STUDY NO. 1 · CRITICAL INFRASTRUCTURE

When the cloud catches fire

European regulatory framework: what NIS2, CER and DORA concretely require, and what they do not cover

Nizar Younes Mqam · ORCID 0009-0008-5549-7620

Study · September 2026, version 1.2 of 9 October 2026 (translated from French version 1.5) · Cahiers de technopolitique, no. 1

On 1 and 2 March 2026, Iranian drones struck two Amazon Web Services data centres in the United Arab Emirates and damaged a third in Bahrain, disrupting more than 109 AWS services in the UAE region alone. This note argues that this was not a cybersecurity incident but a doctrinal shift: for the first time, according to several analysts, a state deliberately targeted commercial data centres in wartime by conventional military means, to the point of rendering some of them unrecoverable.

Executive summary

What happened on 1 March 2026 is not a cybersecurity incident. It is a doctrinal shift. For the first time, according to several analysts, a state deliberately targeted commercial data centres in wartime by conventional military means, to the point of rendering some of them unrecoverable.

The question is no longer whether this can happen. It is whether your organisation would survive the same thing tomorrow morning. This document draws three conclusions. A data centre is no longer a neutral infrastructure: it is a sovereign territory with a physical address and a military strategic value. Any organisation that hosts its critical systems with a single hyperscaler is not renting a service; it is paying tribute to a suzerain on which it depends existentially. And no continuity plan is worth anything if it has never been tested before the crisis.

Three risks documented in these pages deserve immediate attention: the multi-cloud myth, which diversifies computing providers without diversifying physical connectivity; the regulatory blind spot, since NIS2, CER and DORA do not require testing the simultaneous physical destruction of two cloud regions, and only DORA requires the testing of physical scenarios; and digital vassalage, which leaves any organisation without a tested failover plan at the mercy of a foreign operator in a crisis.

A single question is enough to know whether this document concerns you: if your two main cloud regions disappear tonight, how long will it take you to resume your critical operations, and from where?

Methodological note: this analysis describes categories of vulnerability and systemic logics, not specific infrastructure nodes, in accordance with open publication standards in matters of critical security.

Foreword: why this document concerns you

Imagine a Monday morning. Your management application stops, your website disappears, the email system no longer responds. Your teams call IT support, which calls back a few minutes later, dazed: the servers did not go down because of a bug. They were physically destroyed by a military drone. This scenario is no longer fiction. On 1 March 2026, Iranian Shahed drones struck two Amazon Web Services data centres in the United Arab Emirates; the next day, a strike close to an AWS site in Bahrain damaged its infrastructure. Oracle, for its part, was the subject only of an Iranian claim, unconfirmed, on 2 April. In the UAE region alone, more than 109 AWS services were disrupted*, regional banks lost access to their systems, and many users discovered that the cloud they had believed to be immaterial had a physical address, walls, and a very real vulnerability to missiles. It is, according to several analysts, the first time that a state has deliberately targeted commercial data centres in wartime, even if some see them as targets of opportunity. Cyberattacks against civilian digital infrastructure preceded 2026, notably Russia's disabling of the Viasat satellite network in Ukraine in February 2022. What March 2026 changed was the irruption of the missile and the drone into a space that was believed to be protected by its commercial nature and its status as a neutral third party. It will not be the last time.

(*) Source: APIStatusCheck, "AWS Middle East Outage March 2026", 6 March 2026 [7]. AWS keeps the list of affected services, region by region, on its public dashboard (AWS Health).

Laying the foundations. Before building, one must understand. The chapters that follow lay the foundations without which no resilience decision can be taken seriously. A glossary first, because these terms mean very different things depending on whether one sits in the boardroom or in the server room, and because this gap in understanding is costly when crises strike.

A chronology next, because the Iranian strikes of 2026 are not an isolated accident but the culmination of a trend that has been accelerating since 2021. Then an honest deconstruction of the multi-cloud myth, which reassures without protecting. The reason is technical but simple: most multi-cloud architectures share components that are invisible yet common. The same DNS provider, the same content delivery network, the same backbone transit routes, the same regional internet exchange points. Two companies using AWS and Azure respectively can see their services paralysed simultaneously if the internet exchange point that connects them to the rest of the world is destroyed or saturated. Multi-cloud diversifies computing providers. It does not diversify physical connectivity. The European regulatory framework that turns these recommendations into obligations. And a twelve-month roadmap to move from reading to action.

Laying the foundations

Glossary

Hyperscaler. Refers to cloud operators on a global scale: Amazon Web Services, Microsoft Azure, Google Cloud, Oracle Cloud. What they have in common is that they operate hundreds of data centres spread across several continents and interconnected with one another, and sell to thousands of companies the computing power, storage and digital services they need without having to build their own infrastructure.

Cloud region. A geographical area in which a hyperscaler operates one or more interconnected data centres. AWS ME-SOUTH-1, for example, designates Amazon's Bahrain region. When a company chooses to host its data "in the cloud", it is in fact choosing a specific region, that is to say one or more physical buildings in a specific country, subject to the laws of that country and exposed to its geopolitical risks.

Multi-cloud. A strategy consisting of using the services of several hyperscalers simultaneously, for example AWS and Azure in parallel. Often presented as protection against outages and dependencies, this approach is insufficient if the two providers share the same physical connectivity infrastructure, which this document demonstrates in detail.

Internet exchange point (IXP). A physical infrastructure, generally a building, where several network operators connect with one another to exchange internet traffic. It is the invisible node that links data centres to the rest of the world. Destroying or saturating a regional IXP can simultaneously isolate companies that nonetheless use different cloud providers.

Backbone. Refers to the main transit arteries of the global internet: very high-capacity cables, often submarine, which carry data between continents and between major cities. A data centre whose backbone is cut remains physically intact but becomes digitally inaccessible.

DNS (Domain Name System). The system that translates human-readable addresses such as "mybank.ma" into numerical addresses that machines understand. If the DNS server that handles these translations goes down, the application or website becomes inaccessible even if the data are intact, exactly like a telephone switchboard that has stopped working: the lines exist, but nobody can be put through.

CDN (Content Delivery Network). A geographically distributed network of servers whose role is to bring digital content closer to end users in order to speed up its loading. A website or application that depends on a single CDN shares a common vulnerability with all the other customers of that CDN.

Control plane. The part of the cloud infrastructure that orchestrates and manages all operations: starting servers, routing traffic, managing access, allocating resources. It is the brain of the system. If the control plane is destroyed or isolated, all the infrastructure it drives becomes unusable, even if the data are physically intact.

Air-gapped. Said of a computer system that is physically disconnected from any network, internet or intranet. An air-gapped backup is a backup that cannot be reached remotely, neither by a cyberattack nor by a network outage. It is the most robust form of protection against remote attacks, but it requires rigorous physical management.

RTO - Recovery Time Objective. The maximum acceptable time between a service interruption and its full resumption. If your RTO is 4 hours, this means that your organisation has decided not to tolerate more than 4 hours of interruption for that service. Defining an RTO without regularly testing whether the backup infrastructure can actually meet it is one of the most frequent errors in continuity plans.

RPO - Recovery Point Objective. The maximum volume of data that an organisation accepts to lose in the event of an incident, expressed as a duration. An RPO of 1 hour means that the data of the last

60 minutes can be lost without this being judged fatal. Together, RTO and RPO define the real level of resilience of an organisation, not its declared intentions.

SLA - Service Level Agreement. A contract by which a cloud provider commits to a level of availability for its services, generally expressed as an annual percentage. An SLA of 99.9% allows about 8.7 hours of unavailability per year. This document shows why this type of contractual commitment does not cover physical destruction scenarios, which generally trigger force majeure clauses that exonerate the provider.

Kinetic strike. A military term designating an attack by conventional physical means: missiles, drones, bombs. As opposed to a cyberattack, a kinetic strike physically destroys the targeted infrastructure. This is precisely what March 2026 introduced into the range of threats bearing on commercial digital infrastructure.

Latency. The delay between the sending of a piece of data and its reception. It is expressed in milliseconds and depends in particular on the physical distance between the sender and the receiver. This is why an orbital data centre, even in low orbit, introduces an irreducible latency that makes it unsuitable for certain critical uses such as real-time financial transactions.

