

ACHIEVE CLOUD RESILIENCE THROUGH SYSTEMATIC AND CHAOTIC TESTING



Modern cloud systems enabled companies to switch smoothly to work-from-home arrangements. Constantly testing the cloud in systematic and chaotic ways will help companies avoid cloud failures in a time when downtime means disconnect.

Professionals fortunate enough to shelter in place and work from home can thank cloud computing for their ability to shift over so seamlessly during the COVID-19 pandemic.

Imagine a lockdown without the cloud. Grocery and online order systems would be up and down, dependent on the servers at their own data centers. Streaming services would stutter and stall in the hour after dinnertime and other peak demand times. Critical personal and professional relationships would be limited to voice calls, rather than rich social media, video conferencing or collaboration apps. Many companies would simply just be closed to business.

Now imagine a COVID-induced work-from-home scenario with a cloud failure. Teams working remotely would fall apart. Deliveries would be scrambled and lost. More importantly, emergency services, already experiencing huge call volumes would see bigger surges with fewer tools to manage them. Medical research teams around the globe,

currently sharing research and data on the new coronavirus would lose contact with each other, delaying and slowing efforts to treat and cure the COVID-19 virus.

It may seem that the cloud was built just in time to help with our current predicament. But actually, the concept and foundational systems of cloud computing go back more than 50 years. The popular, practical application cloud tools developed more rapidly in the past 20 years, starting with the early iterations of Salesforce.com and more recently with web-based enterprise applications such as Google's G Suite and Microsoft's Office 365.¹ It's the simplicity of these relatively newer tools that have led to their broad adoption. But this surface layer masks a mind-bogglingly complex undergirding that must be fully understood and tested in order to be maintained.

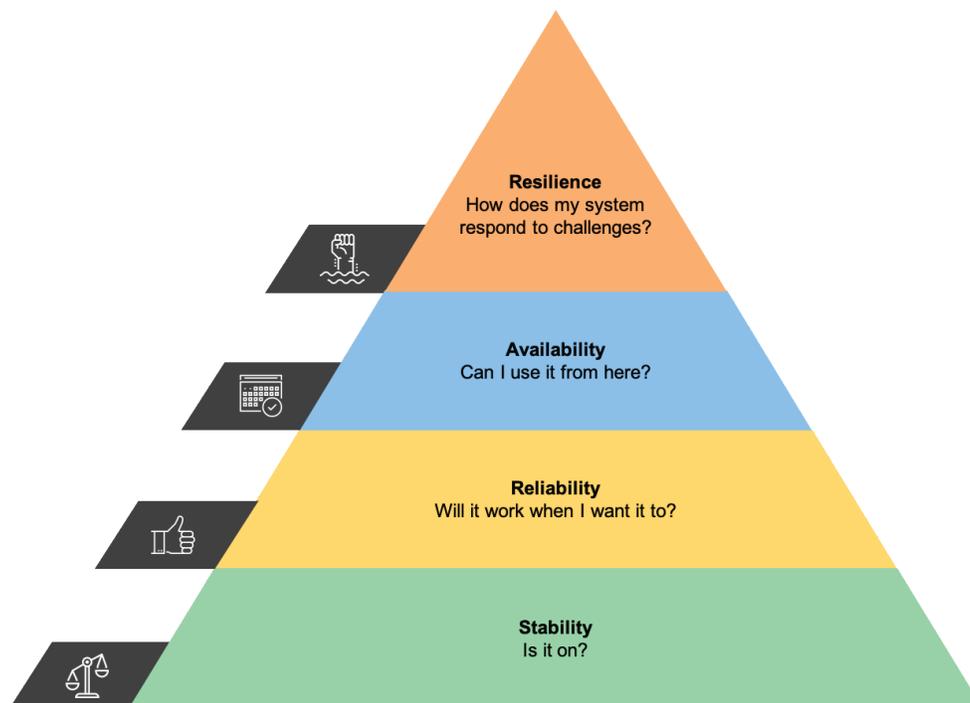
To avoid cloud failure, companies must build resilience into their cloud structure by testing it in continuous

and chaotic ways. Resilience represents a higher standard for computing systems – advancing beyond stability, availability and reliability.

Companies should progressively test their systems as they migrate to cloud services for stability, availability, reliability and, ultimately, resilience.

