



Amazon Nova Reel

AWS AI Service Cards



AWS AI Service Cards: Amazon Nova Reel

Copyright © 2025 Amazon Web Services, Inc. and/or its affiliates. All rights reserved.

Amazon's trademarks and trade dress may not be used in connection with any product or service that is not Amazon's, in any manner that is likely to cause confusion among customers, or in any manner that disparages or discredits Amazon. All other trademarks not owned by Amazon are the property of their respective owners, who may or may not be affiliated with, connected to, or sponsored by Amazon.

Table of Contents

Amazon Nova Reel

Overview

Intended use cases and limitations

Design of Amazon Nova Reel

Deployment and performance optimization best practices

Further information

Glossary

1

1

2

6

13

14

15

Amazon Nova Reel

An AWS AI Service Card explains the use cases for which the service is intended, how machine learning (ML) is used by the service, and key considerations in the responsible design and use of the service. A Service Card will evolve as AWS receives customer feedback, and as the service progresses through its lifecycle. AWS recommends that customers assess the performance of any AI service on their own content for each use case they need to solve. For more information, please see [AWS Responsible Use of AI Guide](#) and the references at the end. Please also be sure to review the [AWS Responsible AI Policy](#), [AWS Acceptable Use Policy](#), and [AWS Service Terms](#) for the services you plan to use.

This Service Card applies to the release of Amazon Nova Reel that is current as of July 16, 2025.

Overview

Amazon Nova Reel is a proprietary multimodal foundation model (FM) designed for enterprise use cases. Amazon Nova Reel generates a novel video from a descriptive natural language text string and an optional reference image (together, the “prompt”). Customers can use Amazon Nova Reel to create content within advertising, branding, product design, and social media workflows. This AI Service Card applies to the use of Amazon Nova Reel via [Amazon Bedrock Console](#) and [Amazon Bedrock API](#). Typically, customers use the Console to develop and test applications, and the API for production loads at scale. Each Nova model is a managed subservice of Amazon Bedrock; customers can focus on executing prompts without having to provision or manage any infrastructure such as instance types, network topology, and endpoints.

An Amazon Nova Reel <text prompt, <optional image prompts>, generated video> triple is said to be "effective" if a skilled human evaluator decides that the generated video: 1/ has the content requested by the input prompts (the combination of text and optional image prompts); 2/ makes reasonable assumptions about elements not specified in the input prompts (for example, if asked for a video of a kitchen, a refrigerator and microwave are present and not a couch or a tiger); 3/ is free from defects or image composition errors (for example, human body parts are attached in the correct places and objects are not warped); and 4/ is consistent with the standards of safety, fairness, and other properties valued by the evaluator. Otherwise, a triple is said to be "ineffective." A customer's workflow must decide if a generated video is effective using human judgment, whether human judgement is applied on a case-by-case basis (as happens when the Console is used as a productivity tool by itself) or is applied via the customer's choice of an acceptable score on an automated test.

The “overall effectiveness” of any traditional or generative model for a specific use case is based on the percentage of use-case specific inputs for which the model returns an effective result. Customers should define and measure effectiveness for themselves for the following reasons. First, the customer is best positioned to know which triples will best represent their use case, and should therefore be included in an evaluation dataset. Second, different video generation models may respond differently to the same prompt, requiring tuning of the prompt and/or the evaluation mechanism.

As with all ML solutions, Amazon Nova Reel must overcome issues of intrinsic and confounding variation. Intrinsic variation refers to features of the input that the model should generate, for example, knowing the difference between the text prompts *'a cute cat'* and *'a cute dog'*. Confounding variation refers to features of the input that the model should ignore, for example, understanding that the text prompts *'a jumping cat'* and *'the jumping cat'* should return the same video, since there should be no semantic difference between 'a' and 'the'. The full set of variations encountered in the input text prompts include language (human and machine), slang, professional jargon, dialects, expressive non-standard spelling and punctuation and many kinds of errors in prompts, for example, with spelling, grammar, punctuation, logic, and semantics.

Intended use cases and limitations

Amazon Nova Reel serves a wide range of potential application domains and offers the following core capabilities:

- Reel 1.0
 - Generate videos of up to 6 seconds given a text prompt.
 - Generate videos of up to 6 seconds given a reference image and a text prompt.
- Reel 1.1
 - Generate videos of up to 120 seconds given a text prompt.
 - Generate videos of up to 120 seconds given a reference image and a text prompt.

