

Service Name

AWS IoT Greengrass

Service Overview

AWS IoT Greengrass is an open-source IoT edge runtime and cloud service that extends AWS IoT Core capabilities to edge devices. It enables devices to act locally on generated data, run machine learning models, filter and aggregate device data, and communicate securely with other devices and AWS IoT Core. The service provides local compute, messaging, data caching, and ML inference capabilities for connected devices, allowing them to operate even with intermittent cloud connectivity. AWS IoT Greengrass V2 introduces modular applications called components that can be deployed using Lambda functions, Docker containers, native OS processes, or custom runtimes.

Key Use Cases

- **Industrial IoT and Edge Computing:** Deploy edge applications for predictive maintenance, factory automation, and equipment monitoring in manufacturing environments
- **Smart Cities and Infrastructure:** Enable autonomous operation of traffic systems, smart lighting, and environmental monitoring with local data processing
- **Connected Vehicles and Transportation:** Support autonomous vehicle operations, fleet management, and real-time vehicle diagnostics with offline capability
- **Agriculture and Environmental Monitoring:** Implement precision agriculture solutions, greenhouse automation, and environmental sensing with intermittent connectivity
- **Healthcare and Remote Monitoring:** Deploy patient monitoring systems, medical device management, and healthcare analytics at the edge with secure local processing

Features

- **Local Processing:** Run AWS Lambda functions locally, support containers, AWS IoT Device Shadows, local messaging, and local resource access
- **Machine Learning Inference:** Deploy and run ML models at the edge with optimized performance using Amazon SageMaker and Neo compiler integration
- **Stream Manager:** Manage and process data streams from multiple sources with local buffering and cloud synchronization
- **Over-the-Air Updates:** Deploy software updates, configuration changes, and new components remotely to edge devices
- **Security Integration:** Hardware Security Module integration, secrets management, and end-to-end encryption for secure edge operations

- **Client Device Support:** Act as secure MQTT endpoint for local devices without direct cloud access, with local messaging and shadow synchronization
- **Component System:** Modular software architecture using pre-built components for extending device functionality and building applications
- **Offline Operation:** Continue local processing, messaging, and data caching when internet connectivity is intermittent or unavailable

Value Proposition

AWS IoT Greengrass provides unique edge computing capabilities that reduce latency, bandwidth costs, and dependency on cloud connectivity while maintaining security and scalability. It offers seamless integration with AWS IoT Core and other AWS services, enabling hybrid cloud-edge architectures:

- The service supports both lightweight (nucleus lite) and full-featured (nucleus) runtimes, accommodating resource-constrained to powerful edge devices
- Key advantages include local ML inference with optimized performance, offline operation capabilities, centralized fleet management, and extensive hardware ecosystem support
- Organizations benefit from reduced operational costs through local data processing, improved responsiveness with sub-millisecond latency for critical applications, and enhanced reliability with autonomous edge operation during connectivity issues

Service Integration

AWS IoT Greengrass integrates deeply with AWS IoT Core for device management, messaging, and shadow synchronization, serving as the primary cloud connection point. It connects with AWS Lambda for serverless edge computing, Amazon S3 for component artifacts and data storage, and AWS Secrets Manager for secure credential management:

- The service integrates with Amazon SageMaker and AWS IoT Analytics for ML model deployment and data analysis
- CloudWatch provides logging and monitoring capabilities, while AWS IAM handles authentication and authorization
- Additional integrations include Amazon Kinesis for data streaming, AWS Step Functions for workflow orchestration, Amazon DynamoDB for edge data storage, and AWS Systems Manager for configuration management
- The service also works with AWS IoT Device Management for fleet provisioning and AWS CloudFormation for infrastructure deployment

Onboarding and Offboarding

Getting started with AWS IoT Greengrass requires an AWS account, creating an AWS IoT thing for the core device, and installing the Greengrass Core software on a Linux or Windows device with minimum 512MB RAM and 1GB storage. The automated device setup

script simplifies initial configuration by installing dependencies, downloading certificates, and creating the Greengrass group:

- Alternatively, customers can follow the step-by-step tutorial covering environment setup, core device configuration, component development, and deployment testing
- Prerequisites include Python 3.6+, Java Runtime Environment, and basic familiarity with IoT concepts
- The service provides comprehensive documentation, video tutorials, and hands-on labs to accelerate adoption
- Fleet provisioning capabilities enable large-scale deployments with zero-touch device onboarding

Getting Started Guide:

<https://docs.aws.amazon.com/greengrass/v2/developerguide/getting-started.html>

Data Removal

When offboarding from AWS IoT Greengrass, customers must delete components and deployments from the cloud service, uninstall the Greengrass Core software from edge devices, and remove associated AWS IoT things, certificates, and policies. Local data stored on edge devices should be manually deleted, including component artifacts, logs, and cached data in the Greengrass work directories. Cloud-side data removal involves deleting component versions, deployment configurations, and associated CloudWatch logs. Customers should revoke device certificates and remove IAM roles and policies created for Greengrass operations to ensure complete data cleanup.

