

Aevral Security Questionnaire Lite

Version 1.0. Posture manifest as of 2026-10-01; verified against the running worker f20ba97a2 on 2026-10-02.

Questions follow CAIQ v4 domains and SIG Lite topics. Answers cite the security posture manifest by claim id. Full CAIQ or SIG on request.

An operational summary. The Data Processing Agreement (section 1.7, <https://trust.ismscopilot.com/dpa>) and the Aevral sub-processor page (<https://trust.ismscopilot.com/en/aevral/subprocessors>) are binding. Live version: <https://aevral.com/security>

Audit and assurance (CAIQ A&A)

aaa-01 Is the vendor certified to ISO/IEC 27001 with this service in scope?

No. Better ISMS EURL, which makes Aevral, is working toward ISO 27001 certification with Aevral in the intended scope. It is not certified yet.

Source: policy.company, gap.certification

aaa-02 Does the vendor have a SOC 2 Type I or Type II report?

No. There is no SOC 2 report.

Source: policy.company

aaa-03 Has the service had an independent penetration test?

No. Aevral has not had an independent penetration test yet.

Source: gap.pentest

aaa-04 Does the vendor hold any other security certification or attestation for the service?

No. Aevral has no certification yet.

Source: gap.certification

aaa-05 How are the vendor's published security statements kept accurate?

Each statement on aevral.com/security is either checked against the engine code by automated tests, read from the running service, or marked as a policy statement that points to its binding document. The page shows the commit it was verified against.

Source: app.permissions.live, flags.runtime_state, model.rail_today

Application and interface security (CAIQ AIS)

ais-01 Are inbound webhooks authenticated?

Yes. Every GitHub webhook is verified with an HMAC-SHA256 signature over the raw body, compared in constant time.

Source: webhooks.verified

ais-02 Are request sizes limited before input is processed?

Yes. Webhook bodies are size-capped before they are read, and scans and pull request reviews have size limits on the code they read.

Source: webhooks.verified, code.access_and_processing

ais-03 Which external events does the application act on?

Installation, repository selection and pull request events from GitHub, and issue comment events only for the @aevral review command. Any other issue comment is ignored.

Source: webhooks.verified, review_command.behavior

ais-04 Is customer code executed by the service?

No. Customer code is read and analyzed, never executed.

Source: code.access_and_processing

ais-05 Are API keys for the service's own API stored securely?

Yes. Aevral API keys are stored only as SHA-256 hashes and looked up by hash.

Source: data.api_keys_hashed

ais-06 Are new features enabled by default?

No. New capabilities ship switched off in code and stay off unless turned on separately for the running worker: the MCP connection, MCP scans, the review command and EU-only processing. The review command is live; the running worker reports the state of each switch publicly.

Source: flags.defaults_off, flags.runtime_state, review_command.behavior

ais-07 Can customers verify which optional capabilities are switched on?

Yes. The running worker reports the state of each switch publicly.

Source: flags.defaults_off, flags.runtime_state

ais-08 Who can trigger a review from a pull request comment?

Only the repository owner, a member of its organization or a collaborator, as GitHub marks them, never a bot. Commenting @aevral review at the start of a new comment in the pull request conversation asks for a review of the whole pull request. The same commit can be re-reviewed at most 3 times, at least 10 minutes apart, not while a review runs, and is not counted again.

Source: review_command.behavior

Artificial intelligence (SIG Lite)

ai-01 Which AI model processes customer code, and who hosts it?

Scans and pull request reviews run on GLM-5.3 or GLM-5.3 Flash, open-source models, through OpenRouter, depending on the job and plan, pinned to the hosting providers on the Aevral sub-processor page. In practice requests are processed in the US.

Source: model.rail_today

ai-02 Are model requests restricted to approved hosting providers?

Yes. Every OpenRouter request is restricted to a closed list of hosting providers. The list is never widened without a sub-processor notice.

Source: model.provider_pin

ai-03 Is customer code used to train AI models?

No. Your code is not used to train models.

Source: policy.no_training

ai-04 Do model providers retain or collect customer code?

Each model request denies data collection and requires zero data retention. The account-level setting is attested by the provider.

Source: model.provider_pin, policy.no_training

ai-05 Does the AI model have tools, shell access or the ability to run code?

No. The model has no tools or shell.

Source: code.access_and_processing

ai-06 How is prompt injection from repository content handled?

Repository content is treated as untrusted data in prompts, model output is neutralized before it is posted, and every finding must match the code verbatim.

Source: prompt.untrusted_input

ai-07 Can AI output change customer repositories?

