The MovieLabs AI Communication Framework

Version: 1.0

October 13, 2026

Legal Information

This documentation is intended as a guide for companies developing or implementing products, solutions, or services for the future of media creation. No effort is made by Motion Picture Laboratories, Inc. to obligate any market participant to adhere to the recommendations contained in this documentation. Whether to adopt these recommendations in whole or in part is left to the discretion of individual market participants, using independent business judgment. Each MovieLabs member company shall decide independently the extent to which it will utilize, or require adherence to, these recommendations. All questions on member company adoption or implementation must be directed independently to each member company.

© 2026 Motion Picture Laboratories, Inc.

This documentation is licensed under the Creative Commons Attribution 4.0 International License.

Creative Commons Attribution 4.0 International License (CC BY 4.0) – Summary

You are generally free to copy and redistribute the material in any medium or format, and remix, transform, and build upon the material for any purpose, even commercially, under the following conditions:

- **Attribution:** You must give appropriate credit to MovieLabs, provide a link to the license, and indicate if you made changes to the material. You may do so in any reasonable manner, but not in any way that suggests that MovieLabs endorses you or your use of the materials.
- **No Additional Restrictions:** You may not apply legal terms or technological measures that legally restrict others from doing anything the license permits.

Note: The above text is a human-readable summary of (and not a substitute for) the license, which can be found at: creativecommons.org/licenses/by/4.0/.

Table of Contents

1. Introduction 4
   - 1.1 Motivation 4
   - 1.2 Scope and Purpose 5
   - 1.3 Potential Future Work 7
   - 1.4 Notice and Disclaimers 7
2. Glossary 8
3. The Framework 10
   - 3.1 Overview of the Framework 10
   - 3.2 Use Case 12
   - 3.3 Tool 15
   - 3.4 AI Provenance 16
   - 3.5 Operational Characteristics 18
   - 3.6 Tool Capability 20
4. Framework in Action 23
   - 4.1 Using the Framework in Approval Flows 23
   - 4.2 Using the Framework for Tool Vendor and Tool Documentation 26
   - 4.3 Using the Framework to Communicate Studio Guidelines 27
   - 4.4 Using the Framework to Document AI Use in Production 28
5. Conclusion 30

1. Introduction

1.1 Motivation

Artificial intelligence (AI) has become an increasingly important part of modern media creation efforts. AI is being used in a growing number of production tools and workflows that support creative, technical, and operational work across the lifecycle of a project. AI-enabled capabilities are helping productions accelerate localization, improve restoration and clean-up of media plates, assist with media tagging and analysis, expand previz and look development options, and offer production alternatives improving talent and crew safety. Overall, AI has demonstrated the potential to help artists and production teams tell stories in settings, at scales, or under production constraints that would otherwise be difficult, cost-prohibitive, or in some cases not feasible.

At the same time, AI introduces a wide range of questions that make clear communication between several stakeholders essential and non-negotiable. Questions may arise about the provenance of AI models and training data, whether model outputs may increase copyright or infringement risk, whether the use of an AI capability affects talent or labor considerations, how the system behaves operationally, what data is collected or retained, and what role AI is actually playing in the transformation from inputs to outputs. Additionally, the phrase *generative AI* is not usually sufficient for governance purposes. In some circumstances the distinction between generative and non-generative AI may be clear. But in most cases, the boundary is less obvious, especially as generative models perform non-generative tasks at a superior level than their traditional counterparts. And different regulatory bodies may end up taking different views on what generative AI is.

Risk tolerance, business objectives, production context, applicable agreements, and governance structures differ across studios, production companies, ad agencies, and even across use cases within a given organization. What is needed, therefore, is not a universal policy, but a common way for those involved in decision making, production technologists, servicing vendors, and those providing AI solutions to communicate the aspects of AI solutions and AI use that matter most to evaluating risk and benefit in media creation.

This document presents such a framework. Its purpose is to identify and bring to focus the major communication categories that repeatedly arise when stakeholders discuss AI-enabled tools, capabilities, and their use in media creation. By defining and elaborating these categories and the associated terminology with greater precision, the framework aims to reduce ambiguity, improve comparability, and help stakeholders get into alignment faster.

A common communication framework is valuable not only because it identifies the emerging key aspects of AI solutions and AI use that need attention, but also because it allows different stakeholders to speak the same language now and a foundation to expand on the framework in the future as we move past the nascent and rapid experimentation phase and into the mature

adoption and adaptation phase. At the moment, tool vendors, studios, production teams, and service vendors often deliberate on overlapping issues, but through different organizational lenses and describe them using different terminology. That fragmentation can slow AI solution approvals, complicate studio-vendor discussions, and make internal coordination more difficult. A shared framework improves communication across organizations and also within them, including between technologists, creative teams, legal teams, security teams, production leadership, and other governance participants.

One especially important benefit is that tool vendors can use the framework to document upfront their AI-enabled solutions and capabilities as part of their standard product offer in a way that can immediately support studio review. When a studio technologist, production team member, or a creative user proposes an AI tool for use in a production, the approval process benefits from the head start. Documentation prepared by tool vendors in the language of the framework can serve as a reusable starting point for evaluation, reducing clarification cycles and providing a more informed basis for governance decisions.

The framework as we designed it is foundational, not the endpoint, and it is devised to evolve as AI capabilities emerge and our understanding of AI use develops.

One key disclaimer is that this framework was neither designed to cover every aspect of communication necessary to support governance and policy making around AI use in productions nor to provide any policy guidance. The focus strictly remained on surfacing only technical and operational aspects that would be relevant for governance at individual organizations while keeping legal, ethical, labor, and commercial considerations out of it. Sections 1.2, 1.3, and 1.4 discusses what is in scope, what could be considered in future versions, and what persistently remains outside of scope.

Here is how this document is structured. The remainder of Chapter 1 discusses the scope and purpose of the framework specifically the initial version, identifies areas where the framework could be extended in future, and clarifies what remains permanently outside the scope. Chapter 2 provides a short glossary to establish a common baseline to interpret subsequent chapters. Chapter 3 presents the framework in detail, including necessary definitions of terms that are introduced. Chapter 4 illustrates how we anticipate the framework to be used in representative communication scenarios. Chapter 5 concludes the document.

1.2 Scope and Purpose

The MovieLabs AI Communication Framework (Framework) supports clear and consistent communication about AI-enabled tools, tool capabilities, and AI use in media creation workflows. It helps stakeholders document and discuss the information that is most relevant when evaluating, governing, approving, constraining, or documenting the use of AI in production-related contexts.

The Framework does not attempt to capture every possible fact about a tool or workflow. Instead, it focuses on the communication aspects that are most salient to the evaluation of risk and benefit of AI use in media creation. The emphasis is on practical and technical clarity: what stakeholders need to understand about an AI solution, its capabilities, its provenance, the conditions under which it operates, and the context in which it is proposed to be used.

The Framework is intentionally designed to be neutral so that it can be used by multiple stakeholder groups:

- **Governance and policy providers** at studios, ad agencies, and other media organizations may use it to express AI approval requirements, perform evaluations, and provide AI usage guidance and constraints.
- **Production technologists and creative users** may use it to express proposed AI use, business opportunities, and stakeholders who would benefit from such use.
- **AI solution providers and Tool vendors** may use it to document their AI tools and capabilities in a reusable and comparable way.
- **Service vendors, production teams, and other tool users** may use it to document actual AI use to demonstrate alignment with studio guidelines.

