MEMO

July 2026

Digital Security and Artificial Intelligence: Evolving Threats and Perspectives

Cigref
SUCCEED
WITH DIGITAL

Cigref Memo

Digital Security and Artificial Intelligence

Evolving Threats and Perspectives

July 2026

TABLE OF CONTENTS

INTRODUCTION

1.1 ARTIFICIAL INTELLIGENCE IS ACCELERATING THE TRANSFORMATION OF THE CYBER THREAT

For several years now, Cigref has been continuously analysing developments in the cyber threat landscape, drawing on feedback from its 160 member organisations and work carried out with the main French, European and international authorities and expert communities. This analysis is based on operational reality. Large companies and public administrations are now facing a sustained increase in the frequency of attacks, a rise in their complexity and a continuous reduction in the time available to detect, contain and remedy them.

The first few months of 2026, however, mark a shift of a different nature. A series of closely spaced events has highlighted an acceleration of the transformations already observed over the past several years. Reports published by leading cybersecurity organisations confirm the collapse in the time between the disclosure of a vulnerability and its exploitation, the industrialisation of attacks aided by artificial intelligence, and the increase in the volume of incidents that security teams must deal with on a daily basis. At the same time, European financial regulators have stepped up their vigilance. The European Central Bank has asked the eurozone's leading banks to update their action plans to address risks related to artificial intelligence, whilst the European Systemic Risk Board has raised its assessment of cyber risk.

News from the summer of 2026 also revealed a development that extends beyond the realm of cybersecurity alone. The decision by the US authorities to subject Anthropic's state-of-the-art artificial intelligence models to export restrictions on the grounds of national security, followed by the temporary suspension of access for users not classified as *'US persons'*, demonstrated that certain artificial intelligence capabilities could now be subject to control measures comparable to those applied to other strategic technologies. Although these restrictions were lifted on 30 June 2026 and access to Anthropic's services was restored from 1 July 2026, this sequence of events serves as a reminder to European organisations that digital resilience also depends on continued access to the technologies underpinning their ability to protect themselves.

These developments are further reflected in the **Action Plan on Cybersecurity and Artificial Intelligence** presented by the European Commission in July 2026. This plan explicitly recognises that cutting-edge artificial intelligence models are profoundly changing the conditions under which cybersecurity is practised. It proposes several measures designed to strengthen Europe's capabilities for the assessment, testing, deployment and development of these technologies for the benefit of critical infrastructure, public administrations and businesses.

The events of recent months do not call into question the analyses conducted by Cigref over the past several years. Rather, they confirm these analyses and lend them new significance. The transformations observed are no longer solely a matter of evolving attack techniques. They now affect the conditions under which large organisations will be able to maintain, over the long term, their capacity to protect their information systems.

This paper offers an analysis of these developments, focusing on eight shifts which, taken together, signal a lasting change in the cyber threat landscape. Each shift is examined on the basis of observations gathered from Cigref members, analyses published by the main competent authorities and leading organisations, and the most recent events, in order to shed light on the consequences of these shifts for the resilience of large organisations.

1.2 A STRUCTURAL DEPENDENCE ON CRITICAL DIGITAL CAPABILITIES OUTSIDE EUROPE

Before setting out the shifts we are observing, it is worth recalling the economic and technological backdrop against which they are unfolding. France and Europe are experiencing a severe economic dependence on the structural dominance of US-Israeli players in the cybersecurity industry, which is increasingly curbing our strategic autonomy. This observation is, moreover, consistent with the overall conclusions of the study carried out by the consultancy firm Asterès on behalf of Cigref, and published in April 2025. This study estimates that, within the European Union, approximately 83 per cent of purchases of software and cloud services in the B2B market are made from the US digital industry. Financially speaking, the North American and Israeli ecosystems attract more than three-quarters of global venture capital investment in the sector, relegating Europe – which attracts between 5 and 15 per cent depending on the year – to the status of a technology consumer market.

To illustrate this, the technology companies occupying the leading positions in the European cybersecurity market are mostly non-European. Microsoft has become one of the world's leading players in cybersecurity and dominates, in particular, the identity, EDR, SIEM and cloud security markets. Alongside it are Palo Alto Networks, Fortinet, Check Point Software Technologies, CrowdStrike, Zscaler and SentinelOne.

This vulnerability is now being exacerbated by the financialisation of the sector, as exemplified by the activities of the US private equity fund Thoma Bravo. The latter is aggressively consolidating the global market through external growth, successively acquiring key companies in the access protection, messaging and detection sectors. The cybersecurity industry is now characterised by increasing concentration, which tends to impose its own economic and technological conditions. This concentration of capital encourages vertical integration of product offerings and reinforces the technological dependence of European users.

As early as June 2024, CESIN had publicly expressed concern about the trend towards extreme concentration of Anglo-Saxon cybersecurity providers within the Thoma Bravo fund. On that occasion, CESIN noted that many companies, which were implementing diversification strategies to avoid dependence on a single cybersecurity provider, were ultimately finding themselves trapped by these successive acquisitions.

This dependence no longer concerns only cybersecurity software or cloud services. The events of June 2026 show that it now extends to the most advanced artificial intelligence models. These are gradually becoming essential components of large organisations' protection, detection and remediation systems. Their availability is now a key factor in the resilience of information systems.

This industrial and financial concentration is leading major European organisations to depend on a limited number of players for capabilities that have become essential to their own protection. This development represents a major challenge to resilience, to which recent developments in artificial intelligence lend new significance.

Faced with this highly structured technological and financial ecosystem, where cutting-edge innovations are continually funded and absorbed by US industry, Europe is suffering from the gradual marginalisation of its own solutions. It is against this backdrop of industrial lock-in and the growing vulnerability of our supply chains for digital security products and services that the following eight shifts take on their full critical significance.

1.3 A SECURITY POSTURE THAT HAS BECOME ILL-SUITED TO THE PACE OF THE THREAT

The doctrinal frameworks, defence instruments and organisational structures on which our national security currently rests were designed for an environment that no longer exists – that of low-intensity conflict, where the time required for remediation and the compartmentalisation of crisis information were still viable options. Today, systematic observation of adversaries' modus operandi clearly demonstrates, against a backdrop of growing strategic asymmetry, that it is no longer possible to respond to tomorrow's threats using yesterday's paradigms. Merely making adjustments to the capabilities or budgets of our current systems will inevitably come up against the mathematical and operational reality of how the threat has evolved.