No. Output is posted only as advisory check runs and pull request review comments. Suggested fixes are text only; Aevral never commits, pushes, merges or approves.

Source: `writes.advisory_only`

ai-08 Can customers see which model and inference path are in use?

Yes. The running worker reports its rail and model ids publicly.

Source: `model.rail_today`

ai-09 Is AI processing available in the EU only?

Not yet. An EU-only option is announced: an organization setting, changed only by an owner or admin, that sends scans and pull request reviews started after the change only to a model provider in the EU. When EU processing is unavailable, jobs wait and retry in the EU for up to 6 hours, then end not reviewed at no charge; they never fall back to the US. Findings returned through an AI tool you connect follow that tool's own processing.

Source: `gap.eu_only`, `announced.eu_only_processing`, `announced.model_eu_path`

ai-10 Are AI findings validated before they reach the customer?

Findings must match the code verbatim before they are posted. They are suggestions in an advisory Check, not blockers.

Source: `prompt.untrusted_input`, `writes.advisory_only`

Business continuity and resilience (CAIQ BCR)

bcr-01 Does the vendor have a business continuity arrangement for the service?

Yes. Aevral is made by a founder-led company with a documented continuity and backup arrangement. Details on request.

Source: `policy.team`

bcr-02 How is key-person dependency addressed?

The company is founder-led. A documented continuity and backup arrangement covers key-person dependency; details on request.

Source: `policy.team`, `gap.team`

Change control and configuration (CAIQ CCC)

ccc-01 Are security-relevant changes communicated before they ship?

`aevral.com/security` lists planned changes under Announced changes, each with the date it was announced, and records changes to the posture in a dated changelog. The review command was listed there before it went live. Currently announced: the standard-sync pull request, EU-only processing and design review of plan documents.

Source: `review_command.behavior`, `announced.contents_write_standard_sync`, `announced.eu_only_processing`, `announced.model_eu_path`, `announced.design_review`

ccc-02 Does the service make changes to customer code or repository settings?

No. It writes only advisory check runs and pull request review comments. It never commits, pushes, creates branches, merges or approves.

Source: `writes.advisory_only`

ccc-03 Does the App hold permissions that no feature uses yet?

Yes, one. Contents write was added on 2026-10-01 for the announced standard sync, and no code uses it yet; no token Aevral mints carries it, and it is for a pull request on the branch `aevral/standard-sync` only. Issues read, added the same day, is used by the live `@aevral` review command, which acts on issue comment events; no Aevral token carries it.

Source: `app.permissions.live`, `tokens.capped_below_app`, `announced.contents_write_standard_sync`, `review_command.behavior`

ccc-04 Are new capabilities rolled out behind switches?

Yes. They ship switched off and the running worker reports each switch publicly.

Source: flags.defaults_off, flags.runtime_state

Cryptography and key management (CAIQ CEK)

cek-01 Are credentials or tokens stored in plaintext?

No. GitHub tokens are never stored; they live in memory only. Aevral API keys are stored only as SHA-256 hashes.

Source: tokens.per_job_scoped, data.api_keys_hashed

cek-02 Is customer data encrypted in transit and at rest?

Yes. Data is encrypted in transit with TLS and at rest in the database, as committed in the security measures of the Data Processing Agreement (section 2.3). The worker is served over HTTPS only; plain HTTP is redirected.

Source: data.encryption

cek-03 How long do GitHub access tokens live?

Tokens are minted per job and expire after GitHub's default of one hour.

Source: tokens.per_job_scoped

Data location and hosting (CAIQ DCS)

dcs-01 Where is customer data stored?

In a Supabase database in the EU (Frankfurt).

Source: data.database

dcs-02 Where is customer code processed?

The worker processes code in memory on a server in Paris. Model inference runs through OpenRouter on pinned providers and is in practice processed in the US.

Source: code.access_and_processing, model.rail_today

dcs-03 Is all processing confined to the EU?

No. Storage is in the EU and the worker runs in Paris, but model inference is in practice processed in the US today. An EU-only option is announced, not live.

Source: data.database, model.rail_today, gap.eu_only

dcs-04 Where are the vendor's sub-processors listed?

On the Aevral sub-processor page of the trust center, which is binding together with the Data Processing Agreement.

Source: model.rail_today, policy.retention_dpa

Data security and privacy (CAIQ DSP)

dsp-01 What customer code does the service read?

Scans read one archive of the repository at a single fixed commit. Pull request reviews read the diff, plus full files at the head commit for some changed files. Both have size limits.

Source: code.access_and_processing

dsp-02 Is customer code kept after processing?

