Kinetic-Cyber Convergence: Physical Destruction of Cloud Infrastructure and the New Resilience Imperative

How the Dissolution of the Physical-Digital Boundary Demands a Fundamental Rethinking of Cloud Security Architecture

Unofficial AI-assisted Research

Cloud Security Alliance AI Safety Initiative

Executive Summary

On March 1, 2026, Iranian drone strikes destroyed two Amazon Web Services data center facilities in the United Arab Emirates and damaged a third in Bahrain, causing outages across 91 cloud services and disrupting banking, payments, logistics, and enterprise software for millions of users across the Gulf region [1][2][3]. These strikes were not directed at military installations or government facilities. They were aimed at commercial cloud infrastructure—and they succeeded. The incident represents the first confirmed instance in history of a major cloud provider's Availability Zones being physically destroyed by ballistic means, and it signals a fundamental shift in how adversaries project power in the digital age [4].

This event did not emerge without warning. For nearly a decade, military theorists, intelligence analysts, and cybersecurity researchers have described an accelerating convergence between kinetic (physical) and cyber operations in modern warfare. What has changed is that this convergence has now unambiguously extended to the commercial cloud infrastructure that undergirds global commerce, government services, healthcare delivery, and financial markets. Cloud architects and security professionals can no longer evaluate resilience exclusively through the lens of software faults, cyber intrusion, or natural disaster. They must now incorporate the possibility of deliberate physical destruction as a design parameter.

The implications are broad and urgent. Organizations operating critical workloads in cloud regions located within or adjacent to geopolitical flashpoints face a materially different risk profile than they did eighteen months ago. Regulators are responding: the European Union's Digital Operational Resilience Act (DORA), which entered into effect in January 2025, already requires financial institutions to demonstrate that critical functions are not unduly concentrated with a single provider or within a single geography [5]. The U.S. Cybersecurity and Infrastructure Security Agency released updated Cybersecurity Performance Goals (CPG 2.0) in December 2025 emphasizing resilience across IT and OT environments [6]. Yet most enterprise cloud architectures were not designed with kinetic threats in mind.

This whitepaper analyzes the kinetic-cyber convergence threat landscape, characterizes the physical attack surface of cloud infrastructure, and provides practical guidance for organizations seeking to build resilience architectures commensurate with the current threat environment. It connects these recommendations to existing Cloud Security Alliance (CSA) frameworks, including the Cloud Controls Matrix (CCM), Zero Trust guidance for critical infrastructure, and the CCSK body of knowledge.

1. Introduction: The Dissolution of the Physical-Digital Boundary

Cloud computing was conceived, marketed, and largely understood as a dematerialized resource—compute, storage, and networking delivered as abstract services from a notional "somewhere else." This abstraction was a feature. It freed developers from the operational burden of physical infrastructure and enabled organizations to build systems of global scale without maintaining global fleets of hardware. The abstraction was also, in retrospect, a cognitive hazard. It encouraged a view of cloud services as inherently non-physical, and therefore immune to the full range of physical risks.

The reality is that cloud infrastructure is deeply, inescapably physical. The services behind every API call ultimately depend on concrete, steel, and silicon: data centers built in specific geographies, connected to specific power grids, cooled by specific HVAC systems, and accessible by specific physical pathways. The scale of the world's hyperscale cloud providers has, if anything, intensified this physical dependence. AWS, Microsoft Azure, and Google Cloud collectively operate hundreds of data centers worldwide, many clustered in geographic regions that can be targeted as coherent entities. When threat actors—state or non-state—apply force directly to these facilities, the logical abstraction collapses.

The concept of kinetic-cyber convergence describes the systematic integration of physical military or paramilitary action with cyber operations to achieve strategic objectives. In its mature form, convergence works in both directions: cyber capabilities support kinetic targeting (reconnaissance, target acquisition, deception), and kinetic capabilities support cyber objectives (destroying backup communications, seizing physical infrastructure, disrupting defensive coordination). Critically, this model increasingly applies not only to government and military networks but to the commercial technology infrastructure that nations depend upon for economic function [7].

Understanding this convergence is prerequisite to designing cloud architectures, incident response plans, and organizational resilience programs that are fit for the current threat environment. The sections that follow trace the doctrinal development of this threat, examine its recent manifestations, characterize the physical attack surface of cloud data centers, and offer a framework for organizational response.

---

2. The Doctrine of Kinetic-Cyber Convergence

2.1 Theoretical Origins and Military Doctrine

The integration of electronic, cyber, and kinetic operations is not new in military doctrine. Russian theorists formalized the concept of "new generation warfare" in the early 2010s, describing a model in which information warfare, economic pressure, cyber operations, and conventional military force operate in coordinated combination rather than in discrete phases [8]. Western military planners subsequently adapted and extended this framework. By the late 2010s, NATO had formally elevated cyber to parity with land, sea, air, and space as a recognized domain of operations, acknowledging that critical infrastructure attacks could trigger collective defense provisions [7].

What has changed since approximately 2022 is the pace and ambition of real-world implementation. Adversaries are no longer practicing or demonstrating kinetic-cyber integration; they are executing it at operational scale. The Russia-Ukraine conflict, which began as a full-scale invasion in February 2022 and has continued through the writing of this document, logged more than 650 parallel cyber events timed with artillery, missile, and drone strikes in its first two years alone [9]. KA-SAT satellite communications were degraded via cyberattack hours before ground forces crossed the Ukrainian border. The Sandworm group—a unit of Russian military intelligence (GRU)—deployed Industroyer2 malware targeting Ukraine's power grid in a manner timed to coincide with kinetic pressure on the same geographic areas [10]. Physical access to data centers in Russian-occupied Ukrainian territory was exploited to compromise corporate networks through direct hardware manipulation [10].

