Availability: The use of AI in improving service resilience

Summary

As we[1] have developed the ITLF Availability series of papers[2], it has become clear that AI has an important potential role in improving Service Resilience, particularly of complex, 24/7 operational systems.

In this paper, we describe the role of AI under three headings which cover aspects of improving service resilience in complex operational systems.

- Architecture: covering both design principles for new software and systems, and the application of architectural principles to existing operation systems;
- Anticipation and Planning for improved Resilience through preventing or limiting the impact of failures.
- Recovery from Failure, where the aim is to recover services to the users quickly and completely.

We also describe some of the known limitations of different types of AI application.

Further, we stress that AI is written in software. Software has defects. So, AI has the same potential for error, vulnerability and failure characteristics as software in other domains[3]. This is in addition to the known AI specific characteristic errors of LLMs[4].

The purpose of this paper is to stimulate cross-organisational interest in improving service resilience through use of AI. It is a snapshot of what is a rapidly developing field. We hope that the framework of Architecture, Anticipation and Recovery supports wider understanding and implementation of measures to improve Service Resilience.

Appendix A summarises some of the AI terminology used.

Complex, 24/7, operational systems

Characteristics of complex 24/7 operational systems

This paper explores AI's role in addressing the management of these systems through targeted applications, drawing on emerging practices. They are difficult to manage for a number of reasons:

1. Complexity of Real-World Environments*:*

Operational systems have diverse user and operator behaviors, variable network traffic, multiple hardware types, many sources of software and services, and geographic dispersion.

2. *Non-Deterministic Behaviour*:

Complex tightly coupled systems with concurrent operations and real-time interactions may exhibit unpredictable behaviour.

3. *24/7 operation*

Many systems delivering digital services are required 24/7. Delivering 24/7 operation requires several factors to be in place. These include:

a. Ability to do maintenance and apply upgrades with little or no interruption in service levels;

b. A trained and responsive staff with specialist skills and understanding of business objectives;

c. A supply chain for complementary and linked services that can operate 24/7;

d. Automation of alerts of service outages or unexpected behaviour that connects when necessary to humans able to take action. AI is increasingly being used to automate testing[5] and has obvious application in filtering alerts and taking routine decisions such as engaging back-up systems;

e. Systematic scheduling of testing – including stress testing for impact tolerances and monitoring of installation of new releases;

f. Problem resolution or effective holding patterns when specialist staff are not available."

4. *Complex supply chain*

Unlike traditional supply chains, the software supply chain is dynamic and changes every day. In addition to new dependencies being added daily, pipelines are constantly updated through automation, and software applications are released multiple times a day.

The supply chain can be broken down into two categories: the people who create the software (including vendors) and the automated systems used to develop the software.

5. *Interdependency and Integration*:

Modern systems rely on third-party services, API's[6], and integrations. Identifying whether failures originate from external (3rd party) systems or within the core product can require new forensic analysis skills. Integration of components built using different architectures or standards can cause unexpected disruption.

6. *Service Level Agreements (SLA's)*.

The use of strict SLAs allowing zero or minimal downtime for new releases necessitate techniques like canary releases[7], blue-green deployments[8], and alpha/beta testing[9].

7. *Collaboration and Communication:*

Operational environments involve multiple stakeholders (developers, testers, product owners, operators, users), who may have different interpretations of requirements that lead to gaps in understanding system behaviours.

8. *Managing Data:*

Processes in which data is accumulated over long periods of time will require explicit consideration of their management. Additionally, test data must now not create biases or ethical challenges within AI enabled systems.

9. *Longevity*

Many organisations have applications containing legacy software dating back to the 1970's. A legacy system is an older (but not always obsolete) part of the IT landscape that's still in use today.

10. *Skills and Expertise:*

The people who designed and built the systems are no longer available, and in many cases have left limited documentation, so the decisions that were made in previous years may not be documented, identified or capable of comprehension.

Managing this complexity

There have been many significant shocks in recent years, from the global CrowdStrike[10] software failure in July 2024, to the power outage that closed London Heathrow[11] in 2025, the blackout that left most of Spain and Portugal without electricity for days[12] in 2025, to major cloud disruptions at AWS[13] and Cloudflare[14] in 2025.

Whilst these all have different causes, they have exposed how fragile and tightly coupled our digital infrastructure has become. Single points of failure cascade across sectors and geographies. Private enterprises, public services and citizens operate ever more fully inside what Luciano Floridi terms the "infosphere".[15].

