Bevel AI

CUSTOMER OVERVIEW

Bevel AI Security, Privacy and Responsible AI

VERSION	ORGANISATION	AUDIENCE	STATUS
1.3	Bevel AI Oy, Espoo, Finland	Prospective and existing customers	Updated 24 November 2026



Overview

Bevel AI Oy ("Bevel AI") operates Bevel Studio ("Studio"). Studio helps people create content with AI assistance and human editorial control. This overview explains the service, customer responsibilities and the choices available for custom Enterprise solutions.

Capabilities described below reflect current product functionality; availability depends on enabled features and the selected deployment. Customer agreements define contractual scope and commitments; this document does not amend them or create a delivery commitment or service-level agreement.

Documents and evidence

Resource	Scope and current status
Customer Terms of Service	Markdown copy of the initial SaaS draft. Commercial terms, billing and exit provisions remain under review.
Privacy Policy	Markdown copy of the initial SaaS draft. Privacy roles, providers, retention and related disclosures need alignment with actual processing.
Data Processing Agreement (DPA)	Relevant data-processing practices are described in this document. A formal SaaS DPA is not included in this package; Custom Enterprise arrangements are agreed with the customer.
Provider and processing information	Section 4 identifies the current hosting, storage and AI providers and explains when they process customer data. The formal subprocessor register and provider-account review are being completed.
Security and recovery evidence	Bevel AI uses automated type, build and test checks in development. This package does not include an independent penetration-test or recovery-test report; customer-specific evidence requirements can be discussed during the assessment.

The standard SaaS terms are designed for Studio SaaS. Custom Enterprise solutions are covered by separate agreements defining scope, pricing, data processing and operational responsibilities.

1 Service scope and offerings

What does Bevel Studio do?

Studio supports creating, editing, localising, versioning and exporting presentations, learning courses, videos and related media. AI assists with content and layout, with image, video, speech and transcription functions depending on enabled services. Users can inspect, edit and reject results.

Which offering should we consider?

Studio SaaS is our recommended starting point when its standard capabilities fit your needs. Custom Enterprise is appropriate to assess when workflow, integration or deployment requirements justify additional tailoring.

Offering	Benefit	Trade-off and responsibility
Studio SaaS	A standard hosted service operated by Bevel AI; the simplest starting point.	Customers use the available service settings and provider arrangements, and manage their users, content and publication decisions. Current providers and data flows are described in section 4.
Custom Enterprise, managed by Bevel AI	A tailored solution with Bevel AI operating the agreed components.	Additional design, setup and commercial agreement. Dedicated infrastructure, integrations and operational scope depend on the selected solution.
Custom Enterprise, customer-managed	Greater customer control over deployment and operation.	Requires customer IT/operational capacity. Hosting, access, patching, monitoring, backups and incident responsibilities must be allocated explicitly.

For customer-managed Enterprise, we recommend a structured handover and induction with a designated member of the customer's IT or operations team. Ongoing Bevel AI support and maintenance are subject to negotiation. Enterprise options are matters to scope together; they do not commit Bevel AI to deliver every configuration.

2 Responsible AI and human review

Who is responsible for reviewing the output?

Customers are responsible for reviewing and approving AI-assisted content before publishing, distributing or using it. Industrial, technical and safety-related training must be reviewed by a qualified person against approved procedures.

Studio is designed for human-led authoring and does not require a general "generated with AI" icon. However, some content or uses may require disclosure under applicable law. Customers should assess the intended use and apply any required disclosures.

AI can produce errors, omissions, bias or unsuitable or infringing content. Automated checks do not guarantee accuracy or suitability for a particular purpose. Studio assists with content creation and must not be used for safety-critical or consequential decisions about individuals, such as employment, credit or admission.

Users must have the right to use their source material and must not use Studio for unlawful or harmful purposes, including deceptive impersonation, discrimination, surveillance or infringement.

3 Customer data and model training

What data is processed?

Depending on the features used, processing can include account/workspace information, documents, prompts and chat context, images, audio, generated content, project history, usage/cost records, and support or billing information. AI tasks send relevant instructions, context and media to the configured services. A connected customer agent introduces its own processing path.

Is our content used for model training?

Bevel AI does not use customer inputs or outputs to train or fine-tune models. Studio SaaS and providers for Enterprise solutions managed by Bevel AI are configured to prevent customer content from being used for training and to use the shortest available retention settings. For Enterprise solutions using customer-owned provider accounts, the customer configures these settings as part of the agreed setup. Limited request metadata may still be retained for security, billing and service operation.