The eight shifts documented in this briefing note highlight the structural obsolescence of our traditional defensive postures. They demand a fundamental overhaul of our national digital security doctrine in order to restore resilience – both at the individual level within businesses and public administrations, and at the collective level across society and its economy – in the face of the speed and scale of the threat as it is already materialising.

To make this brief easier to read and to ground our analysis in the reality faced by our members, we have chosen to structure the presentation of each of these eight shifts according to a single analytical framework:

- The nature of the shift: a factual description of the technical or geopolitical disruption, as we can observe it in its temporal dynamics.
- Illustration: putting the phenomenon into context using statistical data or recently documented case studies.
- The impact on our collective stance: the practical reasons why our traditional defensive approaches are now proving inadequate in the face of this new reality.

This approach aims to share an objective assessment, an essential prerequisite for jointly developing a national response commensurate with these challenges. Events since spring 2026 confirm this analysis. The decisions taken by European financial supervisors, the initiatives launched by the European Commission and the developments observed in artificial intelligence models all point to the same reality: organisations must now adapt their security posture to a threat whose capabilities are evolving at an unprecedented rate.

2 DESCRIPTIVE ANALYSIS OF THE EIGHT SHIFTS

2.1 COLLAPSE IN 'TIME TO EXPLOIT' AND THE OBSOLESCENCE OF PATCH MANAGEMENT

Nature of the shift

The time taken between the discovery of a cyber vulnerability and its active exploitation by malicious actors – commonly referred to *as 'Time to Exploit'* (TTE) – has literally plummeted in recent years, becoming negative since 2024. The traditional defensive paradigm relied on the time it took for defenders to identify the vulnerability, test the patch and deploy it. However, the implementation of *patch management* strategies usually requires non-regression testing, which can be time-consuming and complex to ensure that critical business applications remain operational. As the actual patch deployment cycle within a large organisation takes an average of 60 to 150 days, perimeter defence based on a reactive approach to vulnerabilities has become mathematically obsolete in the face of attackers who now operate at machine speed.

Illustration

According to the *M-Trends 2026* report published by Mandiant, Google Cloud's cybersecurity subsidiary, the average time to exploit a vulnerability was 63 days in 2018. By 2024, this had turned negative, with a vulnerability being exploited, on average, one day before a patch was released. In 2025, experts estimate that the TTE is negative by 7 days. According to the same *M-Trends 2026* report, the time taken within the organised cybercrime chain to pass on initial access from a vulnerability *broker* to a ransomware affiliate has fallen from over 8 hours in 2022 to less than 30 seconds today. The *IBM X-Force Threat Intelligence Index 2026* report confirms this widespread trend, noting a 44 per cent increase in attacks exploiting publicly exposed applications in 2025. Furthermore, 56 per cent of disclosed vulnerabilities require no authentication to be exploited.

This finding is corroborated by the German Federal Office for Information Security (BSI), which has in recent days warned of a dramatic reduction in the time between the discovery of a vulnerability and attempts to exploit it, making patch cycles of several weeks both critical and dangerous. In this context, ANSSI itself fears what its Director-General describes as losing 'the battle against software vulnerabilities'.

As early as 2021, Cigref and its European partners were warning of the crushing burden of this imbalance. We were already highlighting the constant deployment of dozens of full-time equivalents and the millions of euros spent annually by large companies to mitigate the security flaws of software publishers. Today, with a negative TTE, this model – whereby software and digital services publishers externalise the technical and financial risk onto their customers – has reached its breaking point.

Impact on our collective posture

The speed at which threats evolve demonstrates that security can no longer rely solely on victims' ability to apply security patches with increasing urgency. The digital industry must urgently be subject to *'Security by Design'* standards comparable to those in other critical sectors. This has long been a key campaign for Cigref, which formally alerted the Prime Minister as early as November 2020, drawing on

the work of the General Council for the Economy (CGE). As early as 2018, the CGE recommended that *'security by default must become the norm in product* design' and that maintaining it should not require any *'explicit action on the part of the user*'.

Yet, years after this assessment – which was shared by government bodies and backed by the OECD's recommendations published in February 2021 in the report *'Enhancing the digital security of* products' – the regulatory framework is struggling to adapt. Whilst the European *Cyber Resilience Act* (CRA) represented a first step, it remains notoriously ill-suited to addressing the current collapse of the TTE. The legislation remains insufficiently prescriptive regarding the requirement for genuine native resilience and lacks adequate mechanisms to enforce real-time remediation on an industrial scale.

2.2 HYBRID WARFARE AND STATE-SPONSORED COVERT OPERATIONS

Nature of the shift

The digital domain has become the primary theatre for high-intensity conflict described as 'hybrid warfare'. This offensive doctrine, particularly employed by Europe's strategic adversaries, aims to destabilise a target state whilst keeping operations below the threshold of open armed conflict. To ensure 'plausible deniability' and conceal the political attribution of these attacks, we are witnessing a deliberate and systematic erosion of the traditional boundaries between state intelligence services and the cybercriminal ecosystem. States such as Russia or Iran can now outsource some of their espionage, sabotage or influence operations to digital mercenaries, private offensive security firms, hacktivist groups or ransomware affiliates. These actors act as genuine *cyber proxies*, agents who create a technological and organisational smokescreen and complicate the process of attributing attacks.

Illustration

ANSSI's *2025 Cyber Threat Landscape* report provides a clear picture of the scale of this state-sponsored outsourcing. Against the backdrop of European geopolitical tensions and the war in Ukraine, pro-Russian hacktivist groups, such as *NoName057*, are carrying out massive distributed denial-of-service campaigns against European institutions, acting as proxies for low-intensity destabilisation operations. At the same time, recent leaks of internal data have confirmed the existence of close personal links and de facto protection afforded by Russian intelligence services to major cybercriminal groups, such as the BlackBasta ransomware group.

This logic of state-sponsored mercenary activity also underpins China's offensive cyber ecosystem through a thriving private offensive cyber warfare industry (LIOP). Massive leaks of internal documents from commercial service providers in 2024 and 2025 have conclusively demonstrated that companies were operating directly on behalf of the Chinese state's intelligence services. The Islamic Republic of Iran adopts similar asymmetric strategies, relying on hybrid ecosystems to target the economies of Western countries under the guise of misusing legitimate administrative tools.