The Framework is also intentionally designed to be modular and is organized into five major categories, each broken down into sub-categories and fields to enable effective communication between stakeholders. It is not a monolith that must always be used in full. Different stakeholders may use different portions of it depending on the communication need. A tool vendor may use certain categories from the Framework to document their AI solutions. A production technologist, relying on the vendor documentation, may use most categories to communicate the use case and surface relevant technical aspects of the tool or capability being proposed for use. A studio governance or policy body, when weighing the proposed benefits against risks, may want to pull into focus only certain aspects of the received documentation and the Framework helps them with sifting through the documentation to locate the relevant information quickly.

One important clarification: Use case and tool approval processes at studios, ad agencies, and other media organizations can take many different forms. The rigor applied may differ across organizations, across business scenarios, and across production contexts within the same organization. Some use cases may require emphasis on legal or IP considerations, while others may require focus on privacy and copyrightability. The Framework does not replace those processes. Rather, it provides a precise and common way to communicate the information such processes depend on.

In that sense, the Framework is best understood as a communication foundation for governance. It helps stakeholders describe AI solutions and proposed uses with enough precision that each organization can make decisions that align with its own governance, policies, and business judgment.

1.3 Potential Future Work

This document reflects an initial version of the Framework and meant to be foundational work for future extensions. It intentionally limits its scope so that the core communication needs at the current moment of AI adoption may be addressed first. Several adjacent topics, as well as future needs that arise as AI evolves and our understanding of AI use develops, may warrant additional work, whether as extensions of this Framework or as separate but related efforts.

Potential future work includes:

- **AI security and safeguards.** A future effort may address communication needs around AI-specific vulnerabilities, security posture, and operational safeguards.
- **Assessment questionnaires.** The Framework already provides the basis for studios and media creation outlets to create their questionnaires for tool vendors. A future effort may document the common set of questionnaires applicable across such organizations, while remaining separate from any individual organization's approval process.
- **Controlled vocabulary expansion.** This version of the Framework provides communication terms and definitions with examples. A later version may define a more detailed controlled vocabulary for select terms where additional precision or interoperability would be valuable.
- **Agentic AI and agentic workflow behavior.** Future versions may address governance needs, if any, around workflows in which planning, orchestration, tool and capability selection, or multi-step task execution is performed dynamically by AI-driven systems.
- **Worked examples and sample documentation patterns.** A few examples of how the Framework can be used to document a use case or AI tool are included. Additional examples could be developed documenting a wider range of tools, use cases, and guideline scenarios.

1.4 Notice and Disclaimers

This document and the Framework are not intended to establish policies, prescribe values, or assess the impact of any policy. Terms, language, and choices related to ethical, legal, labor, business, and other governance considerations that guide policy and compliance remain outside the scope and are for individual organizations to determine.

This document is intended solely as a guide for organizations with needs to communicate about the use of AI tools in media creation, providing a menu of items from which each may choose to use. Each organization shall decide independently which parts of this framework, if any, to utilize. All questions on the actual use of the contents of this document must be independently directed to organizations using this document. MovieLabs shall not be liable for any exemplary, incidental, proximate or consequential damages or expenses arising from the use of this document. This document defines only one approach to characterizing AI solutions and their use, and other approaches may be available.

Any examples included in this document and the Framework are illustrative only. They are intended to show how the Framework may be used in practice. They should not be interpreted as endorsing particular tools or capabilities, prescribing specific policies, or suggesting that any specific use is acceptable or compliant.

The Framework also does not attempt to resolve the risks associated with AI. Rather, it aims to improve the communication of the facts and the context that organizations may need in order to conduct those evaluations.

2. Glossary

This chapter intentionally defines only the terms *Artificial Intelligence (AI)*, *AI Model*, *AI Tool*, and *Tool Capability* as these are concepts useful to interpret the rest of the document. **The remaining terms used by this document are introduced and defined within Chapter 3 as part of the Framework itself**, where they are presented in the context of the categories and fields to which they belong.

There are varying industry, standards body, and government definitions of Generative AI (GAI). Some of those definitions refer to GAI as a subset of artificial intelligence that learns patterns from data and produces content based on those patterns. That framing of GAI, while may be suitable for certain needs, does not capture the nuance around what content is produced, what patterns are utilized, and what aspects of the inputs supplied during AI use are utilized for that content production. This kind of nuance is necessary for governance. Therefore, this Framework avoids defining GAI, but provides two different ways to frame GAI, as you can see in the Tool Capability section in Chapter 3.

While the Framework does not currently define the GAI term, individual studios and other organizations can of course choose to do so.

For purposes of this document, we defined Artificial Intelligence (AI) in this Framework because the Framework considers several flavors of AI including GAI. The definition for AI below casts that wide net and is not a definition from any known organization.

Artificial Intelligence (AI)

> *Computational system that produces outputs from inputs using methods not fully specified by fixed logic or rules, but instead learned or optimized from examples, feedback, or reward signals during a training or adaptation phase. This includes "traditional" machine learning and reinforcement learning based systems as well as deep learning architecture-based systems such as transformer and diffusion systems.*

AI Model

A specific instance of AI, typically deployed on one or more machines.

AI Tool or AI Solution

A product, service, platform, plug-in, API, or workflow environment made available for use that leverages, embeds, or is integrated with an AI Model or AI Models to provide one or more functions or features. If an AI Model is made available directly for use, this Framework treats that AI Model itself as an AI Tool. The term AI Solution is used synonymously with AI Tool in this document.

Tool Capability

A distinct feature within an AI Tool such as upscaling, denoising, inpainting, or text-to-video generation.

From a tool vendor perspective, an AI Tool is typically the broader product or service being offered to customers, and that product may expose one or more distinct features such as upscaling, denoising, inpainting, voice cloning, or text-to-video generation. In this Framework, those specific AI features are treated as Tool Capabilities.

From a studio, production, or service vendor perspective, however, evaluation and approval often center not on the whole product in the abstract but on the particular feature or function being tested, approved, constrained, or used in a workflow. In practice, stakeholders may therefore refer to that specific feature itself as the "tool". This difference in usage can create confusion when tool vendors and tool users are discussing the same product at different levels of granularity.

This Framework takes the position that the broader product, service, platform, or directly usable model is the AI Tool, while the specific functions within it are Tool Capabilities. This distinction is important because it allows tool vendor documentation to describe a product at the tool level, while still allowing studios and production teams to focus on the particular capabilities that are actually being evaluated or used. The distinction also helps clarify that governance, approval, and workflow documentation may apply differently to different capabilities within the same tool. For example, a single AI Tool may include both lower-novelty capabilities such as denoising or transcription and higher-novelty capabilities such as object generation or voice cloning, and those may warrant different treatment even though they are part of the same tool from a tool vendor perspective.

For clarity and consistency, this document therefore uses AI Tool to refer to the broader thing made available for use, and Tool Capability to refer to the specific feature or function within that tool. Tool Capability should therefore be treated as a separate defined term in this document.

3. The Framework

3.1 Overview of the Framework

The Framework is designed to help stakeholders communicate effectively by surfacing the most important technical and operational elements about proposed AI use as well as about the AI tools and tool capabilities that support such use. The Framework is also designed to help communicate guidelines for productions that use AI as well as document the actual use of AI within productions. The objective is to provide a practical communication guide that can be used repeatedly in discussions between studios, tool vendors, service vendors, and production teams.