IT is critical infrastructure. Its ongoing performance needs appropriate business continuity planning, high-availability design and system resilience engineering. Creating these is becoming beyond human capacity.

This is what makes AI-enabled operations and automation (AIOps, autonomous remediation and AI-assisted resilience engineering) essential levers for decreasing the likelihood of an incident/outage safeguarding and constraining the "blast radius" of future incidents.

The use of AI in architecture design for operationally resilient systems

The applications of AI can be split into two areas;

- architecting new applications and systems,
- improving the resilience of existing applications.

There is already a school of thought that it will be faster, and more cost effective to use AI to rewrite existing applications. There has been much said about the ability of Claude[16] to support COBOL modernisation: this is unpacked, and clearly articulated in a Thoughtworks post[17]. The reality of resilience is that it still has to be engineered into any system, considering the architecture and operational aspects as well as the code. Rewriting legacy applications can be part of a plan, but is not a complete plan.

Architecting new applications or systems

In summary, AI's contribution to operationally resilient architecture is by keeping resilience at the forefront of design, using data to predict where things will break, designing for failure, and through bringing structure to architecture. AI is very much seen as augmenting the architect, and imposing architectural discipline.

AI can support the architect to:

1.Design to the non-functional requirements,

deliver guardrails around reliability, scalability, observability and security, enabling testing of the nonfunctional requirements at the design phase, and understanding the design choices.

2. Design for failure,

using AI assistants to enforce resilience patterns (timeouts, retries with backoff, circuit breakers, bulkheads, graceful degradation) in design reviews, IaC, and code templates. In addition, AI can generate alternative "degraded" flows (smaller models, cached answers, read-only modes) so services can continue in a reduced but safe mode under stress.

3. Design modular, observable architectures

Supporting the design of logically separate data, model, and application layers with clear interfaces and standardised APIs, so faults can be isolated and contained. Additionally, AI can apply observability requirements, using AI to ensure that observability is embedded, enhancing the use of logs, traces, and metrics for anomaly detection and root-cause hints; designing the architecture so every dependency is traceable and measurable.

4. Design to governance, security and change resilience frameworks;

embedding AI-driven policy checks (least privilege, encryption, data residency) into CI/CD so every change is evaluated for its impact on resilience and cyber-posture, as well as using AI tooling to map services and data flows to business services and impact tolerances, ensuring architecture decisions support overall operational resilience objectives.

Improving the resilience of existing complex operational systems.

AI can improve resilience in existing complex operational systems by helping you understand what you really have, then to find the weak points, and so change the system in a safer, more incremental, data-driven way. This is dependent on the observability of the systems and the data available.

Few complex operational systems today exist in isolation. The dependencies may be cross organizational and/or extend into long and complex supply chains. These structures may limit the data available.

The contribution that AI can make can be categorized:

1. *Making the invisible architecture visible*

a. Code and dependency mining: AI can scan large legacy codebases, interfaces, and config to build maps of services, data flows, and dependencies, revealing hidden coupling and single points of failure you can't see from documentation alone. In addition, an expert system[18] guided by human experience can be used to set standards.
b. Hotspot and technical-debt detection: AI can highlight modules with high change/failure concentration, complex dependency fan-in/fan-out, and outdated technology, giving you a prioritized list of resilience risks to address first.

c. Learning from incidents and operations: Incident analytics: NLP over tickets, incident logs, and post-mortems can surface recurring root causes, brittle integration points, and "near misses" that architecturally weaken resilience.
d. Operational risk signals: ML on operational data (alerts, capacity, vendor performance) can identify leading indicators of disruption and inform which parts of the architecture need buffering, decoupling, or redundancy.

2. Scenario testing and complex-system simulation

a. AI-enhanced scenario planning: AI tools can generate "extreme but plausible" disruption scenarios (key provider down, data corruption, cascading job failures) and simulate business and technical impact.
b. Control and optimization of complex systems: For systems like grids, logistics, or trading platforms, AI control agents can propose parameter changes, routing strategies, or throttling policies that keep the system stable under stress. The sheer scale of the number of suppliers involved in these complex systems means that help from AI in keeping track of assets, versions, etc can generate alerts to prioritise interventions, see for instance [19].

3. Safer, incremental modernization for resilience

a. Refactoring and strangler patterns: Generative AI can help design and implement strangler-fig refactors around fragile legacy cores, suggesting service boundaries and target interfaces that reduce blast radius.
b. Automated tests and impact analysis: AI-generated tests plus impact analysis on dependency graphs allow you to change or replace components with higher confidence, reducing the risk that resilience improvements themselves cause outages.
c. Data Management and governance: AI-driven visualization tools for data mapping and analysis are in wide use, and there are many free tools[20]. AI can be used to support robust data governance, ensuring data accuracy, integrity, and availability.