These Ukrainian precedents established that the physical and digital attack surfaces of national infrastructure are effectively inseparable under conditions of conflict, and that well-resourced adversaries will exploit whichever attack vector is more accessible at any given moment.

2.2 Cyber-Enabled Kinetic Targeting

One of the most operationally significant aspects of kinetic-cyber convergence is the use of cyber operations to enable, guide, and refine physical strikes. Amazon's threat intelligence organization documented two particularly illustrative cases in its March 2026 threat intelligence disclosure [4]:

In December 2021, the Iranian threat group designated Imperial Kitten compromised the Automatic Identification System (AIS) platform of a commercial maritime vessel—gaining persistent access to real-time positional and operational data. Amazon's threat intelligence team identified and assisted in

remediating the compromise. In February 2024, U.S. Central Command reported that Houthi forces struck the exact vessel that Imperial Kitten had been tracking [4]. The cyber intrusion had served as a targeting system for a missile attack.

In May 2025, the Iranian threat group MuddyWater provisioned server infrastructure for cyber network operations. The following month, that infrastructure was used to access a compromised server hosting live CCTV streams from Jerusalem, providing real-time visual intelligence that could be used to observe and time attacks against targets within the city [4]. The pattern is consistent with a broader doctrine in which cyber reconnaissance feeds kinetic strike planning.

These cases demonstrate that the relationship between cyber and physical operations is not merely parallel—it is frequently serial. Cyber access precedes and enables physical destruction. This has direct implications for how cloud providers and enterprise customers should interpret network intrusion attempts. Reconnaissance that appears motivated by data theft may in fact be the precursor to physical targeting.

2.3 Kinetic Operations Enabling Cyber Objectives

The reverse relationship—physical action enabling or amplifying cyber effects—is equally relevant. Physical destruction of backup communications, power infrastructure, or redundant data paths can force systems onto degraded paths that are more vulnerable to cyber attack, create chaos that masks intrusion activity, or simply destroy evidence of prior cyber operations. In conflict zones, the physical seizure of storage media and hardware by military forces represents a form of data exfiltration with no digital signature. Russian forces operating in Ukraine reportedly exploited physical access to data center hardware in occupied areas to compromise networks that had otherwise resisted remote intrusion [10].

The combined effect of these bidirectional techniques is an attack surface that spans both domains simultaneously. Defenders who evaluate physical security and cybersecurity in isolation—which describes the overwhelming majority of enterprise security programs—are working from an incomplete model of their threat environment.

---

3. Cloud Infrastructure as a Theater of War: Key Incidents

3.1 Ukraine-Russia Conflict: Establishing Precedents

The Russia-Ukraine conflict established several precedents that shaped the subsequent trajectory of kinetic-cyber convergence. First, it confirmed that large-scale, sophisticated cyber operations would be conducted in coordination with conventional warfare rather than as a substitute for it. Second, it demonstrated that civilian infrastructure—power grids, communications systems, financial networks—would be targeted without restraint. Third, it revealed that organizations operating in conflict-adjacent geographies face risks that conventional enterprise continuity planning does not contemplate.

The CyberPeace Institute documented more than 2,000 cyberattacks linked to the conflict across entities in more than 50 countries [9], illustrating that the effects of geographically localized conflict can propagate far beyond the conflict zone through supply chain, service provider, and partner dependencies. Cloud service providers with infrastructure or operations in affected regions faced both direct physical risk and the secondary challenge of managing the security posture of a globally distributed supply chain under active attack conditions.

For cloud customers, the Ukraine conflict's most important lesson was that physical proximity to conflict —even through a cloud provider rather than directly—creates real operational exposure. Customers with workloads in AWS's Europe (Milan), Europe (Frankfurt), and European edge regions encountered elevated latency, rerouting, and in some cases service degradation as the provider managed both the direct cyber threat environment and the operational strain of migrating Ukrainian customer workloads to protected regions [11]. These disruptions were manageable. The March 2026 events in the Middle East were not.

3.2 The 2026 Middle East Escalation: A Watershed Event

On February 28, 2026, combined U.S. and Israeli forces conducted coordinated strikes on Iranian leadership compounds, IRGC facilities, and nuclear-related infrastructure in an operation designated Roaring Lion / Operation Epic Fury. The operation resulted in the death of Supreme Leader Ayatollah Ali Khamenei [2][3]. Iran's retaliatory campaign, which began within hours, combined widespread cyber

operations—reducing Iran's own internet connectivity to approximately 4% of normal levels due to counter-cyber operations [12]—with drone and missile strikes across Israel, the Gulf states, and U.S. and allied military infrastructure throughout the region.

On March 1, 2026, at approximately 04:30 UTC, Iranian drones struck AWS data center facilities in the United Arab Emirates. Two facilities in the UAE's ME-CENTRAL-1 region (Dubai) were directly struck; a third facility in Bahrain's ME-SOUTH-1 region (Manama) sustained damage from a drone that landed in proximity to the facility [1][2][3]. The strikes caused structural damage, disrupted power delivery systems, and initiated fire suppression responses that resulted in additional water intrusion damage [1]. Within hours, AWS confirmed that 91 services were impacted in ME-CENTRAL-1, including EC2 networking APIs, EBS storage volumes, S3 object storage, Lambda function invocations, and managed database services including RDS and DynamoDB [1][13].