The Framework is organized into five major categories and designed to be modularly used to fit specific stakeholder communication needs:

1. **Use Case** category helps communicate the business and workflow context in which the AI use is proposed and tracked.
2. **Tool** category helps communicate details about the AI solution provider (tool vendor) and how the AI tool or AI model is packaged.
3. **AI Provenance** category helps communicate how the AI is trained and enabled in general at the AI tool level or granularly at the capability or even at the model level.
4. **Operational Characteristics** category helps communicate the conditions around AI deployment, data residency, data collection, and AI fine-tuning support.
5. **Tool Capability** category helps communicate how, and how much, a capability may infuse novelty and creative elements into produced media (instead of industry-agnostic or AI model architecture based framing).

Taken together, these categories help studios communicate AI use guidelines to production teams and servicing vendors, and in turn such users to document their AI use to demonstrate alignment with guidelines.

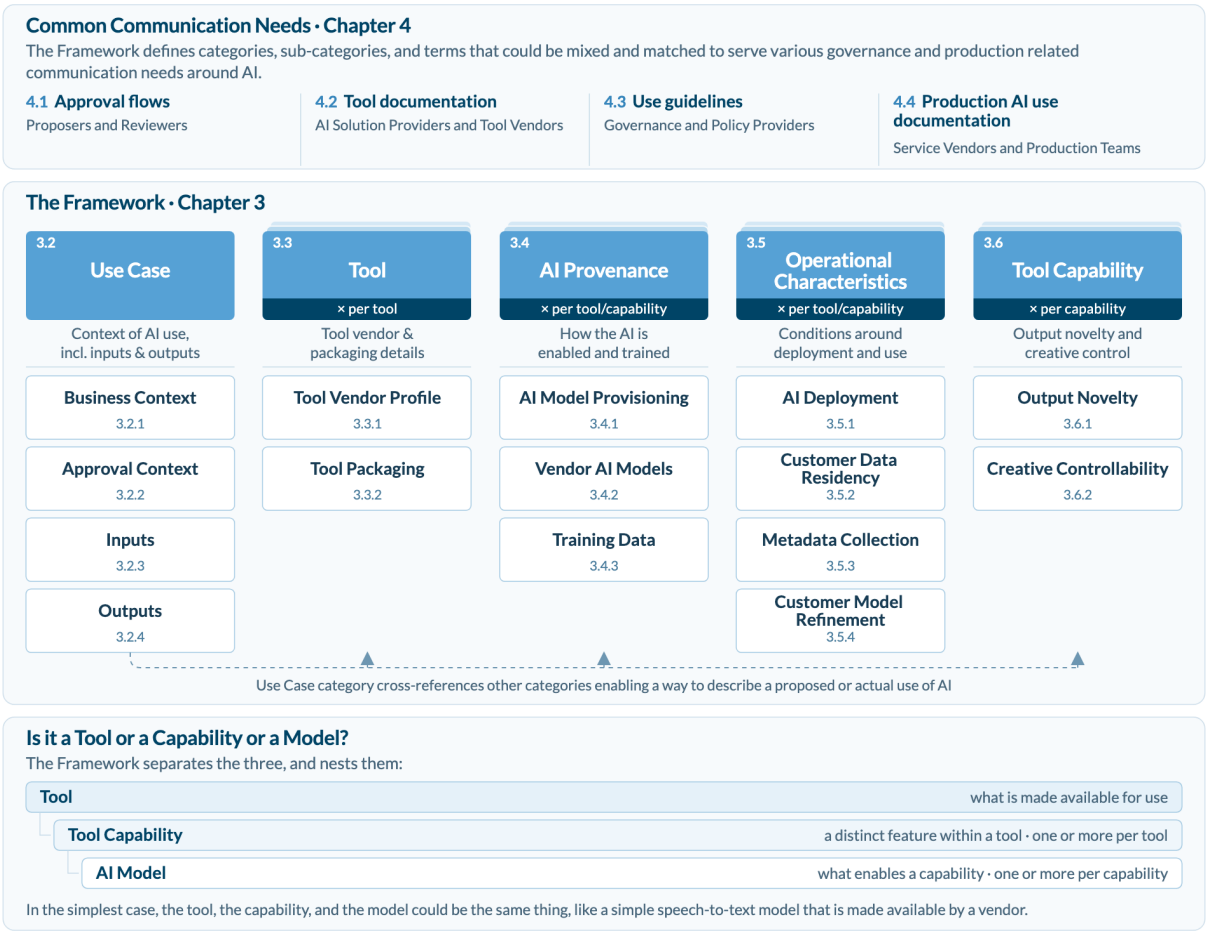


Figure 1: The MovieLabs AI Communication Framework

The Framework diagram presents these categories and sub-categories side by side to show that different stakeholders may enter the discussion from different points. Tool vendors may begin by documenting the tool, its AI provenance, and its operational characteristics. Studios may begin from a use case or tool vendor approval need. Service vendors may focus on documenting a workflow and the particular capabilities in scope.

The categories are also not rigidly tied to one level of granularity. Some information may apply at the tool level, some at the level of an individual capability, and some at the level of a model used to enable that capability. For example, AI Provenance may need to be documented at the tool level in a general product document, but at the capability or model level when a particular feature relies on a distinct model with distinct training data or provenance characteristics. Operational Characteristics may similarly apply differently across a whole tool, a deployment mode, or a specific AI-enabled capability.

What follows describes each category in turn, starting with the Use Case category.

3.2 Use Case

The Use Case category captures the context in which AI use is being proposed. It is not focused on the tool in the abstract, but on the intended use of AI in a particular business and workflow setting. This category is especially important because governance decisions are often use-dependent. The same capability may be treated differently depending on the project visibility, production function, input assets, output role, and intended benefit.

Some of the sub-categories from Use Case may also be used to communicate production use guidelines and in turn document the actual use.

Use Case is divided into four sub-categories: Business Context, Approval Context, Inputs, and Outputs.

3.2.1 Business Context

Business Context describes the organizational and production setting in which the AI use would occur. It helps situate the proposed use within the practicality that shape governance decisions.

Fields in this subcategory include:

- **Project Visibility**: the visibility of the project or activity in which the approved AI use would occur. Examples are internal testing or public visibility.
- **Release and Legal Status of Media Content**: the release and legal status of the media content involved in the use in either inputs or outputs. Examples are public-domain material, pre-release concept art image with copyrighted character, storyboard image with trademarked product, or audio with licensed and likeness-protected voice. Sub-categories Inputs and Outputs in Section 3.2.3 and 3.2.4 can be used to be more precise about the release and legal status.
- **Commercial Nature**: the commercial nature of the project or activity. Examples are commercial release, other commercial use (for example marketing or promotion), or non-commercial use such as demonstration.
- **Business Benefit**: the organizational or production value sought from the use. Examples are accelerating localization, improving restoration quality, supporting creative exploration, improving talent or crew safety, or evaluating tool fit.
- **Use Case Maturity**: how mature or speculative the proposed use is. Examples are engineering effort, short-term research, or long-term research.
- **Stakeholder Function**: the production function or department in which the approved AI use applies. Examples are previz, VFX, sound, localization, archive, or marketing.
- **Primary User**: the profile of the primary user of the approved AI. Examples are a servicing vendor, a technologist within a department, or a creative artist.

This subcategory matters because the center of an approval decision is often an evaluation of the anticipated benefit relative to the relevant risks, and those risks are highly dependent on the business context in which the AI use is proposed.

