Practicality of Reference Architecture Applied to Cyber-Physical Systems for Cyber-Resilience

TNO Public ⟩ TNO 2025 R13183

24 December 2025

ICT, Strategy & Policy
www.tno.nl
+31 88 866 50 00
info@tno.nl

TNO 2025 R13183 – 24 December 2025

Practicality of Reference Architecture Applied to Cyber-Physical Systems for Cyber-Resilience

Author(s)	Sezen Acur, Bas van der Leeuw
Classification report	TNO Public
Title	TNO Public
Report text	TNO Public
Number of pages	15 (excl. front and back cover)
Number of appendices	0
Sponsor	NWO
Project name	INTERSCT
Project number	060.37563

All rights reserved
No part of this publication may be reproduced and/or published by print, photoprint, microfilm or any other means without the previous written consent of TNO.

© 2026 TNO

Abstract: In today's digital landscape, organizations are increasingly seeking to introduce cybersecurity measures within their systems. The recent cyber-attacks fuelled by the geopolitical tensions indicated the true importance of being cyber-resilient. However, many organizations face significant challenges in starting this process. The complexity of cybersecurity frameworks, high costs of implementation, and uncertainty about where to begin can leave organizations feeling overwhelmed. This paper provides in depth concerns of industry from round table sessions and how creating reference architectures may alleviate potential problems.

This paper is intended for an industrial audience interested in how to apply the cyber-resilient systems engineering methodology, created by TNO. Its guidance, methods, and recommendations assume that readers already have a clear understanding of their own systems and processes. In particular, the effective use of this methodology requires the presence of a reference architecture or comparable high-level representation of the system.

Contents

1 Introduction 4
1.1 Findings from the round table 4
1.1.1 Implementing cyber-related standards and regulations 4
1.1.2 Incident management, vulnerabilities, and ownership after deployment 5
1.1.3 Finding the balance between qualities and requirements 5
1.1.4 Organisational struggles with roles and readiness 6
1.2 Reference Architecture 6
1.2.1 Reference Architecture: An e-bike example 7
1.3 Methodology in Summary 12
1.4 Conclusion 13
2 References 15

Table of Figures

Figure 1: Microsoft Cybersecurity Reference Architecture (MCRA) [6] 7
Figure 2: Core e-bike elements per company in system definition diagram 8
Figure 3: E-bike operation goals and capabilities 9
Figure 4: E-bike safety and security capabilities 10
Figure 5: Sequence diagram of safety issue 11
Figure 6: Cyber-resilient system methodology 13

1 Introduction

Within the WP2 of the INTERSCT project, the synergy between systems engineering and cybersecurity has been explored [1]. In our 2023 paper, SOS! Ensuring safety and security in an expanding system of systems landscape, we explored the elements of uncertainty, system of systems thinking within the changing landscape, and developed a methodology [2]. After publishing this report, we have conducted two sets of roundtables in which we have had participants from the Dutch industry, universities as well as government entities. From these round tables, we received several questions from the attendees. Overall, the questions were about the adaptation of standards, how cyber-resilience and cybersecurity are managed and how the cyber-physical systems can be secured especially if there are third-party applications involved.

In this paper, we are going to explore further on the findings from the round tables and how the concerns from the industry partners regarding cyber-resilience can be represented in reference architectures.

1.1 Findings from the round table

In the round table sessions, we gathered information from the participants and listened to their pitches on how they manage security and what challenges they face when they want to implement cyber-resilience measures.

The top concern that resonated mostly with industrial participants is that the customers are worried that if the security measures are added for cyber-resilience, then there is a potential for disruptions in the system. Other concerns were that the customers would like no downtime unless intended for maintenance purposes and some customers do not believe that they will be threatened with potential cyber-attacks and therefore they do not see why they need cyber-resilience measures from commercial aspect.

We collected differences and categorized familiar topics to the following categories shown below:

1.1.1 Implementing cyber-related standards and regulations

With respect to the existing standards and regulations, the participants have a top-level awareness. The customer awareness and importance of cyber measures on their products widely vary from being keen on applying measures to thinking that overly secured products may lead to competitive disadvantage. Some examples may be the cost of the product with security included and another is the ease of use. Extra login procedures, tokens or two-factor authentications hamper ease of use.

Security is considered as a quality and implied in the standards in such a way. One question that came from the participant was “how can security be phrased in the business language?”

Certain products have ISO standards and customers also prefer ISO compliance; however, customers may not recognize security related standards and regulations, and they may not follow why there is a need to tailor or adapt to such standards.

