

Data Processing Agreement

Offer signed by the Processor, Better ISMS EURL (RCS Paris 943 594 648), trading as ISMS Copilot, heyGRC and Aevral

Version	2026-10-26-r2
Effective date	2026-10-26
Source fingerprint (SHA-256)	fed4cacb8b5a5dd21ef8c6fa5dc02fbb180963a591c08bf784f7b150f2909c27
Template revision	2
Approved for issue on	2026-10-26
Verify at	https://trust.ismscopilot.com/en/dpa/versions

Products the Customer uses (for information; this DPA covers every product listed here that the Customer uses, ticked or not):

- ☐ ISMS Copilot
- ☐ heyGRC
- ☐ Aevral

Customer

Legal name	
Registered address	
Registration number	
Account or organization name in the service (optional)	

How to accept: tick the products, complete the Customer fields here and on the signature page, sign, and send the signed copy to privacy@ismscopilot.com. It takes effect when we receive it (DPA §7.4(a)). The agreed terms are the DPA text identified by the source fingerprint and the cover and signature terms identified by the template revision, both published at the address above. Completing the designated fields does not change them. This offer lapses as set out in DPA §7.4(f).

Data Processing Agreement (DPA): ISMS Copilot

Overview

This Data Processing Agreement ("DPA") forms part of the Terms of Service between you (the "Customer" or "Data Controller") and Better ISMS EURL (RCS Paris 943 594 648), trading as ISMS Copilot, heyGRC and Aevral, represented by its gérant ("ISMS Copilot", the "Processor" or "Data Processor"), for the use of the ISMS Copilot, heyGRC and Aevral services. This DPA complies with Article 28 of the General Data Protection Regulation (GDPR) and governs the processing of personal data on behalf of the Customer.

Effective Date: 2026-10-26. For the revision history of this document, see Appendix B (Change Log). This DPA automatically applies to all ISMS Copilot customers processing personal data through the platform. This DPA applies through acceptance of the Terms of Service and use of the Services, without a separate signature. A Customer may also countersign the copy signed by the Processor; §7.4 sets out when that copy takes effect and how it fixes the version that applies to that Customer.

Who This Is For

This Data Processing Agreement is for:

- Organizations using ISMS Copilot to process personal data
- Compliance consultants handling client data through the platform
- Data Protection Officers conducting vendor assessments
- Legal and procurement teams evaluating data processing arrangements
- Auditors reviewing GDPR Article 28 compliance

Definitions

- "Customer" or "Data Controller": The organization or individual subscribing to ISMS Copilot services and determining the purposes and means of processing personal data.
- "Processor" or "Data Processor": Better ISMS EURL (RCS Paris 943 594 648), trading as ISMS Copilot, heyGRC and Aevral ("ISMS Copilot"), processing personal data on behalf of the Customer.
- "Customer Personal Data": Any personal data processed by ISMS Copilot on behalf of the Customer, including conversation content, uploaded documents, and associated metadata.
- "Sub-processor": Any third-party processor engaged by ISMS Copilot to process Customer Personal Data.
- "Data Subject": The identified or identifiable natural person to whom Customer Personal Data relates.
- "Processing": Any operation performed on personal data, including collection, storage, use, disclosure, or deletion.
- "Personal Data Breach": A breach of security leading to accidental or unlawful destruction, loss, alteration, unauthorized disclosure of, or access to Customer Personal Data.
- "Default AI Routing Path": AI processing for a request when Advanced Data Protection Mode is not enabled, served within the non-ADP routing envelope described in §3.1 (Anthropic and the OpenRouter allowlist, with Mistral as circuit-breaker failover). It does not refer to any single provider that is the current default within that envelope.
- "Advanced Data Protection Mode" (ADP): The account- or organization-level setting that routes all AI processing to Mistral AI (EU) with zero retention, bypassing the OpenRouter routing path and the Default AI Routing Path. Zero retention refers to the AI provider: Mistral keeps no copy of the requests. ISMS Copilot still stores Customer conversations in its own database according to the configured retention period (§4.4; minimum 1 day; memories, workspace instructions and text extracted from files

are not covered by that setting), and deleted conversations can remain in encrypted database backups for up to 91 days (§2.8).

1. Scope and Applicability

1.1 Application of DPA

This DPA applies to all processing of Customer Personal Data by ISMS Copilot in the course of providing the platform services described in the Terms of Service. It also applies to heyGRC (the heyGRC GitHub App and the heyGRC console), as described in §1.2, §1.5, §1.6, §2.4 and §2.8, and to Aevral, the Better ISMS code security product at app.aevral.com, as set out in §1.7.

1.2 Subject Matter of Processing

ISMS Copilot processes Customer Personal Data to provide AI-powered compliance assistance, including:

- Processing user queries and generating AI responses
- Storing conversation history and context
- Analyzing uploaded compliance documents
- Maintaining workspace configurations and custom instructions
- Automated content moderation of chat messages (see §2.8 for retention)
- Classifying whether an eligible authenticated request requires current external information and, where permitted by Customer settings and policy, retrieving request-scoped web evidence or the exact public URL supplied by the user
- For paid Customers whose organization owner has installed the heygrc Slack bot: receiving messages addressed to the bot from the Customer's Slack workspace, processing them through the same AI pipeline as web chat, and posting AI responses back to the workspace
- For Customers who install the heyGRC GitHub App (a distinct product from the heygrc Slack bot above): receiving pull-request diffs and metadata from the Customer's authorized GitHub repositories, redacting secret-shaped content and whole sensitive files before any model call, generating an automated compliance-review comment and posting it back to the pull request, using the Customer's company profile and commitments as context. Diffs are used only to produce the review and are not persisted as a record; the generated review (summary and findings, which quote the changed lines a finding points at) and the pull-request and commit identifiers are stored (see §1.5, §1.6, §2.4 and §2.8). /heygrc replies are processed the same way and not stored.
- For Customers who install the Aevral GitHub App: scanning the repositories the Customer authorizes for security weaknesses, reviewing the Customer's pull requests, and posting advisory results back to GitHub (§1.7)
- For heyGRC Customers whose organization owner or admin has enabled the heyGRC context layer: receiving compliance context objects (policy sections, controls with framework mappings, vendor and third-party records, risks, data categories and processing activities) that the Customer's own coding agent sends to the heyGRC API, storing them with their revision history, compiling them into short obligations with verified quotes, and using those obligations to review the Customer's pull requests

1.3 Duration of Processing

Processing continues for the duration of the Customer's active subscription and according to the Customer's configured data retention period (1 day to 7 years, or "keep forever"). Upon termination, all Customer Personal Data is deleted within 30 days unless longer retention is required by law or by §2.8 (moderation retention exception).

1.4 Nature and Purpose of Processing

- Nature: Automated processing using AI models, database storage, file processing, and content

moderation

- Purpose: Provide compliance guidance, document analysis, policy generation, and knowledge management as instructed by the Customer

1.5 Categories of Data Subjects

- Customer's employees and authorized users
- Customer's clients and end-users (when mentioned in uploaded documents or queries)
- Individuals referenced in compliance documentation
- Security incident subjects
- Slack workspace users (heygrc bot only): members of the Customer's Slack workspace who interact with the heygrc bot. These users typically do not hold an ISMS Copilot account; their messages addressed to the bot are processed under the Customer's organization (Slack-originated traffic rolls up to the organization owner for billing and quota purposes, and inherits the organization's Advanced Data Protection setting). Customer remains responsible for informing its Slack workspace users that messages addressed to the bot are processed by ISMS Copilot.
- Individuals appearing in pull requests (heyGRC GitHub App only): contributors and other individuals whose names, e-mail addresses, usernames or other personal data appear in pull-request titles, descriptions or changed lines (for example inside code comments, configuration, test fixtures or documentation) in repositories where the Customer has installed the heyGRC GitHub App, and GitHub users who ask /heygrc questions. heyGRC reviews code changes, not people, and does not fetch commit author metadata; any personal data present is incidental to the diff. The Customer is the controller of that content, determines its lawful basis and informs its contributors; ISMS Copilot processes it as processor on the Customer's documented instructions (installation and configuration of the App).
- Aevral only: the Customer's developers and contributors, including outside contributors to an installed repository, and any individual whose personal data appears in the Customer's code, pull requests or configuration.
- Individuals named in heyGRC compliance context (heyGRC context layer only): the Customer's staff and contractors, and contact persons of its vendors and third parties, where the Customer's policies, controls, vendor records or risks name them. Their presence is incidental; heyGRC does not need personal data to work. The Customer is responsible for informing these individuals (Article 14 GDPR).

1.6 Categories of Personal Data

- User account information (email addresses, authentication credentials)
- Conversation content and AI interactions
- Uploaded document content (policies, procedures, audit reports)
- Workspace configurations and custom instructions
- Usage metadata and timestamps
- Special category data (Article 9 GDPR) and personal data relating to criminal convictions and offences (Article 10 GDPR) may be incidentally present, but are excluded from instructed processing on the Default AI Routing Path; intended workloads require Advanced Data Protection Mode (§4.2)
- Slack integration data (heygrc bot only, paid Customers who install): Slack workspace metadata (team ID, team name, bot user ID), the OAuth bot token issued by Slack, the installer's ISMS Copilot user ID (audit trail), Slack message content of messages addressed to the bot (DMs to heygrc or @heygrc channel mentions only — no other workspace messages are read), and Slack user identifiers (slack_user_id) of users who interact with the bot
- Pull-request review data (heyGRC GitHub App only, Customers who install): the content of pull-request diffs (file patches), file paths, pull-request number / title / description, an optional repository-provided .heygrc.md context file, the Customer's company profile and commitments entered in the

heyGRC console, the text of /heygrc questions and the code hunk GitHub attaches to them, and GitHub organization / repository / installation identifiers. Secret-shaped content and whole sensitive files are stripped before any model call, and model output is redacted again before it is posted or stored. The diff is used only to produce the review and is not persisted as a record; the generated review (summary and findings, which may quote the changed lines a finding points at) and the pull-request and commit identifiers are stored. This is distinct from the Slack integration data above (a different product).

- heyGRC compliance context (Customers who enable the context layer only): the text and small structured fields of the context objects the Customer's coding agent sends (policy sections, controls, vendor and third-party records, risks, data-category descriptions), their source, upstream identifier, version and classification, and the obligations compiled from them. Personal data in this content is incidental (names, roles and business contact details). Signature and approval records are not sent by heyGRC's setup skill, and objects the Customer's source classifies as secret are refused. Objects without a classification are accepted and treated as confidential (never quoted).
- Aevral only: repository content the Customer authorizes (source files, diffs, pull request titles and bodies, file paths), which may incidentally contain personal data such as names, email addresses or usernames; findings derived from it, including short code excerpts and suggested fixes; GitHub organization and installation identifiers.

Special category data (Article 9 GDPR) and criminal-offence data (Article 10 GDPR) are excluded from the instructed processing on the Default AI Routing Path (see §4.2). Customers with special-category or criminal-offence workloads must use Advanced Data Protection Mode (Mistral, EU). This instruction concerns the ISMS Copilot chat platform; for the heyGRC GitHub App, see the next paragraph; for Aevral, which has no Advanced Data Protection Mode, see §1.7. Customer is responsible for not submitting such data to the Default AI Routing Path and for any Article 9(2) condition or, for Article 10 data, the Article 6(1) basis and legal authorization with appropriate safeguards required by Article 10, applicable to its own processing.

Special-category (Article 9) and criminal-offence (Article 10) data may appear incidentally in a diff; heyGRC has no mechanism to detect or exclude it, and the Customer should not connect repositories where such data is expected in changed lines. The Default AI Routing Path exclusion in §4.2 applies to heyGRC's default path.

1.7 Aevral

This section applies when the Customer uses Aevral. Where it differs from the rest of this DPA, this section prevails for Aevral.

- Instructions. Installing the Aevral GitHub App, selecting repositories, starting or scheduling a scan, and leaving pull request review enabled are the Customer's documented instructions. Aevral reads repository content and pull request diffs, sends redacted file contents, pull request diffs and changed files to the AI models listed on the Aevral sub-processor list (<https://trust.ismscopilot.com/en/aevral/subprocessors>), stores findings, and posts advisory checks and comments to the Customer's GitHub repositories. Aevral does not commit, push or merge code. Installing the Aevral GitHub App on a GitHub organization or account is use of the Services and constitutes acceptance of the Terms of Service and this DPA on behalf of the person or entity that controls that GitHub organization or account (the Customer); the person installing confirms they can act for it. Reviews may start on the next pull request.
- Customer authorization. The Customer confirms it has the right to have Aevral process the repositories it authorizes, and that those repositories may incidentally contain personal data, including personal data of contributors who are not Aevral users. The Customer is responsible for informing those individuals where required. Aevral is not designed to process special category or criminal-offence data; the Customer must not rely on Aevral for repositories whose purpose is to hold such data.
- Sub-processors. Aevral uses the sub-processors on the Aevral sub-processor list (<https://trust.ismscopilot.com/en/aevral/subprocessors>). §2.4 applies to changes to that list. Advanced Data Protection Mode and the chat routing rules in this DPA do not apply to Aevral.
- AI tools the Customer connects. Connecting an AI tool to Aevral through the Aevral MCP server, by signing in with Aevral or with an agent key, is the Customer's instruction to send that tool the Aevral

data it requests for the organization chosen: organization and repository names, scan and pull request review results with their identifiers, commits, times and progress, findings including code excerpts, and plan and usage counts. The Customer chooses the tool and its provider, which receives that data under the Customer's own agreement with it and is not a sub-processor of ISMS Copilot; the Customer is responsible for that choice, including any transfer outside the EEA by that provider. Disconnecting the tool, or revoking the key, ends its access at once.