Chronology: how digital infrastructure became a military target

The Iranian strikes of 1 March 2026 did not come out of nowhere. They are the culmination of a trend that has been accelerating since 2021, each episode crossing a further threshold in the nature and intensity of the threats bearing on the world's digital infrastructure. Understanding this progression is essential if March 2026 is not to be treated as an isolated accident.

March 2021. The OVHcloud fire in Strasbourg. An accidental fire destroys OVHcloud's SBG2 data centre and severely damages SBG1. Several million websites suddenly disappear, including public services and companies that believed their data were backed up. The shock is not military, but it reveals for the first time on a large scale a truth that many refused to admit: the cloud burns, literally, and the continuity plans of most customers had never been tested against a scenario of total physical destruction. It is the first massive wake-up call about the physical vulnerability of commercial data centres.

February 2022. The disabling of the Viasat network in Ukraine. About an hour before the Russian invasion, a sophisticated cyberattack disables the KA-SAT satellite network operated by Viasat, disrupting Ukrainian military communications and cutting off tens of thousands of civilian modems across Europe. It is the first large-scale demonstration that a commercial digital infrastructure can be deliberately targeted as the first act of a military operation, even before the tanks cross the border. The distinction between civilian infrastructure and military target begins to fade.

February 2024. The submarine cables severed in the Red Sea. The dragging anchor of a cargo ship hit by a Houthi missile severs three major submarine cables in the Red Sea, immediately disrupting 25% of digital traffic between Asia, Europe and the Middle East. One of these cables needs five months of repairs, because the specialised vessels cannot safely access the conflict zone. This episode reveals a systemic vulnerability that few decision-makers had anticipated: submarine cables, which carry more than 99% of intercontinental data traffic, are as fragile as an overhead power line in a storm.

November 2024 - January 2025. The cable war in the Baltic Sea. Seven submarine cables, including one power cable, are damaged between November 2024 and January 2025 in the Baltic Sea. Most of these incidents are suspected of being sabotage by state actors, with Russia at the top of the list, even though at least one cut was judged accidental by the Swedish judicial authorities. What distinguishes this episode from the previous ones is repetition: in a time of formal peace, the boundary between accident, sabotage and act of war becomes undecidable, and it is this very ambiguity that serves the aggressor.

1 March 2026. The Iranian strikes on AWS and the Gulf data centres. Iranian Shahed drones strike two AWS data centres in the United Arab Emirates and then, on 2 March, damage an AWS site in Bahrain through a strike close by. It is, according to several analysts, the first deliberate strike by conventional military means against commercial data centres. More than 109 AWS services are disrupted in the UAE region. Regional banks lose access to their systems. Many users discover that the cloud they had believed to be immaterial had a physical address, walls, and a very real vulnerability to missiles. The doctrine changes.

What this chronology demonstrates is simple and worrying: in five years, the threats bearing on digital infrastructure have progressed from the industrial accident to the strategic cyberattack, then to presumed sabotage in peacetime, and finally to the deliberate kinetic military strike. Each step was treated as a surprise when it had been foreseeable since the previous step. The next step is just as foreseeable.

The multi-cloud myth: why diversifying your providers is not enough

The multi-cloud strategy is today presented by almost all consulting firms and IT managers as the natural answer to the risk of dependence on a single provider. The idea is intuitively appealing: if you use AWS and Azure simultaneously, the failure of one should not affect the other. This is true in a scenario of internal technical failure. It is false in a scenario of physical destruction or of a regional infrastructure outage. Understanding why is essential before taking any resilience decision.

The fundamental error of multi-cloud reasoning is to confuse the diversity of computing providers with the diversity of physical connectivity. These are two entirely different things. A computing provider is the entity that rents you servers and storage. Physical connectivity is the set of infrastructures that allow those servers to communicate with your users and with the rest of the world. And this connectivity rests on physical equipment which is not, for its part, duplicated by the fact of having two different cloud providers.

Let us take a concrete example. A company uses AWS for its business applications and Azure for its backups, believing itself protected. Its two providers route their data through the same regional internet exchange point, the physical building where the region's networks connect with one another. If this exchange point is destroyed by a strike, saturated by an attack, or simply deprived of electricity by an outage on the surface, AWS and Azure simultaneously become inaccessible from that region. The company loses its business applications and its backups at the same time, despite its two separate providers.

The same reasoning applies to other invisible layers of the infrastructure. Two different cloud providers may use the same DNS network to resolve domain names, the same CDN to distribute their content, the same backbone transit routes to carry data to end users. These common

dependencies are rarely documented in contracts or technical architectures, precisely because they are regarded as stable commodities. They no longer are in a context of hybrid conflict. The Iranian strikes of March 2026 have not yet illustrated this limit: only AWS was hit, and Azure, Google Cloud and Oracle reported no outage linked to the incident. But all of them operate in the same strike zone, and often share the same exchange points and the same cables. Multi-cloud diversifies servers. It does not, on its own, diversify the routes to the world.

Some organisations have anticipated this problem by contractually requiring entirely independent physical routes from their providers, with distinct internet exchange points and backup satellite connectivity. These advanced architectures partially reduce the risk of internal technical failure. They nevertheless remain a small minority in practice, and none of them had been designed to withstand the simultaneous physical destruction of two cloud regions in two different countries, precisely the scenario of 1 and 2 March 2026.

The operational conclusion is therefore the following. A genuine physical resilience strategy is not measured by the number of cloud providers used. It is measured by the ability to maintain or restore connectivity between your systems and your users over entirely independent physical routes, passing through different internet exchange points, different cables and, if possible, geographical areas exposed to distinct geopolitical risks. This is a very different level of requirement from the one that the commercial discourse around multi-cloud implies. And it is precisely what the following chapters make it possible to build methodically.

The five pillars of real physical resilience

Resilience is not a state that one reaches. It is an architecture that one builds deliberately, layer by layer, knowing exactly what each layer protects and what it does not protect. The five pillars that follow are not generic recommendations. They are the minimum conditions without which no continuity plan can withstand a scenario of physical destruction of the cloud infrastructure.

First pillar: identify the brain of your infrastructure. Every digital architecture has a single point of failure, the place whose destruction or isolation renders the rest unusable. It is often the control plane, sometimes a critical DNS server, sometimes a central database. The first step of any resilience approach is to map this place precisely, to name it, and to check that your continuity plan survives its total physical disappearance. A continuity plan that does not answer the question "what happens if this specific building burns down tonight" is not a resilience plan. It is a reassuring document.

Second pillar: diversify physical connectivity, not only providers. As demonstrated in the previous chapter, using several hyperscalers offers no protection if these providers share the same transit routes, the same internet exchange points or the same submarine cables. Real diversification requires that your critical systems can reach your users by at least two entirely independent physical routes, passing through distinct internet exchange points located in geographical areas exposed to different geopolitical risks. This is a level of requirement that most current commercial cloud architectures do not meet.

Third pillar: protect backups through total isolation. A backup hosted in the same data centre as the primary data, or in a data centre connected to the same networks, is not a physical resilience backup. It is a copy. Real protection requires air-gapped backups: physically disconnected from any network, stored at a geographically distinct site, and whose restoration has been tested in real conditions. The

minimum rule recognised in the industry is known as 3-2-1: three copies of the data, on two different media, one of them off-site and disconnected. In a context of hybrid conflict, the requirement is now added that this offline site be in a geopolitically distinct jurisdiction.

Fourth pillar: define and test failover plans before the crisis. An RTO and an RPO defined on paper but never tested are intentions, not guarantees. Failover to a backup infrastructure, under conditions of real stress, systematically brings to light problems that no one had anticipated: invisible dependencies between systems, configurations that have not been updated, access rights that no longer work, teams that no longer know exactly what to do because the procedure has not been rehearsed for eighteen months. Operational resilience requires complete, documented and regularly repeated failover tests, at least once a year for critical systems.

