

INDUSTRY REPORT · OCTOBER 2026

The State of Resilience in Agentic AI Adoption

AI changes the operating model. Resilience assumptions need to change with it.

What resilience leaders revealed about critical AI dependencies, recovery decisions, and operational resilience.

Executive Summary

The adoption of AI across the financial services sector, is one of the most significant technological developments that this generation has seen in industry. With significant rewards promised, comes significant change to operating models, workforce construct, the vendor ecosystem and ultimately the institutions operational resilience posture.

Regulators worldwide have recognized this step change and have subsequently released policies to mature the governance and the safe operationalization of AI in critical services. A common thread of intent across these policies, is the need to strategically design for operational resilience from the outset, not after a significant risk as materialized.

To support financial institutions strategically assess the broader resilience considerations associated with the deployment of agentic AI, iluminr launched an industry wide agentic AI microsimulation for resilience professionals to respond to. Through a severe but plausible event, participants were required to assess the impacts and recovery options available to institutions in the context of both agentic AI deployment, and a workforce reduction.

The findings are eye opening.

Marcus Vaughan

Cofounder and Chief Growth Officer, iluminr

Three clear insights appeared:

1 Confidence in maintaining minimum service levels in a critical service, is low.

Only 16% of resilience professionals could demonstrate confidence in maintaining minimum service levels in critical service delivery after the loss of an agentic AI vendor.

2 People strategies are the contingency strategy of choice after AI fails.

In attempt to recover critical services, people largely turned to people, as the most reliable option of contingency. This elevates the criticality of workforce considerations in resilience alongside AI deployment.

39% of participants selected workforce solutions to recover.

3 Design for resilience, not recovery.

Maintaining a resilient delivery of critical services to customers, came down to decisions made at time of AI deployment, not at the time of the incident.

1 scenario. 7 decisions. A clearer view of operational resilience.

In September 2026, iluminr invited resilience professionals to navigate the failure of a critical agentic AI provider supporting a bank's fraud verification and sanctions screening operations.

Participants faced 2 video updates and 7 decision points, each designed to reveal how they would assess the disruption, engage stakeholders, and prioritize recovery.

The simulation captured individual decisions, supporting rationale, and confidence, providing a view of how practitioners approached the same operational problem.

Who took part. 31 practitioners from 25 organizations in 6 countries, most from banking and insurance. They work in the roles that keep critical services running and step in when they fail: operational resilience, business continuity, crisis and incident management, operational risk and third-party risk. Almost half are director level or above. Each answered in a personal capacity, drawing on their own professional judgment, not on behalf of their organization.

Scenario Overview

An agentic AI vendor running fraud verification and sanctions checks goes into administration, a year after the bank cut headcount in those processes by 70%.

2 video updates, 7 decision points, every decision captured.

3 clear objectives

Assess the stakeholders required to respond during a critical service disruption involving agentic AI.

Partially achieved

Assess the robustness of contingency strategies during a critical agentic AI vendor outage.


Partially achieved

Assess the operational resilience of agentic AI Governance during a critical event

Partially achieved

Partially achieved The cohort addressed the objective, but responses showed gaps in coverage or wide variation in approach.

Cohort operational resilience at each decision point



Readiness: how effectively participants demonstrate the judgment and capability required to respond, scored by iluminr AI from the decisions made in the session.

AI adoption is creating new resilience dependencies

As organizations delegate more operational activity to AI agents, they are also changing the conditions under which critical services must be maintained.

Workforce structures are evolving. Third-party dependencies are increasing. Processes that once relied on human judgment and intervention are becoming increasingly automated.

These changes have consequences for operational resilience, particularly when the assumptions underlying continuity and recovery arrangements were established before AI deployment.

Regulators are beginning to address these dependencies more explicitly. Recent guidance from APRA, the Monetary Authority of Singapore and Canada's Office of the Superintendent of Financial Institutions emphasizes the importance of credible contingency arrangements, third-party risk management, and testing whether recovery options are operationally feasible.

Agent Down was designed to examine how resilience professionals would navigate these challenges when a critical AI provider failed.

