

Data Protection at Libration Accessibility

How the plugin processes data — information for controllers and data protection officers

Version 1.0 · Plugin 3.1 Date 19 August 2026 Libration libration.io

Who this document is for

You are using Libration Accessibility, or you are examining whether you can use it. That makes you the controller under data protection law, and you need to know what the software does with which data. This document answers that — with enough technical precision for your data protection officer to work with, and with a question-and-answer section at the end covering everything that comes up in practice.

We also name the points that require explanation, and the counter-arguments to our own assessment. A data protection information document that contains only the agreeable parts is of no help to you — your reviewing body will find the rest anyway.

The principle

Libration Accessibility draws a strict distinction between what happens when someone visits your website and what happens in your editorial team.

In visitor operation, no data leaves your server — not a single item. All content that requires processing by a service provider — easy-to-read versions, audio recordings, alternative texts — is created beforehand in your editorial team and lies ready in your own installation. A website visit triggers no processing, neither at the visitor's end nor at ours nor at any third party's.

For the most frequent question, this means: the plugin does **not** belong in your consent tool. The reasoning is set out below and is structured so that you can adopt it in your own documentation.

The three areas

First, what computes on your server. The accessibility widget, the corrections in the source code (skip links, focus indication, ARIA markup, tab order), the checklist, the alternative text audit and the contrast assistant. These functions run entirely within your own installation and establish no outbound connection.

Second, what triggers outbound connections. Alternative texts, easy-to-read versions, audio recordings, the external accessibility scan, the additional modules for subtitles and for PDF documents — and the operation of the licence itself, that is, licence and update checks as well as a notice banner in the WordPress backend. All of these operations run server-side and are initiated by logged-in persons or by scheduled tasks, never by a website visit. The complete list, with the trigger and the data transmitted, appears below.

Third, the frontend — what arrives in your visitors' browsers. Exclusively files from the plugin directory of your own WordPress installation: scripts, stylesheets, icons. In addition, the prepared easy-to-read versions from your database and the audio recordings from your own upload directory. No external address, no font loaded in from elsewhere, no tracking pixel, no embedded third-party resource.

The third area is the decisive one for your review: everything that reaches your visitors comes from your own domain.

What happens in visitor operation

No outbound connection

In the frontend, the plugin loads exclusively files from the plugin's own directory within your WordPress installation. It embeds no fonts from external servers, no icon library, no script from a content delivery network, no tracking pixels and no embedded third-party content. The frontend files delivered contain not a single external address.

Easy-to-read versions are delivered from your WordPress database. Audio recordings are held as audio files in your own upload directory and are delivered from your own domain. For both, no processing is required at the time of the visit — it has taken place beforehand.

If a prepared audio recording is missing for a text passage, the widget falls back on the speech synthesis built into the visitor's browser. It sounds more mechanical, but it does not leave the device.

The setting **"No live requests in the frontend"** is active by default and makes this state binding: where it is set, the plugin is technically unable to establish an external connection in the visitor context.

The cookies

The widget stores the display settings that a user has selected themselves in fourteen first-party cookies with the prefix `libration_`. Each contains nothing but the value selected:

- Switch states as `true` or `false`: invert colours, monochrome, dark contrast, light contrast, low saturation, high saturation, highlight links, highlight headings, screen reader, reading mode
- Percentage values: content scaling, font size, line height, letter spacing

Lifetime 30 days, `path=/`, `SameSite=Lax`. No cookie contains an identifier, a user ID, a session ID or any other value by which a terminal device could be recognised. Using the "Reset all" button in the widget, they can be deleted in full at any time.

Writing takes place only at the moment when someone activates or deactivates a function in the widget or moves a slider — nothing is set on a mere page view. **Reading** takes place on every page view, in order to restore an existing setting. We state this expressly, because Section 25(1) TDDDG (the German implementation of the ePrivacy Directive's Article 5(3)) covers not only storage but also access to information already stored. Both operations must fall under the exemption, not just one of them.

Why, in our assessment, this does not require consent

Section 25(2) no. 2 TDDDG exempts from the consent requirement whatever is "strictly necessary in order for the provider of a digital service to provide a digital service expressly requested by the user". The Orientierungshilfe Digitale Dienste (guidance on digital services) of the German Datenschutzkonferenz (DSK, the joint body of the German supervisory authorities), version 1.2, as at November 2024, interprets this exemption. Three points from it are relevant here.

The granular approach. According to para. 72 et seq., it is not the website as a whole that is decisive but the individual function; a service is to be understood as the sum of various functions. The service expressly requested here is not your website but "display the page permanently in high contrast". The user expresses this wish by clicking — that is, by precisely the objective act from which the guidance, at para. 70, infers the user's wish.

