McKinsey & Company

Cybersecurity and Risk & Resilience Practices

A technology survival guide for resilience

Resilience means understanding the criticality of a business process, the capability of the underlying technology, the business impact if the technology fails, and the organization's risk tolerance.

by Jim Boehm, Wolfram Salmanian, and Daniel Wallance

March 2023

It's no secret that in highly competitive business environments, the demand for organizations to grow and increase revenue and profit continues to rise. While meeting the demand and staying current through digitalization, organizations must remain mindful to be efficient, maintain or reduce costs, and keep employee spending in line.

Moving forward in those two areas is difficult enough, but moving in those directions adds stress on corporate technology systems across the technology stack, from data to applications and network infrastructure. Technology constraints include capacity limitations, system uptime, data quality, and the ability to recover from a catastrophic technological, physical, or cyber event.

Resilient technology is critical in maintaining uninterrupted services for customers and servicing them during peak times. This requires a resilient infrastructure with heightened visibility and transparency across the technology stack to keep an organization functioning in the event of a cyberattack, data corruption, catastrophic system failure, or other types of incidents.

Resilient technology needs to be agile, scalable, flexible, recoverable, and interoperable. In addition, resilience needs to exist not only in the architecture and design but also through deployment and ongoing monitoring.

Understanding criticality

To achieve resilience, an organization needs to understand the criticality of a given process, evaluate the underlying technology, recognize the corresponding business impact, and know the risk tolerance of the organization and external stakeholders. To get there, an organization needs to understand where and what its resilience is today and be able to answer the question: Could we recover and rebuild after a catastrophic event?

In a 2022 McKinsey survey on technology resilience that assessed the cybersecurity maturity level of more than 50 leading organizations across North America, Europe, and other developed markets, 10 percent of respondents indicated they have been forced to rebuild from bare metal (for example, due to a catastrophic event), with 2 percent stating that they have already attempted to recover from bare metal but were unsuccessful (for example, deliberate testing).

Additionally, 20 percent of respondents indicated they had already attempted to recover from bare metal and were successful, 8 percent attempted to recover from bare metal, 18 percent noted they had plans to attempt to recover from bare metal, while 36 percent stated there were no plans to recover from bare metal.

Technology resilience is the sum of practices and foundations necessary to architect and deploy technology safely across the technology stack (see sidebar "McKinsey technology resilience principles"). Technology resilience prepares organizations to overcome challenges when their technology stack is compromised, reducing the frequency of catastrophic events and enabling them to recover faster in the case of an event.

In the McKinsey survey, when asked what the recovery time objective was for their highest critical applications, 28 percent of respondents said immediate, while 34 percent said it was less than an hour, 14 percent said less than two hours, and 20 percent said less than four hours. One of the respondents in the survey stated, "Critical systems and applications down for a significant amount of time can cost financial institutions billions of dollars."

Resilience capabilities fall on a maturity spectrum from simple redundancy to duplicate servers through to advanced capabilities with resilience built into architecture by design.

- *Architecture and design:* Mature organizations incorporate technology resilience into enterprise design and architecture. Resilient designs incorporate elements of lessons learned from operations, incidents, and industry trends to make risk-informed technology investments.
- *Deployment and operations:* Resilient operations should consider not only operational contingencies, such as disaster recovery or performance demands that increase exponentially, but also the root cause of

McKinsey technology resilience principles

The following are five principles that we see as foundational for maintaining resilient technology:

— Applications, systems, platforms, and the IT workforce itself are flexible and scalable. On an ad hoc basis, the enterprise can scale up or down services to support changing availability, capacity, or performance demands as business requirements shift.

— Data sets, applications, and network technology infrastructure are fully visible to owners of data and applications and are traceable within the environment. Owners are empowered to raise problems and prevent outages before they occur.

— Data sets and applications are built to be agile and mobile. They must not be tied to a single platform or environment but rather can be rapidly moved between or across platforms as needed.

— The architecture of applications, data platforms, network environment, and the IT workforce is resilient by design—that is, the architecture was built to compensate for probable failures (at lower maturities) and recursively inform future designs (at higher maturities).

— Systems are interoperable and leverage standard API schemas that are defined and well-architected both internally and between and among third-party services.

incidents that arise during business as usual to improve procedures, training, and technology solutions.

— *Monitoring and validation:* This consists of reactive or backward-looking metrics at lower maturity levels. At higher maturity levels, organizations shift to more proactive (and ultimately predictive) measures to stress-test solutions prior to rollout or drill preplanned responses and contingency plans for the most likely eventualities.

