The American Journal of
Engineering and Technology

ISSN 2689-0984 | Open Access

TYPE Original Research
PAGE NO. 15-22
DOI 10.37547/tajet/Volume07Issue09-03

Check for updates

Practices For Planning and Operating Multisite Network Environments Under Crisis Conditions

Alexander Andreyev

Solution Engineer Independent IT Contractor, Seattle, USA

OPEN ACCESS

SUBMITED 08 August 2025
ACCEPTED 19 August 2025
PUBLISHED 09 September 2025
VOLUME Vol.07 Issue 08 2025

CITATION

Alexander Andreyev. (2025). Practices For Planning and Operating Multisite Network Environments Under Crisis Conditions. The American Journal of Engineering and Technology, 7(09), 15–22. https://doi.org/10.37547/tajet/Volume07Issue09-03

Abstract: This paper provides a systematic approach to planning and operating multisite network environments in the face of economic, natural, and cyber crises. The goal of the study is to identify methodological approaches and technological solutions that ensure maintaining business processes via a geo-distributed network and synthesize them. This includes RTO and RPO requirements and topologies (active–active, active–passive, geo-distributed clusters) definition as well as a hybrid-cloud scenario economic feasibility assessment. The relevance of this work is grounded in the increasing sensitivity of the digital economy to downtime, rising cyber risks, and socio-economic shocks (such as migration waves and climate catastrophes), which demand strategic responses to multidimensional threats based on distributed infrastructure and international standards. The novelty of the research lies in a comprehensive content analysis of 18 sources—from reports by Uptime Intelligence, IBM, and UNHCR to the ISO 22301, DORA, NIST SP 800-34, and ISO/IEC 27001 regulations. A cascading approach to calculating RTO/RPO involves criticality classification of services, dependency modeling, regular drill-over tests, and the development of financial justifications for CAPEX and OPEX flows under TCO and risk bonus. Key findings show that with multisite architecture plus a hybrid cloud, resilience to failures is ensured by automatic failover, regular DRP/BCP testing in operation monitoring MTTR/MTBF continuous error-budget management; international standards integration resulting in exhaustive checklists for Business Impact Analysis and verification of target metrics; SOAR automation application plus predictive maintenance which mitigates staff shortage problem speeding up recovery;

network segmentation and multilayer encryption inside Zero-Trust framework. This article will be of use to IT infrastructure managers, network system architects, and specialists in business continuity and information security.

Keywords: Multisite network environments; business continuity; crisis conditions; Business Impact Analysis; RTO; RPO; active–active; active–passive; hybrid cloud; SOAR automation; DRP; BCP; Zero Trust.

Introduction: The modern digital economy has become so sensitive to downtime that even brief service interruptions result in significant financial losses: more than half of operations professionals acknowledged that their last major outage cost their company over $100,000, and in 16% of cases the bill exceeded one million—figures reflecting lost transactions, SLA penalties and erosion of customer trust (Uptime Intelligence, 2024). Against this backdrop, multisite topology is no longer viewed as unnecessary complexity: geographically distributed sites enable load distribution, incident containment and the maintenance of business activity even amid partial infrastructure failures.

A second powerful driver toward a distributed architecture is the rise in cyber risks. The average cost of a single data breach in 2024 reached $4.88 million, setting a new historical high; the bulk of these expenses stem from downtime and operational recovery, rather than the technical remediation work itself (IBM, 2024). The more complex and costly the attack, the more critical it becomes to swiftly redirect critical workloads to backup sites, thereby reducing SOC response time and curbing the scale of damage.

For this article, crisis conditions encompass economic shocks, natural disasters, social migration waves, and cyber threats, each imposing distinct demands on network resilience. Social instability is also intensifying: according to the UNHCR, by the end of 2024, the number of forcibly displaced persons had surpassed 123 million, leading to local overloads of telecom and power infrastructures in host countries (UNHCR, 2024).

Thus, multisite networks emerge not merely as a technical option but as a strategic response to multidimensional threats. The following analysis will demonstrate how to properly design such an architecture—budgeting, allocating resources and establishing processes—to preserve business continuity and user trust under any of the aforementioned crisis scenarios.

