

CMDB Data Governance and Business Continuity: Identifying Research Gaps in AI-Driven IT Operations

¹Ashutosh Palia

¹Senior Web Application Developer, MathWorks Inc

Abstract: The Configuration Management Database (CMDB) is widely positioned as the authoritative record of IT infrastructure and its relationships to business services, yet industry surveys consistently report accuracy levels around 60 percent and find that a majority of CMDB implementations fail to meet their intended objectives. This gap between the CMDB's intended role and its practical reliability has been studied from several angles: general data governance frameworks, machine learning approaches to configuration data quality, and, increasingly, artificial intelligence integration within IT service management platforms. This paper reviews that literature and argues that one consequential intersection remains comparatively under-examined: the relationship between CMDB data governance and Business Continuity and Disaster Recovery (BCDR) outcomes, particularly as AI-driven and increasingly agentic IT operations begin to act on CMDB data with reduced human review. Drawing on recent work in CMDB-focused machine learning, AI integration in IT service management, general data governance frameworks, business continuity and cyber-resilience research, and the longer-standing literature on trust in automation, this review identifies a structural gap: existing work addresses CMDB data quality in general terms, or business continuity in general terms, but rarely connects specific CMDB governance dimensions to specific continuity and recovery outcomes. This is a targeted, narrative literature review rather than a systematic one, and its contribution is a research agenda rather than a new empirical study: an explicit characterization of this gap, expressed as a set of governance dimensions and framework requirements capable of addressing it, intended as the basis for continued research.

Keywords: Configuration Management Database, CMDB, Data Governance, Business Continuity, Disaster Recovery, IT Service Management, ITIL, Artificial Intelligence, AI-Driven IT Operations, Agentic IT Operations, Trust in Automation.

1. Introduction

Organizations increasingly depend on the Configuration Management Database as the authoritative source of truth for what infrastructure and applications exist and how they relate to one another and to business services. In principle, the CMDB underlies incident management, change impact analysis, security operations, and business continuity planning. In practice, CMDB data quality remains a persistent and well-documented weakness of enterprise IT operations. Industry analysis places average CMDB accuracy at approximately 60 percent across organizations [1], and a majority of CMDB implementations are reported to fail to deliver on their intended objectives [2], [3], typically because governance and reconciliation processes have not kept pace with the scale and rate of change in modern IT

environments.

This is not a new observation, and a substantial body of practitioner and academic work has already examined pieces of the problem: how data governance should be structured generally [4], how machine learning can improve specific CMDB quality problems such as naming conformity or relationship reconciliation [5], [6], and how artificial intelligence is being integrated into IT service management platforms more broadly [7], [8]. What has received comparatively little direct attention is the specific downstream consequence of CMDB data quality for business continuity and disaster recovery (BCDR) operations, where the cost of a governance gap is not a reporting inconvenience but a structurally incomplete or incorrectly sequenced recovery plan.

This gap has become more consequential, not less,

with the emergence of AI-driven and increasingly agentic IT operations. Historically, a human consulting an imperfect CMDB record could apply judgment to compensate for gaps or staleness in the data. As remediation, orchestration, and recovery actions move toward automation, that compensating layer of human judgment is reduced, and the reliability of the underlying CMDB data becomes a more direct determinant of operational outcomes.

This review is organized around three research questions that Section 5 returns to explicitly: (RQ1) which CMDB data-quality dimensions are most relevant to business continuity and disaster recovery outcomes specifically, as opposed to CMDB use cases in general; (RQ2) how CMDB governance practices influence recovery outcomes, given that dependency mapping is already established as a prerequisite for recovery planning; and (RQ3) how AI-driven and increasingly agentic IT operations change the risk associated with inaccurate or poorly governed CMDB data. Addressing these questions also requires distinguishing more carefully among terms that this literature often uses loosely: the CMDB as a data repository (the stored configuration item records themselves), the configuration management process (the operational discipline of creating, verifying, and maintaining those records), service mapping (the representation of applications and the technical components that compose them), and dependency mapping (the broader representation of how systems, applications, and data relate to the business functions and continuity requirements they support). This paper uses these four terms in that specific sense throughout, and treats CMDB governance as encompassing all four rather than the data repository alone. This paper also distinguishes AI-driven IT operations, the general use of AI or machine learning techniques to support IT operations tasks, from agentic IT operations, the narrower and more consequential subset in which such systems execute actions with reduced or no human review; because the trend runs from the former toward the latter, this paper consistently uses the paired phrase “AI-driven and increasingly agentic IT operations” when referring to that trend, reserving “agentic IT operations” alone for statements specific to the higher-autonomy subset.