— *Response and recovery:* Organizations with high technology resilience not only respond as incidents occur but they also continuously feed lessons from their own operations, industry trends, and catastrophic events back into the design, operation, monitoring, and planning for their enterprises.

Understanding the components behind the life cycle allows an organization to chart what its technology resilience journey looks like through four maturity levels. Levels one and two are foundational capabilities, while levels three and four are more advanced (Exhibit 1).

Level one consists of basic capabilities where resilience is left to individual users and system owners, and monitoring involves users and customers reporting system outages.

Level two consists of passive capabilities where resilience is through manual backups, duplicate systems, and daily data replication. There is also monitoring at the platform or data center level for system outages.

Level three consists of active resilience through failover. Resilience exists through active synchronization of applications, systems, and databases, and active monitoring at the application level for early indicators of performance and stability issues.

Level four consists of inherent resilience by design. Resilience is architected into the technology stack from the start through inherent redundancy and active monitoring at the data level, which includes anomaly detection and mitigation.

From a life cycle standpoint, the range for architecture and design goes from limited visibility

Exhibit 1

A technology resilience journey is one of evolving complexity and maturity.

The resilience journey by level

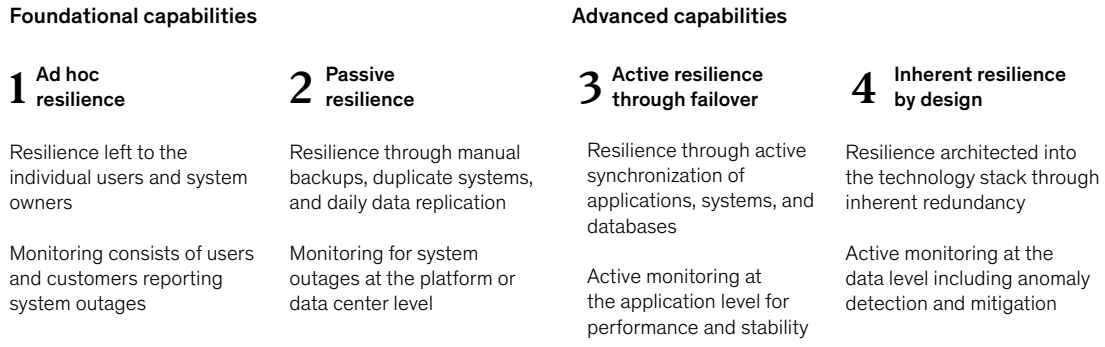

McKinsey & Company

of dependencies for critical and noncritical applications in level one, to dependencies and data flows built in for resilience from initial design for critical and noncritical apps in level four.

For deployment and operations, regular system outages in level one take the place of resilience tests, and in level four, random, in-production failover tests validate resiliency.

In the case of monitoring and validation, in level one, users monitor their own systems for outages, whereas in level four, monitoring and alerting is built in by design, allowing for proactive response.

For response and recovery, responses to incidents in level one are ad hoc and based on best judgment, while in level four, detailed and diverse "break glass" procedures are drilled in by design.

Resilience spectrum

At the most basic level, resilience is left to the individual system owners and users. The database administrator is responsible for backups of organizational data, and individual employees must back up their own data. Moving along the maturity scale, organizations rely on centralized resilience capabilities managed by IT or a resilience function. Such an organization provides for centralized backup solutions, maintains redundant core systems, and monitors for system outages and application failures.

Resilience can be achieved passively by conducting manual backups daily. Shifting to an active approach involves monitoring for early indicators of data corruption or anomalous system behavior and taking preemptive action. Those indicators include an increasing volume of corrupt data, an unusually high number of brief network outages, and a greater than usual number of servers that require reboots. Active resilience further occurs through the continual synchronization of applications, systems, and databases such that redundancy is always maintained. Periodic failover tests are also conducted to validate resilience.

The most advanced level of resilience consists of inherent resilience. The primary differentiator is that resilience is built into the technology stack by design. Inherent resilience includes capabilities such as duplicate processing across systems, modular redundancy, and automatic fault tolerance within systems. True inherent redundancy enables the ability to conduct random in-production failover tests to validate resiliency. Only the technology that enables an organization's most critical business processes needs to be inherently resilient by design.

Most organizations fall within the passive-to-active resilience capability spectrum while making a continual shift toward active resilience.