4. Building an adaptive, learning system

a. Continuous resilience assessment: AI services can act as a "resilience advisor", continuously scanning telemetry, changes, and incidents to

update a live risk view and recommend architectural or configuration adjustments.

b. Feedback into governance: Insights from AI analytics feed architecture and operational-resilience forums, so design standards, risk appetite, and modernization roadmaps evolve with how the system actually behaves in production.

The use of AI to improve operational resilience

Areas where AI can provide support to reduce the number of operational failures, whether the systems are new or existing, are in

- process, people and management
- planning
- testing
- problem anticipation and mitigation.

Process, people and management

1. *Business Continuity planning*

    a. AI can assist with the development of business continuity options and planning for the management of incidents. A BCI report[21] covers the use of AI in compliance, risk and impact assessment, operational costs, incident detection and response, and continuous monitoring.
    b. AI-enabled analytics, simulation and cloud automation can be used to strengthen business continuity plans.
    c. McKinsey uses AI-driven analytics and scenario modelling to strengthen strategic resilience and business continuity across critical operations[22].
    d. KPMG applies AI and advanced analytics to stress-test supply chains and continuity plans, and to govern AI systems safely[23].
    e. PwC combines AI with AWS cloud automation to maintain secure, resilient infrastructure, reduce downtime and support rapid recovery[24].
    f. For some business-critical processes, AI-driven digital twins may be developed for simulation of recovery processes[25].

2. *Collaboration and communication*

Good communication across the organisation is a plank of effective service resilience. AI platforms can capture the body of knowledge on people as well as functions, suppliers, etc which reduce the barriers to planning system changes and managing outages.

3. *Service Level Agreements and Contract management*

This is an area well served by AI-assisted systems, see for instance[26] a podcast introducing methods and tools. Strategy powerhouses and the Big Professional

Services Networks view AI as essential for automating contract lifecycles and enhancing legal knowledge management.

a. McKinsey and Deloitte highlight how AI streamlines complex agreements and mitigates risk,

b. PwC and KPMG utilise these tools to define performance accountability and ensure compliance within legal managed services[27].

4 Regulatory compliance

a. AI can streamline regulatory compliance by automating checks, monitoring changes and providing evidence of due diligence[28]. AI-based compliance extends its utility into real-time data analytics, which is critical for maintaining up-to-date and accurate compliance reporting. By continuously analysing transactions and operations, AI tools can immediately identify discrepancies or non-compliance issues as they arise.

b. Many consultancies use AI, specifically Generative AI, for IT Systems Mapping primarily within the context of Legacy Modernisation and Technical Debt Remediation. They refer to this capability as "Discovery," "Dependency Mapping," or "Codebase Understanding[29].

5 Coding

Sometimes, legacy modules are incompletely documented and there is limited experience and expertise on those systems remaining within the organisation. AI can support the re-coding of the system, regenerating documentation, analysing the code, inputs and outputs, functionality and regenerating code. This is an area which is evolving rapidly[30] with huge focus across the industry on not just cyber security, but identification of code vulnerabilities and code quality.

Planning

AI supports planning for resilience in complex existing systems by giving you better foresight, richer simulations, and continuous evidence about where to invest change effort.

1. *Stronger scenario planning*

a. AI mines internal and external data (incidents, threats, providers, environment) to propose "extreme but plausible" disruption scenarios instead of a few manually imagined ones.
b. It helps quantify each scenario's likelihood and impact, so resilience planning can prioritise the combinations that actually matter (e.g. cyber + vendor outage + surge demand)

2. *Deep simulation of complex behaviour*

a. ML (see Appendix A) enhanced system-dynamics or agent-based models let you simulate how a complex socio-technical system will behave under stress, including cascades and feedback loops you would not see analytically.
b. Synthetic data and "digital twins" of operations allow safe stress-tests of rare events, so you can test options (buffering, routing, throttling, failover) before touching production.

3. *Continuous risk sensing and early warning*

a. Predictive analytics over telemetry, incidents, vendors, and external signals can act as an early-warning system, surfacing patterns that historically precede material disruption.
b. These models feed resilience planning with dynamic risk views, so your scenarios and playbooks update as the system and its environment evolve.

