American Academic Scientific Research Journal for Engineering, Technology, and Sciences
ISSN (Print) 2313-4410, ISSN (Online) 2313-4402
https://asrjetsjournal.org/index.php/American_Scientific_Journal/index

Methodological Approaches to Emergency Recovery of Information Systems in Multi-Contour Environments

Alexander Andreyev*

Solution Engineer Independent IT Contractor,Seattle, USA
Email:alpha@rri.kiev.ua

Abstract

The article examines a methodological approach to the emergency recovery of information systems in multi-contour environments. The increasing complexity of digital platforms underscores the relevance of this study, the need to segment infrastructures into autonomous technological contours, and the requirement to ensure high recovery speed while maintaining the logical and physical isolation of trust zones. Downtime losses and more burdensome regulations (NIS2, DORA, ISO 22301, IEC 62443) are new pressures on the art of Disaster Recovery. The paper aims to develop and justify methodological principles and algorithms for synchronizing disaster recovery plans across different contexts, while keeping them isolated in terms of target RTO/RPO metrics. Therefore, it includes a review of industry reports, regulatory documents, and practical guides that integrate Business Impact Analysis with threat-scenario modeling, following the NIST SP 800-30 methodology. It lives in two dimensions of novelty: the comprehensive integration of four architectural pillars as a unified framework for multi-contour environments and multi-level failover orchestration, together with contour-aware replication that enables seamless switching. Results are keyed to the facts that this method enables the attainment of the desired RTO and RPO with complete security, lowers the risk of cascading failure, and maintains all paths within the bounds of regulatory requirements. Inside each contour, there is synchronous replication, while between contours, there is asynchronous, encrypted replication, combined with the use of the 3-2-1-1+ backup approach. The involvement of DRaaS providers guarantees business-process continuity, preserving the logical isolation of segments. This article will find its readership among IT infrastructure managers and specialists, cybersecurity engineers, Disaster Recovery Solution architects, and information security auditors.

Keywords: multi-contour environments; emergency recovery; Disaster Recovery; Zero Trust; Infrastructure as Code; microservices; DMZ; RTO; RPO; backup 3-2-1-1+; DraaS.

Received: 7/16/2025
Accepted: 9/16/2025
Published: 9/26/2025

---

* *Corresponding author.*

1.Introduction

Digital platforms are increasingly built as assemblies of isolated technological contours, each servicing a specific class of data and level of trust. Such a multi-contour environment represents a logically and physically segmented infrastructure in which traffic exchange is conducted through strictly controlled gateways, and access, encryption, and monitoring policies are applied in a differentiated manner. Characteristic examples include the separation of industrial OT systems and corporate IT, the use of dedicated segments for personal data as mandated by GDPR requirements, and the designation of zones according to the ISA/IEC 62443 model for critical assets. Segmentation enhances resilience against attack propagation, yet complicates operational recovery, since the failure of one contour must neither compromise the integrity of others nor delay the restoration of services to the target RTO/RPO.

Under these conditions, disaster recovery evolves from a local task of bringing a server back online into a systemic discipline that requires synchronizing DR plans across domains without compromising their isolation. Design must account for several factors, including separate backup processes, independent encryption-key stores, limited replication channels, and the requirement for user-transparent switchovers. The significance of a properly structured DR is confirmed by incident economics: a survey of 3,215 enterprises across 13 industries revealed that a typical industrial operator loses approximately $ 125,000 for each hour of unplanned downtime, and a shift-long outage can cost around $1 million [1]. At the same time, only one quarter of organizations report that they have already implemented full asset visibility and network segmentation at the basic first maturity level. However, just one year ago, that figure was only 20% [2]. These data underscore a dual challenge: downtime costs are rising, while infrastructure preparedness for rapid and secure recovery remains limited.

To address this challenge, this paper aims to develop and justify a methodological framework for emergency recovery in multi-contour environments. The proposed approach integrates four key architectural pillars into a cohesive orchestration flow designed to synchronize recovery plans across isolated domains. The objective is to provide a systematic method for achieving target RTO and RPO metrics without compromising the logical and physical isolation that underpins the security of segmented infrastructures.