The findings offer an early indication of how practitioners assess emerging dependencies, where recovery approaches diverge, and which strategic decisions deserve greater scrutiny before deployment.

Regulatory notices on AI Risk

CBI: 26 February 2026

OCC / FDIC: 17 April 2026

APRA: 30 April 2026

NYDFS: 21 May 2026

FCA: 6 July 2026

OSFI: 13 July 2026

MAS: 7 October 2026

4 findings that deserve closer attention

CONFIDENCE

Limited

16%

Recovery remains uncertain

Only 16% were confident they could maintain minimum service levels within impact tolerances. No participant reported being very confident. The responses suggest considerable uncertainty about whether existing recovery arrangements would be sufficient.

READINESS

Moderate

39%

The human fallback has changed

While only 13% of participants selected rehiring as their preferred recovery strategy, 8 of the 10 participants who proposed an alternative strategy relied on some form of workforce intervention. When analyzing the response data holistically, human capacity remained central to recovery, even after the operating model had been redesigned around AI.

ALIGNMENT

Strong

77%

Productivity changes the exposure

77% anticipated breaching impact tolerances sooner following the loss of the AI agent. Several participants questioned whether existing thresholds adequately accounted for contingency capacity and changes in throughput. The findings point to a need to revisit how impact tolerances are established as operating models evolve.

BLIND SPOT

Elevated

5 of 31

The original decision matters

Only 5 participants explicitly connected the disruption to the earlier workforce scale back decision (12 month prior to the incident). Most concentrated on managing the immediate incident. The broader question is whether resilience considerations are sufficiently integrated into decisions about AI investment, organizational design, and workforce transformation.

WHAT MAKES THESE FINDINGS IMPORTANT

Every participant encountered the same scenario. Their decisions, reasoning, and levels of confidence provide a view of where professional judgment converges, where it differs, and which capabilities warrant further testing.

Regulators expect the fallback to work

REGULATION	WHAT IT EXPECTS	WHAT AGENT DOWN SHOWED
APRA (Australia) 30 April 2026 Letter to industry on AI, setting out findings from its review of AI use by regulated entities.	"Where AI supports critical operations, credible fallback processes are required."	Only 16% were confident of holding minimum service levels. The fallback most participants reached for was the workforce that had been cut.
FCA Mills Review (UK) 6 July 2026 Review into AI and the future of retail financial services.	Agentic AI moves firms from assistance to delegation. Staff shift from decision-makers to supervisors who monitor and intervene.	The scenario's 70% headcount cut is that shift. When the agent failed, the remaining team could not absorb the volume. A select few participants flagged the importance of resilience being considered at the point of AI deployment and workforce augmentation.
OSFI (Canada) 13 July 2026 Technology risk bulletin on generative and agentic AI, aligned to Guidelines B-13, E-21 and B-10.	Institutions can "test AI failure and outage scenarios and establish manual fallbacks and continuity measures for AI-supported business processes".	The manual fallback depended on people, and 70% of them were gone. 8 of the 10 participants who wrote their own recovery strategy built it on people.
MAS (Singapore) 7 October 2026 Guidelines on AI risk management for financial institutions.	Contingency plans for "potential failures, unexpected behaviour of third-party AI, or discontinuing of support by vendors".	The scenario trigger was exactly this: the vendor went into administration. No shared recovery strategy emerged. Low confidence levels of maintaining Minimum Service Levels may reflect a gap in testing for this scenario context.

Sources: APRA letter to industry; FCA Mills Review; OSFI technology risk bulletin; MAS AI Risk Management Guidelines, para 5.11(e). Agent Down responses, n=31.

THE SCENARIO

When a critical agent fails, what capacity remains?

1 The agent


An agentic AI platform performs fraud verification and sanctions checks across the bank's transactions. It was one of the earliest agentic AI success stories in a critical service, performing 300% above the prior workforce functions.

2 The workforce decision

12 months earlier, headcount for those processes was cut by about 70%. The bank was an early mover, and the decision was contentious.

