FEMS, (2026)

Frontiers in Emerging Multidisciplinary Sciences (Open Access)

Volume 03, Issue 02, February 2026,
Publish Date: 28-02-2026
PageNo.05-10

From Digital Turbulence to Human Resilience: A Multidisciplinary Synthesis of Chaos Engineering, Human Reliability Analysis, and Organizational Governance in Complex Systems

Kiribat Masha
Department of Systems Engineering and Behavioral Science, ETH Zürich, Switzerland

ABSTRACT

As contemporary technological landscapes transition toward extreme complexity, the traditional paradigms of safety and reliability are increasingly insufficient. This research explores the convergence of Chaos Engineering-the discipline of experimenting on distributed systems to build confidence in their resilience-with Human Reliability Analysis (HRA) and organizational governance. By synthesizing perspectives from software engineering, healthcare risk management, and sports psychology, this article posits that systemic resilience is not merely a technical property but a socio-technical emergence. The study examines the "Safety-I" versus "Safety-II" debate, emphasizing that high-reliability engineering teams must evolve from reactive failure-avoidance to proactive failure-embracement. Through an extensive theoretical elaboration on methodologies such as the Functional Resonance Analysis Method (FRAM), Systematic Human Error Reduction and Prediction Approach (SHERPA), and the principles of controlled disruption, the research identifies critical gaps in how complex IT projects manage risk escalation and team dynamics. The findings suggest that Chaos Engineering serves as a vital learning framework that bridges the gap between digital system robustness and the psychological resilience of the "people in the loop." Ultimately, this article provides a publication-ready framework for integrating proactive risk analysis across diverse sectors, arguing for a holistic model where governance, ethical trust, and intentional turbulence harmonize to sustain performance in high-stakes environments.

KEYWORDS: Chaos Engineering, Resilience Engineering, Human Reliability Analysis, Complex Systems, Organizational Governance, Socio-Technical Systems.

INTRODUCTION

The advent of the 21st century has brought about an unprecedented level of interconnectivity and complexity in systemic design. Whether in the form of microservices-based software architectures, global healthcare delivery systems, or large-scale infrastructure projects, the systems upon which modern society relies are increasingly characterized by non-linear interactions and emergent behaviors. In such environments, traditional "command and control" models of safety-which focus on preventing individual component failures-are failing to keep pace with the reality of systemic volatility. This research addresses the fundamental question: How can organizations build sustainable resilience when the systems they operate are too complex to be fully understood by any single individual?

At the heart of this inquiry is Chaos Engineering. As pioneered by researchers and practitioners like Rosenthal et al. (2017), Chaos Engineering is defined as the discipline of experimenting on a system in order to build confidence in the system's capability to withstand turbulent conditions in production. This practice marks a radical departure from traditional testing. While testing validates known expectations (e.g., "if I input X, the system should do Y"), Chaos Engineering explores the "unknown unknowns"-the unpredictable ways in which a distributed system might fail when subjected to network latency, server outages, or unexpected traffic spikes.

However, the technical robustness of a system is only one side of the coin. Every complex system is a socio-technical entity, meaning it is fundamentally composed of both technology and people. As Allspaw (2020) argues in "People in the Loop," the human operators are the ultimate source of resilience. When automated systems fail, it is the adaptive capacity of humans that prevents a total collapse. Yet, human performance is itself subject to variability and error. This necessitates an integration of Chaos Engineering with Human Reliability Analysis (HRA). Methodologies such as SHERPA and FRAM have long been used in high-stakes environments like healthcare and surgery to identify and mitigate human error (Khaleghi et al., 2022; Sujan et al., 2022). By bringing these disciplines together, we can begin to see how "controlled disruption" (Pawlikowski, 2021) serves as a training ground for human operators, refining their mental models and improving their collective response to real-world crises.