Code is processed in memory. Findings keep code excerpts, which are shown in the Aevral console. Customer data is kept while Aevral is installed and for up to 12 months after uninstall.

Source: code.access_and_processing, data_flow.mcp_excerpts, policy.retention_dpa

dsp-03 Are secrets in customer code protected from being sent to the model?

Partly. Scans leave out files whose paths look like secrets and redact credential patterns. Pull request reviews hold back sensitive files and name them in the Check as not sent. Detection is pattern-based and can miss secrets that do not match a known pattern.

Source: secrets.pattern_based, gap.secret_detection

dsp-04 What is the data retention period?

Customer data is kept while Aevral is installed and for up to 12 months after uninstall. The Data Processing Agreement, section 1.7, is binding.

Source: policy.retention_dpa

dsp-05 Is data deleted on request?

Yes, within 30 days, as set out in the Data Processing Agreement, section 1.7.

Source: policy.retention_dpa

dsp-06 Is data deleted automatically after the customer uninstalls?

Yes. Twelve months after the App is uninstalled or a repository is removed, that repository's scan reports, pull request review content, findings and repository name are deleted or scrubbed automatically, and the installation record loses the account and installer identifiers (the installation number and dates stay). Non-content organization, billing and scan metadata (dates, commit ids, status; no code) is kept until the account is deleted. Deletion is also honored on request within 30 days.

Source: retention.automatic, policy.retention_dpa, gap.post_uninstall_metadata

dsp-07 Which records are deleted automatically today?

Worklist findings 24 months after they were last seen; twelve months after a repository is removed or the App is uninstalled, that repository's scan reports, pull request review content, findings and repository name, and the account and installer identifiers on the installation record; audit log entries 12 months after the organization is deleted, and every entry after 24 months; disconnected AI tool connections 90 days after disconnect.

Source: retention.automatic

dsp-08 Do scan reports and pull request review records have an automatic retention limit?

Yes. Twelve months after the repository is removed or the App is uninstalled, the content of scan reports and pull request review records is scrubbed automatically. While the repository is installed, the contractual retention applies: kept while installed and up to 12 months after uninstall, with deletion on request within 30 days.

Source: retention.automatic, policy.retention_dpa, gap.post_uninstall_metadata

dsp-09 Is customer data isolated between customers?

Yes. Row level security is on for every Aevral table exposed through the database API and limits access to members of your organization. Internal tables sit in a schema the API does not expose.

Source: data.database

dsp-10 Is customer data shared with tools outside the service?

Code excerpts in findings are returned to an AI tool only if you connect one through MCP. Where that tool processes them is set by the tool you choose.

Source: data_flow.mcp_excerpts

dsp-11 Is a Data Processing Agreement available?

Yes. The Data Processing Agreement is published on the trust center; section 1.7 covers Aevral and is binding.

Source: policy.retention_dpa

Governance, risk and compliance (CAIQ GRC)

grc-01 Does the vendor run an information security management system?

Better ISMS EURL is working toward ISO 27001 certification with Aevral in the intended scope. It is not certified yet.

Source: policy.company

grc-02 Does the vendor disclose known gaps in its controls?

Yes. aevral.com/security lists what we do not have yet, including SSO, console MFA, an independent penetration test and certification, and that non-content organization metadata is kept until the account is deleted.

Source: gap.sso, gap.mfa, gap.post_uninstall_metadata, gap.pentest, gap.certification

grc-03 Which document is binding if this questionnaire and the contract differ?

The Data Processing Agreement, section 1.7, and the Aevral sub-processor page. This questionnaire is an operational summary.

Source: policy.retention_dpa

grc-04 Who is the legal entity behind the service?

Better ISMS EURL, Paris, France.

Source: policy.company

Identity and access management (CAIQ IAM)

iam-01 Which GitHub permissions does the service hold?

The App holds Contents read and write, Issues read, Metadata read, Checks write, and Pull requests read and write. Every token Aevral mints is capped below that: Contents read, Metadata read and Checks write, plus Pull requests write for pull request reviews. No minted token carries Contents write or Issues access. No code uses Contents write yet; Issues read serves the review command, which acts on issue comment events.

Source: app.permissions.live, tokens.capped_below_app

iam-02 Does the service hold administration, secrets, workflows, deployments, Issues or Contents write permissions?

No administration, secrets, workflows or deployments permission. The App holds Issues read, added on 2026-10-01 for the review command, and Contents write, added the same day for the announced standard sync, which no code uses yet. No token Aevral mints carries either.