This paper reviews the relevant literature across four areas: general data governance frameworks, machine learning and AI approaches to CMDB data quality,

business continuity and cyber-resilience research, and the trust-in-automation literature that informs how automated systems should be designed to rely on data of uncertain quality. Section 2 describes the review methodology. Section 3 presents this review. Section 4 synthesizes the review into an explicit characterization of the gap between current practice and research. Section 5 discusses implications for practitioners and researchers, including the research questions posed above. Section 6 outlines limitations, and Section 7 concludes with directions for the framework-building work this review is intended to motivate.

2. Review Methodology

This is a targeted, narrative literature review rather than a systematic review conducted against a pre-registered protocol such as PRISMA, and it is labeled as such rather than implying a more exhaustive process than was actually followed. Sources were identified through targeted searches of academic databases and search engines, including IEEE Xplore, ACM Digital Library, SpringerLink, and Google Scholar, using search terms built around the paper's core concepts (“CMDB,” “configuration management database,” “data governance,” “business continuity,” “disaster recovery,” “AIOps,” “IT service management,” and “trust in automation”), supplemented by citation chasing from an initial set of directly relevant papers and by the author's professional familiarity with enterprise IT service management and ServiceNow-based CMDB environments.

Inclusion was governed by relevance rather than by a fixed date range or a minimum venue ranking: sources were selected for direct bearing on the intersection of CMDB governance, CMDB-focused machine learning, business continuity and disaster recovery, and trust in automation, spanning foundational work from 1997 through current 2026 publications. Because academic coverage of CMDB governance specifically is thin, as Section 3 documents, this review also draws on industry and vendor publications denoted throughout as industry evidence where they supply measurement data or practitioner claims not otherwise available in peer-reviewed sources; these are treated as a distinct evidentiary category from peer-reviewed academic work throughout, and are never used in place of peer-reviewed evidence where such evidence exists. Section 6 discusses the limitations of this targeted, non-exhaustive approach

in more detail.

3. Literature Review

3.1. CMDB Data Quality: Scope and Persistent Gaps

Industry measurement of CMDB data quality reported through vendor and practitioner publications rather than peer-reviewed research, and treated here as industry evidence rather than as equivalent to peer-reviewed findings consistently points to the same conclusion from different vantage points. ReadyWorks [1] reports that CMDB accuracy across most enterprise environments sits close to 60 percent, and attributes this largely to reconciliation processes that cannot keep pace with the volume of infrastructure and application change in modern environments. ScienceLogic [2] frames the same problem from an operational tooling perspective, arguing that real-time synchronization between monitored infrastructure and the CMDB record is necessary but rarely achieved in practice, leaving organizations to operate from a CMDB that is perpetually out of date relative to the live environment. Virima [3] proposes a composite CMDB health score intended to make this gap visible and measurable rather than anecdotal, combining completeness, accuracy, and relationship quality into a single assessable metric.

Taken together, these sources establish that CMDB data quality is a widely recognized and reasonably well-quantified problem at the aggregate level. What they do not address is where, specifically, the inaccuracy concentrates, or which classes of data quality failure matter most for which downstream use cases. An aggregate accuracy figure of 60 percent could reflect a fairly uniform shortfall across the environment, or a much sharper shortfall concentrated in a specific, high-criticality subset of infrastructure and the practical consequences of those two scenarios differ substantially, particularly for continuity planning, which depends disproportionately on the accuracy of a relatively small set of Tier-0 and Tier-1 dependency relationships.