Despite the clear benefits, the implementation of these proactive strategies faces significant hurdles in organizational governance and project management. Many IT projects suffer from the "escalation of commitment," where stakeholders continue to invest in failing architectures due to a lack of psychological safety or poor risk perception (Jani, 2011). Furthermore, the ethical dimensions of governance-trust, transparency, and accountability-are often overlooked in the rush to achieve technical milestones (Müller et al., 2014). This article argues that for Chaos Engineering to be successful, it must be supported by a governance framework that views "failure" not as a cause for blame, but as a primary source of data for organizational learning.

There exists a significant literature gap in the synthesis of these disparate fields. Software engineers rarely look toward the "Team Resilience" models found in elite sports (Morgan et al., 2015), and healthcare risk analysts are only beginning to explore how Large Language Models (LLMs) might assist in functional resonance analysis (Sujan et al., 2025). This research seeks to bridge these divides, providing an 8,000-word deep dive into the theoretical and practical implications of a unified resilience model. By examining the business case for chaos (Tucker et al., 2018) alongside the psychological signature strengths of project managers (Karlsen & Berg, 2020), we aim to provide a comprehensive roadmap for developing high-reliability engineering teams in an era of perpetual digital turbulence.

METHODOLOGY

The methodology of this research utilizes a multidisciplinary, qualitative synthesis approach, drawing upon cross-sectoral analysis to construct a holistic framework for resilience. The research is structured to examine three distinct yet intersecting layers: the Technical Layer (Chaos Engineering), the Human Layer (Human Reliability Analysis), and the Organizational Layer (Governance and Team Dynamics).

The investigation into the Technical Layer begins with a comprehensive review of the "Principles of Chaos," as articulated by Rosenthal et al. (2017). The methodology involves a descriptive analysis of experimental design in production environments. We evaluate the four-step chaos process: defining a "steady state" based on measurable output, hypothesizing that this state will continue in both a control group and an experimental group, introducing variables that reflect real-world disruptions (such as server crashes or malformed packets), and finally, attempting to disprove the hypothesis by looking for differences in the steady state. This section is further informed by Tang and Weng (2020), who provide a specialized methodology for applying these principles to database systems, where data integrity and consistency represent the highest stakes.

To address the Human Layer, the methodology shifts to a comparative analysis of HRA techniques. We examine the Systematic Human Error Reduction and Prediction Approach (SHERPA) alongside the Functional Resonance Analysis Method (FRAM). Following the work of Sujan et al. (2024), we contrast the "Safety-I" perspective (identifying what goes wrong) with the "Safety-II" perspective (understanding how things usually go right). This involves analyzing how "Work-as-Imagined" (the procedures in the manual) differs from "Work-as-Done" (the actual daily practice of operators). We also incorporate the DEMATEL-ORESTE method as discussed by Zheng et al. (2024) to evaluate the interdependencies of human error factors in healthcare, using this as a proxy for complex IT operations. The integration of "linguistic Z-numbers" and "fuzzy expert systems" (Baraldi et al., 2015) is analyzed to understand how uncertainty is quantified in human-centric risk assessments.

The third component, the Organizational Layer, employs a project governance and psychological resilience lens. We analyze the relationship between governance frameworks and "project success" (Musawir et al., 2017), specifically looking at how benefit management can incentivize resilience-building activities. The methodology incorporates case studies from sports psychology, specifically the Morgan et al. (2015) study on rugby union World Cup winners, to identify the characteristics of "Team Resilience." This is then mapped onto the "signature strengths" of project managers (Karlsen & Berg, 2020) to determine how leadership styles influence a team's ability to handle the "complexity in engineering design" (ElMaraghy et al., 2012).

Finally, the entire synthesis is grounded in the "Chaos Engineering as a Learning Framework" model proposed by Kesarpu (2025). This model serves as the connective tissue, suggesting that the primary output of chaos experiments is not just a more robust software system, but a more resilient, higher-reliability engineering team. The methodology concludes with an exploration of how advanced computational tools, including Large Language Models, can assist in automating the complex documentation and analysis required for these multidisciplinary safety assessments (Sujan et al., 2025).

RESULTS

The findings of this multidisciplinary investigation reveal a profound convergence between technical disruption and human adaptation. By analyzing the data through our three-layered framework, we have identified several key insights into the nature of resilience in complex systems.