The compromise campaign targeting Salesforce environments in the summer of 2025, via the exploitation of the third-party tool Drift, illustrates this overlap between cybercrime and strategic destabilisation objectives. The nature of the targets – whether world-renowned French luxury brands or global security leaders such as Zscaler or Cloudflare – and the systematic targeting of strategic data suggest a manoeuvre that goes beyond mere extortion. This attack appears to be firmly rooted in a

hybrid warfare narrative in which criminal groups act as *proxies*, and illustrates the objective convergence of interests given that they explicitly claim to be targeting brands representative of Western interests. Under this guise, these actors carry out operations designed to undermine confidence in Western digital infrastructure, thereby effectively fulfilling the destabilisation objectives of hostile powers.

Impact on our collective posture

The growing use of *cyber proxies* is blurring the traditional distinction between responding to an attack under ordinary law – which falls within the remit of the judiciary – and countering state-sponsored aggression, which falls within the remit of diplomacy or the armed forces. Faced with this combination of military and non-military courses of action, our national defence doctrine can no longer operate within watertight administrative silos.

As the new *National Cybersecurity Strategy 2026–2030* quite rightly emphasises, France must, as a matter of urgency, mobilise all its judicial, technical, diplomatic, military and economic resources, under the auspices of the Cyber Crisis Coordination Centre, to counter this global threat. However, Cigref points out that the State cannot wage this fight alone. As large private companies find themselves de facto on the front line against these *proxies* acting on behalf of foreign powers, there is an urgent need to fully integrate the private sector into the sharing of technical intelligence on the threat. Maintaining administrative barriers in the face of hybrid cyber operations would amount to leaving the national economy defenceless. Whilst it is legitimate for the state to safeguard its operational security, it is now essential to establish a framework for sharing sensitive and actionable information in real time with private-sector operators.

2.3 THE MILITARISATION OF AI AND THE SHIFT TO ‘MACHINE SPEED’

Nature of the transformation

Generative and agent-based artificial intelligence systems are radically transforming the threat ecosystem by shifting it to ‘machine speed’. AI is no longer merely a subject of technological foresight. It has already been weaponised by offensive actors, whether state-sponsored or cybercriminals, and acts as a force multiplier. Algorithms now make it possible to industrialise and automate entire stages of the attack chain, notably target reconnaissance, the dynamic development of malicious code, and the generation of hyper-personalised social engineering campaigns on a very large scale. The asymmetry thus becomes cognitive, and the time required for planning, chaining vulnerabilities and executing an attack is dramatically reduced, effectively overwhelming human analytical capabilities and rendering our traditional perimeter defences ineffective.

Illustration

Mandiant’s recent *M-Trends 2026* report confirms this operational reality. Experts have observed the emergence of malware families that actively query running AI models in order to adapt their behaviour in real time and evade detection systems. Other malware, such as credential stealers, directly exploit AI command-line interfaces present on compromised systems to target and extract trade secrets.

In Germany, the BSI describes this shift as ‘profound’ for the protection of critical infrastructure, emphasising that AI makes existing attack methods faster, less costly and easier to scale up. CERT-FR

points out, however, that generative AI currently acts primarily as a 'performance and scaling accelerator' for advanced attackers, rather than as a fully-fledged, autonomous hacker.

Even more worrying from a strategic perspective, this acceleration is illustrated by the emergence of artificial intelligence models whose capabilities are becoming directly relevant to cybersecurity. The *preview* of Anthropic's state-of-the-art Mythos model, presented in early 2026, marked a significant milestone by revealing advanced capabilities for identifying vulnerabilities and assisting in the construction of exploitation chains. Faced with these dual possibilities, Anthropic had initially chosen to restrict access to these capabilities within the framework of the Glasswing project, reserved for a select group of US stakeholders.

Over the following months, the company had begun a gradual opening-up to certain non-US stakeholders, notably in Europe, as part of an effort to assess the conditions for wider access whilst maintaining appropriate control mechanisms. The simultaneous release of the Fable 5 and Mythos 5 models was intended to mark a new stage, based on the assumption that technical protection measures and algorithmic safeguards would enable their distribution to be properly managed.

However, this sequence came to an end on 12 June 2026. At 5.21 pm, the US administration decided to subject these foundation models to export restrictions on the grounds of national security. As it did not immediately have the technical means to identify and distinguish users falling within the category of *'US persons*', Anthropic temporarily suspended access to these models for users not falling within this category, whilst it strengthened its control mechanisms. The US Department of Commerce lifted these export restrictions on 30 June 2026, and Anthropic restored access to both models from 1 July, having strengthened its mechanisms for distinguishing between categories of users.

This episode sets an important precedent. It demonstrates that the most advanced artificial intelligence models are no longer solely a matter of commercial digital products. Some of their capabilities may now be regarded as strategic and subject to control mechanisms comparable to those applied to other sensitive technologies. For user organisations, this development introduces a new dimension to their resilience : the ability to maintain reliable, long-term access to the artificial intelligence technologies necessary for their protection.

Impact on our collective posture

This transformation is leading to an accelerated deterioration in the overall security of the digital meta-structure that underpins all the operational processes of society and its economy, and to a profound change in the conditions under which public and private organisations will be able to ensure their cybersecurity. As with other critical technologies, their availability and continuity of access will become key determinants of resilience. Given the European economy's systemic dependence on the dominant US and Chinese technology industries, and the intensifying economic rivalries, the status quo is becoming untenable.

Moreover, this status quo is being denied to us, to our own detriment, by our main competitors with unapologetic ruthlessness. The state and large companies can no longer hope to counter algorithmic threats using conventional analytical tools or by relying exclusively on external partners who deny us access to their own defences. It is therefore absolutely essential and a matter of the utmost urgency that France and Europe develop an autonomous AI capability for cyber defence. The development of specialised models for cybersecurity will require considerable investment that no single company – and probably no single European state acting alone – will be able to bear on its own. It therefore calls

for a collective effort involving end-user companies, governments and industry players in order to develop capabilities tailored to the protection needs of Europe's critical infrastructure.