1.1.2 Incident management, vulnerabilities, and ownership after deployment

Participants that deploy products to the customers notice that they are not sure how to keep up with increasing number of vulnerabilities. They would like to understand the impact on product architecture where SW and HW updates are implemented. What if there are vulnerabilities and how they can be mitigated? That also brings up a point on when to deploy effective automatic updates, so the business continuity is not disrupted. For instance, there could be an update incoming while there is a presentation or experiment ongoing. This can also change the user interface due to the updates (e.g., Apple IOS updates).

Another aspect is on design, operations, and testing with respect to third-party applications. There is a growing uncertainty on types and levels of security risks that third parties pose. How to manage third party applications in a security posture? What kind of service-level agreements (SLAs) are needed to resolve this issue?

1.1.3 Finding the balance between qualities and requirements

The participants noticed that security and resilience do not go hand in hand in their experience. Resilience affects service and maintenance and in some cases the personnel as well. Security is mostly perceived as a quality for product while resilience affects products and projects over the whole life cycle. We may make it secure, but we may hinder resilience once we secure the product. For instance, denial of service (DoS). The system and data might be secure but perhaps the system becomes overly strict or consumes resource usage which can overload the system or slow down the operation. This becomes a concern especially after the deployment of a product. Maintenance, service, patching and updates in the field become difficult to perform. Given these concerns, one question that came to mind is “where is the trade-off between security and system resilience? “

1.1.4 Organisational struggles with roles and readiness

Moving companies towards considering security into the business aspect rather than a mere quality is difficult. Early considerations of embedding security is not established among industrial entities. Plans regarding how to service after deployment is not often included in the contracts or during the project phase.

Raising security awareness in the organizations and awareness of 3rd party and suppliers in the product are not entirely included in the company trainings. Checklists and assessments are some means and not entirely complete. Methods of mitigation however, can change globally. Certain suppliers can be outside of the EU and they may not address EU regulations that are critical.

> Industrial participants struggle to manage increasing vulnerabilities, assess impacts on product architecture, and deploy updates without disrupting operations, especially when third-party components introduce unclear risks and SLA needs. Security is often treated as a product quality while resilience spans the full lifecycle, creating tensions in post-deployment maintenance, patching, and service. Organisational readiness for security is rarely embedded early in projects, limited training on supplier and third-party risks.

The results from the round table discussions confirm a pattern, that is, business owners like product suppliers, service providers and asset owners tend to focus on the solutions they provide, or what they believe to be the best solutions. The benefit of creating a reference architecture is to help defining the problem and recognise a set of solutions from the architecture. The following section describes what a reference architecture is in systems engineering and in the cyber-security context and provides an example of reference architecture of an e-bike.

1.2 Reference Architecture

In systems engineering, a reference architecture (RA) is reusable, standardized template that captures business, application, and technical layers (including their inter-relations) in an architecture. The focus here is on how the business goals and strategy feed into the architectural decisions and how that can reflect the business-driven reasoning. The major benefits of creating and using an RA not only helps with the alignment of business and technology or improved decision making but also cost reduction and interoperability [3].

In cybersecurity, a reference architecture is an architecture that becomes a foundation, a blueprint for building and/or modernizing a security strategy of an organization, creating standardized framework for components or best practices to protect systems and relevant data. The Amazon Web Services, Microsoft, or Department of Defence (DoD) have their own templates adapting to zero trust principles which is “never trust, always verify” [4] [5] [6] [7].

