Cyber-DRIVE: Dynamic Risk Versatile Engine for Cyber Security Risk Analysis, Situational Awareness and Incident Response

Stefano Solari[1,†], Umberto Pedrotti[1,†], Enrico Castelli[1,†], Gianluca Ceccoli[1,†], Alessandro Oneto[1,†] and Enrico Russo[2,*,†]

[1]Leonardo SpA, Cyber and Security Solutions Division, PTI, Italy
[2]DIBRIS, University of Genoa, Italy

Abstract

In dynamic and regulated critical environments, continuously assessing cyber risks is essential for effective situational awareness and decision-making. While methodologies like Security Risk Analysis (SRA) and Threat Assessment and Remediation Analysis (TARA) provide structured approaches, they often lack dynamic updates, real-time threat modeling, and advanced visualization capabilities. This paper introduces Cyber-DRIVE, a novel framework for dynamic risk analysis, incident response, and situational awareness. By separating technological and business-critical asset models, Cyber-DRIVE enables cyber analysts to evaluate likelihoods and incidents, while business decision-makers assess the impact of Confidentiality, Integrity, and Availability (CIA) degradations. The framework integrates real-time Cyber Threat Intelligence (CTI), supports dynamic "what-if" scenarios using MITRE D3FEND techniques, and provides advanced visualizations of complex infrastructures. Cyber-DRIVE addresses limitations in existing SRA and TARA tools by improving risk estimation accuracy, enabling faster response to changing threat landscapes, and supporting interoperability with Security Operations Centers (SOC) and orchestration platforms (SOAR). A case study demonstrates its applicability in modeling realistic infrastructures and evaluating cyber-attack scenarios.

Keywords

Dynamic Risk Analysis, Situational Awareness, Cyber Threat Intelligence, Cybersecurity Ontology, Security Mechanisms, Threat Propagation, Kill Chain Generation, Markov Decision Process

1. Introduction

In dynamic and regulated critical environments, *continuously* assessing current risk levels is essential to support situational awareness for technical and strategic decision-makers. Within the current approaches to cyber risk management—from risk identification to mitigation and continuous monitoring—methodologies such as Security Risk Analysis (SRA) and Threat Assessment and Remediation Analysis (TARA) [1, 2] play a crucial role. SRA solutions, e.g., MAGERIT/PILAR [3], ITSRM [4], FAIR [5], NIST 800 [6], provide structured frameworks to identify, assess, and manage risks, focusing on evaluating the likelihood and impact of threats to critical assets. Similarly, TARA tools emphasize identifying adversary tactics and vulnerabilities, mapping risk scenarios to actionable mitigation strategies.

By assessing the existing tools following both SRA and TARA principles, we argue that they need to be improved in updating risk assessments dynamically, incorporating real-time threat intelligence, modeling complex cyber environments, evaluating hypothetical scenarios, and offering advanced visualization for situational awareness and decision-making.

To overcome such limitations, in this paper, we propose a novel framework for cyber security risk analysis, situational awareness, and incident response, namely Cyber-DRIVE. It incorporates

Joint National Conference on Cybersecurity (ITASEC & SERICS 2025), February 03-8, 2025, Bologna, IT
*Corresponding author.
†These authors contributed equally.
stefano.solari@leonardo.com (S. Solari); umberto.pedrotti@leonardo.com (U. Pedrotti); enrico.castelli@leonardo.com (E. Castelli); gianluca.ceccoli@leonardo.com (G. Ceccoli); alessandro.oneto@leonardo.com (A. Oneto); enrico.russo@unige.it (E. Russo)
0000-0002-1077-2771 (E. Russo)

Table 1
Comparison of TARA Features, RSA Methods/Tools, and Cyber-DRIVE Framework.