The downstream impact was substantial. Consumer and enterprise services dependent on the affected regions reported service degradation or complete outages. Careem, the regional transportation and delivery platform, lost service across multiple markets. Regional payments processors Alaan and Hubpay experienced transactional failures. Major UAE financial institutions, including ADCB and Emirates NBD, reported disruptions to customer-facing banking services. Enterprise software provider Snowflake, with significant customer base in the Gulf, confirmed service degradation [3]. For organizations that had assumed cloud availability zones provided sufficient resilience, the incident was a sharp demonstration of the limits of that assumption.

Amazon issued guidance urging customers with workloads in the affected regions to consider migrating those workloads to other AWS regions, and to ensure data was adequately backed up in out-of-region locations [1][13]. This guidance, issued mid-incident, illustrated the fundamental gap between the recovery options available to organizations that had pre-positioned multi-region architectures and those that had not.

The strikes were, in Amazon's characterization, collateral damage from Iran's broad retaliatory campaign rather than deliberate targeting of AWS as an entity [4]. The distinction matters less than the result: commercial cloud infrastructure was physically destroyed by military force. Whether future incidents involve deliberate targeting of cloud providers as strategic objectives—rather than as incidental targets—is a question that security planners should treat as a matter of when, not if.

---

4. Physical Attack Surface of Cloud Infrastructure

4.1 Data Center Vulnerabilities

The physical attack surface of a commercial data center is substantially larger and more complex than the footprint of the computing equipment it contains. The facility systems required to operate server infrastructure at scale—power distribution, cooling, physical access control, fire suppression, communications backhaul—each represent distinct attack vectors that can disrupt services without directly damaging compute hardware.

Power infrastructure is among the most critical. Data centers typically rely on utility power supplemented by uninterruptible power supplies (UPS) and diesel generators for backup. Physical attack or sabotage targeting utility substations serving a facility, diesel fuel lines or storage, or the facility's electrical distribution systems can cause outages even when computing hardware is undamaged. The October 2025 AWS northern Virginia event—in which a power system glitch triggered more than 6.5 million website outages affecting banking, logistics, and government operations [14]—demonstrated how fragile this dependency can be even under non-adversarial conditions.

Cooling infrastructure is similarly critical and physically exposed. Modern hyperscale data centers consume enormous quantities of cooling capacity, typically delivered through chilled water systems, mechanical cooling towers, or direct liquid cooling arrays. Disruption or destruction of cooling infrastructure causes servers to overheat and triggers automatic shutdowns, effectively achieving the result of direct server destruction without requiring access to the data halls themselves. Building Management System (BMS) platforms that control HVAC and cooling components are increasingly IP-connected and often inadequately secured, relying on industrial protocols including BACnet, Modbus, and OPC UA that prioritize interoperability over authentication and integrity controls [15][16].

Communications backhaul—the fiber, microwave, and satellite connections that link data center facilities to the internet and to other regions—represents a third attack surface. Physical cutting of submarine or terrestrial fiber cables has precedent as a tool of state-sponsored disruption, and aerial drones can threaten above-ground cable infrastructure. The AWS Middle East incident interrupted connectivity at the application layer, but the physical destruction of fiber infrastructure serving a regional campus would present a qualitatively different recovery challenge, as restoration times for physical cable plant are measured in days to weeks rather than hours.

4.2 Building Management Systems as an Attack Vector

Building Management Systems occupy a particularly concerning position in the data center attack surface because they bridge the cyber and physical domains bidirectionally. BMS platforms control the physical environment—temperature, humidity, power routing, access, fire suppression—through software interfaces that are increasingly network-accessible. A cyber intrusion into a facility's BMS can achieve physical destruction without kinetic action: manipulating temperature set points to induce server thermal damage, disabling fire suppression to allow fire propagation, or triggering false fire events to activate water-based suppression systems over live hardware [16][17].

The protocols most commonly deployed in BMS environments were designed for reliability and interoperability in the era before ubiquitous network connectivity, and most lack native support for authentication, encryption, or access control. BACnet, deployed in the majority of large commercial building management systems, has no mandatory authentication in its baseline specification [15]. Modbus, widely used for power and HVAC control, transmits commands in plaintext without verification [16]. The convergence of IT and OT disciplines in modern data center operations has brought these systems into network environments they were not designed to inhabit, and the resulting security posture is frequently poor.

For adversaries seeking to maximize destructive impact from a cyber intrusion into a data center network, BMS access represents a high-value target. Combined with kinetic action—even relatively unsophisticated kinetic action—BMS vulnerabilities can amplify physical damage substantially. An attacker who disables fire suppression prior to a physical strike, or who manipulates cooling systems to force thermal shutdowns before infrastructure failure, achieves more destruction with fewer resources than a purely kinetic approach would allow. This represents the kinetic-cyber convergence model operating at the facility level.

4.3 Geographic Concentration Risk

Cloud providers achieve operational efficiency in part through geographic concentration: clustering multiple data halls within a region, co-locating interconnects and edge nodes, and routing traffic through regional hubs. This model minimizes latency, simplifies capacity management, and reduces operational complexity. It also creates a physical concentration of infrastructure that presents an attractive aggregate target. Destroying or disabling a well-defined geographic cluster of facilities achieves greater disruption than attacking equivalent distributed capacity.

