Failure Cause Analysis

From Reactive RCA to Proactive MLCA

Avi Harel

Ergolight

October 2025

Abstract

Accidents in complex socio-technical systems rarely result from a single failure. They typically begin with a minor initiating trigger and escalate through interacting weaknesses across technical, human, and organizational layers. This article examines four analytical frameworks: RCA identifies triggers; Proactive RCA shifts the focus toward anticipating and disabling potential triggers; MLCA expands the analysis to reveal escalation enablers that amplify the event; and Proactive MLCA integrates both perspectives, aiming to anticipate and neutralize escalation enablers within the socio-technical system. Together, they represent a multidimensional capability for error protection, culminating in a proactive and integrative approach to resilience. Using the 1979 Three Mile Island (TMI) nuclear accident as a case example, the paper illustrates how event's small failures were amplified by exception enablers.

---

1. Introduction

Accidents in complex socio-technical systems rarely result from a single failure. They typically emerge from the interaction of multiple weaknesses distributed across technical, human, and organizational layers (Reason, 1997; Leveson, 2011). These interactions can transform a minor initiating event into a cascading failure, where defenses erode sequentially or simultaneously—a process often described as the "Swiss cheese" effect of layered vulnerabilities (Reason, 2000). Understanding how these interactions occur, and how they can be anticipated and prevented, is essential to developing resilient systems capable of operating safely under uncertainty (Hollnagel, 2014; Woods, 2015).

Traditional approaches to accident investigation have focused on identifying a "root cause," emphasizing linear causality and discrete fault isolation. Root Cause Analysis (RCA) has long served as a foundation for post-incident learning by tracing the sequence of technical or procedural failures that precipitated an event (Rooney & Vanden Heuvel, 2004). However, in highly coupled and tightly integrated systems, such as nuclear power, aviation, and healthcare, accidents cannot be fully explained through linear cause-and-effect reasoning alone (Perrow, 1999). Instead, they emerge from complex, adaptive interactions that evolve dynamically over time (Rasmussen, 1997; Dekker, 2011).

Building on these insights, Multi-Level Causal Analysis (MLCA) expands the analytical scope of RCA by recognizing that accident causation operates simultaneously across multiple system layers—ranging from physical components to organizational decision-making and regulatory oversight (Le Coze, 2019). MLCA provides a framework for examining weaknesses at one level, such as design flaws, inadequate procedures, or cognitive overload, can combine and interact across other levels, creating pathways for escalation.

In parallel, the evolution from reactive to proactive safety management has introduced two further developments: Proactive RCA and Proactive MLCA. Proactive RCA retains the causal logic of traditional RCA but shifts its orientation toward anticipation, identifying and neutralizing potential triggers before they manifest in operation (Hollnagel, 2018). Proactive MLCA, in turn, integrates the multi-level perspective with this anticipatory stance, focusing not only on preventing initial faults but also on disabling escalation enablers, the conditions that allow small anomalies to grow into major failures (Woods, 2015; Leveson, 2020).

The 1979 Three Mile Island (TMI) nuclear accident provides a revealing case study for illustrating the complementary roles of these four analytical perspectives. The initiating trigger in that event was a feedwater malfunction that interrupted heat removal. Escalation followed when the backup feedwater pump—disabled for maintenance—was unavailable, and when the Pressure-Operated Relief Valve (PORV) stuck open, causing coolant loss. These immediate failures were compounded by ambiguous instrumentation and operator misinterpretation of system indicators (Kemeny Commission, 1979; Reason, 1997).

Through the TMI case, the present article explores how RCA, MLCA, Proactive RCA, and Proactive MLCA together represent an evolving capability to protect complex systems from error. Each framework embodies a different mode of inquiry: diagnostic, systemic, anticipatory, and integrative, that contributes to building resilient performance. Their synthesis demonstrates progression from reactive learning toward proactive foresight,

marking a critical shift in how safety is conceived, managed, and sustained in modern socio-technical systems.

2. Conceptual Foundations

Understanding failures in complex socio-technical systems requires analytical approaches that go beyond identifying a single cause. Four frameworks are discussed in this article: Root Cause Analysis (RCA), Multi-Level Causal Analysis (MLCA), Proactive RCA, and Proactive MLCA. Together, they offer complementary lenses for examining accidents and enhancing system resilience. Each framework contributes distinct capabilities for protecting systems from error, and together they form an integrated toolkit for both reactive diagnosis and proactive prevention.