In the early days of business computing, IT managers aimed for stability, but computers and servers often crashed. So tech managers rebooted terminals or servers on a regular basis. Availability or system uptime could be managed by distributing applications on different servers in different locations and balancing the load between those. Reliability required the system to be both available and to function as expected. Resilience includes the expectation that something will go wrong, and that the system itself has

Figure 1. Resilience is at the top of the IT uptime hierarchy.



Source: Infosys Knowledge Institute

been structured and tested to respond and repair to what has gone wrong.² As companies migrate to cloud services, they should test their new system for all these attributes: stability, availability, reliability and resilience.

Conventional testing fundamentally ensures that key applications have been migrated or set up in the cloud and are working properly. This is a great and proven way to ensure that cloud systems meet design considerations and don't change how applications perform.

However, conventional testing does not probe for unexpected situations. To do that, advanced operations turn to chaos testing, a technique pioneered by Netflix during its own migration to the cloud in 2011.³ Chaos testing creates turbulent situations that will bring points of failure to light, and influence design.

Modern cloud systems bring hardware and software together in such complex and fluid ways that "check your work" testing will never suffice. By engaging in systematic and chaotic testing, companies can develop greater resilience in the cloud and through the entire technology system.

Clouds showing strain

The shift to work-from-home has fueled a massive upsurge in demand for cloud. In the April-June quarter, Microsoft reported 50% revenue growth from its Azure cloud products.⁴ Once businesses, industries and societies shift to cloud computing systems, cloud failure will create wide-ranging disruptions. To avoid this, companies migrating to the cloud must test the systems to their breaking points, and use the results of those tests to redefine and redesign their systems.

During a pandemic, the critical value of resilient cloud systems comes down to two points. First, they must operate

smoothly and glitch-free even when they receive an unexpected surge in online traffic. Second, the shift to most people working from home multiplies the number of endpoints outside the network firewall. A resilient and thoroughly tested system will be able to manage that extra congested traffic in a secure, seamless and stable way.

Clouds are strained right now. IBM endured a cloud outage of roughly two hours on June 9, 2020. Network monitoring company ThousandEyes told tech news site Fierce Telecom that the global nature of the outage suggested a control-plane issue rather than a physical failure such as a fiber cut or router failure.⁵

Microsoft's Azure Cloud Data Center capacity in Europe has shown signs of stress as companies shift to remote work and usage of its communication and collaboration platforms, including Teams balloons. To relieve that stress, Microsoft restricted access to free and trial accounts to limit the impact of performance problems for existing customers and ensure that emergency and critical services like health care receive priority.

Google Cloud's Kubernetes platform and networking services on parts of the US East Coast went dark for a few hours on June 29. The outages affect a range of services including Google Cloud Networking, Google Compute Engine and Kubernetes for varying periods from less than 90 minutes to four hours and 46 minutes.⁶

How testing leads to resilience

Chaos testing is not a substitute to actual system testing - however chaos helps uncover some system anomalies which normally do not manifest during development or testing. These tests introduce an anomaly at any of seven system layers to measure its impact on the system's resilience

as a whole. Following the success of Netflix's Chaos Monkey suite of tools, software engineers developed additional toolsets such as Gremlin, which can be leveraged for chaos engineering. Chaos testing requires careful planning and design, and must be conducted in coordination with the entire IT organization.

While monkeys and gremlins evoke a wild and uncontrolled element, the software tool versions can be controlled in a routine or planned situation such as gameday testing. Gameday testing involves simulating a COVID-like situation where suddenly 90% of the workforce is working from home or customers are all accessing the mobile app at the same time.

This involves the entire engineering team's participation and is usually attempted on a production-like environment where it is run at scale. Testers capture a postmortem report to document the learnings and review the behavior.

To achieve resilience, IT managers must find vulnerabilities not revealed in normal "happy path" testing scenarios. Seeking resilience provides a glimpse into a systems' performance and recoverability by executing a carefully planned set of disasters. Such a disaster sequence could include:

- a. Pulling down the network connection for 10 seconds
- b. Taking a service or server offline
- c. Choking middleware and watching for anomalies

Microsoft's Azure unit is testing its systems in all those fashions at small scale and large scale. This testing is so valuable to the unit that Chief Technology Officer Mark Russinovich wrote in his 2019 blog about the importance of bringing a new quality engineering team to work with the existing site reliability engineering team on testing for failures more rigorously and injecting faults to ensure system reliability.⁷

Figure 2. A sampling of cloud outages in the last 19 months.



