AI Capabilities in Bria

Bria documentation using the MovieLabs AI Communication Framework v1.0

Document version 0.5 | Vendor: Bria AI | Contact: Doug Bloom, doug@bria.ai

Section numbering. Numbered sections and headings in this document are this document's own. References to the MovieLabs AI Communication Framework are always prefixed, as in Framework 3.3. Where a category or field name appears with a number in parentheses, that number is the Framework section defining it.

Summary

Bria is a visual generative AI platform. Bria trains and owns the foundation models that generate and alter media, and those models are trained on commercially licensed data. The platform is available as a hosted service, as software and model weights deployed in a customer's own environment, and as applications and plug-ins used inside creative pipelines.

What this document is. Vendor-prepared documentation of Bria's AI-enabled capabilities, written in the vocabulary and categories of the MovieLabs AI Communication Framework v1.0. Its purpose is the one set out in Framework 4.2: a reusable package a studio technologist, production team, or service vendor can incorporate into a tool or use case approval request. Questions should be directed to the contact above.

What this document covers. Four of the Framework's five categories: Tool (Framework 3.3), AI Provenance (Framework 3.4), Operational Characteristics (Framework 3.5), and Tool Capability (Framework 3.6). It does not address Use Case (Framework 3.2), which describes the business and production context of a particular proposed use and can only come from the organization proposing it.

Two model layers. Bria's platform uses two categories of model with different provenance, and Section 2 documents them separately. Media models generate or alter images and video; they are trained by Bria on licensed data. Auxiliary language models assist with authoring and interpreting the structured specification that drives the media models; they are derived from third-party open weights, or customers can use their own approved closed-weights model. No auxiliary model generates or alters media. This is provided for in Framework 3.4, which notes that provenance may need to be associated with a particular model where a capability relies on one with distinct characteristics.

Deployment mode is a required qualifier. Several Operational Characteristics fields have different answers depending on how a customer deploys. Custodian, Data Location, Update Rollout Control, and Refinement Location cannot be answered once. Section 3 names the deployment modes first and answers each field per mode, as anticipated by Framework 3.5.

Depth of coverage. Section 4 indexes every AI-enabled capability and classifies each on both Tool Capability dimensions. A subset spanning the full range of Output Novelty is documented in detail; the remainder are classified and described briefly.

1. Tool (Framework 3.3)

1.1 Tool Vendor Profile (Framework 3.3.1)

Tool Vendor Profile	
Vendor	Bria AI. Bria is the contracting entity and the developer, trainer, and owner of the media models described in this document.
Headquarters Location	United States. Incorporated in Delaware, headquartered in New York. Legal jurisdiction is the United States.
Operating Locations	New York, Los Angeles, and Israel. The Israel entity is a subsidiary of the US parent.
Contact	Doug Bloom, doug@bria.ai, for governance, approval, and diligence enquiries.
AI Posture	**Two postures apply, one per model layer.** **Vendor providing its own models, for all media generation and editing.** Bria designs and trains these models and assembles the training corpus under commercial licence. They are not wrapped, resold, routed to, or derived from a third-party model. **Using third-party models, for the auxiliary language layer.** Prompt enhancement, authoring and manipulation of the structured specification, and description of a supplied reference image are handled by language and vision-language models that Bria tunes and hosts, or that Bria routes to where a customer has approved its own. These models do not produce or alter media. See Section 2.2 of this document.

1.2 Tool Packaging (Framework 3.3.2)

Tool Packaging	
Packaging	**One platform, several delivery packages.** Bria Build is the platform: the models together with the pipelines, orchestration, and API surface built around them. The other packages are delivery mechanisms over Build, or the models on their own. **Bria Build.** API-driven service covering image generation and editing and a set of video operations. Also available as an MCP endpoint, SDKs, and embeddable components. **Bria Create.** Browser-based application over Build. **Bria Craft.** Craft Studio, a browser-based application, together with plug-ins for Nuke, Houdini, and Griptape. All run Build beneath the interface. **Source code and model weights.** The models licensed on their own, for the customer to deploy and run in an environment it controls. This package is the models without the surrounding Build pipelines. **Cloud marketplace and managed platform listings.** AWS, Azure, and Microsoft AI Foundry.