Root Cause Analysis (RCA)

RCA provides a foundational perspective. Traditionally, RCA focuses on identifying the initiating trigger of a failure. Its strength lies in the clarity it provides by pinpointing the immediate cause, organizations can implement corrective actions to prevent the same fault from recurring. Yet, RCA is inherently reactive. While it can stop the same trigger from causing future events, it offers limited insight into how minor disturbances might evolve into large-scale accidents through system interactions.

Multi-Level Causal Analysis (MLCA)

MLCA extends the focus of analysis beyond the immediate trigger, examining the network of factors: technical, cognitive, procedural, organizational, and regulatory, that can enable a fault to escalate. By revealing these latent escalation enablers, MLCA provides a broader understanding of system vulnerabilities and offers opportunities for corrective action that go beyond the initial trigger. Unlike RCA, which is largely linear and event-focused, MLCA emphasizes systemic interactions and the conditions under which minor faults can propagate into major incidents.

Proactive RCA

Proactive RCA expands the temporal dimension of analysis. Rather than waiting for a failure to occur, Proactive RCA anticipates potential initiating faults and seeks to prevent them before they manifest. By focusing on known vulnerabilities in design, procedures, or operations, it enhances the system's ability to protect itself from error. However, its scope remains primarily at the trigger level; Proactive RCA does not systematically address conditions that might allow a fault to escalate once it occurs.

Proactive MLCA

Proactive MLCA represents the integration of these two extensions. It combines the systemic breadth of MLCA with the anticipatory perspective of Proactive RCA, aiming not only to prevent initiating triggers but also to neutralize escalation enablers before they contribute to accident propagation. In this way, Proactive MLCA provides the most comprehensive error protection, strengthening resilience across technical, human, procedural, and organizational dimensions.

Summary

Taken together, these frameworks represent a multidimensional approach to understanding, preventing, and containing accidents in complex socio-technical systems, offering progression from reactive diagnosis toward proactive, integrative resilience.

3. The Three Mile Island Case

The 1979 Three Mile Island (TMI) Unit 2 accident remains one of the most studied nuclear incidents, providing a rich illustration of how minor initiating faults can escalate into major system failures. The event began with a relatively small disturbance: a malfunction in the secondary feedwater system interrupted heat removal from the reactor core. Alone, this disturbance was manageable. However, the system's design and operational context allowed the situation to escalate rapidly, demonstrating how interacting weaknesses across multiple levels can amplify the consequences of minor faults.

Root Cause Analysis (RCA)

Applying Root Cause Analysis (RCA) to TMI highlights the initial trigger: the feedwater interruption. RCA identifies this fault clearly and allows engineers to implement corrective measures to prevent the same specific malfunction from recurring. Yet RCA alone does not explain why this minor disturbance escalated into a partial core meltdown. It provides insight into what went wrong at the moment the system first deviated from normal operation, but it leaves the broader systemic vulnerabilities unexamined.

Multi-Level Causal Analysis (MLCA)

Multi-Level Causal Analysis (MLCA) extends the lens to reveal these vulnerabilities. In the TMI case, the backup feedwater pump, which could have compensated for the primary pump's failure, had been disabled for maintenance and was unavailable when needed. Later, the Pilot-Operated Relief Valve (PORV) stuck open, allowing coolant to escape from

the reactor pressure vessel. MLCA exposes these and other escalation enablers: unclear instrumentation, alarm overload in the control room, operator misinterpretation, and latent organizational and procedural weaknesses. By mapping the network of technical, human, and organizational factors that allowed the initial disturbance to propagate, MLCA explains why a manageable fault escalated into a serious accident.

Proactive RCA

While RCA and MLCA are reactive, Proactive RCA demonstrates how anticipation can prevent similar triggers from arising in the first place. For example, improved monitoring and maintenance procedures for the feedwater system, along with more rigorous preemptive inspections and redundancy checks, could have reduced the likelihood of the initial feedwater interruption. By addressing vulnerabilities before a fault occurs, Proactive RCA enhances the system's protective capability at the trigger level.

Proactive MLCA

Proactive MLCA represents the most comprehensive approach, integrating both systemic insight and forward-looking prevention. In the context of TMI, Proactive MLCA would address the availability of the backup feedwater pump, improve control room instrumentation to prevent misinterpretation, implement hierarchical alarm management to reduce operator overload, and strengthen training and organizational procedures to limit the propagation of errors. By neutralizing escalation enablers before a failure occurs, Proactive MLCA enhances resilience not only by preventing triggers but also by limiting the system's susceptibility to cascading failures.

