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 summarise potential applications of AI in

- Managing complexity
- Testing in an operational environment
- Problem anticipation and mitigation
- Recovery

We also describe some of the known limitations of different types of AI application, and stress that AI is written in software, and so has the same error, vulnerability and failure characteristics as software in other domains[3].

The purpose of this paper is to stimulate cross-organisational interest in improving service resilience through use of AI, and to provide a framework for consultation with AI experts in BCS and the wider expert network.

An Appendix summarises some of the AI terminology used.

We focus on 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:

- Complexity of Real-World Environments: Operational systems have diverse user and operator behaviors, variable network traffic, multiple

[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

hardware types, many sources of software and services, and geographic dispersion.

- Non-Deterministic Behaviour: Complex tightly coupled systems with concurrent operations and real-time interactions may exhibit unpredictable behaviour.
- Interdependency and Integration: Modern systems rely on third-party services, APIs, 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.
- Service Level Agreements. The use of strict SLAs allowing zero or minimal downtime for new releases necessitate techniques like canary releases, blue-green deployments, and alpha/beta testing.
- 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.
- 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.

Managing complexity

System mapping and documentation

The use of AI to map the architecture, such as interactions between systems modules, APIs and external services, can support documentation of assets. An expert system[4] guided by human experience can be used to set standards.

AI-driven visualization tools for data mapping and analysis are in wide use, and there are many free tools[5]. AI can be used to support robust data governance, ensuring data accuracy, integrity, and availability. It can also help organizations test and validate their digital strategies across multiple dimensions of resilience: system, cyber, informational, organizational, operational, and people.

[4] https://www.techtarget.com/searchenterpriseai/definition/expert-system
[5] https://www.reddit.com/r/datavisualization/comments/1fmxgmg/what_is_the_best_free_ai_tool_for_data/

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 [6].

AI can streamline regulatory compliance by automating checks, monitoring changes and providing evidence of due diligence[7]. 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.

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.

Service Level Agreements and Contract management

This is an area well served by AI-assisted systems, see for instance[8] a podcast introducing methods and tools.

Coding

Sometimes, legacy modules are incompletely documented and the only person who can reliably alter them has left. Here, documentation of inputs and outputs can be used for AI based recoding of the system. This new code is of course subject to the known error rate for code[9], and in addition code generated by AI systems Is known to often contain vulnerabilities to cyber-attacks[10].

[6] https://www.ey.com/en_us/insights/consulting/ai-enhances-it-asset-management-automation-and-security
[7] https://www.caseiq.com/
[8] https://go.contractpodai.com/rs/203-FSR-576/images/A%20Practical%20Guide%20to%20Legal%20GenAI%20Solutions%20-%20ContractPodAi.pdf?version=0
[9] https://www.bcs.org/media/9679/itlf-software-risk-resilience.pdf
[10] https://socprime.com/blog/cve-2025-32711-zero-click-ai-vulnerability/

Testing in an operational environment

Inputs to the testing process

Inputs into testing include specifications, metrics on software quality, feedback from customers, Customer Support, Project Managers, the Product Team and others. AI can help analyse inputs; for example, tools such as ChatGPT[11] can be asked to find inconsistencies in specifications.

Automation of testing

Running automated tests in Continuous Integration on every code change creates confidence to release software changes. AI tools can help teams automate tests. This will help teams that already use automated tests for regression testing, and help teams that still do manual regression testing to automate their regression tests. AI Tools such as Cursor[12] and Copilot[13] can help automate tests.

AI can also 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.

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. AI agents can adjust test scripts based on previous results, focusing on high risk areas.

Managing Test Data and Environments

AI can be used to generate test data, for instance analysing historical user behaviour to test edge cases and hidden bugs. It can mitigate bias in generated data. 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.

Specialised testing can use AI-driven visual testing tools to drive user experience testing. AI Bots can simulate real user interactions, uncovering interactions.

[11] https://www.chatgpt.com
[12] https://cursor.com/en
[13] https://copilot.microsoft.com

Managing the testing programme

AI-enabled platforms can automate repetitive tasks and provide insight into potential issues. They can improve collaboration between QA and development teams as part of CI[14] Pipelines.

Feedback from the testing process

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. AI tools, such as ChatGPT, can review feedback before it is given.

Problem Anticipation and Mitigation

Anticipation

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.

Resilient IT systems need to anticipate and mitigate potential threats. AI-driven systems can continuously monitor operations and detect real-time anomalies or deviations from standard patterns. 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.[15]

AI models can anticipate possible defects by analysing past test cycles, allowing testers to focus on high risk modules. They may also initiate agents to monitor or close down subsystems, geographies or user groups.

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

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

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. AI can also provide analysis of controls and recommendations for control improvements[16].

Experiments indicate that RNN[17]s 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[18]. 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.[19]. In practice it may be that AI generates options for human decision making.

Predictive Analytics for Risk Mitigation

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

Business Continuity planning

AI can assist with the development of business continuity options and planning for the management of incidents. A BCI report[20] covers the use of AI in compliance, risk and impact assessment, operational costs, incident detection and response, and continuous monitoring.