3 The failure

Tuesday, 8:00am. The vendor is placed into administration. The application goes offline and the vendor stops responding.



ON THE INCIDENT CALL

Incident Team Leader	AI Technology Lead
SVP Human Resources	VP Commercial
VP Risk	VP Third-Party Risk

Each participant joins as the operational resilience specialist and makes every call alone.

2 video updates, 7 decision points

08:00

Application offline

Vendor placed into administration. No response from relationship managers.

08:30

Video update 1: incident call

MIM team activated. HR confirms headcount for these processes was cut 70%.

Team to reconvene in 30 minutes.

09:30

Video update 2: no recovery coming

Thresholds hours away. Burnout risk. Hard to rehire. A piloted alternate vendor lost on price. A BPO, with other banks competing for it.



Agent Down

Welcome to Agent Down, a forward-thinking microsimulation that explores the potential impacts an...

🕒 September 10, 2026 - October 6, 2026

OVERALL PROGRESS

32/45



● Completed ● In Progress ○ Not Started

DECISIONS 1-4

Assess

Who else to engage. Whether they are documented. Which impacts to prioritize. Whether 300% throughput means faster breaches, and why.

DECISIONS 5-7


Recover

Highest-priority recovery strategy, why, and its risks. Confidence in holding minimum service levels. 1 priority action.

The stakeholder gap extends beyond technology

4 → Priority 4: Understand the human dependencies

Named as a stakeholder to engage for Response



Are these stakeholders formally documented?



The stakeholder responses reveal an important disconnect. Although the scenario was shaped by a 70% workforce reduction, only 5 of 31 participants identified HR as an additional stakeholder. Just 1 identified the sanctions screening function affected by the outage.

The findings suggest that stakeholder identification may be oriented toward the immediate incident rather than the full operating dependencies involved in recovery.

Customer impacts dominated. Other exposures received less attention.

2 → Priority 2: Reassess impact tolerances

What are the key impacts you would be raising with the team to assess and prioritize?



Top impacts: % of participants citing



BARELY RAISED

2 of 31

Impact tolerances

4 of 31


Fraud exposure

Source: Agent Down Microsimulation responses, 10–30 September 2026, n=31. Percentages rounded.
Free-text impacts (Q22) coded to themes; word size reflects total mentions.

Have impact tolerances kept pace with AI-driven productivity?

2 → Priority 2: Reassess impact tolerances

Will 300% agent throughput breach thresholds faster?



What this tells us

Impact tolerances established around a predominantly human operating model may no longer reflect the volumes, dependencies, or recovery capacity of an AI-enabled operation.

As productivity changes, organizations need to understand whether the loss of automated capability would alter how quickly disruption becomes intolerable.

19% called out alternative positions, dependent on contingency and context.

"Impact cannot be measured by time alone, Contingencies must be taken into account"


"You have to consider if we lost that 300%, do we fall over, or do we fall back to normal thresholds - if we fall over, then we're into time to recover, regulatory timeframes, BC deployment. If it's above threshold, then we back to our normal benchmark - so it's monitoring will the disruption have throw any curve balls that will drop capability below that benchmark."

Recovery decisions revealed no clear consensus

3 → Priority 3: Establish recovery decisions in advance

Three key response strategies are emerging for consideration. Which of these would you suggest receives the highest priority?

Highest-priority recovery strategy



BPO

"We likely already have a relationship with this type of organization, meaning the due diligence has been already done."

Alternate agentic AI

"The firm has already assessed the alternative supplier, who came out strong. The only downside was the commercials."

Rehire

"The organization will have level of direct control and quality over the remediation and recovery process."

The recovery strategies differed. The dependence on people persisted.

4 → Priority 4: Understand the human dependencies

The workforce remains part of the recovery strategy

While only 13% selected rehiring qualified professionals as their preferred recovery approach, an additional 8 out of 10 response from alternative strategies, relied on people. Together, this highlights the strongest concentrated response of 39% of participants selecting workforce solutions to recover. The suggestions included recalling retired employees, redeploying staff, introducing temporary capacity, and using existing institutional knowledge.

The distinction matters.