Fifth pillar: do not depend on a single digital suzerain for critical decisions. This pillar is the least technical and the most strategic. It concerns governance, not architecture. Any organisation whose operational continuity depends entirely on a contract with a single American provider, with no alternative that can be activated within 24 hours, has outsourced a sovereign decision to a foreign entity subject to its own legal, geopolitical and commercial constraints. Diversifying providers, hosting jurisdictions and fallback architectures is not only a technical decision. It is a governance decision that belongs to executive management and boards of directors, not to IT teams alone.

Together, these five pillars do not guarantee invulnerability. No architecture is invulnerable. They guarantee that if one of your systems is struck, destroyed or isolated, the damage remains contained, recovery is possible within a known and tested time, and no external actor alone holds the keys to your continuity.

In recent years the European Union has produced three major texts that turn digital resilience from an optional good practice into a binding legal obligation. Understanding what each requires, and above all what none of the three covers, has become essential for any executive or regulator operating in Europe or with European partners.

NIS2: the directive on the security of network and information systems. Applicable since October 2024, NIS2 considerably extends the scope of its predecessor by widening the sectors concerned and tightening the obligations. It now covers operators of critical infrastructure in eighteen sectors, including energy, transport, health, water, digital infrastructure and cloud services.

In concrete terms, it requires these organisations to put in place cybersecurity risk-management measures, to notify significant incidents within strict deadlines, 24 hours for an early warning, 72 hours for an incident notification, one month for the final report, and to demonstrate that their management has followed adequate training on digital risks.

The penalties for non-compliance must be able to reach at least 10 million euros or 2% of total worldwide annual turnover, whichever is higher, for so-called essential entities.

CER: the directive on the resilience of critical entities. Adopted in parallel with NIS2, in force since January 2023 and, like NIS2, applicable since October 2024, CER goes further on the physical side. It obliges operators of critical infrastructure to carry out risk assessments that explicitly include physical threats: natural disasters, terrorist acts, sabotage. It requires a resilience plan comprising business continuity measures, documented recovery procedures, and reinforced cooperation with the

competent national authorities. It is the text that comes closest to the issues addressed in this document, since it explicitly recognises that the threat to a critical infrastructure can be physical and not only digital.

DORA: the regulation on digital operational resilience for the financial sector. Applicable since January 2025, DORA is aimed specifically at financial institutions: banks, insurers, asset managers, market infrastructures. It imposes detailed requirements on the management of risks related to third-party providers of digital services, including cloud providers. In concrete terms, a European bank must now precisely map its dependencies on its hyperscalers, regularly test its operational resilience, and ensure that its cloud contracts include audit rights and continuity guarantees. DORA also introduces the notion of a "critical ICT third-party service provider", which means that entities such as AWS or Azure can be directly subject to the oversight of European financial regulators.

The blind spot that these three texts share. Despite their ambition and their reach, NIS2, CER and DORA leave open a point that has become a systemic vulnerability since March 2026. DORA goes furthest: its 2024 delegated regulation already obliges banks to test the loss of data centres, physical attacks, political instability in their provider's jurisdiction and large-scale power outages. CER requires the assessment of all risks, terrorism included, but not testing. None of these texts, however, explicitly requires testing the simultaneous physical destruction of two cloud regions in two different countries, which is precisely the scenario that the Iranian strikes materialised, and outside the financial sector nothing requires the testing of a physical scenario. An organisation can therefore be fully compliant while being incapable of surviving this type of event. This is the regulatory work that the events of March 2026 have opened up, and that neither regulators nor operators can continue to ignore. What this concretely means for an executive is simple: regulatory compliance is necessary but not sufficient. It protects against penalties. It does not protect against a strike. The two logics must coexist, and this is precisely what the roadmap in the following chapter makes it possible to build step by step.

A 12-month roadmap: from reading to action

Reading this document without drawing from it a dated action plan with assigned responsibilities turns a strategic analysis into an intellectual exercise. What follows is a guiding sequence, to be adapted to your sector, your IT maturity and your own regulatory constraints. It is not a turnkey implementation guide: the realities of a fifty-person SME and of a regional bank will not call for the same answers. It is a prioritisation framework designed so that each organisation can identify, in its own situation, where to start and in what order to act.

Months 1 to 3: map before deciding

The first mistake of organisations that react to a crisis is to take infrastructure decisions without having mapped what they are trying to protect. The absolute priority of the first three months is therefore a complete map of your critical digital dependencies. This means identifying precisely which systems are hosted where, by whom, in which geographical region, with which dependencies on third-party services such as DNS, CDN and internet exchange points. It also means identifying your single point of failure: the place whose destruction renders the rest unusable. This work must be carried out jointly by the technical teams and the business units, because the former know where the systems are, and the latter know which ones are truly critical to business continuity. The gap between

these two perceptions is almost always wider than one thinks. The deliverable of this phase is a document of one to two pages, accessible to executives without technical training, which answers three questions: where is the brain of our infrastructure, what happens if this brain disappears tonight, and how long does it take us to resume our critical operations in that scenario.

Months 4 to 6: test what you think you have

Most organisations already have continuity plans, backup contracts and multi-site architectures. The problem is not their absence. It is that they have never been tested in real conditions. The second phase therefore consists of organising a complete failover exercise on at least one critical system, simulating the total disappearance of the primary infrastructure in a realistic way. Not a technical test in the server room. An exercise involving operational decision-makers, business teams and, if possible, the external providers concerned. This exercise will systematically reveal problems that no one had anticipated: obsolete configurations, expired access rights, invisible dependencies between systems, procedures that no one knows how to carry out under pressure any more. That is precisely why it must be done now, and not at the moment of a real crisis. The deliverable of this phase is an exercise report documenting the gaps between the contractual RTO and RPO and the times actually observed during the test, together with a prioritised remediation plan.

Months 7 to 9: correct the vulnerabilities identified

The results of the failover exercise set the agenda for this phase. The corrections to be made are rarely all of the same kind. Some are technical and can be dealt with quickly: updating configurations, activating backup routes, restoring access. Others are contractual and take longer: renegotiating SLAs, incorporating physical resilience clauses into cloud contracts, requiring audit rights over the infrastructure of your critical providers. Others still are structural and involve investment decisions: migrating certain data to geographically distinct jurisdictions, setting up air-gapped backups, diversifying physical connectivity routes.

Prioritisation must follow a simple criterion: start with the vulnerabilities whose materialisation would be fatal to the organisation, regardless of their estimated probability. What March 2026 taught the world is that scenarios judged improbable do materialise. The prioritisation criterion is therefore not probability, but impact.

Months 10 to 12: embed resilience in governance

Physical resilience cannot remain a one-off project entrusted to IT teams. It must become a permanent attribute of the organisation's governance, in the same way as financial compliance or the management of legal risks. This final phase consists of institutionalising what has been built. This means putting the annual review of continuity plans on the agenda of the governing bodies, training non-technical decision-makers in the fundamental concepts documented in this glossary, and appointing a person explicitly responsible for the physical resilience of the infrastructure, with a clear mandate and dedicated resources.

It also means updating the organisation's risk register to include explicitly the scenario of a physical strike on cloud infrastructure, with an assessment of its potential impact and a description of the mitigation measures put in place. Under NIS2 and CER, this documentation is not optional. It constitutes the proof of compliance that regulators can request at any time.

At the end of these twelve months, your organisation will not be invulnerable. No organisation is. But it will have accomplished what the great majority of its peers have not yet done: it will know precisely what it is protecting, how it fails over, where it resumes from, and how long it takes. It is this difference, between those who know and those who assume, that March 2026 has made strategically decisive.

The geostrategy of digital infrastructure: a new disciplinary field. What we experienced on 1 March 2026 in Abu Dhabi will not be recorded in the annals as a business continuity incident. It will be the moment when a concept that had been circulating for years in academic circles and think tanks received its most brutal and most visible confirmation: the geostrategy of digital infrastructure exists, it is real, and it kills banking services.

The geostrategy of digital infrastructure

From classical geopolitics to techno-geopolitics

Classical geopolitics studies how physical geography, mountains, rivers, straits, natural resources, determines the balance of power between states. The Red Sea was already a strategic strait under the Pharaohs. What is radically new is that this same Red Sea today carries oil and data simultaneously. Seventeen submarine fibre-optic cables pass through it, carrying the bulk of digital exchanges between Europe, Asia and Africa. Cutting these cables produces the same kind of strategic damage as blocking the tankers. The geography has not changed. What has changed is what moves through it.