3.2.2 Approval Context

Approval Context describes what is being put forward for approval and what bounded use is being requested. Parts of it may be used to communicate studio guidelines for production teams and servicing vendors.

Fields in this subcategory include:

- **Approval Subject**: what is being submitted for approval. Examples are AI Tool, AI Model, AI Workflow, or AI Fine-Tuning Workflow. If the intended use is AI Workflow, this field includes details about the workflow stages. If the intended use is AI Fine-Tuning Workflow, this field includes details about the fine-tuning approach as well as how the fine-tuned AI is intended to be used.
- **Proposed Use**: the bounded set of uses being requested. Examples are fixed-frame restoration only, metadata tagging only, voice cloning for narrowly defined circumstances, or internal previz only.
- **Output Reviews**: the reviews that may be needed for output assets created using the approved AI or fine-tuned AI. Examples are technical review only, creative review in rare cases, creative review required, or compliance check.
- **Tool(s) in Scope**: the tool or tools brought inside the approval request scope. The Tool category described in section 3.3 can be used to express the specifics of the tools.
- **Tool Capabilities in Scope**: the capability or capabilities brought inside the approval request scope. Examples are upscaling, voice cloning, inpainting, script breakdown, or full tool. Capabilities can be expressed using the Tool Capability category described in section 3.6.
- **AI in Scope**: the AI Models and their provenance relevant for processing the approval. The AI Provenance category described in section 3.4 can be used to express the specifics.
- **AI Interaction**: how the Primary User (3.2.1) intends to interface with AI directly. The Creative Controllability subcategory described in section 3.6.2 characterizes the controls that a tool capability exposes. AI Interaction, here, describes how the user intends to use those controls. One example is “the video generation capability offers conditional inputs; we intend to guide the capability using text prompt, depth map, start frame, end frame, and camera poses expressed using DCT-based tokens, and iterate to derive precise video rendering and camera movement”.
- **Workflow Provenance Tracking**: describes the measures that the Primary User (3.2.1) intends to take to record the workflow and asset provenance, where such provenance will be stored (embedded in assets, captured in specific production facility, or both), how sensitive aspects of the workflow such as critical story lines are managed, and if cryptographic measures will be used for ensuring records are tamper-evident.

Provenance may include prompts, seeds, nature and degree of human contributions, AI tools and capabilities used, workflow configurations, and inputs and outputs from the workflow.

Note that if the Approval Subject is AI Fine-Tuning Workflow, then the rest of the fields in this subcategory as well as the fields in the Inputs and Output sub-categories refer to the conditions around the use of AI after it is fine-tuned.

Overall, this subcategory recognizes that approvals may differ across studios and scenarios, and that many groups may be involved, including technologists, production leadership, legal, compliance, and creative stakeholders. This subcategory helps annotate precisely what is being requested so the relevant stakeholders can evaluate it.

3.2.3 Inputs

Inputs describes the types of media content provided to the approved AI or fine-tuned AI.

Fields in this subcategory include:

- **Input Types**: the kinds of media content or material that enters the workflow. Examples are storyboard, concept art, script, audio dialogue, music, on-set plate, 3D, metadata, embeddings, or point cloud data.
- **Release Status**: the release status of content if in-house. Examples are pre-release, released, or current production.
- **Legal Status**: the legal position associated with the input content or elements of the input content. Examples are copyrighted, trademarked, patented, confidential, likeness-protected, or licensed.
- **Legal Subject**: the specific element of the input content that has legal status. Examples are product, character, voice, or image.

Input characterization is often essential because the risk profile of a use case depends heavily on what material enters the workflow. The same tool capability may be treated very differently depending on whether it is operating with public-domain material, pre-release concept art image with copyrighted character, storyboard image with trademarked product, or audio with licensed and likeness-protected voice.

Production guidelines can also use this subcategory to communicate how AI is expected to be used while the workflows can document the inputs they supply to AI.

3.2.4 Outputs

Outputs describes the media content that is produced by the approved AI or fine-tuned AI and the downstream roles such content may play.

Fields in this subcategory include:

- **Output Types**: the kinds of media content that is produced. Examples are concept art, storyboard, background plate, CG asset, synthetic voice track, proxy render, metadata, or embeddings.
- **Output Novelty**: whether the output is a mere extraction of information from inputs or includes newly created elements or media. Output Novelty can be expressed using these classes: information extraction, modality conversion, simple transformation, correction, creative modification, or guided creation. See Tool Capability category for more information on Output Novelty.
- **Intended Output Role**: the role the output is expected to play when produced. Examples are non-media assets (like a report or summary, script breakdown, metadata tags) or working / WIP media asset.
- **Final Use Potential**: whether the output is likely to remain internal and supportive or may cross into final media. Examples are none, layered with other assets for final use, or direct final use (with minor transformations).
- **Talent Likeness**: how talent voice, image, gait, aesthetics, or tag lines are manifested in outputs. Examples are none, voice, aesthetics, or performance-related manifestations.
- **Brand Perception**: how a brand is manifested in the outputs, where relevant. Examples are none, original brand in emotional context, or damaged item. Brand, in this context, means a product or a service with public perception or reputation and most likely has a legal status (e.g., an associated trade name or mark).
- **Provenance**: if, what, and how the asset genealogy is preserved. Genealogy includes workflow inputs, prompts, configurations, outputs, AI tools and capabilities used, and so on. Preservation method examples are visible/invisible watermark, cryptographic/plain-form embedded metadata, or storage in specific production facility. This characterization should align with the Workflow Provenance Tracking plan stated in the Approval Context (3.2.2).

This subcategory helps distinguish between supportive outputs that remain internal and outputs that may affect final media, audience experience, and related considerations.

3.3 Tool

The Tool category captures information about the AI solution provider and how the relevant tool is packaged for use. In this Framework, a tool may be an AI model or a product, service, platform, plug-in, API, workflow environment. A tool offers one or more capabilities for use in media creation.

From a tool vendor perspective, a tool may offer multiple capabilities, only some of which are AI-enabled. From a studio perspective, however, the specific feature of interest may effectively be treated as the relevant tool for governance purposes. The framework addresses this by

distinguishing between **Tool** and **Tool Capability** while allowing a capability to be documented both as part of a larger tool and as the focal unit of evaluation.

3.3.1 Tool Vendor Profile

Tool Vendor Profile describes the organization providing the tool or AI model.

Fields in this subcategory include:

- **Headquarters Location**: the primary country of the vendor (and legal jurisdiction if they are different). Examples are USA, China, or France.
- **Operating Locations**: the regions or jurisdictions in which the vendor operates.
- **AI Posture**: the tool vendor's AI development thesis. Examples are vendor providing its own models, customizing third-party models, using open-source models, or relying on third-party AI services.

This subcategory helps provide organizational context that may matter to customers when evaluating vendor documentation or situating the tool within broader procurement and governance processes.

3.3.2 Tool Packaging

Tool Packaging describes what is being licensed and how the tool is packaged for use.

Fields in this subcategory include:

- **Packaging**: how the tool is packaged and offered as a single unit. Examples are desktop software, SaaS, API-driven service, plug-in, node for use in node-driven workflows, workflow platform, or AI model aggregator/router.
- **Separability**: whether separable AI capabilities can be locked down or selectively governed, such as customer configurable, vendor configurable, user discretion subject to guideline, or not possible because the AI capabilities are deeply integrated within one another.