Tool Packaging	
Separability	**Customer configurable, enforced by the platform.** Capabilities are exposed as discrete endpoints rather than as a fused feature set, and Bria or an account administrator can restrict which endpoints a given key or user may call. A studio can therefore approve background removal without approving generative capabilities. Data retention is configurable at the same granularity. See Section 3.2 of this document.
Integration Interfaces	REST API, SDKs, MCP endpoint, embeddable components, and plug-ins for Nuke, Houdini, and Griptape. Where Bria is deployed in a customer-controlled environment, integrations are pointed at that installation by changing the endpoint address; no other change is required.

2. AI Provenance (Framework 3.4)

Two layers of model are involved, with different origins and different training data. They are documented separately because the distinction is material. Section 2.1 covers the media models, which generate and alter images and video. Section 2.2 covers the auxiliary language models, which operate only on text and structured specifications and on descriptions of supplied reference images.

The dividing line is that no model with undisclosed training data generates or alters media. Everything that produces or modifies a pixel is a Bria-trained model on licensed data.

2.1 Media models

AI Model Provisioning (Framework 3.4.1)	
Vendor Enabled AI Models	**Fibo model family.** Fibo Generation, Fibo Edit, and Fibo Lite. Trained natively on VGL (Visual Generation Language), a structured JSON specification of visual attributes that is the models' own input and output representation rather than a layer placed in front of them. The VGL schema is public and open. **Task-specific image models.** Background removal, object removal, resolution increase, creative enhancement, background generation, image expansion, and background blur. **Task-specific video models.** Background removal, resolution increase, and object removal. **Customer-tailored models.** Fine-tuned by a customer from a Bria base model. See Section 3.4 of this document. All are designed and trained by Bria. No third-party model generates or alters media anywhere in the platform.
Customer BYO AI Models	No BYO support for media models. A customer's own fine-tuned model is a derivative of a Bria base model rather than a third-party model, and is documented

AI Model Provisioning (Framework 3.4.1)	
	under Customer Model Refinement. BYO is supported for the auxiliary layer only; see Section 2.2 of this document.

Vendor AI Models (Framework 3.4.2)	
Model Origin	Vendor-trained. Trained by Bria from the ground up, not obtained from or derived from open weights or another commercial model. Model weights are licensed to customers, which is a licensing property rather than a statement about origin.
Model Provider	Bria AI. Bria is both the tool vendor and the model provider for this layer.
Model Provider Headquarters	United States, as in Section 1.1 of this document.

Training Data (Framework 3.4.3)	
Source Types	**Commercially sourced under licence, and algorithmically generated.** **Licensed.** The corpus is licensed from rights-holding data partners under agreements permitting training. Bria does not ingest web-crawled, scraped, or unlicensed material. **Algorithmically generated.** Synthetic data that is part of the corpus is generated by Bria models that were themselves trained only on licensed data, so its lineage terminates in the licensed corpus. No synthetic data is derived from a model of unknown provenance.
Named Datasets	The corpus is licensed directly from rights holders rather than assembled from named public research corpora, so the relevant disclosure is the set of licensors. Examples include Getty Images, Alamy, Envato, DepositPhotos, and Freepik. The full partner list is available on request from the contact above.
Source Commitment Duration	**Term-based licences with automatic renewal.** No data partner has exited the program to date. **More directly relevant to a production: the rights secured to train a given model persist with that model.** If a partner were to leave, a model already trained and already licensed to a customer is unaffected, and a customer running that model in its own environment continues to do so. What lapses is the right to extend that data into future training or new models.

2.2 Auxiliary language models

A language model and a vision-language model support the media models. Their functions are prompt enhancement, authoring and manipulation of the VGL specification, and producing a text description of a reference image where one is supplied. They emit text and structured specifications. They do not generate, alter, or output media.

