

## Highlights

- A lifecycle approach with people, processes, information, and technology
- End-to-end analysis of applications and workload landscapes
- Service mapping to enable active services with failure transparency
- Site Reliability Engineering and a cultural shift to always-on

# Kyndryl services for always-on workloads with Google Cloud Platform

Modernize the reliability, resiliency, and availability of your mission-critical workloads

Today's information-infused society hinges on cloud, analytics, mobile, and social. On-the-go consumers are almost always within reach of an internet-connected device. Users expect quick, consistent access to their mission-critical applications and business-critical data around the clock. That's why outages—whether they're planned or unplanned—can negatively affect revenue, consumer trust, and your reputation. In an increasingly digital world, your business needs an always-on platform that doesn't compromise the end-user experience or allow for service disruption.

What do we mean by an always-on platform? The idea of "always-on" is about achieving continuous operations so the business service being consumed is not interrupted, making downtime transparent to the end user.

**An always-on platform should be able to:**

- Withstand component failures affecting cloud zones, such as service components, networking infrastructure, operating systems, and even applications
- Withstand catastrophes such as cloud region failures, hurricanes, tornadoes, floods, and other natural disasters
- Introduce changes—for example, scheduled blue/green deployments, application releases, infrastructure updates, and security patches—without disruption

The challenges of always-on are in planning, designing, implementing, and managing the platform, and the right approach involves the full spectrum of people, processes, applications, and technology.

Kyndryl can help ensure your mission-critical workloads are always-on, working with you to modernize the reliability, resiliency and availability of your mission-critical workloads while migrating or transforming your applications on the Google Cloud Platform.

Our approach relies on a set of tried-and-tested principles and clear, well-defined methods to assess weaknesses and their potential impact to the reliability of a mission-critical workload by analyzing the end-to-end business service.

**This analysis encompasses:**

- Infrastructure
- Messaging
- Platforms
- Culture
- Applications

## A lifecycle approach with people, processes, information, and technology

As your organization considers today's business drivers, it's important to look closely at your own requirements.

**Some key questions to consider include:**

1. Are you modernizing your mission-critical workloads to Google Cloud Platform?
2. Do you want zero-outage changes to avoid planned maintenance windows?
3. Are you looking for disaster avoidance or fast failover capabilities instead of disaster recovery for your critical services?
4. In the event of a catastrophic disaster, do your customers expect your services to remain always-on?
5. Do you require the ability to continue uninterrupted service to your customers?
6. Are component failures causing extended downtime and disruption to your business services?
7. Are your digital channels growing, increasing your digital presence, and driving revenue?
8. Are lines of business (LOBs) driving increased changes within the IT platform and causing operational disruption?

If any of these questions apply to your organization, it's time to change your approach. Kyndryl provides an always-on lifecycle methodology that brings together people, processes, information, and technology in a structured manner, designed to help maintain consistent synchronization of IT, cloud transformation and your business.

## End-to-end analysis of applications and workload landscapes

Our consultants begin by working with customers on classifying the criticality of applications and workload landscapes that are moving to Google Cloud Platform and assigning service-level objectives (SLOs) to each service. The always-on methodology tackles services that have aggressive SLOs, such as 99.99% and 99.999% availability.

Next, we perform an end-to-end analysis to break workloads into layers, components, dependencies, and topologies. This process is crucial for documenting weaknesses and understanding if services are fault-tolerant at all levels.

## Google Cloud Platform services for your always-on architecture



### Google Cloud Spanner

Achieve regional consistency with a cloud-native service that provides unlimited scale, consistency, and 99.999% availability.



### Google Cloud Content Delivery Network (CDN) and Cloud Load Balancing

Distribute traffic and workloads across regions and ensure user proximity with high-performance, edge-caching infrastructure.



## Service mapping to enable active services with failure transparency

As workloads are migrated or transformed to Google Cloud Platform, the next critical step is mapping services required by each application to services offered by Google Cloud Platform. We assess and enable multi-active and geo-replication capabilities while modernizing applications for statelessness, autoscaling, and impact containment. This enablement helps in building parallel active services that can adopt or bypass failure transparency instead of disaster recovery.

Another vital aspect of application modernization is integration transformation, as well as the adoption of reactive and event-driven architectures that offer asynchronous reliability of data and application integration.

## Site Reliability Engineering and a cultural shift to always-on

After workloads are re-engineered and modernized to be always-on, the culture of your enterprise comes into play. It is important to align your operation teams toward the business services that require the highest levels of uptime and availability. By adopting Google's Site Reliability Engineering (SRE) principles, you can infuse observability and accountability into your organization and maintain the SLOs required for always-on workloads. In this culture shift, SRE assumes the responsibility for balancing error budgets with velocity and continuous releases that might impact availability and service uptime.

To be effective, SRE must also provide the confidence and evidence that always-on workloads can sustain failures and planned outages. To that end, it is important that reliability tests take place by injecting faults and failures and validating that the fallback mechanisms throughout all layers of the workloads are functioning properly. This practice, known as "chaos engineering," is an essential, culture-changing element that SRE must adopt for building and managing always-on workloads.

## Why Kyndryl?

Kyndryl has deep expertise in designing, running, and managing the most modern, efficient, and reliable technology infrastructure that the world depends on every day. We are deeply committed to advancing the critical infrastructure that powers human progress. We're building on our foundation of excellence by creating systems in new ways: bringing in the right partners, investing in our business, and working side by side with our customers to unlock potential.

## For more information

To learn more about how Kyndryl can help ensure your mission-critical workloads are always-on with Google Cloud Platform, please visit our website:

[kyndryl.com/about-us/alliances](https://kyndryl.com/about-us/alliances)



© Copyright Kyndryl, Inc. 2021

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.