This subcategory helps clarify what is actually being adopted or proposed for use from a studio perspective and what is being made available from a tool vendor perspective. The packaging of a tool often shapes deployment, data handling, update control, access patterns, and the practical ability to constrain or approve only selected AI-enabled capabilities.

3.4 AI Provenance

The AI Provenance category captures how the AI relevant to the tool or tool capability is enabled and trained. This category may apply at the tool level or the tool capability level depending on the communication need. In some cases, it may also need to be associated with a

particular model when a specific capability relies on a distinct model with distinct provenance characteristics.

AI Provenance is particularly important because many of the questions that arise in governance discussions are not about a tool's user interface or marketing description, but about the models behind it: where they came from, how they were incorporated, and what kinds of training data or source commitments are involved.

3.4.1 AI Model Provisioning

AI Model Provisioning describes what AI models are made available within or alongside the tool.

Fields in this subcategory include:

- **Vendor Enabled AI Models**: models designed, bundled, owned, or pre-configured by the tool vendor for the capability, including cases where none are vendor enabled, where vendor-owned models are used, or where specific third-party models are supported. It is likely that multiple cases may apply.
- **Customer BYO AI Models**: customer-supplied or user-selectable models that the tool supports. Examples are no Bring Your Own (BYO) support, some specific supported models, or any model with certain characteristics. Again, multiple cases may be simultaneously true.

This subcategory is useful because governance questions often differ depending on whether the AI is fixed by the tool vendor, selectable from an approved set, or open to customer-provided models.

3.4.2 Vendor AI Models

Vendor AI Models describes the origin and provider context of the models relevant to the tool or a given capability.

Fields in this subcategory include:

- **Model Origin**: whether the model was vendor-trained, obtained as open weights, derived from open weights, commercial, or derived from a commercial model.
- **Model Provider Headquarters**: the provider headquarters of the model source (and legal jurisdiction if they are different).

Where multiple models are involved, more than one origin or headquarters may apply, unless details about each model are spelled out.

3.4.3 Training Data

Training Data describes the broad source types used in training or adapting the relevant models, and the duration of commitment to those source arrangements.

Fields in this subcategory include:

- **Source Types**: how is the training data sourced by the model provider. Examples are commercially sourced, open-source datasets, public domain, web crawl, vendor-owned, algorithmically generated, or distilled from another model.
- **Named Datasets**: enumeration of specific datasets that are used. Examples are Common Crawl August 2026 Crawl Archive (CC-MAIN-2026-34), Fineweb v1.4.0, and so on.
- **Source Commitment Duration**: what is the commitment of the sources. Examples are tool lifetime, long-term support period, or a defined date-based commitment.

This subcategory surfaces the broad provenance conditions that often underlie several governance concerns.

3.5 Operational Characteristics

Operational Characteristics captures the conditions surrounding deployment and operational use. Like AI Provenance, this category may apply at the tool level or the tool capability level depending on the communication need, and in some cases it may need to be associated with particular models or deployment modes.

3.5.1 AI Deployment

AI Deployment describes where the Tool and AI Models are deployed and who controls its rollout. Note the tool itself may be deployed in one location and the AI model used by the tool in another. Or the tool may rely on multiple AI models (for example, when the tool routes requests and aggregates results).

Fields in this subcategory include:

- **Deployment Artifact**: what is being deployed. Examples are a specific tool that includes the AI Model, a specific tool that uses AI Model deployed separately, or the AI Model.
- **Host Environment**: where is the artifact deployed. Examples are on-device, on-prem, customer-controlled data center, custom-controlled network on cloud, vendor-controlled data center, or vendor-controlled network on cloud.
- **Deployment Location**: where is the geographical region, if the location is vendor-controlled data center or cloud. Examples are US, Ireland, Singapore, and so on.
- **Update Rollout Control**: whether rollout timing and updates of specific tool or AI versions are controlled by the customer or vendor.

This information can highlight security and privacy posture as well as operational predictability.

3.5.2 Customer Data Residency

Customer Data Residency describes where the inputs (including prompts), outputs, workflows, and configurations made by the customer reside operationally and who acts as their custodian.

Fields in this subcategory include:

- **Custodian**: who has operational control over the customer data of a tool or model. Examples are user (the artist using the tool), customer (such as a studio), or the tool vendor.
- **Data Location**: which geographical region is the data located, if the custodian is the tool vendor. Examples are US, Ireland, Singapore, and so on.
- **Data Collection**: whether the vendor collects data for non-operational reasons, such as for making product improvements. Examples are not collected, collected with opt-out, or collection is mandatory.
- **Collection Data Types**: if data is collected, the kind of data that is collected. Examples are prompts, input assets, output assets, workflow configurations, and so on.
- **Collection Retention Period**: if data is collected, how long is the data retained. Examples are until deletion request, until account termination, or indefinite.
- **Vendor Use of Collected Data**: if data is collected, what is the collected data specifically used for. Examples are tool improvement for all customers or customer-specific capability improvement.

This subcategory supports communication around enterprise data control and vendor-customer trust boundaries around data handling.

3.5.3 Metadata Collection

Metadata Collection describes what the tool vendor collects from customer use and for how long.

Fields in this subcategory include:

- **Metadata Collection**: whether the vendor collects data for non-operational reasons, such as for making product improvements. Examples are not collected, collected with opt-out, or collection is required.
- **Collection Metadata Types**: if metadata is collected, the kind of metadata that is collected. Examples are usage timestamps, used AI models or workflows, and so on.
- **Collection Retention Period**: if metadata is collected, how long is the metadata retained. Examples are until deletion request, until account termination, or indefinite.
- **Vendor Use of Collected Metadata**: if metadata is collected, what is the collected metadata specifically used for. Examples are tool improvement for all customers or customer-specific capability improvement.

This subcategory supports communication around privacy, enterprise controls, and vendor-customer trust boundaries around metadata handling.

3.5.4 Customer Model Refinement

Customer Model Refinement describes whether customers may fine-tune or otherwise refine models using their own data and where such refinement occurs.

Fields in this subcategory include:

- **Refinement Allowed**: whether customers can fine-tune the model, yes or no.
- **Refinement Location**: where can such fine-tuning occur. Examples are customer-controlled data center, customer-controlled network on cloud, vendor-controlled data center, or vendor-controlled network on cloud.
- **Vendor Access**: whether vendors can access such refined models, yes or no.
- **Vendor Usage**: whether vendors use such refined models for their use or other customers' use, yes or no.

This becomes relevant where fine-tuning, adaptation, or customer-specific model behavior is part of the proposed solution or workflow.

3.6 Tool Capability

The Tool Capability category classifies a capability based on two governance-relevant dimensions: **output novelty** and **creative controllability**.

This approach intentionally avoids model architecture framing. Model-centric distinctions (such as whether it is diffusion or transformer architecture based) are increasingly less useful as a communication foundation because model architectures, training paradigms, and product features evolve rapidly. The same underlying AI architectures may support both non-generative tasks like metadata tagging and highly generative functions. For governance communication, it is often more useful to describe what the capability does to media outputs and how much control the user has over shaping those results.

3.6.1 Output Novelty

Output Novelty asks how the inputs to the AI are processed, specifically whether the capability primarily extracts information from the input media, enhances existing media, corrects limited portions of existing media, creates substantially new creative elements, or creates altogether a new media.