Source: CRN.com, TechRadar, ThousandEyes.com



Russinovich further describes a vision to allow customers to leverage the mechanism of injecting failures into Azure and validating the resilience of their own setups. "Our plan is to eventually make these fault injection services available to customers so that they can perform the same validation on their own applications and services," he stated.

How to start testing

Enterprises just embarking on their migration to cloud, must first study and choose a service model that fits best. Firms looking to migrate a whole workload to the cloud should choose the Infrastructure-as-a-Service (IAAS) option. This allows team to quickly orchestrate their test environments on the cloud and quickly set up systems for storage, back-up and recovery.

For instance, a leading European Bank leveraged AWS to move 90% of its test environment to cloud for faster release cycles and to minimize wait times. In IAAS, while cloud service provider is responsible for functioning of cloud infrastructure, the firm is responsible for ensuring correct configuration of the service.

If an enterprise needs to create an application quickly without the effort required in managing the underlying infrastructure, it should choose the Platform-as-a-Service (PaaS) option. Firms leverage the PaaS option to build applications for mining data and developing analyses that provide insights. For instance, a leading fashion retailer leveraged Google Cloud Platform to build an analytical model that helps maintain optimal and just-in-time inventory.

Cloud migration benefits organizations in many ways, but companies must also consider what applications and systems should be retained on-premises or placed at the edge of the cloud.

To determine this, consider:

- a) Lifecycle of the application – those systems which are going to be sunset should be retained on-premises.
- b) Ring-fenced applications and data required by regulations or for compliance should be retained on-premises.

For instance, a multi-national insurer modernizing its tech landscape by re-hosting applications and systems

on Cloud. It retained old insurance systems on-premises where it maintained existing products. New onboarding of customers and products happened directly on the cloud.

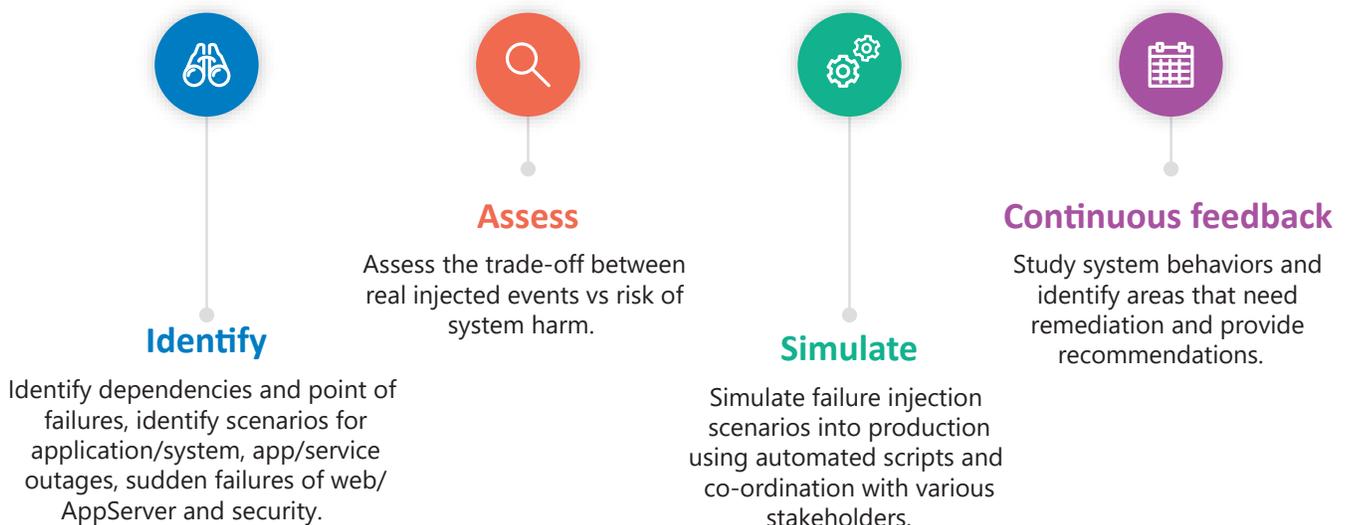
Operating in the cloud like many modern business practices is a cycle, not a linear progression.

