

Cybersecurity

Recommendations for cybersecurity

Reinforce software supply-chain cybersecurity

Reinforcing software supply-chain cybersecurity is crucial given the wide impact of attacks spread through the supply chain, which is all the more important given the large number of components in the next computing paradigm (NCP). Develop code and component analysis technologies for cybersecurity that scale up and support trusted orchestrators, services and communications.

Comprehensive safety, security, and performance coupling requires standardized software vulnerability representation. Increased interconnectivity requires new technologies to isolate threats and proactive cyber-risk management. Develop secure software package and service management that balances usability with strong security.

Use AI for cybersecurity

To enhance NCP cybersecurity in a scalable way, develop i) advanced artificial intelligence (AI) models, including large language models (LLMs), for threat detection and ii) autonomous systems for mitigation (e.g. isolating compromised NCP components, patching vulnerabilities, or restoring services). Utilize federated AI for its decentralized, privacy-preserving and scalable models in the NCP massively interconnected context. Rely on EU-based open AI models and datasets to strengthen EU cybersecurity, sovereignty, and competitiveness.

Reinforce the cybersecurity of AI

Secure AI training methodologies and validation procedures, as well as adversarial defences, are needed. LLM prompt injection attacks must be a major concern, addressed by the development of tools to detect and secure against these, and by establishing benchmarks for prompt injection prevention and response. AI security standards should be established by developing certification procedures to guarantee that LLMs and AI systems adhere to stringent security standard, possibly requiring security audits for AI systems. These efforts should rely on EU-based open AI models.



Background

Cybercrime is known to have been increasing dramatically over the last few years, and this trend is expected to continue, as the following figure from Statista shows:




Figure 1: Cybercrime costs are expected to continue dramatically increasing [Statista-CostCybercrime]

The next computing platform (NCP) orchestrates numerous components and services from cyber-physical systems, the internet of things (IoT), clouds, digital twins, etc. This federated, highly connected and dynamic computing continuum thus offers a particularly large attack surface to cyber villains, and can only be expected to suffer from the aforementioned increasing cybercrimes.

Protecting the NCP thus requires strong, scalable analyses to detect and fix vulnerabilities across all its levels, domains, interconnected components, (micro)services, orchestrators, and their communications and interactions. With all these components and services, at production stage, supply-chain cybersecurity and NCP source code become ever more crucial. NCP cyber defences must also encompass detection and mitigation of attacks when in operation.

As in many domains, AI is being used by cyber villains to help them produce and automate cyberattacks. AI has thus become necessary to cope with this increased threat and with the massive complexity the NCP brings, by scaling up and automating cybersecurity tasks at all levels.

AI, especially the booming LLMs used in the NCP context, also faces crucial and often specific cybersecurity issues that must be addressed for its widespread usage to be secured.

Regulatory measures are also necessary for the cybersecurity of the NCP, and societal preparedness must be reinforced for EU security.

By addressing these points, the EU can establish the NCP as a continuum with strong cybersecurity, thereby maintaining confidence and dependability in it, establishing itself as leader in cybersecurity innovation while protecting the EU cybersecurity and sovereignty. To this end, we make the following three main recommendations, followed by additional recommendations.

Reinforce software supply-chain cybersecurity

Software vulnerabilities present a very significant risk to EU. Reports highlight that over 75% of applications contain at least one flaw, nearly 25% of these being classified as high-severity issues, and that even more alarmingly 26% of organizations still face exposure to vulnerabilities exploited by well-known attacks like WannaCry, years after patches have been released [Qualys2024] [comparitech2024]. Identified common vulnerabilities and exposures (CVEs) follow an ever-increasing trend, as shown in the following graph from [CVEdetails.com]:




Figure 2: The number of identified CVEs keep increasing quickly.

Such statistics underscore the importance of continuous, automated, and scalable vulnerability detection and fixing methods and tools.

However, the NCP introduces very significant complexity through the integration of myriad components and services, ranging from a variety of domains like cyber-physical systems,

the internet of things, clouds, digital twins, etc. all of which work together thanks to powerful, AI-based, orchestration mechanisms. Securing this federated, highly connected and dynamic computing continuum requires powerful and scalable analysis technologies able to detect and address vulnerabilities throughout all the levels and domains of this system architecture.

