

# The Artificial Intelligence Registrar (AIR)

Version 1.0 (2026/04/07)

## Mission Statement

To help steer the private sector AI as it evolves. The purpose of the AI registrar would be to shepherd in the technology in a way that is productive whilst having a minimal impact on innovation and sector growth.

In other words, the purpose of the entity is not to impede or stifle the industry but to ensure the technology is deployed responsibly and in a way that minimises the risk to the public whilst maximising profits in the private sector.

AIR would have the following remit:

Protection of our capitalist system
Proper licensing
Content Identification for the consumer
Protection of our environment (impact assessment)
Complaints handling and management
Evidence Authentication (legal and other cases where appropriate eg. Insurance claims)

AIR would therefore have the following departments:

<a href="#">Licensing and Compliance</a>
<a href="#">Content Identification</a>
<a href="#">Impact Assessment</a>
<a href="#">Complaints</a>
<a href="#">Evidence Authentication</a> > Legal > Private Companies
Services > Website - for providing information to humans and to serve as an interface between the project and the people

> API (Application Programming Interface) - Mostly for allowing automated feedback from licensees

Obviously AIR requires the additional departments necessary to run any organisation - HR, Payroll, Legal, Finance (for invoicing) etc. I am focusing here on the main public service provision.

### Licensing and Compliance

The licensing department is concerned with providing both a free AI licence and a paid version. It is concerned with the upkeep and compliance of those licenses and seeks to identify loopholes (eg. applying for multiple developer licences rather than using a full licence) and offer remedies where appropriate.

The licence department is responsible for making sure as far as possible that the licence holder is properly complying with any AI legislation.

### AI Content Units

Each AI output generated is equal to 1 content unit.

To help track the environmental impact, AIR would be looking to establish an average cost per content unit both monetarily and environmentally.

One must consider that generating a text output is substantially less processor intensive than generating 10 seconds of video for example so there is some flexibility about what exactly should constitute 1 unit.

An alternative metric that measures x processing cycles or x resources might be better suited for this purpose.

Licence types are:

<b>Provider</b> An AI service that has data centres (Paid)
<b>Application</b> A live application (has customers) that is built on a Provider platform (Paid)
<b>Developer</b> An application in development that has no customers and is not yet live (Free)

Provider and Application licence types are paid and as such need to be properly verified eg. Company address and business information checked to ensure integrity. The provision will need to be assessed and approved before a full licence is granted.

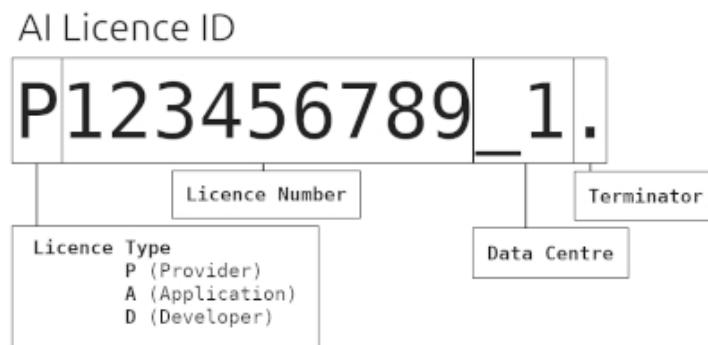
There will also be an expectation that feedback on the AI content units generated and the associated energy and water costs will be integrated with the AIR API.

Failure to provide accurate reporting about these metrics should have serious legal consequences and also include potential licence revocation.

Developer licences are provided free but are restricted to a certain amount of AI content unit generation per amount of time (per month / per day TBD). The purpose of this licence type is to encourage innovation but the expectation is that the licence be upgraded once an application has been created and made live.

## Licence ID

The licence ID is formatted as follows:



## Content Identification

Every AI output must be readily identifiable to make sure that the consumer can tell what is real and what is not.

As the content output may confuse the consumer, there should be strong legal consequences for non-compliance.

