

Service Name

Amazon Rekognition

Service Overview

Amazon Rekognition is a cloud-based image and video analysis service that enables developers to add advanced computer vision capabilities to their applications without requiring machine learning expertise. It uses deep learning technology to detect and analyze objects, scenes, concepts, faces, text, celebrities, and inappropriate content in billions of images and videos daily. The service offers both synchronous image analysis through Amazon Rekognition Image and asynchronous video analysis through Amazon Rekognition Video, supporting both stored videos and real-time streaming analysis.

Key Use Cases

- **Searchable media libraries:** Enable content discovery and organization through automated tagging and metadata extraction
- **Identity verification and face-based authentication:** Verify user identities through face matching and liveness detection
- **Content moderation:** Automatically detect and filter inappropriate, explicit, or unsafe content in images and videos
- **Security and surveillance:** Monitor premises through real-time detection of people, objects, and activities
- **Media analysis and workflow automation:** Streamline video production with shot detection, segment identification, and frame-accurate analysis
- **Personal protective equipment (PPE) detection:** Ensure workplace safety compliance by detecting safety equipment usage
- **Text extraction and document processing:** Extract and analyze text from images for document digitization and processing

Features

- **Object and Scene Detection:** Detects thousands of objects, scenes, concepts, and activities with confidence scores and bounding boxes
- **Facial Analysis and Recognition:** Analyzes facial attributes, emotions, landmarks, and enables face comparison and search capabilities
- **Text Detection and Extraction:** Detects and extracts text from images and videos with word-level confidence scores
- **Celebrity Recognition:** Identifies thousands of celebrities with additional metadata including names and IMDb IDs
- **Content Moderation:** Detects explicit, suggestive, or violent content to enable automated content filtering

- **Custom Labels:** Creates custom models to detect specific objects, scenes, or concepts unique to business needs
- **Video Streaming Events:** Processes real-time video streams for low-latency object detection including people, pets, and packages
- **PPE Detection:** Identifies personal protective equipment including face covers, head covers, and hand covers
- **Media Analysis:** Provides frame-accurate segment detection including shot changes, black frames, and color bars
- **Face Liveness Detection:** Verifies presence of live users during authentication to prevent spoofing attacks

Value Proposition

Amazon Rekognition eliminates the complexity and cost of building custom computer vision solutions by providing pre-trained, highly accurate deep learning models accessible through simple APIs. It offers unmatched scalability to process billions of images and videos, integrates seamlessly with other AWS services, and requires no machine learning expertise:

- The service provides enterprise-grade security and compliance including HIPAA eligibility, while offering cost-effective pay-per-use pricing with generous free tiers
- Customers benefit from continuous model improvements, responsible AI practices, and the ability to customize models for specific business needs through Custom Labels

Service Integration

Amazon Rekognition integrates deeply with Amazon S3 for storing and processing images and videos, AWS Lambda for serverless processing workflows and event-driven automation, Amazon SNS for notifications when analysis jobs complete, Amazon SQS for decoupling processing workflows, Amazon Kinesis Video Streams for real-time video analysis, Amazon CloudWatch for monitoring and metrics, AWS CloudTrail for API logging and auditing, Amazon DynamoDB for storing analysis results and metadata, and Amazon Augmented AI (A2I) for human review workflows. It also supports integration with Amazon EventBridge for event-driven architectures and IAM for fine-grained access control.

Onboarding and Offboarding

Getting started requires creating an AWS account and setting up appropriate IAM permissions for Amazon Rekognition access. Customers then install and configure the AWS CLI and SDKs for their preferred programming languages:

- The service offers a web-based console for testing capabilities, comprehensive developer guides with code examples in multiple languages, and step-by-step tutorials for common use cases

- Amazon provides ready-to-use CloudFormation templates, sample applications, and integration patterns with other AWS services to accelerate adoption
- The free tier allows customers to process 1,000 images and 60 minutes of video monthly at no cost

Getting Started Guide: <https://docs.aws.amazon.com/rekognition/latest/dg/getting-started.html>

Data Removal

Customers retain full control and ownership of their content processed by Amazon Rekognition. Image and video inputs can be deleted at any time through the console or APIs. Face collections and associated metadata can be removed using DeleteCollection and DeleteFaces operations. For complete offboarding, customers should delete all face collections, remove stored analysis results, and revoke IAM permissions. AWS automatically deletes temporary processing data, and customers can use AWS CloudTrail to audit all data access and deletion activities.

Backup/Restore and Disaster Recovery

Amazon Rekognition is a managed service where AWS handles infrastructure backup and disaster recovery. Face collections and user vectors are automatically replicated across multiple Availability Zones for high availability. Customers are responsible for backing up their own analysis results, face collection metadata, and application configurations. Analysis results stored in other AWS services like DynamoDB or S3 can leverage those services' backup capabilities including AWS Backup integration.

Backup Capabilities

Amazon Rekognition does not directly integrate with AWS Backup as it is a fully managed API service. However, customers can backup face collections by exporting face vectors and metadata, and analysis results stored in integrated services like DynamoDB or S3 can utilize AWS Backup. The service maintains automatic replication of face collections for high availability, and customers should implement backup strategies for their analysis workflows and results.

Disaster Recovery

Amazon Rekognition provides built-in disaster recovery through multi-AZ replication of face collections and service infrastructure. The service does not directly integrate with AWS Elastic Disaster Recovery, but customers can implement disaster recovery for their applications using Rekognition by replicating face collections across regions, backing up analysis results, and designing multi-region architectures. AWS manages the underlying service resilience and availability.

Recovery Objectives

Amazon Rekognition supports low RPO and RTO objectives through its multi-AZ architecture and high availability design. Face collections are automatically replicated, enabling near-zero RPO for face data. For comprehensive disaster recovery, customers should replicate face collections across regions, backup analysis results, and design applications with failover capabilities. The service's API-based architecture enables rapid recovery and regional failover when properly architected.

Service Constraints

Service Quotas: <https://docs.aws.amazon.com/rekognition/latest/dg/limits.html>

FAQ: <https://aws.amazon.com/rekognition/faqs/>

Technical Requirements

Service Documentation: <https://docs.aws.amazon.com/rekognition/latest/dg/what-is.html>

User Guide: <https://docs.aws.amazon.com/rekognition/latest/dg/getting-started.html>

API Reference:

<https://docs.aws.amazon.com/rekognition/latest/APIReference/Welcome.html>

Pricing Overview

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Public Pricing: <https://aws.amazon.com/rekognition/pricing/>

Cost Optimization

Amazon Rekognition offers several cost optimization opportunities including tiered pricing that decreases cost per image with higher volumes, free tier providing 1,000 images and 60 minutes of video monthly, ability to set confidence thresholds to process only relevant content, custom moderation adapters to reduce false positives, bulk analysis capabilities for batch processing efficiency, and regional pricing variations. Customers can optimize by choosing appropriate AWS regions, implementing intelligent filtering, using batch processing, and leveraging Custom Labels for specific use cases to avoid processing irrelevant content.

Savings Considerations

Primary cost savings come from leveraging the generous free tier limits, implementing volume-based tiered pricing benefits, using batch processing for large datasets, setting optimal confidence thresholds to reduce unnecessary processing, utilizing Custom Labels

to focus on business-specific objects, choosing cost-effective regions, integrating with AWS Lambda for event-driven processing to avoid continuous polling, and implementing intelligent pre-filtering to process only relevant media. Organizations can achieve significant savings by designing efficient workflows, avoiding redundant analysis, and taking advantage of the pay-per-use model with no upfront commitments or minimum fees.