- GitHub. GitHub, Inc. hosts the Customer's repositories and receives Aevral's checks and comments under the Customer's own agreement with GitHub. In that role GitHub is not a sub-processor of ISMS Copilot; the same applies to the heyGRC GitHub App.
- International transfers. Aevral's AI processing may take place in the United States. The first transfer, to OpenRouter, relies on OpenRouter's Standard Contractual Clauses (Module Two); the onward transfers to Together AI and Fireworks AI rely on those providers' own Standard Contractual Clauses. Supplementary measures: a closed three-host pin, zero data retention and no data collection required on every request, and TLS in transit. OpenRouter does not pin the region of each request.
- Retention and deletion. Raw repository content is processed in memory and not stored. Findings and review results are stored in the EU. On the Customer's request to privacy@ismscopilot.com, at any time, ISMS Copilot will, within 30 days and at the Customer's choice, return the Customer's stored findings and review results, delete the Customer's Aevral data, or both. If the Customer uninstalls the Aevral GitHub App, ISMS Copilot keeps the Customer's Aevral data for 12 months so that a returning Customer finds its history, then deletes it. Backup copies follow §2.8. ISMS Copilot keeps the record of the installation and of the acceptance of these terms as evidence of the contract for the applicable limitation period, and retains data longer only where Union or Member State law requires it.

2. Processor's Obligations (Article 28(3) GDPR)

2.1 Processing Instructions

ISMS Copilot shall process Customer Personal Data only on documented instructions from the Customer, including:

- Instructions provided through the platform interface (queries, document uploads, workspace configurations)
- The user's express request to search, the Customer's web-search policy, and the user's selected Beyond URL scope
- Data retention settings configured by the Customer
- Advanced Data Protection Mode selection (EU-only vs. default AI processing)
- Deletion requests submitted through the platform or to privacy@ismscopilot.com
- For the heyGRC context layer: the Customer's act of creating a heyGRC API key with the `context:write` scope (only an organization owner or admin can do so) and each request the Customer's coding agent makes with that key. The Customer instructs ISMS Copilot to process the content its agent sends. The agent, and the systems it reads from (for example the Customer's compliance platform or Google Drive), are chosen and operated by the Customer; ISMS Copilot does not engage them, does not fetch from them on this path, and is not responsible for the agent's access to them. ISMS Copilot supplies the setup skill; the Customer runs it, reviews its pre-push manifest and decides what is sent. This path is separate from the heyGRC Google Drive connection (§2.4), under which heyGRC itself reads files the Customer picks. The Customer controls who holds the key; any request made with a valid key is treated as the Customer's instruction until the key is revoked.

Independent-controller disclosures (safety and legal reporting). Notwithstanding the instruction-only rule above, where ISMS Copilot becomes aware through content moderation of a credible threat to a person's life or safety, a suspected serious crime, or specific high-severity illegal content that meets its reporting threshold (such as child sexual abuse material or terrorism), it may disclose limited personal data to law enforcement, judicial authorities, or emergency services on its own initiative, as required or permitted by law. For such disclosures ISMS Copilot acts as an independent controller, not as the Customer's processor.

These disclosures are rare, decided by a person, minimised to what is necessary, and governed by an internal authority-reporting procedure. See also the moderation retention exception in §2.8.

Prohibited Processing. ISMS Copilot enforces the following prohibitions through layered contractual and account-level controls:

(a) No training of AI models on Customer Personal Data. Enforced via Anthropic's and Mistral's commercial API terms (which prohibit training on Customer Content), and via account-level "Free Training Disallowed" and "Paid Training Disallowed" flags at OpenRouter, applied to all eight allowlisted underlying providers (Inception, DeepInfra, Cerebras, Google Vertex, Together AI, Fireworks AI, Nebius, and xAI via the "xAI (ZDR)" endpoint), layered on xAI's published API no-training default for the Grok route.

(b) No publication of model outputs derived from Customer Personal Data. OpenRouter "Free Publication Disallowed" is set at the account level.

(c) Retention by AI sub-processors is minimized to the request lifetime where contractually possible. OpenRouter Zero Data Retention is mandatory at the account level - per OpenRouter's published policy, ZDR-mandatory accounts can only route to endpoints with a Zero Data Retention policy. Existing approved Mistral inference routes use the documented commercial zero-retention posture. That statement does not extend to Mistral's web-search feature. Mistral web search runs on Mistral's stateful Conversations/Agents endpoint, which Mistral excludes from its zero-retention posture, and uses Brave Search (Brave Software, Inc., United States) as its web-search sub-processor. A non-ADP authenticated search therefore sends a minimized latest-request query from Mistral (EU) to Brave (United States); that query hop is not EU-only and not zero-retention, and standard (non-zero) retention applies. For this reason authenticated web search is not available under Advanced Data Protection Mode: it fails closed. Under Anthropic's standard commercial API terms, inputs and outputs are ordinarily deleted within approximately 30 days, content flagged by its safety systems may be retained up to 2 years, and safety-classification scores up to 7 years; this data is not used for model training. Customers requiring confirmed EU processing and zero retention can enable Advanced Data Protection Mode to route permitted AI inference through Mistral's EU regional endpoint; authenticated web search then fails closed, while exact user-selected page/document fetch remains available through EU Fly.

(d) PRC-jurisdiction provider blocklist. Alibaba Cloud International, Baidu Qianfan, DeepSeek, Moonshot AI, Xiaomi, and Z.AI are all blocked at the OpenRouter account level. The control is jurisdiction-based and enforced through provider selection and the account-level blocklist: based on the reviewed published deployment documentation, none of the allowlisted providers was found to operate PRC or Hong Kong infrastructure, though OpenRouter does not pin the per-request inference region. This is a Schrems II-style supplementary measure aligned with EDPB Recommendations 01/2020.

If ISMS Copilot believes an instruction violates GDPR or other data protection laws, we will immediately inform the Customer and have the right to suspend processing until the instruction is confirmed or modified. If Customer confirms an instruction that ISMS Copilot reasonably believes violates applicable data protection law, ISMS Copilot may refuse to execute the instruction and, if the disagreement cannot be resolved, terminate the affected processing activities with 30 days notice.

2.2 Confidentiality of Processing

ISMS Copilot ensures that all persons authorized to process Customer Personal Data:

- Are subject to confidentiality obligations (contractual or statutory)
- Receive appropriate training on data protection
- Access data only on a need-to-know basis
- Follow documented data handling procedures

2.3 Technical and Organizational Measures (Article 32 GDPR)

ISMS Copilot implements appropriate technical and organizational measures to ensure a level of security appropriate to the risk, including:

Access Control Measures:

- Row-level security in database preventing cross-user data access
- User authentication required for all protected resources
- Workspace isolation preventing cross-contamination of client data
- Multi-factor authentication (MFA) support
- Automatic session timeout controls

Encryption Measures:

- TLS 1.3 encryption for data in transit
- Database encryption at rest
- Password hashing using industry-standard algorithms (irreversible)
- Encrypted file storage in Supabase

Data Minimization Measures:

- Only essential data collected (email, messages, files)
- No unnecessary demographic or contact information collected
- Analytics configured with `sendDefaultPii: false`
- Customer-controlled retention periods with automated deletion
- Web-search queries are derived only from the latest user request, bounded to the selected provider's limits, and exclude files, memories, workspace context, account identifiers, and user/device identifiers; credentials and signed URLs fail closed
- Mistral's generated search prose, raw discovery response, and request-scoped fetched page text are not persisted in conversation history, Beyond run records, detector telemetry, or application logs

Availability and Resilience:

- Automated database backups, including a daily offsite copy made on Fly.io (EU, Paris) and encrypted there with public-key encryption before upload (the private decryption keys are kept separately and are never on the backup machine or at Backblaze), then stored in Backblaze B2 (EU), where each copy is kept for 90 days, then hidden and deleted one day later
- Disaster recovery procedures
- 24/7 monitoring and alerting via Sentry
- Real-time uptime monitoring via self-hosted status monitors: 14 checks at 60-second intervals across public endpoints covering all product surfaces, with outage and recovery events sent to Sentry for alerting
- Public status page for transparency (status.ismscopilot.com)
- Multi-provider AI failover: Anthropic → Mistral via circuit breaker for any non-ADP traffic served by Anthropic; for the paid Fast and Think routes now defaulting to GLM-5.3 via OpenRouter (Together AI, Fireworks AI and Inceptron), the pre-first-token failover order is GLM-5.3, then Anthropic (Opus for Think, Sonnet for Fast), then Mistral (EU) if Anthropic's circuit breaker is open; the Beyond assistant retries on the same GLM-5.3 model and does not fail over to Anthropic; OpenRouter aggregator-level failover across the allowlisted providers for OpenRouter-served traffic (the Essential subset is Google Vertex and Cerebras; other non-ADP plans span the eight allowlisted providers)

Testing and Evaluation:

- Regular security assessments
- Continuous error monitoring and logging
- Automated data deletion testing
- Access control verification procedures

For detailed technical and organizational measures and the per-activity Article 30 processing inventory, refer to our Register of Processing Activities (RoPA) (<https://trust.ismscopilot.com/ropa>).

2.4 Sub-processor Engagement

General Authorization. Customer provides general authorization for ISMS Copilot to engage the sub-processors listed in this §2.4 and in the sub-processor list of the Trust Center (<https://trust.ismscopilot.com>), including providers marked "standby" there, for the processing of Customer Personal Data, subject to the conditions in this section. A standby provider is a listed provider that ISMS Copilot does not currently use for Customer Personal Data but is genuinely likely to use, for example as a backup or as an underlying provider reached through a listed intermediary such as the OpenRouter allowlist; each is listed with its location, function and safeguards. ISMS Copilot may allocate processing among listed sub-processors, active or standby, and fail over between them, and keeps their status current in the Trust Center. A provider that is not yet listed is added, including as standby, only under "Changes to Sub-processors" below.

Active Sub-processors. The following sub-processors actively process Customer Personal Data:

Sub-processor	Role / when invoked	Location	DPA / Transfer mechanism
Supabase (PostgreSQL + Storage)	Database and file storage (always)	EU (Frankfurt)	GDPR-compliant
AWS	Underlying infrastructure for Supabase	EU-Central-1 (Frankfurt)	GDPR-compliant
Backblaze (B2 Cloud Storage)	Encrypted offsite database backups (always). Each daily copy is encrypted by ISMS Copilot before upload; Backblaze does not hold the key	EU (Amsterdam, Netherlands). Backblaze, Inc. is a United States company	Backblaze DPA with SCCs. Encryption with the key withheld is the supplementary measure for any access from the United States. Each copy is kept for 90 days, then hidden and deleted one day later
Anthropic Claude	AI processing for any non-ADP plan (ADP off), as one of the permitted non-ADP providers, and the circuit-breaker failover context. A non-ADP request may instead be served by the OpenRouter allowlist (see the OpenRouter row), for any chat mode. Current role: Anthropic is not a default destination for any plan; it is the pre-first-token failover (Opus for Think, Sonnet for Fast) for the paid Fast and Think modes, whose current default is GLM-5.3 via OpenRouter (Together AI, Fireworks AI and Inception), and it is not used for web search; ISMS Copilot may serve any non-ADP paid request, in any mode, via any of the eight allowlisted providers or Anthropic and may move cohorts or modes between them over time (allocation among listed sub-processors under §2.4: no advance notice)	United States	SCCs; no training under Anthropic's commercial API terms; retention under Anthropic's standard commercial API terms (ordinary ~30-day deletion; safety-flagged content up to 2 years; safety-classification scores up to 7 years; not training). Customers needing zero retention can enable Advanced Data Protection.
OpenRouter (routing aggregator)	AI processing for free / null-plan users with ADP off (currently z-ai/glm-5.3-flash pinned to Together AI,	United States	OpenRouter's role is account-level enforcement (mandatory ZDR, training-disallowed, allowlist, PRC-

	Fireworks AI and Inceptron, within the allowlisted underlying providers below); for the Essential plan with ADP off (restricted to a closed two-provider subset: Google Vertex and Cerebras); and for paid plans (Plus, Standard, Pro, Business) with ADP off, where GLM-5.3 (z-ai/glm-5.3-flash for Fast, z-ai/glm-5.3 for Think and the Beyond assistant), pinned to Together AI, Fireworks AI and Inceptron, is the current default destination, with Anthropic retained as the automatic pre-first-token failover for Fast and Think, and ISMS Copilot may serve any non-ADP paid request, in any chat mode, via any of the eight allowlisted underlying providers or Anthropic, and may move any cohort or mode between these destinations over time (such moves among listed sub-processors are allocation under §2.4: no advance notice)		blocklist, see below); legal transfer mechanism for Customer Personal Data leaving the EU is anchored at the underlying-provider layer (per-provider rows below)
↳ Inceptron (Inceptron AB)	OpenRouter underlying provider (allowlisted)	Sweden (EU)	SCCs (EU HQ, but OpenRouter does not pin inference region, so treated as a possible transfer); ZDR + no-training enforced via OpenRouter account config
↳ DeepInfra	OpenRouter underlying provider (allowlisted)	United States	SCCs; ZDR + no-training enforced via OpenRouter account config
↳ Cerebras	OpenRouter underlying provider (allowlisted)	United States	SCCs; ZDR + no-training enforced via OpenRouter account config
↳ Google Vertex	OpenRouter underlying provider (allowlisted)	United States	SCCs + EU-US Data Privacy Framework certification; ZDR + no-training enforced via OpenRouter account config
↳ Together AI	OpenRouter underlying provider (allowlisted, added 2026-05-25)	United States (default routing to North America inference data centers per Together docs; region pinning not exposed by OpenRouter)	SCCs; ZDR + no-training enforced via OpenRouter account config; SOC 2 Type II; published DPA
↳ Fireworks AI	OpenRouter underlying provider (allowlisted, added 2026-05-25)	United States (multi-region fleet: US, EU Frankfurt + Iceland, APAC Tokyo only; no PRC or Hong Kong infrastructure identified in Fireworks' published deployment docs; region pinning not exposed by OpenRouter)	SCCs; ZDR + no-training enforced via OpenRouter account config; SOC 2 Type II; published DPA
↳ Nebius	OpenRouter underlying provider (allowlisted, added 2026-05-25)	Netherlands HQ; primary inference in Finland (EU) with US secondary per Nebius docs (region pinning not exposed by OpenRouter)	SCCs; ZDR + no-training enforced via OpenRouter account config; published DPA + sub-processor list