## Output Identification Patterns

The Licence ID is presented in different ways dependant on the type of content. Examples are as follows:

### Text

The first text in the output must display the licence ID as follows:

AI licence ID: P123456789.

The '.' indicates the end of the ID. This ID would then be followed by the

generated content.

**AI licence ID: P123456789.**

Lorem ipsum dolor sit amet, consectetur adipiscing elit, s  
ut labore et dolore magna aliqua. Ut enim ad minim veniam,  
ullamco laboris nisi ut aliquip ex ea commodo consequat. D  
reprehenderit in voluptate velit esse cillum dolore eu fug  
sint occaecat cupidatat non proident, sunt in culpa qui of  
est laborum.

## Image

The licence ID must appear in the top left, right or centre of the output (this is because when displayed the top of the image is prioritised over the bottom which is often cropped) and be no less than 5% of the longest side of the image.

The font must be Arial and must be easy to read eg. plain black background with white text such that a consumer can easily interpret it.

The licence ID must be in the format:

AI licence ID: P123456789.

Note that the '.' indicates the end of the identifier.

As it is a file type the licence ID must also be included in the EXIF data.

Parameter name: **AI licence ID**

Parameter value: < **licence ID eg. P123456789.** >

**AI licence ID: P123456789.**



## Audio

The licence ID must be the first audible sound on the content and may be read out by either a human or text to speech.

The licence ID must take at least 5 seconds of the output every 60 seconds such that a listener can easily interpret it.

If the output is an audio book or some other type of content where repeated licence ID announcement may become infuriating to the listener, exceptions to the 5 seconds every 60 seconds rule may be permitted. For example in the case of an audio book, announcement

once at the beginning of playback may suffice if on an electronic device that can read the Licence ID from the EXIF data or if recorded, once per chapter may be adequate.

Exceptions must be requested when providing the example output for assessment.

The audio must be spoken as follows:

AI licence ID P123456789

As it is a file type the licence ID must also be included in the EXIF data.

If the audio is a filter of some kind then the licence ID must be repeated periodically (eg. once per minute)

Parameter name: **AI licence ID**

Parameter value: **< licence ID eg. P123456789. >**



## Video

The licence ID must appear in one of the four corners of the output and be no less than 5% of the longest side of the image. It must be identifiable on every frame of the output.

The font must be Arial and must be easy to read eg. plain black background with white text such that a consumer can easily interpret it.

The licence ID must be in the format:

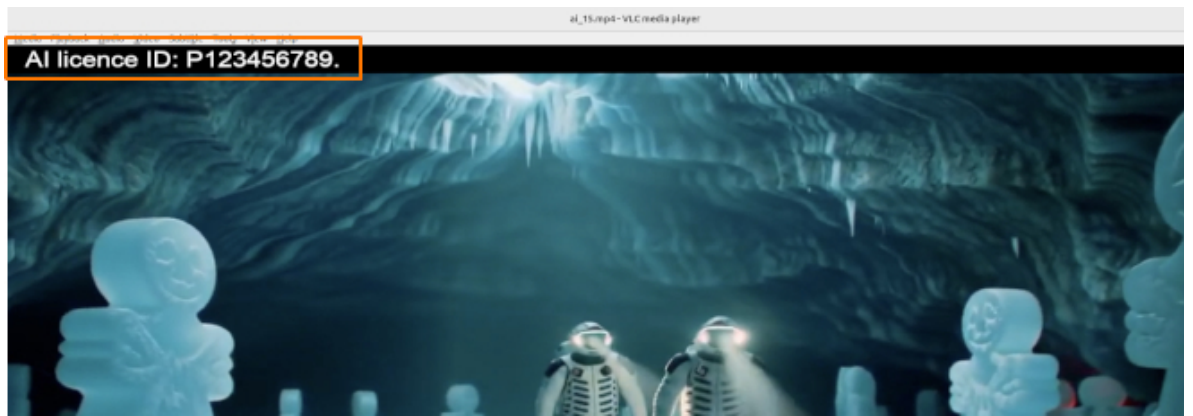
AI licence ID: P123456789.