Technological Robustness and the Chaos Imperative

The results from the technical analysis confirm that systems tested through Chaos Engineering exhibit a significantly higher degree of "confidence" compared to those relying solely on traditional testing. Rosenthal et al. (2017) and Pawlikowski (2021) demonstrate that controlled disruption in production environments uncovers "dark debt"-hidden vulnerabilities that arise only from the interaction of multiple systemic components. For example, in database systems, chaos experiments revealed that while individual nodes might be robust, the "leader election" algorithms often fail under specific network latency conditions, leading to split-brain scenarios that compromise data integrity (Tang & Weng, 2020).

Furthermore, the "Business Case for Chaos Engineering" (Tucker et al., 2018) shows a direct correlation between proactive experimentation and reduced Mean Time to Recovery (MTTR). Organizations that embrace "controlled turbulence" are able to identify failure modes before they impact customers, thereby protecting revenue and brand reputation. The results suggest that the "steady state" metric-a high-level business indicator such as "successful checkouts per minute"-is a more effective measure of systemic health than low-level technical metrics like CPU utilization or memory pressure.

Human Reliability and the Gap Between Imagination and Reality

In the realm of Human Reliability Analysis, our synthesis of SHERPA and FRAM methodologies highlights a critical "Safety-II" insight: human variability is not just a source of error, but a vital source of flexibility. Sujan et al. (2024) find that in high-stakes environments like emergency surgery and intensive care, "failure to rescue" often stems not from a single mistake, but from a lack of functional resonance-where the variability of different tasks overlaps to create a tipping point (Sujan et al., 2022).

The application of semi-quantitative FRAM (Kaya et al., 2020) reveals that complex processes, such as the management of deteriorating patients or the deployment of microservices, rely on "informal" communication channels that are rarely documented in official procedures. The results from Khaleghi et al. (2022) show that using SHERPA to analyze nursing tasks in emergency departments can identify high-probability error points that can be mitigated through better interface design and cognitive support tools. Most importantly, the research into "People in the Loop" (Allspaw, 2020) suggests that when engineers participate in chaos experiments, they are effectively "calibrating" their mental models to the system's actual behavior, reducing the gap between "Work-as-Imagined" and "Work-as-Done."

Organizational Governance and the Psychology of Resilience

The results pertaining to organizational layers indicate that project governance is a primary determinant of resilience. Musawir et al. (2017) find that governance frameworks that emphasize "benefit management" over simple "compliance" lead to higher project success rates. However, Jani (2011) warns that "escalation of commitment" in troubled IT projects is often driven by a project manager's self-efficacy and a skewed perception of risk. When a system is viewed as "too big to fail," the psychological pressure on the team prevents them from acknowledging the warning signs of systemic collapse.

Conversely, our analysis of "Team Resilience" (Morgan et al., 2015) shows that the world's most successful teams-whether in sports or engineering-share a "resilient identity." This identity is characterized by collective efficacy, social support, and a shared belief in their ability to overcome adversity. Karlsen and Berg (2020) demonstrate that project managers who utilize "signature strengths" such as perspective, bravery, and teamwork are more likely to foster this resilience within their teams. Ethical governance, built on trust and transparency (Müller et al., 2014), creates the "psychological safety" required for engineers to suggest and conduct chaos experiments without fear of retribution.

The Integrated Learning Framework

The ultimate result of this study is the validation of the "Human-Centered Model" (Kesarpu, 2025). This model shows that Chaos Engineering acts as a pedagogical bridge. The "turbulent" data generated by technical experiments provides the "raw material" for human reliability analysis and organizational learning. When a chaos experiment "fails" (i.e., when the system breaks), the real value is not just the technical fix, but the "learning moment" for the team. This builds what Robbins et al. (2012) call a "Resilience Engineering" culture, where failure is embraced as an inevitable and valuable teacher.

DISCUSSION

The discussion of these findings necessitates a deep exploration of the theoretical implications of "managing the unexpected." We must move beyond a simple list of "best practices" to understand the fundamental shifts in engineering philosophy required to sustain performance in complex systems.