2.4 THE WIDESPREAD EXPLOITATION OF ZERO-DAY VULNERABILITIES TARGETING CRITICAL INFRASTRUCTURE

Nature of the shift

The exploitation of IT vulnerabilities unknown to their own developers – commonly referred to as 'zero-day vulnerabilities' – has reached a critical milestone. Historically, zero-day vulnerabilities were the preserve of an elite group of leading state actors, due to the colossal costs and cutting-edge expertise required to discover them. Today, this weapon has tragically become commonplace. This widespread adoption can be attributed to the spectacular rise of a private, deregulated and highly lucrative industry specialising in commercial surveillance and offensive cyber operations. These commercial providers of intrusion services now operate with a financial and technological firepower that rivals, or even surpasses, that of state intelligence agencies. They methodically target the blind spots in our defences, particularly edge devices and critical infrastructure, where conventional detection tools are often ineffective. The asymmetry is total and allows attackers to exploit these vulnerabilities invisibly, even before a protective mechanism can be devised.

Illustration

Combined data from ANSSI's *2025 Cyber Threat Panorama* and Mandiant's *2026 M-Trends* report provide an alarming objective picture of this capability gap. The exploitation of vulnerabilities remains the primary vector for initial intrusion, accounting for nearly a third of attacks. The real turning point lies in the collapse of *the 'Time to Exploit'* (see trend 1). Attacks thus systematically precede the release of security patches by the software vendor.

Global threat analyses published in 2026 by Google's teams highlight that commercial spyware providers have, for the first time in history, overtaken states – including China and Russia – in the initial exploitation of these zero-day vulnerabilities. ANSSI has documented the widespread compromise of security equipment and enterprise applications via *zero-day* vulnerabilities by these private actors, who then supply turnkey intrusion tools to their clients.

Impact on our collective security posture

The industrialisation of these arsenals by the offensive private sector renders traditional cycles based exclusively on vulnerability management and reactive remediation structurally inadequate and definitively obsolete. This shift is exacerbated by the institutionalisation of private-sector offensive operations, as evidenced in particular by President Trump's new cyber strategy, published in March 2026. This strategy explicitly provides for the creation of incentives to enable the private sector to identify and disrupt adversarial networks.

In response to this privatisation of digital conflict, France has committed itself to the Pall Mall Process. Initiated with the United Kingdom at the bilateral summit in March 2023 and officially launched in London in February 2024, this process aims to establish an international framework to combat the proliferation and irresponsible use of commercial cyber-intrusion capabilities. Despite the support of 27 states for a code of best practice in August 2025, the ubiquity of zero-day vulnerabilities exploited on an industrial scale demonstrates that our national approach, based on post-disclosure response,

has reached a critical limit. The stability of our economic infrastructure is now exposed to a private offensive threat capable of continuously overwhelming the conventional defence capabilities of our organisations.

2.5 SYSTEMIC SATURATION THROUGH HYPER-VOLUMETRY

Nature of the shift

The massive industrialisation of cyberattacks, driven by automation and artificial intelligence, is in many cases overwhelming defence capabilities. This transformation is no longer based solely on technical sophistication, but on a hyper-volume of malicious events that is exhausting human resources within *Security Operations Centres*. This deluge is all the more undetectable as it now relies on the hijacking of authenticated sessions within cloud environments and via APIs. It is now estimated that 95 per cent of attacks originate from seemingly legitimate sessions, making it virtually impossible for a human operator to distinguish between legitimate and malicious traffic without extensive algorithmic assistance.

Illustration

Mandiant's *M-Trends 2026* report highlights the critical phenomenon of alert fatigue, where the massive volume of events – even those of low severity – leads to such saturation amongst cybersecurity teams that major alerts end up being ignored or addressed too late. This industrialisation is confirmed by the collapse in the average time taken to transmit a malicious access credential between two criminal groups, which has fallen from over 8 hours in 2022 to less than 30 seconds today.

ANSSI's *2025 Cyber Threat Landscape report* confirms this trend, noting that the threat is now systemic and affects the entire economic fabric. The Agency has observed a surge in data exfiltration resulting from the theft of access credentials. In 2025, although the majority of organisations still detect intrusions internally, the speed with which attackers gain administrative control of an environment is now measured in hours, outpacing, in most documented cases, the capacity for human response.

Impact on our collective posture

Hyper-volumetry renders the current cyber defence paradigm, based on real-time human analysis, structurally inadequate. As Cigref had already pointed out to the Prime Minister as early as November 2020, the growing human and financial resources devoted to security are increasingly being diverted away from companies' capacity for innovation, which is needed to counter the ever-evolving effectiveness of an industrialised threat.

This impasse is now recognised internationally. The United States' new cyber strategy of March 2026 explicitly provides for the adoption of AI-powered cybersecurity solutions – to use the term adopted – to ensure defence on the necessary scale. For France, this saturation means that the protection of the nation can no longer rely solely on administrative vigilance or the individual commitment of cyber analysts. It requires a shift towards an augmented and automated defence system, capable of dealing with the threat at 'machine speed', lest our economic fabric be permanently overwhelmed by attacks that exploit our own trusted infrastructure to blind us.

2.6 THE SHIFT TOWARDS COGNITIVE WARFARE AND STRATEGIC PARALYSIS

Nature of the shift

Cyberattacks are no longer limited to purely technical compromise, espionage or financial extortion. They are now part of a broader strategy falling within the realm of cognitive warfare. This subversive dimension directly targets economic cohesion, the European defence industry and critical infrastructure. The aim of these campaigns is no longer merely to destroy or extort, but to destabilise society and its institutions. By instilling a climate of constant and elusive threat, our adversaries – and sometimes our economic competitors – seek to impose strategic and psychological paralysis. The aim is to erode the confidence of citizens and economic actors in the ability of the state and major companies to protect them.

Illustration

ANSSI's *2025 Cyber Threat Panorama* highlights this growing overlap between cyberattacks and operations aimed at influencing and destabilising. The repeated targeting of hospitals, local authorities, digital service providers and operators of essential services by hybrid groups or *cyber proxies* goes far beyond mere criminal motives. The immediate media coverage of these operational disruptions, such as the recent attack targeting the ANTS, acts as a sounding board. Every successful attack on a link in our value chain contributes to the construction of a narrative of systemic Western vulnerability, which is meticulously exploited by adversarial powers to weaken our institutions and discourage innovation.

This narrative of our vulnerabilities, particularly in France, has also recently been exploited, in a 'low-key' manner, by major US technology companies, notably within European institutions in Brussels. They portray French aspirations for digital sovereignty as a costly 'whim' whose price is paid through increased vulnerability. By arguing that only solutions from the American giants are capable of guaranteeing the security of European institutions and the European economy, these players are shifting the battle onto the terrain of political legitimacy. Every successful attack thus becomes an argument in favour of greater technological dependence, transforming a security breach into a potential surrender of our strategic autonomy.