Backup/Restore and Disaster Recovery

AWS IoT Greengrass provides built-in resilience through local data caching, persistent sessions with AWS IoT Core, and automatic reconnection capabilities. For backup and recovery, the service supports local storage destinations for critical data, component artifact backup to S3, and configuration backup through Infrastructure as Code. Disaster recovery solutions include multi-region device registry replication, automated failover using Amazon Route 53, and cross-region component deployment capabilities. The AWS Disaster Recovery for AWS IoT solution provides automated replication of device shadows, certificates, and policies to secondary regions for comprehensive DR coverage.

Backup Capabilities

AWS IoT Greengrass includes native backup capabilities through local data storage, component artifact versioning in S3, and configuration backup via CloudFormation templates. The service supports local storage caching for resilience during connectivity issues and maintains persistent sessions for message delivery:

- While not directly integrated with AWS Backup, customers can backup device data using custom components and leverage S3 versioning for component artifacts

- Stream Manager provides local buffering and data persistence, while local shadows maintain device state information for recovery purposes

Disaster Recovery

AWS IoT Greengrass supports disaster recovery through multi-region deployment capabilities, automated device registry replication, and cross-region failover mechanisms. The service works with Amazon Route 53 health checks for automatic endpoint switching and supports AWS Application Recovery Controller for coordinated disaster recovery orchestration. While not directly integrated with AWS Elastic Disaster Recovery, customers can implement comprehensive DR strategies using AWS Step Functions, Lambda, and DynamoDB Global Tables for device registry replication across regions with automated failover capabilities.

Recovery Objectives

AWS IoT Greengrass helps customers achieve low RPO and RTO objectives through local data persistence, automatic reconnection capabilities, and multi-region deployment support. The service's offline operation capabilities minimize data loss during outages, while persistent MQTT sessions ensure message delivery upon reconnection. Stream Manager's local buffering reduces RPO to near-zero for critical data, and automated deployment mechanisms enable rapid service restoration. Cross-region replication and automated failover capabilities support RTO objectives of hours to minutes depending on the DR strategy implemented.

Service Constraints

Service Quotas: <https://docs.aws.amazon.com/general/latest/gr/greengrassv2.html>

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

Technical Requirements

Service Documentation:

<https://docs.aws.amazon.com/greengrass/v2/developerguide/what-is-iot-greengrass.html>

User Guide: <https://docs.aws.amazon.com/greengrass/v2/developerguide/getting-started.html>

API Reference:

<https://docs.aws.amazon.com/greengrass/v2/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/greengrass/pricing/>

Cost Optimization

AWS IoT Greengrass offers unique cost optimization through edge data processing that reduces cloud computing and data transfer costs. Local ML inference eliminates cloud API calls for real-time predictions, while Stream Manager's data aggregation and compression minimize bandwidth usage:

- The service's offline operation capabilities reduce dependency on continuous cloud connectivity, lowering data transfer costs
- Component-based architecture enables selective deployment of only required functionality, optimizing resource usage
- Free tier includes three Greengrass cores for one year, and pay-per-active-core pricing model aligns costs with actual usage
- Local data filtering and preprocessing reduce cloud storage and processing costs by sending only relevant data to AWS services

Savings Considerations

Primary cost savings with AWS IoT Greengrass come from reduced data transfer costs through local processing and selective data transmission to the cloud. Edge ML inference eliminates expensive cloud API calls for real-time predictions, while local data aggregation reduces cloud storage and processing costs:

- The service minimizes bandwidth requirements through data compression and filtering at the edge, particularly beneficial for remote deployments with expensive connectivity
- Operational savings result from centralized fleet management, automated updates, and reduced need for on-site technical support
- Free tier offering and usage-based pricing model allow cost-effective scaling
- Local processing capabilities reduce dependence on cloud resources for time-sensitive operations, optimizing overall infrastructure costs while maintaining performance requirements