TARA Features	RSA methods/tools	Cyber-DRIVE framework
Attack Vectors (AVs)	Fixed taxonomy	Public ATT&CK Framework
Countermeasures (CMs)	Fixed taxonomy	Public D3FEND Framework
Scoring models to quantitatively assess AVs and CMs	Lookup from predefined tables	Scored Public D3FEND Framework
Threat Matrix	Defined by applicability of AVs and CMs	Cyber-DRIVE Kill Chain Generation
TTP/CM Mapping processes	Asset Types predefined by Taxonomy	Public D3FEND Framework Ontology
Assets Model	Asset Types predefined by Taxonomy, Asset Containers defined by User	Cyber-DRIVE Custom Ontology
Dependency model	Generally Hierarchical, Nested, User definable C,I,A dependencies	Cyber-DRIVE Tech vs Primary Dependency Layer
Risk Scoring, Weighted Risk Scoring	Risk Scoring, Weighted Risk Scoring	Scored Public D3FEND Framework

concepts from ITSRM, most notably the separation between the technological (focused on systems and infrastructure) and "primary" assets (focused on business-critical resources) models. This separation enables a clear abstraction and segregation of roles and competencies between two user categories: (*i*) *cyber technology analysts*, who estimate likelihoods and validate or reject hypotheses based on detected indicators of compromise and attacks (and investigation when cyber incidents occur), and (*ii*) *business or process operations analysts*, who understand the value of assets and estimate the impact of Confidentiality, Integrity, and Availability (CIA) degradation on those assets.

Our approach aims to advance the capabilities of current SRA and TARA tools by enhancing risk estimation accuracy, response times in case of model changes, and information updates on cyber threat behaviors and used tactics and techniques. As a result, it enables the realization of risk management, situational awareness, and decision support tools for a cyber response against security incidents by interoperating with Security Operations Centers (SOC) and orchestration platforms (SOAR), and by being fed with Cyber Threat Intelligence (CTI). Moreover, it supports dynamic "what-if" evaluations of active countermeasures based on techniques described in the MITRE D3FEND framework [7]. Finally, it provides comprehensive visual representations of complex cyberspaces, including software-defined networking, virtualization, and containerization.

The main contribution of the paper can be summarized as follows.

- We present the Cyber-DRIVE framework, designed for dynamic cyber risk identification, analysis, and assessment.
- We define a flexible cybersecurity ontology and propagation rules, enabling the customization of the framework as the core engine for a wide range of domain-specific Dynamic Risk Management tools.
- We propose an asset modeling methodology that surpasses traditional dependency trees or graphs used in classical SRA. It employs Markov Decision Process (MDP) modeling and optimization, with transition probabilities automatically estimated using scores derived from the widely recognized MITRE D3FEND framework.

Paper Structure. The paper is structured as follows. Section 2 presents related work. Section 3 details our methodology. Section 4 describes our reference case study. Section 5 discusses our solution, and Section 6 draws the conclusions.

2. Related Work

Table 1 compares the Cyber-DRIVE framework against three mainstream RSA methods (i.e., MAGERIT/PILAR, FAIR, ITSRM) and features defined by the TARA metamodel.

TARA methodologies use catalogs of Attack Vectors (AVs) to ensure consistent threat identification and risk scenario definition. However, RSA tools often rely on fixed taxonomies, limiting adaptability

to evolving threats. Cyber-DRIVE addresses this by integrating the dynamic ATT&CK Framework, providing up-to-date representation of Tactics, Techniques, and Procedures (TTPs) for emerging threats.

TARA methodologies use predefined Countermeasures (CMs) mapped to threats via consistent catalogs. RSA tools rely on static taxonomies, limiting adaptability to specific contexts or new vulnerabilities. Cyber-DRIVE improves this with the D3FEND Framework, offering detailed and flexible countermeasure mapping for better response planning .

TARA methodologies rank AVs and CMs using scoring models based on risk and cost-effectiveness, aiding decision-making. RSA tools rely on rigid lookup tables, limiting adaptability. Cyber-DRIVE uses dynamic, context-aware scoring from the D3FEND Framework for flexible risk evaluation and mitigation prioritization.

RSA tools use static threat matrices, limiting effectiveness in dynamic cyber environments. Cyber-DRIVE addresses this with real-time kill chain generation, modeling adversarial strategies to identify likely attack paths and impacts for adaptive threat analysis.

TARA methodologies and RSA tools use asset taxonomies to map TTPs to countermeasures, but these mappings are often predefined and static. This requires manual updates to adapt to changes in the threat landscape. Cyber-DRIVE automates this process by relying on a dynamic ontology based on the D3FEND framework, ensuring flexible and consistent mapping of threats to countermeasures.