Summary

Viewed together, the TMI accident illustrates how each analytical framework contributes to system protection. RCA identifies triggers, MLCA maps escalation enablers, Proactive RCA prevents known initiating faults, and Proactive MLCA integrates both dimensions to preemptively mitigate both triggers and pathways of escalation. This case demonstrates the value of a multidimensional approach: while reactive frameworks provide learning from past events, proactive and multi-level analyses equip organizations with tools to anticipate, contain, and ultimately prevent accidents in complex socio-technical systems.

4. Comparative Analysis of Error Protection

The Three Mile Island (TMI) case demonstrates that analytical frameworks differ not only in what they identify but also in how they protect systems from error. Each framework offers a

distinct dimension of insight, and together they provide a multidimensional strategy for understanding and preventing accidents.

Root Cause Analysis (RCA)

RCA focuses squarely on the initiating fault. By identifying the feedwater malfunction at TMI, RCA provides the means to correct specific failures and prevent them from recurring. Applying Root Cause Analysis (RCA) to TMI focuses on identifying all these immediate triggers. RCA clarifies that the accident hinges on several faults:

1. The feedwater interruption,
2. The unavailable backup pump,
3. The faulty PORV sensor, and
4. The misleading instrumentation

All these faults served as triggers that directly contributed to the initial deviation from safe operation. By isolating these initiating faults, RCA provides a reactive foundation for corrective action (Leveson, 2011; Rooney & Vanden Heuvel, 2004). By identifying all triggers, RCA provides a comprehensive diagnostic foundation, enabling corrective measures to prevent recurrence of the same faults. The strength of RCA lies in clarifying the immediate causal sequence, enabling targeted interventions to prevent recurrence of the same faults. However, RCA's retrospective and linear approach cannot fully explain why multiple minor faults interacted to produce a major incident (Perrow, 1999; Reason, 1997).

Multi-Level Causal Analysis (MLCA)

The Three Mile Island (TMI) accident illustrates how a combination of technical, human, procedural, and organizational factors enabled escalation from what began as a manageable disturbance. Multi-Level Causal Analysis (MLCA) allows us to trace how each key trigger was amplified through these interacting layers of weakness, turning a sequence of small faults into a major system failure.

Feedwater Malfunction

The first trigger was not catastrophic in itself. Loss of feedwater is a design-basis event that operators are trained to manage. However, escalation occurred because the technical and procedural environment did not support timely detection or control. The monitoring systems lacked the sensitivity and immediacy to alert operators before reactor parameters began to drift. The operators, accustomed to routine feedwater transients, initially

underestimated the severity of the situation. Procedurally, the absence of clear, prioritized diagnostic steps led to hesitation, and organizationally, the maintenance and workload arrangements left the operating team without sufficient capacity to interpret multiple concurrent alarms. In this sense, the escalation enablers were embedded not only in the equipment but in the cognitive and procedural environment surrounding its use.

Disabled Backup Feedwater Pump

The second trigger reflected deeper systemic weaknesses. The technical design permitted both feedwater trains to be unavailable simultaneously, without built-in interlocks or configuration safeguards. Human and procedural factors compounded this vulnerability: maintenance personnel assumed that temporary isolation of the backup pump was low risk, and no mechanism required operations to reassess overall system readiness after the maintenance action. The organizational and regulatory layers reinforced these weaknesses, as scheduling practices prioritized efficiency over redundancy and oversight processes tolerated temporary loss of redundancy without immediate compensatory measures. The result was a latent escalation pathway: when the feedwater malfunction occurred, the safety margin had already been quietly eroded.

Pressure-Operated Relief Valve (PORV) State Sensor Malfunction

The Pressure PORV created another layer of escalation. Technically, the valve performed its initial function correctly: it opened in response to rising pressure, but failed to reseat when the pressure dropped. The control-room indicator, based on electrical logic rather than direct valve position, falsely showed the valve as closed. This technical shortcoming was compounded by human reliance on the displayed signal and by procedural ambiguity regarding how to verify valve state under uncertain indications. The organization's instrumentation philosophy, simplified, single-channel indications deemed sufficient for operator awareness, created the conditions for misinterpretation. The regulatory environment at the time did not require redundant or independent verification of critical valve positions, reinforcing vulnerability. When combined, these layers produced a powerful escalation enabler: a continuing coolant loss masked by misleading information.

