

Service Name

AWS Resilience Hub

Service Overview

AWS Resilience Hub is a centralized service that enables organizations to manage, assess, and improve the resilience posture of their applications running on AWS. The service provides a comprehensive platform for defining resilience goals, continuously assessing application weaknesses against Recovery Time Objectives (RTO) and Recovery Point Objectives (RPO), and implementing recommendations based on AWS Well-Architected Framework best practices. AWS Resilience Hub integrates with AWS Fault Injection Service to conduct chaos engineering experiments, automatically discovers and maps application resources, and provides actionable insights through drift detection, resiliency scoring, and operational recommendations. The service helps organizations transform resilience management from a reactive process into a proactive strategy by providing automated assessments, detailed recommendations, and clear resilience metrics across multi-account and multi-region architectures.

Key Use Cases

- **Application Resilience Assessment:** Continuously assess applications against defined RTO/RPO targets to identify weaknesses and ensure compliance with resilience policies across different failure scenarios including application, infrastructure, AZ, and regional outages
- **Proactive Disaster Recovery Planning:** Define and validate disaster recovery strategies by simulating various disruption scenarios using AWS FIS experiments and implementing standard operating procedures (SOPs) for efficient recovery
- **Multi-Account Resilience Management:** Centrally manage resilience across distributed, multi-account cloud architectures using a hub-and-spoke model to ensure consistent resilience policies and assessments
- **CI/CD Pipeline Integration:** Integrate resilience validation into software development lifecycle to identify and resolve resilience weaknesses before applications are released into production
- **Operational Readiness and Compliance:** Maintain regulatory and compliance requirements by tracking resilience posture, implementing recommended alarms, and ensuring applications meet defined business continuity objectives

Features

- **Resilience Dashboard:** Visual dashboard providing comprehensive overview of application resilience status, compliance, drift detection, and resiliency scores across multiple applications

- **Application Discovery and Management:** Automatic discovery and mapping of AWS resources from CloudFormation stacks, Resource Groups, EKS clusters, and Terraform state files into logical application components
- **Resiliency Policy Management:** Define and manage customizable resilience policies with specific RTO and RPO targets for different disruption types including software, hardware, AZ, and regional failures
- **Automated Assessments:** Continuous resilience assessments with scheduled daily evaluations, drift detection capabilities, and compliance reporting against defined policies
- **AWS FIS Integration:** Integrated chaos engineering capabilities with tailored fault injection experiments to validate application resilience and verify recovery procedures
- **Operational Recommendations:** AI-powered recommendations for infrastructure improvements, alarm configurations, and standard operating procedures based on AWS Well-Architected Framework
- **Multi-Account Support:** Hub-and-spoke architecture enabling centralized resilience management across multiple AWS accounts with cross-account role assumptions and trust relationships
- **Resiliency Scoring:** Quantitative scoring system measuring application readiness to handle disruptions based on adherence to recommendations, alarms, SOPs, and successful tests

Value Proposition

AWS Resilience Hub provides unique value through its comprehensive, centralized approach to application resilience management that transforms reactive disaster recovery into proactive resilience engineering. Unlike fragmented solutions, it offers unified visibility across multi-account architectures with automated resource discovery and continuous assessment capabilities:

- The service uniquely integrates with AWS Well-Architected Framework to provide evidence-based recommendations and incorporates AWS FIS for validated chaos engineering testing
- Customers benefit from AI-powered insights using Amazon Bedrock for assessment summaries, cost-optimized resilience recommendations with multiple optimization strategies (cost, minimal changes, best AZ/region recovery), and seamless integration with existing CI/CD pipelines
- The platform reduces operational overhead through automated drift detection, scheduled assessments, and standardized operating procedures while ensuring regulatory compliance through detailed reporting and metrics
- Organizations can achieve faster time-to-resilience with pre-built policies, experiment templates, and actionable recommendations that eliminate guesswork in disaster recovery planning

Service Integration

AWS Resilience Hub deeply integrates with core AWS services to provide comprehensive resilience management. It works natively with AWS CloudFormation and Terraform for infrastructure-as-code resource discovery, Amazon EKS and ECS for containerized workload assessment, and AWS Resource Groups for logical resource organization:

- The service integrates with AWS Fault Injection Service for chaos engineering experiments and AWS Systems Manager for standard operating procedures automation
- It leverages Amazon CloudWatch for alarm management and monitoring, AWS SNS for drift notifications, and Amazon S3 for metrics export and storage
- The platform connects with AWS Backup and AWS Elastic Disaster Recovery services for backup and disaster recovery capabilities assessment
- Integration with AWS Trusted Advisor provides additional resilience checks, while AWS CloudTrail enables audit logging of resilience activities
- For enterprise deployments, it works with AWS Organizations through IAM cross-account roles and supports myApplications for unified application management
- The service also integrates with Amazon Bedrock for AI-powered assessment summaries and recommendations, providing intelligent insights across the entire AWS ecosystem

Onboarding and Offboarding

Getting started with AWS Resilience Hub requires setting up IAM permissions and roles, particularly the `AWSResilienceHubAssessmentRole` for resource access. Customers begin by adding applications through resource collections from CloudFormation stacks, AWS Resource Groups, EKS clusters, or Terraform state files:

- The process involves creating or selecting resiliency policies with appropriate RTO/RPO targets, publishing the application version, and running initial assessments
- For multi-account deployments, organizations establish cross-account trust relationships using `AWSResilienceHubCrossAccountRole`
- The service offers a six-month free trial for the first three applications, allowing customers to evaluate capabilities before committing
- Setup includes configuring SNS topics for drift notifications, scheduled assessments, and implementing recommended improvements
- AWS provides free online training, self-guided labs, GitHub repositories with implementation examples, and comprehensive documentation to accelerate adoption
- Customers can leverage pre-built resiliency policies and experiment templates to quickly establish baseline resilience management

Getting Started Guide: <https://docs.aws.amazon.com/resilience-hub/latest/userguide/getting-started.html>

Data Removal

When offboarding from AWS Resilience Hub, customers can delete applications, assessments, and associated data through the console or API operations. Deleting an application removes all versions, assessments, and associated metadata from the service. The DeleteApp API operation permanently removes applications and cannot be undone. Customers should first export any required metrics or reports before deletion. Assessment data, recommendation templates, and past assessments are retained for 365 days by default before automatic deletion. To completely remove data, customers must explicitly delete individual applications, policies, and exported S3 buckets created by the service. The service automatically cleans up resources and stops charging when applications are deleted.

Backup/Restore and Disaster Recovery

AWS Resilience Hub assesses and recommends backup and disaster recovery capabilities by evaluating integration with AWS Backup and AWS Elastic Disaster Recovery services. The service checks whether applications have appropriate backup plans, cross-region replication, and disaster recovery configurations. It provides recommendations for implementing backup strategies across supported resources like RDS, EFS, EBS, and DynamoDB. The platform evaluates disaster recovery readiness across different failure scenarios including AZ and regional outages, recommending multi-AZ deployments, cross-region replication, and standby resources to meet defined RTO/RPO targets.

Backup Capabilities

AWS Resilience Hub itself does not provide backup capabilities but extensively evaluates and recommends AWS Backup integration for applications. The service assesses whether resources have appropriate AWS Backup plans configured and provides recommendations for backup coverage gaps:

- It analyzes backup configurations for services like Amazon RDS, EFS, EBS, DynamoDB, and other supported resources to ensure they align with defined RPO requirements
- The platform generates specific recommendations for implementing backup strategies that meet resilience policy objectives

Disaster Recovery

AWS Resilience Hub supports disaster recovery planning by assessing application readiness for various failure scenarios and integrating with AWS Elastic Disaster Recovery service. The platform evaluates cross-region disaster recovery configurations, multi-AZ deployments, and standby resource availability:

- It provides specific recommendations for implementing disaster recovery strategies including warm standby, pilot light, and backup-and-restore approaches
- The service assesses recovery capabilities against defined RTO targets and recommends improvements to meet disaster recovery objectives across different disruption types

Recovery Objectives

AWS Resilience Hub helps customers meet RPO and RTO objectives by providing continuous assessment against defined targets, identifying gaps in current configurations, and recommending specific improvements. The service evaluates actual recovery capabilities across different failure scenarios (application, infrastructure, AZ, regional) and provides cost-optimized recommendations to achieve target objectives. It offers multiple optimization strategies including best AZ recovery, best region recovery, least cost, and minimal changes approaches. The platform validates recovery procedures through chaos engineering experiments and provides quantified estimates of achievable RTO/RPO based on current configurations.

Service Constraints

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

FAQ: <https://aws.amazon.com/resilience-hub/faqs/>

Technical Requirements

Service Documentation: <https://docs.aws.amazon.com/resilience-hub/latest/userguide/what-is.html>

User Guide: <https://docs.aws.amazon.com/resilience-hub/latest/userguide/>

API Reference: <https://docs.aws.amazon.com/resilience-hub/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/resilience-hub/pricing/>

Cost Optimization

AWS Resilience Hub offers unique cost optimization through its multi-dimensional recommendation engine that provides four distinct optimization strategies: cost optimization, minimal changes, best AZ recovery, and best region recovery. The service analyzes current configurations and recommends the most cost-effective approaches to

achieve resilience targets, including right-sizing resources, optimizing backup strategies, and selecting appropriate disaster recovery patterns:

- Customers benefit from the six-month free trial for first three applications, followed by a simple \$15 per application per month pricing model with no upfront costs
- The platform helps avoid over-provisioning by providing precise recommendations based on actual RTO/RPO requirements rather than generic disaster recovery approaches
- Cost transparency is built-in with detailed cost estimates for recommended changes, enabling informed decision-making about resilience investments
- The service prevents costly outages by identifying weaknesses before they impact production, providing significant ROI through proactive resilience management

Savings Considerations

Key cost savings with AWS Resilience Hub include avoiding expensive downtime through proactive weakness identification and prevention of over-engineered disaster recovery solutions through precise, requirements-based recommendations. The service eliminates the need for separate resilience management tools and reduces operational overhead through automated assessments and drift detection:

- Organizations save on compliance and audit costs through built-in reporting and documentation capabilities
- The platform's integration with existing AWS services eliminates additional tooling costs while providing comprehensive resilience management
- Customers benefit from optimized resource utilization recommendations that reduce waste in backup and disaster recovery configurations
- The predictable pricing model (\$15 per application per month after free trial) enables accurate budgeting compared to complex, usage-based alternatives
- Early identification of resilience gaps prevents costly emergency remediation and reduces the risk of business disruption costs
- Multi-account management capabilities reduce administrative overhead and ensure consistent, efficient resilience policies across organizations