3.2. Data Governance Frameworks and Their Applicability to CMDB

Beyond CMDB-specific measurement, a broader data governance literature offers structured frameworks for reasoning about data quality generally. The DAMA Data Management Body of Knowledge [4] remains the most widely referenced general-purpose framework,

organizing data management into a set of interconnected knowledge areas with data governance positioned at the center, and defining data quality along dimensions including accuracy, completeness, consistency, timeliness, validity, and uniqueness. This framework is deliberately technology- and domain-agnostic, which is both its strength and its limitation when applied to CMDB governance specifically: it supplies vocabulary and structure but offers no domain-specific guidance about which dimensions matter most for a CMDB supporting IT operations, security, or continuity use cases, or how those priorities should shift depending on the downstream consumer of the data.

This generality is a recurring pattern across the broader data governance literature: frameworks are built to apply uniformly across data domains, from customer records to financial data to operational telemetry, and CMDB governance is typically treated as one instance of a general data quality problem rather than as a domain with its own distinct risk profile. This paper argues in Section 4 that this generality is precisely where the literature's coverage of CMDB governance is thinnest.

3.3. Artificial Intelligence and Machine Learning Approaches to CMDB Quality

A more recent and rapidly growing strand of work applies machine learning and AI techniques directly to CMDB data quality problems. Dande et al. [5] provide one of the first structured reviews connecting artificial intelligence to the ITIL configuration management process specifically, observing that despite growing vendor interest in AI-augmented IT service management tooling, academic and industry research on the CMDB process itself remains limited, and arguing that this is an underexamined area given how central the CMDB is to IT service management as a whole.

Darreddy and Munagala [6] address this gap with a more technically developed contribution, proposing a graph machine learning approach to CMDB health management and service graph reconciliation for large-scale ServiceNow IT Operations Management deployments. Their approach combines transformer-based entity resolution for duplicate detection, graph neural networks for relationship anomaly detection, and reinforcement learning for autonomous configuration item remediation, and reports measurable improvements in relationship accuracy, completeness, and duplicate reduction at enterprise scale.

Niewiadomski and Mzyk [7] address a narrower technical problem within the same space, applying machine learning techniques, including feature extraction and neural network classifiers, to conformity checking in CMDB-like data for example, validating that hostnames and application names conform to expected naming conventions and reference data. Their results demonstrate that even relatively mechanical data quality problems benefit from a learned rather than purely rule-based approach.

Chinthapala [8] takes a broader, practitioner-oriented view, surveying how artificial intelligence is being integrated with ServiceNow and the CMDB across a range of organizational use cases and arguing that this integration represents a significant driver of operational efficiency and data-driven decision-making.

Across this strand of literature, the emphasis is technical and platform-specific: better algorithms for detecting duplicates, validating naming conventions, or reconciling relationships. This work is directly relevant to improving CMDB accuracy, but it is largely agnostic to what the resulting data will be used for downstream. None of these contributions examines whether the specific dimensions of CMDB quality they improve are the dimensions that matter most for continuity and recovery outcomes specifically, as opposed to general reporting or incident management use cases.

Two more recent contributions situate this technical work within the broader, fast-moving AIOps literature rather than CMDB-specific tooling narrowly. Zhang et al. [9] survey how large language models are being integrated across the AIOps lifecycle anomaly detection, root cause analysis, and automated remediation and Zota, Bărbulescu, and Constantinescu [10] propose a practical framework for developing agentic AIOps systems capable of autonomous operational decision-making. Relatedly, addressing a structurally similar problem to Darreddy and Munagala's [6] CMDB relationship reconciliation, L. Li [11] applies graph neural networks to service dependency modeling and failure-propagation prediction in distributed systems. None of these three treats CMDB-specific data governance, or business continuity and disaster recovery outcomes specifically, as a distinct design concern, but together they illustrate how rapidly the autonomy assumption underlying Section 3.5's trust-calibration discussion is becoming operational reality

rather than a future scenario.