Impact on our collective stance

In the face of this cognitive warfare, the handling of cyber-attacks can no longer be relegated solely to the status of a digital news item or a crisis confined strictly to IT and the legal sphere. Dealing with these incidents in technical silos amounts to ignoring the broader strategic manoeuvre of which they are merely a vector. The state and major economic players must strengthen the communicational and psychological dimensions right from the outset of their crisis management.

This entails breaking with a certain culture of secrecy, opacity and institutional silence, which only serves to fuel anxiety and rumours. Only a stance of committed transparency, channelled information-sharing and coordinated communication at the highest levels of government and among economic decision-makers will restore confidence and thwart the strategic and political paralysis sought by both the attackers and all those who stand to gain from exploiting the effects of successive cyber crises.

2.7 SYSTEMIC LATENT THREATS

Nature of the transformation

Visible cyberattacks, such as ransomware, large-scale data breaches or denial-of-service attacks, form a constant background noise. Some of these may be diversionary tactics designed to mask a threat of an entirely different nature: that of silent infiltration and the long-term pre-positioning of state or quasi-state actors at the heart of our critical technological infrastructure. This media and operational 'smokescreen' diverts the attention of defenders towards managing immediate crises, which may, in certain circumstances, leave the field open for preparatory operations involving espionage or sabotage. The aim is not immediate disruption, but the establishment of a capability for coordinated surveillance or paralysis, which can be activated at the adversary's discretion, particularly in the event of a major geopolitical shift.

Illustration

Mandiant's *M-Trends 2026* report confirms this extreme persistence. Whilst the global median detection time is 14 days, intrusions for espionage purposes have a median presence of 122 days before being identified. Even more striking is the fact that groups such as UNC6201, linked to China, maintain persistent access for over 390 days by targeting the invisible layers of the infrastructure, such as virtualisation servers and edge devices, which are often lacking conventional detection tools.

These attackers prefer to use legitimate administrative tools already present in the systems to blend in with normal business activity and evade malware signatures. ANSSI's *2025 Cyber Threat Landscape* also warns of this risk of pre-positioning within critical infrastructure in the energy, telecoms and transport sectors. For its part, the US cyber strategy of March 2026 acknowledges this reality by now prioritising the active tracking of adversaries even before they launch their offensive phase.

Impact on our collective posture

The possibility of latent pre-positioning of offensive capabilities within our infrastructure constitutes a major strategic risk. Whilst it remains inherently complex to gauge the exact scale of this at a national level, the convergence of analyses by specialist bodies necessitates the integration of this systemic threat into our doctrine and our digital security posture.

ANSSI has already recognised the scale of this challenge by conducting proactive and targeted operations to detect compromises. However, given the high probability of such dormant threats, the challenge now lies in scaling up this detection effort across the entire critical economic infrastructure.

To detect these stealthy presences before they are activated, many large companies are expressing a need for access to weak signals. Consequently, the withholding of information concerning emerging modus operandi – often justified by the need to safeguard long-term investigations or by a strict culture of classification – risks keeping us collectively in the threat's blind spot. Only the seamless and proactive sharing of technical indicators between the state and certain key private sector actors, within a framework of trust yet to be defined, will enable us to strengthen our collective capacity to detect these silent infiltrations and prevent their effects.

2.8 THE EROSION OF IMPLICIT TRUST AND THE INTENSIFICATION OF SYSTEMIC RIVALRIES BETWEEN ALLIES

Nature of the transformation

The paradigm of automatic digital solidarity between historic allies is giving way to a logic of systemic competition, both geopolitical and economic. In this digital space devoid of physical borders, the distinction between mutual protection and the pursuit of strategic interests has become blurred. This shift is not the result of a choice made by France or Europe – which have long worked towards stable multilateral cooperation – but rather a reality they have been forced to accept. Technological dependence is no longer merely an operational risk. It now constitutes a lever of influence and pressure wielded by our own allies. 'Implicit trust', once the cornerstone of Western security, is proving obsolete in the face of national doctrines that openly prioritise technological dominance and aggressive economic intelligence, transforming our networks into arenas of power rivalry.

Illustration

President Trump's cyber strategy for America, published in March 2026, sets out this doctrinal shift with unprecedented clarity and bluntness. This strategy is no longer confined to the mere protection of networks, but embraces the use of unrivalled non-kinetic capabilities to defend national interests. The call to move away from foreign suppliers in order to promote exclusively American technologies, as well as the encouragement given to the private sector to carry out offensive actions to disrupt adversarial networks, mark a profound break with the spirit of traditional transatlantic collaboration.

This weakening of the transatlantic relationship is particularly evident in the field of defence AI. The announcement by the US company Anthropic of its Mythos model, equipped with autonomous capabilities to exploit software vulnerabilities, was accompanied by the creation of the *Glasswing Project*. Although the European Union remains a key strategic partner and its market accounts for just over 30 per cent of the US digital industry's turnover, European manufacturers have been deliberately excluded from this closed military-industrial consortium. This deliberate exclusion, in favour of an exclusively American ecosystem – potentially open to the *Five Eyes* alliance – is a serious blow to the trust-based interoperability between Europe and the United States.

This observation of a breakdown in traditional alliances featured prominently in the debates at the Davos Forum in January 2026. Faced with the erosion of the multilateral leadership exercised by the Western powers under the aegis of the United States, Mark Carney called for a strategic response from moderate and reasonable middle powers, and for the formation of a renewed coalition amongst them. This momentum could be harnessed to refocus the ambitions set out in the Paris Appeal, launched by President Macron on 12 November 2018, in favour of trust and security in the digital sphere. This convergence of views is already taking shape at the Franco-German level. The BSI now recommends, amongst its immediate measures, the creation of a European capacity to assess the most powerful models and direct state access to the best trust tools, so as to no longer be dependent on opaque or exclusively non-European solutions.

Impact on our collective stance

The intensification of these rivalries, including between allies, requires us to reassess the framework of our national security doctrine. We can no longer ignore the dimension of economic competition and

the capture of strategic data by some of our partners. Whilst we collectively regret this weakening of historic solidarity, we cannot, however, simply maintain the *status quo*.

This new reality demands that we build second-generation alliances, based no longer on trust by default, but on reciprocity, technical transparency and respect for the sovereignty of each partner. The resilience of our economic fabric requires that France and Europe transform this constraint into a driving force for developing autonomous defence capabilities, which alone can guarantee our strategic freedom of action.