These techniques must address the level of individual components, to detect and identifying vulnerabilities in specific software libraries or application programming interfaces (APIs), especially those that are extensively used throughout the continuum. The open-source nature of these components, which facilitates reuse and modularity at the software engineering stage, is also a facilitator for these analyses, since the source code will be available, enabling source-code level analyses and comparison with the generated binaries.

These, or complementary, techniques must also address the interconnections between components and services, which are a key aspect of the NCP. They must be able to evaluate the potential security vulnerabilities stemming from interactions among components, such as unsecured data transmissions or inadequately established protocols, that could leave opportunities for attacks.

Techniques must also address the NCP orchestration level, assessing and scrutinizing the conduct of these AI-driven orchestration systems to guarantee they do not create security vulnerabilities, such as misconfigurations or unintentional privilege escalations. In addition to this scrutiny, the code used for the orchestrators themselves, although not specifically more vulnerable than any component code, should be considered as particularly critical: it will be a favourite target of attackers, given that by controlling orchestration they could obtain tremendous system-wide effects.

The same is true for agents making decisions in the NCP. Although from a technical point of view they are similar to other agents, from an attacker's point of view decision agents may be more interesting than e.g. sensor agents. However, while stealthier, sensor agent attacks may have strong domino effects as well, which can make them attractive. An attacker would thus likely target either sensor agents or decision agents, based on the specifics of each case – the orchestrator, of course, remaining a prized target.

For these analyses to be effective and usable in practice, it is necessary to improve the way software vulnerabilities are represented and modelled. Indeed, the current diversity in the documentation, cataloguing, and resolution of software vulnerabilities continues to hinder cybersecurity.

Currently, vulnerabilities still often are documented in an informal, human-readable way that is not ideal for automation and tooling, with CVEs being represented in a semi-formal way, despite ongoing improvements [CVE-MITRE][CVE-ORG][CISA2024]. Implementing and advocating standardized formats for vulnerability representation and modelling is thus necessary, not only to mutualize efforts but also to improve the interoperability of cybersecurity analyses and (possibly) mitigation tools for the NCP ecosystem.

This clear definition of vulnerabilities should be augmented to encompass the distinctive attributes of NCP components and their interconnections. This should further facilitate the creation of innovative techniques and tools to model and query interrelations among vulnerabilities, threats, and mitigations. Such standardization of representations will enhance automated vulnerability identification and fixing as well as collaboration among developers, security professionals, and organizations.

In addition to securing the code and components, it is necessary to address the security of the chain that supplies them. This is crucial given the extremely wide impact of software supply-chain attacks (i.e. attacks spread through the software supply chain) and the fact that this is becoming one of the most exploited cyberattack vectors.

High-profile supply chain attacks, like those targeting SolarWinds and Log4j, have clearly shown the dire consequences of vulnerabilities being present in software dependencies and distributions. In 2023, supply-chain cyberattacks surged by 200% compared to 2022, with malevolent actors using critical infrastructures and widespread software libraries to disseminate malware [CISA2023][comparitech2024][Ladisa2023a].

Such cyberattacks generally exploit vulnerabilities in third-party code and libraries, using dependencies to stealthily infect software applications. Thus, dependency management, or a software bill of materials (SBOM) [SBOM-NTIA][Dalia2024], is crucial to monitor the interdependencies and vulnerabilities in supply chains, especially due to the federated, decentralized and dynamic nature of the services provided in the NCP and the components underlying them. These cyberattacks also target tools used to develop software, altering development or building tools to embed malicious code (i.e. malware, backdoors or vulnerabilities) in the software produced, or hijacking package managers, updates, or repositories to spread malicious components at distribution time. This is especially important in the NCP, which requires numerous components from various providers.

Research in secure package and component management systems is thus crucial to alleviating these risks. The security features of these package managers must, among others, include integrity verification, to confirm that each package or component is cryptographically signed and validated prior to being used [Sigstore] and dependency security management, to detect and address vulnerabilities in transitive dependencies, an aspect often neglected yet responsible for over 60% of problems in software projects [comparitech2024]. Above all, these package managers should be devoid of mechanisms that make it possible to execute arbitrary code on the target system, which is far from the current situation [Ladisa2023b].