The timing. Para. 80 requires that cookies for additional functions be set only once the corresponding user interaction has actually taken place. This is exactly how the widget behaves.

The content. Para. 81 criticises cookies bearing a unique identifier and contrasts that criticism with how it can be done without such identifiers:

"For this, no a [sic] unique identifying feature such as a unique user ID is required; rather, it suffices to store an item of information that is not identifying in each case, such as 'background-color: black' or 'language: de'."

The cookies of Libration Accessibility follow this pattern. `libration_dark_contrast=true` and `libration_font_size_pct=120` are non-identifying items of information of exactly the kind the supervisory authorities describe here. We do not wish to read more into it than it says: the paragraph concerns the substantive dimension and does not grant a widget blanket exemption. It does, however, dispose of the objection most frequently raised in relation to preference cookies.

In our assessment, the exemption therefore applies to both storage and access. An entry in your consent tool is not required — and would in substance be counterproductive, because a consent banner placed in front of an accessibility function erects an additional hurdle for precisely the people it is intended for.

The point on which we can be contradicted

We regard the 30-day lifetime as the most assailable part of this assessment, and we set out the counter-arguments openly.

Para. 79 requires all dimensions to be reduced to the necessary minimum, expressly including the lifetime. Para. 80 goes further: in the case of individualised cookies, validity is said to be frequently required for one session only, and regular use can in principle be assumed only for services requiring a login. From this it can be inferred that 30 days is too long.

Two things speak against that. The wording in para. 80 refers to **individualised** cookies, that is, those bearing an identifier — ours bear none. And for a person with a visual impairment, it is not a convenience but a precondition that the display once configured is still in place on the next visit; having to set it afresh on every visit would itself be a barrier.

If your reviewing body assesses this differently, that is no obstacle. The lifetime is configurable: 30 days, 7 days, session only — or no storage at all, in which case the settings apply only to the page currently open.

What happens in the editorial team

All operations in this area run server-side: your WordPress server contacts our service, never the browser of a person involved and never the browser of a website visitor.

Operation	When	What is transmitted
Easy-to-read version	When editing or saving a page, on request or in a batch run	Published text of the page, language code, licence key
Audio recording	When a page is approved or on request	Text passages of the page, licence key
Alternative texts	On request in the media library; on upload only if enabled	Image file, file name, title of the attachment, language code, licence key
External accessibility scan	At the touch of a button or on a schedule	Address of your website, addresses of the pages checked, licence key where present
Subtitles and transcripts (additional module)	On request	Video or audio file, title, glossary, address of your website
PDF checking and remediation (additional module)	On request	PDF document, title, language code
Licence and update check	Regularly in the background	Address of your website, licence key, version
Notice banner in the backend	On calling up the WordPress backend	Language and version, can be switched off

Two points we anticipate here, because a thorough review will find them anyway: the external scan and a trial run of the PDF check work to a limited extent **even without a licence key**. The statement "without a key, nothing leaves your server" would therefore go too far, and we do not make it. None of these operations takes place in the visitor context, and none of them transmits data about your website visitors.

Every result is stored in your installation: easy-to-read versions in the WordPress database, audio recordings in the upload directory, alternative texts in the regular alternative text field of the media library. The same content is never transmitted twice.

Sub-processors and transfers to the United States

Who processes what

The service to which your server hands over the content runs at **Hetzner Online GmbH** in Germany. Processing by language models does not take place there, but at two providers:

- **Anthropic Ireland, Limited** — alternative texts, easy-to-read language, linguistic refinement and translation of subtitles, alternative texts in documents.
- **OpenAI Ireland Ltd.** — audio recordings as well as transcription and word-level timestamps for the read-aloud function and subtitles.

The contracting parties within the European Economic Area are in each case the Irish companies; the processing is carried out by

affiliated companies in the United States.

The external accessibility scan and the technical checking of PDF documents run exclusively on our servers in Germany. No further provider is used for this.

Legal basis for the transfer

Both providers are **not** certified under the EU-U.S. Data Privacy Framework. We verify this in the official participant list of the US Department of Commerce and record the current position in the data processing agreement. The European Commission's adequacy decision for the United States is therefore **not applicable** to them.

The transfer is accordingly based exclusively on the EU standard contractual clauses pursuant to Article 46(2)(c) GDPR, Implementing Decision (EU) 2021/914, Modules 2 and 3. A transfer impact assessment should therefore form part of your review; we provide the technical information needed for it.

What happens to the data at the providers

- **No training.** Both providers contractually exclude training models on content from the application programming interface.
- **Retention period.** Anthropic deletes inputs and outputs regularly within 30 days. OpenAI retains 30 days of logs for abuse detection for the endpoints used and stores no application data.
- **Processing on behalf of the controller.** Contracts pursuant to Article 28 GDPR are in place with both.