3.4. Business Continuity, Disaster Recovery, and Cyber-Resilience

Business continuity and disaster recovery constitute a substantial literature in their own right, and a foundational treatment of the field is useful for establishing where CMDB data properly fits within it before examining how thinly that connection has actually been studied. Snedaker and Rima [12], in a widely used practitioner reference for IT business continuity and disaster recovery planning, formalize the now-standard distinction between business continuity planning (BCP), which addresses how an organization keeps critical business functions running during a disruption, and disaster recovery (DR), which addresses how technology and data are restored to meet agreed recovery time and recovery point objectives (RTO/RPO). Their treatment makes clear that dependency mapping identifying which systems, applications, and data support which business functions is a foundational input to both planning processes, since recovery sequencing and failover prioritization cannot be defined without first knowing what depends on what.

Alshammari, Alwan, and Al-Shaikhli [13] provide a systematic review of the research literature on disaster recovery and business continuity in cloud computing environments, examining studies published between 1995 and 2015 across several electronic databases. Their review characterizes disaster recovery primarily as a mechanism for reducing downtime following software or hardware failure and business continuity as an organizational requirement for maintaining operations, and surveys the technical approaches backup architectures, replication strategies, failover mechanisms that the literature has proposed to support these goals. Their review, like the broader cloud-DR literature it summarizes, is organized around infrastructure-level recovery mechanisms rather than around the configuration data governance that determines whether those mechanisms are pointed at the right systems in the first place.

More recently, Ibrahim [14] examined the impact of cloud computing specifically on business continuity and disaster recovery outcomes through a desk-research review of existing studies, reports, and statistics. Notably, that review's own stated findings explicitly identify "a contextual and methodological gap relating to cloud computing on business continuity and disaster recovery"

[14], while its substantive findings focus on how cloud adoption reduces recovery times and operational downtime relative to traditional infrastructure. This is a useful data point in its own right: even a 2024 review focused specifically on cloud computing and business continuity does not engage with configuration data governance as a variable, despite explicitly acknowledging that gaps remain in how this area has been studied.

Hernández-Bejarano, Rodríguez, and Merseguer [15] approach the field from a cyber-resilience rather than a pure continuity-planning angle, reviewing established standards and frameworks including the NIST framework, ENISA guidance, and ISO/IEC 27032, and arguing for a new cyber-resilience framework that leverages machine learning techniques to improve business continuity outcomes. Their work is representative of how business continuity and cyber-resilience research tends to be framed at the level of organizational standards, response processes, and increasingly, machine-learning-assisted detection and response, rather than at the level of the underlying configuration data that any recovery or resilience process ultimately depends on to know what needs to be protected, failed over, or restored.

A more recent conceptual review by Tzavara and Vassiliadis [16] traces the historical evolution of cyber resilience as a field, similarly situating the concept at the level of organizational capability, standards, and response processes rather than the underlying configuration or dependency data that any resilience process depends on to know what it is protecting reinforcing the pattern identified across this subsection.

Notably, the connection between dependency mapping and continuity planning that Snedaker and Rima [12] establish as foundational is echoed in current industry practitioner material rather than in the peer-reviewed literature. Virima [17], writing for a practitioner audience, argues directly that dependency mapping is "the prerequisite" for disaster recovery planning and that a credible recovery plan cannot be built without an accurate, current map of system dependencies underneath it a claim consistent with Snedaker and Rima's foundational framing, but offered without the empirical or peer-reviewed grounding that would allow it to be evaluated rigorously against actual recovery outcomes.

Across this literature, then, a consistent pattern emerges. Foundational and systematic academic treatments of business continuity and disaster recovery [12], [13] acknowledge that dependency data underlies recovery planning but do not treat the governance or quality of that data as a subject of sustained inquiry in its own right. A recent, narrowly scoped review [14] explicitly flags an unresolved gap in this space without identifying configuration data governance as one of its dimensions. Cyber-resilience research [15] engages with machine learning and standards frameworks but similarly treats underlying configuration data as a given rather than a governed asset. And the clearest articulation of the direct link between dependency mapping and disaster recovery readiness [17] currently exists only in practitioner literature, not in peer-reviewed research which is itself evidence of the gap this paper is identifying, rather than a source that closes it.