Note that the '.' indicates the end of the identifier.

As it is a file type the licence ID must also be included in the EXIF data.

Parameter name: **AI licence ID**

Parameter value: **< licence ID eg. P123456789. >**



## Code

The first comment text in the output must display the licence ID as follows:

```
AI licence ID: P123456789.
```

The '.' indicates the end of the ID. This ID would then be followed by the generated content.

## Software

If AI has been used in the production of software then the AI licence ID must be the first comment in the code, the first comment in the manual if it's a text based application or prominently displayed on the About dialog box if a GUI as follows:

```
AI licence ID: P123456789.
```

```
# AI licence ID: P123456789.  
  
OIFS=$IFS  
IFS='@'  
em=$1  
read -ra arr <<< "$em"  
domain=${arr[1]}  
  
name=$(echo -n "$em" | md5sum | cut -d' ' -f1)  
contents=$(echo -n "$domain" | md5sum | cut -d' ' -f1)  
dnsrecord="${name}.${domain}"
```

## Output Assessment

For a full licence, the licensee must provide the appropriate output types, application link (Internet) or application name (App Store), example text for generation and an example output in their licensee account.

AIR will then review, assess and approve the application or provide

feedback if not compliant.

*App Name or Internet Link*

*Example generation text*

*Example output [ **browse for file** ]*

Text Minimum 100 characters Maximum 500 characters
Image Maximum 50mb file
Audio Maximum 50mb file
Video Maximum 100mb file
Code Minimum 100 characters Maximum 500 characters
Software Maximum 100Mb zip file

### Impact Assessment

Monitoring Data Centres and keeping a close eye on their effects to both the environment and the local community will allow for a good understanding of their cost over time. Also, as technology advances, the costs per output unit of content will likely vary. Where a Data Centre generates its own energy for example, there would be different environmental impact if a fossil fuel is used versus a renewable solution.

This is potentially quite a complex subject because there are many factors at play, especially once you start considering the effect to a local community. The additional water and energy grid supplies alone will affect local residents and until those effects can be observed in real time, we cannot speculate as to their impact.

### Private Company Concerns

The purpose of AI regulation is not to harm the profits made by private companies. The purpose is to enable profits to be made but decrease the risk to the general public so that there is no harm caused to our fundamental capitalist system.

Private company officers will remember that the capitalist system is what they rely on to operate and that damaging it will threaten their

long term prospects.

Ignoring the long term implications in favour of a short term strategy could very seriously lead to civil war.

You can't spend money when you're dead. You can't make money without a stable system.

Private company input and constructive criticism will lead to legislation that is protective of both the private sector and the general public's interests.

Responsible commentary by the private sector is to be welcomed.

## Reporting Responsibilities

All reporting will be done by integrating with the AIR API (Application Programming Interface). The suggested strategy would be periodic batch reporting rather than giving feedback as part of the AI content generation process to avoid impeding the speed of creation with unnecessary network calls.

Provider Licensees (Data centre owners) will be expected to report upon the type of AI content unit generated and some kind of processing metric.

Energy and water usage would be reported periodically and used to calculate an average for the output.

Application Licensees will be expected to report the type of AI unit generated and the parent provider licence ID so that their output may be attributed accordingly.

By collecting data about the energy and water usage, the types of energy usage and the respective processing power used it should be possible to calculate a normalised 'Impact Score' which will allow all AI content generators to be plotted on the same graph and give insight into who is providing the most optimised and least environmentally impactful AI service.

By examining the implementations of those services it will be possible to provide suppliers with solid advice on how they can improve and optimise their own service, potentially a very valuable side effect of having to maintain a licence. They also get some insight into exactly what their competitors are doing.