RSA tools and TARA methodologies typically use static, predefined asset taxonomies that do not account for the complexity of modern infrastructures. Cyber-DRIVE introduces a customizable asset ontology, which allows for the modeling of diverse types of assets, including physical, logical, and cyber persona assets. This flexibility supports the representation of complex infrastructures, such as virtualized environments and containerized systems, while maintaining compliance with a baseline ontology.

RSA tools often rely on hierarchical or nested dependency models, which can oversimplify the relationships between assets and their dependencies. Cyber-DRIVE improves this by introducing a layered dependency model that distinguishes between technological and primary assets. This enables more accurate modeling of how risks propagate through a system.

RSA tools and TARA methodologies typically apply static risk scoring or weighted scoring models. These approaches lack adaptability and are often insufficient for dynamically evolving cyber environments. Cyber-DRIVE integrates dynamic scoring models derived from the D3FEND Framework, improving accuracy and enabling real-time updates to reflect changing risks and vulnerabilities.

3. Methodology

The architecture of Cyber-DRIVE consists of three basic components: (*i*) **Knowledge Base (KB)**, an ontology-based, fine-grained knowledge base containing models of user-defined ICT infrastructures (extensible to other domains), (*ii*) **Kill Chains Generation (KGC)**, an optimization algorithm based on Markov Decision Process (MDP) modeling, utilizing Reinforcement Learning (RL) techniques [8], and (*iii*) **Risk Extraction (REX)**, a risk calculation and aggregation procedure that utilizes the most likely kill chains previously generated.

Cyber-DRIVE evaluates potential cybersecurity attack paths in near real-time by leveraging the D3FEND framework (which complements the ATT&CK matrix) and extending its use through a proprietary ontology stored in its KB. The KB enables the modeling of realistic ICT processing and networking stacks, enriching the technological asset landscape with security mechanisms, measures, and CVEs. Users can freely model their infrastructures while ensuring compliance with the baseline ontology in the KB, which forms the foundation for accurate and explainable risk evaluations.

During the KCG phase, the Cyber-DRIVE engine uses the enriched KB to generate adversarial kill chains based on the most threat-rewarding paths. It propagates attack likelihoods from the technological asset layer to the primary asset layer, projecting Confidentiality, Integrity, and Availability (CIA) degradation for primary assets.

Finally, the REX phase then aggregates the results of the kill chain analysis, calculating the expected

levels of degradation (likelihood) and the impact for each cybersecurity risk dimension. This process allows the system to provide actionable insights on residual risks and enable dynamic "what if" assessments for exploring the effectiveness of mitigation solutions.

3.1. Knowledge Base

The Knowledge Base (KB) contains the user-defined model of technological assets and the corresponding ontology, which guides model creation and validation. Technological assets are classified into physical (e.g., servers, switches, devices), logical (e.g., OS, software, artifacts), cyber persona (e.g., credentials, email accounts), or human world entities (e.g., administrators, end-users). The ontology is partially proprietary (covering technological asset entities and their relationships) and partially standard, based on the D3FEND Ontology Framework. Pertinence links connect D3FEND artifacts to proprietary technological asset models, weighted to estimate the effectiveness of attack tactics/techniques and the defensive resistance of D3FEND techniques. These scores feed directly into evaluating transition probabilities for the MDP problem that generates the potential kill chains (see below).

The KB also includes *Security Mechanisms* and *Security Threats*. Users can instantiate Security Mechanisms to represent preventive, mitigative, or protective technologies defending specific assets. These mechanisms establish "defends" relationships with technological assets, defining where and how security measures can be applied, adding realism to "what-if" analysis. Additionally, Security Mechanisms are modeled as implemented within technological assets, enabling the system to account for risks affecting both the protected assets and the assets hosting the protective mechanisms. This feature allows further research and potential algorithm extension to encompass cases where risk affects technological assets that also implement in some way the Security Mechanism protecting other assets.

Security threats justify the need for Security Mechanisms. These threats establish "attacks" relationships with technological assets, parametrized using ATT&CK tactics and techniques. This integration ensures that the impact of Security Threats on the asset model is incorporated into residual risk calculations, similar to the effects of cyber incidents.