Auxiliary layer	
Function	Prompt enhancement and interpretation; composing and refining VGL specifications; describing a supplied reference image so that its content can be expressed in a specification. **The vision-language model does see customer imagery** where a reference image is supplied for a generation. Its output is a description of that image. No media is produced by this model.
Model Origin	For enterprise deployment Bria uses tuned models derived from open weights, or customer-approved models that have API support.
Model Provider Headquarters	Not applicable in the ordinary sense. For models that are open weights hosted and operated by Bria, there is no runtime dependency on, control by, or data flow to the originating provider. The specific base models are disclosed on request. Where a customer routes the function to its own approved provider, the provider and its jurisdiction are the customer's choice and are governed by the customer's agreement with them.
Training Data	Not disclosed by the model provider. These are third-party open-weights models and Bria makes no representation about the data used to train them. This is the reason the layer is documented separately: the licensed-data position stated in Section 2.1 of this document applies to the media models and does not extend to this layer.
Customer BYO AI Models	**Supported.** A customer may route the language or vision-language function to a provider of its choosing using its own API credentials, so the customer's own contract, security review, and data-handling terms with that provider govern the call. Google Gemini is in use by customers today; open-weights models such as Llama or Gemma are equally workable.
Bypassable	The auxiliary layer can be avoided entirely. The VGL schema is public and open, and direct authoring is already in use by customers. Where a user authors the specification directly and supplies no reference image, the generation path runs only on Bria-trained models. This gives a studio a way to scope approval to paths that do not invoke a model with undisclosed training data. The AI Interaction field in Approval Context (Framework 3.2.2) is where that scoping would be recorded.

3. Operational Characteristics (Framework 3.5)

Four deployment modes are in production use. The fields below are answered per mode where the answers differ.

Mode	Description	Host Environment (Framework 3.5.1)
Bria-Hosted	Operated by Bria. The customer points at Bria's endpoint and authenticates with a token or SSO.	Vendor-controlled network on cloud
Private Cloud	Inference services deployed in the customer's own cloud environment, managed by the customer. Same plug-ins and interfaces, pointed at an internal endpoint.	Customer-controlled network on cloud
Local	A single machine or an on-premises node, offline capable. Includes air-gapped installations.	On-device, or customer-controlled data center
Hybrid	Backend selected per user or by studio policy, with routing between local, private cloud, and Bria-hosted. Used for burst capacity and mixed on-site and remote teams.	Varies by active backend

Applies to all packages except Bria Create, which is available hosted or in private cloud but is not locally deployable. Artist-facing tooling is identical across modes; the only change is the endpoint address.

3.1 AI Deployment (Framework 3.5.1)

AI Deployment	
Deployment Artifact	**Bria-Hosted.** A vendor-controlled tool that includes the AI models. **Private Cloud and Local.** Either the platform with its models deployed into the customer environment, or the models and weights licensed on their own. No call is made to Bria infrastructure at inference time and there is no licensing callback. **One item does return to Bria periodically.** The attribution report described in Section 3.3 of this document. It carries irreversible vectors that indicate which data partners to compensate. It does not carry prompts, specifications, input assets, or output assets, and it says nothing about what was created.
Host Environment	As in the table above. All four modes are in production with customers today, including customers running licensed source and weights.
Deployment Location	**Bria-Hosted.** Deployment to any region is supported. Deployments exist today in the United States and the European Union. **EU AI Act.** Bria is compliant with the EU AI Act. **Private Cloud, Local, Hybrid.** Determined by the customer. Bria has no involvement in where a customer deploys.

AI Deployment	
Update Rollout Control	**Bria-Hosted: controlled by Bria, with customer controls.** A customer can pin a model version so that a production is not moved onto a new checkpoint mid-show. Deprecation notice terms are defined at contract time. **Private Cloud and Local: controlled by the customer.** A deployed version runs until the customer replaces it. Bria cannot push an update into a customer-controlled environment.
Terms and Conditions	IP indemnity applies across all deployment modes where an enterprise agreement is in place.