How to become resilient

It's one thing to lay the groundwork and point out the issues behind resiliency, but just how does one get there? There are three keys to establishing and growing a more resilient technology environment:

1. *Blame-free culture:* When problems arise, teams and managers don't look for whom to blame. They focus on fixing the problem and preventing recurrences. Teams celebrate members who expose vulnerabilities and weaknesses as necessary to build more resilient technology.

2. *Metric-driven approach:* Teams relentlessly measure their own performance and focus on which incidents they created (for example, from releases or patches) or repeat incidents that have the same root cause.

3. *Rehearse the outage:* Teams anticipate problems and iteratively build up and train to respond to complete system outages. They build from individual applications to systems to products (systems of systems) to entire services.

When asked in the McKinsey survey how often they test critical applications, slightly more than 60 percent of respondents said they tested at least quarterly. Of those, 14 percent said they tested weekly, 26 percent test monthly, and 26 percent test quarterly. Overall, 28 percent said they test every six months, while 6 percent indicated they test annually. One respondent said, "There are quarterly tests. The most critical systems will be tested each time, less critical systems are spread out to every other test cycle or annual at a minimum."

Risk-based resilience

Companies are moving to risk-based technology resilience (see sidebar "A European bank works toward technology resilience"). The approach recognizes that not all assets are created equal, nor can they be equally protected in today's all-encompassing digital environment.

Some capabilities and underlying assets are more critical to a company and its business than others. In the case of a large electric utility, for example, these include the technology systems that enable the delivery of electricity and natural gas to customers. In the case of a global financial-services institution, the trading platforms and those that support customer transactions are most critical. The digital business model is, in fact, entirely dependent on trust and the ability to continuously provide customer-facing services. Ensuring resilience over those assets is at the heart of an effective strategy to protect against catastrophic events.

Three levers to build technology resilience

Reaching high maturity levels of technology resilience requires building the necessary capabilities and processes, using three levers as guidance.

1. *Prioritize services:* Not all business services and systems should be treated equally when deploying technology resilience capabilities. Rather, organizations should define their most critical services. These comprise the crucial services needed to fulfill obligations to customers, business partners, regulators, and society.

   After identifying and obtaining cross-business agreement on these services, understanding the underlying technology landscape is essential, including which applications and systems enable the most critical business services, their dependencies, and how they are interconnected.

   Having visibility and transparency into the most critical services and underlying applications, systems, and dependencies allows for assessing the current resiliency level and prioritizing the target resiliency on an application-by-application and system-by-system basis.

   In the McKinsey study on resilience, respondents were asked, "How long did it take

A European bank works toward technology resilience

Understanding technology resilience is an ongoing process, and by employing the three levers—prioritization, assessments, and remediation—organizations can find success. When it comes to technology resiliency, one European bank with traditional data centers recognized it needed to understand its deficiencies to be able to withstand any type of incident it might face, whether technologically based or cyber-related.

Regulatory findings and recent crises such as the COVID-19 pandemic, geopolitical conflicts, energy crises, and flood risks led management to evaluate and strengthen its technology resilience and crisis capabilities.

Understanding that its technology landscape consisted of a mainframe and server environment that was largely on premises in data centers, the bank analyzed how it could enhance resiliency—particularly by leveraging the cloud for out-of-region recovery—and flexible scaling of resources and related services.

The bank included cybersecurity and data privacy efforts to harmonize application and infrastructure requirements as one of the key levers for efficient implementation.

An Asian fintech leverages cloud for resilience

A fintech with a cloud-only infrastructure landscape launched its business and immediately faced security, performance, and scalability challenges. Through a review of the cloud configuration, the fintech identified gaps and set a path to enhance its resiliency. This was done primarily via the setup of regions/availability zones, load balancing, data mirroring, and snapshots/backups combined with testing. The initiative enabled continuous service delivery in the face of outages and cyberattacks and supported the hypergrowth of customers using its services.

Oil and gas company transforms in face of acute threats

A large oil and gas provider faced frequent cyberattacks, even as it undertook a digital transformation that had the potential to increase the exposure of its critical systems. A successful attack on its assets had the potential to affect the economy of an entire nation.

The organization started by identifying and protecting its "crown jewels," its most important assets, via a library of controls. This was supported by building capabilities and addressing silos (for example, between information technology and operational technology capabilities). The organization outlined and implemented its plan for a holistic cybersecurity transformation, including a three-year implementation program with prioritized initiatives, an estimated budget, and provisions to integrate technology resilience in its digitization effort.

you to get all your highest critical applications in line with recovery time objectives?" Here, 26 percent of respondents said less than a year, while 28 percent said less than two years, and 26 percent said less than three years.