2.Materials and Methodology

The research is grounded in the analysis of 18 sources, including academic articles, industry reports, regulatory documents, and practical guides. The theoretical foundation comprised works that described segmentation principles and trust-zone interactions. Flå and his colleagues presented a model of zones and conduits by the IEC 62443 standard Reference [5], and Commvault detailed a cascading approach for calculating Recovery Time Objective (RTO) and Recovery Point Objective (RPO) across varying data-criticality levels [7]. Risk analysis and threat-scenario modeling principles were adopted from the NIST SP 800-30 guide [8]. At the same time, primary data on economic losses from downtime were drawn from the ABB report, which states an average unplanned downtime cost of $125,000 per hour [1].

For the comparative analysis of architectural solutions, both quantitative and qualitative data from industry

surveys and studies were utilized. Downtime costs and network-segmentation maturity levels were assessed by ABB [1] and Fortinet [2]. Progress in specialized OT-segment monitoring was reported by the SANS Institute in 2024 [3, 6]. The HashiCorp State of Cloud Strategy confirmed the efficacy of modular and microservices architectures survey [10] and the IBM Cost of a Data Breach report, which noted reductions in IT-team response times and cost savings in incident remediation through automation of recovery processes [11].

The systematic review of regulatory and standards requirements encompassed the EU NIS2 Directive and the DORA act [4], along with the provisions of ISO 22301 and IEC 62443 [5]. This review identified key obligations: documentation and auditing of backup plans, ensuring invisible switching between contours without violating logical isolation, and regular testing of recovery scenarios with OT-specific tools.

However, prior studies offer important building blocks but address them primarily in isolation. Research aligned with IEC 62443 formalizes zones and conduits, clarifying how conduit protection levels bound risk propagation across trust boundaries, but it does not prescribe end-to-end disaster-recovery orchestration across multiple contours with heterogeneous RTO/RPO targets [5]. Methodological guides for risk and impact—such as NIST SP 800-30 and industry BIA practices that define RTO/RPO—provide robust lenses for assessing threats and setting recovery objectives. Still, they remain agnostic to the operational mechanics of synchronous intra-contour versus asynchronous inter-contour replication under strict key-management separation [7, 8]. Surveys and posture reports further document visibility gaps and lateral movement risks in OT/ICS and segmented networks, underscoring the need for designs that prevent cross-contamination escalation during failover rather than treating recovery as a uniform enterprise process [3, 6, 9].

Complementary streams examine enabling technologies, including IaC-driven immutable infrastructure, microservice decoupling, Zero Trust with just-in-time privileges, multi-region patterns, and DRaaS adoption, which demonstrate faster incident response and lower remediation costs when automation and consistent configuration are present [10, 11, 15, 17, 18]. However, these works stop short of unifying segmentation theory, BIA-driven objectives, and failover mechanics into a contour-aware DR method. The present study positions itself at that junction by integrating four architectural pillars with a concrete orchestration flow that maintains logical isolation, employs synchronous replication only within contours, and encrypted asynchronous movement across contours. These choices are tied to auditable compliance, thereby extending prior art from component practices to a cohesive framework for multi-contour environments.To address the identified gap, the modeling and validation methodology comprised four complementary stages. First, an inventory of assets and zone boundaries was performed based on IEC 62443 recommendations [5] and CISA reports on inter-contour traffic vulnerability analysis [9]. Second, a cascading Business Impact Analysis was conducted using the Commvault methodology to assess maximum losses within each contour and model the effects of partial connectivity loss [7]. Third, threat-scenario analysis was conducted according to NIST SP 800-30 [8], accounting for lateral movement through one-way channels and contour-bound encryption keys. Finally, the reliability of recovery orchestration was tested via failure emulation and chaos experiments, drawing on AWS Prescriptive Guidance for multi-level failover scenarios [15]. The backup-practice evaluation applied the 3-2-1-1+ standard, using data from Arserve [16], while the DRaaS market-trend analysis referenced TechTarget [17] and Markets and Markets [18].

3.Results and Discussion

Multi-Contour architecture implies a strict separation of trust zones: an open public contour, a corporate IT contour of moderate criticality, and closed segments in which highly confidential data are processed or critical-infrastructure operations are performed. Isolation is achieved through a combination of network gateways, firewall policies, and dedicated key-management systems. When the load is evenly distributed, this approach reduces the likelihood of cascading failures. However, in the event of an incident, it introduces additional complexity, as each contour must recover autonomously without compromising the integrity of adjacent zones or violating the target RTO/RPO defined for the end-to-end business process.