3.2 Customer Data Residency (Framework 3.5.2)

Customer Data Residency	
Custodian	**Private Cloud and Local: the customer.** Inputs, prompts and specifications, outputs, workflow configuration, and any data used for model refinement remain in the customer environment. **Bria-Hosted: Bria,** for the period set by the customer's retention configuration. **Hybrid: depends on the backend handling the request.**
Data Location	**Private Cloud and Local.** Wherever the customer deploys. **Bria-Hosted.** Generally the United States, configurable to other regions as required. See Deployment Location above.
Data Collection	**Not collected for product improvement in any mode.** Customer content is not used to train or improve Bria's models. Bria's training corpus is licensed data and customer content is not added to it. **Bria-Hosted retention is operational and customer-controlled.** Data is retained for support purposes for 7 days by default. The customer can set this to zero, in which case data is deleted as soon as the task completes and results are returned. The retention period is configurable per endpoint, so a customer can set zero retention on sensitive capabilities and retain the default elsewhere. **Private Cloud and Local: nothing is retained by Bria,** because no customer content reaches Bria. See Metadata Collection below for what does.
Collection Data Types	Not applicable for non-operational collection. For operational retention under the default setting, the data held is that required to support the request.
Collection Retention Period	Bria-Hosted: 7 days by default, configurable to zero, set per endpoint. Private Cloud and Local: not applicable.
Vendor Use of Collected Data	Customer data is not used for model training, for tool improvement, or for the benefit of any other customer, in any deployment mode.

3.3 Metadata Collection (Framework 3.5.3)

Metadata Collection	
Metadata Collection	**Attribution reporting is required in all deployment modes, including customer-controlled and air-gapped ones.** It is a condition of the commercial licence. This is the only information that leaves a customer-controlled environment, and it is described in full below rather than summarized, because a reviewer evaluating an isolated deployment needs to know precisely what crosses the boundary. **What is reported.** Irreversible vectors computed from generated outputs. Their sole function is to identify which data partners' licensed material contributed to an output so that those partners can be compensated. Prompts, specifications, input assets, and output assets are not included and are not transmitted. The vectors cannot be used to reconstruct the content they were derived from. **How it is reported.** An Attribution Agent runs wherever inference occurs. In customer-controlled deployments the report can be sent to Bria monthly; an open network port is not required. **Why it is not optional.** Attribution is the mechanism by which Bria compensates the rights holders whose material it licenses, under revenue-share agreements between Bria and its data partners. It is what the licensed-data position and the indemnity rest on, and it is what allows per-output provenance to be produced. A customer cannot disable it under the commercial licence. **It is not usage metering.** The report determines how Bria allocates payment among its data partners out of its own revenue share. It does not price the customer, and generation volume does not change what a customer pays.
Collection Metadata Types	Irreversible attribution vectors. Optionally, customer-defined sub-identifiers for project, department, or budget tracking, where the customer chooses to use them for its own internal allocation.
Collection Retention Period	Attribution reports are retained per the terms agreed with each data partner, since their purpose is to substantiate compensation under those agreements.
Vendor Use of Collected Metadata	Compensation of data partners. Not used for model training, product improvement, or any customer-specific purpose.
Usage metadata	**Bria-Hosted.** Endpoint usage and volume are collected at organizational level only, not per individual user. This is the same data required to bill for use of the hosted service, so it is collected for as long as the hosted service is used. The dashboard control governs access to the data rather than whether it is collected. **Private Cloud and Local.** No telemetry path other than the attribution report described above.
Related: generation metadata held by the customer	The creative-application integrations write a generation record for each operation, comprising the model and operation invoked, the prompt or full VGL specification, the seed, generation parameters, and output details. It is stored as a sidecar file or embedded in the host application's project file.

Metadata Collection	
	This record stays with the customer and Bria does not receive it. It supports the Workflow Provenance Tracking field in Approval Context (Framework 3.2.2), which a proposer completes.

3.4 Customer Model Refinement (Framework 3.5.4)

Customer Model Refinement	
Refinement Allowed	**Yes, on the Fibo model family.** Customers can fine-tune a Fibo base model on their own data, and the resulting model is the customer's. Because model weights are licensed to the customer, a refined model can be held, inspected, and deployed without dependence on Bria's service. **Where training happens.** Training is performed in Bria Build. A refined model can then be used across the other products. **Methods and scale.** Options range from LoRA adaptation, which can work from as few as five images, through fine-tuning and deep fine-tuning on hundreds or thousands of images. The appropriate method depends on the dataset available and on what the model is being taught.
Refinement Location	**Customer-controlled data center and customer-controlled network on cloud.** A customer with licensed weights trains entirely in its own environment, including air-gapped. Training data does not leave the customer's control. **Vendor-controlled network on cloud.** Hosted training is also offered. Training data is sequestered in storage dedicated to that customer and is not pooled with any other customer's data.
Vendor Access	**Customer-side refinement: no.** Bria has no access to a model trained in the customer's environment or to the data used to train it. **Hosted refinement: not by default.** Training data sits in customer-dedicated storage. Bria can access it only with the customer's permission, and only for support or professional services work the customer has asked for.
Vendor Usage	No, in all modes. Bria does not use customer-refined models, or the data used to produce them, for its own purposes or for any other customer.