↳ xAI (X.AI LLC), "xAI (ZDR)" endpoint, serving the Grok models	OpenRouter underlying provider (allowlisted). Not currently a default destination for any plan or mode. Since the GLM-5.3 rollout (Appendix B, 2026-09-01) it serves the logged-out risk-analysis demo and remains available as the operator rollback route for the paid Fast and Think modes and the Beyond assistant (it was the paid default from 2026-07-21). Not used for free / null-plan routing or under Advanced Data Protection	United States (OpenRouter-published provider footprint; OpenRouter does not pin the per-request inference region)	SCCs at the underlying-provider layer via xAI's published DPA (EU SCCs, Modules 2/3, Irish-law), with OpenRouter Article 28(4) sub-processor flow-down; zero data retention as configured and verified 2026-07-19 (OpenRouter account-level ZDR via the "xAI (ZDR)" endpoint); no-training per the OpenRouter account-level training-disallowed setting layered on xAI's published API no-training default
Mistral AI	AI processing for ADP users (any plan); circuit-breaker failover destination for paid Anthropic; content moderation for all users; conversation compaction; conversation summaries; search-intent classification; and provider-attested source-URL discovery for eligible ADP-off authenticated search. Mistral's generated search prose is discarded. Bounded exact-fetched evidence goes only to the answer provider already applicable to the conversation or Beyond run. Authenticated web search fails closed under ADP; heyGRC EU inference (organization opt-in, EU regional endpoint, account-level zero retention)	Mistral inference, moderation, embedding and summarization requests are sent to Mistral's EU regional endpoint (api.eu.mistral.ai), which processes them in the EU. The one exception is Mistral web search: it runs on Mistral's global Conversations endpoint, which has no regional processing commitment, so web-search discovery is not EU-only. For ADP-off authenticated search, Mistral's web-search discovery runs on Mistral's stateful Conversations/Agents endpoint, which is outside Mistral's zero-retention posture, and Mistral engages Brave Search (Brave Software, Inc., United States) as Mistral's own web-search sub-processor (see Mistral's published sub-processor list). That discovery path is therefore not EU-only and not zero-retention	Existing inference/moderation routes: EU residency, zero retention, no training under Mistral's commercial API terms. Mistral web search is excluded from that zero-retention posture; via Mistral's web-search chain the minimized query reaches Brave (US) under standard, non-zero retention per the applicable Mistral/Brave provider terms, so web-search discovery is disclosed as a US transfer via Mistral rather than an EU zero-retention route. ISMS Copilot does not separately engage Brave for this launch route: Brave is engaged by Mistral, as Mistral's sub-processor (see the sub-processor chain transparency note below).
Stripe	Payment processing	Global (EU DPA)	GDPR-compliant; PCI DSS Level 1
ConvertAPI	Document format conversion	EU endpoint	GDPR-compliant; ISO 27001:2022; signed DPA
Sentry	Error monitoring	Germany	GDPR-compliant
Axiom (Axiom, Inc., United States)	Log storage: application logs of the production services hosted on Fly.io, kept 30 days for security monitoring and incident investigation. The logs carry pseudonymous identifiers and technical metadata, not chat or document content	EU (Frankfurt, AWS eu-central-1). Axiom states that some of its operations may be processed outside the EU	Axiom DPA with SCCs (Module Two)
Chatbase (Chatbase, Inc., United States)	Optional website support chat on the ISMS Copilot and heyGRC websites, documentation, consoles and trust center: visitor messages and an optional lead email for human follow-up. Not used for product chat accounts or pull-request source code	United States	Chatbase DPA with EU and UK SCCs

Vercel	Frontend hosting	Global CDN	GDPR-compliant
Fly.io	Backend API hosting, chat orchestration, request-scoped exact fetching of user-submitted or Mistral-attested public URLs, and the daily database backup (the copy is encrypted on the backup machine before upload to Backblaze)	EU deployment	Executed DPA, EU SCCs
SendGrid (Twilio)	Transactional email	United States	SCCs
Kit (ConvertKit)	Onboarding email	United States	SCCs

Every reference to ZDR / zero retention in the OpenRouter rows above is the account-level Zero Data Retention described in §3.1: no persistent retention of request content beyond serving the request; a provider may hold content transiently in memory for the duration of processing, as is inherent to serving an inference request.

Sub-processor chain transparency (external recipient via Mistral). For the launch authenticated web-search route (ADP off), ISMS Copilot does not separately engage Brave and does not list Brave as one of its own Article 28 sub-processors. Brave Software, Inc. (United States) is engaged by Mistral as Mistral's own web-search sub-processor; the minimized latest-request query reaches Brave only through Mistral's web-search chain. Mistral publishes its own sub-processor list, which identifies Brave as its web-search sub-processor, and customers can review it there. For transparency, we disclose this onward Mistral (EU) to Brave (US) query hop as an external recipient reached via Mistral: it is not EU-only and not zero-retention, and standard, non-zero retention applies per the applicable Mistral/Brave provider terms. This launch route is distinct from the future Brave-direct route (see "Notified external search recipient" below), in which ISMS Copilot would itself call the Brave Search API; that route remains dark and carries its own 30-day advance-notice treatment.

User-directed feature (launch disclosure basis). Authenticated web search is a user-directed feature, not a change to baseline AI processing. It fires only on the user's own search intent, as determined by a conservative intent classifier applied to the user's latest request; it is disabled entirely under Advanced Data Protection Mode (it fails closed); it is controllable at both the organization and personal level through the Allow web search settings; and when a search runs it is shown to the user through an in-product provenance indicator. On that basis, ISMS Copilot discloses the launch (Mistral-mediated) web-search route at launch, through the Trust Center, this DPA, and the in-app changelog once the feature is live, as a user-directed, customer-controllable feature, rather than treating it as a change to baseline AI processing that would require 30 days' advance sub-processor notice. This rationale narrows, but does not eliminate, the interpretive question of whether the onward Mistral-to-Brave query hop should be treated as a new sub-processor path; the 30-day advance-notice mechanism continues to apply to the separate future Brave-direct route.

OpenRouter account-level controls (enforced by Better ISMS as the OpenRouter account holder, applied to every request, applicable to all eight allowlisted underlying providers):

- Zero Data Retention is mandatory at the account level: per OpenRouter's published policy, ZDR-mandatory accounts can only route to endpoints with a Zero Data Retention policy. Zero Data Retention here means no persistent retention of request content beyond serving the request; a provider may still hold content transiently in memory for the duration of processing, as is inherent to serving an inference request.
- Free Training Disallowed and Paid Training Disallowed are both set.
- Free Publication Disallowed is set; the model-publication channel is closed.
- Closed eight-provider allowlist. The seven glm-4.7 hosts (Inception, DeepInfra, Cerebras, Google Vertex, Together AI, Fireworks AI, and Nebius) plus xAI via the "xAI (ZDR)" endpoint for the Grok models may serve our requests. (The seven glm-4.7 hosts were expanded from 4 to 7 on 2026-05-25; xAI was activated for paid Customer-Content on 2026-07-21; see the customer-facing change log (<https://trust.ismscopilot.com/en/changelog>) for the notices.)

- PRC-jurisdiction blocklist — Alibaba Cloud International, Baidu Qianfan, DeepSeek, Moonshot AI, Xiaomi, and Z.AI are all blocked. This is a Schrems II-style supplementary measure aligned with EDPB Recommendations 01/2020.

Configuration-integrity caveat. OpenRouter account-level controls are configured per the policies above. Better ISMS does not currently rely on an OpenRouter API or signed attestation for real-time integrity. Evidence of the configuration is two-fold: (a) this DPA is itself a contemporaneous record of the configured controls as of its effective date, and (b) Better ISMS will demonstrate the live OpenRouter account configuration via a guided dashboard walkthrough on customer request (typically a recorded screen-share session). Ad-hoc screenshots may be captured on specific customer request or when controls materially change.

Notified external search recipient (future Brave-direct route), pending activation. This is a separate route from Brave's role as Mistral's web-search sub-processor above. At launch, `SEARCH_RETRIEVAL_PROVIDER` defaults to Mistral for eligible non-ADP discovery (Mistral in turn queries Brave), and the direct-Brave adapter (in which ISMS Copilot would call the Brave Search API itself, rather than reaching Brave via Mistral) remains dark. No earlier than 2026-08-12, and only after at least 30 days' actual notice plus founder contractual acceptance, eligible authenticated requests with ADP off may send a bounded query derived only from the latest user request directly to Brave Software, Inc. through the Brave Search API (United States; AWS infrastructure). ISMS Copilot would not send files, memories, workspace context, account identifiers, or user/device identifiers. Brave's standard API privacy notice states that query logs may be retained for up to 90 days; the transfer mechanism is the Standard Contractual Clauses. Brave's published DPA expressly excludes "Search Query Data" from its processor scope. We therefore disclose Brave as a limited external recipient and do not represent this query processing as covered by Brave's Article 28 processor terms. The Customer's decision to leave web search enabled under its plan and organization policy, together with the user's latest request or one-shot search preference, would be treated as the documented instruction for this limited disclosure; the conservative classifier would only determine whether that instruction requires current external evidence. We apply the materially-adverse-change notice and objection process below to this future recipient even though Brave's contractual classification for Search Query Data is not that of an Article 28 processor.

Customer-Activated Integrations. The following sub-processors only become active for a Customer's data when that Customer's authorized administrator (e.g., an organization owner) explicitly enables an optional integration in-product. Because activation is contingent on the Customer's affirmative installation step, and because no Customer Personal Data flows to the sub-processor unless and until that step occurs, the prior-notice rule in "Changes to Sub-processors" below does not apply. By installing a Customer-Activated Integration, the Customer simultaneously authorizes ISMS Copilot to engage the corresponding sub-processor for the Customer's data only. Entries in this table are either sub-processors engaged by ISMS Copilot (Slack; Google LLC for the heyGRC Google Drive connection) or, where marked, third-party platforms selected and contracted by the Customer through which ISMS Copilot receives and returns Customer Personal Data on the Customer's instruction (GitHub); the latter are listed for transparency and are not engaged by ISMS Copilot under the general authorisation above.

Integration	Sub-processor	Activated by	Location	DPA / Transfer mechanism
heygrc Slack bot	Slack Technologies, Inc.	Paid-organization owner installs from the Connectors page; OAuth callback rejects free / null-plan installs with a <code>paid_plan_required</code> error. Uninstall hard-deletes integration records.	United States	SCCs
heyGRC GitHub App	GitHub, Inc. (Customer-directed platform, not an ISMS Copilot sub-processor: source of pull-request diffs and metadata; destination for the check run,	Customer's authorized administrator installs the heyGRC GitHub App on selected repositories (GitHub Marketplace or App page), granting read access to contents,	United States (GitHub's own footprint)	Governed by the Customer's agreement with GitHub; GitHub states DPF certification and SCCs (2021/914) for its own transfers

	review and comments)	metadata and issues and write access to checks and pull requests only; heyGRC never writes to repository code. Removing the App, or removing a repository from its scope, stops processing immediately; the Customer's GitHub account name is deleted at uninstall. Stored reviews are retained and deleted per §2.8.		
heyGRC Google Drive	Google LLC (Google Drive API, OAuth, Picker)	Available to every heyGRC organization; inactive until an organization owner or admin connects Google Drive from the heyGRC console. That person may also turn on Open with, so heyGRC appears in their Google Drive Open with menu. Files enter only when that person picks them in the Google Picker, or opens them with heyGRC and confirms. Disconnect revokes the grant at Google (unless the same Google account still connects another heyGRC organization) and deletes the stored token. Scopes <code>drive.file</code> (chosen files only) and, with Open with, <code>drive.install</code> (menu entry only, no file access).	United States (document location per the Customer's Google Workspace settings)	EU-US Data Privacy Framework + SCCs

heyGRC AI processing. Two inference paths exist, both under zero-retention, no-training terms. Default path (reviews and /heygrc replies): the redacted diff runs through the same OpenRouter account and account-level controls described in §2.1 (mandatory Zero Data Retention, training-disallowed, publication-disallowed, closed allowlist, PRC-jurisdiction blocklist), pinned in each request to a named subset of the allowlisted hosts (currently Together AI (US), Fireworks AI (US) and Inception (Sweden, EU); documented rollback set Cerebras and Google Vertex (US)), each request carrying `data_collection: deny` and `zdr: true`. The reviewing model is an open-weights model developed by Z.AI (currently GLM-5.3 Flash; documented rollback GLM-4.7); Z.AI does not host these requests (its first-party endpoint is on the §2.1(d) blocklist) and is not, on that basis, a recipient or sub-processor; consistent with §2.1(d), no Customer Personal Data transits infrastructure under PRC jurisdiction. EU inference path (Customer opt-in per organization; compliance reviews and linked-document and context compiles; not /heygrc replies): reviews and compiles are sent to Mistral AI on its EU regional endpoint (`api.eu.mistral.ai`) under the zero-retention and no-training configuration of ISMS Copilot's Mistral account and Mistral's Data Processing Addendum already listed in the Active Sub-processors table, with no fallback to the default path; the model served there is currently an open-weights model developed by Z.AI, hosted by Mistral; Z.AI does not receive Customer Personal Data on that path either. /heygrc replies use the default path regardless of the EU setting. No AI sub-processor beyond those already listed is engaged for heyGRC; the

Mistral AI row's purpose is extended accordingly.

When a new Customer-Activated Integration is offered for installation (or replaced), ISMS Copilot will document the integration in this section and announce its availability through normal product-update channels. The prior-notice rule in "Changes to Sub-processors" below applies to additions to the Active Sub-processors table above (which run automatically on Customer Personal Data without a per-Customer activation step), not to integrations that require explicit Customer-side activation.

Reserved Sub-processors (code paths exist but are not invoked from any user-facing flow). The following providers have integration code in the platform but are not currently used to serve any user request:

Sub-processor	Code path purpose	Status
OpenAI	Direct OpenAI API path	Reserved — not invoked from any current user-facing flow
X.AI (Grok)	Direct X.AI API path	Reserved — not invoked from any current user-facing flow
Google Gemini	Direct Gemini API path	Reserved — not invoked from any current user-facing flow

Activation of any Reserved sub-processor for live processing of Customer Personal Data requires customer notice under the change-of-sub-processor procedure in this §2.4 before any Customer Personal Data is processed.