Can we upload confidential or sensitive information?

Studio can process confidential business content when the selected service and provider configuration meet the customer's requirements. Customers should minimise unnecessary personal data and consult Bevel AI before uploading regulated, special-category or highly sensitive information. Enterprise solutions can address additional hosting, provider and access requirements.

What are the privacy roles?

Customers generally act as the data controller for personal data included in their content, while Bevel AI acts as the processor and handles that data according to customer instructions. Bevel AI acts as a controller for its own purposes, such as account administration, billing and security. The applicable privacy terms and Enterprise agreements define these roles for the selected service.

4 Hosting, providers and processing locations

How does customer data flow through Studio?

The diagram shows the main content and AI-processing paths for Studio SaaS. Arrows identify the data exchanged when the relevant feature is used.

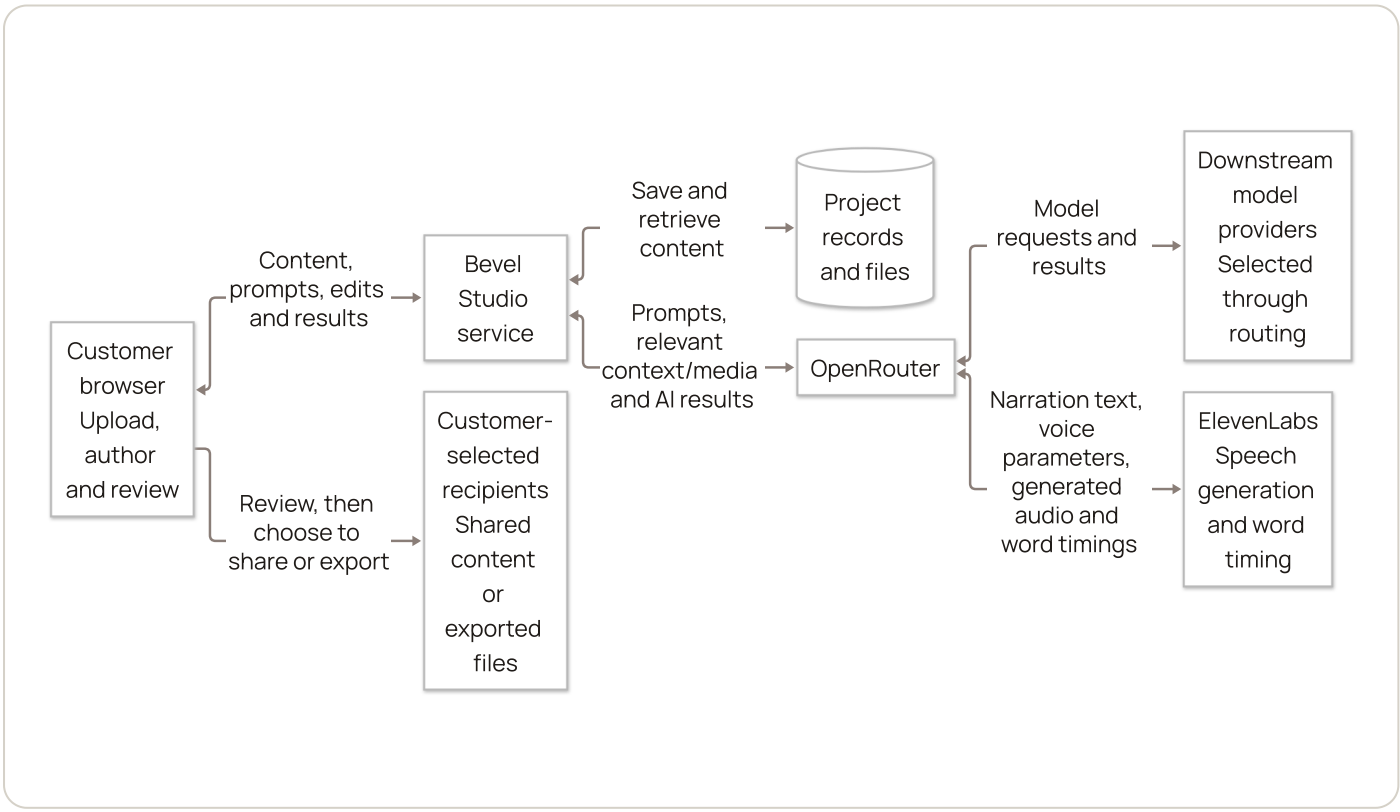

The production hosting and storage arrangement is described below. OpenRouter and ElevenLabs are confirmed active AI providers; Studio reaches ElevenLabs only through OpenRouter, as a downstream provider for speech. Requests also involve service and usage metadata. The diagram summarises the main paths; optional identity, email, billing, diagnostics, remote resources and connected agents are described below. Processing locations and provider handling depend on the services and configuration in use.