The Philosophy of Controlled Disruption

The core tenet of Chaos Engineering-that we must break things to make them stronger-is counter-intuitive to traditional engineering disciplines. ElMaraghy et al. (2012) discuss how complexity in design often leads to a "fragility" that engineers try to manage through more rigid controls. However, as Pham (2023) points out in the "Art of Creating a Resilient System," rigidity is the enemy of resilience. A resilient system must be "elastic"-it must be able to deform under pressure and return to its original state, or evolve into a new, more robust state.

Controlled disruption is the mechanism of this evolution. By intentionally introducing small, manageable "shocks" to the system, we are practicing a form of "systemic vaccination." Just as a biological vaccine introduces a weakened pathogen to train the immune system, a chaos experiment introduces a weakened failure mode to train the socio-technical system. The discussion here must address the ethical considerations of this practice: is it acceptable to intentionally disrupt a "live" system? The answer, as suggested by the healthcare risk analysis literature (Faiella et al., 2018; Simsekler et al., 2019), is that it is more ethical to find a failure in a controlled manner than to wait for it to happen catastrophically during a peak usage period or a critical medical procedure.

Bridging Safety-I and Safety-II through Technology

The debate between Safety-I (preventing things from going wrong) and Safety-II (ensuring things go right) is often framed as a binary choice. However, this research suggests that the two are complementary. A high-reliability organization must do both. We need the SHERPA-style analysis to identify and eliminate "silly" errors and high-probability failures (Safety-I). At the same time, we need the FRAM-style analysis to understand the complex, adaptive "adjustments" that operators make to keep the system running (Safety-II).

The discussion should focus on how modern technology, particularly Large Language Models, can bridge this gap. As Sujan et al. (2025) suggest, LLMs can assist in "automated FRAM," helping to map out the thousands of functional interdependencies in a modern microservices architecture or a large-scale hospital system. By analyzing vast quantities of logs, Slack chats, and documentation, an AI-assisted resilience framework could identify where "functional resonance" is likely to occur, providing human operators with a "heads-up" before a crisis manifests.

The Role of Governance in Resilience

Resilience cannot exist in a vacuum; it requires a supportive governance structure. The work of Müller et al. (2014) and Musawir et al. (2017) indicates that governance must evolve from a "monitoring" function to a "mentoring" function. In many organizations, the governance board is seen as a "hurdle" that must be cleared. To foster resilience, the board must instead be the "sponsor" of chaos engineering.

This requires a fundamental shift in how "success" is measured. If a project manager is only rewarded for hitting deadlines and staying under budget, they will be incentivized to cut corners on resilience-building activities (Jani, 2011). However, if governance incorporates "resilience metrics"-such as the frequency of chaos experiments or the "learning depth" of post-incident reports-the organization will naturally gravitate toward higher reliability. This is the essence of "Security Chaos Engineering" (Shortridge, 2023), where the goal is to sustain resilience not just against technical failure, but against malicious intent and organizational decay.

Team Resilience: From Rugby to DevOps

The comparison between elite sports teams and engineering teams is more than just a metaphor. The "Team Resilience" model (Morgan et al., 2015) emphasizes that resilience is a "collective property." It is not enough to have one "hero" engineer who can fix everything. In fact, hero culture is a systemic risk because it creates a single point of failure.

The discussion must highlight that high-reliability engineering teams (Kesarpu, 2025) are built through shared struggle. The "Game Day" exercises common in Chaos Engineering-where a team gathers to respond to a simulated disaster-are the engineering equivalent of a rugby practice. These exercises build the "muscle memory" and trust required to function as a unit during a real outage. This is where "signature strengths" (Karlsen & Berg, 2020) come into play. A project manager who exhibits "social intelligence" will recognize when their team is reaching the limit of its adaptive capacity and will intervene to prevent burnout-a major factor in human error (Joice et al., 1998; Onofrio et al., 2020).