Sub-processor Requirements. ISMS Copilot ensures all sub-processors:

- Provide sufficient guarantees of GDPR compliance
- Are bound by a written contract that imposes the same data protection obligations as this DPA (Article 28(4) GDPR)
- Implement appropriate technical and organizational measures
- Remain subject to ISMS Copilot's supervision and audit rights

Changes to Sub-processors.

(a) Allocation and removal. No advance notice is required to remove a sub-processor, or to allocate processing among, or fail over between, listed sub-processors (active or standby) under "General Authorization" above. ISMS Copilot publishes these changes in the Trust Center and in Appendix B. Where this DPA, the Terms of Service, the privacy policy or the Register of Processing Activities call such a move "control-neutral", it means allocation under this paragraph; a new sub-processor legal entity is never control-neutral and always follows (b). The Customer may raise a concern about an allocation at any time with privacy@ismscopilot.com; the parties will discuss it in good faith, including any available configuration such as Advanced Data Protection Mode.

(b) Prior notice of a new sub-processor. ISMS Copilot gives at least 14 days' prior notice before any new sub-processor legal entity processes Customer Personal Data, including a new provider reached through an intermediary or added to an allowlist (such as the closed OpenRouter allowlist), and including a provider added to the list as standby. For a materially adverse change the notice period is at least 30 days. Notice is given by email to the privacy contact the Customer has designated or, if none is designated, to the owner of the Customer's account or organization, and is also published in the Trust Center and the customer-facing change log. The notice names the sub-processor's identity, location and function, its transfer mechanism and safeguards, the effective date, and how to object.

(c) Materially adverse changes. A change is materially adverse where it materially changes the nature of processing, introduces a new category of Customer Personal Data processed, materially weakens retention, training, security, transfer, or residency controls, or introduces a materially different jurisdiction or transfer-risk posture. For materially adverse changes to this DPA or to the sub-processor framework, ISMS Copilot gives at least 30 days' advance notice by in-app announcement and email, unless a shorter period is required by law.

(d) Opt-in early access. An optional feature or model that needs a new sub-processor may be activated for a Customer before the notice period in (b) has run only by that Customer's workspace or organization administrator, after an in-product disclosure that names the provider. That activation is the Customer's

specific authorization of that sub-processor, for that feature and for that Customer's data only. Customer-Activated Integrations keep the rule set out under that heading above.

(e) Emergency replacement. Where a listed sub-processor fails or poses an immediate security risk and no listed sub-processor can reasonably take over, ISMS Copilot may engage a replacement that offers at least equivalent protection (no training on Customer Personal Data, zero data retention where available, and a valid transfer mechanism) before the notice period in (b) has run. Emergency use is permitted only to the extent and for as long as strictly necessary, and ends once a listed sub-processor can reasonably take over. ISMS Copilot notifies the Customer as set out in (b) within 72 hours of the replacement, and the Customer then has the objection rights in (f) as if notice had been given in advance, with the notice period running from that notification.

(f) Objection. The Customer may object in writing to privacy@ismscopilot.com within the notice period (14 days, or 30 days for a materially adverse change) on reasonable data protection grounds. If the Customer does not object within that period, the change is authorized. If the Customer objects, ISMS Copilot will, within 30 days of the objection, either not process that Customer's Personal Data with the sub-processor concerned or offer a reasonable alternative, such as Advanced Data Protection Mode (which routes all AI processing to Mistral AI in the EU regardless of plan and is available to every user). If ISMS Copilot does neither, the Customer may terminate the affected Services without penalty, and ISMS Copilot refunds fees prepaid for the unused period of those Services.

(g) Not sub-processors under this DPA. A third-party feature through which the data subject gives personal data directly to the third party (for example a testimonial or review widget), and a provider that Better ISMS uses as a controller for its own purposes, are not sub-processors under this DPA. Where applicable they are still listed in the Trust Center.

2.5 Data Subject Rights Assistance

ISMS Copilot will assist the Customer in fulfilling data subject rights requests, including:

ISMS Copilot will respond to Customer requests for data subject rights assistance within the timeframes specified below. Customer remains responsible for meeting GDPR's one-month response deadline to data subjects (Article 12(3)).

Right of Access (Article 15):

- Self-service access to all conversations and files through the platform
- Self-service complete data export in JSON format via Settings → Data Protection (available to all plans)

Right to Rectification (Article 16):

- Self-service updates to account settings
- Email-assisted email address changes (privacy@ismscopilot.com, within 30 days)

Right to Erasure (Article 17):

- Self-service account deletion via Settings → Data Protection (available to all plans)
- Email-mediated path for deletion within flagged threads (see §2.8 — moderation retention exception)
- Complete data deletion within 30 days, subject to §2.8

Right to Data Portability (Article 20):

- Machine-readable JSON export including all Customer Personal Data, available self-service in Settings → Data Protection

Right to Restrict Processing (Article 18) and Right to Object (Article 21):

- Email-mediated to privacy@ismscopilot.com, processed on a case-by-case basis
- Response within 30 days

Customer is responsible for verifying data subject identity before requesting data access or export. ISMS

Copilot provides the tools and processes, but Customer maintains primary responsibility for responding to data subject requests.

2.6 Data Breach Notification

In the event of a Personal Data Breach affecting Customer Personal Data, ISMS Copilot will:

Detection and Assessment:

- Continuously monitor for security incidents via Sentry and automated alerting
- Conduct security incident review within 24 hours of detection
- Assess risk and potential impact on Customer Personal Data

Notification to Customer:

- Notify Customer within 48 hours of confirming a Personal Data Breach affects Customer Personal Data
- For suspected breaches under investigation, provide preliminary notification within 24 hours with updates as information becomes available
- Provide description of the breach, including categories and approximate numbers of affected data subjects
- Describe likely consequences of the breach
- Outline measures taken or proposed to address the breach and mitigate its effects
- Provide contact point for further information

Cooperation:

- Cooperate with Customer's investigation and remediation efforts
- Provide reasonable assistance for Customer's notification to supervisory authorities and data subjects
- Document all breaches and remediation measures

Customer remains responsible for determining whether notification to supervisory authorities (within 72 hours per Article 33) and data subjects (Article 34) is required. ISMS Copilot provides information to support Customer's decision and obligations.

2.7 Data Protection Impact Assessment (DPIA) Support

ISMS Copilot will provide reasonable assistance when Customer conducts a Data Protection Impact Assessment or prior consultation with a supervisory authority, including:

- Providing the Register of Processing Activities for reference
- Describing technical and organizational measures implemented
- Clarifying data flows and sub-processor arrangements
- Answering specific questions about processing operations

2.8 Deletion and Return of Data

Upon termination of services or Customer request, ISMS Copilot will:

Standard Deletion (Default):

- Delete all Customer Personal Data within 30 days of termination
- Delete backup data within 91 days
- Provide written confirmation of deletion upon request

Data Export Before Deletion:

- Customer may export their data self-service via Settings → Data Protection at any time
- Export provided in JSON format

Legal Retention Exceptions:

- Anonymized billing records retained for 7 years (tax and accounting compliance)
- heyGRC audit records, kept by ISMS Copilot as controller for up to 12 months (see heyGRC audit records below); they are ISMS Copilot's own security records, not Customer Personal Data
- Anonymized analytics data may be retained
- Data required to be retained by applicable law will be isolated and protected until the legal retention period expires
- Slack integration data (heygrc bot only): if the Customer disconnects the integration (uninstall via Slack's app management UI, or removing the integration from the ISMS Copilot Connectors page), all Slack-specific records are hard-deleted within seconds via the `app_uninstalled` event handler — the OAuth bot token, workspace metadata in `slack_integrations`, and Slack-thread-to-ISMS-thread mappings in `slack_threads` are removed via cascading delete. Conversation content stored in the main messages and threads tables is preserved (it belongs to the Customer's organization and follows the Customer-configured retention setting), but the link from a Slack thread back to a specific Slack workspace user is severed at uninstall.
- Pull-request review data (heyGRC GitHub App only): diffs are not persisted as a record (they are held in the review worker only for the duration of the request; posted findings and stored findings may contain file paths and quoted changed lines). Generated reviews and the associated organization, repository and installation records are retained while the App is installed. Uninstalling the App stops processing, deletes the Customer's GitHub account name and starts a 30-day window after which reviews and remaining records are deleted automatically (subject to records ISMS Copilot must keep for legal, tax or accounting purposes); earlier deletion on Customer instruction is completed within 30 days. Reviews posted into the Customer's GitHub repositories remain under the Customer's control on GitHub. Backups are overwritten on the backup rotation described in §2.3.
- Google Drive connection (heyGRC, Customer-Activated Integration only): when an organization owner or admin disconnects Google Drive in the heyGRC console, ISMS Copilot asks Google to revoke heyGRC's access and then immediately hard-deletes the stored, encrypted Google refresh token, whatever Google answers. The Google account email and account identifier recorded for the connection are cleared at the same time. The same happens automatically, normally within about 30 minutes, when the organization uninstalls heyGRC or when the person who connected Drive is no longer an owner or admin of the organization. Revocation at Google is skipped only while the same Google account still connects another heyGRC organization, because Google holds one grant per account and app; the token held for the disconnected organization is still deleted. Documents already linked from Drive keep their last extracted text as ordinary linked-document records: they are deleted when the Customer removes the document, and deleted about 30 days after the organization uninstalls heyGRC, by a daily purge job. A disconnect audit record (who, when, reason; no token, no document content) is retained for up to 12 months (see heyGRC audit records below). If access is removed at Google instead, the stored token is deleted the next time heyGRC calls Google; the account email and identifier remain until the connection is disconnected in the console or the organization is purged.
- heyGRC compliance context (context layer only): The current revision of each object is kept while the organization is active. Objects removed in the heyGRC console (a removed source, or an approved removal of an object deleted at the source) keep their current revision for up to 12 months so they can be restored; their text is then erased (once no stored review cites it), or sooner on the Customer's documented request (within 30 days) or about 30 days after the organization uninstalls heyGRC. A superseded revision is kept while a stored review cites it; its text is erased within a day of that review's deletion. A superseded revision that no review cites has its text erased 12 months after it was superseded, or once its rules are no longer in force if that is later (identifiers, hashes and timestamps kept). On the Customer's documented request, ISMS Copilot erases the text of named revisions within 30 days; erasing a revision still in force takes that object out of force. This erasure does not reach review comments heyGRC has already posted on the Customer's GitHub pull requests; they stay in the Customer's repository, where the Customer can edit or delete them. An object the Customer's sync no

longer lists is held for an owner or admin to approve its removal, and stays in force until then. When an object is later sent again with a secret classification, ISMS Copilot erases its text from every stored revision and withdraws its obligations. All context data is deleted about 30 days after the organization uninstalls heyGRC, by the same daily purge job, and immediately if the organization is deleted; heyGRC audit records follow the next bullet. Records of held-change decisions (who decided, row state only) are kept on the Customer's behalf for up to 12 months, or until the organization is deleted if that comes first. The Customer can ask for deletion of specific objects at privacy@ismscopilot.com

- heyGRC audit records (ISMS Copilot as controller): Audit records of heyGRC API and console actions (who acted, when, which key and route, source IP address and user agent, and identifiers, hashes, counts and settings changed; never document content) are kept for up to 12 months from creation and then deleted automatically, including after the Customer's organization uninstalls heyGRC or the user's account is deleted. ISMS Copilot keeps them as controller for the security of the service (Art. 6(1)(f) GDPR), and shows them to the organization's owners and members while it is active.

Moderation Retention Exception (Article 17 limitation). When a chat message is flagged by our automated content moderation system, a `moderation_events` record is retained containing only metadata — message identifier, thread identifier, abuse categories, and timestamp. The full message content is not stored in the moderation record itself. In addition, the affected thread is locked from user-initiated deletion to prevent destruction of abuse evidence. This is a security measure required to prevent abuse evidence from being destroyed.

Customer Content within a flagged thread is still subject to deletion on a verified Article 17 erasure request submitted to privacy@ismscopilot.com; we evaluate each such request against the legitimate-interest balancing test (Article 17(3)(e) and recital 47) and respond within 30 days. The audit metadata (no content) is retained for up to 12 months from creation, after which it is automatically purged.

2.9 Audit Rights

Customer has the right to audit ISMS Copilot's compliance with this DPA, subject to reasonable limitations:

Documentation Review:

- Customer may review publicly available compliance documentation in our Trust Center (<https://trust.ismscopilot.com>), including the Register of Processing Activities (RoPA) (<https://trust.ismscopilot.com/ropa>)
- Request additional documentation through privacy@ismscopilot.com (e.g., sub-processor agreements, security policies, or a guided live walkthrough of the OpenRouter account configuration per §2.4)

On-Site Audits:

- Customer may conduct on-site audits with 60 days advance written notice
- Maximum of one audit per year unless necessitated by a data breach
- Audits must be conducted during business hours and not interfere with operations

Customer is responsible for audit costs unless the audit reveals non-compliance that: (a) constitutes a personal data breach, or (b) involves systematic failure to implement documented security measures, or (c) results in regulatory enforcement action. In such cases, ISMS Copilot will bear reasonable audit costs incurred after the non-compliance was identified.

Results remain confidential and may not be shared except as required by law.

Third-Party Certifications:

- Better ISMS is pursuing ISO/IEC 27001:2022 for the ISMS described at <https://trust.ismscopilot.com/en/iso-27001>. This is a management-system programme, not a product certification. Better ISMS is not certified today.
- Once a certification is issued, reports may be shared upon request subject to NDA
- Customers may rely on third-party certifications of our hosting providers, and on any Better ISMS certification once issued, in lieu of conducting their own audits

3. International Data Transfers

3.1 Data Transfer Mechanisms

ISMS Copilot processes Customer Personal Data in accordance with Chapter V of the GDPR. Transfer mechanisms apply per Active sub-processor. Two framing notes: (i) for the OpenRouter allowlist, the Standard Contractual Clauses operate at the underlying-provider layer through the OpenRouter aggregator (OpenRouter-intermediated); ISMS Copilot does not sign separate direct SCCs with each underlying host. (ii) Where a provider holds EU-US Data Privacy Framework certification, that is an Article 45 adequacy mechanism, distinct from the Article 46 Standard Contractual Clauses relied on elsewhere.