The AWS Middle East architecture at the time of the March 2026 incident comprised multiple Availability Zones within the ME-CENTRAL-1 and ME-SOUTH-1 regions. While AWS's multi-AZ model is designed to tolerate the failure of a single availability zone through redundancy across the others within a region, it

was not designed to tolerate the simultaneous physical attack of multiple AZs within the same regional cluster—nor could it reasonably be, given the scale of force required to achieve that outcome in a non-conflict environment. When two UAE facilities and one Bahrain facility were simultaneously struck, the regional redundancy model was exhausted.

The geographic concentration problem is compounded for enterprises that have adopted single-region cloud architectures. Such organizations gain the redundancy benefits within a region under the failure modes the platform was designed to handle—hardware faults, software faults, localized power events—but inherit the full geographic risk without mitigation. In the context of kinetic-cyber convergence, this architectural choice becomes a material resilience liability when operations are concentrated in or near geopolitical flashpoints.

---

5. Threat Actor Taxonomy

5.1 Nation-State Actors

Nation-state actors represent the most capable and the most doctrinally sophisticated segment of the kinetic-cyber threat landscape. Their defining characteristics are the integration of intelligence collection, cyber operations, and kinetic capability within unified command structures, the application of long planning horizons, and access to resources that enable sustained, multi-vector campaigns. The cases described in Sections 2 and 3—Imperial Kitten, MuddyWater, Sandworm—reflect well-resourced state intelligence and military units operating with strategic direction and operational patience.

For cloud infrastructure, nation-state actors present several distinct threat profiles. Pre-positioning operations involve persistent access to cloud management planes, provider networks, or facility control systems that can be activated during a crisis to amplify kinetic effects or deny defensive response. Reconnaissance-to-targeting pipelines, as demonstrated by the Imperial Kitten maritime case, involve using cloud telemetry or provider network access to support physical strike targeting. Direct physical attack, as in the Iran-AWS case, involves applying military or paramilitary force directly to data center facilities when strategic objectives warrant. Intelligence preparation operations involve compromising cloud-hosted government or defense contractor systems to inform kinetic military planning.

Amazon's March 2026 threat intelligence disclosure documented a broader campaign of nation-state actors—including Russian, Iranian, and North Korean affiliated groups—systematically bridging cyber intrusion capabilities with kinetic operational planning [4]. This campaign predates the March 2026 strikes by years and reflects a deliberate strategic evolution rather than an opportunistic response to the specific conflict.

5.2 Non-State Actors

Non-state actors, including proxy forces, hacktivist collectives, and criminal organizations operating under national direction, represent a lower-capability but higher-frequency threat to cloud infrastructure in conflict-adjacent environments. The Houthi maritime campaign—which leveraged Iranian-provided intelligence including, in at least one documented case, cyber-collected targeting data—demonstrates that non-state actors can execute kinetic attacks against commercial infrastructure with effects comparable to those of state actors when equipped with precision munitions and quality targeting intelligence.

Hacktivist groups have historically focused on Distributed Denial of Service attacks, website defacement, and data exfiltration—attacks that are disruptive but not destructive. The increasing availability of commercial drone technology, including autonomous systems capable of precise delivery of small payloads, lowers the technical barrier to physical infrastructure attacks for motivated non-state actors. As kinetic-cyber convergence doctrine diffuses through hacktivist communities—particularly those ideologically aligned with state actors who openly practice this doctrine—the physical attack surface of cloud infrastructure expands to a threat population with greater numbers and less predictable targeting logic than organized state actors.

5.3 Criminal Organizations

Ransomware and cybercriminal organizations have not, to date, demonstrated interest in the physical destruction of cloud infrastructure—destruction eliminates the asset from which ransoms can be extracted. However, there is a plausible pathway through which criminal organizations could be engaged to support physical infrastructure reconnaissance or access facilitation as part of hybrid operations orchestrated by state actors. Iranian cyber operations have historically demonstrated a pattern of using criminal tactics—including ransomware—in conjunction with or as cover for state-directed destructive operations [18]. The distinction between criminal and state-sponsored activity in the kinetic-cyber domain may be less meaningful than it appears from organizational structure.

---

6. Cascading Failure and AI-Intensified Dependencies

The March 2026 AWS incidents illustrated a dynamic that will intensify as AI-driven automation extends deeper into enterprise technology stacks: the cascading failure effect of cloud region disruption. In 2019 or 2020, an organization whose cloud region experienced extended outage would face degraded service delivery and manual workaround burdens, but human-operated systems would continue to function in degraded mode. By 2026, the penetration of AI agents, automated decision systems, large language model APIs, and cloud-native automation into operational processes means that regional cloud failure often means complete process failure rather than degraded function.

The 2025 Data Center Frontier analysis of that year's major outage events documented this dynamic explicitly: as AI-driven systems become tightly coupled to cloud services, control-plane failures cascade immediately into operational paralysis rather than controlled degradation [14]. Banking operations that automated fraud detection, loan processing, and customer authentication on cloud AI services found those automations unavailable during the UAE outage—and backup manual processes had been atrophied or eliminated entirely during the preceding years of automation investment. This coupling amplifies the strategic value of cloud infrastructure disruption from an adversary's perspective and raises the resilience requirements for organizations that have adopted deep AI integration.

Cybersecurity Performance Goals 2.0, released by CISA in December 2025, explicitly addresses the cascading failure risk in IT and OT-converged environments, including guidance on architectural separation of AI decision systems from single-provider dependencies and requirements for graceful degradation planning [6]. Regulated industries are beginning to internalize these requirements; the broader enterprise landscape has not yet done so at scale.