4. *Prioritising resilience investments*

a. AI can estimate how changes (extra redundancy, refactoring a hub system, new playbook) would shift key metrics like time to detect/respond/recover, helping you compare options and pick the highest-leverage interventions.
b. It supports portfolio-level planning by clustering systems and processes by criticality and fragility, guiding which parts of the architecture to harden first.

5. *Tracking execution and demonstrating progress*

a. AI tools can track remediation actions from exercises and reviews, send reminders, and maintain a live view of completion status against resilience roadmaps.
b. They generate resilience "scorecards" and trend metrics from repeated simulations and real incidents, evidencing to boards and regulators how the complex system's resilience is maturing over time.

Testing

The use of AI to provide testing is becoming widespread and broad, covering many of the associated processes.

1. *QA*

The big consultancies leverage AI to transform Quality Assurance for system resilience.

a. Deloitte deploys AI agents to autonomously generate executable test scripts[31]
b. McKinsey uses agentic AI to automate the testing lifecycle and eliminate technical debt.
c. PwC shifts to dynamic, intelligence-driven validation systems, allowing AI to interpret real-time system changes[32].

2. *Inputs to the testing process.*

a. Inputs into testing include specifications, metrics on software quality, feedback from customers, Customer Support, Project Managers, the Product Team and others.
b. AI can help analyse inputs; for example, tools such as ChatGPT[33] can be asked to find inconsistencies in specifications.

3. *Automation of testing*

a. Running automated tests in Continuous Integration on every code change creates confidence to release software changes.
b. AI can help teams automate tests eg Tools such as Cursor[34] and Copilot[35].

4. *Growing Expertise*

a. AI can be used as a “learning assistant” to help engineers resolve errors in automated tests, for example, an automation engineer can ask ChatGPT to explain a coding error.

5. *Risk based automation of testing.*

a. AI can be employed to manage the balance between manual and automated regression testing. Prioritising automation for frequently changed systems can reduce change-related incidents.
b. AI agents can adjust test scripts based on previous results, focusing on high risk areas.

6. *Managing Test Data and Environments*

a. AI can be used to generate test data, for instance analysing historical user behaviour to test edge cases and hidden bugs.
b. It can mitigate bias in generated data.
c. It can also be used to document changes to existing datasets or data models, and oversee the integrity of existing data and processes with long running data accumulation, particularly in identifying anomalous events.
d. Specialised testing can use AI-driven visual testing tools to drive user experience testing. AI Bots can simulate real user interactions, uncovering interactions.

7. *Managing the testing programme*

a. AI-enabled platforms can automate repetitive tasks and provide insight into potential issues.
b. They can improve collaboration between QA and development teams as part of CI[36] Pipelines.

8. *Feedback from the testing process*

a. AI can be used to review the effectiveness of testing processes by using AI to analyse post-deployment metrics. The testing process provides feedback to engineers, Project Managers, the Product Team, Customer Support and others. The feedback can consist of bug reports, questions, performance metrics, predictive analytics, and comments.
b. Generative AI tools can review feedback before it is published.

Problem anticipation and mitigation

AI supports both problem anticipation and recovery from partial or total failure of complex systems by speeding up detection, improving diagnosis and

increasingly automating much of the remediation. Given the close alignment of anticipation and remediation in IT operations, Artificial Intelligence for IT Operations (AIOps) is becoming both a practice and a set of tools and processes to increase availability and systems reliability.

AIOps utilises AI and machine learning to analyse the vast telemetry from modern IT estates, correlating alerts, detecting anomalies, identifying root causes, and increasingly automating remediation across complex, hybrid environments.

The use of AI to anticipate and remediate system reliability issues can be seen in areas such as:

1. *Resource Optimisation:*

AI can dynamically allocate resources (such as compute, storage, or network bandwidth) based on real-time demand and predicted needs, ensuring optimal performance and minimizing bottlenecks during peak loads or disruptive events.

2. *Operational signalling*.

a. AI-driven systems can continuously monitor operations and detect real-time anomalies or deviations from standard patterns.
b. Machine learning models can rapidly process vast amounts of data to highlight emerging risks. For example, Amazon uses AI to forecast demand spikes and supply bottlenecks, ensuring availability while reducing excess inventory by 15 to 20 percent.[37]

3. *Risk and control frameworks*

a. AI can compare risks and controls in the organisation with industry-/sector-wide data to identify what common risks and controls appear to be missing from the organisation's own risk and control registers.
b. AI can also provide analysis of controls and recommendations for control improvements[38].