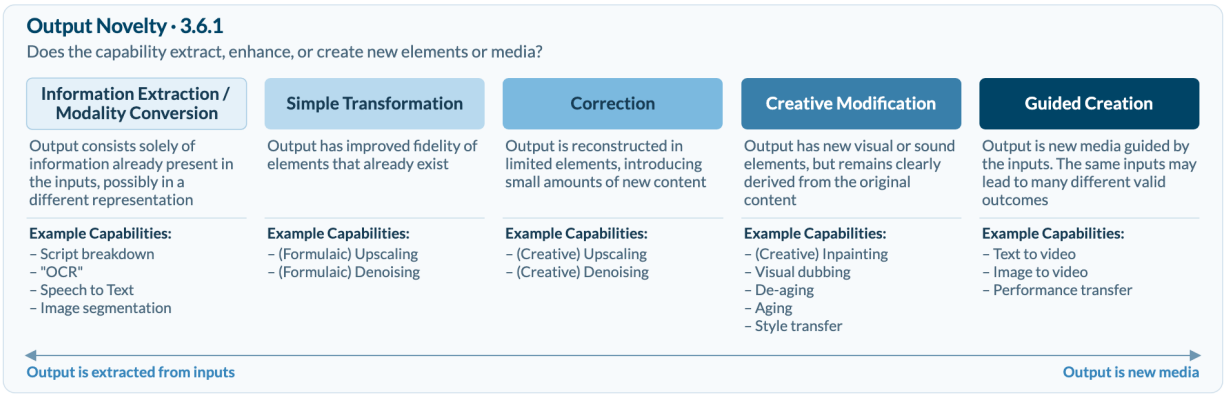


Figure 2: Tool Capability – Output Novelty Dimension

A capability may fall along a spectrum including:

- **Information Extraction / Modality Conversion**: where information already present in the input is extracted or converted into another representation. Examples are script breakdown, document parsing (vision model version of "OCR"), speech to text, or image segmentation.
- **Simple Transformation**: where fidelity is improved or media is transformed without introducing materially new elements. Examples are formulaic upscaling or formulaic denoising.
- **Correction**: where limited new elements are introduced to locally repair or complete an otherwise existing asset. Such repairs could be performed on different layers in the media and may be composited back with other non-AI produced layers or may be performed directly on the completed media. An example is "creative upscaling where the character's facial hair shows more detail or facial skin shows more wrinkles".
- **Creative Modification**: where new visual, audio, or structural elements are introduced in a way that materially changes the media, but the progression from original can be visually spotted or discovered. Examples include visual dubbing, de-aging, and aging.
- **Guided Creation**: where substantially new media is generated, often with many different valid outcomes from similar inputs. Examples are text-to-video, image-to-video, or AI outputs driven by conditioning assets.

This dimension helps communicate the degree to which a capability extracts, preserves, repairs, modifies, or creates output media.

3.6.2 Creative Controllability

Creative Controllability asks how much of the AI-based transformation recipe is defined by, or inherent, to the AI versus explicitly driven by the user.

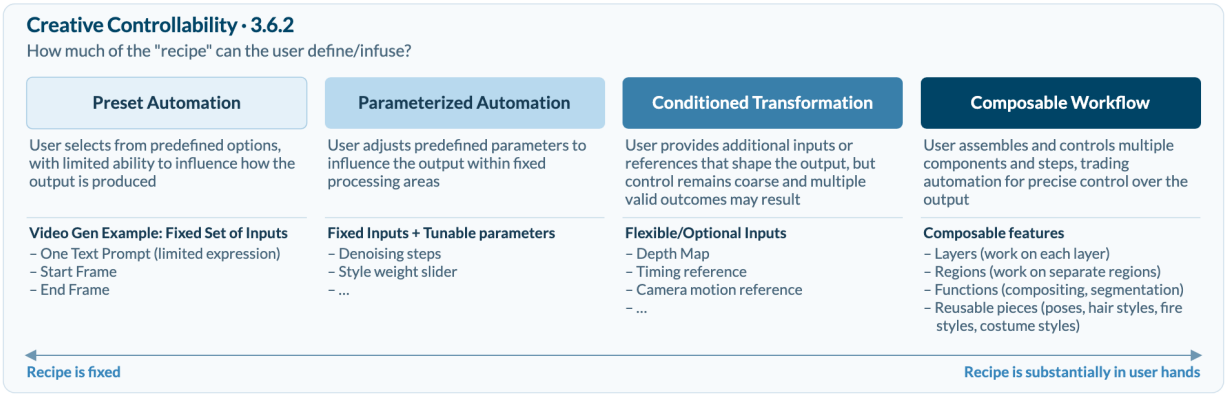


Figure 3: Tool Capability – Creative Controllability Dimension

A capability may fall along a spectrum including:

- **Preset Automation**: where the user selects from predefined operations and the output is largely determined by the tool. In a video generative example, this could be when the capability exposes only a fixed set of input controls to the user, such as a single text prompt (limited tokens), start frame, and end frame.
- **Parameterized Automation**: where the user adjusts predefined parameters to influence the output, but within fixed processing areas. In a video generative example, this could be when the capability exposes a fixed set of input controls to the user plus some limited tunable parameters such as control over the grain in the media (by providing a control on the number of denoising steps).
- **Conditioned Transformation**: where inputs and selected controls meaningfully shape the result within a constrained recipe, and multiple valid outcomes may exist. In a video generative example, this could be when the user has some freedom to control the output using, say, depth maps and camera motion track. In this hypothetical, while the user has the freedom to provide different depth maps and camera motion tracks, they neither can control how the AI uses them to produce the outputs nor can they control the camera angle.
- **Composable Workflow**: where the user assembles and controls the workflow significantly and therefore has substantially more influence over the outcome. In the video generative example, this could be when the user can produce an output by deciding on the number of layers in the media and which layer to target in what order, by having the ability to work on specific region within each layer separately and to use several AI features such as segmentation, compositing, and inpainting on specific targeted regions and layers to produce the media.

This dimension is important because governance may consider not only by what kind of output is produced, but also by how much creative agency, variability, and user-directed composition is involved in producing it.

4. Framework in Action

This section illustrates how the Framework can be used in practice. Its purpose is not to prescribe a universal process or to suggest that all approvals, reviews, or governance bodies operate in the same way. Individual studios are likely to already use a variety of internal review models, escalation paths, and approval checkpoints, and the rigor applied may vary depending on the proposed use, the nature of the content involved, the stage of production, and whether the intended use is limited to internal testing or is intended for public release.

What follows is therefore intended only to provide intuition for how the Framework may support real communication flows among the relevant stakeholders. In particular, the Framework helps those stakeholders communicate using the same categories, sub-categories, and fields when proposing an AI-enabled use, documenting an AI solution, requesting additional information, expressing studio guidance, or documenting actual use in production.

The sections below describe four ways in which the Framework may be used: as part of an approval flow, as a reusable documentation package from tool vendors, as a way for studios to express AI use guidelines, and as a way for production teams and service vendors to document actual AI use in workflows.

4.1 Using the Framework in Approval Flows

In many settings, some form of governance or review body signs off before an AI-enabled use, tool vendor, or tool capability is approved. The composition and authority of that body may differ across organizations. It may include representatives from production technologists, production leadership, legal, compliance, security, privacy, and creative stakeholders, among others. Regardless of the exact structure, these reviewers often weigh anticipated benefit against different forms of risk, including risks related to IP, copyrightability, privacy, security, labor, talent, brand, operational reliability, or public-facing use. The Framework does not replace that evaluation. The Framework is intended to support the evaluation by helping the proposers and reviewers describe relevant facts with greater precision.