Your obligations as controller

1. **Data processing agreement** pursuant to Article 28 GDPR. It automatically becomes part of the contract upon conclusion of the contract; a signature is not required under Article 28(9) GDPR. The version in force is available in the customer area as a PDF.
2. **Entry in your records of processing activities**, insofar as you use the functions involving external processing. We supply the technical information.
3. **Addition to your privacy policy**, using the wording below.
4. **No entry in your consent tool**, on the assessment set out above.

Wording for your privacy policy

The first passage always applies. You only need the second if you use functions involving external processing.

Accessibility tools

On our website we use the plugin Libration Accessibility (provider: Libration, Holger Klopmeier, Germany). It provides accessibility controls with which you can adapt contrast, font size, line spacing and further display characteristics to your needs.

If you actively make one of these settings, the plugin stores your selection in cookies on your terminal device and reads them out again on your next visit so that the display is retained. These cookies contain nothing but the value you have selected and no identifier by which you could be recognised. They are set only after a corresponding action on your part and are deleted after 30 days. Using the "Reset all" button in the control panel, you can remove them yourself at any time. The legal basis is Section 25(2) no. 2 TDDDG, since storage and access are strictly necessary in order to provide the display adjustment expressly requested by you.

The plugin does not load any content in from external servers and does not transmit any data about your visit to third parties.

Simplified versions and speech output

For the "easy-to-read language" and "read aloud" functions we hold prepared versions of our content available. These are produced editorially in advance; in the course of this, the content of our website concerned is processed by our service provider Libration, which uses Anthropic Ireland, Limited and OpenAI Ireland Ltd. as sub-processors for this purpose. The processing takes place in part in the United States on the basis of the EU standard contractual clauses pursuant to Article 46(2)(c) GDPR. Only the content published on our website is transmitted. No transmission takes place when a page is called up; your IP address and other data about your visit are not passed on to third parties. The legal basis for this use is our legitimate interest in an accessible offering pursuant to Article 6(1)(f) GDPR.

Adjust the statement of the lifetime if you have the cookie lifetime changed.

Frequently asked questions

Consent and consent tool

Do we have to include the plugin in our consent tool?

No. No cookies are set on a mere page view, no identifiers are stored and no connections to third parties are established. The values stored are display settings that a user has selected themselves. The reasoning under Section 25(2) no. 2 TDDG is set out in full above and can be adopted in your documentation.

Our data protection officer nonetheless insists on consent. What now?

Then we switch the cookie lifetime to session only or disable storage entirely — which removes the point of contention. But bear in mind: a consent banner placed in front of an accessibility function affects precisely the people the function is intended for. We advise against it and are happy to set out the arguments in a conversation.

Does our consent tool block the plugin as long as no consent has been given?

Some consent tools block all scripts across the board unless they have been expressly released. If yours is configured that way, please add the plugin's scripts to the "strictly necessary" group. Otherwise the accessibility controls are only available once consent has been given — which is precisely what undermines accessibility.

Does this also apply to the read-aloud function and easy-to-read language?

Yes. Both are delivered from your own installation. Calling them up or listening to them creates no connection to a third party.

Contracts and documentation

Do we need a data processing agreement?

Yes, as soon as you use a function involving external processing. It has already been concluded: on taking out your plan, it automatically becomes part of the contract. Under Article 28(9) GDPR the electronic form suffices; a signature is not required.

We need a signed copy for our files.

The PDF in the customer area is the contract document and is sufficient for your files. If your organisation absolutely requires signatures, talk to us — we will find a solution.

Can we obtain a transfer impact assessment from you?

We supply the information you need for it: places of processing, retention periods, encryption, exclusion of training, contractual assurances and the providers' current sub-processor lists. The weighing-up itself is for you as controller to carry out; we cannot do it for you.

Do we have to enter the plugin in our records of processing activities?

If you use only the functions that run on your server, no processing activity of your own vis-à-vis third parties arises. As soon as you use alternative texts, easy-to-read language, read-aloud, the external scan or one of the additional modules, that belongs in your records. We provide the technical information.

Do we need a data protection impact assessment?

As a rule, no: there is no systematic monitoring, no profiling, no automated decision producing legal effects. What is processed is published content from your website. The assessment is nonetheless yours to make — in particular if your content contains special categories of personal data within the meaning of Article 9 GDPR.

How will we be informed about new sub-processors?

At least 30 days before they take effect, in text form. You may object within 14 days on important data protection grounds. If no agreement is reached, you may terminate the service concerned; we will refund fees paid in advance on a pro rata basis.

What is transmitted

Do you see our visitors' IP addresses?

No. Your visitors' browsers never contact us at any point. All connections are established by your own server. What reaches us is your server's IP address, never that of a visitor.

Does the plugin load Google Fonts or similar external resources?