The features differ in the parameters (for example, size or number of reference images) required to invoke them. For more information about these specifications, see [Amazon Nova User Guide](#).

When assessing a video generation model for a particular use case, we encourage customers to specifically define the use case, that is, by considering at least the following factors: the **business problem** being solved; the **stakeholders** in the business problem and deployment process; the

workflow that solves the business problem, with the model and human oversight as components; key system **inputs and outputs**; the expected intrinsic and confounding **variation**; and, the types of **errors** possible and the relative impact of each.

For example, Amazon Nova Reel can be used as a creative tool to help advertisers or brands create video-based assets for their advertising or marketing campaigns. The business goal may be to improve the cost, quality, and productivity required to create a video asset to be used in marketing campaigns. The stakeholders may include the advertiser or brand, who wants to create a functional video ad.

Amazon Nova Reel 1.0 offers only Single Shot mode. Amazon Nova Reel 1.1 offers three modes: 1/ Single Shot, 2/ Multi-shot Automated, and 3/ Multi-shot Manual

In Single Shot mode, customers can create a 6-second video asset using only a text prompt or a text prompt in combination with an input image. The workflow is the customer 1/ provides a text prompt describing the desired video, and 2/ provides an image to be used as the starting frame for the video (optional).

In Multi-shot Automated mode, customers have the ability to create a video asset using a single text prompt. The workflow is the customer 1/ provides a text prompt describing the desired video, and 2/ specifies the desired output video duration. The output will be a long video comprising multiple shots that reflect the text prompt and video duration delivered as 1/ a single stitched video in the specified duration, and 2/ individual 6 second shots stored separately in the customer's storage location. Shots are fixed at 6 seconds. If the input or resulting video violates [our core dimensions of Responsible AI](#), the model will not generate the video

In Multi-shot Manual mode, customers have the opportunity to provide a text prompt and optional input image for each shot within the long video, allowing granular controllability which is essential for precise storytelling and creative direction. The workflow is the customer 1/ provides a series of text prompts, one for each shot, and 2/ provides images to be used as the starting frames for one or more of those shots (optional). The output will be a long video comprising multiple shots that reflect the text prompts and images provided for each shot, delivered as 1/ single stitched video, and 2/ individual 6 second shots stored separately in the customer's storage location. Shots are fixed at 6 seconds. If the input or resulting video violates [our core dimensions of Responsible AI](#), the model will not generate the video. Output videos will contain any images provided as input by the user, depictions of the details mentioned in the prompt, as well as other components that the model fills in. Input variations include all the normal variations in English expression across different individuals, differences in the definition of design concepts/jargon, inaccuracies, misspellings, and undefined abbreviations. The error types, ranked in order of

estimated negative impact on stakeholders, include a/ harmful or otherwise inappropriate content, and b/ misrepresenting the input product.

Text Prompt

"A cute raccoon playing guitar underwater."

Output Video

[Sample video output from Nova Reel](#)

Amazon Nova Reel has a number of limitations requiring careful consideration.

Appropriateness for Use

We make every effort to design, develop, and rigorously test our models to help ensure they produce appropriate outputs based on user inputs, but generative models are by their nature non-deterministic and may occasionally produce unintended or undesirable outputs. We encourage users to report questions and provide feedback [here](#) about our models to help us continuously improve their performance. Customers should evaluate outputs for accuracy and appropriateness for their use case, especially if they will be directly surfaced to end users. Additionally, if Amazon Nova Reel is used in customer workflows that produce consequential decisions, customers must evaluate the potential risks of their use case and implement appropriate human oversight, testing, and other use case-specific safeguards to mitigate such risks. See the [AWS Responsible AI Policy](#) for more information. Customers who use Amazon Nova Reel models are responsible for ensuring that their use of Amazon Nova Reel and the generated video or other output complies with all applicable laws. The model and output may not be used for any prohibited practices under the EU AI Act.

Safety Filters

Amazon Nova Reel is designed to disengage with attempts to circumvent its safety measures through prompt engineering. If a customer's video output generation request is unsuccessful, it may be due to one or more such measures. The safety filters for Amazon Nova Reel cannot be configured or turned off. However, they are periodically assessed and improved in response to feedback.