4.1.1 Explorative Approval Flow

In one pattern, a production technologist, a pipeline lead, or a restoration specialist discovers a promising new AI-enabled tool or capability in the market. But further testing and evaluation is deemed necessary before the capability can be recommended for use in a production. For example, a technologist may identify a creative upscaling capability that appears to be superior to the currently practiced method for restoring existing tent-pole assets. That individual may recognize that the tool vendor is not yet on an approved list (and that the capability has never been evaluated for the intended context). The individual may therefore seek approval for internal testing first.

In this scenario, the process benefits from the Framework. The approval request package would be light on the business context and the purported business benefit, as the idea is to first perform some internal testing on the functionality and reliability of the solution. The proposal begins with documenting the Tool (3.3) and Tool Capability (3.6), then adding the relevant Use Case (3.2) details needed for the studio to determine whether the request is appropriate for testing.

4.1.2 Production Use Approval Flow

In another pattern, a studio may evaluate the proposed AI use in a largely comprehensive manner. A production leadership team or a studio executive believes an AI-enabled capability presents a substantive business opportunity, such as using AI in support of an upcoming or ongoing project, for say, a localization scenario, that significantly widens the theatrical release market. In that situation, the proposal effectively starts being communicated comprehensively with the proposal including the business use case, details about the tool vendors that could support the scenario, and the relevant tools or tool capabilities in one single package upfront, even if the internal risk analysis moves through those elements in a particular order.

A review sequence begins with the Use Case (3.2) category. Reviewers may first consider the Business Context (3.2.1) and Approval Context (3.2.2): what benefit is sought, what kind of project visibility is involved, what is the release and legal status of the media content that is involved, how mature or speculative the use is, what type of output would be produced, how the user intends to interact with AI and document their interaction, and whether the proposed use is bounded tightly enough for evaluation. That use case analysis may then be weighed against the AI solution provider and the tool itself by looking at the Tool (3.3), AI Provenance (3.4), Operational Characteristics (3.5), and Tool Capability (3.6) categories. The reviewers may review, for example, who the tool vendors are, how the tools are packaged, how the AI in the tools is enabled, what training data source types are involved, where the systems are deployed, what data is collected, and what degree of novelty and creative controllability the relevant capability introduces or supports.

4.1.3 Framework Impact

In either the Exploration Flow or the Production Use Flow, the Framework has the potential to help the process tremendously.

The Framework provides many of the necessary fields for the proposer to communicate to reviewers a proposal package. To the proposer, the Framework makes it clear what important details need to be surfaced to the reviewers using what language and categorization. Where tool vendors already document their offerings using the Framework (see Section 4.2 below), the proposer of the request incorporates that vendor documentation directly into the approval package (perhaps after removing unrelated pieces of documentation) and then adds the relevant Use Case (3.2) information for the proposed production context. The time and effort required to therefore create the request package is likely to have significantly reduced.

To the reviewers, the Framework makes it clear how to interpret and contextualize the information package received from the proposer. The categories, sub-categories, and fields along with their basic definitions makes the communication precise and clear.

The review process may then proceed iteratively. Reviewers may weigh the documented benefits against the risks and determine that more information is required. Importantly, the Framework provides a common language not only for the initial request, but also for follow-up questions. Reviewers may ask for additional information using the Framework categories and fields: for example, additional detail about the Training Data (3.4.3) sources, clarification about the Inputs (3.2.3) and Outputs (3.2.4) in scope, greater precision about the Tool Capability (3.6) being proposed, or more specificity about Customer Data Residency (3.5.2) and Metadata Collection (3.5.3). The result is a more focused exchange than would otherwise occur when different parties use different terms for the same underlying concepts.

Ultimately, the request may be approved, approved with constraints, approved for internal testing only, or rejected. The Framework does not determine that outcome. Its value lies in giving both the proposer and the reviewers a common way to communicating the decision-making information.

4.1.4 The Framework During Transition and After Adoption

The Framework is useful both in a future state where it is broadly adopted and in the current transition state where adoption is still incomplete.

In a mature future state, AI solution providers and tool vendors routinely document their offerings using the Framework. Likewise, the studio guidance on the proposal package include several parts of the Framework directly incorporated. In that environment, a production technologist or a creative user or a studio executive proposing an AI-enabled use can begin with an existing tool vendor-prepared package. That package may already describe the Tool (3.3), AI Provenance (3.4), Operational Characteristics (3.5), and Tool Capability (3.6) in a reusable way. The proposer can then add the specific Use Case (3.2) information relevant to the production and submit a more complete request with fewer clarification cycles.

In the current transition state, however, many tool vendors will not yet have documentation prepared in the language of the Framework. In those situations, the Framework is still valuable because it can shape the approval package even when the information must initially be assembled by the proposer rather than by the tool vendor. A studio may provide internal guidance using the Framework about what should be included in a proposal depending on whether the request is for internal testing, a limited pilot, or public release. The proposer may then prepare a request package using the Framework categories and fields.

That package may include information the proposer can readily supply, especially in the Use Case category, while working in other portions through vendor engagement. During review, the studio's governance or review body may identify gaps. Some additional information may be

requested directly from the proposer, particularly where Business Context (3.2.1), Inputs (3.2.3), Outputs (3.2.4), and proposed workflow details are concerned. Other information may need to come from the tool vendor, particularly where the reviewers need more clarity about AI Provenance, or Operational Characteristics. The studio reviewer may obtain that information directly from the tool vendor, or it may ask the proposer to procure it.

In either case, what matters is that the information request and the information response are both framed using the same language from the Framework. The reviewers can ask for specific sub-categories and fields, and the responder can answer within that same structure. The Framework therefore improves not only documentation quality, but also the precision of the information-seeking process itself. It reduces ambiguity by helping the information seeker and the responder use the same language with the same intended interpretation.

4.2 Using the Framework for Tool Vendor and Tool Documentation

The Framework is also designed to be used by tool vendors to prepare reusable documentation packages for their AI-enabled tools and capabilities and include such documentation as part of their product's standard offering. This is one of the most practical applications of the Framework because studios, service vendors, and production teams often need the same core information each time a tool is used or proposed for use. When that information is documented once (or as frequently as there are tool revisions) using the categories and terms defined in the Framework by the tool vendor, it can be reused across multiple studio evaluations, proposal packages, and production contexts.

From a tool vendor perspective, the most relevant categories are typically Tool (3.3), AI Provenance (3.4), Operational Characteristics (3.5), and Tool Capability (3.6). Together, these categories allow the vendor to communicate who is providing the tool, how the tool is packaged, how the relevant AI is enabled, what broad training data source types are involved, where the system is deployed, what data is collected or retained, whether refinement is supported, and what the relevant capabilities actually do to media. This creates a documentation package that is much more useful than ordinary marketing material or feature lists because it is disclosed around the questions that customers and governance bodies are likely to ask.

Such documentation may be prepared at different levels of granularity depending on the product. In some cases, the relevant information applies across an entire tool. In other cases, separate capabilities within the same tool may rely on different models, distinct training data arrangements, different deployment conditions, or different operational controls. In those circumstances, the Framework allows the vendor to document shared information once at the tool level while also documenting capability-specific details where needed. This is especially

important for tools that combine relatively low-novelty functions, such as denoising or metadata extraction, with high-novelty functions such as object generation or voice cloning.