Techno-geopolitics, a field that the analyst Abishur Prakash has notably helped to develop under the name of the geopolitics of technology, goes further. It studies how technology itself becomes a territory, with borders, resources, zones of influence and conflicts of appropriation. In this reading, a data centre is not a factory. It is a sovereign digital territory. Locating this data centre under a Norwegian mountain rather than in Abu Dhabi is exactly the same strategic decision as choosing where to establish a military port.

Three converging disciplines

What the Iranian strikes revealed is the convergence of three intellectual fields that had until now operated in silos and that can no longer be thought about separately.

- **Geopolitics of critical infrastructure.** What it studies: the physical location of digital power: where the cables, the data centres and the critical internet nodes are, and who controls them. Reference author(s): Frédérick Douzet (IFG Paris 8), Castex Chair of Cyberstrategy.
- **Digital sovereignty.** What it studies: decision-making independence: who can cut off access to your data, who controls the algorithms, who holds the encryption keys. Reference author(s): Asma Mhalla, *Technopolitique* (Seuil, 2024); the programme *CyberPouvoirs*, France Inter.
- **Technological geostrategy.** What it studies: inter-state competition for mastery of the foundational technologies, semiconductors, AI, cloud, quantum, of which data centres are the most visible physical expression. Reference author(s): Abishur Prakash, *The World is Vertical* (2021); AI/energy reports.

Technofeudalism as an analytical lens

There is a fourth conceptual framework, an economic one this time, which sheds light on what happened on 1 March 2026 in a way that purely technical analyses do not allow. The economist Cédric Durand developed in his book *Technoféodalisme* (Technofeudalism; Zones, 2020) an analysis of digital rent: the major platforms no longer function as classical capitalist markets but as fiefs that extract value from their captive users without producing directly. His thesis concerns first of all distribution platforms, Google, Apple, Amazon Marketplace. This document extends this logic to the cloud infrastructure itself, an extrapolation that Durand does not formulate explicitly but that the events of March 2026 make difficult to avoid: when a company hosts all of its critical systems with a single hyperscaler, it is not renting a service; it is paying tribute to a suzerain on which it depends existentially. When an Iranian strike on AWS simultaneously takes offline banking and payment services in the Emirates and regional delivery applications, it is exactly this feudal structure that is revealed in all its fragility: a single technological suzerain, thousands of vassals who did not choose it for its geography, and when the castle burns, the whole fief collapses.

The resilience I recommend in this document is not only a technical measure. It is an act of political sovereignty: freeing oneself partially from digital vassalage. We can no longer talk about cybersecurity without talking about geography, and we can no longer talk about geopolitics without talking about digital infrastructure. This logic of vassalage does not apply only to the American hyperscalers. Chinese players such as Alibaba Cloud, Huawei Cloud and Tencent are subject to the 2017 National Intelligence Law, Article 7 of which imposes a legal obligation to cooperate with Beijing's intelligence services.

An organisation that migrated to these Chinese platforms in order to "diversify" its dependence would create a symmetrical vassalage under a different jurisdiction. Real digital sovereignty can only be achieved through an infrastructure whose jurisdiction is politically consistent with the organisation's own interests. The boundary between military attack and cyberattack has become as blurred as the boundary between a data centre and a military target. The strikes of 1 March 2026 in Abu Dhabi officially marked this shift.

The MENA/Maghreb perspective

This analytical lens takes on a particular sharpness for organisations operating between Europe, the Maghreb and the Gulf, because their structural exposure is different from that of European or American companies. It is at once more concentrated and less visible. More concentrated, because the Maghreb has almost no hyperscaler cloud region on its soil. To date, neither AWS, nor Azure, nor Google Cloud operates a region there; only Oracle opened, in early 2026, a public region in Casablanca. For the most part, a Maghreb organisation that adopts a cloud architecture therefore hosts its critical systems abroad and, in almost all cases, the choice has turned towards the Gulf regions, for reasons of latency, cultural proximity and the availability of the sector certifications required by local regulators. This absence of territorial anchoring is not a strategic choice. It is an imposed constraint, rarely named as such.

Less visible, because the dependence does not stop at the hosting of data. It extends to connectivity itself. In Algeria, Law no. 18-04 of 2018 legally obliges all operators to route their international traffic (excluding satellite) through the infrastructure of Algérie Télécom [REF-AT]. The dependence here is

not geographical, since Algerian cables connect directly to Europe (SEA-ME-WE 4 from Annaba to Marseille, ORVAL from Oran to Valencia), but institutional: a single operator holds all the country's exit gates. For Morocco, the international bandwidth used reached 7,351 Gbit/s at the end of 2025, half of the deployed capacity [REF-ANRT]: the margin exists, but there is no public way of knowing through which foreign nodes this traffic transits. Tunisia already had direct routes to Europe (SEA-ME-WE 4 to Palermo and Marseille, Hannibal to Italy). The Medusa cable, which landed at Bizerte in November 2025 and whose entry into service is announced for 2026, adds a direct route to Marseille, and also serves Algeria (Algiers, Collo) and Morocco (Nador, Tetouan). This diversification remains recent and has not yet been incorporated into the continuity plans of most of the organisations that depend on it [REF-MEDUSA].

The Africa Coast to Europe (ACE) cable, whose route runs along Africa's Atlantic seaboard for 17,000 km, linking 24 countries, theoretically offers an alternative route to Europe without passing through the Gulf [REF-ACE], but it does not land in any country of the central Maghreb: reaching it requires terrestrial transit or another cable. But for most Maghreb organisations, this route has not been contracted as an independent backup path. It shares transit nodes with the main routes, recreating exactly the problem described in the chapter on multi-cloud: apparent diversity masks a common physical dependence.

This absence of contracting is not documented in public reports; it is a gap precisely because no organisation has to date published its continuity plans with the details of its backup routes. It is, however, consistent with the level of peering recorded at the CAS-IX exchange point in Casablanca (operated by N+ONE): three registered networks, two of them PCH's, for a capacity of 30 Gbit/s [REF-CASIX].

Casablanca's internet exchange points are the most under-exploited geostrategic asset in the region. They could serve as an alternative transit node for MENA traffic to Europe without passing through the Gulf IXPs, provided their scale is changed. But their current level of peering remains insufficient to absorb an emergency regional failover. They lack direct interconnection agreements with tier-1 cloud operators, the bandwidth capacity to handle a crisis volume and, above all, contractual recognition as a backup route in the continuity plans of the organisations that would need it most. This is not a technical problem. It is a problem of political and regulatory decision.

This is precisely where the geostrategy of digital infrastructure meets regional industrial policy. Strengthening the CAS-IX exchange point, developing direct peering agreements with the hyperscalers on this node, and encouraging Maghreb companies to contract transit routes independent of the Gulf are decisions that belong to the national regulators, the ANRT in Morocco, the ARPCE in Algeria, the ATI in Tunisia, and not to the IT teams of the companies concerned.

Morocco has begun this movement: the cloud qualification framework published in the Bulletin officiel in August 2025 now requires public entities and infrastructures of vital importance to use only cloud providers qualified under this framework, and in July 2025 the government announced a 500 MW sovereign data centre project as part of the Digital Morocco 2030 roadmap [REF-DM2030]. These decisions go in the right direction. They do not, however, resolve the question raised by this document: having a sovereign data centre is not enough if the connectivity routes that lead to it still pass through the same transit nodes.

As long as this question remains without a tested answer, the question addressed to every Maghreb decision-maker remains open: if your main route disappears, have you checked that your backup route does not itself pass through the same node?

These recommendations concern the interdependencies between industrial infrastructure and operational continuity in the MENA area. They do not constitute interference in public policy.

A complete map of digital dependencies sector by sector, with operational recommendations adapted to the regulatory and budgetary constraints specific to each jurisdiction, will be the subject of a dedicated edition.