3.5. Trust, Reliance, and the Human-in-the-Loop Assumption in Automation

A longer-standing literature on trust in automation offers a useful, if indirect, lens on why the emergence of AI-driven and increasingly agentic IT operations changes the stakes of CMDB data quality. Parasuraman and Riley [18] characterize the central challenge of human-automation interaction as one of appropriately calibrated reliance: automation is used, misused, disused, or abused depending on whether an operator's trust in it matches its actual reliability, and mismatches in either direction produce predictable categories of failure. Lee and See [19] extend this into a more developed model of trust calibration, arguing that appropriate reliance depends on automation conveying enough about its own performance, process, and purpose for a human operator to form an accurate assessment of when to rely on it.

This literature was developed primarily around human operators relying on automated systems, not around autonomous agents relying on underlying data with no human trust-calibration step in between. It nonetheless establishes a principle directly relevant to CMDB governance for agentic IT operations: reliance on an imperfect information source is manageable when the relying party can calibrate trust appropriately, and becomes considerably more consequential when that calibration step is absent or unreliable. Applied to CMDB-driven automation, this suggests that the historical tolerance for CMDB data quality gaps has depended,

often implicitly, on a human trust-calibration step that autonomous or semi-autonomous IT operations do not automatically inherit.

4. Synthesis: The Gap Between Practice and Research

Reviewed together, the literature examined in Section 3 covers CMDB data quality measurement, general data governance frameworks, technical AI/ML approaches to CMDB accuracy, business continuity and cyber-resilience frameworks, and the general theory of trust calibration in automation. Each of these areas is reasonably well developed on its own terms. What is largely absent is work that connects them directly: a treatment of CMDB data governance that is organized specifically around business continuity and disaster recovery outcomes, rather than around general data accuracy or general resilience process design.

Three specific gaps follow from the review above.

First, CMDB data quality measurement [1], [2], [3] operates at the aggregate level and does not distinguish which specific governance dimensions matter most for continuity outcomes as opposed to other CMDB use cases such as reporting or incident triage. A governance program calibrated for general accuracy is not necessarily calibrated correctly for continuity risk, but the literature offers little guidance on how those priorities should differ.

Second, the technical AI/ML literature on CMDB quality [5], [6], [7], [8] is largely use-case-agnostic: it improves accuracy, relationship reconciliation, and naming conformity without examining whether the specific errors it corrects are the errors most consequential for business continuity and disaster recovery specifically. A completeness improvement in a low-criticality part of the CMDB does less for continuity readiness than an equivalent improvement concentrated in Tier-0 dependency relationships, but this distinction is not addressed in the reviewed technical literature.

Third, and most directly, the business continuity, disaster recovery, and cyber-resilience literature [12], [13], [14], [15] engages minimally with CMDB data governance as a contributing variable, generally treating the underlying configuration and dependency data as a given input to resilience processes rather than as a governance problem with its own dimensions, failure modes, and remediation strategies. The clearest

statement of the link between dependency mapping and recovery readiness currently comes from practitioner material [17] rather than peer-reviewed research, and even a review specifically scoped to cloud computing and business continuity as recently as 2024 explicitly flags an unresolved gap in this area without identifying data governance as one of its dimensions [14]. This is a significant omission given that recovery sequencing and failover planning are directly downstream of CMDB dependency data accuracy.

The trust-in-automation literature [18], [19], while not focused on CMDB or continuity specifically, supplies a useful general principle for why closing these gaps is increasingly urgent: as IT operations move toward greater autonomy, the informal human trust-calibration process that has historically compensated for CMDB data quality gaps becomes less available, which converts a longstanding but tolerated data quality problem into a more direct source of operational and continuity risk.

