software.

What the tool will not do

- It will not let you approve without a record. A comment is required wherever the decision closes something down.
- It will not delete an asset from your website. Remove the asset is a recommendation addressed to a person, with the steps to carry it out and a verification step to confirm it worked.

4.5 Correct what the tool got wrong

What you are trying to achieve

Replacing a guess with a fact you happen to know — that you commissioned an illustration, that a licence exists, that the place in the photograph is real.

How to do it

1. Open the asset and scroll to Correct the origin.
2. Set Origin, Licence type, Rights holder, Licence document and Consent as far as you can evidence them.
3. Under Subject matter, answer the questions that decide the deep-fake screen: depicts a real person, a real place, a real event; photorealistic; artistic or satirical context; shows a third-party mark; presented as a real reference or customer.
4. Tick Generated by or commissioned by us if that is the case — it moves the Article 50(2) marking duty onto you.
5. Add a comment and press Save.
6. Press Re-run assessment at the top of the page to see the consequence.

The screenshot shows the 'AI Content Compliance' interface. At the top, there's a navigation bar with links: Dashboard, Websites, Assets, Review queue, Rule book, Reports, Models. On the right, there are buttons for 'DE' and 'Sign out'.

The main content area is divided into two sections. The top section, titled 'Used on (1)', contains a table with three columns: URL, Source, and Alt text. The URL is 'https://demo.aica.local/referenzen.de', the Source is '', and the Alt text is 'Raffinerie unseres Kunden'.

The bottom section, titled 'Correct the origin', contains a form for overriding inferred values. It includes fields for Origin (a dropdown menu), Generator hints (a text input), Licence type (a dropdown menu), Rights holder (a text input), Licence document (a text input), and Consent (a dropdown menu with 'Unknown' selected). Below these is the 'Subject matter' section, which contains several dropdown menus for: Depicts a real person (Unknown), Depicts a real place (Yes), Depicts a real event (Unknown), Photorealistic (Yes), Artistic, satirical or fictional context (Unknown), Shows a third-party logo or mark (Unknown), and Presented as a real reference or customer (Yes). At the bottom of the form, there is a checkbox labeled 'Generated by or commissioned by us'.

Used on lists every page the asset appears on. Below it, the declaration form — your entry overrides the inferred value and is recorded with your name.

What you get

A declared origin outranks an inferred one, and the assessment is recomputed against it. Because the assessment is a pure function of the recorded facts and the rule-set version, the same facts will produce the same verdict in two years' time.

What the tool will not do

- It will not silently overwrite your correction on the next crawl. Declared values are yours; inferred values are the tool's.
- It will not let a correction disappear. The override is written to the decision history with your name attached.

4.6 Read the rule book

What you are trying to achieve

Checking the reasoning rather than the conclusion — or handing it to someone whose job is to disagree with it.

How to do it

1. Open Rule book.
2. Note the active rule-set version at the top; every assessment records the version it was produced under.
3. Read a rule: its source text, our interpretation, when it fires and its exceptions.

AI Content Compliance

DE Sign out

Dashboard Websites Assets Review queue Rule book Reports Models

Rule book

Every rule, its source, and our reading of it
Active rule-set: 2026.08.01.23

Decision support, not legal advice. This rule-set has not been reviewed by legal counsel. Treat every assessment as engineering's reading of the law.

BINDING LAW (12)

Machine-readable marking (provider obligation) AIA-50-2-PROVIDER-MARKING@1.0.0

Jurisdiction: EU In force from: 2026-08-02 Severity: Information
Confidence: 0.95 Requires human review: No Reviewed by: **Not yet reviewed by counsel**

Source text
Providers of AI systems ... generating synthetic audio, image, video or text content, shall ensure that the outputs ... are marked in a machine-readable format and detectable as artificially generated or manipulated.
Regulation (EU) 2024/1689, Article 50(2)

Our interpretation
Article 50(2) puts the machine-readable marking duty on the PROVIDER of the generating AI system, not on the website that publishes the result. Where we generated the image ourselves with our own system, we are the provider and the duty is ours.