The principal risk in segmented environments lies not so much in elementary hardware failures as in misconfigurations at contour junctions and insufficient visibility of assets beyond the gateway. The stricter the segmentation, the costlier the downtime, since redundant equipment must be duplicated not across the entire cluster but within each contour separately [3].

Regulatory pressure further intensifies the challenge: the EU NIS2 Directive and the DORA Act require organizations in critical sectors to demonstrate all-hazards resilience, including formalized backup plans, recovery without loss of isolation, and reporting on actual business-continuity tests—penalties for non-compliance being comparable to GDPR fines [4]. Simultaneously, the ISO 22301 and IEC 62443 standards require the documentation of RTO/RPO dependencies about trust levels and the dedicated storage of backups for each zone. Consequently, business requirements and regulations converge on a single point: the DR program must guarantee rapid service restoration within the specified contour, prevent unauthorized data transfers between segments, and be validated by regular, auditable tests.

A clear classification of contours is the first step in managing emergency risks. The international standard IEC 62443 defines a zone as a group of physical or logical assets sharing standard security requirements, and a conduit as a specialized channel through which such zones exchange data; the protection level of a conduit equals the highest security level of the zones it connects, thus establishing a natural boundary for incident propagation [5]. In practical infrastructures, this typically manifests as at least three trust levels: public services, corporate IT, and a closed OT contour—each of which may be further subdivided into sub-zones based on process criticality or data confidentiality.

Asset and boundary inventory begins with the creation of a unified registry of components, interconnections, and controlled gateways. In practice, omissions in the inventory remain the leading cause of blind spots: the SANS 2024 survey showed that only 52% of enterprises have implemented specialized OT monitoring, while only 34% regularly conduct realistic scenario drills—hence nearly half of organizations still fail to observe a significant portion of inter-contour traffic and do not verify recovery plans under segmentation conditions [6], as illustrated in Figure 1.

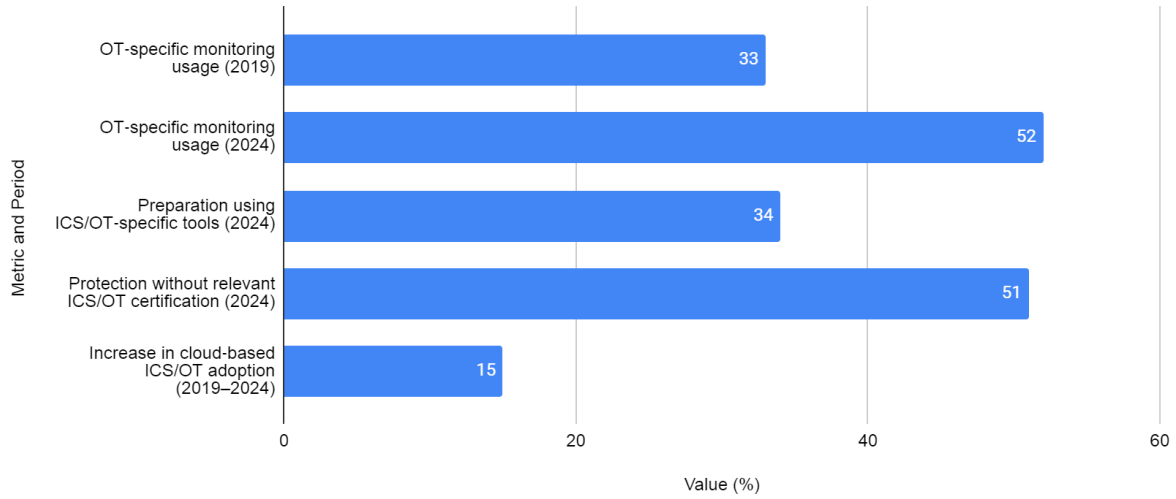

Figure1: Evolution of OT/ICS Security Posture [6]

Comprehensive asset registration must record not only network addresses and firmware versions but also each resource's zone affiliation, conduit type, encryption requirements, and permissible communication vectors.

The next layer of analysis—Business Impact Analysis—calculates the Recovery Time Objective and Recovery Point Objective for each zone. Industry data reveal a gap between theory and practice: even in highly regulated healthcare, requirements of RTO $\leq$ 2 hours with RPO $\leq$ 12 hours are not uncommon, whereas business-critical online platforms strive for minute-level targets [7]. Accordingly, zoned BIA is constructed in a cascading manner: first, the maximum loss within each zone is assessed; then the cascading effect under partial conduit unavailability is modeled; finally, differentiated RPO/RTO targets and capacity reserves are defined for each segment.