5. Discussion and Implications

Returning to the three research questions posed in Section 1: on RQ1, the reviewed literature does not itself rank CMDB data-quality dimensions by BCDR relevance, but the pattern across Sections 3.1 through 3.4 points most consistently toward relationship and dependency integrity, concentrated specifically in Tier-0 and Tier-1 infrastructure, as the dimension most consequential for BCDR outcomes specifically. On RQ2, the foundational continuity literature [12] establishes dependency mapping as a prerequisite for BCDR planning, but no reviewed source examines the governance conditions validation, ownership, freshness under which that dependency data can be trusted well enough to serve that role reliably; this is the central open question a future governance framework would need to answer. On RQ3, the trust-in-automation literature [18], [19] indicates that the human trust-calibration function that has historically compensated for CMDB data-quality gaps is not automatically inherited by AI-driven and increasingly agentic IT operations, which is why the same underlying data-quality problem becomes a more direct source of BCDR risk as automation increases.

The gap identified in Section 4 suggests that CMDB governance research would benefit from being organized around specific downstream use cases, rather than treated as a single, undifferentiated data quality objective. Business continuity and disaster recovery (BCDR)

represent one such use case with particularly direct and well-defined stakes: an incomplete or stale dependency relationship translates into an incomplete or incorrectly sequenced recovery plan, in a way that is more structurally direct than the consequences of the same data gap for general reporting or incident triage.

This suggests a concrete direction for future work: a governance framework that identifies which CMDB data quality dimensions matter most specifically for continuity and recovery outcomes, evaluated against those outcomes directly rather than against general accuracy metrics. Such a framework would need to draw on the general data governance vocabulary established by frameworks such as DAMA-DMBOK [4], the domain-specific technical capabilities demonstrated in the CMDB-focused machine learning literature [6], [7], the continuity and dependency-mapping foundations established in [12] and [17], and an explicit accounting of what changes about data reliability requirements once AI-driven or agentic processes, rather than only human operators, are consuming that data—a concern the trust-in-automation literature [18], [19] makes clear cannot simply be assumed away.

For practitioners, the immediate implication is that CMDB governance programs oriented primarily around aggregate accuracy metrics may be poorly calibrated for continuity risk specifically, and that continuity and disaster recovery planning teams may be operating with less visibility into their underlying data quality than their process-level planning maturity would suggest.

6. Limitations

This review has clear limitations. It draws on a targeted rather than exhaustive set of sources across four literatures, selected for direct relevance to the intersection this paper is examining rather than for comprehensive coverage of any one of those literatures individually. The characterization of a research gap in Section 4 is necessarily an interpretive synthesis rather than a systematic review conducted against a pre-registered search protocol, and a more exhaustive systematic review would strengthen confidence in the specific boundaries of the gap identified here. Finally, this review does not itself propose or validate a governance framework; it is intended to establish the case for one, which is left as the subject of subsequent work.

7. Conclusion

CMDB data quality is a long-recognized and reasonably well-measured problem, and a growing body of technical work is improving specific dimensions of CMDB accuracy through machine learning and AI methods. Separately, business continuity and cyber-resilience research continues to mature, increasingly incorporating machine-learning-assisted detection and response. This review has argued that the connection between these two literatures remains underdeveloped: existing work rarely examines CMDB data governance specifically through the lens of BCDR outcomes, and rarely accounts for how the emergence of AI-driven and increasingly agentic IT operations changes the stakes of that connection. Closing this gap requires a governance framework organized specifically around continuity and recovery outcomes, grounded in the general data governance and CMDB-specific technical literature reviewed here, and informed by the trust-calibration principles established in the automation literature. The author intends to pursue this framework as the direct next stage of this research.