Customers review outputs before use. The review step shown is a customer responsibility, not an automatic publication gate. For Enterprise, hosting, storage, AI providers and operating responsibilities follow the selected solution; disabling built-in AI removes those generation paths, while separately connected agents require their own assessment.

Where is SaaS data hosted and processed?

Since 20 November 2026, Studio SaaS production runs on Microsoft Azure. The Studio web front end (studio.bevelai.com) is served by Azure Static Web Apps. The application API (api.bevelai.com) runs on Azure Container Apps in the West Europe region. The database is Azure Database for PostgreSQL Flexible Server with private network access, and customer projects and uploaded files are stored in private Azure Blob Storage, both in West Europe. Application secrets are kept in Azure Key Vault.

All production data was migrated from the previous Fly.io deployment on 20 November 2026. That deployment is stopped; its database and storage are kept intact as a fallback during a short post-migration observation window and will then be deleted. A separate staging deployment, Studio Next, runs on Azure in the France Central region with test accounts and synthetic content.

Bevel AI selects SaaS providers and available processing locations based on data protection, security, reliability and performance. AI processing uses OpenRouter, including ElevenLabs as its downstream provider for speech, and its location depends on the selected service and routing.

For Custom Enterprise, hosting, storage, AI providers and processing locations are agreed with the customer as part of the selected solution.

Which services can receive data?

Studio uses OpenRouter for AI features; speech requests are served by ElevenLabs through OpenRouter. The table below summarises the services that may process customer data, the data involved and when each service applies. Enterprise provider arrangements depend on the selected solution.

Path	Data involved when used	Applicability
OpenRouter and downstream model providers	Prompts, relevant context/media, responses and request/usage metadata.	OpenRouter is active for Studio AI. Enterprise can use customer OpenRouter credentials. The serving provider depends on routing.
ElevenLabs speech through OpenRouter	Narration text, selected voice parameters, generated audio, word timings and request metadata.	ElevenLabs is active for Studio AI as an OpenRouter downstream provider: it generates speech and transcribes the generated audio to time its words. Studio has no direct ElevenLabs integration. Enterprise applicability depends on the selected solution.
Hosting, database and object storage	Accounts, workspace/content records and stored files, as configured.	Microsoft Azure for Studio SaaS, as described above, or hosting and storage services agreed with the customer for Custom Enterprise.

Path	Data involved when used	Applicability
Identity, email and billing services	Identity exchanges, password-reset messages and billing-related information respectively.	Depending on configuration, Studio may use a customer-selected identity provider, Resend for password-reset emails and Stripe for billing.
Telemetry, remote resources and external agents	Diagnostics, resource requests or workspace context according to configuration and granted access.	Applies only when telemetry, remote resources or external agents are configured. Bevel AI assesses services it selects; customers assess services they select or connect.

What can Enterprise customers control?

Enterprise workspace settings support credentials configured by Bevel AI, customer OpenRouter credentials (which also cover speech), or disabling built-in AI generation. Customer-key mode is not an arbitrary private-model connector. Separately connected, customer-approved agents can use authenticated Model Context Protocol (MCP) access.

Disabling built-in generation leaves deterministic functionality available; a separately connected agent is a separate processing route. Customer-managed installations support configured PostgreSQL and S3-compatible storage endpoints. A private or fully local solution requires models, agents, identity, telemetry, assets and other outbound dependencies to fit the agreed environment.

5 Identity, access and customer isolation

Who can access our workspace and shared content?

Studio uses individual accounts, workspace memberships and role-based access. Customers control who belongs to their workspace and should remove access when it is no longer needed. Each workspace is logically separated, and Enterprise owners control settings such as AI providers and single sign-on.

Customers can share specific projects using revocable links. Anyone with a link can access the content covered by it, so links should be shared only with intended recipients. Revoking a link prevents future access but cannot remove copies already downloaded or distributed.

Can Bevel AI personnel access customer content?