4. *Anomaly detection*:

a. Experiments indicate that RNN[39]s (see Appendix A) have great potential for finding anomalies in networks. They make use of the sequential dependencies existing in network traffic data, for instance-

-observing anomalies that would otherwise have been overlooked and addressing false positives[40]. They are effective with both normal network traffic and malicious intrusions, and are a key component in modern cyber defence. They enable real-time monitoring, anomaly detection, and rapid response, far quicker than manual systems[41].

b. In practice it may be that AI generates options for human decision making.

5. *Predictive Analytics for Risk Mitigation*

a. AI can identify patterns within large data sets, finding patterns and anticipating events much more accurately than people can. By leveraging historical data, real-time information, and advanced algorithms, AI can forecast potential disruptions - such as hardware failures, cyber threats, or environmental hazards—before they escalate. This allows organizations to act pre-emptively.

Recovery from failure of complex systems

AI supports recovery from failure in complex systems by speeding up detection, improving diagnosis, and automating as much of the remediation as can be delegated.

AIOps

As stated earlier, AIOps brings together both anticipation and recovery in a set of holistic capabilities. Many major consultancies are developing these capabilities, as they see AIOps as a core lever for IT systems resilience and availability.

a. McKinsey describes it as part of a holistic monitoring approach that uses machine learning and alert correlation to "predict incidents, reduce false positives, identify root causes, and even perform self-healing", improving uptime and end-to-end observability[42].

b. BCG defines AIOps engines that autonomously handle proactive troubleshooting, upgrades, and service performance improvements, and explicitly links combining Site Reliability Engineering (SRE) with AIOps to enhance resilience through intelligent, automated recovery and maintenance[43].

c. Deloitte, in turn, positions AIOps platforms that aggregate metrics, logs and traces, perform proactive anomaly detection and automate incident resolution as a way to reduce downtime, improve system reliability and ensure resilience in increasingly complex IT environments[44].

Core capabilities for recovery

The core capabilities that AI can provide in a recovery situation are:

1. *Faster detection and triage*

a. AI-driven monitoring can analyze huge volumes of logs, metrics and traces in real time, spotting anomalies and incidents much earlier than threshold rules alone.

b. Incident-response assistants can automatically group related alerts, assess probable severity and impact, and route or escalate to the right teams with context included.

2. *Better diagnosis and decision support*

a. AI models can correlate current signals with past incidents and propose likely root causes, affected components, and blast radius, cutting time spent in initial diagnosis.
b. Generative AI can search runbooks, knowledge bases, and historical tickets to suggest concrete next steps, queries, or configuration checks tailored to the live incident.
c. AI platforms can capture the body of knowledge on people as well as functions, suppliers, etc. which reduce the barriers to managing outages.

3. *Automated and self-healing recovery*

a. Self-healing mechanisms use AI to trigger automated runbooks: restarting services, draining and replacing bad nodes, rolling back recent changes, or failing over to standby systems when confidence is high.
b. Event-driven orchestration with AI decision engines can choose between fully automated, semi-automated (with human approval), or advisory-only actions based on risk and confidence thresholds.

4. *Disaster recovery and continuity*

a. In larger disruptions, AI can manage backup and replication schedules, pick optimal recovery points, and orchestrate orderly failover of applications and data to alternate sites or clouds.
b. During recovery, AI can dynamically prioritise which services to restore first based on business criticality and interdependencies, improving effective RTO for key services

.

5. *Learning to recover better next time*

a. After incidents, AI can mine post-incident data (chats, tickets, telemetry) to identify recurring failure patterns and weak runbooks, then suggest improvements to architecture and playbooks.

b. Over time, this turns recovery into a learning loop, improving automated remediation coverage and reducing manual effort and downtime with each incident.

The wider role of AI in business continuity.

Whilst AIOps focusses on the IT Operations, AI has a broader role in business continuity. AI algorithms can identify a cybersecurity threat, supply chain disruption, or operational failure, and trigger automated response protocols, minimizing downtime and mitigating the impact of emergencies, or generating options for human decision making.

1. Better documentation for operators and users

ServiceNow[45] are pioneering the use of Generative AI to summarise manuals, and recommend actions in the case of an outage. (Some of the most extended outages have been where the initial problems were exacerbated by operators not trained in the correct recovery response, leading to them taking actions which took days to resolve/roll back[46]). ServiceNow AI Agents also use cross-enterprise data to evolve from the more familiar prompt-based activity to deep contextual comprehension, keeping people in the loop for robust oversight and governance[47].