When it applies
Fires when the asset's origin indicates AI involvement at LIKELY_AI or stronger. Reports that the machine-readable marking duty sits with the PROVIDER of the generating system, not with us as publisher — unless we operated the generating system ourselves.

Exceptions

- AI systems performing an assistive function for standard editing
- AI systems that do not substantially alter the input data or its semantics
- Use authorised by law to detect, prevent, investigate or prosecute criminal offences

Provenance marking appears to have been stripped AIA-50-2-MARKING-STRIPPED@1.0.0

Jurisdiction: EU In force from: 2026-08-02 Severity: Medium
Confidence: 0.70 Requires human review: No Reviewed by: **Not yet reviewed by counsel**

Source text
...marked in a machine-readable format and detectable as artificially generated or manipulated.
Regulation (EU) 2024/1689, Article 50(2); C2PA specification

Our interpretation
Provenance assertions are only useful if they survive publication. Image optimisation, CDN transforms and build pipelines routinely strip XMP and C2PA data, which undermines the marking regime and destroys our own best evidence of origin.

When it applies

The rule book. Source text and interpretation are shown separately, and every rule states whether counsel has reviewed it.

What you get

Twenty-three rules in six categories, each with a stable identifier and a version — AIA-50-2-PROVIDER-MARKING@1.0.0, for example. When a rule fires on an asset, that identifier appears on the asset, so a finding can always be traced back to the sentence of law it came from.

The categories are not decoration. Binding law, official guidance, a voluntary code of practice, internal company policy, technical best practice and legal interpretation are different kinds of claim, and the tool never lets a company policy look like a legal requirement.

What the tool will not do

- It will not present its reading as settled law. Every rule shows "Not yet reviewed by counsel" until that changes, and rules resting on an unsettled question are labelled as interpretation and carry a lower confidence.


4.7 Export a report

What you are trying to achieve

Giving a stakeholder, a board or an auditor something they can read, and that can still be verified after you have forgotten how it was made.

How to do it

1. Open Reports.
2. Choose Executive report, Asset report or Audit package.
3. Click the language and format you want: DE or EN, as PDF, XLSX, CSV or JSON.



Three reports, two languages, four formats. The manifest note at the bottom is the part an auditor cares about.

What you get

The executive report carries the counts that matter — websites and pages checked, assets found, confirmed AI content, unclear origin, possible deep-fake cases, missing licence evidence, immediate actions — followed by a 7 / 30 / 90-day plan derived from the triage bands rather than written by hand.

Every generated report is recorded with its timestamp, the rule-set version it was produced under, its manifest, and a SHA-256 of the exported bytes. A reader two years from now can confirm that the file they are holding is the file that was produced, and under which rules.

What the tool will not do

- It will not present the report as a legal opinion. The disclaimer is part of the document, not part of the covering email.

4.8 Decide which language model would answer which question

What you are trying to achieve

Settling — before any model is called — which questions a model may help with, what that would cost on your estate, and what a model is never allowed to conclude.

Note No language model is called anywhere in the current pipeline. Every assessment in version 0.9.0 is produced by deterministic code: metadata parsing, hashing and the rule book. This screen configures which model would answer which question. Nothing is switched on by default and nothing is billed.

How to do it

1. Open Models.
2. Read the boundary statement at the top before anything else.
3. Choose a ranking: Quality, Cost–benefit or Cost.
4. To use one model everywhere, press Apply to all use cases under One model for everything.
5. To choose per question, pick from the Alternatives list in a use case and press Apply next to it.
6. Press Refresh prices now if the catalogue looks stale.

AI Content Compliance

DE Sign out

Dashboard Websites Assets Review queue Rule book Reports Models

Language models

Which model answers which question, what it costs on this estate, and what it is allowed to conclude.