One survey respondent said, "Being clear on which systems are most critical is an ongoing challenge." While another said, "It was during Superstorm Sandy that the bank became very concerned about its robustness, or lack thereof, and this became front and center immediately afterward."

2. *Assess current level of resilience and review past crises:* The next step involves assessing existing technology resilience. Organizations should assess their maturity along the same S-curve of technology resilience, whether they have resilient architecture and capabilities, passive resilience capabilities, active resilience with failover capabilities, or are inherently resilient by design.

Typically, organizations should assess current capabilities across the four dimensions in the technology resilience life cycle. The most mature organizations incorporate technology resilience into application and system architecture by design. In deployment and operations, resilient operations should consider not only operational contingencies but also the root cause of incidents that arise during business as usual to improve procedures, training, and technology solutions. Monitoring and validation involves reactive or backward-looking metrics at lower maturity levels. At higher maturity levels, organizations shift to proactive measures to look

for early indicators of resilience issues and test responses and contingency plans for the most likely eventualities. In response and recovery, organizations with high technology resilience not only respond as incidents occur but they also continuously learn from their own operations, industry trends, and catastrophic events and then feed that back into technology design, operation, monitoring, and planning.

Organizations should also assess past technology-related incidents to identify and uncover common contributing factors that can be addressed to increase technology resilience. Typically, this consists of selecting a broad set of recent incidents of varying duration and impact across business functions to evaluate. It can also include reviewing past incident-response logs, incident reports, and other documents to identify contributing factors, patterns, and insights that can shed light on causes behind the incidents. Meeting with engineers, product or system owners, release managers, and others involved in the incident and response can uncover what happened, what could have been done to prevent the incident, and initiatives that are already under way.

Once completed, it's then possible to identify and ultimately remediate common factors that led to these incidents, which may include the technology environment itself, the architecture of applications, interfaces between systems and third parties, and the way resilience was built into individual applications and systems.

3. *Remediate gaps through cross-functional approach:* Achieving technology resilience requires remediating gaps identified from the assessment of the organization's technology and diagnostic of past incidents. In addition to directly remediating the gaps identified, organizations should take the following specific steps:

   Determine ownership and accountability of technology resiliency activities. Distributed systems can have multiple owners, and developers aren't always incentivized to architect and design for resilience. Applications and systems must have clear ownership, developers need incentives with performance goals tied to the resilience of the applications they build, and third-party contracts must include resilience requirements and clauses. The absence of clear system ownership and responsibility to remediate gaps will adversely affect the resilience of systems and business processes.

   Enhance governance toward resiliency levels. Oversight of resilience must be implemented from the executive level on down. The C-suite needs to communicate its intention and prioritization of resilience down through all levels of the organization with continuous and consistent messaging. Town halls, quarterly newsletters, and webinars are all potential avenues. Likewise, awards and other forms of monetary and nonmonetary incentives may be considered.

   Increase resilience of individual applications and application groups. The resilience of individual applications and systems also needs to be addressed and remediated. Those that have the highest number of incidents and support the most critical business processes need to be prioritized for remediation.

   Strengthen the hosting setup, whether on premises or on cloud. The underlying platforms on which applications reside also need to be designed and architected for resilience. Organizations should work to increase the resilience of their on-premises and cloud platforms through remediating known gaps and addressing contributing factors from past incidents.

   Work with third parties to increase the resilience of third-party platforms on which critical business processes and services depend. There could be incentives for third parties to build resilience into their systems, and contracts must have clear language on performance requirements for resilience.

   Implement regular testing, with a focus on automatic failover capabilities for large-scale environments and selective exercises for testing recovery from backups. Resilience is a continual journey, and systems must be regularly tested

and validated to ensure they meet resiliency requirements. Monthly failover testing of business-critical applications is essential both at the application and platform level. Failover tests should be designed to test not just the expected but also the unexpected, such as through hard shutdowns or introduction of capacity surges that mirror real scenarios. Where resilience is built in by design, applications should be randomly shut off in production to test whether inherent resilience is truly architected and built into the application or system.

In the McKinsey survey, when asked what failover scenarios respondents planned or tested, 92 percent said they tested for a single data center failure and for nonphysical impact, while 52 percent said a dual data center failure, and 83 percent said physical impact (Exhibit 2).

When asked, "Do you run unplanned failover testing" (that is, randomly shut off systems and test the organization's ability to respond/recover), 54 percent said none, while 26 percent said most critical applications only, and 20 percent said they test for all applications (Exhibit 3).

