

kyndryl.

aws

Enhance the resilience of your critical workloads with Reliability Engineering Services



Contents

- 02 Executive summary
- 03 Balancing cost and risk
- 04 Integrating AWS native resilience services with Kyndryl
- 05 Harnessing multi-region architectures
- 05 Selecting the optimal DR approach
- 05 Simplifying DR processes with automation
- 06 Why choose Reliability Engineering Services?
- 06 Connect with us

Executive summary

If your organization suffered a cyber incident, how confident would you be in your disaster recovery and business continuity capabilities?

In this white paper, we explore how Kyndryl helps organizations to build and maintain highly resilient systems on Amazon Web Services (AWS). Our Reliability Engineering Services provide an end-to-end approach to resilience: from initial assessment and design to the implementation and ongoing management of resilient architectures.

As an AWS Resilience Services Competency Partner, Kyndryl helps businesses use the latest multi-region failover capabilities from AWS to automate and streamline disaster recovery and business continuity processes. With Reliability Engineering Services, enterprises can maintain hard-won customer confidence and ensure consistent service delivery in an increasingly digital world.



Balancing cost and risk

Across many sectors, digital channels are now the primary customer touchpoint. For a growing number of businesses, these systems are essential for delivering mission-critical services. Even short periods of disruption or downtime can carry substantial financial, reputational and regulatory risk.

To meet and exceed customer expectations, it has never been more important for businesses to build secure and resilient systems. But when it comes to resilience, a one-size-fits-all approach will rarely deliver the optimal balance between cost and risk.

Businesses need to weigh several factors when architecting resilient digital systems. One of the most important considerations is the business criticality of each service. For some services, periods of downtime may be acceptable. However, truly mission-critical systems will require the highest level of availability at almost any cost. According to

Gartner, large enterprises face downtime costs averaging \$500,000 to \$1 million per hour, while in high-stakes sectors like finance and healthcare, the costs may exceed \$5 million per hour.¹ In these high-stakes sectors, five or even six nines of availability may be deemed essential to keep services available 24/7. Similarly, retailers may devote a considerable portion of their IT budget to ensure the availability of their e-commerce environments.

However, many applications can tolerate lower levels of availability. Three nines, or 99.9% availability, represents just under nine hours of downtime: still a relatively short amount of time over the course of an entire year. Determining the proper number of nines for your applications will also depend on budget constraints, application architectures and operational maturity.



Integrating AWS native resilience services with Kyndryl


For businesses looking to build resilient architectures that can withstand and rapidly recover from disruption, AWS offers a wide range of global infrastructure capabilities. The AWS platform is built on multiple layers of redundancy. Globally, there are 32 AWS Regions, 102 AWS Availability Zones, and an extensive network of AWS Direct Connect and AWS Edge Network Locations. The AWS platform's elasticity allows for immediate capacity scaling, and its pay-as-you-go model makes maintaining standby environments more cost-effective compared to on-premises data centers. Furthermore, AWS's compliance certifications and security controls help organizations meet regulatory requirements for business continuity and disaster recovery (DR).

On top of this strong foundation, AWS provides a range of native services designed specifically for high availability (HA) and DR. These include: Amazon Aurora Global Database, Amazon DynamoDB Global Tables, and Amazon Simple Storage Service Cross-Region Replication, all of which provide built-in replication mechanisms. Amazon Application Recovery Controller (ARC) enables automated failover orchestration.

Reliability Engineering Services combine deep cloud expertise with AWS's powerful native services to deliver enterprise-grade resilience for your most critical workloads. This approach ensures that business continuity and disaster recovery are not just technical features, but strategic advantages.

Key integration highlights:

- **Automated recovery and failover:** Using AWS Application Controller (ARC), Kyndryl automates the process of switching workloads between regions in the event of a disruption. This approach enables rapid recovery with minimal manual intervention.
- **Global data availability:** Using services such as Amazon Aurora Global Database and DynamoDB Global Tables, Kyndryl helps ensure data availability across regions, supporting uninterrupted operations and a consistent user experience.
- **Secure, compliant data replication:** Using Amazon S3 Cross-Region Replication, Kyndryl configures policies to keep business-critical data protected and accessible, meeting regulatory requirements and recovery objectives.



Reliability Engineering Services integrate AWS native services into a comprehensive resilience framework by combining automated failover, continuous monitoring and detailed recovery playbooks. This holistic approach means that every layer of your AWS environment is protected and recoverable.

In the event of a regional outage, Kyndryl's automation triggers rapid failover, validates data replication, and ensures your applications and data remain available. Kyndryl provides real-time dashboards and compliance documentation for peace of mind.

Harnessing multi-region architectures

Businesses that run key systems on AWS have several options when it comes to data resilience. For many, architecting systems that span multiple AWS Regions is the most effective approach for achieving HA and DR objectives. Indeed, AWS itself reports that multi-region architectures are now a top priority for mission-critical workloads, driven by compliance, latency, and disaster recovery needs.²

Organizations typically choose multi-region architectures to address critical requirements such as geographic redundancy for DR, reduced latency and improved application performance for global users, and compliance with data sovereignty regulations. According to AWS, organizations adopting multi-region designs typically reduce recovery time objectives (RTO) from 72 hours to under 2 hours.²

ARC now offers Region switch: a fully managed, highly available capability that enables organizations to simplify multi-region failover operations by efficiently planning, practicing, and orchestrating Region switches. The service orchestrates recovery procedures across AWS Regions while providing real-time visibility through integrated dashboards and detailed logging capabilities. By automating manual processes such as failover execution, dashboard creation, and recovery validation documentation, this new feature reduces operational overheads while strengthening operational efficiency and compliance.