Primary Storage (Always EU).

- Database storage occurs in Frankfurt, Germany (AWS EU-Central-1)
- Conversation history, uploaded files, and account data are stored in the EU
- No adequacy decision required for primary storage

Backups. A daily copy of the database is made in the EU on Fly.io (Paris), encrypted there with public-key encryption before upload, and stored in Backblaze B2 (EU, Amsterdam), where each copy is kept for 90 days, then hidden and deleted one day later. The private decryption keys are kept separately and are never on the backup machine or at Backblaze. This applies to all Customers, including those with Advanced Data Protection Mode enabled. Backblaze, Inc. is a United States company; its DPA with Standard Contractual Clauses applies, and encryption with the key withheld is the supplementary measure for any access from the United States.

AI Processing — per routing path.

- ADP enabled (any plan): permitted AI processing uses Mistral AI's EU regional endpoint with zero retention. Authenticated web search is not available under ADP and fails closed, because Mistral web search is not EU-only and not zero-retention (it runs on Mistral's stateful endpoint and uses Brave (US) as sub-processor). Exact user-selected page/document fetch remains available through EU Fly and subsequent AI processing remains on Mistral.
- Authenticated web search at launch, ADP off: the minimized latest-request query is sent to Mistral AI for provider-attested source-URL discovery. Mistral's web-search discovery runs on Mistral's stateful Conversations/Agents endpoint, which is outside Mistral's zero-retention posture and uses Brave (United States) as its web-search sub-processor, so the query hop reaches Brave (US) under standard (non-zero) retention and is not an EU-only zero-retention route. Mistral's generated search prose is discarded. The EU Fly.io service exact-fetches those URLs, applies SSRF and content/size/time controls, and passes bounded labeled page extracts to the answer model already applicable to the conversation or Beyond run. No additional semantic grounding model receives the complete draft. A discovered public host may be outside the EU and receives ordinary outbound request metadata from the EU fetcher.
- Future Brave-direct route (not before 2026-08-12, and not part of launch): this is a distinct route from Brave's role as Mistral's sub-processor above; here ISMS Copilot would call the Brave Search API itself. Only after the notice and founder-acceptance conditions in §2.4 are complete could a minimized latest-request query be transferred directly to Brave Software, Inc. in the United States under SCCs. Brave standard query logs may be retained up to 90 days, and its published DPA excludes Search Query Data from processor scope. Personal and organization-level Never search the web controls would prevent the transfer.
- Non-ADP AI processing is a single routing envelope (any plan, ADP off). With ADP off, AI requests are served within one non-ADP envelope whose permitted destinations are Anthropic Claude and the OpenRouter allowlist (the eight vetted underlying providers listed below, including xAI; the Essential-plan subset is Google Vertex + Cerebras), with Mistral (EU) as the circuit-breaker failover. ISMS Copilot may serve any non-ADP request via Anthropic or via any allowlisted provider, and may move any plan or cohort between these destinations over time; we do not commit to a fixed per-plan or per-request provider assignment. The per-plan descriptions that follow are the current routing snapshot within this single envelope, not a contractual per-plan limitation. Anthropic-served traffic follows Anthropic's standard commercial API terms (ordinary ~30-day deletion; safety-flagged content up to 2 years;

safety-classification scores up to 7 years; not training); OpenRouter-served traffic is zero-retention (no persistent retention beyond serving the request; transient in-memory caching only); ADP (Mistral, EU) is zero-retention. Change classification (§2.4): allocation among listed sub-processors, no advance sub-processor notice. No new sub-processor is introduced (Anthropic and every allowlisted provider are already disclosed non-ADP sub-processors); only the set of cohorts that may be routed to each already-disclosed destination is broadened. The account-level controls, allowlist, PRC-blocklist, transfer mechanisms, and the Advanced Data Protection (Mistral, EU, zero-retention) guarantee are all unchanged, and ADP remains the contractual EU-only, zero-retention option available to any customer at any time. (Assessment note: broadening the free-tier cohort into a destination whose retention exceeds its prior zero-retention default has a stricter reading under which it could be viewed as materially adverse; ISMS Copilot has assessed this change as not materially adverse on the basis above.) For paid plans (Plus and above) in the Fast and Think modes and the Beyond assistant, and for free / null-plan users, the current default destination is GLM-5.3 via OpenRouter pinned to Together AI, Fireworks AI and Inceptron, and Anthropic is the pre-first-token failover for paid Fast and Think only; in any chat mode ISMS Copilot may serve a non-ADP paid request via any of the eight allowlisted providers or Anthropic and may move any cohort or mode between these destinations over time, which is allocation among listed sub-processors under §2.4 (no advance notice).

- Paid (Plus and above) + ADP off: For the Fast and Think modes and the Beyond assistant the current default is GLM-5.3 via OpenRouter (z-ai/glm-5.3-flash for Fast, z-ai/glm-5.3 for Think and Beyond), pinned to Together AI (US), Fireworks AI (US) and Inceptron (Sweden, EU) (zero data retention; SCCs at the underlying-provider layer), with Anthropic as the pre-first-token failover for Fast and Think; the Beyond assistant retries on the same GLM-5.3 model. Authenticated web search does not change this route: Mistral discovers the sources and the answer stays on the default model, so Anthropic is not used for web search. In any chat mode, ISMS Copilot may serve a non-ADP paid request via any of the eight allowlisted underlying providers or Anthropic, and may move any cohort or mode between these destinations over time (such moves among listed sub-processors are allocation under §2.4: no advance notice). Anthropic Claude (United States) or the OpenRouter allowlist (mostly United States; Inceptron is Sweden (EU) and Nebius is Netherlands (EU), but OpenRouter does not pin inference region, so the SCC safeguard still applies to all eight), as permitted routing destinations for any chat mode. ISMS Copilot may serve a paid request via Anthropic or via any allowlisted provider; Anthropic is also the circuit-breaker failover context. Transfer mechanism for Anthropic-served traffic: SCCs (Module Three, Processor-to-Processor); supplementary measures: encryption in transit; no training under Anthropic's commercial API terms; retention under Anthropic's standard commercial API terms (inputs and outputs ordinarily deleted within approximately 30 days; content flagged by its safety systems retained up to 2 years; safety-classification scores up to 7 years; never used for model training). Transfer mechanism for OpenRouter-served traffic: the underlying-provider DPA / SCC / DPF stack, with OpenRouter account-level controls (mandatory ZDR, training-disallowed, publication-disallowed, closed eight-provider allowlist, PRC-blocklist) as the enforcement layer; that traffic is zero-retention (no persistent retention beyond serving the request; transient in-memory caching only). Fallback within a single request: where a paid Fast or Think request is first attempted on an allowlisted OpenRouter provider and that provider does not begin responding within a short time-to-first-response timeout, ISMS Copilot aborts the attempt and re-serves the request via the failover path. That failover path is Anthropic (Opus for Think, Sonnet for Fast) or, where Anthropic's circuit breaker is open (Anthropic degraded or unavailable), Mistral in the EU (see the failover provisions). In that fallback case a single request transits two disclosed destinations in sequence: the OpenRouter attempt is zero-retention (no persistent retention of Customer Content there beyond serving the request), and the second leg follows its standard posture (Anthropic under SCCs with retention under its standard commercial API terms as described above, or Mistral in the EU with zero retention). No additional sub-processor is involved and no control is weakened; the deepest fallback (Mistral, EU, zero-retention) is the most protective destination. Customers needing zero retention on every path, or contractual EU-only residency, can enable Advanced Data Protection at any time to route all AI processing through Mistral (EU).
- Essential + ADP off: OpenRouter aggregator (United States), restricted to a closed two-provider subset of the allowlist: Google Vertex (multi-region; SCCs and EU-US Data Privacy Framework certification) and Cerebras (United States; SCCs). The same OpenRouter account-level controls (mandatory ZDR, training-disallowed, publication-disallowed, PRC-blocklist) apply as the enforcement layer. Customers needing contractual EU-only data residency can enable Advanced Data Protection at any time to route through Mistral (EU).

- Free / null-plan + ADP off: OpenRouter aggregator (United States), currently serving z-ai/glm-5.3-flash pinned to Together AI, Fireworks AI and Inceptron (Plus users over their usage cap are served on the same route), within the seven allowlisted underlying providers:
 - DeepInfra, Cerebras (United States): SCCs. Inceptron (Inceptron AB, Sweden / EU): SCCs as well, because OpenRouter does not pin inference region.
 - Google Vertex (multi-region; EU residency available at the underlying provider): SCCs and EU-US Data Privacy Framework certification.
 - Together AI (default routing to North America data centers per Together docs; region pinning not exposed by OpenRouter aggregator): SCCs.
 - Fireworks AI (multi-region fleet covering US, EU Frankfurt + Iceland, and APAC Tokyo only; no PRC or Hong Kong infrastructure identified in Fireworks' published deployment docs; region pinning not exposed by OpenRouter aggregator): SCCs.
 - Nebius (Amsterdam HQ; primary inference in Finland (EU) with US secondary per Nebius docs; per-request region pinning not currently exposed by OpenRouter aggregator, so the EU-primary posture is the published default but not contractually guaranteed via OpenRouter): SCCs for any US-routed traffic.
- The transfer mechanism for the free-tier path is the underlying-provider DPA/SCC/DPF stack. OpenRouter account-level controls (mandatory ZDR, training-disallowed, publication-disallowed, closed eight-provider allowlist, PRC-blocklist) act as the enforcement layer ensuring Customer Personal Data only ever lands at endpoints with these mechanisms in place. OpenRouter's aggregator API does not currently expose per-provider region pinning to its account holders; for the three providers added on 2026-05-25 (Together AI, Fireworks AI, Nebius), regional routing follows each provider's default deployment posture. Based on each provider's published deployment documentation reviewed on 2026-05-25, none of those defaults places Customer Personal Data in PRC or Hong Kong infrastructure. Non-PRC default posture is not equivalent to contractual EU-only data residency; customers requiring the latter should enable Advanced Data Protection Mode, which routes every request to Mistral AI's EU regional endpoint and bypasses the OpenRouter aggregator entirely.

Email Communications (US-Based).

- Email addresses transferred to SendGrid and Kit (United States)
- Protected by Standard Contractual Clauses approved by the European Commission
- Customers can minimize transfers by unsubscribing from non-essential emails

Slack Integration (US-Based, paid Customers who install heygrc only).

- Workspace metadata, OAuth bot token, message events addressed to the bot, and AI response posts transferred to/from Slack Technologies, Inc. (United States)
- Protected by Standard Contractual Clauses
- Customer controls activation: the integration only exists if a paid-organization owner installs the bot, and can be removed at any time by uninstalling from Slack or via the Connectors page in the ISMS Copilot app (uninstall hard-deletes all Slack-side records held by ISMS Copilot per §2.8)

Google Drive (heyGRC, Customers who connect only).

- OAuth credential exchange with Google LLC (United States); the documents themselves are read from the Customer's own Google Drive into the EU-hosted heyGRC worker
- Protected by the EU-US Data Privacy Framework and Standard Contractual Clauses in Google's terms
- Customer controls activation: the connection exists only if an organization owner or admin connects Google Drive, and can be removed at any time from the heyGRC console or at myaccount.google.com (disconnect deletes the stored token per §2.8)

3.2 Standard Contractual Clauses (SCCs)

For transfers to the United States, ISMS Copilot relies on Standard Contractual Clauses (Commission Implementing Decision (EU) 2021/914):

- Customer to ISMS Copilot: Module Two (Controller to Processor) applies where Customer acts as data controller
- ISMS Copilot to US sub-processors: Module Three (Processor to Processor) where the sub-processor's terms provide it. OpenRouter's executed terms use Module Two, with ISMS Copilot as data exporter (see §1.7 and §3.1)
- Governing law for SCCs: French law (Clause 17, Option 1)
- Competent supervisory authority: CNIL, France (Clause 13)

Copies of executed SCCs with sub-processors are available on request via privacy@ismscopilot.com.

3.3 Supplementary Measures

ISMS Copilot implements supplementary measures to protect data transferred outside the EU. These measures are layered and apply per routing path:

For all transfers outside the EU:

- Encryption in transit (TLS 1.3) for all data in transit
- Customer ability to control transfer destination via Advanced Data Protection Mode (Mistral, EU)
- Continuous monitoring of legal developments regarding international transfers
- For authenticated web search, latest-turn-only query construction; strict length and word bounds; no files, memories, workspace context, account identifiers, or user/device identifiers; credential and signed-URL rejection; personal/organization hard-off controls; Mistral-prose discard; and EU exact-fetch of only provider-attested URLs

For paid-tier transfers served by Anthropic (currently the pre-first-token failover for paid Fast and Think only):

- No training on Customer Content under Anthropic's commercial API terms
- retention under Anthropic's standard commercial API terms (inputs and outputs ordinarily deleted within approximately 30 days; content flagged by its safety systems retained up to 2 years; safety-classification scores up to 7 years; never used for model training); customers needing zero retention can enable Advanced Data Protection at any time to route through Mistral (EU, no retention)

For OpenRouter transfers (free-tier, Essential, and paid plans Plus and above):

- OpenRouter account-level controls applied to every request: mandatory Zero Data Retention (no persistent retention beyond serving the request; transient in-memory caching only); Free Training Disallowed + Paid Training Disallowed; Free Publication Disallowed
- Closed eight-provider allowlist (the seven glm-4.7 hosts, Inceptron, DeepInfra, Cerebras, Google Vertex, Together AI, Fireworks AI, Nebius, plus xAI via the "xAI (ZDR)" endpoint for the Grok models; the seven glm-4.7 hosts were expanded from 4 to 7 on 2026-05-25, and xAI was activated for paid Customer-Content on 2026-07-21). For Essential plan routing the closed two-provider subset is further restricted to Google Vertex and Cerebras. For paid plans (Plus and above) in the Fast and Think modes and the Beyond assistant, and for the free / null-plan path, the current default is GLM-5.3 pinned to Together AI, Fireworks AI and Inceptron, with Anthropic as the pre-first-token failover for paid Fast and Think; xAI serves the logged-out risk-analysis demo and remains the operator rollback route.
- PRC-jurisdiction blocklist (Alibaba Cloud Int., Baidu Qianfan, DeepSeek, Moonshot AI, Xiaomi, Z.AI all blocked) — a Schrems II-style jurisdiction-based supplementary measure aligned with EDPB Recommendations 01/2020
- Each underlying provider independently confirms zero retention or no-training-on-customer-data in their published policies
- Contractual chain to the underlying providers: ISMS Copilot has an executed Data Processing

Agreement with OpenRouter (accepted via the OpenRouter Terms as a paid commercial customer), which incorporates the EU Standard Contractual Clauses for the ISMS Copilot to OpenRouter transfer, and under which OpenRouter is contractually required to bind each underlying model provider, by written agreement, to data-protection obligations no less protective than its own (GDPR Article 28(4) sub-processor flow-down) and remains liable for those providers' acts and omissions. This flow-down is what anchors the per-provider transfer mechanisms above; customers do not contract with the underlying providers individually.