Text and Image Inputs

Amazon Nova Reel text prompts for Single Shot or Multi-shot Manual modes cannot exceed 512 characters. Prompts for Multi-shot Automated model cannot exceed 4,000 characters. Additionally, input images are limited to specific dimensions (only 1280x720). Reference

images, which are submitted as part of the prompt, can be formatted as either PNG or JPEG. For more information, see [Amazon Nova User Guide](#).

Novel Outputs

Amazon Nova Reel is designed to create novel videos through its own expression in generated videos and not to imitate any particular existing works.

Limited Prior Knowledge

Amazon Nova Reel is trained from images and videos of objects. It does not store explicit 3D models of objects, or model lighting physics. It does not know about every possible object (including living creatures) or every possible variation of an object. Amazon Nova Reel also does not know about all possible arrangements of objects, or about actual distributions of objects (for example, how many cars are appropriate to show near the Arc de Triomphe at given time or the average demographic distribution of soldiers in the French foreign legion in the 1960s).

Limited Specification

Customers can influence the composition of a generated video via detailed prompting. However, customers should not expect to be able to describe all aspects of any desired video within the maximum character length of a text prompt, e.g., there are many possible generated videos that would match a text prompt of '*the dog*.' Amazon Nova Reel "fills in" unspecified information automatically, extrapolating creatively from images and videos. Thus, customers may encounter unexpected elements in generated outputs, such as unrealistic shapes with impossible angles or proportions, inconsistent lighting, or unnatural colorations.

Image Design, Composition, and Stylization

There are a wide variety of video styles (for example, illustration, digital, fine art, anime, caricature, and cartoon), and each style can be interpreted and expressed in many ways. Amazon Nova Reel is designed to produce realistic styles. Furthermore, in the absence of clear instruction in the text, Amazon Nova Reel might produce outputs that have cropped compositional elements (for example, for the prompt '*a red sports car*', the resulting image might be a close up of the wheels or the door handle of the car) and might arbitrarily vary camera angles, lighting, and object pose.

Generating Humans

Amazon Nova Reel is based on an ML technology (diffusion model) that does not explicitly model parts of objects. As a result, it might produce depictions of the human face and body that are anatomically incorrect (for example, noses, fingers and toes). Customers who use Amazon Nova Reel to generate humans are responsible for ensuring that the output and use

of the generated video complies with all applicable laws, including but not limited to laws governing biometric privacy or digital replicas.

On-screen Text

Users should not expect Amazon Nova Reel to generate coherent text within generated outputs.

Languages

Amazon Nova Reel currently only supports video output generation for English prompts. However, if the user prompts the model to embed short non-English text strings (such as Spanish) or small numbers in the generated image, the model may produce those requests. For example, the model can produce acceptable results to a prompt that states: *'a man holding up a sign that says "Hola mundo!"'* or *'a blue backpack with "1845" printed on the front'*.

Output Resolution and Framerate

For a list of supported resolutions and framerates, see [Amazon Nova User Guide](#).

Modalities

Amazon Nova Reel does not currently support audio or 3D content.

Design of Amazon Nova Reel

Machine Learning

Amazon Nova Reel performs video output generation using machine learning, specifically, a text-conditioned [diffusion model](#). At a high level, the core service works by encoding text prompts as numerical vectors, finding nearby vectors in a joint text/image embedding space that correspond to images, and then using these vectors to transform a low-resolution image encoding of random values into several key frame images that captures the information in the prompt. This new key frame images encoding is then expanded into a full high-resolution video. A diffusion model is trained on videos of objects and associated captions, but not directly on 3D models of objects or on real-world physics. The runtime service architecture for Amazon Nova Reel works as follows: 1/ Amazon Nova Reel receives a user prompt (along with desired configuration parameters) using our API or Console; 2/ the model filters the prompt to comply with safety, fairness, and other design goals. If a filter is triggered, then the prompt is rejected and no output is produced; 3/ if no filter is triggered, the prompt is sent to the model and an output is generated; 4/ additional output moderation and filters are applied to further check for safety and other concerns; 5/ lastly, if no filters are triggered, the video is returned to the customer.