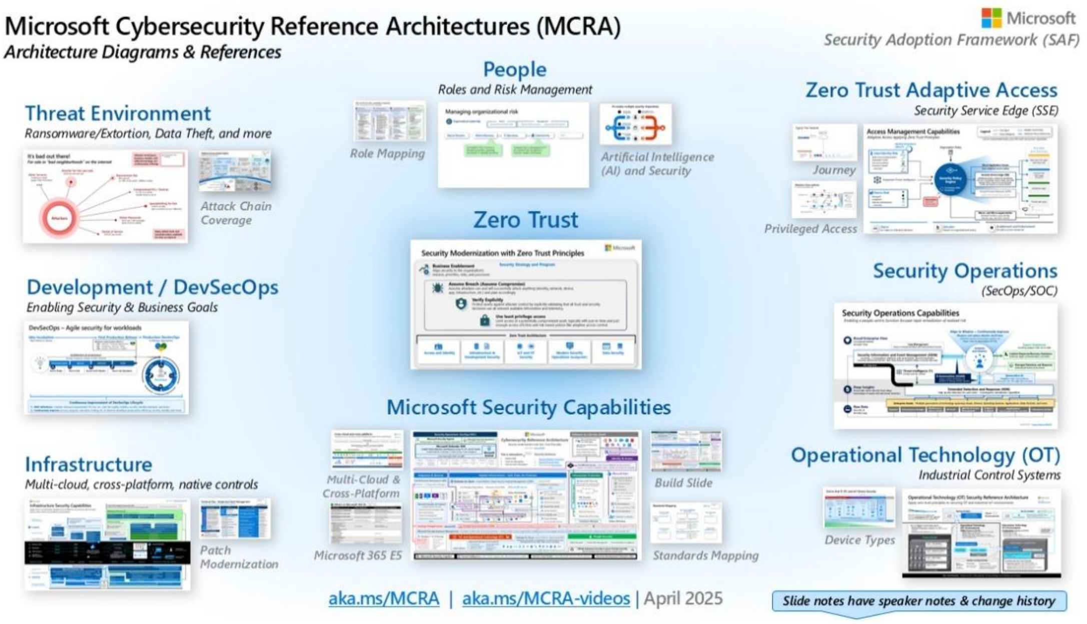
Microsoft Cybersecurity Reference Architectures (MCRA)
Architecture Diagrams & References
Microsoft
Security Adoption Framework (SAF)
People
Roles and Risk Management
Role Mapping
Artificial Intelligence (AI) and Security
Threat Environment
Ransomware/Extortion, Data Theft, and more
Attack Chain Coverage
Zero Trust Adaptive Access
Security Service Edge (SSE)
Journey
Privileged Access
Zero Trust
Development / DevSecOps
Enabling Security & Business Goals
Security Operations
(SecOps/SOC)
Microsoft Security Capabilities
Infrastructure
Multi-cloud, cross-platform, native controls
Patch Modernization
Multi-Cloud & Cross-Platform
Microsoft 365 E5
Build Slide
Standards Mapping
Operational Technology (OT)
Industrial Control Systems
Device Types
aka.ms/MCRA | aka.ms/MCRA-videos | April 2025
Slide notes have speaker notes & change history

Figure 1: Microsoft Cybersecurity Reference Architecture (MCRA) [6]

In both fields, by definition, there are overlaps about what a reference architecture is. If complex cyber-physical systems are being built and if cybersecurity is treated as an integral part of the systems engineering process, these reference architectures (from systems engineering and cybersecurity) shall meet and merge in order to optimize the reusability and scalability aspects of the systems with respect to performance and cyber-resilience.

Merging a system reference architecture and cybersecurity reference architecture (creating a hybrid reference architecture) can result in a synergic library with best practices, frameworks, and protocols. These can be tailored to organizations regardless of size, industry, or risk profile. The organizations have the means for consistency and reliability by following tried and tested architectures, avoiding reinventing the wheel and risks associated with ad-hoc security solutions [8].

Referring to the concerns of the industrial participants, with a hybrid reference architecture, they can identify what they value in their businesses, how they can protect their businesses while interacting with third party suppliers. Based upon that information, they can create a breakdown of their system and determine where security plays a role as well as what roles are needed to mitigate risks associated to security. Next section provides an example of reference architecture for an e-bike modelled in CATIA Magic.

1.2.1 Reference Architecture: An e-bike example

In 2023, a whitepaper, *Reference Architecture in Relation to Business Reasoning*, was published on how reference architecture can capture the business strategies in architectural form [3]. For the sake of demonstrating how a reference architecture can be created, a couple of business strategies for an e-bike were introduced:

Company A: A business intended to create an e-bike for off-road adventure seekers, in all terrain. A sturdy bike frame is core component for this company.

Company B: This business targeted the commuter group, looking for an easy method of transportation. Range of the bike plays a role in the decision-making process of the customers; therefore, battery and motor are key components for Company B.

Company C: The business strategy for this company was in the users who combine biking and workout experience. High quality data of the sensors, providing telemetry data of the biker (heart rate, calories burned) and the trip information (distance, speed) add value to the competitive edge of the market.