However, for these secured systems to be successful, i.e. adopted by developers, they must not trade ease of use and developer-friendliness for security. There lies an important challenge: having package managers that are both secure and easily usable.

Reconciling usability with security is a fundamental problem in securing the supply chain that has been poorly addressed so far, developers often prioritizing accessibility to repositories and tools above rigorous security measures. To address this crucial issue, it is necessary to carry out research on secure package managers that: facilitate smooth integration with common, established development workflows; provide or integrate with developer-friendly tools for vulnerability detection and possibly fixing; and provide developer-friendly dashboards to monitor and manage software dependencies.

Furthermore, for maximum impact and effectiveness, research bodies and industry must foster collaborations to establish unified standards and tools for secure package and component management systems. These efforts could build upon existing initiatives such as the SBOM and frameworks for secure software development, like NIST's Secure Software Development Framework (SSDF) [NIST-SSDF].

Use AI for cybersecurity

In order to handle the sheer volume and complexity of components, interconnections and data in the NCP, and reach the appropriate levels of scalability, it will be both crucial and necessary to develop scalable automated analysis methods—leveraging AI and machine learning (ML) where appropriate. Indeed, human-centric methods for detection and response to cyberthreats and cyberattacks have become inadequate. In 2023 for example, global cybersecurity incidents increased by more than 40%, propelled by the extensive use of networked devices and the increase of AI-assisted cyberattacks [Statista-Breaches2024], like ransomware as a service (RaaS)[IBM-RaaS]. Manual methods have thus become insufficient to match the speed and magnitude of these dangers, hence the need for fast, scalable, automated methods and tools.

AI, especially ML, is transforming cybersecurity. It empowers attackers to create cyberattacks more easily, even for people with a lower level of technicity, hence spreading the fire. However, it also offers many opportunities for cyber defenders.

Indeed, AI-driven algorithms can scrutinize extensive datasets to find deviations from standard behaviour, and identify anomalies very quickly, drastically decreasing threat detection time from months to seconds. AI technologies used in security orchestration, automation, and response (SOAR) systems can triage, analyse, and mitigate threats autonomously, providing automate incident response. They can also even help anticipate and mitigate threats, since predictive analytics can discern nascent assault patterns, enabling and facilitating pre-emptive defence strategies.

Overall, AI-driven automation has the potential to enhance scalability and continuously adjust to emerging threats in near real time, ensuring stronger security as the quantity of devices, components and services increases rapidly in the NCP context.

However, although automation presents significant potential, especially AI-driven automation, its implementation faces several challenges. First, such tools must integrate smoothly across many platforms, services, and physical components within the NCP, which may require significant engineering efforts. In addition, automated systems must combine efficient monitoring with solid privacy rules, such as GDPR. Finally, AI and ML-based systems must be protected against involuntary biases and model vulnerabilities, and against adversarial AI attacks that could compromise their performance, which implies research to investigate into more robust models.

As a consequence, to enhance and scale up automated cybersecurity within the NCP, research must be encouraged and tools developed on i) advanced threat-detection AI models, based on deep learning, graph-based analysis, natural language processing (NLP) and large language models (LLMs) to improve detection and monitoring capabilities, and ii) autonomous threat mitigation systems that can perform automatic actions, such as isolating compromised components of the NCP, applying patches to vulnerabilities, restoring services, etc.

In the continuum of the NCP, federated AI systems should be investigated, as their decentralized AI models can help function across distributed, massively interconnected components and services, in an edge and cloud context, while preserving data privacy and scalability. To this end, the use of EU-based open AI models and datasets such as HuggingFace [HuggingFace] and Mistral [MistralAI] should be favoured, as this can help the EU boost its cybersecurity while preserving its sovereignty and reinforcing its competitiveness.

Reinforced cybersecurity of AI

The use of AI systems, especially LLM-based systems, has become in a few years extremely widespread in almost all domains and applications of computing, hence all across the NCP. The (cyber)security of such systems is thus crucial, yet their rapid adoption brings unique challenges that require urgent attention. Indeed, deployed LLMs are currently susceptible to numerous security vulnerabilities. Malevolent actors can exploit prompts [Pasquini2024, Liu2024] to compel LLMs to produce detrimental or unauthorized content.