Methodology

The study of practices for planning and operating multisite network environments under crisis conditions is based on the analysis of 18 sources, including industry reports (Uptime Intelligence, 2024; IBM, 2024), publications by international organizations (UNHCR, 2024; NCEI, 2025), corporate case studies (Alicke & Foster, 2024; Antithesis, 2024), as well as regulatory documents and standards (ISO 22301 (ISO, 2019); DORA (EIOPA, 2025); NIST SP 800-34 (NIST); ISO/IEC 27001). This broad coverage enabled the formation of a holistic view of the technical, operational, and normative requirements for ensuring the continuity of multisite networks under natural, economic, and cyber crises.

The theoretical foundation of the research comprised works on assessing the impact of downtime and cyber-threats on business: the Uptime Intelligence report shows that more than half of severe outages cost companies over $100,000, and 16 % exceed one million (Uptime Intelligence, 2024), while the average cost of a data breach reached a record $4.88 million, the bulk of which is attributable to downtime losses (IBM, 2024). UNHCR data indicate the social burdens on infrastructure in host countries (UNHCR, 2024), and NCEI information underscores the rise in climate catastrophes causing losses exceeding $1 billion (NCEI, 2025). The economic justification for hybrid-cloud solutions relies on estimates of a 30–40 % reduction in TCO when migrating workloads to the public cloud (Environment Analyst Global, 2020; Squalio, 2025), and supply-chain research confirms the necessity of diversification under instability (Alicke & Foster, 2024).

Methodologically, the research combined several key approaches: comparative analysis of active–active versus active–passive topologies concerning RTO and RPO requirements and redundancy cost calculations (Antithesis, 2024; Amazon, 2025); a systematic review of international standards and regulations—ISO 22301 with a documented risk-assessment process (ISO, 2019), DORA with mandatory ICT operational-resilience testing (EIOPA, 2025), NIST SP 800-34 on the lifecycle of contingency planning (NIST), ISO/IEC 27001 on information security; quantitative analysis of failover practices and

WAN-channel replication stress tests based on data on the frequency of drill-over checks and the actual attainment of target metrics (Security Magazine, 2021; Uptime Intelligence, 2024); and content analysis of SOAR-playbook and MFA-solution implementations with an evaluation of their impact on MTTR and MTBF (Grand View Research, 2024; Okta, 2024).

Results And Discussion

Preliminary risk analysis begins with mapping the spectrum of crises that realistically affect IT infrastructure. Natural–climatic catastrophes remain the most predictable in frequency yet the most unpredictable in geographic distribution: in 2024 alone, 27 individual events with losses exceeding $1 billion were recorded in the USA, as shown in Figure 1, and this figure became the second highest on record, vividly demonstrating how a single season of extreme weather can overload regional data centres and data-transport channels (NCEI, 2025).