The journey to technology resilience in three steps

With an understanding of the three levers to technology resilience, an organization can embark on its technology journey in three steps.

Technology resilience diagnostic

Identify two to three critical business processes and map the underlying data sets, applications, and technology systems that enable the processes. Evaluate the resilience of each component of the value stream. This will lead to uncovering the technology resilience of the data, applications, and systems that underpin critical business processes along with risk-mitigating actions.

Conduct an incident retrospective

Conduct a retrospective on recent technology-related incidents to identify common contributing factors and develop remediation actions to decrease the incident rate and increase the resilience of the technology environment. Interview developers, release engineers, and others involved with the incidents to uncover contributing factors and what could have been done to prevent them. The result will provide a stronger perspective on

Exhibit 2

Single data center failure and nonphysical and physical impact are top of mind in failover testing and planning.

Scenarios planned and tested in failover,[1] % of respondents

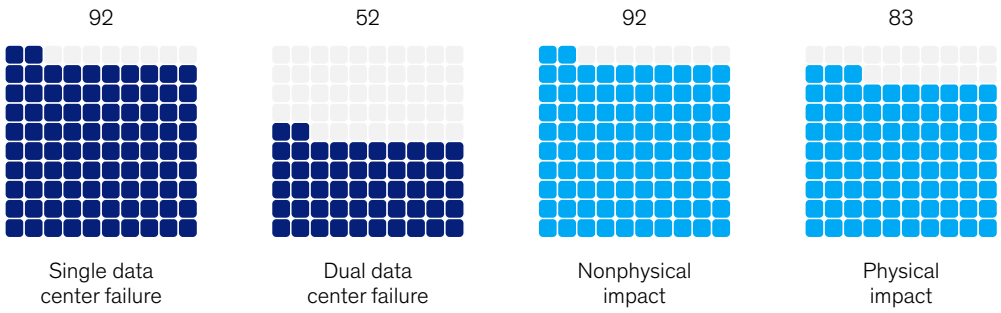

[1]Question: What scenarios did you plan/test failover for?
Source: Survey of BCM leaders at top banks (pilot, n = 50)

McKinsey & Company

Exhibit 3

More than half of survey respondents say they do not perform random failover testing, while only one in five tests all applications.

Unplanned failover testing, by type,[1] % of respondents

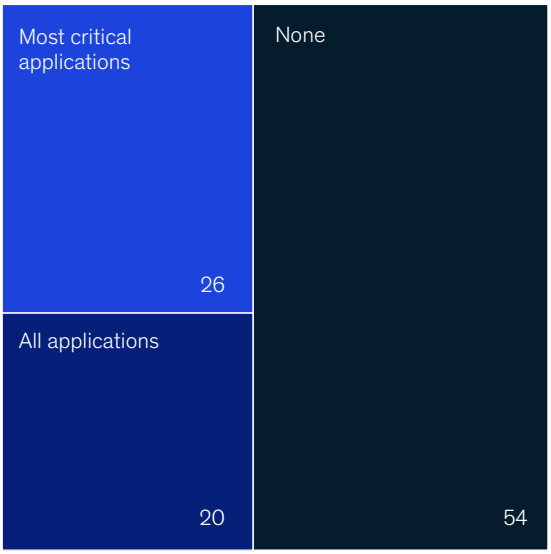

[1]Question: Do you run unplanned failover testing (eg, randomly shut off systems and test the organization's ability to respond/recover)?
Source: Survey of BCM leaders of top banks (pilot, n = 50)

McKinsey & Company

contributing factors that led to the incidents and actions that can be taken to decrease the incident rate and increase technology resilience.

Develop a redundant technology capability
Design a resilient architecture for one or more components of the technology stack and a future-state technology architecture to address the previous diagnostic and incident retrospective. These capabilities should include a transition and implementation plan and requirements for ongoing monitoring, maintenance, and validation. The result should be a resilient technology architecture, transition, and implementation plan along with monitoring and validation requirements.

Achieving resilience is not a one-time activity; rather, it's an ongoing process and capability that will take time to evolve into a solid defense mechanism.

As with all types of protection, it's not "you get what you pay for" but rather "you get what you prepare for." It would be easy to throw money at all forms of resilience, but understanding what you possess and then having visibility and transparency into what you have will bring focus, allowing any organization to remain resilient and either stay up and running or get back to a steady state as soon as possible.

Jim Boehm is a partner in McKinsey's Washington, DC, office, **Wolfram Salmanian** is a consultant in the Munich office, and **Daniel Wallance** is an associate partner in the New York office.