2. Self-healing systems

The industry is shifting its focus toward autonomous remediation and self-healing systems.

a. McKinsey uses AIOps with machine learning to correlate alerts, predict incidents and trigger self-healing workflows that automate ticketing and playbook execution[48].
b. BCG and Deloitte describe AIOps platforms that detect anomalies, triage alerts and execute automated issue remediation[49],
c. KPMG highlights AI-assisted incident response that prioritises alerts and launches automated containment actions to cut detection and response times[50].

Efficient resource allocation is critical for effective response and recovery efforts during emergencies. AI algorithms can optimize resource allocation by analysing factors such as geographic location, the severity of the crisis, and available

resources in real-time. Agents and/or humans can allocate resources where needed most, maximizing the effectiveness of response efforts.

Limitations of AI

Structural issues

There are examples of AI systems which have included code errors, given wrong advice, implemented bias, and deleted user databases[51]. Analysis of the underpinning of LLMs suggests that they have the potential to hallucinate in critical tasks[52] but also reinforces the need for good data, and clear guardrails. Whilst models are being released on a regular basis, and changing rapidly, and the quality and reliability of models, as well as their potential capability is increasing, resilience requires both AI and human capability to be most effective.

We also stress that AI is implemented in software, and so has the same potential for error, vulnerability and failure characteristics as software in other domains[53]. So, we recommend hybrid AI-human systems for managing complex IT systems, and ongoing metrics to identify performance.

Practical issues

1. *Implementation costs:*

   Implementing AI in business involves significant costs across data acquisition, infrastructure, talent/staff training and recruitment, model development, integration, compliance, and maintenance. Businesses must balance cloud vs. on-prem solutions, hiring AI experts vs. outsourcing, and regulatory adherence. AI costs range from £10K for small-scale automation to £10M+ for enterprise AI[54]. The danger is that data is expensively collected without effective plans for its utilisation.

2. *Ongoing high computational costs*

   For real-time 24/7 monitoring, costs can be a factor in determining the approach, depending on the risk appetite of the organisation. Similarly, computation resources constrain new applications.

3. *Scalability and rollout need extensive planning*:

   The UK Government has produced a useful handbook: “By using the evidence-based frameworks and toolkits outlined in these resources, we have scaled *Assist* to 200+ government organisations in less than a year

since cross-government launch, achieving a 70% adoption rate which is increasing every week (as of May 2025). Specific interventions developed based on these frameworks have led to a 180% increase in the completion of AI training, a 23% improvement in users' confidence using AI at work and the de-risking of over 50 uses for Assist with a range of evidence-based mitigations." [55]

4. *Standards:*

Operational systems handling sensitive data are subject to new standards, e.g. IEEE AI Ethics Framework[56] and implementation as *Ethically Aligned Design* [57]. The EU AI Act[58] is a European regulation on artificial intelligence (AI). It assigns applications of AI to three risk categories. First, applications and systems that create an unacceptable risk, such as government-run social scoring are banned. Second, high-risk applications, such as a CV-scanning tool that ranks job applicants, are subject to specific legal requirements. Other applications have been deemed to be low risk and are largely left unregulated. Like the EU's General Data Protection Regulation (GDPR) in 2018, the EU AI Act could become a global standard.

Conclusions and recommendations

Conclusions

AI has significant potential to improve service resilience across the three aspects covered in this document, namely architecture, anticipation and recovery. AI will decrease the likelihood of service failures happening, regardless of root cause, and will decrease the potential impact and blast radius when those service failures do occur.

AI enabled tools and technology will increasingly support and augment the people and processes responsible for service resilience, which will only increase in complexity and importance.
AI is already being used in a number of ways within the service resilience environment, and many players in this industry, whether consultancies, software providers or outsourcers, are already investing in AI tooling and technology.

Recommendations

1. ITLF should incorporate AI in Service Resilience into the autumn Conference & Symposium and invite participants from across BCS and other organisations.

2. ITLF should seek to develop and publicise case studies for the use of AI in improving Service Resilience, with other BCS members and the wider network.

Appendix A: AI Terminology

One type of AI that is well proven is large language models (LLMs) which excel at processing, understanding, and generating human language. These are highly effective for automating text-based tasks, generating human-like responses, providing 24/7 customer support, creating content, analysing data, and powering intelligent assistants. Many cloud service providers use LLM models to identify and predict faults, as well as moving services and data around within the infrastructure[59]. The well-known problems of LLMs such as bias, data privacy concerns and resource used for training and deployment[60], are less important in most service resilience applications.

Augmentation of humans by AI is characteristic of early successful applications of AI[61] for instance in healthcare, where analysis of images is mapped to human decision making. Such systems were called *Expert systems*[62] and are now often referred to as Workforce AI.