Finally, the KB allows users to adjust the relative importance of CIA dimensions for each primary asset, enabling tailored risk assessments.

3.2. Kill Chains Generation

The generation of adversary kill chains for risk evaluation is a key research area in cyber risk management [9]. Our approach relies on a probabilistic attack path generation [10] obtained by preparing, running, and solving an MDP model. The above kill chains are the results of optimal decision strategies in a carefully crafted MDP space of states/actions.

Using an MDP framework ensures efficient handling of the combinatorial complexity in large attack path search spaces, provides a robust foundation for simulating cyber threat actions, and captures the probabilistic and unpredictable nature of advanced threats like APTs. Its state/action abstraction aligns with the logical movements of cyber threats in constrained environments. At the same time, its reinforcement learning basis supports future research leveraging advanced AI techniques for improved scalability and accuracy.

Below, we introduce the overall procedure in terms of defining the input of the MDP problem, representing states and actions in the MDP model, estimating probabilities in MDP transitions, and defining the rewarding mechanism.

Problem input. The MDP problem input is dynamically adjusted/updated according to the risk evaluation setting. From the risk management perspective, it can reflect the need to evaluate CIA-based risk values without any countermeasure (initial/baseline risk), with only default active countermeasures (steady state risk), or otherwise triggered with additional countermeasures to assess Cyber Response Options during Cyber Incident Response (residual risk).

States. A state represents the levels and types of cyber threat advantage, i.e., the obtained foothold over any technological asset of the target system landscape together with information *flags*. Flags keep track of the status of the privilege escalation, persistency, and/or command and control capability obtained from the technological asset. As realistic target infrastructures can have hundreds or even thousands of technological elements, the Cyber-DRIVE framework avoids a possible combinatorial explosion due to the state cardinality by properly segmenting the state representations encoding during the optimization problem preparation.

Actions. Actions in the MDP model represent adversary choices to transition between states, expanding their attack path and/or increasing their threatening posture. These actions are encoded as threat propagation rules, which define how adversaries explore and exploit realistic attack paths. Actions originating from a specific state, such as a technological asset with a predefined compromise status like Elevation of Privilege (EoP) or Command and Control (C2), can semantically correspond to TTPs, CVE exploits, or malicious but system-legitimate activities (e.g., using stolen credentials to access otherwise restricted services or data).

Transition probabilities estimation. Each state/action transition probability is estimated via an evaluation function that depends on the type of action.

For action associated with an ATT&CK Tactic or Technique, the approach is inspired by MITRE recommendations for TARA methods [2, 4.2 Assessing Countermeasure Effects]. Briefly, these recommendations surpass the characterization of countermeasures with a binary state (effective or it is not) and classify their effects by type, i.e., detect, neutralize, limit, and recover, and by magnitude, i.e., low, medium, and high.

Building on this perspective, our approach develops and applies a more sophisticated cyber effects model, fully leveraging the D3FEND ontology artifacts and, particularly, the existing relationship among techniques and artifacts.

For each technological asset X in the model, quantifying this interaction requires calculating the *Threat*, *Defense*, and *Mitigated Intensity Scores*. To enable these calculations, we consider $SA(X)$, the set of artifacts pertaining to each asset. The integer-valued function $Pert(X, \text{artifact} \in SA(X))$ assigns weights to artifacts based on their pertinence to the asset.

To calculate the Threat Intensity Score $TIS(X, T)$ for any attack technique T whose actions act on the set of artifacts $SA(X)$, the following formula is applied:

$$TIS(X, T) = \sum_{a \in SA(X)} \Big[TAS(X, T, a) \cdot Pert(X, \text{a}) \Big]$$

where $TAS(X, T, a)$ represents the contribution of each action of technique T on the specific artifact a, within the context of the technological asset X, namely the Threat Action Score. This score is a framework-defined parameter estimated by experts, reflecting the relative severity of different attack actions.