What a model may and may not do here A model may add evidence with a stated confidence. It may never write a finding, answer either cumulative element of Article 3(60), or clear an asset for publication. Every model-derived fact is marked as machine-inferred and still goes to a human.

Full pass over the estate
US\$0.00

Per month
US\$0.00

Billed this month
US\$0.00

Catalogue
400

No model is in use
Every assessment in this tool is currently produced by deterministic code: metadata parsing, hashing and the rule book. Nothing below is switched on, and nothing is billed.

OPENROUTER_API_KEY is not set. Prices and recommendations work without it; making calls does not.

Catalogue Refresh prices now
Prices synced 8 Aug 2026, 17:08
314 models in the catalogue have no quality rating yet. A qualitative judgement from published capability reports as of 2026-08-03. It is an opinion, not a measurement.
Monthly figures assume 10% of the estate is re-analysed each month and 4 executive reports. Full-pass figures are one analysis of all 5 active assets.
Billed in USD. Euro figures convert at 0.92 as configured on 2026-08-03; they are an indication, not a quote.

Quality Cost–benefit Cost Where paying more stops buying meaningfully better answers.

One model for everything
Applies a single model to all eight use cases and removes any per-use-case overrides.
OpenAI: GPT-5 Codex (batch) openai/gpt-5-codex:batch Quality 89 US\$0.14 / per month - €0.13
Apply to all use cases

The Models screen. Costs are zero because nothing is switched on; the catalogue is current because prices refresh themselves weekly.

What you get

Eight named use cases, each with the fact it would write, the quality floor it must clear and a projected cost for your estate in US dollars with a euro indication. The projection scales with your inventory: it assumes a full pass over every active asset, plus a monthly figure based on ten per cent of the estate being re-analysed and four executive reports.

AI Content Compliance

DESign out

DashboardWebsitesAssetsReview queueRule bookReportsModels

Per use case

Does the image look photographic?

The question that was never asked about the AI portrait published as a real photograph. Feeds realisticAppearance, which decides whether a disclosure obligation is even considered.

Sensitivity: Legally decisiveFloor for this question: 78Writes: AssetFacts.realisticAppearance, as machine-inferred evidence

Deterministic (no model)

RecommendedOpenAI: GPT-5 Codex (batch)openai/gpt-5-codex:batch

Quality 88 · Curated estimate · Needs reviewPer call: US\$0.00203Per month: US\$0.00203 (€0.00186)Full pass over the estate: US\$0.011 calls6 more quality points than Google: Gemini 3 Flash Preview (batch) for 0.00 USD more per month — 0.00 USD in total.

Alternatives

Keep deterministicApply

Does it depict a real person, place or event?

The first cumulative element of the deep-fake definition in Article 3(60). A model may narrow this to 'probably not'; only a human may record 'no'.

Sensitivity: Legally decisiveFloor for this question: 78Writes: AssetFacts.depictsRealPerson / depictsRealLocation / depictsRealEvent

Deterministic (no model)

RecommendedOpenAI: GPT-5 Codex (batch)openai/gpt-5-codex:batch

Quality 89 · Curated estimate · Needs reviewPer call: US\$0.00255Per month: US\$0.00255 (€0.00235)Full pass over the estate: US\$0.011 calls8 more quality points than Google: Gemini 3 Flash Preview (batch) for 0.00 USD more per month — 0.00 USD in total.

Alternatives

Keep deterministicApply

What kind of image is it?

Logo, portrait, industrial facility, decorative background. Currently guessed from the filename; the class decides whether a disclosure recommendation is suppressed as noise.

Sensitivity: HighFloor for this question: 64Writes: AssetFacts.functionalClass, alongside the filename heuristic

Per use case, under the Cost–benefit ranking. Each recommendation states its quality score, where that score came from, and what the next cheaper option would cost you in quality.

Switching ranking changes the recommendation and says why. Under Cost, the tool still refuses to drop below the floor for the question: a cheaper model is offered only if it clears the same bar.