3 LESSONS AND OUTLOOK

A cross-analysis of the eight shifts documented in this paper clearly demonstrates that major organisations, both public and private, are facing a systemic and irreversible disruption of the threat ecosystem. Cyber security and the protection of information systems have definitively moved beyond the sole scope of technical risk management to take centre stage in an intensified geopolitical and economic competition, where the speed of algorithmic execution supersedes human reaction capabilities.

The events that took place during the first half of 2026 are acting as a powerful catalyst in this regard. The emergence of cutting-edge artificial intelligence in the field of cyber-offence is not merely an incremental evolution in adversaries' modus operandi; it is fundamentally altering the very conditions of digital resilience. The integration of these capabilities into attack chains renders our organisations' traditional defence postures mathematically and operationally obsolete. Consequently, having guaranteed and sustainable access to AI models equipped with advanced reasoning capabilities is no longer merely a matter of technological innovation, but has become an essential prerequisite for maintaining critical digital infrastructure and information systems in operational condition.

To fully grasp the shockwaves of this new reality, we must set aside emotional reactions and take a rational look at the strategic reality now facing Europe, its economic fabric and its institutions. Learning the lessons from this shift means facing up to the mechanisms of our dependence, the sovereign nature of our allies' technological decisions, and the urgency of the pooling arrangements required to avoid permanent decline.

3.1 ANTHROPIC: FROM SHOCK TO AN OBJECTIVE ANALYSIS OF A DECISION

The US administration's sudden decision on 12 June 2026 to suspend global access to Anthropic's Fable 5 and Mythos 5 frontier models – despite their restoration three weeks later – provoked strong reactions within the European digital ecosystem. It revealed that a critical cyber capability could be revoked within the hour. However, to fully grasp the shockwaves caused by this event, it is essential to exercise discernment and critical thinking.

An analysis of this emergency measure taken by the Trump administration shows that it follows a certain logic. When publishing these models, the US publisher had relied on the promise of algorithmic filters designed to ensure their safe use. However, the likely rapid demonstration of the fragility of these safeguards and the model's susceptibility to circumvention techniques has changed the situation. Once an advanced reasoning model exhibits vulnerabilities that allow it to be misused, it exposes the international community to an immediate risk of its use for criminal purposes and the proliferation of offensive capabilities to rival or hostile state powers.

Given this prospect, cutting-edge artificial intelligence can no longer be regarded as a mere commercial digital product. The dual nature of its uses means it can effectively be equated with a cyber-weapon, which may in fact be subject to a regime of control comparable to that applied to certain strategic technologies. Consequently, a government's decision to technically and legally block the global

dissemination of such an arsenal is by no means an anomaly. On the contrary, it constitutes the strictest and most traditional exercise of a sovereign responsibility for national security.

Moreover, this national security imperative must be scrupulously distinguished from the current US political climate. Whilst it is undeniable that Donald Trump's presidency is pursuing a cyber doctrine that explicitly favours unilateralism and the exclusive defence of national interests, it would probably be a mistake to conclude that this specific technological restriction is purely ideological, discriminatory or motivated by speculative interests. Behind the political fog, the issue of technological non-proliferation remains the most rational and comprehensible driving force behind this restriction on access, and certainly not – as some observers have been tempted to interpret this event – as punitive discrimination against European allies.

3.2 CHINA'S RESPONSE AND THE PROLIFERATION OF *OPEN-WEIGHT* FRONTIER MODELS

A few weeks after the Trump administration's decision to suspend global access to Fable 5 and Mythos 5 on 16 July 2026, the Chinese start-up Moonshot AI released Kimi K3, a reasoning model with 2,800 billion parameters, whose performance, according to available independent evaluations, ranks just behind Fable 5 and GPT-5.6 Sol. The full release of its weights – that is, the internal numerical values progressively adjusted during its training phase, which determine how it processes information and generates its responses – is scheduled for 27 July 2026, although no licence agreement has yet been published. Consequently, there is no guarantee that this opening strategy will be sustained for future technological cycles. Beijing could well impose export control measures to block the release of future generations of models.

This rapid progress is not without controversy. Anthropic and OpenAI have publicly accused several Chinese research laboratories, including Moonshot AI, DeepSeek and MiniMax, of using unauthorised distillation techniques to develop their models – a method that involves training a model with lower performance using the outputs produced by a more advanced US model, in order to replicate its performance at lower cost and in less time. This debate, which has not yet been formally resolved, calls for methodological caution. The performance demonstrated by Chinese *open-weight* models cannot be analysed in isolation from the question of their origin, nor can it be exploited without reservation to dismiss the significance of these advances out of hand.

An *open-weight* model, such as Kimi K3, makes the trained parameters of the neural network available, enabling anyone with sufficient computing resources to download it, run it locally and adapt it to their own needs. This openness does not necessarily extend to the training data or the model's entire production chain. *Open-weight* thus differs from *open source*. This distinction has a direct consequence, as a user can remove or weaken the safeguards built into the model without needing to understand how it was trained. All they need to do is access the weights and retrain them locally for purposes that no longer need to comply with the conditions set by the original publisher.

It is precisely this factor that distinguishes the trajectory of Kimi K3 from that of Fable 5. Whereas Anthropic has built its cyber safeguards as a control layer integrated into a service accessible via a proprietary interface and revocable by the publisher, the distribution of the weights places these protections – if they exist in a comparable form – beyond the reach of any subsequent revocation. Once the weights have been downloaded by an indeterminate number of users across the globe, no

public policy decision – whether American, Chinese or European – will be able to regulate their use. This is where the structural limitation of the export control logic outlined in the previous chapter lies, as it presupposes the existence of an identifiable point of control over which public authorities can exert their authority. By its very design, the *open-weight* model no longer presents such a point of control once its initial dissemination is complete.

It would nevertheless be premature to simply equate Kimi K3 with a replica of Mythos made available to everyone. To date, no public and independent assessment has established that Kimi K3 possesses offensive capabilities comparable to those that prompted the national security measure of 12 June. What is, however, established – and which is sufficient to justify heightened vigilance – is the combination of the increasing performance of Chinese models published in *open-weight format* and the structurally greater difficulty in maintaining robust safeguards once the weights have been released.