The future of physical infrastructure: from target to fortress. Data centres have become military targets. This sentence, which would have seemed exaggerated five years ago, is today a documented, claimed and repeated fact. It will transform the way these infrastructures are designed, located and protected, not because the industry chooses to, but because it no longer has a choice. What follows is not speculative foresight. These are five trajectories already under way, some operational for years, others under active development, all accelerated by the events of March 2026.

The future of physical infrastructure: from target to fortress

Underground and fortified data centres: a reality already in operation

The idea of placing a data centre under a mountain comes directly from the military doctrine of the Cold War. States built bunkers in granite to house their command centres: rock absorbs shock waves, protects against direct strikes and naturally insulates from outside heat. This logic never disappeared. The Pionen Data Center in Stockholm is its most visible demonstration: dug 30 metres deep into Swedish granite, it is a former civil defence shelter converted into a high-security data centre, with real armoured doors, standby diesel generators, and an architecture designed to withstand conventional strikes. Switzerland has industrialised this model for decades. Companies such as Mount10 or Safe Host sell Swiss political neutrality as much as Alpine granite: two arguments that March 2026 suddenly made very concrete for their customers. This model will be exported massively in the coming years, and not only to neutral countries.

Iceland offers a different but complementary variant. Rather than protection by rock, it offers protection by geographical remoteness, combined with rare energy advantages: underground water sources at a stable temperature provide natural cooling, while hydropower and geothermal energy supply electricity that is almost entirely renewable and cheap. A country without an army, a founding member of NATO, geologically stable and remote from current conflict zones. After 2026, this profile is worth its weight in gold.

Protection by rock is real, but it has a limit that the news of 2025 made brutally visible. At Natanz, the first Israeli strikes first destroyed the power supply on the surface, and the IAEA estimated that this outage alone may have damaged the underground centrifuges, before concluding that the hall itself had probably been hit by penetrating munitions. An underground data centre dependent on the national power grid shares exactly this vulnerability: the rock cushions the direct strike, with no guarantee against penetrating munitions, and offers no protection whatsoever against a power cut 50 kilometres away at a substation that is not buried. Pionen's batteries last twelve minutes at full load;

beyond that, everything rests on two diesel generators and on their fuel supply. In a prolonged conflict, this is not enough.

Real underground resilience therefore requires three simultaneous criteria:

Physical protection by rock, local electricity generation or generation buried deep enough to be protected itself, and a cooling system integrated into the underground architecture and not exposed on the surface.

This is why Iceland offers a rare level of energy resilience: part of the electricity comes from below ground, cooling comes from underground water sources, and part of the energy chain can be buried, even though three quarters of generation remain hydroelectric, and therefore on the surface. This is not an advantage of energy comfort. It is a geostrategic advantage.

A second vulnerability is rarely mentioned: human access. A data centre requires physical interventions that automation cannot entirely replace. In a conflict situation, the underground entrance is an obvious and inexpensive point of attack. The additional American strikes on the Isfahan site did not necessarily destroy the galleries: they targeted the tunnel entrances, which is enough to immobilise the facility in the medium term without physically destroying it. An underground data centre whose accesses are blocked remains intact, switched on, cooled. But it is unusable. Physical resilience must therefore include redundant entrances, buried and approaching from different directions, and access protocols that do not depend on a single point of entry.

The third vulnerability is network connectivity. The fibre-optic cables that connect an underground data centre to the rest of the world come up to the surface exactly as power cables do. Pionen in Stockholm is protected against bombs, but its links to internet exchange points must, at some point, join the city's network.

Cutting these connections does not destroy the data centre: it remains intact, switched on, cooled. But it is cut off from the world, as useful as a library whose door has been locked. Network resilience requires buried, redundant connections arriving from different geographical directions, with backup routes that share no common transit point.

One last systemic fragility deserves to be named frankly:

The paradox of the recommendation. If all critical actors follow the strategy advocated in this document and migrate to Iceland, Norway and Switzerland, these three countries in turn become priority strategic targets. The vulnerability is not eliminated; it is moved geographically. Real systemic resilience is not concentration in safe jurisdictions, but a distribution wide enough that no single strike, however massive, can simultaneously paralyse several critical nodes. In concrete terms, this worldwide distribution must include geographical areas with genuinely distinct risk profiles. Sub-Saharan Africa, notably Kenya and Rwanda, which are actively developing their sovereign digital infrastructure, offers a geopolitical exposure radically different from that of Northern Europe. Latin America, with Brazil and Chile as emerging infrastructure hubs, represents a third credible zone of diversification. It is a map of diversified risks, built on the same principle that this document has defended since its first page: no single strike, however massive, must be able to reach simultaneously nodes located under geopolitically independent jurisdictions.

It is a permanent displacement of the point of weakness towards a place where it is harder to reach, more costly to exploit, and less fatal if it is hit.

Underwater data centres: a commercial reality with displaced vulnerabilities

Air-gapped backups placed in an underground bunker constitute the most robust combination of protection available today. Three simultaneous isolations: logical, because the system is disconnected from the network and therefore unreachable remotely; physical, because the hardened structure absorbs shocks and withstands strikes; geographical, because the site is distinct from the main one and shares no common dependency with it. It is what the military call a hardened fallback position, the only type of infrastructure that withstands at the same time a physical strike, a cyberattack and a prolonged network outage. Between 2018 and 2020, Microsoft conducted the best-documented experiment in the sector with Project Natick: a data centre in a sealed container, submerged 36 metres deep off the Scottish coast for two years, without any human intervention. The results contradicted several initial intuitions. Server reliability was 8 times higher than that of an equivalent land-based data centre. The reason is counter-intuitive: when submerged, the container is filled with inert nitrogen, which eliminates oxygen and thus removes corrosion, the number one enemy of electronic components. The sea also absorbs the heat of the servers passively, removing most air-conditioning needs. The main limit that this project highlighted remains maintenance. When something fails inside, raising the whole container is costly and slow. The architecture must therefore be designed so that the servers operate without human intervention throughout their service life, a constraint radically different from that of land-based data centres. What this constraint implies operationally is not often spelled out: for the whole time a module is being raised, maintained and resubmerged, all of the infrastructure hosted in that module is unavailable. For a commercial operator that has contractually guaranteed its customers a minimum level of availability (SLA), this is a service continuity constraint that remains unresolved to date. Hainan's modular model partially answers this problem: 100 modules make it possible to raise one without affecting the others. But this redundancy has an infrastructure cost that is not yet reflected in economic comparisons with land-based data centres. It is for this reason that Microsoft chose in 2024 not to commercialise Natick: technically proven, but not yet economically viable at scale. What Microsoft did not cross, China crossed. In December 2023, the company Highlander submerged the world's first commercial underwater data centre 35 metres deep off Lingshui, in Hainan province. In February 2025, a new module was added; the whole can now process about 7,000 simultaneous artificial intelligence conversations per second. Ten companies have already signed contracts to use this infrastructure, and the plan ultimately provides for 100 submerged modules in this zone alone. HiCloud claims an overall energy efficiency gain of 40 to 60% for its Hainan modules, according to statements by the operator reported by Data Center Dynamics in 2024. Later announcements report more conservative gains: more than a fifth less consumption for the Shanghai-Lingang project, according to the Chinese government, and more than 30% energy efficiency for Lingshui in 2026. The energy advantage is real; its exact magnitude remains, at this stage, self-declared. This efficiency nevertheless rests on an assumption that deserves to be named: the stability of the water temperature. Coastal underwater data centres are located at shallow depths, precisely where the oceans are warming fastest under the effect of climate change. The announced energy efficiency is calculated on current water temperatures. It will gradually deteriorate, and cooling may become insufficient during episodes of exceptional marine heat. This is a vulnerability on a 10-20 year horizon that is not taken into account in current projections. But the

Chinese experience also reveals what this document has sought to demonstrate since its first page: each new form of infrastructure displaces vulnerability without eliminating it. Hainan lies in the South China Sea, one of the most militarily contested maritime areas in the world. A data centre protected from aerial drones remains exposed to other forms of underwater threat, and these are now documented, operational, and proliferating rapidly.