The initial recovery selections suggested considerable variation in approach. The supporting reasoning revealed a recurring dependency on human capability.

39% of participants selected workforce solutions to recover.

“Contact and recruit possible retired staff on a contractual basis to assist. They will already have knowledge of the workarounds.”


“Identify current employees who are already trained in similar skills, partnering with Learning and Development, and use existing processing documentation to get them up to speed.”

Organizations using AI to reshape workforce structures, need to assess how much institutional capacity can be removed before the organization also loses its ability to recover. A scenario exercise of this nature, at the point of AI deployment will illuminate this insight.


Recovery options offered limited confidence

1 → Priority 1: Test the operating model before committing to it

Confidence in holding minimum service levels within impact thresholds



Mean confidence by chosen strategy (1–5)




Confidence remained low across all 4 recovery strategies. Participants who selected rehiring reported the lowest average confidence, although the small number of respondents in that group limits the conclusions that can be drawn.

The broader finding is that identifying a recovery option did not necessarily translate into confidence that minimum service levels could be maintained.

This distinction warrants attention. A contingency strategy has limited value unless the organization has established that it can execute it within the required time and capacity constraints.

WHO TOOK PART

Practitioners from inside the institutions that would have to act



Resilience belongs in every decision in the AI lifecycle.

The strongest insight from Agent Down concerns decisions made long before the “incident timeline” began.

In this scenario, the organization's dependence on the AI provider was already established. Its workforce had been reduced. The ability to return to manual processing had changed. By the time the vendor failed, the available recovery options were constrained by earlier strategic decisions. In one participant's words “This was a ticking timebomb from the start”.

Several participants recognized this directly, calling out the need for earlier consideration and intervention on various strategies long before the incident began.

This insight elevates the criticality of ensuring that the operational resilience of strategic AI decisions is considered at every decision point. Be it during the business case for AI deployment, the approval process for scaling down teams, or expanding the use cases of already deployed AI. Highlighting the material risks that counter balance productivity rewards require thorough examination and governance scrutiny given the customer, brand and regulatory impacts of maintaining critical services.

“The time to have an alternate solution was long before the crisis. The idea that you could contract with another firm quickly enough is wrong.”

“We created a single point of failure by not leveraging this as a tool instead of a replacement for the previous colleagues conducting the work.”

“Many business leaders focus on the positive outcomes shifting to AI will provide but don't have a concrete plan for when it goes down.”

“The organization cannot allow businesses to deploy AI solutions without considering resilience aspect in the event the AI fails.”

For resilience professionals, scenario exercises remain the strongest tool available in building deep risk awareness, establishing risk reward clarity and ultimately robust governance for each consideration.

From scenario to Capability Intelligence

Agent Down is one example of how iluminr helps organizations build the human capability to work safely alongside AI, and see that capability clearly over time.

1. Practice

Put people inside realistic, role-relevant scenarios.

2. Observe


Capture the decisions they make, where they hesitate and where responses diverge.

3. Compare

Benchmark decisions across teams, roles, locations or peer cohorts.


4. Improve

Use those signals to identify gaps, target future practice and track capability over time.



From individual decisions to organizational capability

Every response creates a signal. Taken together, those signals can show how capability is developing across people, teams, the organization and the wider ecosystem.



IN ILUMINR

iluminr Capability Intelligence brings these signals together over time, turning repeated practice into a clearer view of where capability is strong, where it varies and where more attention is needed.

1

Individual

See how a person responds under pressure, where they are confident, where they hesitate, and how their decisions change with practice.

2

Team

Understand whether teams make consistent decisions, know who holds authority, and can coordinate effectively when conditions change.

3

Organizational

Compare capability across functions, business units, locations or scenarios to identify recurring strengths, gaps and areas for investment.

4

Ecosystem

Extend the view across critical partners, suppliers and other stakeholders to understand how well the wider response system works together.

Turning individual decisions into evidence of capability

The value of simulation extends beyond whether participants complete an exercise.

By examining decisions, reasoning, confidence, and differences across participants, organizations can begin to identify specific capabilities that require attention.