Controllability

We say that a Nova model exhibits a particular "behavior" when it generates the same kind of output for the same kinds of prompts and configuration (e.g., seed). For a given model architecture, the control levers that we have over the behaviors are primarily a/ the training data corpus, b/ different parameters such as seed, and c/ the filters we apply to pre-process prompts and post-process outputs. Our development process exercises these control levers as follows: 1/ we pre-train the FM using curated data from a variety of sources, including licensed and proprietary data, open source datasets, and publicly available data where appropriate; 2/ we adjust model weights via supervised fine tuning (SFT) to increase the alignment between Nova models and our design goals; and 3/ we tune safety filters (such as privacy and toxicity filters) to block or evade potentially harmful prompts and video outputs to further increase alignment with our design goals.

Performance Expectations

In general, we expect implementations of similar video output generation use cases by different customers to vary in their inputs, their configuration parameters, and in how overall effectiveness is measured. Consider two applications A and B, each a version of the home design use case described above, but deployed by different companies. Each application will face similar challenges, e.g., the designer and owner will likely differ in the language they use to express design ideas, and in the degree of verisimilitude they expect/need of the output video and the owner's actual expectation. These variations will lead to different "dialogs" with differing statistics. As a result, the overall utility of Amazon Nova Reel will depend both on the model and on the workflow it enables. We recommend that customers test Amazon Nova Reel both on their own content and with different workflows.

Test-driven Methodology

We use multiple datasets and human work forces to evaluate the performance of Nova models. No single evaluation dataset suffices to completely capture performance. This is because evaluation datasets vary based on use case, intrinsic and confounding variation, and other factors. Our development testing involves automated testing against publicly available and proprietary datasets, benchmarking against proxies for anticipated customer use cases, human evaluation of outputs against proprietary datasets, manual red teaming, and more. Our development process examines Amazon Nova Reel's performance using all of these tests, takes steps to improve the model and/or the suite of evaluation datasets, and then iterates.

- **Human Evaluation:** Human evaluation is a critical step in evaluating the model's outputs. For example, the effectiveness criteria for outputs generated for an advertisement campaign

could differ from those for a fashion design concept, i.e., an ad campaign owner might care more about the realism of the face than a fashion designer, who is more focused on the composition of fabrics and colors of the clothing. Using human judgement is critical for assessing the effectiveness of a video model on more challenging tasks, because only people can fully understand the context, intent, and nuances of more complex prompts and video output generations.

- Independent Red Teaming Network: Independent Red Teaming Network: Consistent with our Frontier AI Safety Commitments on ensuring Safe, Secure, and Trustworthy AI, we partner with a variety of third parties to conduct red teaming against our AI models. We leverage red teaming firms to complement our in-house testing in areas such as safety, security, privacy, fairness, and veracity-related topics. We also work with specialized firms and academics to red-team our models for specialized areas such as Cybersecurity and Chemical, Biological, Radiological, and Nuclear (CBRN) capabilities.

Safety

Safety is a shared responsibility between AWS and our customers. Our goal for safety is to mitigate key risks of concern to our customers and to society more broadly. Our customers represent a diverse set of use cases, locales, and end users, so we have the additional goal of making it easy for customers to adjust model performance to their specific use cases and circumstances. Amazon Nova Reel is designed to block problematic inputs and outputs. In a case where a customer asks Amazon Nova Reel to generate a video and no safety filters are triggered, the model will return a video. In a case where the model cannot complete a prompt, it will not display a video and should generate an error message. The system is designed to prevent the generation of content that may cause physical or emotional harm to a consumer, as well as content that may harass, harm, or encourage harm to individuals or specific groups, especially children. Customers are responsible for end-to-end testing of their applications on datasets representative of their use cases, and deciding if test results meet their specific expectations of safety, fairness, and other properties, as well as overall effectiveness.

- Harmlessness: We evaluate Amazon Nova Reel 1.1's ability to accurately reject potentially harmful prompts using multiple datasets. For example, on a proprietary dataset (1.3k samples) containing prompts that attempt to solicit videos containing harmful content (e.g., abuse, violence, hate, nudity, insults, profanity), Amazon Nova Reel 1.1 correctly blocks 96.4% of harmful prompts. In order to ensure we are maintaining high performance, we augment our training dataset with benign prompts, and we measure our true pass rate for harmless prompts using an internally curated set of common-nouns and phrases with a 92.1% pass rate.