Organizations already in the cloud or partly migrated, must take time to understand their cloud service models. This could prompt the question: Is this model right for my organization in the years ahead?

While operating in the cloud, these companies can start planning gameday or chaos testing scenarios to gauge their systems' resiliency. In doing so, companies should set an objective and time window for running the tests. Then, results of these tests help inform the requirement design phase of their cloud setup.

Chaos testing focuses on breaking the right (inducing failures on production or production like systems) to design a better left (infrastructure environment). Cloud services have allowed countless workforces to shift

Figure 3. Four steps for robust cloud testing.



Source: Infosys Knowledge Institute

from the office to other arrangements for the time being. Well-tested cloud computing systems hold many permanent benefits stemming from the very nature of their resilience.

With all systems and applications in a virtual cloud context, future digital transformations will become more fluid. Systems that cannot fail such as remote monitoring, augmented reality, virtual reality and geo-fencing will rely on robust cloud systems for delivery.

Properly tested clouds also handle demand spikes more readily, and will handle the viral adoption of an

application or a rapid location switch or expansion of workforces.

Using chaos and site reliability engineering delivers resiliences through the enterprise in the forms of:

- Cloud and infrastructure resilience.
- Data resilience via continuous monitoring.
- Resilient cybersecurity by integrating security with governance and control mechanisms.
- Presentation layer resilience by ensuring user interfaces hold up under high stress conditions.

Constant, systematic, and chaotic testing increases the resilience of cloud infrastructure

To become resilient, companies must create resilient IT systems. These systems will rely partly or wholly on cloud infrastructure. Consistent, creative testing reveals the true state of cloud systems and shows how they can be made better.



References

- ¹ [A history of cloud computing](#), Arif Mohamed, *Computer Weekly*, April 9, 2018.
- ² [Availability vs Reliability vs Durability vs Resiliency](#), Aditya, *Medium*, April 23, 2020.
- ³ [The Netflix Simian Army](#), *Netflix Technology Blog*, July 19, 2011.
- ⁴ [Microsoft Cloud Strength Drives Fourth Quarter Results](#), *Microsoft Corp.*, July 22, 2020.
- ⁵ [IBM Cloud suffers widespread outage on Tuesday, but services are now up and running](#), Mike Robuck, *Fierce Telecom*, June 10, 2020.
- ⁶ [Google Cloud partially evaporates for hours amid power supply failure: Two US East Coast zones rattled](#), *The Register*, June 29, 2020.
- ⁷ [Advancing Microsoft Azure reliability](#), *Microsoft Azure blog*, Mark Russinovich, July 15, 2019.

Authors

Venkatesh N Iyengar

*AVP & Group Practice Engagement Manager,
Infosys Validation Solutions
nvenkatesh@infosys.com*

Shishank Gupta

*Service Offering Head,
Infosys Validation Solutions
shishankg@infosys.com*

Sundar Gomathi Vallabhan

*Practice Engagement Manager,
Infosys Validation Solutions
ssubramanian_gv@infosys.com*

Arvind Sundarraman

*Associate Practice Engagement Manager,
Infosys Validation Solutions
arvind01@infosys.com*

Chad Watt

*Infosys Knowledge Institute
chad.watt@infosys.com*

About Infosys Knowledge Institute

The Infosys Knowledge Institute helps industry leaders develop a deeper understanding of business and technology trends through compelling thought leadership. Our researchers and subject matter experts provide a fact base that aids decision making on critical business and technology issues.

To view our research, visit Infosys Knowledge Institute at infosys.com/IKI

For more information, contact askus@infosys.com



© 2020 Infosys Limited, Bengaluru, India. All Rights Reserved. Infosys believes the information in this document is accurate as of its publication date; such information is subject to change without notice. Infosys acknowledges the proprietary rights of other companies to the trademarks, product names and such other intellectual property rights mentioned in this document. Except as expressly permitted, neither this documentation nor any part of it may be reproduced, stored in a retrieval system, or transmitted in any form or by any means, electronic, mechanical, printing, photocopying, recording or otherwise, without the prior permission of Infosys Limited and/ or any named intellectual property rights holders under this document.