---

7. Regulatory and Compliance Landscape

7.1 EU Digital Operational Resilience Act (DORA)

The European Union's Digital Operational Resilience Act entered into application on January 17, 2025, establishing mandatory operational resilience requirements for financial entities operating within the EU and for the critical ICT service providers that serve them [5]. DORA is notable in the present context for several reasons. It explicitly addresses ICT concentration risk—requiring that financial institutions ensure critical and important functions are not unduly concentrated with a single provider or small group of providers—which directly implicates geographic concentration in cloud architecture. It establishes an EU oversight framework for Critical Third-Party Providers (CTPPs), with 19 ICT service providers including AWS, Microsoft Azure, and Google Cloud designated as critical and subject to direct EU supervisory oversight since November 2025 [5]. And it requires contractual arrangements for exit strategies, data portability, and recovery capabilities that financial sector cloud customers must now negotiate and maintain.

The geographic scope of DORA's concentration risk provisions is significant: a financial institution that hosts its core banking systems exclusively in a single cloud provider's Middle East regions would, under DORA's principles, be required to demonstrate that concentration is managed and that recovery alternatives exist. The March 2026 events validate the regulatory concern underlying these requirements.

7.2 CISA Cybersecurity Performance Goals 2.0

Released in December 2025, CISA's Cybersecurity Performance Goals 2.0 (CPG 2.0) updated and expanded the foundational cybersecurity practices recommended for critical infrastructure owners and operators across all sectors [6]. The updated framework introduced expanded guidance on several dimensions relevant to kinetic-cyber convergence: supply chain risk management addressing both digital and physical dependencies; incident response and recovery requirements that encompass physical security incidents as well as cyber events; and IT/OT convergence security that acknowledges the physical control systems underlying digital infrastructure as in-scope security considerations [6].

While CPG 2.0 does not use the term "kinetic-cyber convergence" explicitly, its expansion of recovery and resilience requirements into physical security domains reflects an implicit recognition that the boundary between cyber incident and physical incident has become operationally meaningless in high-

threat environments. Organizations seeking to align with CPG 2.0 should evaluate their recovery plans against scenarios involving simultaneous physical and cyber disruption, not only against the separate failure modes the framework's individual controls address.

7.3 Emerging International Standards and Frameworks

Beyond DORA and CPG 2.0, a broader constellation of regulatory and standards activity is coalescing around physical-cyber resilience requirements for cloud infrastructure. The EU NIS2 Directive, which entered into force in October 2024, expanded the scope of cybersecurity obligations to include physical security of network and information systems for essential and important entities [19]. The UK National Cyber Security Centre has published updated guidance on cloud security that addresses geographic risk for sensitive workloads. International standards bodies including ISO/IEC are actively developing updates to the ISO 27001 and ISO 22301 families to better address the intersection of information security and business continuity in cloud environments.

Organizations operating internationally should anticipate a continued tightening of regulatory requirements in this area, driven by the March 2026 events and the broader recognition that kinetic-cyber convergence has moved from theoretical risk to operational reality.

---

8. Cloud Resilience Architecture for the Kinetic Threat Era

8.1 The Geopolitical Risk Assessment

The first requirement for resilience architecture in a kinetic-cyber threat environment is an accurate geopolitical risk assessment that informs cloud region selection decisions. Cloud architects have historically evaluated region selection on the basis of latency to end users, data residency requirements, available services, and cost. These factors remain relevant, but they must now be supplemented by an analysis of the geopolitical risk environment in each region and the potential for conflict escalation affecting data center facilities.

This assessment should be conducted in partnership with the organization's risk and intelligence functions, and should address several questions: Is the region within a recognized geopolitical flashpoint or adjacent to one? Are there state or non-state actors with demonstrated intent and capability to target commercial technology infrastructure in the region? What is the lead time for workload migration or failover if conditions deteriorate? How dependent are critical business processes on services hosted exclusively in this region? The AWS guidance issued during the March 2026 incident—that customers should consider migrating workloads—was operationally sound but arrived too late for organizations that had not pre-positioned their recovery architecture.

8.2 Multi-Region and Multi-Provider Architecture

The foundational architectural response to kinetic-geographic risk is distribution: spreading critical workloads and data across multiple cloud regions, multiple availability zones, and where warranted, multiple cloud providers. The March 2026 incident demonstrated that AWS's multi-AZ resilience model within a single region—a highly effective design for the software and hardware fault scenarios it was engineered to address—is not sufficient protection against kinetic attacks affecting multiple AZs simultaneously. Protection against regional destruction requires multi-region architecture.

Multi-region architecture exists on a spectrum from passive warm standby to active-active deployment. At the minimum, organizations should ensure that data critical to business recovery is replicated to an out-of-region location on a continuous basis, with recovery point objectives (RPOs) that reflect the business cost of data loss. More robust architectures include active multi-region deployments in which workloads run concurrently across regions, with traffic routing that automatically shifts to unaffected

regions when one is disrupted. The AWS Disaster Recovery as a Service model and equivalent offerings from Azure and Google Cloud provide managed tooling for both approaches; the appropriate tier depends on the organization's recovery time objectives and the criticality of affected services [20].

Multi-provider architecture—distributing workloads across AWS, Azure, Google Cloud, or other providers—addresses a risk that multi-region architecture within a single provider cannot: the provider-wide impact of a catastrophic event or of retaliatory action specifically targeting a single provider's global infrastructure. This architecture introduces operational complexity, requiring that teams maintain proficiency across multiple cloud toolchains and that application architectures avoid provider-specific service dependencies. The complexity cost is real, and many organizations will appropriately conclude that robust multi-region within a single provider is a more practical resilience investment than multi-provider architecture. However, organizations operating genuinely critical national infrastructure, regulated financial services, or systems supporting defense functions should evaluate multi-provider architecture as part of their resilience strategy.