In this example, there are two system contexts; one in E-bike Development and Realisation, and another in E-bike Deployment and Use. To provide most value in the market, the companies must be able to buy parts, assemble parts, develop frame set & wheels, and integrate sub-systems in the development and realisation context. Considering this aspect, we presume that there is not enough demand for company C so the company C changed their business strategy to developing a helmet that can provide a rich telemetry data and is suitable for a variety of fitness apps. This is depicted in the **Error! Reference source not found.** below:

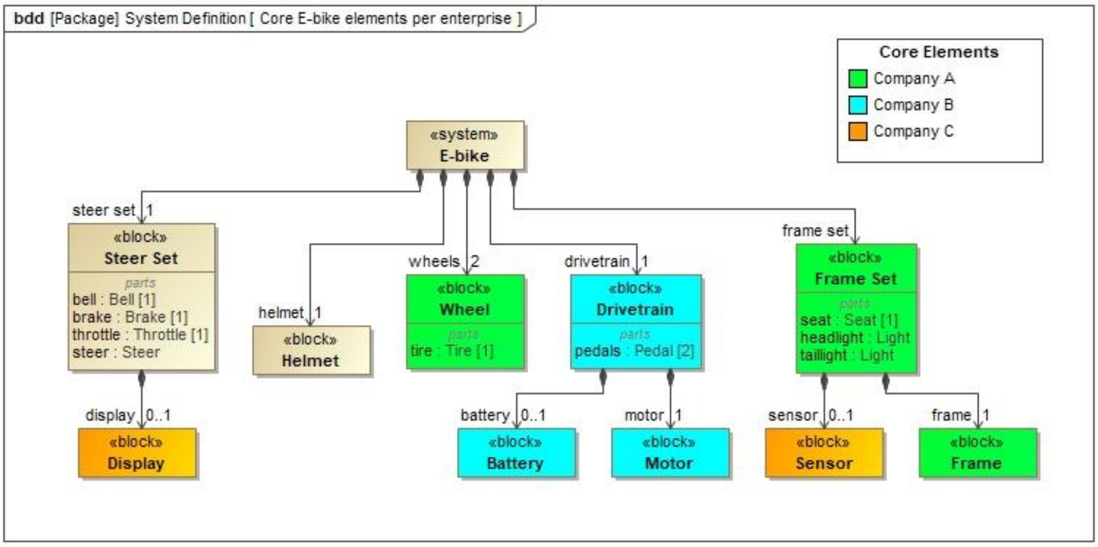

Figure 2: Core e-bike elements per company in system definition diagram

In the other context, e-bike deployment and use, shown in **Error! Reference source not found.**, we translate the company strategies into goals and capabilities. E-bike is considered as the system for all the businesses and the users are defined as explorer, commuter and exerciser respectively to the strategies. The e-bike must provide the capability to ride cycle multiple distances as it is the main capability; going off-road and dodge any obstacle are added capabilities.