Existing protective measures seem to be rather an external layer of LLMs, since relatively simple tricks have been shown to bypass these security measures. Challenges thus exist in completely mitigating these vulnerabilities, which ideally should be done within the LLM's internal behaviour, not at its periphery.

Attackers can also compromise data [Monkam2024] used to train AI models, resulting in skewed outputs and weakened defences. LLMs generating inaccurate or false information, either by mistake or by having been skewed to do so [Wu2024], can then be leveraged to

disseminate misinformation or influence choices, such as elections, all across the NCP, with low cost and high spread. On several occasions, LLMs have been shown to leak sensitive corporate information [Raz2024]. LLMs are also used to help developers code, hence generate source code; the latter can however contain cyber vulnerabilities.

The EU should thus invest in secure AI research focused on secure training methodologies and validation procedures, as well as adversarial defences. LLM prompt injection attacks must be a major concern, addressed by research on detecting and securing against these and establishing benchmarks for prompt injection prevention and response, e.g. in the spirit of the CyberSecEval benchmarks [CyberSecEval13]. Benchmarks and AI security standards should be established by developing certification procedures to guarantee that LLMs and AI systems adhere to stringent security standard, possibly requiring security audits for AI systems. These efforts should rely on EU-based open AI models (see [HuggingFace] [MistralAI]).

Additional recommendations

In addition to the above three main, critical recommendations, additional relevant recommendations can also be made as follows.

Authentication, intrusion and attack detection in massively interconnected systems

It is necessary to support research and tools for intrusion and attack detection in systems with massively interconnected components and services, including authentication mechanisms that scale up within the NCP.

Indeed, the NCP offers a large attack surface, due to its numerous and massively interconnected components and services. Efficient and effective intrusion and attack detection is thus necessary but faces distinct issues from conventional cybersecurity tools.

The volume and velocity of data generated by the continual exchange across interconnected (micro-)services and components produces a tremendous traffic volume, making it difficult to distinguish harmful activities from normal ones. The myriad of NCP components, while facilitating compartmentalization and isolation, also provides attackers with opportunities for concealment, allowing them to use strategies involving long-term infiltration and lying dormant to avoid detection, moving only within the ecosystem of components and services when attacking their real target, a technique which is called "lateral movement". Furthermore, the heterogeneity and dynamicity of the NCP, characterized by the dynamic orchestration of resources, require adaptive and context-sensitive detection techniques and tools.

Research must thus be encouraged to develop automated, big-data-capable intrusion-detection systems (IDS), capable of monitoring and analysing extensive data sets in real-time, to detect attacks and intrusions early, as they develop, not after. The goal is to identify and obstruct malevolent actors in real time, which is very far from the industry average time-to-detection of over 200 days [IBM2021], building on already-reported time savings of 108 days provided by AI-powered tools [IBM2023].

Indeed, AI and ML can be keystones to this end, helping for example with anomaly detection and pattern recognition (in logs or execution traces) associated with cyberattacks. Today, IDS already employ ML algorithms to attain detection rates over 90% in specific circumstances [CISA2024]. Given the sheer amount of data to analyse, it is crucial for usability that false positive rates are kept extremely low, while effectiveness commands that false negative rates remain low as well, which is always a challenge and one that research must address upfront. Federated learning, due to its distributed nature, should be investigated for its scalability.

To prevent intrusion, one specific aspect to address in the security of the NCP is secured authentication solutions that must also scale efficiently to address the NCP's dynamic and distributed characteristics. Authentication solutions exist, but these must evolve to preserve both security and ease of use, minimizing friction for NCP users.

To this end, new multi-factor authentication (MFA) technologies incorporate biometrics, contextual awareness, and behavioural analytics to deliver strong and user-friendly authentication solutions. Zero-trust architectures, whose principles mandate continual authentication and authorization of every entity irrespective of its location, seem essential for the NCP, since they can guarantee secure interactions even in extremely dynamic settings. Blockchain-based decentralized identifiers (DIDs) could also facilitate scalable and secure authentication in an NCP context by removing dependence on centralized authorities. Furthermore, integrating anomaly-based intrusion detection with adaptive authentication is very important for the NCP, because it allows access controls to be dynamically modified in response to identified threats, thus significantly improving security.