AI Content Compliance

DE Sign out

Dashboard
Websites
Assets
Review queue
Rule book
Reports
Models

Per use case

Does the image look photographic?

The question that was never asked about the AI portrait published as a real photograph. Feeds realisticAppearance, which decides whether a disclosure obligation is even considered.

Sensitivity: **Legally decisive** Floor for this question: 78 Writes: AssetFacts.realisticAppearance, as machine-inferred evidence

Deterministic (no model)

Recommended

Google: Gemini 3 Flash Preview (batch)

google/gemini-3-flash-preview:batch

Quality 82

Curated estimate - Needs review

Per call: US\$0.00045

Per month: US\$0.00045 (€0.00041)

Full pass over the estate: US\$0.00223

1 calls

Cheapest option clearing the floor of 78: scores 82 at 0.00 USD per month.

Alternatives

Keep deterministic

▼ Apply

Does it depict a real person, place or event?

The first cumulative element of the deep-fake definition in Article 3(60). A model may narrow this to 'probably not'; only a human may record 'no'.

Sensitivity: **Legally decisive** Floor for this question: 78 Writes: AssetFacts.depictsRealPerson / depictsRealLocation / depictsRealEvent

Deterministic (no model)

Recommended

Google: Gemini 3 Flash Preview (batch)

google/gemini-3-flash-preview:batch

Quality 81

Curated estimate - Needs review

Per call: US\$0.00062

Per month: US\$0.00062 (€0.00057)

Full pass over the estate: US\$0.00308

1 calls

Cheapest option clearing the floor of 78: scores 81 at 0.00 USD per month.

Alternatives

Keep deterministic

▼ Apply

What kind of image is it?

Logo, portrait, industrial facility, decorative background. Currently guessed from the filename; the class decides whether a disclosure recommendation is suppressed as noise.

Sensitivity: **High** Floor for this question: 64 Writes: AssetFacts.functionalClass, alongside the filename heuristic

The same two questions under the Cost ranking. The wording changes from "more quality points for" to "cheapest option clearing the floor of 78".

What the tool will not do

- A model may add evidence with a stated confidence. It may never write a finding, answer either cumulative element of Article 3(60), or clear an asset for publication. Every model-derived fact is marked as machine-inferred and still goes to a human.
- It will not recommend an unrated model for a legally sensitive question. Prices refresh automatically each week; quality does not, because no public benchmark measures these questions. Rated model families carry a dated, curated opinion; the rest are marked unrated and capped below the threshold for any high or legally decisive question. You may still choose one — the tool simply will not recommend it.
- It will not recommend a free endpoint, a model alias, an image generator or a meta-router. A free endpoint is paid for with your imagery; an alias makes an assessment irreproducible; a router's actual model is unknown at the moment of the call.

Reference

Triage bands

Band	Shown as	Means
RED	Review immediately	A fabricated reference, an unauthorised depiction of a real person, or anything routed to legal review.
ORANGE	Address soon	A disclosure that is genuinely required, an unverified licence, or missing alt text.
YELLOW	Improve documentation	The publication is defensible; the paperwork behind it is not.
GREEN	No material action identified	Nothing the rule-set asks for. Not a legal clearance.
GREY	Not sufficiently assessable	Too little is known to place it. Usually means the file could not be fetched or carries no usable metadata.

Origin classes

Class	What it takes to reach it
Confirmed AI-generated	Hard evidence only: a C2PA manifest, an XMP or EXIF field, or a production record.
Declared AI-generated	A person recorded it, with their name.
Mixed or edited	Documented as part real, part generated — ranked with the declared origins, because Article 50(4) reaches manipulation as well as generation.
Likely AI-generated	Softer signals: filename, path, alt text.
Possible AI-generated	A weak single hint.
Undetermined	The default, and an honest answer.
Confirmed non-AI, Stock media, Original photography, Human designed	Recorded by a person or evidenced by a licence document.