- **Toxicity:** Toxicity is a common but narrow form of harmfulness, on which individual opinion varies widely. We assess our ability to avoid prompts and to not generate videos that contain potentially toxic content using automated testing with multiple datasets, and find that Amazon Nova Reel 1.1 performs well on common types of toxicity. For example, on a proprietary toxic image-prompt dataset (3.4k samples), which we classified into sub-categories (e.g., violence, gore, self-harm), Amazon Nova Reel 1.1's end-to-end toxicity guardrails accurately block 95.8% of toxic content.
- **Chemical, Biological, Radiological, and Nuclear (CBRN):** Compared to information available via internet searches, science articles, and paid experts, we see no indications that Amazon Nova Reel increases access to information about chemical, biological, radiological or nuclear threats. Consistent with our voluntary endorsement of the Frontier AI Safety Commitments at the AI Seoul Summit, we continue to test for CBRN risk, and engage with other video generator vendors to share, learn about, and mitigate possible CBRN threats and vulnerabilities.
- **Famous/Public Figures:** Amazon Nova Reel has safeguards to deter the generation of videos of famous or public figures to help prevent the use of generative AI for intentional disinformation or deception.
- **Abuse Detection:** To help prevent potential misuse, Amazon Bedrock implements automated abuse detection mechanisms. These mechanisms are fully automated, so there is no human review of, or access to, user inputs or model outputs. To learn more, see [Amazon Bedrock Abuse Detection](#) in the *Amazon Bedrock User Guide*.
- **Child Sexual Abuse Material (CSAM):** At Amazon, we are [committed](#) to producing generative AI services that keep child safety at the forefront of development, deployment, and operation. We utilize Amazon Bedrock's Abuse Detection solution (mentioned above), which uses hash matching or classifiers to detect potential CSAM. If Amazon Bedrock detects apparent CSAM in user image inputs, it will block the request, display an automated error message and may also file a report with the National Center for Missing and Exploited Children (NCMEC) or a relevant authority. We take CSAM seriously and will continue to update our detection, blocking, and reporting mechanisms.

Fairness

Amazon Nova Reel is designed to generate videos that a diverse set of customers will find effective across the wide range of object categories that customers may wish to depict. Two key design questions are: 1/ should the service be able to render any variation of an object (for example, person, pet, car) if explicitly asked, and 2/ what should the service default to

generating when the user does not provide guidance in the prompt? We have taken steps to address both questions.

1. It is not currently possible to build training datasets that cover all varieties of every object; however, for humans in particular, we aim to combat societal bias and cultural appropriation. We test Amazon Nova Reel ability to moderate these outcomes using a proprietary dataset of aggregated red teaming iterations that depict bias, stereotyping, and hate against individuals and groups. We find that Amazon Nova Reel 1.1 blocks 95.2% of observed bias in generations.
2. When users provide no guidance about the desired attributes of an object or person, it is unclear how to judge output over repeated renderings of the object. For example, for the prompt "basketball players", some users might prefer a team with similar demographic attributes and other might want a distribution of attributes (for example, gender) matching some distribution they have in mind. Given this ambiguity, when there is no information included in the prompt, Amazon Nova Reel is designed to return diverse results, but without specifying a desired distribution. Given the current limits of datasets and technology, and the intrinsic ambiguity of generating videos without guidance, we recommend that customers consider specifying desired object attributes in the text prompt.

Explainability

Customers wanting to interpret the output of an Amazon Nova model can utilize [Titan Multimodal Embeddings](#) to output numerical representations (known as embeddings) of both the prompt text and the generated image produced by the model. These embeddings capture the semantic information present in the prompts and Nova model's outputs and can be compared (using cosine similarity, euclidean distance or some other measure) to verify that the produced output is consistent with the input prompt. Generated output that is more consistent with the prompt will have larger similarity (lower distance) than output that does not contain relevant information.

Veracity

Videos created by diffusion models can contain unrealistic representations of known objects or representations of new objects that are not physically possible in the real world. From a customer's perspective, whether this is an advantage or a disadvantage depends on the use case. Nova was trained to favor generating more "typical" objects (e.g., hands with five fingers) unless otherwise specified in the prompt.