In a nutshell, research and industry must be encouraged to develop, for dynamic and highly interconnected contexts such as the NCP, real-time, scalable, and adaptable technologies that can identify both known and undiscovered threats in extensive systems, as well as decentralized, context-sensitive authentication systems.

Secure critical infrastructure

Critical EU infrastructure, encompassing utilities, healthcare facilities, and transportation systems, constitutes the foundation of contemporary society. The interruption of the services such infrastructure provides can lead to significant societal and economic repercussions. As the NCP consolidates these systems into a cohesive, highly interconnected continuum, enhancing their cybersecurity is literally vital.

Critical infrastructure has been subjected to cyberattacks, ransomware, and state-sponsored cyber assaults for a long time, predominantly affecting the energy, healthcare, and water-management sectors [Zendra2023]. These cyberattacks generally exploit vulnerabilities in legacy systems, and interconnectivity to disrupt critical services or gain influence in geopolitical conflicts.

The EU must continue implementing regulatory frameworks that require the protection of essential infrastructure. This encompasses fundamental needs for cybersecurity measures, regular evaluations, and criteria for secure-by-design elements and services. The EU NIS2 Directive [NIS2-EU], effective in 2024, and the EU Cyber Resilience Act (CRA) [CRA-wiki] [CRA-EU] adopted in October 2024, are steps in the right direction. However, both require solid implementation measures to ensure compliance and effectiveness.

Mass cyberattacks can incapacitate centralized systems. To address this, the EU should make it mandatory that essential infrastructure integrate autonomous "archipelago" systems – i.e. self-sufficient components that can function independently during disruptions. One example is smart grids, which should incorporate localized microgrids capable of autonomously maintaining electricity delivery in their area during an attack.

Thanks to its nature of interconnected yet separate components and services, the NCP is in many ways suitable as a supporting infrastructure for these archipelagos, thanks to which systems should be compartmentalized to avert cascade failures. In a nutshell, EU regulations must mandate the design and implementation of infrastructure elements that can function independently.

In addition to regulation, emergency preparedness plans and drills should be conducted at EU level. Cyberattack simulations and coordinated exercises are crucial to prepare organizations for degraded-mode operations. EU-wide exercises such as Cyber Europe [CyberEU-ENISA] offer opportunity to evaluate resilience and response tactics across national boundaries. Such programmes should be augmented and adapted to address the

particular issues presented by the massively interconnected components and services of the NCP.

Liability

The involvement of software and hardware suppliers is crucial in cybersecurity. Ensuring their accountability for security vulnerabilities encourages higher standards in product security [Zendra2023]. The EU has already moved in that direction with the Cyber Resilience Act (CRA) and the NIS2 Directive, which stress the need for inherently secure technologies, and make providers of information and communication technology (ICT) systems accountable for cybersecurity deficiencies, including insufficient safeguards or unresolved known vulnerabilities.

These efforts must be continued and their proper, concrete and effective implementation ensured.