- Advanced Data Protection remains the account-level EU-residency guarantee for AI processing: enabling it at the organization level routes permitted AI processing to Mistral's EU regional endpoint, suppresses OpenRouter for every user, and disables authenticated web search entirely (it fails closed, because Mistral web search is not EU-only and not zero-retention). Exact user-selected page/document fetch remains available under ADP through EU Fly.

3.4 Transfer Impact Assessment

ISMS Copilot has conducted a Transfer Impact Assessment (TIA) for US-based sub-processors and determined that:

- Standard Contractual Clauses provide appropriate safeguards under GDPR Chapter V
- Supplementary technical and contractual measures (encryption, no-training, zero retention, OpenRouter account-level controls including the PRC blacklist) enhance protection
- Customers have the option to avoid US AI processing transfers entirely by enabling Advanced Data Protection Mode (EU-only permitted AI processing with zero retention; authenticated web search is unavailable and fails closed, because Mistral web search is not EU-only and not zero-retention)
- Email transfers to US providers (SendGrid, Kit) remain regardless of Advanced Data Protection Mode but are protected by SCCs and encryption

The complete Transfer Impact Assessment, including risk assessment methodology and US surveillance law analysis, is available on request via privacy@ismscopilot.com.

Organizations with strict EU data residency requirements should enable Advanced Data Protection Mode to eliminate AI processing transfers and simplify Transfer Impact Assessment obligations.

4. Customer Obligations as Data Controller

4.1 Lawfulness of Processing Instructions

Customer warrants that:

- All processing instructions comply with GDPR and applicable data protection laws
- Customer has a lawful basis for processing all personal data uploaded to the platform
- Customer has informed data subjects about the processing and their rights
- Customer maintains appropriate records of processing activities (Article 30 GDPR)

4.2 Special Category and Criminal-Offence Data (excluded from the Default AI Routing Path)

ISMS Copilot does not authorize the processing of special category data or personal data relating to criminal convictions and offences through the Default AI Routing Path. OpenRouter's data processing terms state that OpenRouter is not required to process Sensitive Data unless expressly agreed and that no such processing is intended. ISMS Copilot has not entered into such an agreement with OpenRouter for OpenRouter-routed traffic. Accordingly:

- (a) Customer, as controller, instructs ISMS Copilot not to submit special category or criminal-offence data to the Default AI Routing Path, and shall not do so;
- (b) any such data that Customer nonetheless includes in a prompt, message, or upload is submitted at

Customer's sole responsibility, on the basis of an Article 9(2) condition (or, for Article 10 data, an Article 6(1) basis and either official-authority control or authorization under Union or Member State law providing appropriate safeguards, as required by Article 10) that Customer has established and can evidence, and subject to Customer's obligations under §4.1 and §5;

(c) for any processing intended to involve special category or criminal-offence data, Customer shall enable Advanced Data Protection Mode, which routes all AI processing to Mistral AI (EU) with zero retention and does not use the OpenRouter routing path; and

(d) ISMS Copilot gives no representation that special-category or criminal-offence content is detected, filtered, or blocked before it reaches a sub-processor, and does not warrant the suitability of the Default AI Routing Path for such data.

Where such data is nonetheless present on the Default AI Routing Path, the sub-processors' baseline Article 28 obligations (confidentiality, Article 32 security, the Article 28(4) sub-processor flow-down, deletion, and the applicable Standard Contractual Clauses safeguards) continue to apply to it as personal data. The exclusion in this section concerns the absence of any enhanced warranty for the special-category subset, not a withdrawal of those baseline protections.

Customer remains responsible for (i) classifying its inputs, (ii) ensuring it does not submit excluded data to the Default AI Routing Path, (iii) any Article 9(2) condition or, for Article 10 data, the Article 6(1) basis and official-authority control or legal authorization with appropriate safeguards required by Article 10, applicable to Customer's own processing, and (iv) conducting a Data Protection Impact Assessment (Article 35) where required.

4.3 Data Subject Rights Management

Customer is responsible for:

- Receiving and responding to data subject rights requests
- Verifying data subject identity before requesting data from ISMS Copilot
- Determining whether to notify supervisory authorities and data subjects in case of breaches
- Ensuring data subjects are informed about ISMS Copilot's role as processor

4.4 Data Retention Configuration

Customer must:

- Configure appropriate data retention periods matching their data protection policies
- Review retention settings periodically to ensure compliance
- Request deletion when data is no longer necessary for the original purpose

4.5 Workspace Isolation

Customer should:

- Create separate workspaces for different clients or data categories
- Avoid mixing personal data from different data subjects in single workspaces
- Delete workspaces when projects are completed and data is no longer needed

5. Liability and Indemnification

5.1 Allocation of Liability

Under Article 82 GDPR:

- Customer and ISMS Copilot are each liable for damages caused by their own GDPR violations
- ISMS Copilot is exempt from liability if it proves it was not responsible for the event giving rise to the damage

- ISMS Copilot is not liable for damages resulting from Customer's unlawful processing instructions

5.2 Indemnification

Customer will indemnify ISMS Copilot against any claims or damages arising from:

- Customer's violation of GDPR or other data protection laws
- Customer's unlawful processing instructions
- Customer's failure to obtain necessary consents or legal basis for processing
- Customer's submission of special category or criminal-offence data to the Default AI Routing Path contrary to §4.2 (Customer must use Advanced Data Protection Mode for such workloads)

Nothing in this §5.2 limits data-subject rights or requires Customer to indemnify ISMS Copilot to the extent a claim or damage is attributable to ISMS Copilot's breach of this DPA or its obligations under applicable data protection law, including Article 28(4) or Article 82 GDPR.

5.3 Limitation of Liability

(a) Subject to (b) and (c), the Processor's total aggregate liability in damages for all claims arising out of or in connection with this DPA, whatever the legal basis, together with its liability under Section 8(iii) of the Terms of Service, is one combined cap. That cap does not exceed the greater of (i) the fees the Customer paid for the Services in the 12 months before the first event giving rise to liability and (ii) EUR 1,000. Related events count as one event. A later version of the Terms of Service does not reduce this cap for a countersigned version.

(b) The cap does not apply to: liability for faute lourde ou faute dolosive; liability under Article 82 GDPR, including recourse under Article 82(5); death or personal injury; or any liability that mandatory law does not allow the parties to limit or exclude.

(c) Nothing in this DPA limits the rights of data subjects against either party or the powers of a supervisory authority.

6. Term and Termination

6.1 Term

This DPA takes effect on the date Customer first uses ISMS Copilot services and continues for as long as ISMS Copilot processes Customer Personal Data.

6.2 Termination

This DPA terminates automatically upon:

- Termination of the Terms of Service
- Completion of all processing activities and deletion of Customer Personal Data

6.3 Effect of Termination

Upon termination:

- ISMS Copilot will delete or return all Customer Personal Data as described in §2.8 (subject to the moderation retention exception)
- Obligations regarding confidentiality, data security, and legal retention survive termination
- Customer's right to audit survives for 12 months after termination

7. Amendments and Updates

7.1 DPA Updates

ISMS Copilot may update this DPA to reflect:

- Changes in data protection laws or regulatory guidance
- Changes to processing operations or sub-processors
- Improvements to security measures or data protection practices

7.2 Notification of Changes

§7.2 and §7.3 do not apply to a countersigned version except as §7.4 provides.

- Materially adverse changes to this DPA (changes that materially weaken Customer's data-protection rights, reduce retention or training protections, introduce a materially different jurisdiction or transfer-risk posture, or otherwise impose materially new obligations on Customer) will be notified at least 30 days in advance by in-app announcement and email, including through ISMS Copilot's regular customer product-update or changelog email.
- Non-materially-adverse changes to this DPA (clarifications, changes that strengthen Customer's protections, alignment with regulatory guidance, etc.; sub-processor changes follow §2.4) will be notified through the Trust Center, the in-app changelog, and the customer-facing change log at <https://trust.ismscopilot.com/changelog>.
- The updated DPA will be posted at <https://trust.ismscopilot.com/dpa> with a new "Effective Date" and "Last Updated" date.
- Continued use of services after the effective date constitutes acceptance of the updated DPA.
- Customer may object during the stated notice period by contacting privacy@ismscopilot.com (see §7.3).

7.3 Objection Rights

- Customer may object to material changes within 30 days of notification
- The simplest in-product alternative is to enable Advanced Data Protection Mode, which keeps all AI processing in the EU at Mistral regardless of plan
- Objections to sub-processor changes, and their outcome, follow §2.4
- For customers under contracts with formal objection rights for other changes, formal objection may be sent to privacy@ismscopilot.com; if the objection cannot be resolved, Customer may terminate the service without penalty

7.4 Countersigned Copies

(a) Offer and acceptance. The Processor publishes at <https://trust.ismscopilot.com/en/dpa> a copy of this DPA signed by the Processor. The signed copy is an offer. A Customer accepts it by completing and countersigning the copy and sending it to privacy@ismscopilot.com. The countersigned copy takes effect when the Processor receives it. Until then, and if the Customer never sends it, the DPA in force under the Terms of Service applies.

(b) What is agreed. The agreed terms are (i) the DPA text identified by the version and source fingerprint (SHA-256) printed on every page and (ii) the cover and signature page terms identified by the template revision printed on every page. Both are published for that version at <https://trust.ismscopilot.com/en/dpa/versions>. The Customer completes only the designated fields: the product boxes, the Customer details and the Customer signature block. Completing them does not change the terms. Any other addition, deletion or alteration, handwritten, typed or electronic, has no effect unless both parties sign a separate written amendment.

(c) Precedence. The countersigned version governs the processing of Customer Personal Data between the parties for every product covered by this DPA that the Customer uses, whether or not it ticked that product's box. It prevails over any later version of this DPA and, to the extent of any conflict, over any

later version of the Terms of Service, including an entire-agreement clause in a later version of the Terms.

(d) Changes that reach a countersigned version. Only the following changes apply to a countersigned version without the Customer's signature:

- (i) changes to sub-processors made under the procedure in §2.4, including prior notice and the right to object; and
- (ii) changes required by mandatory data protection law, or by a binding decision of a competent authority that applies to the Processor. Such a change applies only to the extent necessary and is notified at least 30 days in advance, or sooner where the law requires. The Customer may terminate the affected Services without penalty under §2.4(f).

(e) Later versions. Any other later version applies to the Customer only if the Customer accepts it in writing, including by countersigning a copy of it. The later countersigned version then replaces the earlier one.

(f) Lapse. The offer in a signed copy lapses on the earlier of (i) 60 days after the Processor publishes a later version and (ii) the date the Processor marks that version as withdrawn at <https://trust.ismscopilot.com/en/dpa/versions>. A copy received after the offer has lapsed does not create a countersigned version.

8. Governing Law and Jurisdiction

8.1 Governing Law

This DPA, and any non-contractual obligation arising out of or in connection with it, is governed by French law, without prejudice to the GDPR and to the mandatory provisions of the data protection law that applies to the Customer as controller.

8.2 Jurisdiction

The competent courts of Paris, France, have exclusive jurisdiction over any dispute arising out of or in connection with this DPA, except where mandatory law confers jurisdiction on another court (including for claims under Article L. 442-1 of the French Commercial Code). This clause does not affect the right of data subjects to bring proceedings under Articles 79 and 82(6) GDPR, or the competence of a supervisory authority.

9. Contact Information

9.1 Data Protection Contacts

For DPA-related questions or requests:

- Email privacy@ismscopilot.com from your registered account email address
- Include "DPA Request" or "Data Processing Agreement" in the subject line

9.2 Data Protection Officer

ISMS Copilot has not designated a Data Protection Officer as we do not meet the mandatory designation criteria under GDPR Article 37. For data protection inquiries related to this DPA, contact privacy@ismscopilot.com.

9.3 Supervisory Authority

Commission Nationale de l'Informatique et des Libertés (CNIL)

- Website: <https://www.cnil.fr/en>
- Address: 3 Place de Fontenoy, TSA 80715, 75334 Paris Cedex 07, France
- Phone: +33 1 53 73 22 22

10. United States State Privacy Laws

10.1 Scope

This section applies where a Customer is subject to a United States state privacy law, or where personal data ISMS Copilot processes on the Customer's behalf is subject to such a law. It covers the California Consumer Privacy Act as amended by the California Privacy Rights Act (together, "CCPA"), the Virginia Consumer Data Protection Act ("VCDPA"), the Colorado Privacy Act ("Colorado CPA"), the Connecticut Data Privacy Act ("CTDPA") and the Utah Consumer Privacy Act ("UCPA"), together with any other US state privacy law that imposes equivalent processor-contract requirements ("Applicable US State Privacy Laws"). In this section, "personal information" has the meaning the applicable law gives it.