Authorised Bevel AI personnel may access customer content only when necessary for support, service operation, security incident response or legal obligations. Support access requires customer permission, except where access is required for a security incident or legal obligation, and is limited to the necessary content. For customer-managed Enterprise, access arrangements are agreed with the customer.

Are SSO, MFA and password recovery available?

Enterprise supports OpenID Connect SSO and MFA through the customer's identity provider. Available sign-in methods are agreed during setup to ensure MFA requirements are applied appropriately. For local accounts, passwords are hashed and require at least ten characters. Password recovery uses rate-limited requests and expiring reset links, and successful resets revoke existing sessions. Email recovery is available when configured.

6 Encryption and security engineering

How is customer data protected?

Customer data is encrypted in transit and at rest. Model-provider credentials receive additional application-level encryption and are not exposed through workspace settings. Studio is not end-to-end encrypted because enabled features must process content to provide the service. Enterprise encryption and key-management arrangements depend on the agreed deployment.

How is application security checked?

Bevel AI uses input validation, content integrity and versioning controls, together with automated type, build and test checks. Independent or customer-specific security assessments can be agreed separately. Suspected vulnerabilities can be reported through the main contact.

7 Logging and monitoring

How is activity monitored and investigated?

Studio uses diagnostic logs, request correlation, metrics and traces to monitor operation and investigate issues. Access to diagnostic data is restricted, and credentials and unnecessary customer content are excluded. Project history supports tracing changes to customer content.

For Custom Enterprise, logging destinations, retention and any audit or SIEM export requirements are agreed with the customer.

8 Retention, correction and service exit

How can we manage, export and delete our content?

Studio supports editing, version restoration, export and deletion of projects and Library Documents. Account or workspace deletion and service-exit assistance can be requested through the main contact. Shared workspace content is handled separately from individual user accounts so that removing one user does not remove content belonging to the organisation.

Deleted information may remain temporarily in backups or provider systems until the applicable retention periods expire. Some account, billing, security and support records may be retained where required for operational or legal purposes. Browser data, downloads and distributed copies remain under the customer's control.

Retention periods and privacy rights are described in the **Privacy Policy**. For Custom Enterprise, retention, export, deletion and exit responsibilities are agreed with the customer.

9 Availability, backups and recovery

How are availability and data recovery managed?

Bevel AI manages availability, backups and recovery for Studio SaaS. Backups are designed to cover both database records and stored files, and recovery procedures include restoration testing. For Studio SaaS on Azure, the database can be restored to a point in time within the previous 35 days, and deleted stored files and storage containers can be recovered for 30 days. Project history also allows earlier content versions to be restored, but it is separate from service backups.

Specific availability and recovery commitments are defined in the applicable customer terms. For Custom Enterprise, recovery objectives, backup operation, restoration and continuity responsibilities are agreed with the customer based on the selected deployment.

10 Incident response and support

How can customers contact Bevel AI?

Send privacy requests, service issues, suspected security incidents or vulnerabilities, and concerns about harmful, biased or infringing AI output to info@bevelai.com. The inbox is monitored on Finnish business days.

Include the relevant workspace, approximate time, affected feature or output, and observed impact. Avoid sending sensitive technical evidence in the initial message; an appropriate channel can be agreed when needed. Bevel AI may verify identity or authority before processing privacy, access or deletion requests. Enterprise support and escalation arrangements are defined in the customer agreement.

How are incidents handled?

When Bevel AI becomes aware of a suspected security incident, a responsible team member takes ownership and involves another team member. The response covers assessment, containment, evidence preservation, recovery, customer impact, required notifications and corrective action.

Customers affected by a personal-data breach are notified without undue delay. For customer-managed Enterprise deployments, Bevel AI and the customer coordinate the response according to their agreed responsibilities.

11 Governance and assurance

Who is accountable?

Bevel AI's co-founders are accountable for security, privacy and responsible AI. Bevel AI provides relevant guidance to personnel, reviews access and provider arrangements, assesses material service changes and reviews its policies periodically.

Bevel AI is not currently certified under ISO 27001, SOC 2 or an equivalent independent standard. Certifications held by hosting and AI providers apply only to their respective services. Available evidence or a customer evaluation using synthetic data can be discussed with Bevel AI.

How are customers informed about changes?

Bevel AI maintains the provider information in this document and notifies affected customers of material provider changes according to the applicable customer agreement. Contractual and service changes are communicated under the applicable customer terms.