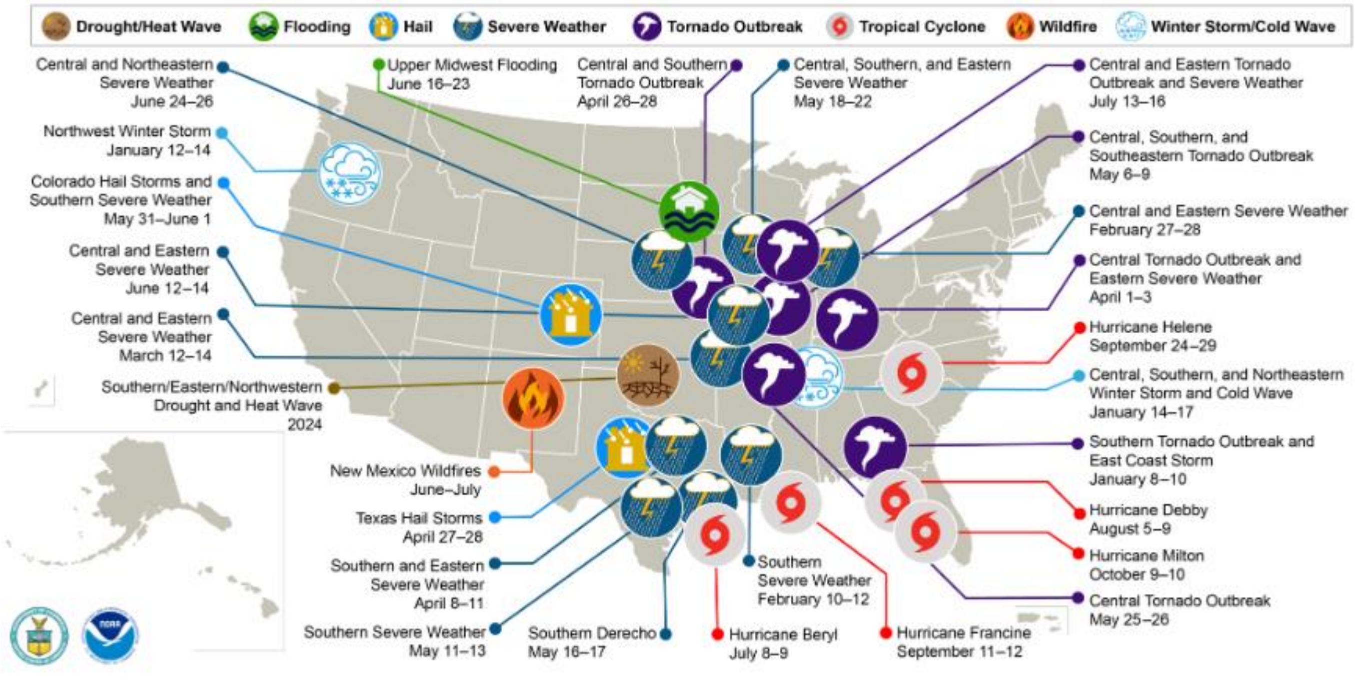
Drought/Heat Wave
Flooding
Hail
Severe Weather
Tornado Outbreak
Tropical Cyclone
Wildfire
Winter Storm/Cold Wave
Central and Northeastern Severe Weather June 24–26
Upper Midwest Flooding June 16–23
Central and Southern Tornado Outbreak April 26–28
Central, Southern, and Eastern Severe Weather May 18–22
Central and Eastern Tornado Outbreak and Severe Weather July 13–16
Northwest Winter Storm January 12–14
Central, Southern, and Southeastern Tornado Outbreak May 6–9
Colorado Hail Storms and Southern Severe Weather May 31–June 1
Central and Eastern Severe Weather February 27–28
Central and Eastern Severe Weather June 12–14
Central Tornado Outbreak and Eastern Severe Weather April 1–3
Central and Eastern Severe Weather March 12–14
Hurricane Helene September 24–29
Southern/Eastern/Northwestern Drought and Heat Wave 2024
Central, Southern, and Northeastern Winter Storm and Cold Wave January 14–17
New Mexico Wildfires June–July
Southern Tornado Outbreak and East Coast Storm January 8–10
Texas Hail Storms April 27–28
Hurricane Debby August 5–9
Southern and Eastern Severe Weather April 8–11
Hurricane Milton October 9–10
Southern Severe Weather February 10–12
Central Tornado Outbreak May 25–26
Southern Severe Weather May 11–13
Southern Derecho May 16–17
Hurricane Beryl July 8–9
Hurricane Francine September 11–12

Fig. 1. U.S. 2024 Billion-Dollar Weather and Climate Disasters (NCEI, 2025)

An unprecedented scale of forced displacements introduces a social risk vector, generating localized spikes in telecommunications and energy demand in host countries, and requiring operators to redistribute traffic among sites rapidly.

Finally, economic disruptions remain a chronic backdrop: 90% of surveyed global supply-chain directors reported at least one significant disruption in 2024, and 73% have already implemented dual-sourcing strategies, confirming that geographic and supply diversification are also key elements of resilience architecture (Alicke & Foster, 2024).

Each of the aforementioned classes of crises affects the IT landscape differently, but collectively they create four critical impact zones: energy availability, network connectivity stability, data integrity, and regulatory compliance. It is for this reason that the methodological foundation of subsequent planning is a structured Business Impact Analysis (BIA). The BIA maps failure scenarios to business processes, assigns recovery priorities, and determines the upper limit of acceptable downtime.