8.3 Sovereign Cloud and Data Repatriation Considerations

Sovereign cloud models—cloud deployments in which data and infrastructure remain within a specific national jurisdiction—have gained significant adoption momentum as a response to data residency and regulatory requirements. In the context of kinetic-cyber convergence, sovereign cloud presents a complex tradeoff. Locating data within a sovereign boundary may provide regulatory assurance and, in some contexts, greater physical protection under national defense arrangements. However, it also concentrates risk within a single geographic and political context. A sovereign cloud deployment in a nation that becomes a conflict participant or a target of adversary action may offer less resilience than a geographically distributed commercial cloud architecture [21].

Organizations evaluating sovereign cloud should apply geopolitical risk assessment with particular rigor, since the regulatory or contractual commitments that drive sovereign cloud adoption may create architectural constraints that are difficult to unwind if the risk environment changes. A hybrid model—sovereign cloud for regulated data categories combined with commercially distributed cloud for operational workloads—may provide the best balance of compliance and resilience in many enterprise contexts.

8.4 Physical Security Integration with Zero Trust Architecture

Zero Trust security architecture, as articulated in CSA's Zero Trust Guidance for Critical Infrastructure [see Section 10], is conventionally understood as a framework for identity-based access control in network environments. Its core principle—never trust, always verify—applies with equal force to physical

access. The kinetic-cyber threat model requires that organizations extend Zero Trust principles into their physical security programs, treating physical and logical access as components of a unified access control system.

This integration has several specific implications for data center and cloud security. Facility BMS and OT systems should be treated as security-critical assets equivalent to network infrastructure, subject to the same access control, logging, segmentation, and vulnerability management requirements. Physical access to data halls should be monitored, logged, and integrated with security information and event management (SIEM) systems such that anomalous physical access generates the same alert responses as anomalous logical access. Network segmentation should extend to BMS and facility control systems, isolating them from general enterprise networks and from the internet with the same rigor applied to OT environments in industrial settings [15][16].

Zero Trust additionally implies continuous verification of the integrity of the systems that control physical infrastructure. A BMS that has been compromised is functionally equivalent to a compromised firewall: it presents a facade of normal operation while enabling adversary objectives. Continuous monitoring of BMS telemetry for anomalous command patterns, unexpected set point changes, or unusual access attempts provides early warning of cyber-enabled physical attacks before they achieve their objectives.

8.5 Continuity Planning for Kinetic Scenarios

Business continuity and disaster recovery planning in most organizations has been developed against scenarios that include hardware failure, software failure, natural disaster, and cyber attack. Kinetic scenarios—physical destruction of data center infrastructure by deliberate human action—are rarely included in tabletop exercises or recovery playbooks. The gap between the scenarios organizations plan for and the scenarios they may face in a kinetic-cyber convergence environment is a resilience liability.

Continuity plans should be updated to include scenarios in which a cloud region becomes unavailable due to physical destruction or ongoing conflict, rather than only due to technical failure. These scenarios have different characteristics than technical failure scenarios: the timeline to restoration may be measured in months rather than hours; the provider may be unable to communicate reliable recovery estimates; adjacent regions may also be at risk, requiring migration to geographically distant alternatives rather than nearby substitutes; and the broader conflict context may create staffing, communications, and regulatory complications that standard DR plans do not address.

Recovery Time Objectives (RTOs) and Recovery Point Objectives (RPOs) established for technical failure scenarios may not be achievable in kinetic scenarios if backup infrastructure is insufficiently distributed. Organizations should evaluate whether their documented RTOs and RPOs assume recovery

resources are available and accessible—an assumption that may not hold if conflict conditions prevail in or near the cloud region where those resources are hosted.

---

9. Recommendations

Organizations operating cloud workloads should treat the convergence of kinetic and cyber threats to cloud infrastructure as an architectural requirement, not as a low-probability scenario to be addressed in contingency planning. The following recommendations are organized by urgency and scope.

Immediate actions for organizations with workloads in or adjacent to geopolitical flashpoints include auditing cloud region dependencies to identify workloads without out-of-region data replication, activating multi-region failover capabilities for critical services and testing them under simulated regional loss conditions, reviewing BMS and facility control system network segmentation and access controls in owned or co-location facilities, and briefing executive leadership and boards on kinetic-cyber convergence risk with a clear articulation of current architecture's geographic risk profile.

Architectural investments that organizations should plan and fund in the near term include implementing continuous out-of-region replication for all data supporting critical business processes, with RPOs validated against kinetic loss scenarios; adopting multi-region active-active or active-warm-standby architectures for critical application tiers, eliminating single-region dependencies for production workloads; extending Zero Trust network access principles to facility OT and BMS environments, including network segmentation, continuous monitoring, and access logging for physical control systems; and conducting geopolitical risk assessments of cloud region portfolios, with review cycles tied to major geopolitical developments.

Strategic program investments include incorporating kinetic disruption scenarios into business continuity and disaster recovery plans, with tabletop exercises that test organizational response to scenarios in which a cloud provider's regional infrastructure is physically destroyed; evaluating multi-provider cloud architecture for workloads classified as critical national infrastructure, regulated financial services, or defense-supporting; engaging with regulatory frameworks including DORA, CISA CPG 2.0, and sector-specific resilience requirements to ensure architectural investments satisfy emerging compliance obligations; and establishing intelligence feeds and geopolitical monitoring capabilities to provide early warning of deteriorating conditions in relevant cloud regions.