Threat modeling completes the risk-management cycle. The NIST SP 800-30 methodology recommends evaluating threats through the lens of three factors—source, vulnerability, and potential impact—and repeating this assessment throughout the system's life cycle [8]. In a multi-contour environment, this yields two classes of scenarios. Traditional scenarios include hardware failures, update errors, and natural disasters, in which a zone typically suffers in isolation. Inter-contour scenarios are more insidious: by exploiting conduit configurations, an attacker can jump from an external segment into the root OT zone, and an emergency failover that ignores zone policy will only expand the attack perimeter. Therefore, the risk model must account for specific lateral movement channels, the presence of unidirectional diodes, contour-bound encryption key assignments, and control of temporary privilege elevation during recovery. Only the integrated combination of inventory, BIA, and threat-scenario analysis provides a foundation for a methodologically rigorous and regulatorily compliant disaster-recovery plan in a multi-contour architecture.

Design of disaster recovery in a multi-contour environment relies on four complementary architectural principles, each addressing gaps identified in inventory, BIA, and threat modeling. The first layer consists of controlled gateways and DMZ segmentation, which limits inter-contour traffic, employs unidirectional channels or VPNs with strict filtering, and thus breaks potential lateral movement chains. A current analysis of 105 federal sites in

the CISA 2023 report showed that eliminating a single DMZ bypass reduced the share of successful internal attacker movements during escalation from 42% to below 10% [9].

Next comes modularity and service decoupling: switching a container or function is significantly simpler than restoring a monolith, so microservice granularity directly translates into shorter RTOs. According to the HashiCorp State of Cloud Strategy 2024 survey, organizations that implemented an automated and consistently configurable environment reported a 49% increase in security and a 45% reduction in IT-team incident-response time, owing to failures being isolated within a single service rather than paralyzing the entire contour [10], as shown in Fig. 2.

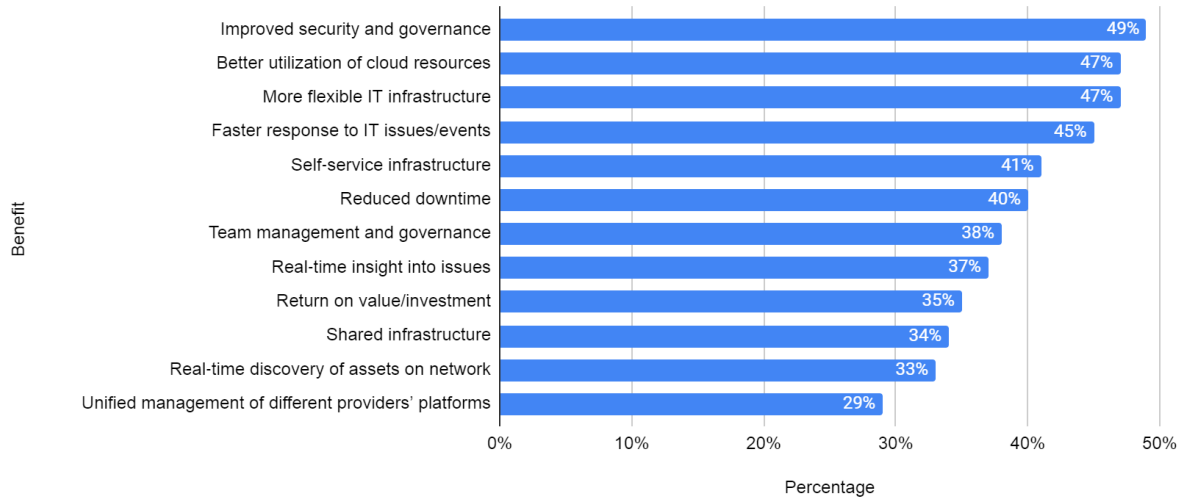

Figure2: Benefits of Cloud-Enabled IT Infrastructure [10]

The third layer is immutable infrastructure, deployed via IaC templates and orchestrated with the same pipelines used in daily releases. This approach eliminates manual incident remediation: the environment is recreated from its declarative description, and automation tools converge it to the target state. The economic impact is confirmed by the IBM Cost of a Data Breach 2024 report: organizations with extensive automation and AI usage spent, on average, 2.2 million USD less on incident remediation than companies relying on manual recovery, primarily due to acceleration of the identify → contain → restore phases [11].