Confusion in Temperature and Pressure Readings

The final trigger demonstrated the human-system interaction failures that characterize complex incidents. Instrumentation inconsistencies and sensor noise created contradictory cues about the reactor's state. Technically, the system lacked a coherent integration of signals that could present operators with a unified picture of plant behavior. Cognitively, the operators were overloaded by alarms and conflicting data, leading to

tunnel vision and misinterpretation. Procedurally, alarm management rules and diagnostic aids were insufficiently structured to guide decision-making under stress. Organizationally, training emphasized adherence to procedural norms rather than adaptive reasoning in ambiguous conditions. These interacting layers transformed uncertainty into misjudgment, allowing escalation to continue unchecked.

Summary

Across all four triggers, MLCA reveals that escalation at TMI resulted not from isolated faults but from the coupling of technical deficiencies with human, procedural, and organizational weaknesses (Le Coze, 2019). The feedwater malfunction was magnified by poor situational awareness; the disabled backup pump by inadequate configuration control; the PORV malfunction by unreliable feedback design; and the instrumentation confusion by cognitive overload and training limitations. Each layer, taken alone, might have been manageable. Together, they created a system that lacked the capacity to absorb disturbance or to detect its own degradation.

MLCA exposes these escalation enablers across technical, human, procedural, and organizational layers, highlighting the systemic pathways through which minor faults propagate into major failures (Woods, 2015). The effects of each trigger were magnified by procedural gaps, unclear instrumentation, alarm overload, and organizational factors that complicated operator decision-making (Dekker, 2011; Rasmussen, 1997). By illuminating these interacting layers, MLCA demonstrates that the real hazard lies not in the trigger itself, but in the system's latent pathways of escalation, shaped by design philosophy, operational practice, and organizational culture. This understanding sets the stage for Proactive MLCA, which extends the same multi-layer perspective forward in time, aiming not merely to identify enablers but to neutralize them before they combine into failure.

Multi-Level Causal Analysis (MLCA), in contrast, analyzes **how triggers interact with systemic weaknesses to amplify consequences** (Le Coze, 2019). Unlike RCA, MLCA does not stop at identifying triggers but explains how interactions between layers create vulnerabilities.

Proactive RCA

Proactive RCA shifts the focus from diagnosing past failures to anticipating and preventing the triggers that initiate unsafe conditions. It integrates both anticipatory and systemic perspectives, seeking to prevent triggers and neutralize escalation enablers simultaneously (Hollnagel, 2018; Weick & Sutcliffe, 2015).

In the Three Mile Island (TMI) accident, several immediate triggers contributed to the deviation from safe operation. A proactive analysis of these triggers highlights measures that could have prevented the incident at its origin.

Feedwater Malfunction

The initiating disturbance was a malfunction in the secondary feedwater system, which interrupted heat removal from the reactor core. Proactive RCA would emphasize preventive maintenance, redundancy, and continuous monitoring. Regular inspection and testing of feedwater pumps, automated flow verification, and rapid automatic transfer to standby systems could have either prevented this failure or mitigated its impact before the reactor parameters deviated from normal operation.

Disabled Backup Feedwater Pump

The backup feedwater pump, designed as a redundancy safeguard, had been disabled for maintenance and was unavailable when the primary feedwater system failed. Proactive RCA would target maintenance planning and procedural control (Leveson, 2020; Dekker, 2011). Preventive measures might include ensuring that critical redundancy is never compromised, applying lockout/tagout rules that require functional overlap, and implementing a configuration management check before returning the unit to operation. Even a simple procedural verification ensuring that at least one feedwater train remained operational would have prevented this condition.

Pressure-Operated Relief Valve (PORV) State Sensor Malfunction

When system pressure rose, the PORV opened as designed but failed to close afterward. The control room indicator, however, erroneously showed it as closed, thus misleading the operators. Proactive RCA would focus on instrument reliability, status verification, and fail-safe feedback design. Using independent pressure feedback signals, redundant valve position sensors, or diagnostic self-tests could have exposed the discrepancy in real time. These measures would prevent false indications and ensure operators are alerted when valve position and pressure behavior diverge.