Review statuses

Discovered · Automatically analysed · Needs information · Needs review · Legal review · Action proposed · Approved · Change prepared · Pull request created · Published · Verified · Rejected · Archived.

Only the first two are reached without a person. Everything from Needs review onwards is a human decision, recorded as one.

Roles

Role	May
Administrator	Everything, including recording partner clearance and changing model settings.
Website owner	Add websites, start crawls, answer EU nexus, decide on assets.
Content reviewer	Answer EU nexus, correct origin, re-run an assessment, decide on assets.
Legal reviewer	Decide on assets, including those escalated to legal review.
Designer, Developer	Named as responsible on recommended actions.
Partner reviewer	Confined to partner-site material.
Auditor	Read-only. Action controls are absent, not merely disabled.

Rules in the active rule-set

Identifier	About
AIA-2-EU-NEXUS	Whether the AI Act reaches this site at all.
AIA-50-2-PROVIDER-MARKING	Machine-readable marking — the provider's duty, not the publisher's.
AIA-50-2-MARKING-STRIPPED	Provenance data destroyed by an optimisation or build pipeline.
AIA-50-4-DEEPFAKE-DISCLOSURE	Visible disclosure where both elements of Article 3(60) are met.
AIA-50-4-DEEPFAKE-UNCERTAIN	One element is unclear — escalate rather than guess.
AIA-50-4-ARTISTIC-LIMITATION	The reduced disclosure for artistic, satirical or fictional work.
AIA-50-5-DISCLOSURE-MANNER	How a required disclosure must be presented.
AIA-DEPLOYER-ROLE-UNSETTLED	Whether a publisher who merely licensed the image is a deployer. Marked as interpretation.
COP-AIA-50-7-MARKING	Content Credentials as voluntary practice.
UCPD-MISLEADING-IMAGERY	Misleading imagery under consumer law, independent of the AI Act.
COPYRIGHT-LICENCE-UNDOCUMENTED	No licence and no rights holder.
COPYRIGHT-AI-OUTPUT-PROTECTION	What you do and do not own in an AI output.
GDPR-IDENTIFIABLE-PERSON	An identifiable person is depicted.
POL-NO-DISCLOSURE-NOISE	Do not label functional or abstract imagery.
POL-NO-FABRICATED-REFERENCES	A synthetic image may not be presented as a real customer or installation, and a label does not fix it.
POL-VOLUNTARY-DISCLOSURE	Voluntary disclosure for photorealistic AI imagery.

POL-ORIGIN-DOCUMENTATION	Record where imagery came from.
POL-INDUSTRIAL-ACCURACY	Industrial imagery must be plausible: PPE, safe working, coherent machinery and piping.
POL-PARTNER-CONFIRMATION	Partner material needs confirmation.

The eight model use cases

Use case	Quality floor
Does the image look photographic?	78 — legally decisive
Does it depict a real person, place or event?	78 — legally decisive
What kind of image is it?	64 — high
Visual generation artefacts	Deliberately weak evidence
How is the image presented on the page?	78 — legally decisive
Draft alt text	Lower — a human edits and approves it
Draft disclosure wording	Lower — reviewed before it reaches a page
Executive report preamble	Lowest risk of the eight

Commands

Command	Does
pnpm dev	The web application on port 3110
pnpm worker	The crawl, fetch and analysis worker
pnpm db:migrate	Applies migrations, then RLS policies and grants
pnpm db:seed	Tenant, administrator, rule-set, seed websites
pnpm db:seed:demo	Five illustrative assets under demo.aica.local
pnpm crawl <domain> [depth] [pages] [delayMs]	Crawl one registered site
pnpm llm:sync [--schedule]	Refresh model prices; --schedule books the weekly repeat
pnpm llm:report [QUALITY BALANCED COST]	Print what the Models screen would recommend, and what it would cost
pnpm test	The test suite; database-backed tests skip themselves if Postgres is down