No. All fonts, icons, scripts and stylesheets are held in the plugin directory of your installation. The frontend files delivered contain not a single external address.

What happens to images showing people?

They are transmitted to our service, and from there to the language model, in order to generate the alternative text, and they are not used there for training. No facial recognition takes place and no biometric data is generated. If you wish to exclude individual images, enter the alternative text by hand — existing alternative texts are never overwritten.

Our pages contain names and contact details of members of staff. Is that a problem?

This information has already been published, and it is processed within the framework of processing on behalf of the controller pursuant to Article 28 GDPR. Legally, that is covered. If you nonetheless wish to exclude individual pages, simply do not generate a simplified version for them — control lies page by page with your editorial team.

How long do you store our content?

We do not store the content itself permanently, only consumption metadata: the time and extent of use. We hold results from the additional modules available for retrieval for a maximum of 14 days. At the language model providers, the 30 days stated above apply.

What happens to our data when the contract ends?

We delete it or return it, as you choose, unless a statutory retention obligation prevents this. We confirm deletion on request. The results stored in your WordPress installation remain available to you — they are held by you, not by us.

Is our content used to train AI models?

No. Both providers exclude this contractually, and we pass this obligation on to you in our data processing agreement.

AI providers and third countries

Does a transfer to the United States take place?

Yes, if you use functions involving language model processing. We do not gloss over this. The contracting parties are Irish companies, the processing takes place in the United States, and the legal basis consists exclusively of the EU standard contractual clauses.

Does the adequacy decision for the United States apply?

No. Neither Anthropic nor OpenAI is certified under the EU-U.S. Data Privacy Framework. We verify this in the official participant list. Providers who refer to the framework across the board here should be asked for the certificate.

Is there a variant without any transfer to a third country?

Yes, if you forgo the language model functions. The widget, the corrections in the source code, the checklist, the alternative text audit and the contrast assistant run entirely on your server. You can also decide function by function: alternative texts yes, easy-to-read language yes, read-aloud no.

Can you offer processing entirely within the EU?

For part of it we are working on this: for speech output, a European region exists at OpenAI. For text processing, EU processing is currently not technically selectable for the model used. If exclusively European processing is a condition for you, talk to us — we will then examine a European-hosted model for your installation.

What happens if the Data Privacy Framework is struck down?

For us, nothing changes, because we do not rely on it. The standard contractual clauses continue to exist independently of it. Such a development may nonetheless affect the risk assessment in your transfer impact assessment.

Operation, configuration and review

Can we verify what the plugin actually does?

Yes. The plugin is licensed under the GPL, the complete source code lies open in your installation and can be examined. The frontend files are the quickest way in: anyone searching there for external addresses will find none. On request, we will go through the source code together with your reviewing body.

Does the plugin also run if our server does not permit outbound connections?

Yes, with limitations. The widget, the corrections in the source code, the checklist, the alternative text audit and the contrast assistant need no outbound connection. The functions involving external processing, as well as the licence and update check, do need one.

Can we switch off the notice banner in the backend?

Yes, in the settings. After that, the plugin no longer retrieves anything when the backend is called up.

Does the plugin set cookies in the WordPress backend?

No. The fourteen [libration](#) cookies belong to the visitor widget. In the backend, the plugin works with the regular login mechanisms of WordPress.

Are visitors' IP addresses processed anywhere?

Not in the default configuration. If you expressly permit live requests in the frontend, the plugin creates a counter to protect against abuse, which holds the IP address exclusively in hashed form and for no more than one hour on your own server. This processing does not leave your server.

Who decides which pages are processed?

Your editorial team, page by page. You determine which content types are covered at all, and you approve every version before it goes public. Nothing is transmitted automatically and unasked.

Can we edit the AI versions before they go online?

Yes. Simplified versions are created as drafts, are initially not visible to visitors and can be reworked in full. Only on approval do they become public; editing and approval can be assigned to different people. Two fields record who checked the version and when — for organisations with their own review group for easy-to-read language, this is the decisive point.

Does the plugin make our website automatically compliant with the law?

No, and we advise caution with providers who promise that. The plugin remedies many barriers automatically, documents the current position and guides you through the remaining steps. Full conformity also depends on content, structure and design.

Note on the binding nature of this document

We are a software provider, not lawyers. The assessment in this document is based on the wording of the statute, on the Orientierungshilfe Digitale Dienste of the Datenschutzkonferenz in its November 2024 version and on a review of our own source code. The technical information is robust and verifiable, and we stand behind it. The legal appraisal is for you and your data protection officer. We are available for follow-up questions, and gladly also in direct conversation with your reviewing body.

Contact: Holger Klopmeier, Libration, contact@libration.io. The data processing agreement, the terms of use and the current version of this document are available at libration.io.