The fourth principle integrates the Zero Trust model and Just-in-Time PAM. Emergency recovery mustn't open additional doors to attackers: privileges for DR operations are granted only after multi-factor verification and automatically expire. An IBM press release on its 2022 study found that critical-infrastructure organizations not adopting Zero Trust incurred an average incident cost of $ 5.4 million , $ 1.17 million higher than companies that had implemented the approach [12]. Together, these four architectural pillars form a framework within which automated DR orchestration can guarantee target RTO/RPO without compromising contour isolation or increasing operational risk.

Contour-aware replication relies on the principle that objects within the same trust level can exchange data

synchronously. In contrast, boundary crossings must occur asynchronously via encrypted, unidirectional channels and with separate root keys. The shift to cloud platforms has accelerated this differentiation: 74% of organizations plan to incorporate public or hybrid clouds into their data-protection schemes by 2025, while 85% of respondents have experienced at least one ransomware attack and now rate replication time as the key factor for service survivability [13]. In a multi-contour architecture, this leads to a two-speed practice: critical transactions within a contour are written synchronously with RPO ≈ 0, whereas off-contour or public-segment replication waits for a checkpoint and occurs with a delay sufficient for inspection and mandatory encryption of data on egress.

The cost of downtime remains the primary economic motivator for multi-level failover orchestration. A joint Splunk and Oxford Economics report showed that unplanned outages cost Global 2000 companies $ 400 billion annually—approximately 9% of their aggregate profits [14]. To reduce these losses, orchestrators implement a three-layer switchover model: local cluster failover groups preserve integrity within a segment; horizontal scenarios migrate services to an equivalent contour at another site; and vertical procedures elevate access only for the duration of recovery, relying on Just-in-Time PAM. The AWS Prescriptive Guidance document emphasizes that such a multi-level scheme must be complemented by regular testing and automatic rollback to the original region; otherwise, operational costs and complexity outweigh the benefits of redundancy [15].

Finally, contour-agnostic backup adheres to the 3-2-1-1+ rule: at least three copies on two different media types, one off-site, one immutable or air-gapped, plus mandatory automatic verification. Despite the method's clarity, 23% of companies still have not implemented the extended 3-2-1-1 variant and therefore risk losing their last line of defense in a complex attack [16]. This gap is bridged by combining immutable copies in a separate trust zone, periodic offline tapes, and DRaaS replicas, whose share grew from 39% to 53% between 2016 and 2019 and continues to rise, spurring companies to more rigorous recovery plan testing [17]. As a result, the symbiosis of contour-aware replication, layered failover, and contour-agnostic backups forms a continuous protective perimeter, in which each failure is localized, and business processes are restored without risking isolation breaches.

Container packaging of services shifts recovery from manual repair to automated reinstantiate-from-template: upon failure, the orchestrator deploys the known image, attaches volumes, and returns the application to its target state without affecting other segments. Periodic chaos experiments reinforce this effect by uncovering hidden dependencies and verifying that backup schemes indeed function under operational conditions. Together, containerization and chaos testing create a self-regulating environment in which a single node failure is rapidly contained and does not evolve into a cascading outage.

A further level of resilience is achieved by off-site replication to a disaster recovery site. Regulations often constrain direct output of critical data to a public cloud, so the architecture is constructed in a hybrid manner: a private zone holds hot copies. At the same time, a managed DR-service provider maintains a warm or cold tier. The market for such services already reflects demand: the global DRaaS market is projected to reach $ 26.5 billion by 2028 at a CAGR of 19.8%, as shown in Fig. 3 [18].