Moreover, this strategy of openess is probably not driven solely by technological philanthropy. It appears to be a deliberate manoeuvre aimed at undermining the business model and technological hegemony of US laboratories. By offering models—whose performance rivals that of the best proprietary offerings—free of charge or at a much lower cost, Chinese tech firms are making it harder for Anthropic, OpenAI and their peers to recoup colossal training investments through a business model based on paid and controlled access. This strategy of circumventing competition through pricing and mass distribution constitutes, in itself, an instrument of industrial and strategic policy designed to challenge US technological power.

This technological dynamic is now accompanied by a diplomatic offensive. On 16 July 2026, China brought together twenty-nine countries in Shanghai to establish the *World AI Cooperation Organisation*, presented as an alternative to the governance frameworks championed by the United States, the European Union and the G7. Alibaba has also stated, in the wake of Kimi K3, that its Qwen family of models now ranks second in several independent global rankings, just behind Fable 5. This dual advance – both technological and institutional – illustrates that the competition is no longer solely about the performance of the models but also about the ability to set the rules for their use on a global scale. It also lends weight to the hypothesis that Beijing may in future adopt a tougher stance on its own technology exports. Several sources report internal discussions between the Chinese authorities and players such as Alibaba and ByteDance regarding possible restrictions on foreign access to the most advanced Chinese models.

For European businesses and public authorities, this development confirms that mastery of these cutting-edge technologies can no longer be based on the assumption of lasting geographical or legal containment. Companies whose application infrastructure relies on legacy architectures, insufficiently patched or poorly monitored systems, will remain the most vulnerable to the increasing automation of intrusion attempts, regardless of the nationality of the vendor behind the model used to carry them out. It is this shift in responsibility – from technological control over the tool to the resilience of the organisation that uses or is subjected to it – that must now shape our collective thinking in the face of the proliferation of *open-weight* frontier models.

3.3 ARTIFICIAL INTELLIGENCE AS A STRATEGIC ASSET

A joint analysis of the US decision and the proliferation of *open-source* models – particularly those from China – leads to a deeper understanding of the very nature of these disruptive technologies. Cutting-edge artificial intelligence is now establishing itself as a critical component in the protection and security of digital architectures, and no longer merely as a driver of economic performance. State-of-the-art models are gradually entering the realm of strategic technologies, access to which is no longer guaranteed by market forces alone.

Recent events have indeed shown that state-of-the-art models equipped with advanced reasoning capabilities are capable of identifying vulnerabilities on a massive scale, detecting previously unknown exploitation chains, and assisting with remediation operations. This development is now inevitable. Once attackers – whether they are part of organised cybercrime or state-sponsored actors – have access to these same automated targeting tools, there will be no viable option for the European economy, its businesses and public administrations, or indeed the sovereignty of its Member States, without guaranteed and sustainable access to a state-of-the-art reasoning model specialising in digital security.

Furthermore, the dynamics of Chinese *open-weight* models present Europe with a choice whose consequences extend beyond the mere issue of comparative model performance. Assuming that European businesses and public administrations become accustomed to basing their cyber resilience on state-of-the-art *open-weight* models due to a lack of equivalent European and autonomous capabilities, they would expose themselves to a major strategic paradox: replacing their dependence on US laboratories with a dependence on the dissemination policy of Chinese laboratories. There is no guarantee that this policy will endure. Should Beijing one day decide that the national interest requires it to stem this flow—in order to reserve its future algorithmic advances for strictly internal use or for sharing with its allies—through a measure mirroring that taken by Washington on 12 June, Europe, having thus exposed itself to a new form of dependency, would see the upgrading of the foundations of its digital security abruptly halted, without possessing the capacity for autonomous, cutting-edge development to take over within the necessary timeframe. The current openness of China's cutting-edge models should therefore not be interpreted as a lasting guarantee, but as a strategic window which could still be closed at any time by a unilateral decision of a third country.

From now on, much like global geopolitical balances, the digital sphere will become polarised between the 'equipped' powers – foremost among which are the United States and China – and the others, by analogy with the dialectic of nuclear deterrence. The latter will find themselves, through usage restrictions or service disruptions, condemned to a long-term dependence on the former to ensure their own basic digital security. In this context, Europe has no choice but to be a 'capable' power.

However, the parallel with nuclear deterrence encounters a major structural and economic limitation here. Whilst nuclear forces have historically been funded and operated solely at the sovereign level for the defence of the nation, no European state today has the capacity to bear, on its own, the financial and capability burden required to develop a dedicated cyber-defence model. The asymmetry here is twofold. On the one hand, the scale of the investment exceeds the scope of any single national budget. On the other hand, the beneficiaries of this algorithmic shield are not the nation as a whole, but each of the companies and operators of critical infrastructure – both public and private – that underpin the real economy.

This analysis echoes the ideas set out in an essay, published on 8 July 2026 in the geopolitical journal **Le Grand Continent*,* written by Tristan Claret-Trentelivres, Raphaël Doan, Aymeric Roucher and Victor Storchan. In this fascinating and thoroughly researched article, this group of young researchers and senior civil servants adopts precisely this parallel interpretation, drawing parallels with General de Gaulle's nuclear arsenal and deterrence programme. They propose a three-year plan for France and its partners to develop a cutting-edge AI laboratory capable of rivalling the best laboratories in the United States, in order to ensure their strategic independence. The article is explicitly titled 'Operation Prometheus: France can win the AI race. What would be the cost?' The authors estimate that it would be necessary to mobilise around $170 billion from 2027, then over $300 billion in 2029, totalling approximately $700 billion over three years, with computing power accounting for the bulk of the cost. As a proportion of French GDP, this effort would represent between 4.5 and 8 per cent, including 1.5 per cent in public investment per year, making it a commitment on a historic scale. The remainder of the funding would be provided by private investment, but also by private or public capital from other middle powers with a stake in the project, whether to derive direct benefits or to escape their own dependence on the Sino-American duopoly, at a time when these two powers might be tempted to restrict their economic or geopolitical rivals' access to cutting-edge AI systems. The key lesson to be drawn from this essay is the scale of the investment required, which exceeds the capabilities of any single European state and calls for a coalition-based approach.

3.4 THE CONDITIONS FOR A COLLECTIVE EUROPEAN RESPONSE

Given the unprecedented cost of such an investment and the systemic risks of access restrictions, the main focus of concern lies not in Washington, but on our side of the Atlantic. The shockwave of June 2026 starkly highlights the chronic inability of European institutions and Member States to stimulate and finance the emergence of a world-class artificial intelligence industry capable of competing with the American giants. Europe's current dependence is a direct reflection of its own industrial failures and its long-standing passivity as a consumer of technology, and certainly not the result of any injustice inflicted upon it from abroad.