Confusion in Temperature and Pressure Readings

Operators faced inconsistent readings from temperature and pressure instruments, leading to incorrect assumptions about system state. Proactive RCA highlights human-system integration and interface clarity as essential to preventing such confusion. Improvements could include harmonized sensor calibration, standardized display hierarchies, intuitive alarm grouping, and simulator-based training that prepares operators

to interpret conflicting data under stress. Such measures reduce cognitive overload and enhance situational awareness.

Summary

Taken together, these proactive RCA interventions would form a first line of defense against system upset by eliminating or mitigating potential triggers. In the TMI context, this approach targets both technical and human vulnerabilities that initiated the sequence of failure, illustrating that prevention must encompass all immediate triggers, not just the first event in the chain. By addressing the origins of faults before they occur, Proactive RCA strengthens operational resilience and reduces the likelihood that small, independent failures evolve into large-scale accidents (Hollnagel, 2014; Woods, 2015).

Proactive MLCA

While MLCA exposes how latent weaknesses across technical, human, procedural, and organizational layers can amplify a disturbance, Proactive MLCA extends this perspective into the domain of foresight. Rather than analyzing escalation after it has occurred, Proactive MLCA seeks to anticipate how combinations of small weaknesses might interact under plausible scenarios. Its aim is not only to prevent initiating triggers, but also to preempt the mechanisms that allow those triggers to grow into crises. Applied to the TMI case, this perspective reveals a path toward systemic resilience that could have neutralized the accident's escalation potential long before the event itself.

Feedwater Malfunction

A proactive multi-layer analysis of the feedwater system would have identified its vulnerability to concurrent maintenance and operational disruptions. From a technical standpoint, modeling of feedwater system dependencies could have exposed how the loss of one train reduces redundancy and how sensor delays affect operator recognition time. Human-factors analysis might have revealed the cognitive gap in how operators perceive feedwater transients, viewing them as routine rather than potentially critical. Procedurally, proactive review could have emphasized diagnostic clarity, ensuring that feedwater anomalies trigger immediate verification steps. Organizationally, a proactive stance would institutionalize periodic cross-functional reviews of maintenance and operations coordination, treating system readiness as a shared, continuously monitored attribute rather than an assumed condition.

Disabled Backup Feedwater Pump

The backup feedwater pump scenario demonstrates how Proactive MLCA could have preempted escalation by embedding resilience into maintenance and scheduling practices. A layered foresight analysis would have explored how routine maintenance interacts with operational readiness, using system-level models to simulate degraded configurations. Such an approach could have revealed that disabling the pump, even temporarily, created a single-point vulnerability with no compensating safeguards. The proactive intervention here would not be limited to issuing new rules; it would involve designing technical and organizational controls, such as automated configuration checks, interlocks preventing full redundancy loss, and decision-support alerts for risk-based maintenance scheduling. Through these mechanisms, potential escalation paths would be closed before they could align.

Pressure-Operated Relief Valve (PORV) State Sensor Malfunction

The Pressure-Operated Relief Valve (PORV) case exemplifies how proactive, cross-layer thinking could have mitigated informational and design weaknesses that later proved critical. From a technical perspective, Proactive MLCA would encourage analysis of the valve's indication logic under failure modes, identifying the misleading nature of the "closed" signal. Human-system integration assessments would have shown that operators were likely to rely on such feedback in high-pressure situations, highlighting the need for redundant or physically verified indications. Procedurally, it would promote the development of verification checklists and simulation-based training that expose operators to ambiguous signal conditions. At the organizational level, a proactive safety culture would have prioritized investment in diagnostic transparency, recognizing that uncertainty in critical indications represents not a nuisance but a direct operational risk.

Confusion in Temperature and Pressure Readings

The confusion over temperature and pressure readings underscores the need for proactive management of cognitive and informational complexity. A Proactive MLCA approach would analyze the human–machine interface as a dynamic system, examining how alarm density, display design, and signal coherence affect situational awareness. Technically, it would call for integrated displays capable of reconciling sensor discrepancies and presenting operators with synthesized state information rather than raw data streams. Procedurally, it would emphasize cognitive resilience, training operators not only to follow rules but to navigate uncertainty through model-based reasoning. Organizationally, it would institutionalize learning cycles that continually refine alarm management, interface design, and simulation fidelity based on operator feedback and near-miss analysis.