To calculate the Defense Intensity Score $DIS(X, D)$ for any defense technique D whose actions act on the set of artifacts $SA(X)$, the following formula is applied:

$$DIS(X, D) = \sum_{a \in SA(X)} \Big[DAS(X, D, a) \cdot Pert(X, a) \Big]$$

where $DAS(X, D, a)$ represents the contribution of each action of defense technique D on the specific artifact a, within the context of the technological asset X, namely the Defense Action Score. Like TAS, it is a framework-defined parameter estimated by experts, reflecting the relative effectiveness of defense actions.

Finally, to calculate the Mitigated Threat Intensity Score $MTIS(X, T, SD(X))$ for any attack technique T applicable to the asset X, and any set of given defense techniques $SD(X)$ acting on $SA(X)$, the following formula is applied:

$$MTIS(X, T, SD(X)) = \sum_{a \in SA(X)} \Big[\big(TIS(X, T) - DIS(X, D) \big) \cdot Pert(X, a) \Big]$$

where $TIS(X, T)$ is the Threat Intensity Score for technique T and $DIS(X, D)$ is the Defense Intensity Score for technique D, both evaluated for the specific artifact a within the context of the technological asset X.

The scheme of MDP transition probabilities is derived by normalizing the Mitigated Threat Intensity Score at the graph level, which serves as the overall scoring function. This function estimates the effectiveness of the given ATT&CK Technique over the target Technological Asset X, while accounting for the defensive resistance provided by the selected mitigations ($SD(X)$) in the form of D3FEND techniques.

The MDP transition probability scheme is derived by normalizing the Mitigated Threat Intensity Score at the graph level. This score serves as the overall function to estimate the effectiveness of a given ATT&CK technique on a target technological asset, while also considering the defensive resistance provided by user-selected D3FEND techniques during the risk evaluation process. The numerical balance is safeguarded through saturation functions (MAX, MIN) and standard normalization operations. This ultimately contributes to estimating the transition probability distribution over MDP states reachable from a given initial state, when a cyber threat action aligns with a specific ATT&CK tactic or technique. Although transition probabilities in the MDP problem are generated automatically, the framework allows users to override these parameters for fine-tuning specific propagation rules.

For actions associated with a CVE exploit, each MDP state/action transition probability is estimated using an evaluation function based partially on the CVSS score system [11] and partially on expert-provided data for new CVEs applicable to the technological asset landscape. Cyber-DRIVE mitigates such exploits through CVE removal (e.g., patching or procedures linked to the CVE description), fully counterbalancing the associated threat intensity score.

For actions representing malicious but system-legitimate use, the transition probability is set to a constant value between 90% and 100%, as these actions are inherently hard to detect and prevent due to their legitimacy.

Rewarding. The rewarding strategy of the MDP problem is designed to prioritize states where the threat gains an advantage over technological assets that are linked to the user-defined primary assets (see the Dependency model layer in Section 3.3). This approach ensures that the generated threat strategies focus on degrading the CIA of technological assets, ultimately impacting the CIA dimensions of the primary assets.

3.3. Risk Extraction

This phase starts by post-processing the outcome of the optimization problem, which is an optimal strategy in the threat advantage state space, to project it into a viable attack path in the technological cyber terrain. During the risk extraction process, the paths originating from the assets declared in the cyber incidents and terminating at the technological assets on which primary assets depend are considered. These paths are derived from the optimal strategy generated as the solution to the MDP problem, which is triggered by the active cyber incident information and the hypothesized set of mitigation measures to be activated as input. During extraction, path probabilities are calculated starting from the transition probabilities used for the MDP problem input.

Other KPIs are gathered during the extraction process based on the attributes stored in each threat action propagation rule, such as the stealthiness/detectability level and the CIA degradation likelihoods for all the assets involved in the threat paths. In particular, the CIA degradation likelihoods are considered for the technological asset at the end of the path, on which primary assets depend. These likelihoods are multiplied by the overall path probability to compute the total CIA degradation likelihood for each technological asset.

The last steps of the processing are: (*i*) the projection and calculation to the primary asset CIA degradation likelihoods through the *dependency model layer*, and (*ii*) the application of risk formula taking into account the primary assets CIA degradation likelihoods and the corresponding loss values.