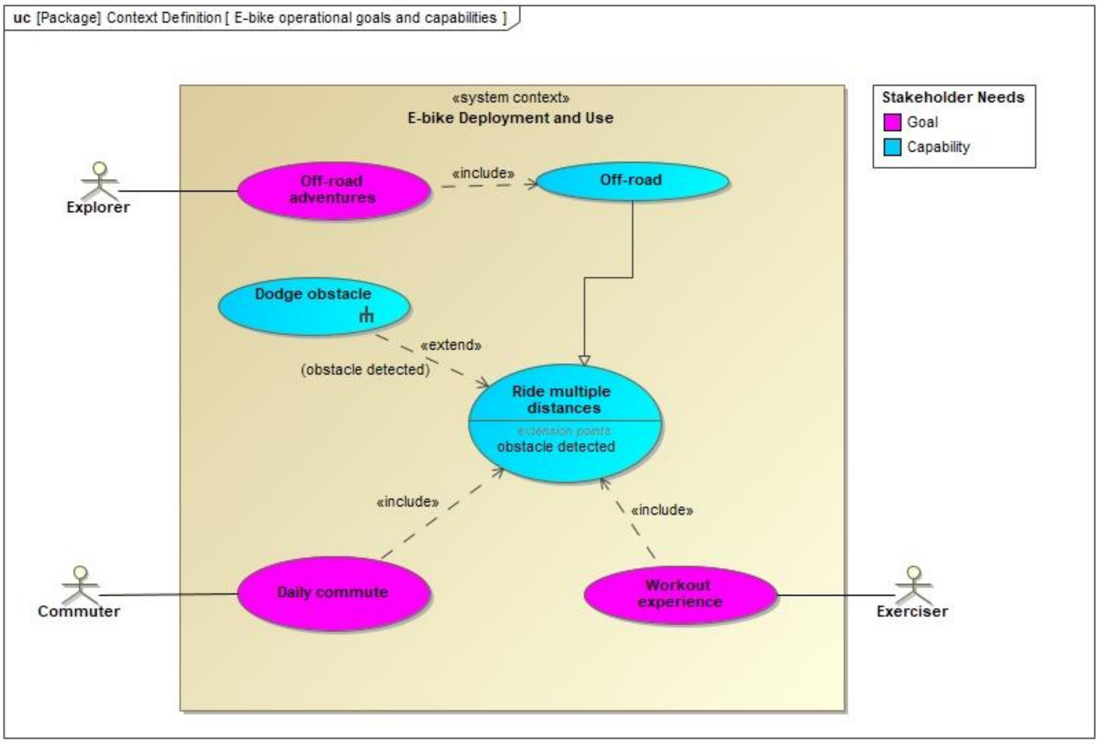

Figure 3: E-bike operation goals and capabilities

In **Error! Reference source not found.**, we are looking at what capabilities and goals the e-bike must achieve to ensure safety and security of the user. From security perspective, the e-bike must be protected from theft and to provide that capability, an anti-theft system must be in place. From the safety perspective, while the user is biking, he or she may see someone on the bike-lane and would like to warn the public road user that he/she is coming. A bell may be a simpler option for such a capability. If it does not work and the user cannot use his/her voice to shout, then the environment (including the user) should be protected from the e-bike. The user is part of the environment while the e-bike remains to be the system itself. In this case, there should be a capability to protect the user’s head from injury if the e-bike diverts away from the path to avoid hitting other people or objects in its environment.

Looking at the diagram in detail, another aspect to consider is what if there was a privacy violation or data breach? If somehow the e-bike has been hampered by the adversaries and they gain access then, can the security breach become a safety concern? If there is a capability add-on for sending messages such as theft warning or that smart helmet is compromised and the components associated with that capability are hampered, then we can see how the security issue becomes a safety issue in this case.

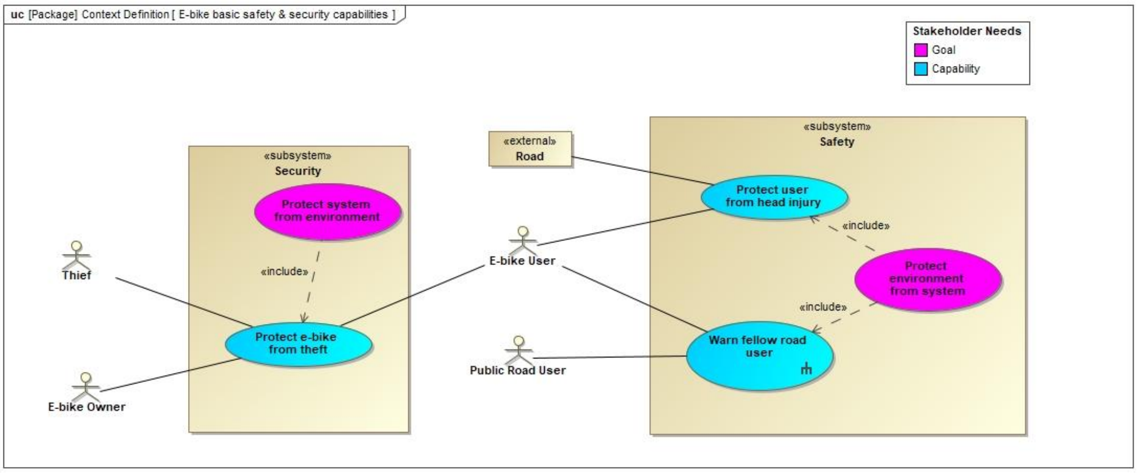
uc [Package] Context Definition [E-bike basic safety & security capabilities]
Stakeholder Needs
🟪 Goal
🟦 Capability
«subsystem»
Security
Protect system from environment
«include»
Protect e-bike from theft
Thief
E-bike Owner
«external»
Road
E-bike User
Public Road User
«subsystem»
Safety
Protect user from head injury
«include»
Protect environment from system
«include»
Warn fellow road user

Figure 4: E-bike safety and security capabilities

A good method to understand how the system operates and what components are needed to fulfil the safety and security measures is creating a sequence diagram. As seen in **Error! Reference source not found.**, this diagram shows the messages going across the actors and the system, referencing the components that play a role in the sequence process. In this sequence diagram, exchanges can continue up until the scenario has been realized. For instance, here is a scenario of how this sequence diagram is carried out. It is a busy city and early in the morning. There are tourists on your e-bike path, and you are going fast on your e-bike. You notice that there is someone walking on the e-bike path, so you ring the bell to warn. You notice the person does not hear the bell, so you yell at him/her to move. The person still does not move. So, you swerve your bike away from the person and your bike heads towards a ditch. You rely on your helmet for your safety as you fall.