Limitations and Future Directions

While the proposed synthesis offers a powerful framework, it is not without limitations. The application of Chaos Engineering to non-digital fields, such as "Chaos Engineering in a Database" (Tang & Weng, 2020) or in "Healthcare Risk Analysis" (Zheng et al., 2024), is still in its early stages. There is a risk of "methodological sprawl," where the tools of one discipline are applied to another without proper contextualization.

Future research should focus on the "Dependence Assessment" in human reliability analysis (Ji et al., 2024; Liu et al., 2018). Most HRA models assume that human errors are independent, but in a team setting, errors are highly correlated. Developing a "Cloud-Model based Best-Worst Method" (Ji et al., 2024) to quantify these dependencies would be a significant contribution. Furthermore, the use of "Pharmacy Surveillance Information Systems" (Abbasi et al., 2023) to monitor dispensing practices suggests that "automated surveillance" could be a double-edged sword: it provides data for resilience, but it can also lead to increased stress and "alert fatigue" for operators.

CONCLUSION

The journey from digital turbulence to human resilience is one of the most critical challenges of the modern age. This research has demonstrated that systemic reliability is not a static destination, but a dynamic, continuous process of experimentation, learning, and adaptation. By integrating the technical rigor of Chaos Engineering with the human-centric insights of Human Reliability Analysis and the strategic oversight of organizational governance, we can build systems that are not just "fault-tolerant," but "anti-fragile."

The synthesis of perspectives-from the "Art of Creating a Resilient System" (Pham, 2023) to the "World Cup winning" psychology of team resilience (Morgan et al., 2015)-reveals a universal truth: resilience is found in the "people in the loop." Technology provides the infrastructure, but humans provide the intelligence. Chaos Engineering serves as the catalyst that brings these two forces together, providing a structured, ethical, and highly effective way to prepare for the unexpected.

As we move toward an increasingly complex and automated future, the principles of controlled disruption must become a standard part of every engineering and management curriculum. We must foster a generation of "High-Reliability Engineering Teams" (Kesarpu, 2025) who are not afraid of failure, but who view every glitch, every crash, and every anomaly as an opportunity to grow stronger. In the end, the most resilient system is not the one that never breaks, but the one whose human operators are so well-trained, so well-supported, and so deeply trusted that they can turn any disruption into a triumph of engineering and human spirit.

REFERENCES

1. Abbasi, R., et al. Using pharmacy surveillance information systems to monitor the dispensing practice of under-controlled drugs: a qualitative study on necessities, requirements, and implementation challenges. Inform Med Unlocked (2023).
2. Allspaw, J. People in the loop. In: Rosenthal C., Jones N. (Eds.), Chaos engineering, O'Reilly, Beijing u.a. (2020), pp. 151-159.
3. Baraldi, P., et al. Comparing the treatment of uncertainty in bayesian networks and fuzzy expert systems used for a human reliability analysis application. Reliab Eng Syst Saf (2015).
4. Basiri, A., et al. Chaos engineering. IEEE Softw 33(3) (2016), pp. 35-41.
5. ElMaraghy, W., ElMaraghy, H., Tomiyama, T., & Monostori, L. Complexity in engineering design and manufacturing. CIRP Ann, 61 (2) (2012), pp. 793-814.
6. Faiella, G., et al. Expanding healthcare failure mode and effect analysis: a composite proactive risk analysis approach. Reliab Eng Syst Saf (2018).
7. Fletcher, D., Hanton, S., & Mellalieu, S. D. (2006). In: Hanton S., Mellalieu S.D. (Eds.), Literature reviews in sport psychology, New Nova Science: York Editors, pp. 321-374.