Source: app.permissions.live, tokens.capped_below_app

iam-03 Are repository access tokens scoped per job?

Yes. Each scan or pull request review gets its own GitHub token, minted for that job, limited to one repository and capped below the App's permissions, with no Contents write or Issues access.

Source: tokens.per_job_scoped, tokens.capped_below_app

iam-04 Are installation-wide tokens used?

Only for one stated purpose: refreshing the list of repositories you selected, with the same capped permissions. No other part of Aevral can obtain an installation-wide token.

Source: tokens.repo_list_exception

iam-05 Does the console support single sign-on (SSO or SAML)?

No. There is no SSO or SAML for the console yet.

Source: gap.sso

iam-06 Does the console support multi-factor authentication?

Not yet. There is no multi-factor authentication in the Aevral console yet.

Source: gap.mfa

iam-07 Who in the customer organization can see Aevral data?

Members of your organization. Row level security limits access to them.

Source: data.database

iam-08 How is the service's own database access restricted?

The worker uses its own database role without row level security bypass, which may only call a fixed list of functions and holds no direct table access. It refuses to start with any other role's key.

Source: data.database

iam-09 Are changes to team membership recorded?

Yes. An append-only audit log records team changes.

Source: data.audit_log

Infrastructure security (CAIQ IVS)**ivs-01 Is the processing environment sandboxed?**

Yes. The worker runs in the Deno sandbox with outbound network limited to a list of hosts and environment access. It has no permission to write files.

Source: code.sandbox_today

ivs-02 Can the processing environment run subprocesses or native code?

No. It has no permission to run subprocesses, load native code or read system information, and it loads no code at run time beyond what the image cached when it was built.

Source: code.sandbox_today

ivs-03 Is outbound network traffic from the processing environment restricted?

Yes. Outbound network is limited to an exact list of hosts in these categories: GitHub, the model providers, the EU database, error monitoring, payments and alerting. A call to any other host is refused by the sandbox.

Source: code.egress_allowlist

ivs-04 Can the processing environment write to disk?

No. The worker has no write access to disk at all, not even to a temporary folder. Repository archives are read in memory.

Source: code.no_file_writes, code.sandbox_today

ivs-05 Is code analysis bounded to a known version of the repository?

Yes. Scans read one archive at a single fixed commit; pull request reviews read the diff and head-commit files. Both have size limits.

Source: code.access_and_processing

Logging and monitoring (CAIQ LOG)**log-01 Is there an audit log of security-relevant events?**

Yes. An append-only audit log records the scan lifecycle and team changes.

Source: data.audit_log

log-02 Is the audit log tamper-resistant?

It is append-only.

Source: data.audit_log

log-03 Is there a retention or purge policy for audit logs?

Yes. Audit log entries of a deleted organization are purged 12 months after the deletion, and every entry is purged after 24 months.

Source: retention.automatic

Security incident management (CAIQ SEF)

sef-01 Will customers be notified of a personal data breach?

Yes. We notify affected customers within 48 hours of confirming a personal data breach.

Source: policy.breach_notice

sef-02 Is the breach notification commitment contractual?

Yes. It is set out in the Data Processing Agreement.

Source: policy.breach_notice

sef-03 Are records available to support incident investigation?

An append-only audit log records the scan lifecycle and team changes.

Source: data.audit_log

Supply chain and transparency (CAIQ STA)

sta-01 Are sub-processor changes notified in advance?

The list of hosting providers for model requests is never widened without a sub-processor notice. Planned changes, such as the EU model path, are listed under Announced changes on aevral.com/security.

Source: model.provider_pin, announced.model_eu_path

sta-02 Which third parties receive customer code?

The model hosting providers listed on the Aevral sub-processor page, through OpenRouter. In practice requests are processed in the US.

Source: model.rail_today, model.provider_pin

sta-03 Is the vendor's team size and continuity disclosed?

Yes. Aevral is made by a founder-led company with a documented continuity and backup arrangement. Details on request.

Source: policy.team, gap.team

sta-04 Does the vendor publish what it does not do yet?

Yes. Gaps and announced changes are published on aevral.com/security with a dated changelog.

Source: gap.sso, gap.eu_only, announced.eu_only_processing

Threat and vulnerability management (CAIQ TVM)

tvm-01 Has the service been tested by an independent third party?

No. Aevral has not had an independent penetration test yet.

Source: gap.pentest

tvm-02 How are secrets in scanned code detected?

With pattern-based detection, which can miss secrets that do not match a known pattern. Aevral is not a secret scanner.

Source: secrets.pattern_based, gap.secret_detection