Generative AI (GenAI) involves AI tools that can generate new content based on specific prompts or instructions. Their use in the creative industries is causing major disruption and challenge to intellectual property regimes[63]. In IT Service Management, GenAI is used to summarise manuals, and recommend actions in the case of an outage[64].

AI agents may be triggered autonomously for instance to undertake stress testing or recovery actions[65]. *Agentic AI* refers to the system underpinning AI agents, from the algorithms to the architecture, linking them together in an ecosystem. It is the framework within which AI agents operate.

An *RNN* (recurrent neural network) is a deep neural network trained on sequential or time series data to create a machine learning (ML) model that can make sequential predictions or conclusions based on sequential inputs.

Links and references

[1] In this text and others in the series of ITLF Availability Papers, "we" refers to the Availability/Service Resilience Working Group of the IT Leaders Forum of the BCS – the Chartered Institute for IT, and colleagues from our network who have provided additional insights.

[2] https://www.bcs.org/membership-and-registrations/member-communities/bcs-it-leaders-forum/papers/

[3] https://www.bcs.org/media/9679/itlf-software-risk-resilience.pdf

[4] https://www.sbs.ox.ac.uk/oxford-experience/blogs/yashwant-aditya/hidden-costs-ai-5-critical-limitations-large-language-models-businesses

[5] How to gradually incorporate AI in software testing | TechTarget

[6] Application Programme Interface

[7] https://sre.google/workbook/canarying-releases/

[8] https://www.ibm.com/think/topics/blue-green-deployment

[9] https://www.linkedin.com/pulse/navigating-alpha-beta-testing-comprehensive-guide-charolia-cd3rf/

[10] https://en.wikipedia.org/wiki/2024_CrowdStrike-related_IT_outages#:~:text=On%2019%20July%202024%2C%20the,Windows%20computers%20running%20the%20software.

[11] https://www.thebci.org/news/the-kelly-review-lessons-from-heathrow-s-power-outage.html

[12] https://www.bbc.co.uk/news/articles/cvg0r3z3lvqo

[13] https://www.bbc.co.uk/news/articles/cev1en9077ro

[14] https://www.theguardian.com/technology/2025/dec/05/another-cloudflare-outage-takes-down-websites-linkedin-zoom

[15] https://www.researchgate.net/publication/30384445_Ethics_in_the_Infosphere

[16] https://claude.com/blog/how-ai-helps-break-cost-barrier-cobol-modernization

[17] https://www.thoughtworks.com/en-gb/insights/articles/claude-code-cobol-modernization-reality

[18] https://www.techtarget.com/searchenterpriseai/definition/expert-system

[19] https://www.ey.com/en_us/insights/consulting/ai-enhances-it-asset-management-automation-and-security

[20] https://www.reddit.com/r/datavisualization/comments/1fmxgmg/what_is_the_best_free_ai_tool_for_data/

[21] https://www.thebci.org/news/solving-the-top-5-iso-22301-challenges-with-ai.html

[22] https://www.mckinsey.com/industries/technology-media-and-telecommunications/our-insights/sovereign-ai-building-ecosystems-for-strategic-resilience-and-impact

[23] https://assets.kpmg.com/content/dam/kpmgsites/cz/pdf/2025/kpmg-future-of-supply-chain-report.pdf.coredownload.inline.pdf

[24] https://www.pwc.com/us/en/technology/alliances/amazon-web-services/ai-solutions.html

[25] https://www.sciencedirect.com/science/article/pii/S277266222300005X

[26] https://go.contractpodai.com/rs/203-FSR-576/images/A%20Practical%20Guide%20to%20Legal%20GenAI%20Solutions%20-%20ContractPodAi.pdf?version=0

[27] https://legalbriefs.deloitte.com/post/102igbe/discover-the-future-of-contract-management-how-generative-ai-is-transforming-con; https://www.mckinsey.com/featured-insights/in-the-balance/in-the-age-of-gen-ai-legal-km-is-more-important-than-ever;

https://kpmg.com/uk/en/services/products/cognitive-contract-management.html; https://www.pwc.com/m1/en/publications/contract-lifecycle-management.html

[28] https://www.caseiq.com/

[29] https://www.bcg.com/x/the-multiplier/how-gen-ai-rewriting-legacy-tech-modernization-rules; https://www.mckinsey.com/capabilities/quantumblack/our-insights/ai-for-it-modernization-faster-cheaper-and-better; https://kpmg.com/kpmg-us/content/dam/kpmg/pdf/2025/from-legacy-to-leading-revolutionizing.pdf