Based on the outputs of the BIA, numeric target metrics for RTO and RPO are established: the former constrains the duration of functional unavailability, and the latter the volume of data that can be lost between backups. In a multisite environment, these parameters are expressed not only in time but also in the required replication bandwidth, which dictates narrow synchronization windows and drives the choice between active–active and active–passive topologies.

To calculate realistic RTO/RPO values, organizations employ a cascaded approach: classifying services by criticality level, modeling dependency chains, and stress-testing backup channels. A typical practice involves assigning three tiers: critical processes (minutes to first revenue loss), important processes (hours), and supporting processes (24 hours or more).

Periodic failover tests complement comprehensive documentation, since actual attainability often diverges from design calculations due to WAN delays and human factors. In a multisite architecture, these tests should be conducted sequentially for each site, covering both remote synchronous and asynchronous replication, to confirm that the recovery point indeed corresponds to the calculated RPO.

The regulatory environment complements technical requirements with normative ones: ISO 22301 provides the framework for a business continuity management system and mandates a documented process for risk assessment, strategy selection, and recovery metrics (ISO, 2019). As of 17 January 2025, the DORA regulation in the EU financial sector has entered into force, requiring mandatory operational-resilience tests for ICT services, provider monitoring, and stringent incident-reporting obligations (EIOPA, 2025).

For global enterprises operating under U.S. jurisdiction, NIST SP 800-34 remains the reference point, describing the lifecycle of contingency planning and the minimum set of response plans adaptable to both physical and cyber threats. Complementing this is ISO/IEC 27001, which focuses on information security, enabling the alignment of continuity processes with data-protection controls. As a result, the requirements of standards and regulations become not external constraints but sources of detailed checklists that help to objectify chosen BIA, RTO, and RPO metrics and ensure the comparability of practices across sites.

Strategic planning of a distributed network begins by translating abstract BIA metrics into concrete engineering solutions. Suppose the RTO for critical transactions is measured in minutes and the RPO approaches zero. In that case, redundancy at the application and data levels must be embedded into the topology long before hardware procurement. Thus, the architecture becomes the exoskeleton of previously conducted business assessments, rather than a set of isolated technologies selected just in case.

When choosing between active–passive and active–active schemes, the decisive factor is less the technical differences and more the cost of downtime. Gartner simultaneously estimates average business losses at $5,600 per minute of downtime—meaning that savings from avoiding complete duplication are quickly eroded within the first hour of a major outage (Antithesis, 2024).

From a resource-allocation perspective, active–active clusters require at least two fully provisioned sites connected by low-latency channels and synchronous database replication. To prevent asymmetric loading, traffic is distributed via global load balancers, and application state is stored in shared in-memory caches or distributed file systems. In an active–passive model, infrastructure costs are lower, but one must accept a failover delay and a wider RPO window. Nevertheless, regular drill-over tests remain as essential as in active–active deployments; otherwise, the passive segment degrades only on paper.

Geo-distributed clusters logically extend the active–active approach by adding a third or fourth fault zone, allowing maneuvering among different legal jurisdictions, energy prices, and climatic risks. Consistency algorithms become particularly critical here: data controllers and message queues must choose between strict synchronization (with increased latency) and eventual consistency, the cost of which is greater complexity in rollback logic.

Practice shows that a cloud hybrid almost always complements corporate multisite strategies: 70% of companies already host workloads in at least one public cloud and one private cloud, forming a natural framework for traffic distribution among sites and providers (Squalio, 2025). This approach reduces vendor lock-in, provides out-of-the-box points of presence, and enables scaling of compute resources during peak periods without requiring capital investment in proprietary data centers.

The economics confirm the rationale of hybrid scenarios: an Accenture study reports a 30–40% reduction in total cost of ownership when migrating workloads to the public cloud, with savings derived from improved server utilization, flexible licensing, and reduced operational costs for power and cooling (Environment Analyst Global, 2020). The benefit becomes even more pronounced when cloud-neutral components—from container orchestrators to infrastructure-as-code—are factored in from the design stage.

Therefore, budgeting for a multisite architecture is built around three interconnected streams: CAPEX for minimally sufficient redundancy, OPEX for continuous replication and testing, and a risk bonus represented by losses averted from downtime. Aligning these streams with calculated TCO effects and costly-outage statistics makes it possible to justify even expensive active–active schemes at the board level and embed them in the company's long-term financial plan.