10.2 Service-Provider and Processor Role

For personal data that a Customer, acting as a business or controller under an Applicable US State Privacy Law, entrusts to ISMS Copilot for processing on its behalf for the Services (including Aevral, §1.7, and heyGRC), ISMS Copilot acts as a service provider under the CCPA and as a processor under the other Applicable US State Privacy Laws.

10.3 Commitments

For personal data processed on the Customer's behalf, the parties agree that:

- the Customer discloses the personal data to ISMS Copilot only for the limited and specified purposes of providing the Services described in this DPA (§1.2, and §1.7 for Aevral; see also Appendix A), and ISMS Copilot processes it only for those purposes and only on the Customer's documented instructions (§2.1);
- ISMS Copilot refrains from selling the personal data, from sharing it for cross-context behavioral advertising (as those terms are defined by the CCPA), and from processing it for targeted advertising, profiling or sale as those terms are defined by the other Applicable US State Privacy Laws;
- ISMS Copilot refrains from retaining, using or disclosing the personal data outside the direct business relationship between ISMS Copilot and the Customer, and from using it for any purpose other than the business purposes specified in this DPA, including for the development, training or improvement of any AI model, save for processing required by Union or Member State law to which ISMS Copilot is subject and the independent safety and legal disclosures reserved in §2.1, for which ISMS Copilot acts as an independent controller;
- ISMS Copilot combines the personal data with personal data received from or on behalf of other persons, or collected from ISMS Copilot's own interactions with any consumer, only as expressly permitted by the applicable state privacy law, and only for the business purposes specified in this DPA;
- ISMS Copilot will comply with all applicable obligations under the CCPA and its regulations, and under the other Applicable US State Privacy Laws, in its processing of the personal data;
- ISMS Copilot will notify the Customer if it determines that it can no longer meet its obligations under any Applicable US State Privacy Law, and, unless legally prohibited, grants the Customer the right, upon notice, to take reasonable and appropriate steps to stop and remediate unauthorized use of the personal data;
- the Customer has the audit and inspection rights in §2.9, so the Customer can monitor compliance with this section;
- ISMS Copilot will assist the Customer in responding to verifiable consumer requests, by making available the personal data in its possession as Processor (§2.5), and with data protection impact assessments and similar risk assessments (§2.7);
- ISMS Copilot will delete or return the personal data at the end of the provision of the Services, at the Customer's choice (§2.8; for Aevral, §1.7, including its 12-month post-uninstall retention so a returning Customer finds its history);
- ISMS Copilot will grant the personal data the same level of privacy protection as required of the Customer by the applicable state privacy law, and will ensure that sub-processors (§2.4) and any person

it authorizes to process the personal data are bound by contracts that comply with the requirements of this section as the CCPA and its regulations require of service providers;

- ISMS Copilot will notify the Customer of any breach of the security of the personal data without undue delay, per §2.6;
- ISMS Copilot certifies that it understands the restrictions in this section and will comply with them.

10.4 Opt-Out Preference Signals

ISMS Copilot does not sell or share personal data, so a sale or sharing opt-out preference signal (including the Global Privacy Control) does not change the processing described in this DPA. Should ISMS Copilot ever wish to sell or share personal information, that change would require an amendment to this DPA and prior written notice to the Customer.

10.5 Precedence

Where this section imposes stricter obligations than the rest of this DPA, this section prevails for the applicable state privacy law. This section does not reduce the GDPR obligations in this DPA.

Appendix A: Processing Details Summary

Subject Matter. Provision of AI-powered compliance assistance platform including conversation processing, document analysis, knowledge management, automated content moderation, and policy-controlled authenticated web research.

Duration. For the term of the Customer's active subscription plus retention period configured by Customer (1 day to 7 years), followed by 30-day deletion window. Moderation metadata retained up to 12 months per §2.8.

Nature and Purpose.

- Nature: Automated AI processing, database storage, file conversion and analysis, content moderation, search-intent classification, and transient web-evidence retrieval
- Purpose: Enable compliance professionals to receive AI guidance, analyze documents, generate policies, manage compliance knowledge, and ground current/company-specific claims in retrieved sources

Categories of Data Subjects.

- Customer employees and authorized platform users
- Customer clients (when referenced in documents or queries)
- Individuals mentioned in compliance documentation
- Security incident subjects
- Slack workspace users interacting with the heygrc bot (paid Customers who install only) — see §1.5
- Individuals appearing in pull-request content reviewed by the heyGRC GitHub App, and GitHub users who ask /heygrc questions (Customers who install only); see §1.5

Categories of Personal Data.

- Contact information (email addresses)
- Authentication credentials (hashed passwords)
- Conversation content and AI interactions
- Uploaded compliance documents
- Usage metadata and timestamps
- Special category data (Article 9) and criminal-offence data (Article 10): excluded from the Default AI Routing Path (§4.2); may be incidentally present in unstructured customer inputs at the Customer's

responsibility

- Moderation metadata (message ID, thread ID, abuse categories, timestamp — no message content) for flagged messages only
- Slack integration data (paid Customers who install heygrc only) — see §1.6
- Pull-request review data (Customers who install the heyGRC GitHub App only): redacted diffs and pull-request metadata, company profile and commitments, /heygrc questions; stored reviews quote changed lines; see §1.6 and §2.8
- heyGRC compliance context (Customers who enable the context layer only), see §1.6
- For web search: limited recent conversation text used by the intent classifier; an ADP-off minimized latest-request query; transient public search results or exact-URL page text; bounded excerpts and source titles/URLs supplied to the already applicable answer/Beyond model; generated answer and source links retained under the conversation setting

Sensitive data (SCC Annex I.B safeguards). Special category data (Article 9) and personal data relating to criminal convictions and offences (Article 10) may be incidentally present in unstructured customer inputs at the Customer's responsibility. Such data is excluded from instructed processing on the Default AI Routing Path under §4.2, and intended workloads require Advanced Data Protection Mode. Applied restrictions and safeguards include strict purpose limitation to providing the services; no training; no publication of model outputs derived from Customer Personal Data; TLS in transit and encryption at rest; mandatory Zero Data Retention, the closed provider allowlist, and the PRC-jurisdiction blocklist for OpenRouter-routed traffic; the Anthropic and authenticated web-search retention and transfer controls disclosed in §§2.1(c) and 3.1; and Advanced Data Protection Mode as an EU-only, zero-retention alternative for AI processing.

Appendix B: Sub-processor Change Log

This appendix tracks all sub-processor additions, removals, and changes since the DPA effective date. From the 2026-10-26 revision, a new sub-processor is notified by email at least 14 days before it processes Customer Personal Data (30 days for a materially adverse change), and allocation among listed sub-processors, removals and customer-activated features follow §2.4.

Effective	Change	Notice issued
2026-10-26 (editorial; no change to processing)	Encryption wording. §3.3 lists the in-transit measure as encryption in transit (TLS 1.3). Data is encrypted in transit and at rest. No change to processing, sub-processors or transfers.	Trust Center and change log
2026-10-26 (non-materially-adverse; strengthens Customer rights)	Sub-processor change procedure and countersigned copies. §2.4: general authorization covers listed sub-processors, including providers marked standby in the Trust Center; at least 14 days' prior email notice before any new sub-processor legal entity processes Customer Personal Data, including a new provider reached through an intermediary or allowlist (30 days if materially adverse); opt-in early access by the Customer's administrator after in-product disclosure; a narrow emergency replacement with notice within 72 hours; objection with a 30-day outcome (no processing with that sub-processor for that Customer, a reasonable alternative, or termination with a pro-rata refund); also: the Processor is named (Better ISMS EURL); new §5.3 liability cap (greater of 12 months of fees or EUR 1,000; faute lourde ou faute dolosive and Article 82 GDPR outside the cap);	Trust Center, change log and in-app changelog (per §7.2, non-materially-adverse)

	"fines" removed from the Customer indemnity; new §7.4 countersigned copies; §8 French law and Paris courts, data subjects keep their own forum. Mistral AI wording (Definitions, §2.1, §2.4, §3.1, §3.3, §4.2): inference, moderation, embedding and summarization requests go to Mistral's EU regional endpoint; Mistral web search runs on Mistral's global endpoint and is not EU-only. No new sub-processor, retention period or transfer.	
2026-10-21 (control-neutral; publication-only per §2.4)	Current AI routing snapshot corrected. §2.3, §2.4 (Anthropic, OpenRouter and xAI rows), §3.1 and §3.3 now describe the routing in production. With ADP off, paid Fast, Think and Beyond requests (including app-managed Plus trials) and free requests are served by GLM-5.3 via OpenRouter, pinned to Together AI, Fireworks AI and Inceptron (the 2026-09-01 rollout below, now complete). Anthropic is the pre-first-token failover for paid Fast and Think only; it is not a default destination for any plan and is not used for web search. xAI is no longer a default destination: it serves the logged-out risk-analysis demo and remains the operator rollback route. Web search is unchanged (Mistral source discovery with ADP off; unavailable under ADP). No new sub-processor, category of personal data or transfer; the disclosed routing envelope and account-level controls are unchanged.	Publication-only (correction of the current routing snapshot): published in the Trust Center and the customer-facing change log, effective 2026-10-21, per §2.4. Customers may object at any time (§2.4).
2026-10-20 (control-neutral; publication-only per §2.4)	Database backups produced in the EU; GitHub Actions removed. The daily database backup is made on Fly.io (EU, Paris), already a sub-processor, and encrypted there with public-key encryption before upload to Backblaze B2 (EU, Amsterdam), where each copy is kept for 90 days, then hidden and deleted one day later. GitHub, Inc. (GitHub Actions) is no longer a sub-processor. Backup retention is stated as the full storage lifecycle: deleted data can remain in encrypted backups for up to 91 days. No new sub-processor, category of personal data or transfer.	Publication-only (removal): published in the Trust Center and the customer-facing change log, effective 2026-10-20, per §2.4.
2026-10-16	Fly.io safeguard stated. The Fly.io row names the executed Fly.io DPA (countersigned 2026-10-08, binding since delivery on 2026-10-12) with the EU Standard Contractual Clauses. No change to processing.	Publication-only (correction)
2026-10-16	Chatbase listed in §2.4 (disclosure correction). Chatbase, Inc. (United States) was already on the Trust Center sub-processor list for the optional website support chat on the ISMS Copilot and heyGRC websites, documentation, consoles and trust center (visitor messages and an optional lead email for human follow-up). The §2.4 table now lists it too. Not used for	Publication-only (disclosure correction)

	product chat accounts or pull-request source code. No change to processing.	
2026-10	PostHog removed (product analytics no longer used). PostHog (EU, Frankfurt) is no longer a sub-processor; no replacement sub-processor is added.	Publication-only (removal): published in the Trust Center and the customer-facing change log, effective 2026-10-13, per §2.4.
2026-10-08 (control-neutral addition; noticed per §2.4)	Axiom added for log storage. Application logs of the production services hosted on Fly.io (chat, API, embed, heyGRC and Aevral) are copied to Axiom, Inc. and kept for 30 days in the EU (Frankfurt) for security monitoring and incident investigation. The logs hold the same pseudonymous identifiers (for example user, thread, message and job ids) and technical metadata that Fly.io and Sentry already process; our logging rule keeps chat and document content out of logs. Axiom's AI features are turned off for our account. Axiom, Inc. is a United States company and processes under its DPA with SCCs (Module Two). Customers may object at any time (§2.4).	Trust Center and change log
2026-10-07 (editorial; no change to processing)	Database backup wording. §2.3 and §3.1 (and the privacy policy and RoPA) describe the backup job as it runs today, on GitHub Actions; the text no longer refers to moving it. Facts unchanged.	Trust Center and change log
2026-10-04 (non-materially-adverse addition; noticed per §7.2)	DPA revision: new §10, United States State Privacy Laws. For personal data a Customer entrusts to ISMS Copilot for processing on its behalf, ISMS Copilot acts as a service provider under the CCPA/CPRA and as a processor under the other Applicable US State Privacy Laws, and the section carries the statutory contract requirements: limited and specified purposes and documented instructions only, no sale or sharing, no use outside the direct business relationship, no independent use including no development, training or improvement of any AI model, combining only as expressly permitted, compliance with all applicable CCPA provisions and regulations, notification and remediation rights, audit rights (§2.9), consumer-request and assessment assistance (§2.5, §2.7), deletion or return (§2.8), the same level of privacy protection with sub-processor flow-down (§2.4), breach notice per §2.6, and a certification of understanding. The DPA section adds restrictions and protections only; no change to processing, sub-processors, retention or transfers.	Trust Center, change log and in-app changelogs (per §7.2, non-materially-adverse)
2026-10-02	Backblaze and GitHub listed for database backups (job in use since 2026-03-09). A daily job on GitHub Actions (GitHub, Inc., United States) copies the production database, encrypts it with AES-256 and uploads it to Backblaze B2 (EU, Amsterdam), where the newest 90 daily copies are kept. Backblaze holds only encrypted copies and never the key. On GitHub's	Trust Center and this change log