Robustness

Amazon Nova Reel is optimized for creativity. Customers can expect that similar prompts will generate similar outputs, in the sense that *"the blue backpack"* and *"a blue backpack"* will both yield videos that contain blue backpacks. However, customers should not expect that semantically identical prompts (as above) will necessarily generate identical outputs, given the goal of novelty. Instead, we prioritize having the focus of the generated outputs align with the focus of the text prompt. We measure this alignment with testing on both public benchmarks and proprietary datasets.

Privacy

Amazon Nova Reel is available in Amazon Bedrock. Amazon Bedrock is a managed service and does not store or review customer prompts or customer video outputs, and prompts and outputs are never shared between customers, or with Amazon Bedrock third party model providers. AWS does not use inputs or outputs generated through the Amazon Bedrock service to train Amazon Bedrock models, including Amazon Nova Reel. For more information, see Section 50.3 of the [AWS Service Terms](#) and the [AWS Data Privacy FAQs](#). For service-specific privacy information, see Security in the [Amazon Bedrock FAQs](#)

- **PII:** The system is designed to prevent the generation of content that contains personally identifying information. If a user is concerned that their personally identifying information has been included in Nova model outputs, the user should contact us [here](#).

Security

All Amazon Bedrock models, including Amazon Nova Reel, come with enterprise security that enables customers to build generative AI applications that support common data security and compliance standards, including GDPR and HIPAA. Customers can use AWS PrivateLink to establish private connectivity between customized Titan models and on-premises networks without exposing customer traffic to the internet. Customer data is always encrypted in transit and at rest, and customers can use their own keys to encrypt the data, for example, using AWS Key Management Service (AWS KMS). Customers can use AWS Identity and Access Management (IAM) to securely control access to Amazon Bedrock resources. Also, Amazon Bedrock offers comprehensive monitoring and logging capabilities that can support customer governance and audit requirements. For example, Amazon CloudWatch; can help track usage metrics that are required for audit purposes, and AWS CloudTrail can help monitor API activity and troubleshoot issues as Amazon Nova Reel is integrated with other AWS systems. Customers can also choose to store the metadata, prompts, and video generations in their own encrypted Amazon Simple Storage Service (Amazon S3) bucket. For more information, see [Amazon Bedrock Security](#).

Intellectual Property

Amazon Nova Reel is designed for generation of new creative content. AWS offers uncapped intellectual property (IP) indemnity coverage for outputs of generally available Amazon Nova models (see Section 50.10 of the [AWS Service Terms](#)). This means that customers are protected from third-party claims alleging IP infringement or misappropriation (including copyright claims) by the outputs generated by these Amazon Nova models. In addition, our standard IP indemnity for use of the Services protects customers from third-party claims alleging IP infringement (including copyright claims) by the Services (including Amazon Nova models) and the data used to train them.

Transparency

Amazon Nova Reel provides information to customers in the following locations: this Service Card, AWS documentation, AWS educational channels (for example, blogs, developer classes), and the AWS Console. We accept feedback through customer support mechanisms such as account managers. Where appropriate for their use case, customers who incorporate Nova models in their workflow should consider disclosing their use of ML to end users and other individuals impacted by the application, and customers should give their end users the ability to provide feedback to improve workflows. In their documentation, customers can also reference this Service Card.

- **Watermarking:** Amazon Nova Reel applies an invisible watermark to all videos it generates, helping identify AI-generated videos to promote the safe, secure, and transparent development of AI technology and helping reduce the spread of disinformation. The detection solution can also check for the existence of the watermark, helping customers confirm whether a video was generated by Nova models. For more information, see the [AWS News launch blog](#), [Amazon Nova product page](#), [Amazon Nova User Guide](#) and [watch the demo](#).
- **Content Credentials:** To increase transparency around AI-generated content, Amazon Nova Reel 1.1 adds Content Credentials to videos it generates. Content Credentials are based on a technical specification developed and maintained by the [Coalition for Content Provenance and Authenticity](#) (C2PA), a cross-industry standards development organization. C2PA metadata includes the model and the platform used to generate the video which allows people to identify the source/provenance of generated images. To inspect the Content Credentials, you can drag and drop videos generated using Amazon Nova Reel 1.1 into publicly available tools such as [Verify](#).

Governance

We have rigorous methodologies to build our AWS AI services responsibly, including a working backwards product development process that incorporates Responsible AI at the design phase, design consultations, and implementation assessments by dedicated Responsible AI science and data experts, routine testing, reviews with customers, best practice development, dissemination, and training.