The transition from architectural blueprint to daily operation requires converting calculated RTO/RPO into actual platform redundancy: only then do engineering solutions become not insurance on paper but mechanisms capable of withstanding a real crisis. The first line of defense is physical and logical channel duplication. Even perfect physics is useless without instantaneous logical switching. The shorter the calculated RTO window, the less room there is for manual procedures: automatic failover must be entirely operator-less. The Uptime report notes that four out of five operators consider their recent outages preventable with better configuration management, effectively turning self-healing mechanics and continuous configuration-drift detection into savings measurable in millions (Uptime Intelligence, 2024). In practice, this means active–active distribution of the leader role within the cluster, continuous L7-level probing, and in-memory state replication for services handling user sessions.

However, automation without a plan is merely an accelerator of chaos; therefore, formal DRP and BCP remain a mandatory framework. Research shows that 54% of companies still lack a documented recovery plan; half of those conducting tests do so less than once a year, casting doubt on any theoretical reliability metrics (Security Magazine, 2021). In multisite environments, best practice dictates quarterly testing of failover scenarios for each node, simulating loss of network connectivity, power, and database; the results are entered into the risk registry and used to adjust the target RPO.

The digital measurement of success is articulated through a set of common availability metrics. From these, the target SLA threshold is derived: for critical workloads, the internal SLO target must not fall below 99.95 %, while for external customers, public clouds set the benchmark; for example, AWS guarantees 99.99 % when deploying instances across two AZs within a single region, attaching proportional penalties for deviations (Amazon, 2025). Continuous export of MTTR, MTBF, and Error Budget metrics to the NOC/SOC dashboard transforms the agreement into a living part of operational economics: deviations from budgets immediately block new releases, while compliance opens up a deployment allowance for new functionality, thus closing the quality loop.

Thus, continuity is formed through backup chains, automated failover, documented DRP/BCP, and measurable SLAs, which mutually reinforce each other. Each facet addresses the vulnerability identified by quantitative risk analysis, and together they render the multisite architecture resilient to the very crises that motivated this study. Resilience of a distributed architecture is impossible without a clear security model; hence, a logical extension of redundancy is strict network segmentation and a Zero Trust approach: nearly two-thirds of organizations have already adopted this access principle in response to the fact that 88 % of web application breaches occur using stolen credentials (Verizon, 2025).

The encryption layer serves as the next layer of protection. Regulatory acts ISO 27001 and PCI DSS 4.0 explicitly require protection at rest and in transit, and the financial motivation is evident: it is the leakage of unprotected data that accounts for the largest share of costs. Counteracting credential theft is reinforced by multi-factor authentication and precise access control. In corporate environments, MFA is already used by 66% of regular employees and 91% of administrators, which reduces the window for phishing attacks and renders password compromise insufficient for infiltrating critical segments (Okta, 2024). Nevertheless, passwords still dominate, with 95% of users, and among MFA methods, push notifications are the most widespread (29%), whereas more modern and robust solutions, such as WebAuthn (3%) or smart cards (0%), are hardly used, as shown in Figure 2.