	virtual machine the copy is unencrypted for the duration of the job, for every Customer, including those with Advanced Data Protection Mode enabled. Customers may object at any time (§2.4).	
2026-10-01 (non-materially-adverse clarification; Customer-Activated; no advance notice under §2.4)	heyGRC Google Drive: Open with. The organization owner or admin who connected Google Drive may turn on Open with, which adds Google's <code>drive.install</code> permission (menu entry only, no file access) so files can be opened with heyGRC from Google Drive instead of being picked in the Google Picker. Linked Drive files are checked for changes normally once a day (previously weekly; a file that keeps changing is checked less often, down to once a week) and re-read only when changed. Same sub-processor (Google LLC), same scope for file access (<code>drive.file</code>), same jurisdiction and transfer mechanism, same retention.	Trust Center + customer-facing change log + product-update channel
2026-10-01 (customer-activated integration and purpose extension)	heyGRC GitHub App disclosed: §1.1 (heyGRC named in the scope of this DPA), §1.2, §1.5, §1.6, §2.4 (GitHub as a Customer-directed platform, not an ISMS Copilot sub-processor; heyGRC AI processing note covering the default OpenRouter host pin and the organization opt-in EU inference path; Mistral AI purpose extended to heyGRC EU inference), §2.8 (pull-request review data retained while installed and deleted automatically 30 days after uninstall).	No new sub-processor is engaged. GitHub is Customer-activated and Customer-contracted, so the 30-day pre-notice rule does not apply; the Mistral AI row is an extended purpose of an already-disclosed sub-processor. Published in the Trust Center and customer-facing change log.
2026-10-01 (no new sub-processor; publication-only). Processing for an organization starts only when its owner or admin creates a <code>context:write</code> key.	heyGRC context layer. The Customer's own coding agent can send compliance context objects to heyGRC (§1.2, §1.6), which heyGRC compiles and uses in pull-request reviews. Inference runs on the sub-processors already listed for heyGRC reviews: OpenRouter currently pinned to Together AI, Fireworks AI and Inceptron (rollback hosts documented in DPA §2.4, heyGRC AI processing) with in-request zero data retention, or Mistral AI (EU) for organizations that enabled EU inference. The Customer's agent and the systems it reads are the Customer's own, not ISMS Copilot sub-processors. Retention and deletion in §2.8.	Trust Center + customer-facing change log + product-update channel
2026-09-28 (non-materially-adverse clarification; no advance notice under §7.2)	Aevral named in the scope of this DPA (§1.1, §1.2, §1.5, §1.6, new §1.7). Aevral uses sub-processors already on the ISMS Copilot list (OpenRouter with Together AI, Fireworks AI and Inceptron; Fly.io; Supabase; Sentry; Vercel; Stripe), published on the Aevral sub-processor list. GitHub is stated not to be a sub-processor for the Aevral and heyGRC GitHub Apps. No new sub-processor, no change for ISMS Copilot chat Customers.	Trust Center and change log
2026-09-28 (Customer-Activated; no advance notice under §2.4)	heyGRC Google Drive connection available to every heyGRC organization (allowlist removed). Still inactive until an owner or admin connects. §2.8 now states deletion on disconnect, uninstall	Trust Center + customer-facing change log + product-update channel

	and connector departure.	
On enabling for the Customer's organization (Customer-Activated; no advance notice under §2.4)	heyGRC Google Drive connection made available as a Customer-Activated Integration (see §2.4). Google LLC (Google Drive API, OAuth, Picker) becomes a sub-processor for a Customer's data only when that Customer's organization owner or admin connects Google Drive and picks documents. Scope <code>drive.file</code> (only picked files), read use, encrypted refresh token, deleted on disconnect, admin departure, or uninstall.	Trust Center + customer-facing change log + product-update channel
2026-09-01 (enabled; gradual rollout)	Chat model swap to the GLM-5.3 family. When enabled: paid chat Fast moves from <code>z-ai/glm-5.2</code> to <code>z-ai/glm-5.3-flash</code> (model-version swap on the SAME disclosed OpenRouter pin: Together AI, Fireworks AI, Inceptron; in-request ZDR and the account-level controls unchanged), and paid chat Think moves from the disclosed xAI (ZDR) endpoint (<code>grok-4.6</code>) to <code>z-ai/glm-5.3</code> on the same Together AI / Fireworks AI / Inceptron pin; xAI remains disclosed and available as the operator rollback route. The free/logged-out surface and the Plus+ over-cap overflow lane move from <code>z-ai/glm-4.7</code> (latency-sort, or the disclosed closed Google Vertex + Cerebras subset for overflow) to <code>z-ai/glm-5.3-flash</code> on that same pin. No new sub-processor is introduced (Together AI, Fireworks AI, Inceptron, and xAI were all disclosed 2026-06-23 / 2026-07-21); the Think and overflow conversions are routing-scope changes within the disclosed allowlist, classified control-neutral. Advanced Data Protection (Mistral, EU) remains available and suppresses the OpenRouter path entirely.	Control-neutral, publication-only under §2.4: Trust Center + in-app changelog at enable time. No 30-day advance notice required; customers may object or enable ADP at any time.
2026-07-23 (control-neutral; strengthening; no advance notice)	DeepInfra removed from the <code>z-ai/glm-5.2</code> provider set; in-request Zero Data Retention added. DeepInfra (a previously-allowlisted OpenRouter underlying provider) is removed from the closed only list for <code>z-ai/glm-5.2</code> across chat, the API, and heyGRC review, leaving Together AI, Fireworks AI, and Inceptron; DeepInfra carried ~0% of <code>glm-5.2</code> traffic and was dropped for lacking a publicly-verifiable SCC-bearing DPA. In the same change, request-level <code>data_collection: "deny"</code> + <code>zdr: true</code> were added to the <code>glm-5.2</code> provider pins (the shared chat <code>glm-5.2</code> route, which the partner premium tier also uses; the API <code>glm-5.2</code> pin; and the heyGRC review <code>glm-5.2</code> pin), so zero-retention and no-training are enforced in-request and no longer rest on the OpenRouter account-level toggle alone. DeepInfra remains on the account allowlist for other models.	Control-neutral / strengthening under §2.4: the recipient set is narrowed (a provider removed) and retention controls are strengthened, so the 30-day advance-notice rule for materially-adverse changes does not apply. Communicated via in-app changelog and the Trust Center on the effective date.
2026-07-21 (control-neutral; no advance notice)	xAI (X.AI LLC), via the OpenRouter "xAI	Published in the Trust Center and

notice; published effective 2026-07-21)	(ZDR)" endpoint (Grok models) - already present on the OpenRouter account allowlist and already serving the logged-out demo - now activated for paid Customer-Content and made the default OpenRouter underlying-provider destination for paid plans (Plus and above), ADP off, for the chat Fast and Think modes and the Beyond assistant, replacing the prior Anthropic default; Anthropic retained as the pre-first-token failover. Classified control-neutral under example (c) of the §2.4 wording then in force: zero data retention (as configured and verified 2026-07-19, stronger than the Anthropic path it replaces), no-training, US (non-PRC) jurisdiction, and the SCC transfer mechanism all continue to apply under unchanged OpenRouter account-level controls; the Advanced Data Protection (Mistral, EU, zero-retention) guarantee remains available to suppress the OpenRouter path entirely.	customer-facing change log, effective 2026-07-21 (control-neutral, publication-only per §2.4). Customers may object at any time or enable Advanced Data Protection.
2026-10-01 (target launch day)	Authenticated web-search Mistral bridge. Mistral performs intent classification for eligible users and provider-attested source-URL discovery when ADP is off. Its generated search prose is discarded. The EU Fly.io service exact-fetches attested or user-submitted public pages, and bounded evidence goes to the already applicable answer/Beyond model. No additional semantic grounding model receives the draft. Authenticated web search fails closed under ADP; exact ADP page/document fetch remains available. The direct-Brave route (ISMS Copilot calling the Brave API itself) is not constructed or called. Raw discovery and page evidence is not persisted by ISMS Copilot.	Trust Center and in-app disclosure update. Mistral and Fly are already-disclosed recipients. Mistral's web search runs on Mistral's stateful Conversations/Agents endpoint (outside Mistral's zero-retention posture) and uses Brave (US) as its web-search sub-processor, so the non-ADP query hop reaches Brave (US) under standard retention and is not EU-only or zero-retention; see the internal launch decision record. This entry is effective only when the feature is published and live.
2026-08-12 (earliest; materially adverse; pending, not the launch default)	Potential future Brave-direct external recipient (distinct from Brave's launch role as Mistral's web-search sub-processor). A future non-ADP route in which ISMS Copilot would call the Brave Search API itself may send only a bounded latest-request query to Brave Search in the United States; no files, memories, workspace context, account identifiers, or user/device identifiers would be sent. Brave standard query logs may be retained up to 90 days; SCCs apply; Brave's DPA excludes Search Query Data from processor scope, so Brave is disclosed as a limited external recipient rather than represented as an Article 28 processor.	Notice package drafted 2026-07-13. 2026-08-12 is available only if notice is actually issued on 2026-07-13; otherwise activation moves to at least 30 days after actual notice. Founder acceptance is also required. Brave remains dark meanwhile.
2025-11	Initial sub-processor list established (Anthropic, Mistral AI, Supabase, Stripe, ConvertAPI, PostHog, Sentry, Vercel, Fly.io, SendGrid, Kit)	Initial publication
2026-04-16	heygrc Slack integration made available to paid organizations as a Customer-Activated Integration (see §2.4 —	Documented in this DPA revision (2026-04-27); the integration's availability for installation went live

	Customer-Activated Integrations). Slack Technologies, Inc. only becomes a sub-processor for a Customer's data when that Customer's organization owner explicitly installs the bot from the Connectors page; uninstall hard-deletes all integration records. Activation gated on paid plan; free / null-plan organizations cannot install. Transfer mechanism: SCCs.	2026-04-16. The 30-day advance-notification rule does not apply to Customer-Activated Integrations because no Customer's data flows to Slack unless and until that Customer's organization owner takes an explicit installation step.
2026-05-27	OpenRouter added as routing aggregator for free-tier and null-plan users with ADP off. Routes initially to four allowlisted underlying providers: Inception, DeepInfra, Cerebras, Google Vertex. Account-level controls enforced: mandatory Zero Data Retention, Free/ Paid Training Disallowed, Free Publication Disallowed, closed 4-provider allowlist, PRC-jurisdiction blocklist (Alibaba Cloud Int., Baidu Qianfan, DeepSeek, Moonshot AI, Xiaomi, Z.AI). Free-tier users previously served by reserved direct-API paths are now served via OpenRouter. ADP path (Mistral) and paid path (Anthropic) unchanged.	2026-04-27 (30-day customer email under prior §2.4 wording)
2026-05-28	Essential plan (\$12/mo) routing. The new Essential plan, with ADP off, routes through OpenRouter restricted to a closed two-provider subset of the existing allowlist: Google Vertex and Cerebras. No new sub-processor is introduced: OpenRouter, Google Vertex, and Cerebras are already on the list above (added 2026-05-27), with the same account-level controls (mandatory ZDR, training-disallowed, publication-disallowed, PRC-blocklist) and transfer mechanisms (SCCs; EU-US DPF for Google Vertex). ADP-on (any plan, including Essential) continues to route to Mistral (EU). This is a routing-scope extension to a new paid plan using already-disclosed sub-processors.	No new sub-processor is added, so the 30-day advance-notice rule for sub-processor additions does not apply. Communicated via in-app changelog and the Trust Center on 2026-05-28.
2026-06-23	OpenRouter underlying-provider allowlist expanded from 4 to 7. Three additions: Together AI, Fireworks AI, Nebius. First control-neutral sub-processor change under the new §2.4 wording (also effective 2026-06-23). Account-level controls unchanged. Each addition evaluated against the same privacy-and-jurisdiction bar as the original four (zero-retention posture, no training on customer data, published DPA, non-PRC jurisdiction, public deployment-region disclosure). Novita AI evaluated under the same review and not added (public-disclosure opacity on infrastructure location prevents evidencing the §3.1.4 PRC-jurisdiction control).	2026-05-26 (in-app changelog + Trust Center publication). Bundled labeled entry in the May 2026 monthly product-changelog email shipping in the first days of June. A more detailed internal change record is available on request.
2026-06-13 (issued; never published)	Amendment B (superseded by Amendment C before publication). Amendment B would have expanded OpenRouter scope to paid-plan overflow routing after the 4-hour token cap, gated	Not published. Superseded by Amendment C before any customer notice took effect.

	by an explicit per-session end-user election, restricted to the Google Vertex / Cerebras subset. It was issued on 2026-06-13 but never published or made operative, and is superseded in full by Amendment C below. Recorded here for historical completeness only.	
2026-06-23 (control-neutral)	OpenRouter allowlist added as a permitted paid routing destination alongside Anthropic, for any chat mode. For paid plans (Plus, Standard, Pro, Business) with ADP off, the closed 7-provider OpenRouter allowlist becomes a permitted AI-routing destination alongside Anthropic, for any chat mode (fast, think, or agentic/Beyond), mirroring the free / null-plan path. ISMS Copilot may serve a paid request via Anthropic or via any allowlisted provider, and may move any paid cohort or mode between them over time without a further sub-processor change. Anthropic is not removed: it remains a disclosed paid provider and the circuit-breaker failover context, no longer the sole paid path. No new sub-processor is introduced (OpenRouter and all seven underlying providers are already disclosed). Account-level controls (mandatory ZDR, training-disallowed, publication-disallowed, allowlist, PRC-blocklist) and the underlying-provider transfer mechanism are unchanged. Retention for any traffic served by Anthropic follows Anthropic's standard commercial API terms (ordinary ~30-day deletion; safety-flagged content up to 2 years; safety-classification scores up to 7 years; not training); traffic served by an allowlisted OpenRouter provider is zero-retention. ADP remains the account-level EU-residency guarantee.	Control-neutral change under §2.4; communicated via in-app changelog and the Trust Center on the effective date. The 30-day advance-notice rule for materially-adverse changes does not apply because no new sub-processor is introduced, the same account-level controls and underlying-provider transfer mechanism apply unchanged, every allowlisted provider is equal-or-stronger than Anthropic on retention / training / jurisdiction, and ADP remains available to keep all processing in the EU.

All future sub-processor changes will be documented here with effective date, sub-processor name and location, nature of change (addition, removal, replacement), processing purpose, and customer notification date.

Signatures

Better ISMS EURL signs this offer of the Data Processing Agreement, version 2026-10-26-r2 (source fingerprint fed4cacb8b5a5dd21ef8c6fa5dc02fbb180963a591c08bf784f7b150f2909c27, template revision 2), through its g rant. The Customer accepts it as set out in DPA  7.4(a). This document is established in electronic form; each party can keep a copy, and the agreed terms remain available at the address on the cover page.

For and on behalf of Better ISMS EURL

Tristan Roth

Electronic signature for this DPA version only

Tristan Roth, G rant

Approved for issue on 2026-10-26

For and on behalf of the Customer

Name: _____

Title: _____

Date: _____

The signatory confirms they are authorised to sign for the Customer.