Remove the demo data with: delete from websites where domain = 'demo.aica.local';

Troubleshooting

Symptom	Likely cause	What to do
Start crawl seems to work, nothing happens	The worker is not running	Start pnpm worker in a second terminal. The job is queued and will run as soon as it does.
A website shows 0 pages and 0 assets	It has not been crawled yet	Open the site and press Start crawl. A row of zeroes means unexamined, not clean.
A crawl finishes with far fewer pages than expected	Max pages or Max depth, or robots.txt disallowing the paths	Raise the limits on the website record. If robots.txt is the cause and the site is yours, the override exists — and writes an audit entry naming you.
A partner site refuses to crawl	No written clearance is recorded	An administrator records the clearance reference on the website record. The refusal is deliberate.
Everything is Undetermined	The images carry no usable provenance metadata	That is the honest answer, not a failure. Use Correct the origin to record what you know.
An assessment looks wrong	The tool inferred a fact you can correct	Correct the origin, save, then press Re-run assessment.
An asset sits in GREY	It could not be fetched, or carries nothing to assess	Check whether the URL is still reachable and whether the file is an image at all.
The Models screen shows US\$0.00 everywhere	No model is bound to any use case	That is the default state. Costs appear once you apply a model.
Prices on the Models screen look stale	The weekly refresh has not run	Press Refresh prices now, or run <code>pnpm llm:sync --schedule</code> to book the weekly repeat.
The application will not start on port 3110	Another process holds the port	Move this service to a free port. Do not stop the other process — check first with <code>lsof -nP -iTCP:3110 -sTCP:LISTEN</code> .
Sign-in fails after a database reset	The seed has not been re-run	Re-run <code>pnpm db:seed</code> with <code>SEED_ADMIN_PASSWORD</code> set.

Glossary

Term	Meaning here
Article 50	The transparency obligations of Regulation (EU) 2024/1689, applicable from 2 August 2026.
Article 50(2)	The machine-readable marking duty. It falls on the provider of the generating system.
Article 50(4)	The visible disclosure duty. It falls on the deployer, and only for deep fakes.
Article 3(60)	The definition of a deep fake: resembles something existing AND would falsely appear authentic. Both, cumulatively.
Provider	Whoever develops or operates the AI system that generated the content.
Deployer	Whoever uses the system's output under their own authority — in this context, whoever publishes the image.
EU nexus	Whether the AI Act reaches a given website, assessed on weighted indicators rather than on the domain suffix.
C2PA / Content Credentials	A signed, tamper-evident record of how a file was made, carried inside the file.
XMP, EXIF, IPTC	Metadata blocks inside an image file. XMP and EXIF fields count as hard evidence of origin; a filename does not.
Asset	One image, identified by the hash of its bytes and stored once per tenant.
Placement	One appearance of an asset on one page. An asset may have many.
Triage band	The five-way priority of an assessment: RED, ORANGE, YELLOW, GREEN, GREY.
Rule-set version	The dated, versioned catalogue an assessment was produced under. Recorded on every assessment and every report.
Hard evidence	Evidence strong enough to support Confirmed AI-generated: a C2PA manifest, an XMP or EXIF field, or a production record.
Quality floor	The minimum quality score a language model must reach before it may be recommended for a given question.
Machine-inferred	A fact produced by software rather than recorded by a person. Always marked as such, and always still reviewed.

Known limits of this version

- The rule-set has not been reviewed by legal counsel. It is engineering's reading of the law until it is.
- Whether Article 50(4) reaches a publisher who merely licensed an AI image is unsettled. The tool presents both readings and escalates rather than deciding.
- Model quality is an opinion, not a measurement. Prices refresh weekly and automatically; quality does not. The authoritative source — running a candidate model over images this team has already labelled — has its schema and its read path, and needs a labelled set that does not exist yet.
- Repository scanning with pull-request creation, and image generation, are designed but not built.