Dependency model layer. We propose a dependency model that enables users to define logical gates representing reasonable projections for the CIA dimensions. This solution applies not only to Availability−−where AND and OR logic is straightforwardly interpreted and utilized to model redundancy, recovery, and primary versus alternative dependencies−−but also to Confidentiality and Integrity.

The dependency model layer relies on three logical gates: *AND*, *OR*, and *MAX*. AND gates are helpful for modeling single points of failure, where a direct degradation of one technological asset fully propagates to the dependent primary asset. OR gates are suitable for redundancy modeling, where the presence of alternative assets can mitigate the degradation of a single technological asset. MAX gates are applied in scenarios where the maximum degradation among multiple technological assets defines the impact on the primary asset.

Risk formula. The key parameters required for risk calculation, which are assigned to each primary asset, as well as the calculations for likelihood, impact, and risk evaluation, can vary depending on the most suitable value modeling approach for the specific application. In Cyber-DRIVE, we chose not to treat the value as an additive and potentially unbounded measure (as might be the case if a value represents monetary amounts). Instead, Cyber-DRIVE adopts the concept of a consumable, finite value measure, which is more suitable for addressing risks associated with assets such as technical performance or operational capabilities. This approach benefits applications where additional domain-specific impact models are applied downstream to Cyber-DRIVE. Examples include evaluating security risks in embedded systems for cyber-physical systems or assessing mission risks and impact indicators in the context of military operations planning.

The Cyber-DRIVE framework evaluates primary asset risk through three components: input parameters provided by the user, intermediate results bridging input to output, and final output parameters summarizing overall risk.

Input parameters are provided by the business-level user and ensure that the evaluation process reflects organizational priorities and dependencies. Considering $X \in \{C, I, A\}$, the parameters include:

- $BASE\_VALUE_{PA}$, intrinsic total value assigned to the primary asset (PA).
- $X\text{-}TOTAL\_LOSS_{PA}$, value loss in case of X total loss for the primary asset (PA).
- $REL\_WEIGHT_{PA,X}$, CIA relative weight (relevance) for the primary asset (PA).
- $LKH_{TA1}, \ldots, LKH_{TA_{MTA}}$, likelihoods associated with technical assets.
- $DEPENDENCY\_MODEL_X(TA1, \ldots, TA_{MTA})$, processing of Cyber-DRIVE generated likelihoods of X degradation for technical assets ($LKH_{TA1}, \ldots, LKH_{TA_{MTA}}$) involved in the cyber incident through user-defined AND_X or OR_X gates with X, giving LKH_X as output.

The intermediate results are derived from the dependency model and input parameters to quantify the likelihood and impact of degradation for each security dimension $X \in \{C, I, A\}$.

$CONS_{PA,X}$ is the relative consequence (or impact) for any security dimension X. It represents the proportion of the potential value loss for the primary asset due to the degradation of X, normalized by the total value of the asset:

$$CONS_{PA,X} = \frac{X\text{-}TOTAL\_LOSS_{PA}}{BASE\_VALUE_{PA}}$$

LKH_X is the likelihood of degradation for the security dimension X. Computed using the dependency model, it aggregates the degradation likelihoods of all related technical assets ($LKH_{TA1}, \ldots, LKH_{TA_{MTA}}$) based on user-defined gates (AND or OR):

$$LKH_X = DEPENDENCY\_MODEL_X(LKH_{TA1}, \ldots, LKH_{TA_{MTA}})$$

$LOSS\_RISK_X$ combines the consequence ($CONS_{PA,X}$) and likelihood (LKH_X) to compute the loss risk for the security dimension X, quantifying the expected degradation impact:

$$LOSS\_RISK_X = CONS_{PA,X} \cdot LKH_X$$

The output parameters represent the aggregated consequences, likelihoods, and risks for each primary asset (PA) after processing the intermediate results. These are defined as follows:

$CONS_{PA}$ aggregates $CONS_{PA,X}$ across all security dimensions X, weighted by their $REL\_WEIGHT_{PA,X}$ for the primary asset (PA):

$$CONS_{PA} = \sum_X (REL\_WEIGHT_{PA,X} \cdot CONS_{PA,X})$$

$LOSS\_RISK_{PA}$ is the overall risk for the primary asset (PA). It is calculated by combining the relative consequence and degradation likelihood for each security dimension X, weighted by their relative importance. It can also be expressed as:

$$\begin{aligned} LOSS\_RISK_{PA} &= \sum_X REL\_WEIGHT_{PA,X} \cdot (CONS_{PA,X} \cdot LKH_X) \\ &= \sum_X REL\_WEIGHT_{PA,X} \cdot LOSS\_RISK_X \end{aligned}$$