Summary

Taken together, these interventions represent more than preventive maintenance of parts or procedures: they amount to the engineering of foresight within the socio-technical system. Proactive MLCA redefines safety as a dynamic capability: the capacity of a system to recognize, absorb, and adapt to weak signals before they combine into threats. By examining how layers interact under stress, it provides a structured basis for designing interventions that are both technically sound and organizationally sustainable.

Applied broadly, Proactive MLCA transforms system protection from a reactive exercise in blame assignment to a strategic discipline of anticipatory design. In the TMI context, it suggests that the accident could have been prevented not merely by correcting the feedwater or valve design, but by cultivating a system-wide awareness of how technical, human, and organizational elements coevolve. Through this lens, resilience emerges not from redundancy alone but from the integration of foresight across all system layers—an essential capability for the sustainability of modern socio-technical operations.

Integrative Summary

The four analytical frameworks: RCA, MLCA, Proactive RCA, and Proactive MLCA, represent a progressive evolution in how complex systems conceptualize and defend against error. Each framework embodies a distinct cognitive stance toward failure: from reaction to anticipation, from linear causality to systemic interaction, and from component reliability to organizational resilience. Examined together through the lens of the Three Mile Island (TMI) incident, they trace a trajectory from understanding failure to engineering foresight.

Toward a Unified Framework for Error Protection

Root Cause Analysis (RCA) begins with the recognition that every event has identifiable triggers. In the TMI case, these included the initial feedwater interruption, the unavailability of the backup feedwater pump, the misleading indication of the Pressure-Operated Relief Valve (PORV), and the confusion created by inconsistent instrumentation. RCA's strength lies in its disciplined identification of these triggers and their immediate effects, allowing organizations to eliminate direct fault sources. Yet, its focus remains confined to the proximate layer of causation, the "what happened" rather than the "how and why the system allowed it to happen."

Multi-Level Causal Analysis (MLCA) expands this perspective by uncovering the cross-layer interactions that enabled escalation. It shifts attention from individual faults to systemic

relationships, how technical design, operator cognition, procedures, and organizational context intertwine under stress. Through MLCA, the TMI incident is reframed as a systemic failure of coupling and communication, where local weaknesses in design and human–machine interfaces resonated with procedural and cultural vulnerabilities to magnify a minor disturbance into a crisis. MLCA thus reframes safety as an emergent property of interdependence, not isolation.

Proactive RCA advances the logic of prevention by focusing on potential triggers rather than past ones. Its objective is to foresee and disable those initiating mechanisms before they activate. Applied to TMI, it translates into ensuring the readiness of feedwater systems, verifying redundancy configurations, improving sensor accuracy, and validating operator interfaces. The preventive mindset inherent in Proactive RCA embodies a transition from learning from accidents to learning before accidents occur.

Proactive MLCA represents the most integrative and mature stage of this evolution. It merges the anticipatory orientation of Proactive RCA with the systemic depth of MLCA. Instead of seeking a single root cause or a specific trigger, it examines how diverse system layers might combine to permit escalation, and how those combinations can be neutralized in advance. In the TMI context, it would involve continuous cross-layer modeling of operational resilience, real-time monitoring of system dependencies, and organizational processes that ensure foresight is both distributed and actionable.

Together, these four frameworks form a continuum of protective capability. RCA ensures learning from failure; MLCA reveals the mechanisms of escalation; Proactive RCA anticipates and disables potential triggers; and Proactive MLCA builds a culture and infrastructure of foresight capable of interrupting escalation before it begins. Their relationship is not hierarchical but evolutionary and integrative: each framework enriches the others, expanding the system's ability to detect, interpret, and respond to weak signals.

In this sense, the movement from RCA to Proactive MLCA marks a transformation in how safety is conceived. What begins as an analytical tool for diagnosing accidents becomes a strategic discipline for sustaining safe performance under uncertainty. The TMI case, revisited through this lens, is not merely a story of technological failure but a demonstration of the need for systemic anticipation, an organizational capability that transforms error protection into resilience.

Conclusions and Implications

The examination of the Three Mile Island (TMI) accident through the combined lenses of RCA, MLCA, Proactive RCA, and Proactive MLCA underscores a crucial evolution in how complex socio-technical systems understand and manage safety. What began as a search for individual technical faults has matured into a systemic inquiry into how organizations can cultivate foresight, adaptability, and resilience in the face of uncertainty (Reason, 1997; Hollnagel, 2014).