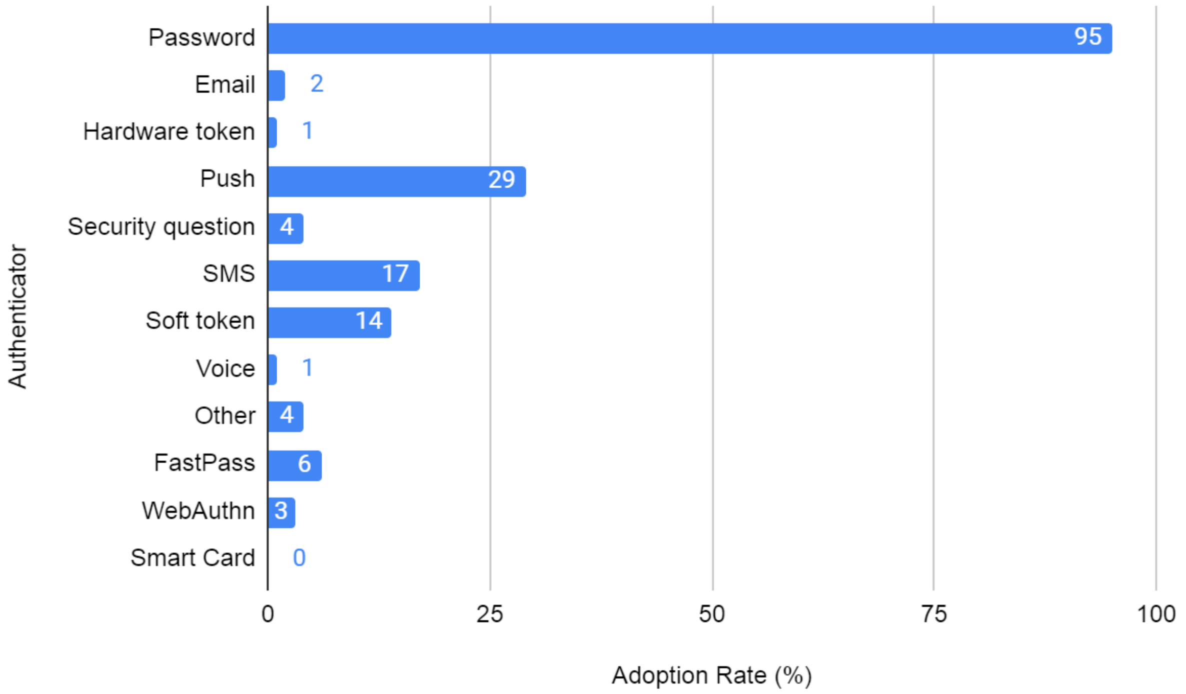

Fig. 2. Adoption Rates of Multi-Factor Authentication Methods (Okta, 2024)

To detect and localize intrusion attempts, distributed sites aggregate telemetry into a SIEM core and deploy SOAR playbooks. Automated scenarios reduce MTTR: research shows that orchestration cuts response time to minutes, diminishing direct losses and reputational risks (Grand View Research, 2024). At the same time, the global security orchestration, automation, and response market size was estimated at USD 1.72 billion in 2024 and is projected to reach USD 4.11 billion by 2030, growing at a CAGR of 15.8 % from 2025 to 2030, as shown in Figure 3.

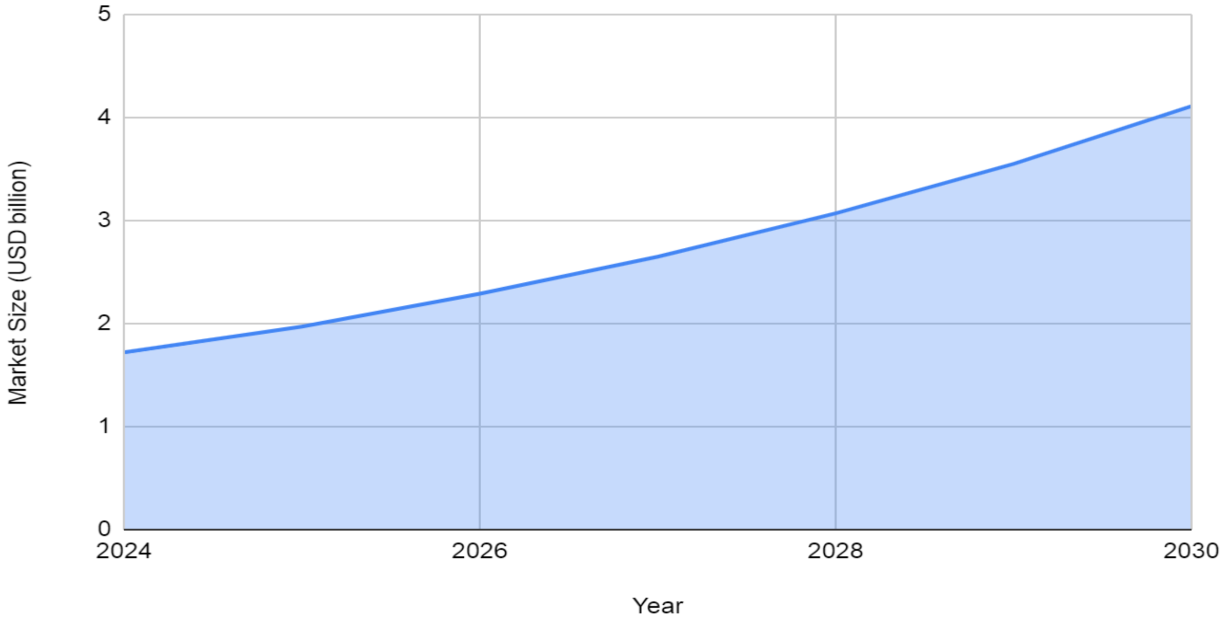

Fig. 3. The global security orchestration, automation, and response market size (Grand View Research, 2024)