---

10. CSA Resource Alignment

The kinetic-cyber convergence threat landscape engages multiple dimensions of the Cloud Security Alliance's established research and framework corpus. Organizations navigating this threat environment should consider the following CSA resources as foundational references.

The **Cloud Controls Matrix (CCM)** provides the primary controls framework for evaluating and improving cloud security posture. In the kinetic-cyber context, the CCM's controls in the Physical and Environmental Security domain (A&A, BCR, DSP) are directly relevant, as are Business Continuity Management and Operational Resilience controls that address recovery architecture and geographic distribution. Organizations should map their multi-region and physical security investments against CCM controls to identify gaps and to communicate security posture to governance and audit functions.

The **Zero Trust Guidance for Critical Infrastructure** publication addresses the extension of Zero Trust principles into IT/OT convergence environments, which directly maps to the BMS and facility control security challenges described in this paper. The guidance's framework for identity-based access control across IT and OT domains provides the foundation for extending Zero Trust into physical security system management. Organizations implementing the physical security integration recommendations in Section 8.4 should reference this guidance for control design patterns.

The **Disaster Recovery as a Service (DRaaS)** CSA publication provides architectural guidance for cloud-native recovery architectures spanning backup-only, pilot light, warm standby, and multi-site active-active tiers. The multi-region resilience recommendations in Section 8.2 align with the multi-site active-active and warm standby patterns documented in DRaaS guidance. Organizations seeking to move from single-region to multi-region architecture should use the DRaaS framework to define RTO/RPO objectives and select the appropriate tier of recovery architecture.

The **Cloud Adversarial Vectors, Exploits, and Threats (CAVEaT) Emerging Threat Matrix** provides a structured taxonomy of cloud-specific threats that extends traditional threat modeling frameworks for cloud environments. Kinetic attacks on cloud infrastructure and BMS exploitation represent emerging threat categories that security architects should incorporate into cloud threat models. The CAVEaT framework provides the vocabulary and structure for formalizing kinetic-cyber threats within existing cloud security programs.

The **CCSK (Certificate of Cloud Security Knowledge) v5** body of knowledge provides the foundational treatment of cloud security architecture principles including multi-region deployment models, shared responsibility frameworks, and resilience design. Cloud architects and security

professionals seeking to build competency in cloud resilience design should reference the CCSK v5 curriculum, particularly its treatment of cloud governance, risk management, and architecture security domains.

The **MAESTRO** framework for agentic AI threat modeling is relevant for organizations that have deployed AI agents and large language model integrations dependent on cloud services. As Section 6 discusses, AI-driven automation intensifies the business impact of cloud region disruption. MAESTRO provides the framework for threat modeling AI agent dependencies on cloud infrastructure and for evaluating the cascading failure risks that AI automation introduces.

---

11. Conclusions

The March 2026 destruction of AWS availability zones in the UAE and Bahrain by Iranian drone strikes should be understood as a threshold event: a demonstration that commercial cloud infrastructure will not be protected from physical destruction by its civilian identity, its economic importance, or its role as neutral communications infrastructure. In the doctrinal framework of kinetic-cyber convergence, cloud infrastructure is not neutral—it supports military communications, intelligence analysis, financial operations that fund conflict, and the social media environments in which information warfare is contested. Adversaries who believe that destroying it serves their strategic objectives will destroy it.

Cloud security professionals, architects, and the organizations they serve cannot fully control the geopolitical environment that creates this risk. They can, however, design systems that survive regional destruction, maintain business function across provider failures, and recover from kinetic disruption on acceptable timelines. The architectural patterns required to achieve this—multi-region replication, active-active or warm standby deployments, physical-cyber security integration, BMS hardening, geopolitical risk-aware region selection—are mature and available. They have been available for years, largely implemented as optional best practices for performance and reliability. The events of 2026 make them necessary for survivability.

The regulatory environment is moving in the same direction. DORA's concentration risk requirements, CISA CPG 2.0's resilience framework, NIS2's physical security obligations, and the growing body of sector-specific resilience regulations are collectively expressing a policy consensus that will increasingly be expressed in enforcement. Organizations that align their cloud architecture with the resilience requirements implicit in these frameworks will find themselves better positioned both operationally and regulatorily as the kinetic-cyber threat landscape continues to evolve.

The physical-digital boundary has dissolved. Cloud infrastructure exists in geography, depends on power and cooling, and can be destroyed. Security programs that treat cloud as a logical abstraction insulated from physical risk are operating on a model that the world has disproved. Rebuilding those programs on a model that integrates physical and cyber risk—and that treats geographic resilience as a core design requirement—is the defining security architecture challenge of this decade.