We might also expect to see AI-driven digital twins for simulation of recovery processes.

[16] Examples are available at Diligent - Clarify risk. Elevate governance.
[17] 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.
[18] https://ieeexplore.ieee.org/document/10467790
[19] https://www.paloaltonetworks.co.uk/cyberpedia/ai-risks-and-benefits-in-cybersecurity
[20] https://www.thebci.org/news/solving-the-top-5-iso-22301-challenges-with-ai.html

Automated Incident Detection and Response

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. In practice, AI may best be used to generate options for human decision making.

ServiceNow[21] 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 recover response, leading to them taking actions which took days to resolve/roll back[22]). 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[23].

Resource allocation

AI can address resource allocation in every day operation, through machine learning models that analyse usage patterns in real-time, such as Google's AI-driven traffic management in data centres. This reduced latency by predicting load spikes and reduced the cooling bill by 40%[24].

Efficient resource allocation is critical for effective response and recovery efforts during emergencies. AI algorithms can optimize resource allocation by analyzing 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.

Intelligent automation helps maintain critical services when human resources are stretched or systems are under pressure. AI can reallocate resources dynamically, while tools like chatbots can offer real-time customer support and

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

[22] 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

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

[24] https://deepmind.google/discover/blog/deepmind-ai-reduces-google-data-centre-cooling-bill-by-40/

reassurance during service disruptions[25]. Reinforcement learning[26] is built into a number of resource allocation systems[27].

Limitations of AI

Structural issues

AI systems are well known to include code errors, give wrong advice, implement bias, and to wipe out databases[28]. Analysis of the underpinning of LLMs suggests that they are prone to hallucinate in critical tasks[29]. This is a concern for crisis management.

GenAI characteristics include hallucination (e.g. of books or articles that do not exist), giving false information causing people to break the law or suffer financial loss, and implementing illegal selection criteria.

We also stress that AI is implemented in software, and so has the same error, vulnerability and failure characteristics as software in other domains[30].

So, we recommend hybrid AI-human systems (aka Workforce AI) for managing complex IT systems, and ongoing metrics to identify performance.

Practical issues

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

[25] https://drj.com/journal_main/business-continuity-management-and-artificial-intelligence/
[26] Reinforcement learning (RL) is a machine learning (ML) technique that trains software to make decisions to achieve the most optimal results. It mimics the trial-and-error learning process that humans use to achieve their goals.
[27] https://www.linkedin.com/advice/3/how-can-you-use-reinforcement-learning-optimize-jekvc#:~:text=Reinforcement%20Learning%20(RL)%20is%20beneficial,adjusting%20strategies%20in%20real%2Dtime
[28] https://www.cio.com/article/190888/5-famous-analytics-and-ai-disasters.html
[29] https://arxiv.org/html/2409.05746v1
[30] As [3]

regulatory adherence. AI costs range from £10K for small scale automation to £10M+ for enterprise AI[31].

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

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." [32]

Standards: Operational systems handling sensitive data are subject to new standards, e.g. IEEE AI Ethics Framework[33] and implementation as *Ethically Aligned Design* [34].

The EU AI Act[35] 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.

---

[31] https://www.walturn.com/insights/the-cost-of-implementing-ai-in-a-business-a-comprehensive-analysis

[32] https://assets.publishing.service.gov.uk/media/683ef8c0d21e8a73d10d32ca/The_People_Factor_A_human-centred_approach_to_scaling_AI_tools.pdf

[33] 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.

[34] https://standards.ieee.org/wp-content/uploads/import/documents/other/ead_v2.pdf

[35] https://artificialintelligenceact.eu/

Recommendations

1. ITLF should work with BCS SGAI to develop the contents of this draft paper, while in parallel looking for real cases from the ITLF, CIO, SM-ITAM and SIGIST Communities.

2. ITLF should seek to develop and publicise metrics for the use of AI in improving Service Resilience, with other BCS members and the wider network. This could include proposing metrics (e.g., downtime reduction, incident response time) to measure AI's impact on resilience.

Appendix – AI Terminology[36]

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[37]. The well-known problems of LLMs such as bias, data privacy concerns and resource used for training and deployment[38], are less important in most service resilience applications.

Augmentation of humans by AI is characteristic of early successful applications of AI[39] for instance in healthcare, where analysis of images is mapped to human decision making. Such systems were called *Expert systems*[40] 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[41]. In IT Service Management, GenAI is used to summarise manuals, and recommend actions in the case of an outage[42].

AI agents may be triggered autonomously for instance to undertake stress testing or recovery actions[43]. *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.

---

[36] Based on work with the perplexity engine https://www.perplexity.ai/

[37] As [24]

[38] https://lumenalta.com/insights/7-surprisingly-powerful-large-language-model-applications

[39] 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 .

[40] https://www.freshworks.com/freshservice/resources/cio-guide-to-modern-itsm-report/

[41] https://www.weforum.org/stories/2025/01/the-impact-of-genai-on-the-creative-industries/

[42] As [21]

[43] https://www.zendesk.co.uk/blog/agentic-ai-in-itsm/#Use%20cases%20of%20agentic%20AI%20in%20ITSM