IBM confirms the additional financial impact: companies employing AI automation in SOCs save, on average, USD 2.22 million per incident (IBM, 2024). The compliance loop is closed through periodic audits. A multisite structure must provide uniform observability across all locations to verify conformity with DORA, GDPR, and local critical-infrastructure criteria. A centralized control log and unified policies enable audits to be independent of site geography.

In transitioning to operations, crisis preparedness is maintained through combined NOC/SOC centers, supplemented by regional on-call teams. A side effect of digitization is a skills shortage: 67% of companies report a cyber-specialist deficit, and 58% believe this directly elevates organizational risk levels (ISC2, 2024). Automation becomes a compensator for the personnel shortage. Infrastructure as code, SDN, and CI/CD pipelines enable configuration management for all sites from a single version-control system. According to Datadog, 80 % of companies already use at least one IaC tool. Nevertheless, unpatched updates remain a convenient entry point: 32% of all ransomware attacks in 2024 began with a vulnerability for which a patch already existed (Datadog, 2024). Therefore, hot-patching using a synchronous repository and canary deployments is scheduled regularly alongside daily checksum verification.

Predictive maintenance completes the picture: correlating metric streams with machine learning helps predict node degradation before actual failure, freeing up time for planned traffic manoeuvres and reducing the burden on on-call staff. The human factor is addressed through organizational measures. Rotas are organized on a sun never sets principle: distributed teams provide a continuous response window, alleviating stress during peak shifts. Regular drills further mitigate the expert shortage: SANS' practical exercises for critical infrastructure have shown that participants who complete the course halve their post-incident stabilization time (SANS Institute, 2024).

Finally, clear communication channels constitute a separate pillar of resilience. Defined escalation routes, out-of-band backup channels, and pre-agreed single-window rules for the business shorten decision-making times and minimize the likelihood of duplicative actions, as every minute of downtime translates into direct losses.

In conclusion, the integration of multi-layered risk analysis and stringent BIA metrics with the selection of an optimal multisite topology—supported by automated SOAR playbooks and hybrid-cloud scenarios—enables achievement of required RTO/RPO and SLA even under natural disasters, cyber-attacks and social upheavals; reliance on international standards ISO, DORA and NIST, regular DRP/BCP testing and continuous MTTR/MTBF measurement transforms the architecture from paper insurance into a living, self-healing system capable of preserving business continuity and user trust in any crisis.

Conclusion

In the face of escalating economic, natural, and cyber crises, multisite network environments demonstrate their indispensable role in ensuring business continuity. The conducted analysis showed that only geographically distributed infrastructure enables the localization of failures—whether due to extreme weather events, social migration surges, or sophisticated cyber-attacks—and maintains the availability of critical services at levels required to meet stringent SLAs and internal SLOs. The foundation of this approach is a structured Business Impact Analysis, which translates failure scenarios into concrete RTO and RPO metrics, as well as into requirements for replication-channel bandwidth and the frequency of failover tests.

The choice between active–active and active–passive topologies, as well as the extension of the cluster into geodistributed fault domains, is determined by the ratio of redundancy costs to the economic damage from downtime. The financial efficiency of hybrid-cloud scenarios is confirmed by a 30–40% reduction in TCO, which, together with regular drills and automated SOAR-playbooks, transforms redundancy from a theoretical concept into a practical mechanism for instantaneous recovery. At the same time, it is critically important to integrate the requirements of international standards (ISO 22301, ISO 27001, DORA, NIST SP 800-34) into the design and operational cycle, thereby creating exhaustive checklists for BIA and mechanisms for verifying target metrics.