In Agent Down, iluminr's Capability Intelligence analysis identified uncertainty about maintaining minimum service levels, gaps in stakeholder documentation, and weaknesses in recovery planning.

These findings were translated into prioritized actions, including clearer contingency activation rules, defined accountability for vendor arrangements, and further testing under more complex conditions.

The result is a more specific basis for deciding what to strengthen, who is accountable, and what should be tested next.

INSIGHTS

HIGH

Staff lacked confidence in their organization's ability to maintain service levels

MEDIUM

Participants knew who to involve, but documentation was incomplete

MEDIUM

Participants understood the immediate problem, but not the recovery steps

ACTIONS

CRITICAL

Document and communicate the agentic AI contingency plans to all relevant staff

HIGH

Clarify the decision rules for when to activate each contingency option

HIGH

Test the contingency plan in a follow-up simulation with overlapping incidents

MEDIUM

Confirm who owns vendor contingency verification and how often it is checked

What individual decisions reveal about organizational capability

Evidence	Survey	Facilitated tabletop	Individual microsimulation
Primary perspective	Stated views and expectations	Collective discussion and response	Individual choices in a defined scenario
Differences in judgment	Can be compared across responses	May be reconciled through discussion	Captured at each decision point
Reasoning	Depends on question design	Captured through discussion and observation	Recorded alongside individual decisions
Repeated measurement	Possible through repeat surveys	Possible across exercises	Structured comparison across simulations

From isolated exercises to a clearer picture of capability

Agent Down illustrates the value of examining how individuals approach the same decisions, including where their reasoning and confidence diverge.

Repeated across teams, functions, and scenarios, this evidence can help organizations identify patterns, assess whether interventions are making a difference, and establish a more consistent view of operational capability over time.

The opportunity is to make scenario testing a continuous source of insight for management and the board.

5 priorities for organizations adopting agentic AI

1

Test the operating model before committing to it

Evaluate the consequences of critical AI failure before making irreversible workforce, technology, or process decisions. Include resilience, risk, technology, and business stakeholders in the assessment.

2

Reassess impact tolerances

Determine whether existing tolerances reflect AI-enabled transaction volumes, dependencies, and available recovery capacity. Test the assumptions under which minimum service levels can be maintained.

3

Establish recovery decisions in advance

Define the conditions that would trigger alternative vendors, manual processing, external support, or workforce intervention. Confirm who has the authority to make those decisions and what capacity can realistically be mobilized.

4

Understand the human dependencies

Identify the skills, institutional knowledge, staffing levels, and intervention capabilities required when automated processes fail. Assess how workforce changes affect the credibility of recovery arrangements.

5

Bring evidence into governance

Use scenario testing to demonstrate whether critical AI dependencies can be managed within the organization's risk appetite and resilience requirements. Give management and boards visibility into the decisions, capabilities, and remaining gaps that matter.

Understand your operational resilience. Strengthen your capability.

For participants

See how your decisions compared with the wider cohort, including where approaches differed and what the results suggest about your own operational resilience.

For resilience leaders

Run Agent Down across your organization to understand how different teams assess critical AI dependencies, prioritize recovery, and make decisions under pressure. Use the findings to target improvements and measure progress.

For risk and governance leaders

Bring evidence from scenario testing into conversations about AI adoption, third-party dependencies, workforce planning, and operational resilience. Understand which assumptions have been tested and which require further assurance.

Actions - What needs to happen next?

• CRITICAL	Document and communicate the agentic AI contingency plan to all staff	< AI ▾
• HIGH	Clarify the decision rules for when to activate each contingency option	< AI ▾
• HIGH	Test the contingency plan in a follow-up simulation with overlapping incidents	< AI ▾
• MEDIUM	Confirm who owns vendor contingency verification and how often it is checked	< AI ▾

[Explore Agent Down at iluminr.io](#)

How we ran the research

The scenario

Agent Down is a single-player iluminr Microsimulation. 2 video updates (08:30 and 09:30) present an agentic AI vendor failure in fraud verification and sanctions screening, 12 months after a 70% headcount reduction. Participants answer decision, free-text and confidence questions; follow-up questions branch on each choice.