References

[1] Amazon Web Services, "AWS Service Health Dashboard — ME-CENTRAL-1 Outage, March 1–4, 2026," AWS Status, March 2026. https://status.aws.amazon.com

[2] CNBC, "Amazon says drone strikes damaged 3 facilities in UAE and Bahrain," CNBC Technology, March 2, 2026. https://www.cnbc.com/2026/03/02/amazon-says-drone-strikes-damaged-3-facilities-in-uae-and-bahrain.html

[3] CNBC, "Banking, payments services disrupted after Amazon UAE data centers hit in drone strikes," CNBC Technology, March 3, 2026. https://www.cnbc.com/2026/03/03/iran-war-uae-drone-strikes-aws-data-centers.html

[4] Amazon Web Services, "New Amazon Threat Intelligence findings: Nation-state actors bridging cyber and kinetic warfare," AWS Security Blog, March 2026. https://aws.amazon.com/blogs/security/new-amazon-threat-intelligence-findings-nation-state-actors-bridging-cyber-and-kinetic-warfare/

[5] European Commission, "Digital Operational Resilience Act (DORA) — Regulation (EU) 2022/2554," Official Journal of the European Union, January 2025. https://www.digital-operational-resilience-act.com/

[6] Cybersecurity and Infrastructure Security Agency, "Cybersecurity Performance Goals 2.0 for Critical Infrastructure," CISA, December 11, 2025. https://www.cisa.gov/news-events/alerts/2025/12/11/cybersecurity-performance-goals-20-critical-infrastructure

[7] HSToday, "The Hybrid Threat Era: Kinetic Warfare, Cyber Escalation, and Global Instability," Homeland Security Today, 2025. https://www.hstoday.us/subject-matter-areas/cybersecurity/the-hybrid-threat-era-kinetic-warfare-cyber-escalation-and-global-instability/

[8] Recorded Future, "Preparing for Russia's New Generation Warfare in Europe," Recorded Future Research, 2025. https://www.recordedfuture.com/research/preparing-for-russias-new-generation-warfare-in-europe

[9] CyberPeace Institute, "Cyber Dimensions of a Hybrid Warfare," CyberPeace Institute, 2024. https://cyberpeaceinstitute.org/news/cyber-dimensions-of-a-hybrid-warfare/

[10] Trustwave SpiderLabs, "The Russia-Ukraine Cyber War Part 3: Attacks on Telecom and Critical Infrastructure," Trustwave, 2025. https://www.trustwave.com/en-us/resources/blogs/spiderlabs-blog/the-russia-ukraine-cyber-war-part-3-attacks-on-telecom-and-critical-infrastructure/

[11] Microsoft, "The hybrid war in Ukraine," Microsoft On the Issues, April 27, 2022. https://blogs.microsoft.com/on-the-issues/2022/04/27/hybrid-war-ukraine-russia-cyberattacks/

[12] CloudSEK, "Situation Report: Middle East Escalation (February 27–March 1, 2026)," CloudSEK Blog, March 2026. https://www.cloudsek.com/blog/middle-east-escalation-israel-iran-us-cyber-war-2026

[13] Data Center Dynamics, "AWS UAE suffers AZ outage after 'objects strike data center' and cause fire, amid Iran attacks," Data Center Dynamics, March 2026. https://www.datacenterdynamics.com/en/news/aws-uae-outage-after-objects-struck-the-data-center-cause-fire-amid-iran-attacks/

[14] Data Center Frontier, "2025 Data Center Failures: What Outages Revealed About Resilience in the AI Era," Data Center Frontier, 2025. https://www.datacenterfrontier.com/hyperscale/article/55340621/2025-data-center-failures-what-outages-revealed-about-resilience-in-the-ai-era

[15] Claroty, "The Power of Building Management System (BMS) Cybersecurity," Claroty Blog, 2025. https://claroty.com/blog/the-power-of-building-management-system-bms-cybersecurity

[16] Data Center Dynamics, "The cybersecurity blind spot in data center building systems," Data Center Dynamics, 2025. https://www.datacenterdynamics.com/en/opinions/the-cybersecurity-blind-spot-in-data-center-building-systems/

[17] Data Center Knowledge, "Cybersecurity Risks Threaten the Physical Infrastructure of Data Centers," Data Center Knowledge, 2025. https://www.datacenterknowledge.com/cybersecurity/cybersecurity-risks-threaten-the-physical-infrastructure-of-data-centers

[18] Halcyon AI, "Iranian Use of Cybercriminal Tactics in Destructive Cyber Attacks: 2026 Updates," Halcyon Research, 2026. https://www.halcyon.ai/ransomware-alerts/iranian-use-of-cybercriminal-tactics-in-destructive-cyber-attacks-2026-updates

[19] European Parliament, "Directive (EU) 2022/2555 on measures for a high common level of cybersecurity across the Union (NIS2)," Official Journal of the European Union, October 2024.

[20] Cloud Security Alliance, "Disaster Recovery as a Service," CSA Artifacts, accessed 2026. https://cloudsecurityalliance.org/artifacts/disaster-recovery-as-a-service

[21] QA Ltd., "When one cloud region falls: The real risk we keep ignoring," QA Resources, 2025. https://www.qa.com/resources/blog/when-one-region-falls-the-real-cloud-risk-we-keep-ignoring/

[22] DefenseScoop, "Commercial data centers emerge as targets in modern warfare after drones hit 3 AWS facilities," DefenseScoop, March 3, 2026. https://defensescoop.com/2026/03/03/commercial-data-centers-drone-warfare-amazon-aws/

[23] Fortune, "Iran's revenge: drones damage data centers for Amazon Web Services, reveal west's Achilles Heel," Fortune, March 3, 2026. https://fortune.com/2026/03/03/irans-revenge-drones-damage-data-centers-for-amazon-web-services-reveal-wests-achilles-heel/

[24] SecurityWeek, "Cyber Insights 2026: Cyberwar and Rising Nation State Threats," SecurityWeek, 2026. https://www.securityweek.com/cyber-insights-2026-cyberwar-and-rising-nation-state-threats/

[25] FPRI (Foreign Policy Research Institute), "Data Centers at Risk: The Fragile Core of American Power," FPRI, November 2025. https://www.fpri.org/article/2025/11/data-centers-at-risk-the-fragile-core-of-american-power/