Selecting the optimal DR approach

In a basic DR configuration, organizations may use one AWS Region for production workloads while maintaining a second AWS Region as a passive standby environment. This approach trades a longer RTO and recovery point objective (RPO) for significantly lower operational cost versus an active-active configuration.

By distributing workloads across multiple AWS Regions simultaneously, businesses can achieve near-zero RTO and RPO. However, the requirement for sophisticated data replication strategies, consistent application state management and intelligent traffic routing drives up cost and management complexity substantially. For this reason,

active-active configurations are particularly suitable for organizations requiring continuous availability and/or serving a global user base, while active-passive setups may suffice for organizations with more flexible recovery time requirements and/or users based in one geographical area.

ARC Region switch supports both active-passive configurations, managing complete failover and fallback procedures, and active-active architectures, where it can orchestrate the removal or addition of Regions from the active pool. This automated approach transforms what traditionally required hours of careful coordination and engineering effort into a streamlined, repeatable process that organizations can execute with confidence. ARC Region switch represents a significant advancement in making sophisticated multi-region resilience strategies more accessible and manageable for organizations of all sizes.

Simplifying DR processes with automation

Fundamentally, the operational success of AWS multi-region resilience strategies depends on building automation frameworks that can consistently execute complex failover and fallback procedures.

Services such as ARC Region switch are a powerful way to orchestrate failover and fallback procedures. However, organizations face significant challenges in maintaining accurate, up-to-date operational runbooks and dependency maps; these can quickly become outdated in dynamic cloud environments. Automation frameworks must address several critical aspects, including continuous validation of infrastructure state across regions, automated capacity scaling mechanisms that respond to regional failures, and orchestrated traffic management.

In an active-active configuration, the automation must constantly monitor regional health, adjust traffic distribution and increase capacity in healthy AWS Regions if another Region shows degradation. For active-passive configurations, it is crucial to maintain infrastructure parity and execute rapid failovers, including database synchronization, application deployment validation, and systematic traffic cutover. Often overlooked but equally critical is the automation of the fallback process. This is crucial to ensure safe return to the original operating state without data loss or service disruption.

Why choose Reliability Engineering Services?

The business case for investing in resilience is clear. Organizations with advanced business continuity and resilience programs report an average 4x ROI, driven by reduced downtime, easier regulatory compliance and improved customer trust.³

For businesses looking to strengthen the resilience of key workloads on AWS, Kyndryl brings deep experience in delivering best-of-breed DR solutions. Our Always On Center of Excellence offers a 99.999%+ uptime SLA from an end-to-end service perspective, guaranteeing uninterrupted and dependable service on AWS. Our five-nines-plus SLA reduces expected annual downtime to five minutes or less.

Reliability Engineering Services bring a comprehensive approach to cloud resilience that combines deep technical expertise with proven methodologies for implementing and managing highly available systems on AWS.

For mission-critical systems, Reliability Engineering Services typically take a multi-active approach. Such an architecture uses multiple AWS Availability Zones and Regions, strengthening reliability, simplifying maintenance, and supporting blue-green deployments. These capabilities make the multi-active approach particularly valuable for applications that demand the highest levels of resilience and availability. For systems that can tolerate longer periods of downtime, Kyndryl is also highly experienced in delivering active-passive solutions — offering simplicity, cost-efficiency, and ease of integration with legacy applications.

As an AWS Resilience Services Competency Partner, Kyndryl has developed a framework that addresses the full spectrum of resilience requirements. The framework runs from initial assessment and architecture design to ongoing operational management and continuous improvement. Reliability Engineering Services are built on four key pillars:

01. Infrastructure resilience through multi-region architectures
02. Application resilience through modern development practices
03. Operational resilience through automated recovery procedures
04. Business resilience through comprehensive continuity planning

Reliability Engineering Services include automated testing frameworks, continuous compliance monitoring and detailed recovery playbooks. These help to ensure that organizations can maintain resilience while meeting regulatory requirements. Reliability Engineering Services make regular testing of these procedures — now a regulatory requirement in many sectors — more manageable and less risky. In this way, organizations can carry out DR exercises more frequently without overwhelming operational teams. As a result, businesses can typically shift from periodic DR exercises to a continuous, automated and verifiable approach.

Connect with us

Reliability Engineering Services combine our extensive experience in managing mission-critical workloads on AWS with the data resiliency capabilities of the AWS platform. The service empowers organizations to move beyond traditional DR models and achieve true operational resilience.

Ready to get started? [Learn more](#) about Kyndryl's AWS Offerings, or find out how to enhance the resilience of your most critical workloads with Reliability Engineering Services by [connecting with us here](#).





© Copyright Kyndryl, Inc. 2025

Kyndryl is a trademark or registered trademark of Kyndryl, Inc. in the United States and/or other countries. Other product and service names may be trademarks of Kyndryl, Inc. or other companies.

This document is current as of the initial date of publication and may be changed by Kyndryl at any time without notice. Not all offerings are available in every country in which Kyndryl operates. Kyndryl products and services are warranted according to the terms and conditions of the agreements under which they are provided.

Sources:

1. <https://www.erwoodgroup.com/blog/the-true-costs-of-downtime-in-2025-a-deep-dive-by-business-size-and-industry/>
2. <https://docs.aws.amazon.com/pdfs/prescriptive-guidance/latest/aws-multi-region-fundamentals/aws-multi-region-fundamentals.pdf>
3. <https://riskconnect.com/en-gb/content-library/bci-continuity-resilience-report-2024/>