References

- CISA2024: 2023 Top Routinely Exploited Vulnerabilities. 12 Nov. 2024. <https://www.cisa.gov/news-events/cybersecurity-advisories/aa24-317a>
- comparitech2024: Cybersecurity vulnerability (CVE) statistics and facts (2019-2024). <https://www.comparitech.com/blog/information-security/cybersecurity-vulnerability-statistics/>
- CRA-EU: EU Cyber Resilience Act. <https://digital-strategy.ec.europa.eu/en/policies/cyber-resilience-act>
- CRA-wiki: Cyber Resilience Act. Wikipedia. https://en.wikipedia.org/wiki/Cyber_Resilience_Act
- CVE-MITRE: <https://cve.mitre.org/>
- CVE-ORG: <https://www.cve.org/>
- CVEdetails.com: <https://www.cvedetails.com/browse-by-date.php>
- CyberEU-ENISA: Cyber Europe - Leading the way in cybersecurity preparedness. <https://www.enisa.europa.eu/topics/training-and-exercises/cyber-exercises/cyber-europe-programme>
- CyberSecEval3: <https://meta-llama.github.io/PurpleLlama/>
- Dalia2024: G. Dalia, C. A. Visaggio, A. D. Sorbo, and G. Canfora. SBOM ouverture: What we need and what we have. ARES '24: Proceedings of the 19th International Conference on Availability, Reliability and Security. Article No.: 116, pp. 1-9. <https://doi.org/10.1145/3664476.3669975>
- HuggingFace: <https://huggingface.co/>
- IBM-RaaS: What is ransomware as a service (RaaS) ? Jim Holdsworth, Matthew Kosinski. 5 September 2024. <https://www.ibm.com/topics/ransomware-as-a-service>
- IBM2021: IBM Report: Cost of a Data Breach Hits Record High During Pandemic. Jul 28, 2021. <https://newsroom.ibm.com/2021-07-28-IBM-Report-Cost-of-a-Data-Breach-Hits-Record-High-During-Pandemic>
- IBM2023: What's new in the 2023 Cost of a Data Breach report? Sarah Villavicencio. July 25, 2023. <https://community.ibm.com/community/user/security/blogs/sarah-dudley/2023/07/25/costofatabreach2023>
- Ladisa2023a: Taxonomy of Attacks on Open-Source Software Supply Chains. Piergiorgio Ladisa, Henrik Plate, Matias Martinez, Olivier Barais. 2023 IEEE Symposium on Security and Privacy (SP), San Francisco, CA, US, 2023 pp. 1509-1526. <https://arxiv.org/abs/2204.04008>
- Ladisa2023b: The Hitchhiker's Guide to Malicious Third-Party Dependencies. Piergiorgio Ladisa, Merve Sahin, Serena Elisa Ponta, Marco Rosa, Matias Martinez, Olivier Barais. 2023 Workshop on Software Supply Chain Offensive Research and Ecosystem Defenses, pp. 64-74. <https://dl.acm.org/doi/pdf/10.1145/3605770.3625212>

Liu2024: Formalizing and Benchmarking Prompt Injection Attacks and Defenses. Yupeí Liu, Yuqi Jia, Runpeng Geng, Jinyuan Jia and Neil Zhenqiang Gong. 33rd USENIX Security Symposium (USENIX Security 24).

MistralAI: <https://mistral.ai/>

Monkam2024: A topological data analysis approach for detecting data poisoning attacks against machine learning based network intrusion detection systems. Galamo Monkam, Michael J. De Lucia, Nathaniel D. Bastian. Computers & Security, vol. 144, 09/2024, Elsevier.

NIS2-EU: Directive on measures for a high common level of cybersecurity across the Union (NIS2 Directive). <https://digital-strategy.ec.europa.eu/en/policies/nis2-directive>

NIST-SSDF: Secure Software Development Framework (SSDF). <https://csrc.nist.gov/Projects/ssdf>

Pasquini2024: Neural Exec: Learning (and Learning from) Execution Triggers for Prompt Injection Attacks. Dario Pasquini, Martin Strohmeier, Carmela Troncoso. 2024 Workshop on Artificial Intelligence and Security (AISec '24).

Qualys2024: 2023 Threat Landscape Year in Review: If Everything Is Critical, Nothing Is. Saeed Abbasi. January 4, 2024. <https://blog.qualys.com/vulnerabilities-threat-research/2023/12/19/2023-threat-landscape-year-in-review-part-one>

Raz2024: The good, the bad, and the ugly: Microsoft Copilot. hack.lu Security Conference 2024. <https://archive.hack.lu/hack-lu-2024/talk/NNFQ3G/>

SBOM-NTIA: <https://www.ntia.gov/page/software-bill-materials>

Sigstore: <https://www.sigstore.dev/>

Statista-Breaches2024: Data breaches worldwide - statistics & facts. Statista. <https://www.statista.com/topics/11610/data-breaches-worldwide/#editorsPicks>

Statista-CostCybercrime: <https://www.statista.com/chart/28878/expected-cost-of-cybercrime-until-2027/>

Wu2024: Fake News in Sheep's Clothing: Robust Fake News Detection Against LLM-Empowered Style Attacks. Jiaying Wu, Jiafeng Guo, Bryan Hooi. 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD '24).

Zendra2023: O. Zendra and B.Coppens. From cybercrime to cyberwarfare, nobody can overlook cybersecurity any more. In M. Duranton et al., editors, HiPEAC Vision 2023, pages 130-144, Jan 2023. DOI: 10.5281/zenodo.7461910