The vendor would need to provide their licence ID, an MD5 hash of the content created (for file identification) and a timestamp within a certain time of the creation. This is so that licence holders can keep a record of

the units created within a time period and submit them as a batch so reporting them doesn't impede the creation process. The MD5 hash is used to later identify content if there's a complaint.

At their own end, licence holders should store the content units generated against a client ID and pass on any costs incurred appropriately.

The API would expect something like.

**Header:**

```
API ID: Some longish API ID
Line of data:
    Date / Time: yyyyymmddss
    Content Type: text, image, audio, video, code or softv
    MD5 Hash: An MD5 of the content (for later use in the
    Units: Processor Content Units TBD
```



**Example JSON:**

```
{id: "api id", ["dt": "yyyyymmddss", "type": "text", "md5": "1234567",
"units": 1], ["dt": "yyyyymmdd", "type": "audio", "md5": "8910111",
"units": 15]}
```

### Complaints

The complaints department would be concerned with:

Helping consumers navigate the complaints process
Management of the complaints process
Referral to third parties where appropriate (eg. Police)
Liaising with Complainants to ensure satisfactory resolutions

The purpose of the complaints process is not to punish the licence holder but to provide them with feedback to help them to tune their application accordingly. It assumes a good faith application by the licence holder.

### Complaints Process

1. The complainant gives details of the content they object to. A licence ID (which should be on the content if from a compliant licence holder), a free text reason and either a link to the content or by uploading a file. The complaints software then attempts to link the offending content to a licence ID by generating an MD5 hash and then searching the database.

2. The complaints officer then reviews and assesses the content and has the option to revoke / suspend the licence.

Revoke completely cancels the licence and is final (this may require secondary approval by a manager to prevent malicious revocations by a bad actor) - the licence holder must completely re-apply, they are forbidden from creating content until they do.

Suspend puts the licence on hold temporarily (changes the status) until the licence holder addresses the complaint. There is also an option to refer the complaint to a third party (eg. the Police) for further investigation (issues a key pair to allow review without login).

3. The complaint is either dismissed, referred to the licence holder or referred to a 3rd party.

Either way, the complainant is kept informed at every stage of the process.

4. TBD dependant on outcome!

## Evidence Authentication

With such powerful technology, we cannot disclude AI use for nefarious purposes such as manufacturing evidence in legal cases.

The main purpose of the Evidence Authentication department is to provide a service to identify when AI has been used to generate content to help rule out its use before presentation as evidence in a court of law.

AI used in this way will make a mockery of our legal system.

### Legal

Although I'm sure that AI usage to generate content can be discerned by identification of the presence of specific algorithms within the content. A superior approach would be to feed an AI model with both AI and real content whilst informing it which is which and then ask it to discern the differences. This would provide a flexible model which would be able to grow with the technology as it improves. Periodic testing would of course be required with the software possibly being limited to giving a percentage chance that a piece of content is real.

A manual process for use in ongoing investigations and an automated process when evidence is submitted to a digital repository would be preferable so that officers can quickly ascertain if evidence is likely real or fake.

If the evidence is already stored in a digital repository then establishing its authenticity should be applied to any existing or new evidence which may trigger alternative outcomes for cases already ruled upon.

Can you see the implications already stacking up?

Both the manual and automated systems would use the same identification software with the final evidence submission result taking precedence over any manual scan previously done during the investigation to take advantage of any recent advances in the detection software.

This is a brief discussion of the legal ramifications of AI usage and its detection before presentation as evidence in a court of law and certainly requires the expertise of legal professionals to understand exactly how it must be applied to comply with existing processes.

The systems used by our legal system should be free.

### **Private Companies**

Insurance companies and other private businesses that need to authenticate evidence will have access to a paid version of the same software again with both a manual and automatic option perhaps with the option to automatically upload scanned content to their cloud services to establish a legitimate chain of custody.

The idea being that the payments made by private companies would effectively subsidise the free legal services.

[benosborne.com](#) | [Site Map](#)

2026