The first is the acoustic threat. The acoustic risk to submerged structures is not theoretical. In November 2024, Thales and FEBUS Optics signed a co-development agreement for an acoustic monitoring solution for critical undersea infrastructure, based on the detection of any vibrational anomaly at any point of a submerged structure. The very existence of this co-development agreement confirms that operators of critical undersea infrastructure now incorporate acoustic risk into their protection architectures (passive monitoring of hydrophone signals). The principles of underwater acoustic warfare, the exploitation of low-frequency vibrations to destabilise sensitive structures or equipment, have for their part been documented in military literature for several decades, and their application to submerged data centres constitutes an attack vector that the operators of these infrastructures now incorporate into their risk analyses. These devices are discreet, inexpensive compared with a missile, and leave no visible signature on the surface. It is a vulnerability that no one had anticipated when these infrastructures were designed, precisely because they were conceived as responses to aerial threats and not underwater ones. The countermeasure exists: it consists of vibration isolation systems for the racks and acoustic sensors for perimeter monitoring. But it is not yet incorporated into the construction standards of current commercial underwater data centres, Hainan included.

The second is the threat of autonomous underwater drones. "The age of seabed warfare, where undersea infrastructure will increasingly be targeted, has arrived," wrote H I Sutton in Naval News in February 2024. In March 2025, China developed a tool capable of cutting armoured submarine cables at a depth of 4,000 metres, mounted on a robotic arm using a diamond-coated grinding wheel rotating at 1,600 revolutions per minute. Russia operates a fleet of ships, submarines and underwater drones through its Main Directorate of Deep-Sea Research, specifically dedicated to operations against critical undersea infrastructure. These vectors are not projects: they are deployed operational capabilities. Between November 2024 and January 2025, seven submarine cables were damaged in the Baltic Sea, in incidents most of which are suspected of being sabotage by state actors. An autonomous underwater drone launched hundreds of kilometres from its target can approach undetected, locate the module by sonar, and neutralise it without leaving an exploitable signature. The response to this threat is under development: the defence ministries of six countries bordering the Baltic and the North Sea created in December 2023 the Seabed Security Experimentation Centre (SeaSEC), which in May 2025 held weeks of trials integrating underwater surveillance drones, distributed acoustic sensors and autonomous patrol vehicles. But offensive capability precedes defensive capability, as always. And connectivity remains the weakest link in the whole system. The cables that connect an underwater data centre to the rest of the network are submarine cables, exactly those whose vulnerability the Red Sea has shown: in February 2024, three cables there were severed by the dragging anchor of a cargo ship hit by a Houthi missile, disrupting 25% of traffic between Asia, Europe and the Middle East, and one of them needed five months of repairs because the specialised vessels could not safely access the area. An underwater data centre whose connections are cut remains intact, switched on, cooled, and completely inaccessible from the rest of the world. Resilience

is not a destination. It is a permanent displacement of the point of weakness towards a place where it is harder to reach, more costly to exploit, and less fatal if it is hit.

Orbital data centres: physical reality and unavoidable constraints

Several companies are working on orbital data centres. Starcloud (formerly Lumen Orbit), in the United States, is one of the most advanced. The European ASCEND study, funded by the European Commission and coordinated by Thales Alenia Space, published its feasibility results in June 2024. In low Earth orbit, between 400 and 2,000 km in altitude, a satellite has continuous solar energy and is beyond the reach of any conventional military strike. No Shahed drone climbs to 500 km. This sentence is accurate, but it calls for an important qualification: while no Iranian drone climbs to 500 km, China, Russia and the United States all have operational anti-satellite missiles capable of reaching low Earth orbit. An orbital data centre escapes the drones of regional hybrid conflicts. It does not withstand a confrontation between major powers. Stratification towards orbit therefore displaces the vulnerability towards a higher level of threat: from a regional conflict to a strategic confrontation between great powers. This displacement may be acceptable depending on the context, but it must be consciously accepted, not ignored. But physics does not negotiate.

The first constraint is latency, and it is irreducible. Even in very low orbit at 500 km, the signal takes 3 to 5 milliseconds for the round trip. This delay is invisible for archiving or scientific processing, but it is prohibitive for a real-time financial transaction or remote surgery. Orbital data centres will not be able to serve every use, only those that tolerate this delay. In October 2024, JAXA and NEC achieved optical inter-satellite communication by laser at 1.8 gigabits per second over a distance of 40,000 km, a world record for data rate in the 1.5 μm band, according to NEC. This is a real and significant advance, notably because a laser beam a few centimetres in diameter aimed precisely at a satellite is almost impossible to intercept or jam, unlike conventional radio communications. From the point of view of critical infrastructure security, this technology is therefore highly relevant. But it changes nothing about latency. Laser light travels at the same speed as radio waves. What the laser improves is the quantity of information carried in the same time, not the speed at which the signal makes the round trip. The physical constraint of latency remains intact, whatever the transmission technology used.

The second constraint is cooling, and it is counter-intuitive. One might think that the vacuum of space, at minus 270 degrees, is an ideal freezer.

In reality, without an atmosphere, the only way to remove heat is infrared radiation. To dissipate the watts produced by high-density servers, radiator panels of several hundred square metres per megawatt of installed power would be needed, a surface impossible to deploy with current launchers.

The third constraint is maintenance. Sending a technician to repair a server in orbit costs tens of millions of dollars per mission. The architecture must therefore be designed from the outset to operate without any human intervention for 10 to 15 years, which requires extreme hardware redundancy and fundamentally limits the computing density that can be carried on board.

A fourth constraint deserves to be added, because it escapes any geopolitical analysis: a Carrington-type event, whose probability of occurrence scientific estimates place between less than 2% and 12% per decade depending on the models used (the work of Riley, Space Weather, AGU, 2012, puts it at about 12%, while Moriña et al., Scientific Reports, Nature, 2019, give a range of 0.46% to 1.88%),

represents a low-probability but potentially catastrophic risk for any digital infrastructure not protected against large-scale electromagnetic surges. Current commercial satellites are not designed to withstand this level of exposure. An orbital data centre would be vulnerable to a natural phenomenon that no state actor can either provoke or prevent, and against which there is not yet any economically viable technical solution. It is the only threat identified in this document that is entirely independent of geopolitics.

Cold data, that is to say data that must be kept but are rarely consulted, legal archives, genomic sequences, digital cultural heritage, long-term scientific records, will move to orbit within a horizon of 15 to 20 years, provided that launch costs continue to fall at the pace dictated by launcher reusability.

This is not speculative foresight. It is a trajectory already under way, each step of which is conditioned by measurable economic and technological variables.

It nevertheless remains subject to a fundamental geopolitical condition:

Low Earth orbit is out of reach of regional conflicts, but not of confrontations between great powers. Storing legal archives or cultural heritage in orbit makes sense in a world where the main threat remains the drone or the theatre missile.

That sense disappears if the context tips towards a direct confrontation between nuclear powers. It is a strategic choice to be made consciously, not a technical self-evidence.

In simple terms: underground for the critical, underwater for the sensitive, orbit for the cold. Provided that the sky remains a sufficiently preserved space for this last layer to keep its promises.

This stratification is an architecture of destination. It is not put in place at the moment of the crisis: it must exist beforehand. This point is crucial and often misunderstood. Critical data are not migrated to an underground bunker in response to an imminent strike. Migration takes time and bandwidth, and presupposes that the decision to fail over was taken and tested long before the threat arrives.

Real operational resilience therefore requires that each layer of this stratification be accompanied by a defined failover plan: the RTO (Recovery Time Objective, the maximum acceptable time before service resumes) and the RPO (Recovery Point Objective, the maximum acceptable data loss) for each type of data and each type of incident.

An underground data centre that has never tested failover from the primary data centre is not a fortress. It is a backup room that no one knows how to use under pressure.

Towards a stratification of infrastructure according to risk

The time horizons indicated are conditional estimates, not predictions. The Concrete signal column anchors each trajectory in facts observable as of April 2026. The ⚠ marks indicate the residual vulnerabilities after the threat has been displaced.