For true resilience of a multisite architecture, one-off investments are insufficient; continuous operational discipline is required, including ongoing monitoring of MTTR and MTBF, Error Budget management, automatic L7-level failover, and configuration drift control. Network segmentation and a Zero Trust model, along with multilayer encryption, MFA authentication, and centralized telemetry in SIEM/SOAR, provide not only reactive but also proactive protective measures. Predictive maintenance, based on machine learning and infrastructure-as-code, mitigates the human factor and personnel shortages, transforming multisite operations into an autonomous, self-healing process.

Thus, the systematic combination of multilayered risk analysis, precisely defined BIA metrics, a well-considered topology, and security automation forms a living, high-availability ecosystem within multisite networks. With regular DRP/BCP testing, strict regulatory compliance, and continuous measurement of key metrics, the architecture becomes not merely a technological solution but a strategic asset capable of withstanding any crisis challenge, preserving user trust and the company's financial stability.

References

1. Alicke, K., & Foster, T. (2024, October 14). Supply chains: Still Vulnerable. McKinsey & Company. https://www.mckinsey.com/capabilities/operations/our-insights/supply-chain-risk-survey
2. Amazon. (2025). Amazon Compute Service Level Agreement. Amazon. https://aws.amazon.com/ru/compute/sla/
3. Antithesis. (2024). How much does an outage cost? Antithesis. https://antithesis.com/resources/cost_of_outages/
4. Datadog. (2024, April 17). State of DevSecOps. Datadog. https://www.datadoghq.com/state-of-devsecops/
5. EIOPA. (2025). Digital Operational Resilience Act (DORA). EIOPA. https://www.eiopa.europa.eu/digital-operational-resilience-act-dora_en
6. Environment Analyst Global. (2020, September 23). Accenture quantifies the green benefits of migration to the cloud. Environment Analyst Global. https://environment-analyst.com/global/105947/accenture-quantifies-green-benefits-of-migration-to-the-cloud
7. Grand View Research. (2024). Security Orchestration, Automation & Response Market Report 2030. Grand View Research. https://www.grandviewresearch.com/industry-analysis/security-orchestration-automation-response-market-report
8. IBM. (2024). Cost of a data breach report. IBM. https://www.ibm.com/reports/data-breach
9. ISC2. (2024). Employers Must Act as the Cybersecurity Workforce Growth Stalls and Skills Gaps Widen. ISC2. https://www.isc2.org/Insights/2024/09/Employers-Must-Act-Cybersecurity-Workforce-Growth-Stalls-as-Skills-Gaps-Widen
10. ISO. (2019). ISO 22301:2019. ISO. https://www.iso.org/standard/75106.html
11. NCEI. (2025, January 9). Assessing the U.S. Climate in 2024. NCEI. https://www.ncei.noaa.gov/news/national-climate-202413
12. Okta. (2024). The Secure Sign-in Trends Report. Okta. https://www.okta.com/sites/default/files/2024-10/Secure%20Sign-in%20Trends%20Report%202024.pdf
13. SANS Institute. (2024). Critical Infrastructure Strategy Guide for 2024: A Call to Action for Securing ICS/OT Environments. SANS Institute. https://www.sans.org/press/announcements/sans-institute-unveils-critical-infrastructure-strategy-guide-2024-call-to-action-securing-ics-ot-environments/
14. Security Magazine. (2021). Only 54% of organizations have a company-wide disaster recovery plan in place. Security Magazine. https://www.securitymagazine.com/articles/95521-only-54-of-organizations-have-a-company-wide-disaster-recovery-plan-in-place
15. Squalio. (2025). Flexera's 2025 State of the Cloud Report. Squalio. https://squalio.com/news/flexeras-2025-state-of-the-cloud-report
16. UNHCR. (2024). Global Trends Report 2024. UNHCR. https://www.unhcr.org/global-trends-report-2024
17. Uptime Intelligence. (2024). Executive summary of Uptime Intelligence report. Uptime Intelligence. https://datacenter.uptimeinstitute.com/rs/711-RIA-145/images/2024.Resiliency.Survey.ExecSum.pdf
18. Verizon. (2025). 2025 Data Breach Investigations Report. Verizon. https://www.verizon.com/business/resources/reports/dbir/