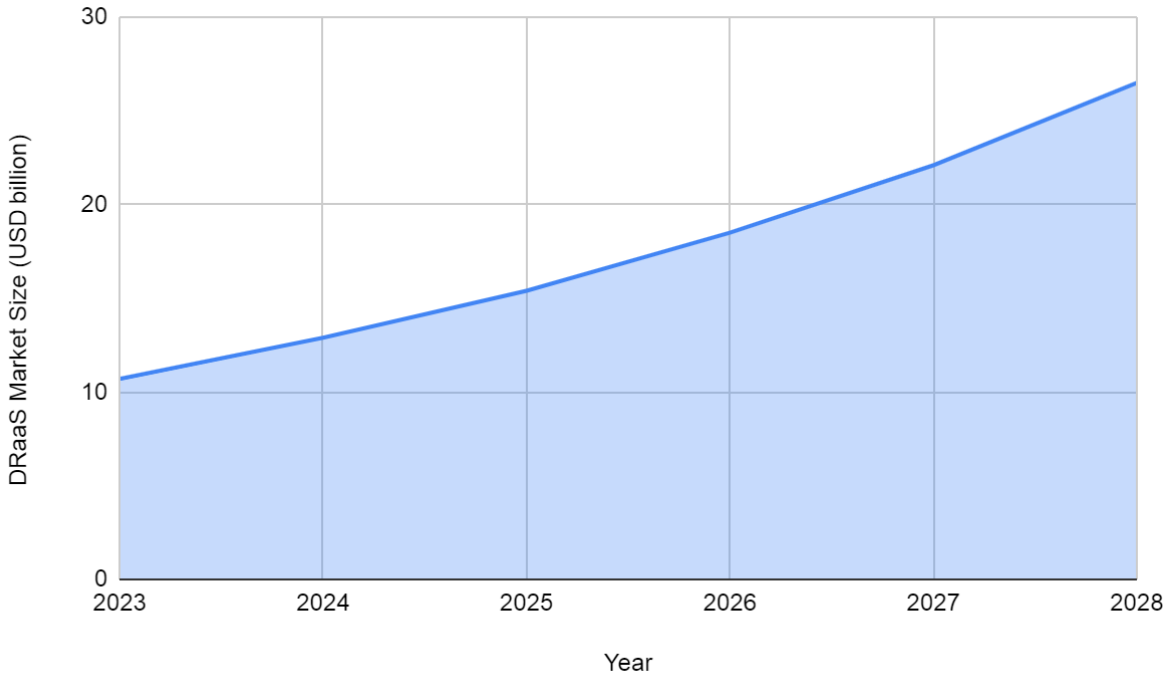

Figure3: Disaster Recovery as a Service Market Size [18]

Such segregation enables compliance with isolation requirements while still allowing for flexible resource scaling and failover to an external site in the event of a large-scale outage.

To ensure that a suite of isolated contours remains compliant with regulatory norms, recovery processes incorporate continuous compliance through declarative policies. Checks embedded directly in build and deployment pipelines automatically detect violations, generating an audit trail without human intervention and reducing approval delays. Policies are described as code alongside infrastructure templates, ensuring a single source of truth and eliminating discrepancies between formal documentation and the actual state of the environment.

A critical pillar of the methodology remains team preparedness. Developers and operators undergo regular training in secure coding and incident response, and procedural instructions are maintained in machine-readable formats. Such runbook files are versioned, automatically linked to specific pipelines, and can be invoked by the orchestrator upon detecting an event. Thus, individuals, procedures, and instruments create a continuous cycle: issues are swiftly managed, rules are promptly checked for adherence, and an organized group is restored while maintaining both quickness and safety.

Therefore, setting up a comprehensive DR plan in a multi-contour environment requires a combination of clear area marking and controlled entry points, a modular structure, and a stable setup, along with Zero Trust methods utilizing Just-in-Time PAM and contour-aware copying. This comprehensive scheme not only minimizes the risks of cascading failures and isolation breaches during incidents but also ensures attainment of target RTO/RPO without compromises between business continuity and cybersecurity.

The framework's effectiveness rests on four interlocking mechanisms that reinforce each other to achieve resilient recovery without weakening isolation boundaries. Controlled gateways and DMZ segmentation limit lateral movement by ensuring that inter-contour traffic only flows through conduits with equal or stronger protections. Service decoupling and containerization further localize faults, enabling selective recovery of only the impacted functions rather than restoring entire monoliths. Immutable infrastructure expressed through Infrastructure as Code ensures consistency and eliminates drift by replacing ad-hoc fixes with deterministic redeployment. Finally, Zero Trust principles, combined with Just-in-Time privileged access, keep recovery operations tightly bound, temporary, and fully auditable, thereby preventing privilege escalation. Together, these mechanisms explain how the framework meets target recovery objectives while maintaining strict contour isolation.