Fieldwork

Responses collected 10–30 September 2026.

Sample

31 external participants submitted responses to all core decision questions. 1 iluminr staff response excluded.

Recruitment

Invited through GFMI, social media and email. Participation was voluntary, so the cohort is self-selected.

Scoring

Operational resilience, confidence, alignment and blind-spot signals generated by iluminr Capability Intelligence. Every answer option was accepted; there is no right or wrong answer.

Free text

Coded thematically by iluminr. Quotes are verbatim, lightly trimmed, and attributed by role only.

Personal views

Responses reflect individual judgment, are reported in aggregate, and are not attributed to any organization.

What these findings show

- Patterns in how experienced practitioners judge the same AI failure
- Where decisions converge, where they diverge, and where confidence is low
- Which capabilities warrant further testing inside an organization

What they don't show

- Industry-wide operational resilience: 31 self-selected participants are not a representative sample
- How any named organization would perform in a live incident
- Statistically significant differences between subgroups

DISCLAIMER

About this report

iluminr's microsimulations are designed for training and awareness purposes only. They are not professional advice and should not be relied on as a substitute for legal, regulatory, financial, or other specialist guidance.

The scenarios are fictional but may draw on real-world events. They can include realistic images, video, audio, and simulated news or communications. Any resemblance to actual organizations or incidents is coincidental.

Readiness, Alignment, Confidence and Blind Spot signals, and the thematic coding of free-text responses in this report, are generated with the assistance of iluminr AI from the decisions participants made in the session and reviewed by iluminr. They are indicative, intended to support discussion, and should be read alongside the underlying responses.

Content is provided for learning purposes and may not reflect the full complexity of real situations. While care is taken in its development, we do not guarantee that all information is complete or suitable for every context. Decisions in your organization should be based on your own policies, expert advice, and verified information.

To the extent permitted by law, Catalyst Technologies Pty Ltd, iluminr inc, and their related entities are not responsible for decisions made based on these microsimulations or for any outcomes that may result from their use.

If you have concerns about the content or its impact, you should stop the activity and seek appropriate guidance within your organization.

APPENDIX

Appendix

Participant responses

The one action participants would prioritize

“What is the one action that you believe should be prioritized to increase confidence in maintaining minimum service levels for this scenario?”

Every free-text response, coded by iluminr to the theme of its main action. Each response sits in 1 theme. The following pages list all 31 responses verbatim.



Process

8 of 31

"What is the one action that you believe should be prioritized to increase confidence in maintaining minimum service levels for this scenario?"

Prioritization, workarounds, procedures and service levels

"Reprioritize work based on company SLO."

Business Continuity Lead

"Know what the SLA actually is. The data from the business team of actual throughput and backlog, plus estimate in how long this volume can be maintained before becoming overwhelmed"

Group Director of Operational Resilience

"Continue servicing what we can and begin preparations with contingency plans."

VP

"aligning and updating current respond and recovery procedures, including steps and relevant key resources, and knowledge required. Finally exploring alternatives for key/critical vendors"

Domain Manager

"Validating the proposed workarounds - can additional controls be used so that processing can occur and do we have an alternative provider we can readily outsource processing to?"

Manager, Operational Resilience

"Internal mitigations"

Group Head of Major Incident Management

"To ensure recovery statues are activated"

Enterprise Crisis Management Director

"Have a portfolio of levers"

Participant

People

7 of 31

"What is the one action that you believe should be prioritized to increase confidence in maintaining minimum service levels for this scenario?"

Capacity, skills, redeployment, staffing and team wellbeing

"Some additional capacity in the professional "human in the loop side" that may be somewhat idle but able to support if needed. Perhaps they are cross trained in a couple areas."

Head of Risk

"Immediately engage with HR to determine if offering overtime employees on the effected team, current staffing, identify other employees with existing skillsets, and determine feasibility of expediting onboarding skilled human professionals"

Responsible AI Specialist

"Redistribution of internal staff will increase confidence"

Specialist Business Continuity Management

"create teams within the resources you have, and ensure downtime for each team once you weather the initial 24 hours. As outsourced resources are added, ensure this cadence remains in place."