This development now calls for a sudden surge of collective clarity among European stakeholders, even as the political sphere attempts to formulate an initial regulatory and financial response. The Action Plan on Cybersecurity and Artificial Intelligence, presented by the European Commission in July 2026, formally acknowledges this paradigm shift. Building on existing legislative frameworks such as *the AI Act*, the *Cyber Resilience Act* and the NIS2 Directive, this plan aims to equip the Union with medium-term protection mechanisms. The institutional ambition is thus to stimulate the ecosystem through the deployment of *AI Factories* and the mobilisation of equity capital via the *Scale-up Europe Fund*.

However, given the protracted nature of public inertia and the day-to-day operational urgency surrounding resilience, the way forward now appears to lie in a pragmatic reversal of the traditional institutional approach. Since there is no viable short-term alternative without guaranteed access to a specialised digital security framework, the outcome depends on the emergence of a coalition of leading companies capable of creating, on their own, the essential critical mass. The establishment of such an industrial core, based on the pooling of resources and shared governance amongst major economic users, would enable a concrete dynamic to be set in motion, which the public sector could then join and support, particularly with regard to the allocation of the necessary computing capacity.

This grassroots approach would also provide the best means of fostering direct cross-border cooperation, such as the strategic partnerships that are essential at the Franco-German level.

Moreover, the trajectory imposed by the ECB in its capacity as regulator for the DORA regulation confirms this operational urgency. Although the first assessment of major incidents relating to information and communication technologies, published in early June 2026, indicates that cyberattacks account for only 10 per cent of cases within the European financial sector – with the remainder largely attributable to system failures – the authorities are taking a very clear-eyed view of the speed at which the danger zone is drawing nearer. Consequently, just five weeks after this report, the European Systemic Risk Board (ESRB) raised its cyber risk assessment to 'severe'. On the same day, the ECB issued a directive to the 110 largest banks in the eurozone, giving them less than four months to submit an action plan to prepare for cyberattacks driven by artificial intelligence.

This institutional shift stems from an operational reality that has been thoroughly documented by recent assessments of AI models classified as offensive. Artificial intelligence does not necessarily invent entirely new vulnerabilities, but it massively industrialises the exploitation of vulnerabilities that were previously tolerated, such as exposed servers, delayed patches and dormant service accounts. Whereas older-generation models produced a handful of functional exploits, the cutting-edge models of summer 2026 now generate hundreds of them against a single target. What the ECB is now calling for – recognising that the timeframe for mass exploitation is now measured in hours, or even minutes, rather than weeks – is the accelerated and flawless implementation of security fundamentals. The financial regulator thus confirms that, faced with a rigorously compartmentalised and up-to-date system, the asymmetry of the algorithmic threat can and must be contained, heralding the definitive end of the era of tolerance for residual or undocumented vulnerabilities.

4 A CONSTANTLY SHIFTING LANDSCAPE

By its very nature, a briefing note captures reality at a given moment, but the exceptional speed of the changes documented here means that this exercise is destined to become obsolete particularly quickly. In an era of 'machine speed' and sudden geopolitical shifts, the key trends observed in the summer of 2026 are likely to shift in response to the next waves of edge computing deployments or new export control directives. This brief should therefore not be read as a static snapshot of the landscape, but rather within the dynamic context it suggests. The analyses produced today will need to be continually re-examined, given that the blurred boundaries between technological innovation, cybercrime and the interests of major powers create an environment of unprecedented instability.

The immediate future, between now and the end of 2026, will serve as a crash test to gauge the resilience of major organisations in the face of this new era. The first critical milestone will centre on the 31 October 2026 deadline imposed by the ECB. The requirement for Europe's largest banking institutions to submit an action plan to counter AI-driven attacks will act as a powerful litmus test and, we hope, as a catalyst. It will indeed prompt the entire sector to assess, in practical terms, its ability to patch and segment systems at industrial scale, thereby creating de facto new standards for other sectors not subject to DORA. The lessons learnt from this banking standardisation exercise will reveal whether the security market is able to absorb the cognitive asymmetry imposed by offensive models.

At the same time, the coming months will provide an opportunity to assess the viability of the structural responses initiated on both sides of the Atlantic. It will be a matter of observing whether the European Commission's Action Plan manages to move beyond the stage of regulatory intentions to trigger the operational roll-out of its test platforms and funding mechanisms. Similarly, the ability of the European economic fabric to foster private-sector pooling initiatives capable of competing with the Sino-American duopoly will be the true indicator of the expected resurgence.

The observations we have gathered during the first half of 2026 lead to a simple conclusion. Artificial intelligence is becoming an essential component of large organisations' cybersecurity systems. The models with the highest performance are already contributing to their protection. From an activity that was gradually enhanced by technology, cybersecurity has become an algorithmic activity supervised by humans. The conditions governing access to these capabilities are determined as much by industrial decisions as by national security considerations. Artificial intelligence now constitutes a critical part of the cybersecurity infrastructure. Digital resilience therefore also depends on organisations' ability to secure, over the long term, guaranteed access to the models and computing capabilities that make this protection possible.

Cigref is a network of major French companies and public administrations whose mission is to develop its members' ability to integrate and master digital technologies. Through the quality of its thinking and the representativeness of its members, Cigref is a unifying force in the digital society. Cigref was founded in 1970 under the French law of 1901, and does not engage in any profit-making activities.
To achieve its mission, Cigref relies on three core businesses that make it unique.

Membership
Cigref embodies the collective voice of France's leading companies and government agencies on digital issues. Its members share their experience of technology use within working groups, to help identify best practices.

Intelligence
Cigref participates in collective reflection on the economic and societal challenges of information technologies. Founded nearly 50 years ago, Cigref is one of the oldest digital associations in France, and draws its legitimacy from both its history and its mastery of technical subjects, the foundation of skills and know-how that underpin the digital world.

Influence
Cigref promotes and respects the legitimate interests of its member companies. As an independent forum for exchange and production between practitioners and players, it is a benchmark recognized by its entire ecosystem.

www.cigref.fr
21 av. de Messine, 75008 Paris
+33 1 56 59 70 00
cigref@cigref.fr


Cigref
SUCCEED
WITH DIGITAL