[30] https://www.anthropic.com/glasswing

[31] Deloitte QA

[32] https://www.deloitte.com/nz/en/services/consulting/perspectives/ai-assisted-software-engineering.html; https://www.mckinsey.com/capabilities/quantumblack/our-insights/ai-for-it-modernization-faster-cheaper-and-better; https://www.pwc.com/us/en/industries/health-industries/library/computer-system-validation.html

[33] https://www.chatgpt.com

[34] https://cursor.com/en

[35] https://copilot.microsoft.com

[36] https://aws.amazon.com/devops/continuous-integration/#:~:text=Continuous%20integration%20refers%20to%20the,for%20a%20release%20to%20production .

[37] https://neontri.com/blog/ai-demand-forecasting/

[38] Examples are available at Diligent - Clarify risk. Elevate governance.

[39] See Appendix A

[40] https://ieeexplore.ieee.org/document/10467790

[41] https://www.paloaltonetworks.co.uk/cyberpedia/ai-risks-and-benefits-in-cybersecurity

[42] https://www.mckinsey.com/capabilities/tech-and-ai/our-insights/tech-forward/boosting-it-resilience-efforts-through-application-performance-monitoring

[43] https://web-assets.bcg.com/img-src/BCG-Ready-Or-Not-AI-Is-Coming-to-IT-Operations_August-2019_tcm9-227502.pdf; https://www.bcg.com/publications/2021/unlock-value-with-sre-and-ai-in-it-operations

[44] https://www.deloitte.com/de/de/services/consulting/perspectives/aiops.html

[45] https://www.servicenow.com/blogs/2024/generative-ai-use-cases

[46] https://www.reuters.com/article/world/british-airways-it-outage-caused-by-contractor-who-switched-off-power-times-idUSKBN18T0L6/#:~:text=LONDON%20(Reuters)%20%2D%20A%20contractor,Times%20newspaper%20reported%20on%20Friday

[47] https://m.digitalisationworld.com/news/68430/servicenow-releases-its-most-comprehensive-set-of-new-ai-innovations

[48] https://www.mckinsey.com/capabilities/tech-and-ai/our-insights/tech-forward/boosting-it-resilience-efforts-through-application-performance-monitoring

[49]

[50] https://www.mckinsey.com/capabilities/tech-and-ai/our-insights/tech-forward/boosting-it-resilience-efforts-through-application-performance-monitoring; https://web-assets.bcg.com/img-src/BCG-Ready-Or-Not-AI-Is-Coming-to-IT-Operations_August-2019_tcm9-227502.pdf; https://www.deloitte.com/de/de/services/consulting/perspectives/aiops.html; https://kpmg.com/sa/en/insights/ai-and-technology/making-ai-a-cybersecurity-ally.html

[51] https://www.cio.com/article/190888/5-famous-analytics-and-ai-disasters.html

[52] https://arxiv.org/html/2409.05746v1

[53] As [3]

[54] https://www.walturn.com/insights/the-cost-of-implementing-ai-in-a-business-a-comprehensive-analysis
[55] https://www.deloitte.com/uk/en/forms/report/state-of-ai-in-enterprise.html https://assets.publishing.service.gov.uk/media/683ef8c0d21e8a73d10d32ca/The_People_Factor_A_human-centred_approach_to_scaling_AI_tools.pdf
[56] The key principles of the IEEE AI Ethics Framework are: Human rights: Respecting and protecting internationally recognised human rights. Well-being: Prioritising the overall well-being of humanity and the environment. Accountability: Ensuring that designers and operators are responsible and accountable.
[57] https://standards.ieee.org/wp-content/uploads/import/documents/other/ead_v2.pdf
[58] https://artificialintelligenceact.eu/
[59] As [24]
[60] https://lumenalta.com/insights/7-surprisingly-powerful-large-language-model-applications
[61] https://clanx.ai/glossary/human-ai-colaboration#:~:text=%E2%80%8D-,What%20is%20an%20example%20of%20Human%2DAI%20collaboration%3F,medical%20images%20and%20patient%20records .
[62] https://www.freshworks.com/freshservice/resources/cio-guide-to-modern-itsm-report/
[63] https://www.weforum.org/stories/2025/01/the-impact-of-genai-on-the-creative-industries/
[64] As [21]
[65] https://www.zendesk.co.uk/blog/agentic-ai-in-itsm/#Use%20cases%20of%20agentic%20AI%20in%20ITSM