

Cloud Workload Resilience

SPONSORED BY



AS ENTERPRISES CONTINUE to shift their workloads into the cloud, the notion of resiliency has taken center stage, supplementing traditional metrics such as uptime and mean time to restore. Resiliency is not merely an extension of these concepts; it represents a shift in our expectations from the infrastructure and applications that power digital businesses.

Resilience is a multifaceted attribute, indicative of the ability of cloud workloads to not only withstand but also adapt to unexpected or rapidly changing conditions, ensuring operational effectiveness through diverse challenges. High traffic volumes, cyberattacks, and unforeseen corner cases — the resilience of cloud infrastructures and applications is being tested in ways that demand innovative approaches to design, build, test, secure and manage cloud workloads. The resiliency of cloud workloads encompasses the robustness of the infrastructure as well as the adept management of applications, data, and their interfaces across dispersed cloud environments. It signifies the capacity of these systems to continue their operations and successfully complete processes even under new or extreme circumstances.

While not all cloud workloads will be modernized to utilize contemporary software architectures, cloud workloads benefit from using a cloud-native design approach built using containers, microservices, container orchestration and more adaptable cloud infrastructure. Using cloud native, applications are designed to be more easily scalable, and also more resilient under failure conditions.

How do you define cloud resiliency for cloud workloads?

There's a growing awareness of reactive measures and continuous service availability in resiliency for cloud workloads.

THE ABILITY TO RECOVER FROM FAILURES, HIGH LOADS AND CYBERATTACKS



20%

MAINTAINING UNINTERRUPTED SERVICE IN THE FACE OF SOFTWARE, INFRASTRUCTURE AND NETWORK FAILURES OR DISRUPTIONS



14%

CONTINUE SERVICING WORKLOAD REQUESTS DURING THE RECOVERY OF FAILED COMPONENTS OR SERVICES



13%

IMPLEMENTING SECURITY MEASURES TO PROTECT CLOUD WORKLOADS FROM CYBERSECURITY THREATS LIKE DDOS



12%

DESIGNING CLOUD WORKLOADS TO ADAPT TO SUDDEN CHANGES IN DEMAND AUTOMATICALLY



12%

USING GEOGRAPHICALLY DISTRIBUTED OR REDUNDANT CLOUD DATA CENTERS



9%

REGULARLY PERFORMING BACKUPS TO PREVENT DATA LOSS AND ENABLE QUICK RECOVERY IN CASE OF FAILURES



8%

DESIGNING CLOUD WORKLOADS SO THEY CAN MORE EASILY MOVE BETWEEN CLOUD LOCATIONS AND SERVICE PROVIDERS



8%

RESILIENCY IS JUST ANOTHER MARKETING BUZZWORD



4%

The resilience of cloud infrastructures and applications is being tested in ways that demand innovative approaches to design, build, test, secure and manage cloud workloads.

In late 2023, Techstrong Research polled our community of cloud, cloud-native, cybersecurity and infrastructure readers and viewers to take their pulse on cloud workload resilience. While there are multiple definitions of resilience, 33% of respondents defined resiliency as the ability to recover from failures, high loads and cyberattacks while continuing to service workload requests during the recovery of failed components or services. 61% of respondents said cloud workload resilience is in the top three priorities for their organization.

TECHSTRONG RESEARCH ANALYST VIEW

The industry's pivot towards resilience in cloud workloads reflects an understanding that operational continuity under even extreme operational conditions is paramount in today's complex digital environment. Organizations have transcended beyond the traditional metrics of uptime and MTTR to ensure robustness against diverse challenges like traffic surges, security threats, and previously unexperienced system anomalies. This change is propelled by the necessity for cloud services to remain reliable and responsive beyond operational norms, underpinning the critical operations of businesses in a non-stop global economy.

In what ways does cloud-native architecture improve the resiliency of workloads deployed in the cloud?

There is a preference toward smaller, self-contained microservices, as they don't impact overall availability. Containerized software is also improving the resiliency of cloud workloads with its scalability and ability to recover from failures.