Sr. Director - Customer GRC and DR

"Manual processing but in bulk categorization, repurpose resources within the organization by training temporary resources to support the remediation. Where possible workload shift parts of the process to other teams to prioritize recovery."

Senior Manager BCIM

"breaches should be highly prioritized, to maintain minimum services it is important to stay collective and calm as a team."

Broker Assistant

"Executive support for the plan and call out for the amazing work the team will achieve by recovering from this incident."

Vice President of Enterprise Business Resilience

Vendor

7 of 31

"What is the one action that you believe should be prioritized to increase confidence in maintaining minimum service levels for this scenario?"

Alternate providers, BPO capacity, contracts and third-party testing

"Contractual capacity with existing BPOs"

SVP, Operational Risk

"Conducting third-party risk stress testing of service level agreements on vendor dependent systems in an appropriate exercise will increase confidence."

Director - Enterprise Business Resiliency

"As a reflection after the fact? Identify key risks such as this and go after executive support for solutions. The classic example is a sole source contract to save money. It may be less costly over the long run to have at least two vendors that are set with enough work to keep them in the loop and understand their ability to upscale quickly, if required."

Director Enterprise Resilience

"Fast tracking the new agentic provider. it will address the longer term issue the company is faced with and will most likely be the key factor in returning them to a normal operating state, or a least a new normal operating state."

Resilience Manager

"Testing with alternative provider, service protocols, documentation / SLA, data transfer testing, etc"

Director

"Training of the new AI"

CMCB Manager CLC

"A multi-layered approach is required. While contract staffing can, hopefully, restore minimum services levels in the short term, we should also work a medium term strategy of either using BPO or an alternate provider of AI processing."

Director, Enterprise Resilience

Governance

6 of 31

"What is the one action that you believe should be prioritized to increase confidence in maintaining minimum service levels for this scenario?"

AI dependency, planning before deployment, testing and oversight

"balance risk and resilience with velocity and volume. the increase in volume and reduction in staffing was a poor choice. better balancing for this operation is required going forward. better monitoring of the agents performance is also required. this was a ticking time-bomb from the start"

Head of Enterprise Resilience

"Having these strategies in place before completely shifting to AI rather than being reactive during the incident. Many business leaders focus on the positive outcomes of the day to day shifting to AI will provide but don't have a concrete plan for when it goes down."

Incident Manager

"Contingency planning where there is a significant tech dependency."

Operational Resilience & Continuity Manager

"I think we would need to have a clear picture in our company where there is AI dependency and concentration risk."

Director

"Testing"

IT Business Continuity Manager

"Teams to be allowed some flexibility around governance (which could be swept up at a later point). The whole focus should be minimizing the impact to the customer."

Operational Resilience Manager

Communication

3 of 31

"What is the one action that you believe should be prioritized to increase confidence in maintaining minimum service levels for this scenario?"

Clients, impacted colleagues, regulators and the recovery roadmap

"That we should be communicate a service disruption to the impacted clients, while not mentioning the exact issue specifically, provide the clients with resources or escalation paths, and to discuss the increased work with impacted colleagues and setting expectations that this would be temporary until we implement the correct recovery strategy."

Director, Global Business Resiliency

"Communication to both staff on what the firm is facing, the roadmap of how we all get out of the situation, assistance in helping onboard any new staff and checks on the work undertaken and support in the BPO and alternative provider for the mid to long term. Also communication to regulators or customers to understand what the company faces in a slow down in processing."

Senior Principal Analyst - Crisis Management

"I think active communication is important especially to the teams the most impacted that it is being addressed. Most times we work in silos and forget about the impacted parties that might be directly affected and are taking on additional workload. They are mentioning the workload and the teams on the call but is communication with them and perhaps their viewpoint on how to handle the impact being taken into account? The people on the 'ground' so to speak sometimes have better ideas because they are living in the everyday work."

2LOD Operational Resilience NA Lead