HORIZON	TYPE OF INFRASTRUCTURE	OPTIMAL USE	CONCRETE SIGNAL	LEVEL OF REALITY
Today	**Underground / fortified bunker** In a geopolitically neutral country	Critical services, sovereign data, air-gapped backups.	**Pionen, Stockholm**: former civil defence shelter, converted, operational since 2008. **Mount10 and Safe Host**, in Switzerland, sell Swiss neutrality as much as Alpine granite, two arguments that March 2026 made very concrete.	**Operational. Argument reinforced by the AWS strikes of March 2026.** ⚠ Limit: power supply and cooling come up to the surface (cf. Natanz, 2025). Diesel autonomy: a few days only. Iceland (underground geothermal energy) offers a higher level of resilience.
2026 to 2030	**Coastal underwater** Submerged at a depth of 30 to 40 m	Distributed storage, hot data, reduced regional latency.	**Hainan, China**: first commercial underwater data centre (Highlander / HiCloud), operational in December 2023, extended in February 2025. **Project Natick** (Microsoft): reliability 8 times higher, proven over two years.	**Commercial viability proven. Deployments under way.** ⚠ Identified threats: acoustic warfare (low-frequency vibrations), autonomous underwater drones (China developed a cutting tool for 4,000 m, March 2025), connectivity cables exposed at the surface. Warming of coastal waters: energy efficiency to be monitored over a 10 to 20-year horizon.
2030 to 2035	**Orbital edge computing (LEO)** Micro-nodes embedded on satellites	Satellite processing, observation data, embedded AI, maritime surveillance.	**SpaceX Starlink**: computing integration under evaluation. **Starcloud (formerly Lumen Orbit)**: active roadmap and fundraising. **JAXA and NEC**: 1.8 Gbps inter-satellite laser link demonstrated in October 2024 over 40,000 km.	**Active roadmaps. First successful demonstrations.** ⚠ Out of reach of regional conflicts (drones, theatre missiles). Vulnerable to the anti-satellite (ASAT) missiles of China, Russia and the United States. This displacement towards orbit must be consciously accepted.
2035 to 2045	**High-density orbital data centre** Heavy infrastructure in low Earth orbit	Cold archives, massive scientific computing, very long-term heritage data.	**Conditional on three measurable variables:** 1. launch cost below USD 200/kg (SpaceX Starship under way); 2. solving the problem of thermal radiators (several hundred m² per megawatt); 3. geopolitical stability of low Earth orbit (ASAT risk).	**Conditional foresight. Trigger variables identified and measurable.** ⚠ Additional risk: extreme solar events (Carrington-type) liable to destroy unshielded electronics in LEO. No economically viable solution to date.

Conclusion: resilience, a philosophy as much as a technique

The cloud is not an abstract and immaterial infrastructure. It is concrete, cables, batteries, generators, and therefore a physical target, like a bridge, a power station or an airport. The Iranian strikes on AWS in March 2026 officially put an end to an illusion two decades old: the idea that data centres are neutral commercial infrastructures, outside geopolitical and military equations. Three frameworks now make it possible to read this reality with precision. The geostrategy of digital infrastructure, first, as a discipline that names what many sensed without being able to formulate it. Technofeudalism next, as an economic lens that explains why the destruction of a single digital castle can paralyse thousands of vassals on five continents. The stratification of physical infrastructure, finally, from the

underground bunker to the orbital data centre, as a twenty-year horizon of transformation that is already under way, not still to come.

Together, these three dimensions make it possible to think of resilience no longer only as a business continuity plan, but as a choice of sovereignty. This choice must be made with eyes open. What this document advocates is not a single fortress, however sophisticated. It is a federation of distributed sovereign nodes, each physical layer of which responds to a distinct level of threat. The underground protects against the conventional strike and remains under direct territorial control. The underwater escapes aerial conflicts while maintaining physical connectivity controlled by identified operators. The orbital removes cold data from any regional reach, provided that the launch jurisdiction remains politically consistent with the organisation's interests.

These three layers do not operate in silos. They operate as the nodes of a federated architecture: distributed so that no single strike paralyses the whole, but controlled so that sovereignty over each node remains clearly assigned. It is not a decentralised blockchain without authority. It is a sovereign federation in which trust is redistributed among actors who control one another, exactly like member states that cooperate without dissolving.

This document has demonstrated it section by section: each proposed solution displaces vulnerability without eliminating it. The underground data centre withstands the bomb better, but not the power cut on the surface. The underwater data centre escapes aerial drones but not underwater drones or acoustic warfare. The orbital data centre escapes regional conflicts but not the anti-satellite missiles of the great powers. A regulatory blind spot must also be named: outside the financial sector, no European text requires the testing of a physical destruction scenario, and even DORA, which requires the testing of the loss of a data centre or of a physical attack, does not address the simultaneous destruction of two cloud regions in two different countries.

The events of March 2026 have revealed open regulatory work that neither regulators nor operators can ignore. Resilience is not an impregnable fortress. It is an architecture of conscious trade-offs: choosing where to displace the risk, to what level of threat to expose it, and making sure that if that level is reached, the damage remains bearable. That is why it is a philosophy as much as a technique.

Real resilience is not technology. It is diversity of infrastructure, the ability to fail over quickly, and the political decision to no longer depend on a single digital suzerain. Because the castle can burn. The fortress, for its part, must hold.

References

All sources were consulted and checked between 1 March and 4 April 2026, then checked again against the primary sources on 7 October 2026. The numbering is that of the document.

Iranian strikes on AWS: facts and chronology

[1] Rest of World, "Iranian drone strikes at Amazon sites raise alarms over protecting data centers", 4 March 2026. https://restofworld.org/2026/iran-amazon-data-center-strikes/

[2] The Conversation, "Why Iran targeted Amazon data centers", 1 April 2026. https://theconversation.com/why-iran-targeted-amazon-data-centers-278642

[3] Fortune, "Iran's attacks signal a new kind of war", 9 March 2026. https://fortune.com/2026/03/09/irans-attacks-on-amazon-data-centers-in-uae-bahrain-signal-a-new-kind-of-war-as-ai-plays-an-increasingly-strategic-role-analysts-say/

[4] Bloomberg, "How Amazon Data Centers Became a Casualty of Iran War", 5 March 2026. https://bloomberg.com/news/articles/2026-03-05/how-amazon-data-centers-became-a-casualty-of-iran-war

[5] Data Centre Magazine, "Reports: Iran Missiles Hit AWS Cloud Site in Bahrain", 2 April 2026. https://datacentremagazine.com/news/iran-missiles-hit-aws-cloud-infrastructure-in-bahrain

[6] CNBC, "Data centers become military targets", 6 March 2026. https://cnbc.com/2026/03/06/iran-war-data-centers.html

Technical impact: AWS Health Dashboard and analyses

[7] APIStatusCheck, "AWS Middle East Outage March 2026", 6 March 2026. https://apistatuscheck.com/blog/aws-middle-east-drone-strike-outage-march-2026

[8] Network World, "Amazon waives March AWS charges", 30 March 2026. https://networkworld.com/article/4151880

Submarine cables and global connectivity

[9] Rest of World, "U.S.-Iran war threatens Gulf AI infrastructure", March 2026. https://restofworld.org/2026/us-iran-war-gulf-ai-submarine-cables/

[10] TeleGeography, "Submarine Cable Infrastructure and Hormuz", March 2026. https://resources.telegeography.com/submarine-cable-infrastructure-strait-hormuz

[11] Submarine Networks, "War in the Gulf Severs the World's Digital Arteries", March 2026. https://submarinenetworks.com/en/nv/insights/war-in-the-gulf-severs-the-world-s-digital-arteries

[12] German Marshall Fund, "Stuck in Hostile Waters", March 2026. https://gmfus.org/news/stuck-hostile-waters

Underwater threats and seabed warfare

[13] Naval News, "Cable Attack: New Undersea Threat Is Starting To Reshape Naval Wars", February 2024. https://navalnews.com/naval-news/2024/02/cable-attack-new-undersea-threat-is-starting-to-reshape-naval-wars/

[14] Asia Times, "PLAN's big underwater drones push undersea power toward US shores", December 2025. https://asiatimes.com/2025/12/plans-big-underwater-drones-push-undersea-power-toward-us-shores/