4. Tool Capability (Framework 3.6)

The Framework classifies each capability along two dimensions rather than by model architecture. Output Novelty (Framework 3.6.1) asks what the capability does to media. Creative Controllability (Framework 3.6.2) asks how much of the transformation recipe is fixed by the tool and how much the user defines.

A note on Creative Controllability

For most tools, Creative Controllability is a fixed property of a capability: the tool exposes a set of controls and the user works within them. For Bria's Fibo capabilities it is a range, and the position within that range is set by how the user chooses to work rather than by what the capability permits.

The reason is that the VGL specification is the models' native input. A user may send a short natural-language prompt, in which case the system composes the specification and the interaction resembles Preset Automation. The same user may instead author the specification directly, declaring object positions, colours, camera optics, lighting, and composition, and compose those specifications across layers and regions in a node-based application. That is Composable Workflow, on the same capability and the same model.

The practical consequence for governance is that the Creative Controllability of a Bria generative capability cannot be determined from the capability name, so a guideline written against the capability alone will not be precise. The Framework already provides the right instrument in the AI Interaction field of Approval Context (Framework 3.2.2), which asks how the Primary User intends to use the controls a capability exposes. For these capabilities that field carries more of the governance weight than the classification does, and the classifications below are stated as ranges for that reason.

4.1 AI-Enabled Capabilities at a Glance

Capability	Output Novelty (Framework 3.6.1)	Creative Controllability (Framework 3.6.2)	Model
VGL extraction from image	Information Extraction / Modality Conversion	Preset Automation	Auxiliary VLM
Background removal	Simple Transformation	Preset Automation	Task-specific
Upscale / resolution increase	Simple Transformation	Preset Automation	Task-specific
Creative enhancer	Correction	Preset to Parameterized Automation	Task-specific
Background blur	Simple Transformation	Preset Automation	Task-specific
Object removal (eraser)	Correction	Conditioned Transformation	Task-specific
Image expansion (outpainting)	Correction to Creative Modification	Conditioned Transformation	Task-specific
Background generation / replacement	Creative Modification	Conditioned Transformation	Task-specific

Capability	Output Novelty (Framework 3.6.1)	Creative Controllability (Framework 3.6.2)	Model
Fibo Edit	Correction to Creative Modification	Conditioned Transformation to Composable Workflow	Fibo Edit
Fibo Generation	Guided Creation	Preset Automation to Composable Workflow	Fibo Generation / Lite
Generation from a customer-tailored model	Guided Creation	Preset Automation to Composable Workflow	Customer-refined
Video background removal	Simple Transformation	Preset Automation	Task-specific
Video upscale	Simple Transformation	Preset Automation	Task-specific
Video eraser	Correction	Conditioned Transformation	Task-specific

Generation is images only. The video capabilities are the three listed above. Their models are trained by Bria on licensed data on the same basis as the image models, and they compute attribution in the same way.

Fine-tuning is not listed here. Model refinement is a training process rather than a generation-time transformation of media and does not classify on either dimension. It is documented in Section 3.4 of this document. Where a production proposes to fine-tune, the Framework directs that to the Approval Subject field in Approval Context (Framework 3.2.2), which provides for AI Fine-Tuning Workflow as a distinct subject of approval.

4.2 Fibo Generation

Generation of a new image from a VGL specification, from a natural-language prompt, or from both, optionally with reference material. Fibo Lite is a lower-latency variant taking the same specification without modification.

Output Novelty: Guided Creation. Substantially new media is produced and many valid outcomes exist for similar inputs. This holds regardless of how tightly the user constrains the specification. A detailed specification narrows the range of outcomes but the output remains newly created media rather than a transformation of an existing asset.

Creative Controllability: Preset Automation to Composable Workflow.