Operationally, recovery follows a structured, contour-aware orchestration flow. After detecting and classifying the event, the system stabilizes writes and takes a consistent checkpoint before re-instantiating services from golden images. Healthy pipes allow safe outside copying. Bad pipes or stopped leaving rules start a careful switch to a ready disaster help level. During this, all top actions are checked by multi-step, short-time passes, and setups are always checked with rule-as-code fences. Auto return ensures safety when breaks occur and ends steps by matching logs, renewing danger plans, and improving work harm checks. This closely planned method combines speed, holding in, and checking for various types of problem cases.

The provided examples demonstrate how it tunes itself in various running environments. External-facing applications can be designed to achieve minute-level recoveries through stateless frontend horizontal failover and asynchronous replication of session state. Corporate IT implements less than thirty-minute recovery, no data loss in all normal conditions based on containerized applications with replicated databases. Near-zero RPO, quick local recovery is made possible in high-criticality OT scenarios by using synchronous writing combined with a hot-standby pair, while deliberately delaying external replication to maintain isolation. Such differentiation targets occur across the broader industry landscape, where online services strive for near-real-time recoveries, whereas regulated domains adhere to much more conservative bounds. The proposed framework separates synchronous intra-contour recovery from inter-contour replication, which appears to occur asynchronously, and enforces this balance transparently.

4.Conclusion

In the context of increasing complexity and criticality of digital platforms implemented according to multi-contour segmentation principles, the formulation and implementation of methodologically grounded approaches to emergency recovery become an integral component of protecting business processes and ensuring cybersecurity. Recovery-related reliability and speed depend on the clarity of separation between trust zones, the completeness of asset inventory, and the accuracy of each level of business-impact calculation. It is precisely a systematic combination of these elements that minimizes the risk of blind spots and allows for differentiation and determination of RTO/RPO per contour area, with no unintended interaction between the segments.

The key outcome of the study is the identification of four complementary architectural principles that address gaps in inventory, business-impact analysis, and threat-scenario modeling. Controlled gateways and DMZ

segmentation limit inter-contour traffic and reduce the likelihood of lateral attacker movement; a modular microservice architecture shortens switch-over times and localizes failures; immutable infrastructure deployed via Infrastructure as Code eliminates manual operations and maintains environment consistency; and integration of Zero Trust models with Just-in-Time PAM ensures that privileges during DR operations remain strictly limited and controlled.

Particular attention is paid to contour-aware replication and multi-level orchestration of failover scenarios, where synchronous replication within a contour is combined with asynchronous, encrypted channels to subordinate zones. Adoption of the 3-2-1-1+ backup approach and engagement of external DRaaS providers provide additional guarantees of physical and logical isolation of copies. Meanwhile, containerization and regular chaos testing validate the effectiveness of recovery schemes under real-world conditions.

It is important, however, to acknowledge the limitations of this study. The research relies primarily on the analysis of regulatory documents and industry reports, which may lack a representative base of practical cases with quantitative verification of achieved RTO/RPO or controlled comparative experiments. Consequently, the conclusions are methodological rather than empirically confirmed. The generalizability of the approach is limited by the assumptions of a mature level of segmentation, automation (IaC/CI/CD), and the availability of DRaaS. In environments with legacy monoliths, strict prohibitions on data removal, or a different trust contour model, the stated benefits and target metrics may be unachievable without significant modifications and additional costs.

Finally, the complete disaster-recovery cycle includes automated compliance control via declarative policies embedded in CI/CD pipelines, as well as continuous team preparedness through regular exercises and updates to machine-readable runbook files. This integrated methodology enables organizations not only to restore services rapidly within defined RTO/RPO constraints but also to preserve contour isolation intact—an essential requirement of modern regulatory standards and economic resilience.

References

[1] ABB survey reveals unplanned downtime costs $125,000 per hour, ABB, Oct. 11, 2023. https://new.abb.com/news/detail/107660/abb-survey-reveals-unplanned-downtime-costs-125000-per-hour (accessed Jun. 22, 2025).

[2] OT Cybersecurity Risk Elevates within Executive Leadership Ranks, Fortinet, 2025. https://www.fortinet.com/corporate/about-us/newsroom/press-releases/2025/fortinet-report-ot-cybersecurity-risk-elevates-within-executive-leadership-ranks (accessed Jun. 23, 2025).