[15] Internationale Politik Quarterly, "The Underwater Battlefield: Protecting Submarine Critical Infrastructure", January 2025. https://ip-quarterly.com/en/underwater-battlefield-protecting-submarine-critical-infrastructure

[16] University of Washington / Jackson School, "Baltic Sea Undersea Cable Security", blog post by Sophie Himka, undergraduate student, 9 July 2025. https://jsis.washington.edu/news/baltic-sea-undersea-cable-security/

[17] Blueye Robotics / SeaSEC, "Surveillance of Critical Underwater Infrastructure using modern underwater technology", Oda Ryggen, July 2025. https://blueyerobotics.com/blog/surveillance-of-critical-underwater-infrastructure-using-underwater-technology

[18] Thales & FEBUS Optics, "Thales and FEBUS Optics sign strategic co-development agreement to protect critical undersea infrastructure", 6 November 2024. https://www.thalesgroup.com/en/worldwide/defence-and-security/press_release/thales-and-febus-optics-sign-strategic-co-development

Underground facilities: the Natanz empirical reference

[19] Institute for Science and International Security, "Post-Attack Assessment of the First 12 Days of Israeli and U.S. Strikes on Iranian Nuclear Facilities", David Albright and Spencer Faragasso, 24 June 2025. https://isis-online.org/isis-reports/post-attack-assessment-of-the-first-12-days-of-israeli-strikes-on-iranian-nuclear-facilities

[20] Arms Control Association, "Iran's Nuclear Facilities: Status Updates", June 2025. https://armscontrol.org/blog/2025-06-25/irans-nuclear-facilities-status-updates

[21] New Civil Engineer, "What we know about Iran's underground nuclear facilities", June 2025. https://newcivilengineer.com/latest/what-we-know-about-irans-underground-nuclear-facilities-19-06-2025/

European regulatory framework: NIS2, CER, DORA

[22] European Commission, "CER and NIS-2 Directives enter into application", October 2024. https://ec.europa.eu/newsroom/cipr/items/859754/

[23] Bird & Bird, "Understanding NIS2, CER, and CRA", March 2025. https://twobirds.com/en/insights/2025/understanding-key-eu-cybersecurity-legislative-acts-nis2,-cer,-and-cra

[24] nFlo, "CER: six months to full implementation", November 2025, updated February 2026. https://nflo.tech/knowledge-base/cer-six-months-critical-infrastructure-implementation/

Data centre projects of the future

[25] Microsoft Research, Project Natick. https://natick.research.microsoft.com/

[26] Data Centre Dynamics, "Project Natick: Microsoft's underwater voyage of discovery", January 2021. https://datacenterdynamics.com/en/analysis/project-natick-microsofts-underwater-voyage-discovery/

[27] Highlander / HiCloud, Hainan Undersea Data Center. https://www.highlander.com.cn

[28] Data Center Dynamics, "China's HiCloud reports commercial underwater data center modules are working", 2024. https://www.datacenterdynamics.com/en/news/chinas-hicloud-reports-commercial-underwater-data-center-modules-are-working/

[29] Lumen Orbit (now Starcloud), Orbital computing. https://www.starcloud.com

[30] Bahnhof AB, Pionen Data Center Stockholm. https://bahnhof.se/foretag/colocation/datacenter/pionen/

[31] JAXA and NEC, LUCAS inter-satellite laser communication, 1.8 Gbps, January 2025. https://nec.com/en/press/202501/global_20250123_01.html

Intellectual framework: digital geostrategy and techno-geopolitics

[32] Asma Mhalla, *Technopolitique* [Technopolitics]. Seuil, 2024.

[33] Cédric Durand, *Technoféodalisme. Critique de l'économie numérique* [Technofeudalism: A Critique of the Digital Economy]. Zones, 2020.

[34] Frédérick Douzet, Chaire Castex de Cyberstratégie [Castex Chair of Cyberstrategy] (2013-2018), IHEDN / Institut Français de Géopolitique (IFG) [French Institute of Geopolitics], Université Paris 8. Author's page (GEODE): https://geode.science/frederick-douzet/ IFG Paris 8: https://www.univ-paris8.fr/Institut-francais-de-geopolitique-IFG

[35] Abishur Prakash, techno-geopolitical analyses. https://www.abishurprakash.com

Carrington probability

[36] Riley, P. (2012). "On the probability of occurrence of extreme space weather events". Space Weather, American Geophysical Union. https://doi.org/10.1029/2011SW000734

[37] Moriña, D., Serra, I., Puig, P., & Corral, Á. (2019). "Probability estimation of a Carrington-like geomagnetic storm". Scientific Reports, Nature Publishing Group. https://doi.org/10.1038/s41598-019-38918-8

The MENA/Maghreb perspective

[38] [REF-AT-1] Loi n° 18-04 [Law no. 18-04], article 126, Journal officiel algérien [Official Journal of Algeria]. [REF-AT-2] Algérie Télécom, *Catalogue d'Interconnexion 2025-2026* [Interconnection Catalogue 2025-2026]. https://www.algerietelecom.dz/docs/document/autres/CATALOGUE-D-INTERCONNEXION-D-ALGERIE-TELECOM-2025-2026-14776.pdf

[39] [REF-MEDUSA] Tunisie Telecom, Medusa submarine cable (Bizerte-Marseille, 22 Tbps), partnership signed February 2025, landing at Bizerte on 1 November 2025. https://www.leaders.com.tn/article/37431-tunisie-telecom-celebre-l-atterrissement-du-cable-sous-marin-medusa-a-bizerte-un-tourna

[40] [REF-ACE] Africa Coast to Europe Submarine Cable (ACE), 2021. https://ace-submarinecable.com/le-systeme-de-cable-ace-relie-desormais-leurope-lafrique-de-louest-et-lafrique-du-sud/

[41] [REF-ANRT] ANRT, *Observatoires ANRT, Le Mémo T4-2025* [ANRT Observatories, The Memo, Q4 2025], March 2026. https://anrt.ma/sites/default/files/2026-03/Observatoires%20ANRT%20-%20le%20Memo%20-%20T4-2025.pdf

[42] [REF-CASIX] Internet Society Pulse / PeeringDB, Casablanca Internet Exchange Point, CAS-IX, November 2025 data. https://pulse.internetsociety.org/fr/ixp-tracker/ixp/760/

[43] [REF-DM2030-1] Arrêté du Chef du gouvernement n° 3-17-25 du 1er août 2025 fixant le référentiel des exigences de qualification des prestataires de services cloud [Order of the Head of Government no. 3-17-25 of 1 August 2025 setting out the framework of qualification requirements for cloud service providers], Bulletin officiel no. 7432, DGSSI Morocco (regulatory). [REF-DM2030-2] Agence Ecofin / Reuters, "Cloud hybride et souveraineté : le Maroc annonce un nouveau datacenter de 500 MW" [Hybrid cloud and sovereignty: Morocco announces a new 500 MW data centre], July 2025. + Bulletin officiel no. 7432, cloud qualification framework, 21 August 2025. https://www.agenceecofin.com/actualites-numerique/1107-129990 https://www.digital.ma/articles/le-maroc-impose-un-referentiel-cloud-pour-proteger-sa-souverainete-numerique

To cite this study

Mqam, N. Y. (2026). *When the cloud catches fire European regulatory framework: what NIS2, CER and DORA concretely require, and what they do not cover.* Cahiers de technopolitique, no. 1. https://doi.org/10.5281/zenodo.23270771

Cahiers de technopolitique, no. 1 · Study · September 2026, version 1.2 of 9 October 2026 (translated from French version 1.5)
Publisher and publication director: Nizar Younes Mqam · Place of publication: Tourcoing, France
DOI: 10.5281/zenodo.23270771
Contact: y.nizar@proton.me · https://technopolitique.eu/cahiers/n1/en/
Author's ORCID: https://orcid.org/0009-0008-5549-7620
© 2026 Nizar Younes Mqam. Published in open access under the Creative Commons Attribution – NonCommercial – NoDerivatives 4.0 International licence (CC BY-NC-ND 4.0). Sharing and citation permitted with attribution; modification and commercial use prohibited. creativecommons.org/licenses/by-nc-nd/4.0/