LKH_{PA} is the overall likelihood of degradation for the primary asset (PA). It is derived as the ratio of the overall loss risk to the aggregated consequence:

$$LKH_{PA} = \frac{LOSS\_RISK_{PA}}{CONS_{PA}}$$

4. Case Study

To test and generate benchmark output data for Cyber-DRIVE, we modeled a ICT infrastructure using the Cyber-DRIVE ontology, as shown in Figure 1, and simulated a realistic company network for evaluation. Briefly, at the top of the diagram, a simulated Internet connection includes routers and infrastructure services such as DNS servers. In the middle, the company backbone network connects the DMZ, core infrastructure, and client areas. Each element in the diagram represents a technological asset, which provides the functionality of primary assets, such as data or services that are targets of risk calculation. These assets correspond to physical workstations or network devices. On top of these physical elements, we modeled software components like operating systems, application software, services, network interfaces, user accounts, and credentials. The Cyber-DRIVE ontology defines all possible relationships between components and their attributes. For instance, attributes include traffic type (e.g., user sessions, email, web), encryption level, storage information (e.g., local, cloud), and authentication types (e.g., local, domain, two-factor).

Kill chain of the modeled attack. We modeled an attacker sending an email containing malware to an engineering user, who opens a malicious macro that downloads and persists a file from the attacker's C&C server. The malware then moves laterally to disrupt critical services by exploiting a vulnerability.

Cyber-DRIVE encodes these steps as follow: (1) movement from a mail server software to the computer system hosting it; (2) movement from the mail server to the client computer system that receives the mail; (3) the malicious payload from the client is available in the mail software; (4)

movement from the client mail software, the attached malicious file is downloaded onto the client; (5) on the client, the user executes a malicious file; (6) the client opens a reverse shell to allow C2C from an attacker; (7) a port monitor is opened on the client to gain persistence; (8) connection from the client to a remote SSH server suffering from a public CVE; (9) the SSH server authenticates as root user; (10) the software running on the computer system stops working; (11) the software loses data.

Table 2 represents the translation of the modeled kill chain into Cyber-DRIVE propagation rules. The Propagation Rule and Description columns specify the rule and its concise description governing how the attack progresses based on the interaction between the asset's properties and the threat advantage. The Tech Asset Type categorizes the affected assets (e.g., servers, clients, or specific software). It can specify both the source and the destination assets. The Properties of the Technological Asset column lists relevant characteristics of the asset involved in each action, such as traffic profiles, software attributes, or storage types. Finally, the Threat Advantage State specifies the attacker's pre-existing threat advantages necessary to execute the step and the resulting advantage gained. In particular, we specify a multi-attribute advantage state in terms of gained privilege level, obtained persistence, obtained external communication (and command reception) capability, and obtained availability/readiness to execute a payload.

For example, step (1) of the kill chain aligns with the propagation rule related to ATT&CK technique T1566.001 [12], involving a spearphishing attachment. This assumes an email with a malicious payload is received by a mail server managing mailboxes, and the propagation rule makes the email accessible to connected mail clients through standard mail messaging protocols via existing L3 network connections. In this scenario, the rule represents a foothold movement without any immediate attacker advantage. Step (2) of the kill chain corresponds to a file download action involving the email client and the operating system ("computer system") handling the file. Step (3) reflects the technique T1204.002 [13], where executing the file initiates infection. This step does not propagate to other assets but provides the attacker with a foothold by enabling the payload, whose effects will propagate in subsequent steps.

Moreover, Figure 1 illustrates the visualization capabilities of the framework. In particular, Cyber-DRIVE highlights the expected attack path in red, identifying the client machine where the email is received, the assets involved in lateral movements, and the final target system.