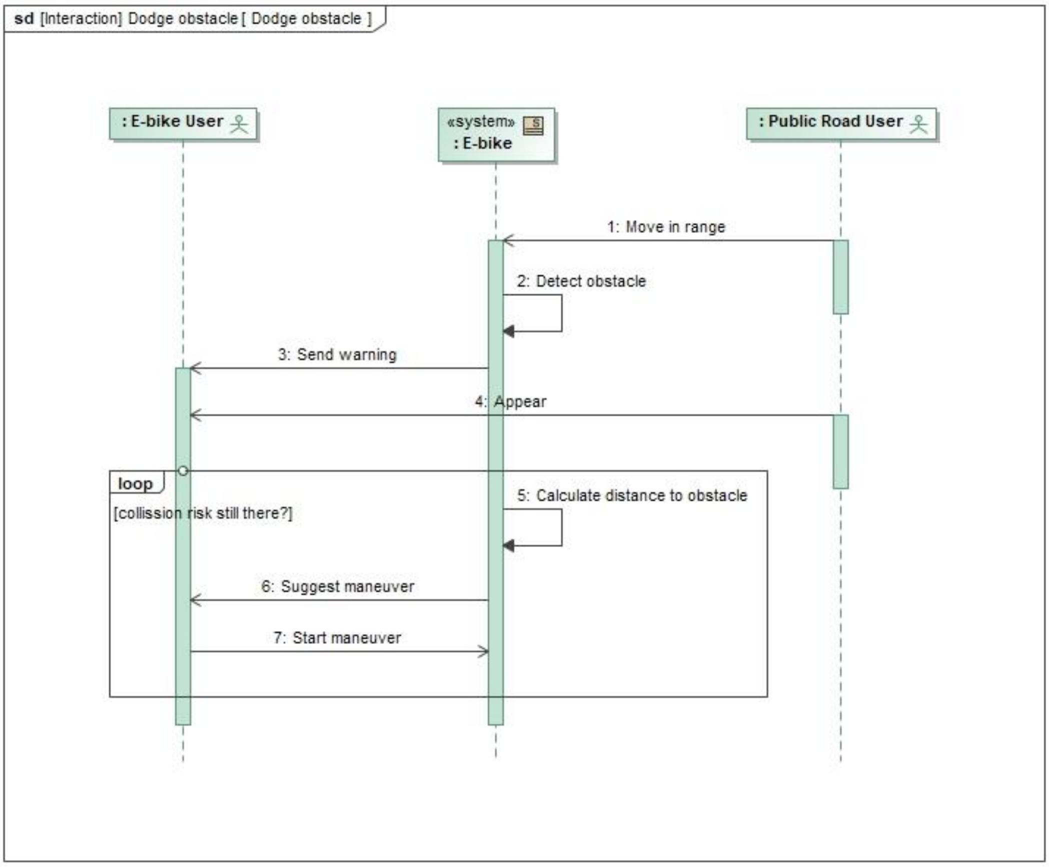

Figure 5: Sequence diagram of safety issue

The e-bike example shows why the reference architecture matters; strategy must drive architecture, otherwise there is a chance to lose the competitive edge in the market. The e-bike companies have different target customers, but if their system architecture doesn't reflect those differences, they risk building the wrong product. The example shows how architectural decisions choose different components; they either support or undermine the business goals.

We also noticed that Company C changed its business strategy from a telemetric bike to a smart helmet. The architecture changes as the business strategy changes. The reference architecture helps show exactly what must change in system components and capabilities to support the new strategy.

Reference architecture clarifies how the system must behave in the real world, not just on the drawing board. This is especially true when we consider security and safety as qualities. By mapping goals to capabilities, the framework shows that theft protection, user safety, and environmental safety all require specific system features (e.g., anti-theft systems, alerts, helmet).

Lastly, by modelling, we can represent our mental model but also get the chance to see hidden dependencies. Sequence diagrams and capability models highlight interactions between components and users, making it easier to detect design gaps or conflicts early.

1.2.1.1 Future Works

To extend this reference architecture, we can incorporate cyber-resilience as a set of architectural drivers, ensuring security requirements are considered across development, deployment, use, and recovery phases. We can map regulatory requirements such as NIST or CRA to system capabilities and design decisions, ensuring compliance is built into the architecture rather than bolted on later. We can also extend the capability model to implement cyber-resilience capabilities such as authentication, anomaly monitoring, cryptographic key management, incident response and recover or data protection. Finally, interaction diagrams between the system and the environment of the system can help us identify where to place cyber-resilient controls and how the interaction can help us detect anomalies or show safety or security weaknesses in the system.