- **Preset Automation.** A short natural-language prompt. The system composes the specification and the user does not adjust it.
- **Parameterized Automation.** Prompt plus exposed parameters such as seed, aspect ratio, and guidance settings.
- **Conditioned Transformation.** Generation conditioned on supplied guidance material, or driven by a partially authored specification.
- **Composable Workflow.** The user authors the full specification and assembles generation into a node graph, targeting regions and layers and chaining operations.

Inputs and outputs. Input is a VGL specification, a prompt, or both, optionally with reference imagery. Output is a generated image together with the specification that produced it. Where a reference image is supplied, the auxiliary vision-language model described in Section 2.2 of this document is invoked to describe it.

Human control. The specification is a human-readable artifact that can be reviewed, edited field by field, stored alongside the deliverable, and re-executed. The same specification and seed reproduce the same output.

Model. Fibo Generation or Fibo Lite. Specifications are portable between them without translation.

4.3 Fibo Edit

Modification of an existing image, guided by a specification, a natural-language instruction, or supplied masks and reference material.

Output Novelty: Correction to Creative Modification. This capability spans two classes and should not be recorded with a single label.

- **Correction** where limited new elements locally repair or complete an existing asset, for example removing an object and reconstructing the area behind it.
- **Creative Modification** where new elements materially change the image while the progression from the original remains discoverable, for example replacing a background, altering signage, or changing the time of day.

Which class applies is a property of the operation, not of the capability. A production documenting actual use under Framework 4.4 should record the class for the operation performed. A guideline permitting Correction but not Creative Modification can be expressed precisely against this capability provided the operation is specified.

Creative Controllability: Conditioned Transformation to Composable Workflow. A masked edit driven by an instruction is Conditioned Transformation. Working inside a node-based or layer-based application, targeting regions and layers in a chosen order and editing the specification between steps, is Composable Workflow.

Model. Fibo Edit.

4.4 VGL extraction from an image

Analysis of an existing image to produce a structured description of it: the objects present and their positions, the lighting, camera characteristics, palette, and composition.

Output Novelty: Information Extraction / Modality Conversion. The capability reads what the image already contains and converts it into a structured representation. It produces no media and alters nothing. This is the lowest-novelty class on the Framework's dimension.

Creative Controllability: Preset Automation. The user supplies an image and receives a description. No parameters shape the analysis.

Model. This capability is performed by the auxiliary vision-language model described in Section 2.2 of this document, not by a Fibo model. It is the one capability in this document in which customer imagery is processed by a model derived from third-party open weights. Where a customer routes the vision-language function to an external provider under its own credentials, the reference image is processed by that provider under the customer's agreement with them.

Inputs and outputs. Input is an image. Output is a VGL specification in JSON. No media asset is produced and the input is not modified.

4.5 Remaining capabilities

Classified in the table above and described briefly here. Each uses a Bria-trained task-specific model.

- **Background removal.** Simple Transformation, Preset Automation. Foreground is separated from background and returned with transparency. Nothing is invented; the result is a subset of the input. Available for images and video.
- **Upscale and resolution increase.** Simple Transformation, Preset Automation. Resolution is increased without introducing detail that was not present in the source. Available for images and video.
- **Creative enhancer.** Correction, Preset to Parameterized Automation. A separate capability from upscale, and classified differently because it does introduce detail that was not in the source. This is the distinction the Framework draws in its own Correction example, and the two are kept separate here for that reason.
- **Background blur.** Simple Transformation, Preset Automation. A depth-aware treatment of existing pixels.
- **Object removal (eraser).** Correction, Conditioned Transformation. Removes a designated object and reconstructs the occluded region. The user designates the region; the reconstruction is generated. Available for images and video.
- **Image expansion (outpainting).** Correction to Creative Modification, Conditioned Transformation. Extending a frame to complete a composition is Correction. Extending it far enough that substantial new scene content appears is Creative Modification. The class should be recorded against the operation performed.
- **Background generation and replacement.** Creative Modification, Conditioned Transformation. The foreground subject is preserved unaltered while a new background is generated around it.
- **Generation from a customer-tailored model.** Guided Creation, with the same controllability range as Fibo Generation. Refinement does not change the classification: a tailored model still produces substantially new media. What changes is that the space of likely outputs is constrained toward the customer's own material. Provenance for a tailored model combines Bria's base-model provenance in Section 2.1 of this document with the customer's own training data, and the latter is the customer's to document.