[3] A. Ribeiro, SANS Institute 2024 survey reveals progress and gaps in ICS/OT cybersecurity for critical infrastructure, Industrial Cyber, Oct. 08, 2024. https://industrialcyber.co/threats-attacks/sans-institute-2024-survey-reveals-progress-and-gaps-in-ics-ot-cybersecurity-for-critical-infrastructure/ (accessed Jun. 24, 2025).

[4] E. Hancock, Compliance is evolving — Is your resilience ready?, TechRadar, 2025.

https://www.techradar.com/pro/compliance-is-evolving-is-your-resilience-ready (accessed Jun. 26, 2025).

[5] L. Flå, M. Lundteigen, F. Gratte, and M. Gilje, Implementation of Zones and Conduits in Industrial Control and Automation Systems, Jaatun, 2024. Accessed: Jun. 27, 2025. [Online]. Available: https://jaatun.no/papers/2024/Zones-and-Conduits.pdf

[6] SANS 2024 ICS/OT Cybersecurity Survey Reveals Significant Progress While Highlighting That Major Gaps Remain in Securing Critical Infrastructure, SANS Institute, 2024. https://www.sans.org/press/announcements/2024-ics-ot-cybersecurity-survey-reveals-significant-progress-while-highlighting-that-major-gaps-remain-in-securing-critical-infrastructure (accessed Jun. 28, 2025).

[7] RTO (Recovery Time Objective) and RPO (Recovery Point Objective), Commvault, Jun. 05, 2024. https://www.commvault.com/glossary-library/rto-rpo (accessed Jun. 30, 2025).

[8] NIST, Guide for Conducting Risk Assessments, NIST, Sep. 2012. Accessed: Jul. 01, 2025. [Online]. Available: https://nvlpubs.nist.gov/nistpubs/legacy/sp/nistspecialpublication800-30r1.pdf

[9] CISA Analysis: Fiscal Year 2023, Risk and Vulnerability Assessments, CISA. https://www.cisa.gov/sites/default/files/2024-09/FY23_RVA_Analysis_508.pdf (accessed Jul. 02, 2025).

[10] HashiCorp, HashiCorp State of Cloud Strategy Survey, HashiCorp, 2024. https://www.hashicorp.com/en/state-of-the-cloud (accessed Jul. 03, 2025).

[11] IBM, Cost of a Data Breach Report, IBM, 2024. Accessed: Jul. 05, 2025. [Online]. Available: https://wp.table.media/wp-content/uploads/2024/07/30132828/Cost-of-a-Data-Breach-Report-2024.pdf

[12] IBM, IBM Report: Consumers Pay the Price as Data Breach Costs Reach All-Time High, IBM, Jul. 27, 2022. https://newsroom.ibm.com/2022-07-27-IBM-Report-Consumers-Pay-the-Price-as-Data-Breach-Costs-Reach-All-Time-High (accessed Jul. 07, 2025).

[13] Veeam, 2023 Data Protection Trends, Veeam, 2023. https://www.veeam.com/infographics/data-protection-trends-report-infographic-global-2023_wp.pdf (accessed Jul. 09, 2025).

[14] Splunk, Splunk Report Shows Downtime Costs Global 2000 Companies $400B Annually, Splunk, 2024. https://www.splunk.com/en_us/newsroom/press-releases/2024/conf24-splunk-report-shows-downtime-costs-global-2000-companies-400-billion-annually.html (accessed Jul. 10, 2025).

[15] AWS multi-Region fundamentals AWS Prescriptive Guidance, Amazon, 2025. https://docs.aws.amazon.com/pdfs/prescriptive-guidance/latest/aws-multi-region-fundamentals/aws-

multi-region-fundamentals.pdf (accessed Jul. 12, 2025).

[16] State of Data Resilience in the Enterprise, Arserve, 2024. Accessed: Jul. 14, 2025. [Online]. Available: https://goto.arcserve.com/rs/431-WBH-895/images/Executive-Report-State-of-Data-Resilience-Enterprise.pdf?version=0

[17] DRaaS guide: Benefits, challenges, providers and market trends, Tech Target, 2020. https://www.techtarget.com/searchdisasterrecovery/DRaaS-guide-Benefits-challenges-providers-and-market-trends (accessed Jul. 16, 2025).

[18] Disaster Recovery as a Service (DRaaS) Market, Markets and Markets, 2023. https://www.marketsandmarkets.com/report-search-page.asp?rpt=disaster-recovery-as-a-service-draas-market (accessed Jul. 19, 2025).