1.3 Methodology in Summary

Referring to the white paper that was written about the cyber-resilient systems engineering methodology from the introduction, we extended the systems engineering life cycle to include more comprehensive roles, inputs, outputs, recursive and iterative loops in the decision making, see **Error! Reference source not found.**. This revised methodology takes a part of the existing system of systems, makes it cyber-resilient and repeats the process until a resilient system of systems is achieved. It starts with understanding what the system of systems is supposed to do and what role sub-system plays. Even though these systems already work, they are treated as “brownfield” cases because they are not yet resilient against cyber threats. The teams, in coloured boxes shown in **Error! Reference source not found.**, then examines what qualities the system needs (such as safety, security, and resilience) and checks for operational and capability weaknesses. Experts such as the system operator, business manager, and safety/security engineers identify risks, propose solutions, and decide together whether the system is resilient enough or needs further improvement.

Once a resilient system concept is defined, engineers set detailed requirements and perform safety, security, and system tests to make sure the system meets the expected standards. If the system does not pass, it loops back for refinement. When it finally meets the necessary quality levels, it is validated, transitioned into operation, and confirmed to deliver the expected capabilities reliably and safely. The overall process involves roles, including systems engineers, business managers, operators, and security specialists, all working together to ensure the system becomes and remains cyber-resilient.

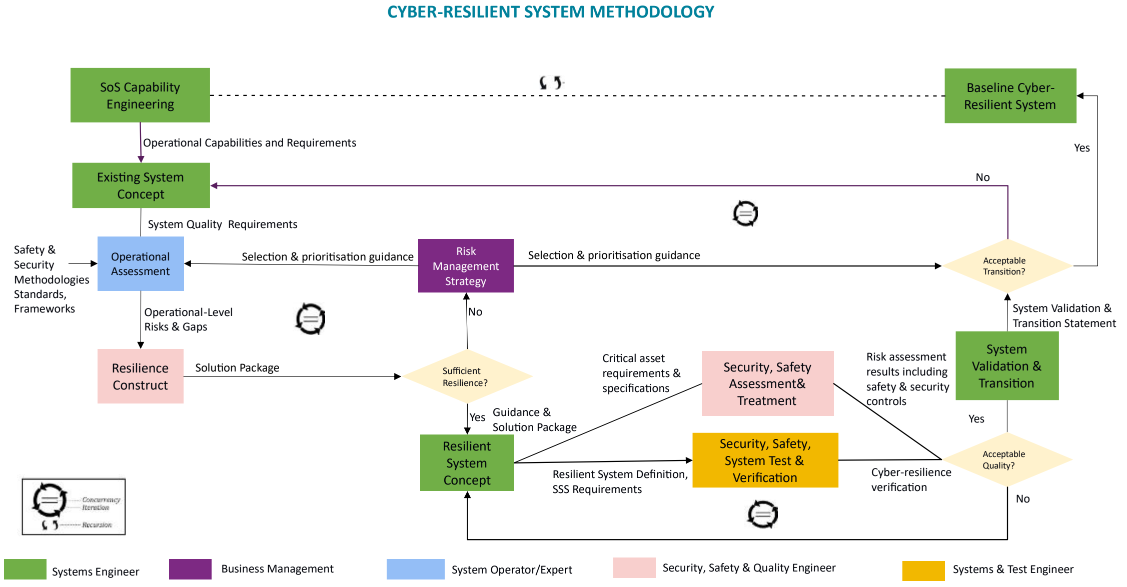
CYBER-RESILIENT SYSTEM METHODOLOGY
SoS Capability Engineering
Baseline Cyber-Resilient System
Operational Capabilities and Requirements
Yes
No
Existing System Concept
System Quality Requirements
Safety & Security Methodologies Standards, Frameworks
Operational Assessment
Selection & prioritisation guidance
Risk Management Strategy
Selection & prioritisation guidance
Acceptable Transition?
No
System Validation & Transition Statement
Operational-Level Risks & Gaps
Resilience Construct
Solution Package
Sufficient Resilience?
Critical asset requirements & specifications
Security, Safety Assessment& Treatment
Risk assessment results including safety & security controls
System Validation & Transition
Yes
Guidance & Solution Package
Yes
Resilient System Concept
Resilient System Definition, SSS Requirements
Security, Safety, System Test & Verification
Cyber-resilience verification
Acceptable Quality?
No
Concurrency
Iteration
Recursion
Systems Engineer
Business Management
System Operator/Expert
Security, Safety & Quality Engineer
Systems & Test Engineer

Figure 6: Cyber-resilient system methodology

A breakdown of the methodology including workflows in the process can be reviewed in the Cyber-secure by Design ERP paper [9]. Within the ERP process, this methodology was introduced to different industry participants and nonetheless, they were unsure where to start from the methodology.
Therefore, there is a need to create a reference architecture for their systems and decide what they need in terms of security and business resilience. They may decide to apply the entire methodology or a part of the methodology that suit their business.