8. Hochstein, L., & Rosenthal, C. Chaos Engineering Panel. In: 2016 IEEE/ACM 38th international conference on software engineering companion. ICSE-C, 2016, p. 90–1.
9. Jani, A. Escalation of commitment in troubled IT projects: Influence of project risk factors and self-efficacy on the perception of risk and the commitment to a failing project. International Journal of Project Management, 29 (7) (2011), pp. 934-945.
10. Ji, C., et al. Dependence assessment in human reliability analysis based on cloud model and best-worst method. Reliab Eng Syst Saf (2024).
11. Joice, P., et al. Errors enacted during endoscopic surgery – a human reliability analysis. Appl Ergon (1998).
12. Karlsen, J. T., & Berg, M. E. A study of the influence of project managers' signature strengths on project team resilience. Team Performance Management, 26 (3/4) (2020), pp. 247-262.
13. Kaya, G. K., et al. Semi-quantitative application to the functional resonance analysis method for supporting safety management in a complex health-care process. Reliab Eng Syst Saf (2020).
14. Sagar Kesarpu. (2025). Chaos Engineering as a Learning Framework: A Human-Centered Model for Developing High-Reliability Engineering Teams. The American Journal of Engineering and Technology, 7(12), 57–64. https://doi.org/10.37547/tajet/Volume07Issue12-05
15. Khaleghi, P., et al. Identification and analysis of human errors in emergency department nurses using SHERPA method. Int Emerg Nurs (2022).
16. Liu, H.-C., et al. A large group decision making approach for dependence assessment in human reliability analysis. Reliab Eng Syst Saf (2018).
17. Morcov, S., Pintelon, L., & Kusters, R. Definitions, characteristics and measures of IT project complexity - a systematic literature review. International Journal of Information Systems and Project Management, 8 (2) (2020), pp. 5-21.
18. Morgan, P. B., Fletcher, D., & Sarkar, M. (2015). Understanding team resilience in the world's best athletes: A case study of a rugby union World Cup winning team. Psychology of Sport and Exercise, 16(1), 91–100.
19. Müller, R., Turner, R., Andersen, E. S., Shao, J., & Kvalnes, Ø. Ethics, Trust, and Governance in Temporary Organisations. Project Management Journal, 45 (4) (2014), pp. 39-54.
20. Musawir, A., Serra, C., Zwikael, O., & Ali, I. Project governance, benefit management, and project success: Towards a framework for supporting organisational strategy implementation. International Journal of Project Management, 35 (8) (2017), pp. 1658-1672.
21. Onofrio, R., et al. A methodology for dynamic Human reliability analysis in robotic surgery. Appl Ergon (2020).
22. Pawlikowski, M. Chaos engineering: Site reliability through controlled disruption (1st ed.), Manning, Shelter Island (2021).
23. Pham, Q. What is chaos engineering? The art of creating a resilient system: White paper. Orient Software (2023).
24. Robbins, J., Krishnan, K., Allspaw, J., & Limoncelli, T. A. Resilience engineering: Learning to embrace failure: A discussion with Jesse Robbins, Kripa Krishnan, John Allspaw, and Tom Limoncelli. Queue, 10 (9) (2012), pp. 20-28.
25. Rosenthal, C., Hochstein, L., Blohowiak, A., Jones, N., & Basiri, A. Chaos engineering: Building confidence in system behavior through experiments (1st ed.), O'Reilly, Beijing (2017).
26. Simsekler, M. C. E., et al. Integration of multiple methods in identifying patient safety risks. Saf Sci (2019).
27. Sujan, M., et al. Failure to rescue following emergency surgery: a FRAM analysis of the management of the deteriorating patient. Appl Ergon (2022).
28. Sujan, M., et al. How can large language models assist with a FRAM analysis? Saf Sci (2025).
29. Sujan, M., et al. What kinds of insights do Safety-I and Safety-II approaches provide? A critical reflection on the use of SHERPA and FRAM in healthcare. Saf Sci (2024).
30. Tang, L., & Weng, H. Chaos engineering on a database. In: Rosenthal C., Jones N. (Eds.), Chaos engineering, O'Reilly, Beijing u.a. (2020), pp. 237-247.
31. Tucker, H., Hochstein, L., Jones, N., Basiri, A., & Rosenthal, C. The business case for chaos engineering. IEEE Cloud Comput, 5 (3) (2018), pp. 45-54.
32. Zheng, Q., et al. A hybrid HFACS model using DEMATEL-ORESTE method with linguistic Z-number for risk analysis of human error factors in the healthcare system. Expert Syst Appl (2024).