Deployment and performance optimization best practices

We encourage customers to build and operate their applications responsibly, as described in [AWS Responsible Use of AI Guide](#). This includes implementing Responsible AI practices to address key dimensions including controllability, safety, fairness, veracity, robustness, explainability, privacy, security, transparency, and governance.

Workflow Design

The performance of any application using Amazon Nova Reel depends on the design of the customer workflow, including the factors discussed below:

- 1. Effectiveness Criteria:** Customers should define and enforce criteria for the kinds of use cases they will implement, and, for each use case, further define criteria for the inputs and outputs permitted, and for how humans should employ their own judgment to determine final results. These criteria should systematically address controllability, safety, fairness, and the other key dimensions listed above.
- 2. Configuration:** In addition to the required text prompt, Amazon Nova Reel has various required and optional configuration parameters to help customers achieve the best results. For more information, see [Amazon Nova User Guide](#).
- 3. Prompt Engineering:** Prompt engineering refers to the common practice of optimizing the text inputs of FMs to obtain desired responses. High-quality prompts condition the model to generate more desirable videos. Some key aspects of prompt engineering include word choice, style and tone, structure and length, guiding details that narrow the scope, iteratively refining the prompt, and drawing inspiration from examples. For more detailed guidelines, see [Amazon Nova User Guide](#). Here are some specific tips to consider when constructing prompts:
 - a. Prompt style and tone:** To get the best results, prompts should read like video captions (*'a black train moving through a lush mountain range'*), not like chat messages (*'I want a*

black train with rlly sick mountains') or commands (*'generate an image of black train, lush mountain range'*). Effective prompts tend to be detailed but not overly long, providing key visual features, styles, emotions or other descriptive elements. Prompts should not include negating language like *'no cats'* or *'not brown'*.

- b. **Prompt details:** When crafting a prompt, customers should focus their writing on details they want included in the video and avoid superfluous language (for example, *'award winning'*, *'4K'*, *'ultra-high resolution'*, *'best image'*). These statements have little to no impact on the quality of the output.
- c. **On-screen text:** When trying to display text elements in the image generation, Amazon Nova Reel produces better results when provided with double quotes in the prompt. For example, *'an image of a boy holding a sign that says "success"'* instead of *'an image of a boy holding a sign that says success'*.
4. **Human Oversight:** If a customer's application workflow involves a high risk or sensitive use case, such as a decision that impacts an individual's rights or access to essential services, human review should be incorporated into the application workflow where appropriate.
5. **Performance Drift:** A change in the types of prompts that a customer submits (for example, asking for photo-realistic generations instead of animated generations) to Amazon Nova Reel might lead to different outputs. To address these changes, customers should consider periodically retesting the performance of Amazon Nova Reel and adjust their workflow if necessary.
6. **Updates:** We will notify customers when we release a new version, and will provide customers time to migrate from an old version to the new one. Customers should consider retesting the performance of the new Nova model version on their use cases when changing to the updated model.

Further information

- For service documentation, see [Amazon Nova](#), [Amazon Bedrock Documentation](#), [Amazon Bedrock Security and Privacy](#), [Amazon Bedrock Agents](#), and [Amazon Nova User Guide](#).
- For details on privacy and other legal considerations, see the following AWS policies: [Acceptable Use](#), [Responsible AI](#), [Legal](#), [Compliance](#), and [Privacy](#).
- For help optimizing workflows, see [Generative AI Innovation Center](#), [AWS Customer Support](#), [AWS Professional Services](#), [Ground Truth Plus](#), and [Amazon Augmented AI](#).
- If you have any questions or feedback about AWS AI service cards, please complete [this form](#).

Glossary

Controllability: Steering and monitoring AI system behavior.

Privacy & Security: Appropriately obtaining, using and protecting data and models.

Safety: Preventing harmful system output and misuse.

Fairness: Considering impacts on different groups of stakeholders.

Explainability: Understanding and evaluating system outputs.

Veracity & Robustness: Achieving correct system outputs, even with unexpected or adversarial inputs.

Transparency: Enabling stakeholders to make informed choices about their engagement with an AI system.

Governance: Incorporating best practices into the AI supply chain, including providers and deployers.