References

1. ReadyWorks, "CMDB accuracy issues: Time for a new strategy," ReadyWorks Blog, 2025. [Online]. Available: <https://www.readyworks.com/blog/cmdb-accuracy-woes-its-time-for-a-new-approach>
2. ScienceLogic, "Achieve real-time CMDB accuracy," ScienceLogic Resources, 2023. [Online]. Available: <https://sciencelogic.com/product/resources/achieve-cmdb-accuracy-with-real-time-sync-of-monitored-environment>
3. Virima, "CMDB health score: What it is, how to measure it," Virima Blog, 2026. [Online]. Available: <https://virima.com/blog/cmdb-health-score>
4. DAMA International, DAMA-DMBOK: Data Management Body of Knowledge, 2nd ed. Basking Ridge, NJ: Technics Publications, 2017.
5. F. Dande, X. Li, M. A. Shofoluwe, and A. McLeod, "Artificial intelligence integration in IT service management: An ITIL configuration management process review," in Proc. Industrial Engineering and Operations Management (IEOM) World Congress, 2024.

6. J. K. R. Darreddy and M. K. Munagala, "AI-powered CMDB health and autonomous service graph reconciliation: Graph machine learning for configuration data quality in large-scale ServiceNow ITOM deployments," *International Journal of Computer Information Systems and Industrial Management Applications*, vol. 18, no. 5s, pp. 727–737, 2026.
7. S. Niewiadomski and G. Mzyk, "ML support for conformity checks in CMDB-like databases," in *Artificial Intelligence and Soft Computing (ICAISC 2023)*, Lecture Notes in Computer Science, vol. 14126, L. Rutkowski et al., Eds. Cham: Springer, 2023, pp. 366–376.
8. Y. Chinthapatla, "AI integration with ServiceNow and CMDB: Revolutionizing industries and society," *Journal of Science & Technology*, 2024. [Online]. Available: <https://www.researchgate.net/publication/379261632>
9. L. Zhang, T. Jia, M. Jia, Y. Wu, A. Liu, Y. Yang, Z. Wu, X. Hu, P. S. Yu, and Y. Li, "A survey of AIOps in the era of large language models," *ACM Computing Surveys*, 2026, early access, doi: 10.1145/3746635.
10. R. D. Zota, C. Bărbulescu, and R. Constantinescu, "A practical approach to defining a framework for developing an agentic AIOps system," *Electronics*, vol. 14, no. 9, art. 1775, 2025.
11. L. Li, "Service dependency modeling and failure propagation prediction in distributed systems based on graph neural networks," *Discover Artificial Intelligence*, vol. 6, no. 1, art. 512, 2026.
12. S. Snedaker and C. Rima, *Business Continuity and Disaster Recovery Planning for IT Professionals*, 2nd ed. Waltham, MA: Syngress, 2013.
13. M. M. Alshammari, A. A. Alwan, and I. F. Al-Shaikhli, "Data recovery and business continuity in cloud computing: A review of the research literature," *International Journal of Advancements in Computing Technology*, vol. 8, no. 5, pp. 30–36, 2016. https://www.researchgate.net/publication/312166397_Data_recovery_and_business_continuity_in_Cloud_computing_A_Review_of_the_Research_Literature
14. O. Ibrahim, "Impact of cloud computing on business continuity and disaster recovery," *Journal of Technology and Systems*, vol. 6, no. 5, pp. 16–28, 2024.
15. M. Hernández-Bejarano, R. J. Rodríguez, and J. Merseguer, "A vision for improving business continuity through cyber-resilience mechanisms and frameworks," in *Proc. 16th Iberian Conf. Information Systems and Technologies (CISTI)*, 2021, pp. 1–5.
16. V. Tzavara and S. Vassiliadis, "Tracing the evolution of cyber resilience: A historical and conceptual review," *International Journal of Information Security*, vol. 23, no. 3, pp. 1695–1719, 2024.
17. Virima, "Business continuity dependency mapping and IT resilience," *Virima Blog*, 2026. [Online]. Available: <https://virima.com/blog/boost-resilience-with-business-continuity-dependency-mapping>
18. R. Parasuraman and V. Riley, "Humans and automation: Use, misuse, disuse, abuse," *Human Factors*, vol. 39, no. 2, pp. 230–253, 1997.
19. J. D. Lee and K. A. See, "Trust in automation: Designing for appropriate reliance," *Human Factors*, vol. 46, no. 1, pp. 50–80, 2004.