Traditional Root Cause Analysis (RCA) remains an indispensable tool for learning from failure, but it is inherently retrospective and reductionist (Leveson, 2011). It seeks closure by identifying discrete triggers, which, while informative, rarely capture the interdependent nature of modern system operations. Multi-Level Causal Analysis (MLCA) extends the field of view, illuminating how interactions across technical, human, procedural, and organizational layers transform localized faults into system-wide failures (Le Coze, 2019). This layered understanding reframes safety not as the absence of error, but as the dynamic management of complexity within socio-technical systems (Perrow, 1999; Rasmussen, 1997).

Proactive RCA and Proactive MLCA together shift this paradigm from diagnosis to prevention. Proactive RCA anticipates initiating triggers, seeking to disable faults before they occur, aligning with the principles of predictive and learning-based safety management (Hollnagel, 2018). Proactive MLCA, in turn, anticipates the enablers of escalation, those latent configurations and cross-layer couplings that allow small disturbances to grow. In doing so, it transforms system safety from a reactive defense into an active capability: the capacity to sense, interpret, and neutralize weak signals before they accumulate into failure (Weick & Sutcliffe, 2015; Woods, 2015).

The TMI accident illustrates how each framework contributes a distinct layer of protection. RCA explains why the event began; MLCA explains how it grew; Proactive RCA shows how it could have been prevented at its origin; and Proactive MLCA reveals how it could have been contained by disrupting the mechanisms of escalation. Together, they form an integrative model for error protection that unites reactive learning and proactive foresight within a single continuum of organizational resilience (Dekker, 2011; Hollnagel, 2014).

The implications for modern engineering are profound. As systems grow more tightly coupled, adaptive, and AI-augmented, the limits of retrospective analysis become more pronounced. Safety management must evolve beyond compliance and error counting toward anticipatory governance, a mindset that treats foresight as a core design function

(Leveson, 2020). Proactive MLCA provides such a pathway, embedding anticipation into the design, operation, and governance of socio-technical systems.

Ultimately, protecting systems from error requires more than eliminating faults; it demands an organizational and cognitive transformation. It calls for systems that continuously learn, simulate, and self-examine; for engineers and operators who understand not only how systems fail, but how they drift; and for leadership cultures that reward curiosity, questioning, and reflection as much as technical precision. The evolution from RCA to Proactive MLCA thus represents not just a methodological advance, but a paradigm shift, from reacting to what has gone wrong, to preparing for what could (Hollnagel, 2018; Woods, 2015).

References

Dekker, S. (2011). *Drift into Failure: From Hunting Broken Components to Understanding Complex Systems.* Ashgate.

Hollnagel, E. (2014). *Safety-I and Safety-II: The Past and Future of Safety Management.* Ashgate.

Hollnagel, E. (2018). *Safety-II in Practice: Developing the Resilience Potentials.* Routledge.

Kemeny Commission. (1979). *Report of the President's Commission on the Accident at Three Mile Island.* U.S. Government Printing Office.

Le Coze, J.-C. (2019). *Post Normal Accident Thinking: Complexity, Epistemology and Safety. Safety Science, 118*, 430–443.

Leveson, N. G. (2011). *Engineering a Safer World: Systems Thinking Applied to Safety.* MIT Press.

Leveson, N. G. (2020). *Systems-Theoretic Accident Model and Processes (STAMP) Handbook.* MIT Press.

Perrow, C. (1999). *Normal Accidents: Living with High-Risk Technologies.* Princeton University Press.

Rasmussen, J. (1997). Risk management in a dynamic society: A modelling problem. *Safety Science, 27*(2–3), 183–213.

Reason, J. (1997). *Managing the Risks of Organizational Accidents.* Ashgate.

Reason, J. (2000). Human error: Models and management. *BMJ, 320*(7237), 768–770.

Rooney, J. J., & Vanden Heuvel, L. N. (2004). Root cause analysis for beginners. *Quality* Weick, K. E., & Sutcliffe, K. M. (2015). *Managing the Unexpected: Sustained Performance in a Complex World* (3rd ed.). Wiley.

Weick, K. E., & Sutcliffe, K. M. (2015). *Managing the Unexpected: Sustained Performance in a Complex World* (3rd ed.). Wiley.

Woods, D. D. (2015). Four concepts for resilience and the implications for the future of resilience engineering. *Reliability Engineering & System Safety, 141*, 5–9.