The value of this documentation is not limited to initial approval discussions. Once prepared, it may be reused repeatedly whenever the tool is considered for different productions, departments, or use cases. A studio technologist proposing a new use may be able to incorporate the vendor's documentation directly into an approval request and then add only the relevant Use Case information for the particular production context. A service vendor may do the same when documenting AI use in a workflow. In this way, the Framework helps shift part of the documentation burden upstream to the party best positioned to describe the tool accurately.

The Framework also helps improve follow-up communication between tool vendors and their customers. When a studio or production team needs more information, it can ask for clarification using the same categories and fields from the Framework. A vendor can then respond with greater precision, for example by clarifying which capability relies on which model, whether customer data is retained, or whether deployment differs by packaging mode. This reduces ambiguity and avoids the situation in which each customer asks substantially similar questions in different language.

In that sense, tool vendor documentation prepared using the Framework becomes more than a product description. It becomes a reusable communication asset that can support procurement, approval, guideline interpretation, and workflow documentation. The Framework therefore gives tool vendors a practical way to describe their AI-enabled solutions in a form that is immediately useful to the broader media creation stakeholders.

4.3 Using the Framework to Communicate Studio Guidelines

Another important use of the Framework is in expressing studio guidance to production teams, service vendors, and other users about how and where AI may or may not be used.

Studios often need to communicate that some uses are allowed, some are disallowed, and some require case-by-case review. Those distinctions are frequently difficult to express clearly when guideline language remains too general, such as when it refers only to "generative AI" or "AI tools". The Framework provides the precision and richness to express guidelines.

For example, a studio can use the Use Case category (3.2) to define the workflow boundaries within which AI use is acceptable. It can specify guideline language in terms of Inputs (3.2.3), Outputs (3.2.4), specifically Intended Output Role, and Final Use Potential. The same studio can then use Approval Context (3.2.2) to explain what kinds of bounded uses are acceptable or may require additional escalation.

The studio can also combine that use case framing with selected fields from Tool (3.3), AI Provenance (3.4), Operational Characteristics (3.5), and Tool Capability (3.6) to communicate what kinds of AI solutions or capabilities are acceptable in those contexts. For example, a guideline may distinguish between capabilities that fall under Simple Transformation versus Guided Creation.

A studio can also use the Framework to express whether only certain tool vendors, packaging models, deployment environments, or data handling conditions are acceptable for particular classes of production work. A guideline may, for example, be expressed in terms of an AI-enabled restoration capability deployed in a customer-controlled environment and the bounded use expressed using the classes in the Output Novelty (3.6.1) and Creative Controllability (3.6.2) dimensions of the Tool Capability.

The Framework is also useful because it allows guidelines to be written at different levels of specificity. Some studio guidance may remain broad and apply across many productions, such as allowing or disallowing certain classes of input materials or requiring a particular deployment and data residency conditions. Other guidance may be narrower and tied to a department, a type of output, a stage of production, or a particular business scenario. The categories and fields of the Framework allow that specificity to be expressed in a consistent way without forcing every guideline into a single template or level of detail.

Another benefit is that the Framework helps make escalation paths clearer. A studio may communicate that some combinations of Inputs, Outputs, Final Use Potential, and Tool Capability are pre-approved, while others require additional review. Rather than relying on broad and potentially ambiguous phrases, the studio can identify more precisely what triggers additional scrutiny. This makes guideline communication more actionable for those designing workflows and reduces the likelihood that production teams or service vendors misinterpret what is or is not permitted.

In this way, the Framework helps studios express AI guidance not merely as general policy language, but as operationally meaningful communication. It allows studios to describe acceptable use boundaries, constrained approval conditions, escalation triggers, and prohibited uses using the same structure that tool vendors and production teams may already be using to document tools and workflows. The result is greater precision, better comparability, and a stronger foundation for alignment across all involved stakeholders.

4.4 Using the Framework to Document AI Use in Production

The Framework is equally useful for production teams, service vendors, and other users to document the actual use of AI in production workflows. This use is closely related to the guideline scenario described above, but the purpose is different. Studio guidance typically describes the range of uses, tools, capabilities, inputs, outputs, and conditions that are permitted or prohibited or require escalation. Production documentation, by contrast, records

the specific AI-enabled tools, capabilities, inputs, outputs, and workflow conditions that were actually used in a given case.

A guideline may express multiple acceptable ways to perform a task, multiple approved tools, or several bounded capability classes. But in documenting an actual production workflow, only the specific tool, capability, deployment condition, input type, and output role that were used should be recorded. The Framework supports that specificity by allowing documentation to be narrowed to the exact categories and fields that are relevant to the implemented workflow.

In practice, a production team or service vendor can use the Use Case category (3.2) to document the usage Context, Inputs, and Outputs associated with the actual use. It can also reference the relevant Tool documentation already prepared by the tool vendor and identify the specific Tool Capabilities in scope. Where necessary, the documentation can also capture the AI Provenance (3.4) and Operational Characteristics (3.5) that are relevant to the particular deployment mode or workflow instance rather than to the tool in the abstract.

This kind of documentation can help demonstrate alignment with studio guidance. If a studio permits AI only for specific bounded purposes, the production documentation can show that the actual use remained within that approved boundary. If the guidance permits certain uses only when specific controls are present, such as customer-controlled deployment or required output review, the production documentation can record those actual conditions as being met (if met) as part of the workflow record.

The Framework is also useful where multiple outputs or multiple workflow branches exist. Different versions of an asset may be created using different permissible AI-enabled methods, and those differences may matter. In such cases, the documentation can associate specific outputs or output sets with the specific documented AI use that produced them. This makes it possible to distinguish, for example, between one version created using a restoration capability and another version created using a more generative capability that required different review.

One thing to note is that while the Framework can be used to document which approved classes were used and applied in production, at the tool capability, workflow, and output level, along with specific ways the user interacted with AI, the Framework, by design, does not provide the nuanced ontological structure needed to capture asset level information, pipeline details, per-asset linkage, event history, or contributor identities. It provides a governance level entry point with the help of AI Interaction and Workflow Provenance Tracking fields in Approval Context (3.2.2) and Workflow Output Provenance (3.2.4) field for asset-level provenance schemes. A technical solution to capture can use this as a foundation to extend and capture the full breadth of the asset provenance / genealogy relationships.

5. Conclusion

AI is becoming increasingly embedded in the tools and workflows used across media creation, and that trend is likely to continue. As a result, clear communication about AI-enabled tools, capabilities, and uses has become essential for organizations involved in media creation including studios, ad agencies, production teams, service vendors, and tool vendors alike.

This Framework is intended to provide a practical and shared guidance for that communication. By highlighting the major categories of information that repeatedly arise in governance, approval, guideline, and workflow discussions, it helps reduce ambiguity, improve comparability, and support more efficient alignment across stakeholders. It does not prescribe policy or replace organization-specific decision making. Rather, it provides a common foundation that different organizations can use to communicate the technical and operational facts that their decisions depend on.

Specifically, the Framework completes the communication loop. Tool vendors document their tools and capabilities using Framework language. Governance teams evaluate Framework-based proposals put forward by technologists or creatives or executives. Studios and agencies communicate their guidance and approval expectations using Framework language. Production teams and service vendors then document the actual AI use in workflows using the same language. This shared foundation improves clarity across the full lifecycle of proposing, evaluating, approving, constraining, and recording AI use in media creation.

Over time, broader adoption of the Framework helps the industry move from fragmented, ad hoc AI discussions towards more consistent and reusable communication and documentation practices.