SMALLER, SELF-CONTAINED MICROSERVICES FAIL INDEPENDENTLY WITHOUT IMPACTING OVERALL AVAILABILITY

25%

CONTAINERIZED SOFTWARE PROVIDES ISOLATION AND CONSISTENCY, MAKING IT EASIER TO SCALE AND RECOVER FROM FAILURE

22%

STATELESS DESIGN PATTERNS INCREASE SCALABILITY AND CAN FAIL INDEPENDENTLY WITHOUT IMPACTING OTHER PARTS OF CLOUD APPLICATIONS

19%

SERVERLESS DESIGN PATTERN ALLOWS EVENTS TO INITIATE THE OPERATION OF THE DISCRETE FUNCTIONS AS NEEDED

14%

CLOUD-NATIVE ARCHITECTURE SIGNIFICANTLY INFLUENCES THE RESILIENCY OF CLOUD-DEPLOYED APPLICATIONS

13%

CLOUD-NATIVE ARCHITECTURE PROVIDES LIMITED OR NO RESILIENCY IMPROVEMENT

6%

How important is improving the resiliency of cloud workloads for your organization in 2024?

With a majority of organizations prioritizing the resiliency of their cloud workloads in 2024, there is a growing consensus on the importance of its role in cloud strategies.

THE NUMBER ONE PRIORITY

24%

IN THE TOP THREE PRIORITIES

37%

IN THE TOP TEN PRIORITIES

22%

IMPORTANT BUT NOT A PRIORITY

11%

NOT IMPORTANT


6%






Which of the following do you use to improve the resiliency of cloud workloads?

Microservices are at the forefront of improving cloud workload resilience, however organizations are still taking a balanced, multifaceted approach.



What factors in cloud-native (microservices) workloads increase resiliency?

Workload distribution and portability are among the biggest factors in increasing resilience, making flexibility and resource optimization in cloud environments non-negotiable.




Cloud-native architectures are instrumental in this resilience-centric shift. Utilizing containers, microservices, serverless, and orchestrated automation, they allow for the creation of resilient systems with great scalability and increased tolerance under failure conditions. The flexible nature of modern cloud infrastructure complements this by providing dynamic, scalable resources that can be adjusted in real-time to meet changing demands. Stateless design patterns also increase resilience as they may independently fail, be restarted and even moved to other cloud resources without collateral damage to elements of the cloud workload. These resilient frameworks are designed to ensure performance and availability are maintained, even under demanding conditions.

Techstrong Research's data shows that the cloud native approach is one of several options to increasing cloud workload resilience. These include workload portability, distributing workloads, edge computing, Infrastructure as Code and failure testing. The greatest challenge to improving cloud workload resilience are the lack of skills and know-how, legacy or existing application architectures and large technical debt. The good news is these problems don't have to be tackled by one group or job function alone.

New roles such as Site Reliability Engineers (SREs) and platform engineers are relatively new options where resources can be directed toward improving resilience. SREs bring a software engineering discipline to system operations, crafting automated solutions for maintaining and scaling system reliability. Platform engineers focus on creating robust platforms that empower developers to create and deploy resilient applications efficiently. Together, these professionals embody the resilience ethos, integrating reliability from application code to the underlying cloud workload infrastructure. Additionally, organizations can work with cloud providers and their partners who can bring new learnings and approaches shown to have an impact on resilience, whether for new applications or existing.


What is the size of your business or organization?

Among the size categories, the smallest (1-100 employees) and largest (over 5000 employees) make up the majority of the respondent pool.



What are the most significant challenges to improving the resiliency of cloud workloads or cloud-native applications?

There is a critical skills gap in enhancing cloud-native applications, which poses the biggest challenge to improving the resiliency of cloud workloads. However, legacy architecture of applications and technical debt are also major hurdles for organizations.



How much of your cloud workload is cloud native?

54% of respondents have more than half of their cloud workloads as cloud native – 22% of which have all cloud-native workloads – underscoring the shift toward modern cloud practices.