1.4 Conclusion

Organizations are keen to introduce cybersecurity measures but are often unsure where to start due to costs, complexity, and unclear guidelines. From the round table session, we observed that the industrial participants have concerns. They are unsure how to handle vulnerabilities, how to keep products safe after deployment, interact with third-party suppliers and their components or prepare for organisational readiness.

By following a structured, step-by-step approach, starting with a reference architecture, defining what they value in business and how that reflects on their system, organisations can start building a cybersecurity framework that is both effective and resilient. From then onward, applying the cyber-resilient system methodology becomes easier as the organisation can define where they feel the priority and need from risk prioritization and business continuity.

The e-bike example demonstrates that **reference architecture is a strategic tool,** not just technical documentation. It ensures that *what the business wants*, *what the system must do*, and *how the system is built* all align so companies avoid misaligned designs, wasted

investment, and lost competitive advantage. They can also visualise how to introduce safety and security in the system and how that can impact or change the system operation.
To extend this reference architecture into cyber-resilience, we can map cyber-resilience related regulatory and resilience requirements onto system goals, capabilities, and lifecycle stages. Cyber-resilience becomes an architectural driver, influencing capability models (e.g., secure updates, anomaly detection, key management), system contexts (development, deployment, operation), and sequence diagrams that show secure interactions across the ecosystem. This ensures that the system is not only aligned to business strategy but is also engineered to withstand, detect, respond, and recover from cyber-attacks in accordance with emerging regulatory expectations.

2 References

[1] INTERSCT, “INTERSCT,” 2020. [Online]. Available: https://intersct.nl/. [Accessed 3 December 2025].

[2] S. Acur, T. Hendriks, Y. Meijjard and C. van der Vliet-Hameeteman, “SOS! Ensuring safety and security in an expanding system of systems landscape,” 28 September 2023. [Online]. Available: https://ris-data.tno.nl/bibliotheek/sv-015068/TNO/Rapporten/2023/TNO-2023-R11859.pdf. [Accessed 1 September 2025].

[3] S. Acur and T. Hendriks, “Reference Architecture in Relation to Business Reasoning,” 17 April 2023. [Online]. Available: https://ieeexplore.ieee.org/document/10131128/.

[4] Amazon, “Amazon Web Services Architecture,” Amazon Web Services, 2025. [Online]. Available: https://aws.amazon.com/architecture/. [Accessed 3 September 2025].

[5] NSC, “What about zero trust?,” 12 September 2020. [Online]. Available: https://www.ncsc.nl/actueel/weblog/weblog/2020/what-about-zero-trust. [Accessed 13 September 2025].

[6] Microsoft, “Microsoft Cybersecurity Reference Architectures (MCRA),” Microsoft, 2025. [Online]. Available: https://learn.microsoft.com/en-us/security/adoption/mcra . [Accessed 3 September 2025].

[7] D. o. D. (DOD), “Department of Defense Cybersecurity Reference Architecture,” DOD, 30 January 2023. [Online]. Available: https://dodcio.defense.gov/Portals/0/Documents/Library/CS-Ref-Architecture.pdf. [Accessed 9 September 2025].

[8] G. Muller, R. Cloutler, D. Verma, R. Nilchiani, E. Hole and M. Bone, “The Concept of Reference Architecture,” 22 January 2009. [Online]. Available: https://incose.onlinelibrary.wiley.com/doi/10.1002/sys.20129. [Accessed 11 September 2025].

[9] S. Acur, S. K. Das, B. v. d. Leeuw, A. Vasenev and P. Goosen, “Cyber-resilient system design methodology,” 30 April 2025. [Online]. Available: https://repository.tno.nl/SingleDoc?find=UID%200181a4c4-cd2d-4eae-a12d-f2e639b969eb. [Accessed 1 November 2025].

ICT, Strategy & Policy

High Tech Campus 25
5656 AE Eindhoven
www.tno.nl

TNO innovation for life