5. Discussion

The Cyber-DRIVE framework achieves dynamic cyber risk management by integrating advanced probabilistic modeling, ontology-based knowledge representation, and real-time adaptability. The adoption of MDPs enables the dynamic generation of optimal kill chains by analyzing threat propagation paths and incorporating evolving cyber incident data, hypothesized countermeasures, and KB inputs. Furthermore, the dependency model layer facilitates the seamless propagation of risk impacts from technological assets to primary assets, supporting real-time updates of risk metrics as new information becomes available. This dynamic capability ensures that Cyber-DRIVE remains effective in scenarios where rapid decision-making and adaptation are crucial.

For risk analysis, Cyber-DRIVE leverages threat and defense intensity scores derived from the stat-of-the-art MITRE ATT&CK and D3FEND frameworks. These scores are dynamically applied to evaluate the likelihood and impact of attack techniques on technological assets. Resulting evaluations are projected onto primary assets through the dependency model layer, mapping technical risks to business impacts using logical gates (AND, OR, MAX). This feature ensures that the analysis bridges both technical and operational perspectives.

In terms of risk assessment, Cyber-DRIVE combines likelihoods and impact metrics through a structured risk formula, producing clear and actionable insights for each primary asset. By supporting "what-if" analysis, the framework enables users to evaluate the effects of mitigation measures and adjust risk estimates in realtime. This capability ensures informed decision-making for effective and cost-efficient incident response strategies. Moreover, it is worth noting that the framework fits the cognitive needs of business process experts and decision-makers who focus on asset value and potential value

loss while abstracting the technological details of cyber threat propagation. These details are managed by cyber protection service providers, ensuring that risk assessments bridge technical evaluations with actionable business insights.

6. Conclusions

This paper presents the development of the Cyber-DRIVE framework, guided by key requirements and insights from RSA and TARA methodologies. The proposed methodology, with its potential applications and functional components, demonstrates a robust approach to cyber risk management. A realistic case study of a complex cyber-attack benchmarks the effectiveness of the kill chain generation approach. These contributions position Cyber-DRIVE as a promising tool for advancing cyber risk management by addressing the evolving challenges of cybersecurity in dynamic and critical environments.

Acknowledgments

This work was partially funded by the European Defence Industrial Development Programme (EDIDP) project "European Cyber Situational Awareness Platform" (ECYSAP) and the NextGenerationEU project "Security and Rights in CyberSpace" (SERICS).

Declaration on Generative AI

The author(s) have not employed any Generative AI tools.

References

[1] J. Wynn, J. Whitmore, G. Upton, D. Spriggs, R. McKinnon, R. McInnes, L. Graubart, J. Clausen, Threat assessment & remediation analysis (tara) methodology description version 1.0, Bedford, MA (2011).

[2] J. Wynn, J. Whitmore, G. Upton, L. Spriggs, D. McKinnon, R. McInnes, R. Graubart, L. Clausen, Threat assessment and remediation analysis (tara), MITRE Corporation: Bedford, MA, USA (2014).

[3] Mit java wordnet interface (jwi) user's guide, https://www.ccn-cert.cni.es/es/pdf/guias/series-ccn-stic/guias-de-acceso-publico-ccn-stic/2841-ccn-stic-470i1-pilar-manual-de-usuario-v7-1/file?format=html, 2018. Accessed: 2024-12-14.

[4] E. U. A. for Cybersecurity., Risk management standards: analysis of standardisation requirements in support of cybersecurity policy., Publications Office, LU, 2022. URL: https://data.europa.eu/doi/10.2824/001991. doi:10.2824/001991.

[5] O. F. Dan Blum, CISSP, P. Laura Voicu, Case study: How fair risk quantification enables information security decisions at swisscom, ISACA Journal 5 (2020) 38–47.

[6] Integrated enterprise-wide risk management, https://csrc.nist.gov/csrc/media/events/ispab-july-2009-meeting/documents/ispab_july09-ross_harmonization-sp800-53-rev3.pdf, 2009. Accessed: 2024-12-14.

[7] P. E. Kaloroumakis, M. J. Smith, Toward a knowledge graph of cybersecurity countermeasures, The MITRE Corporation 11 (2021) 2021.

[8] R. S. Sutton, A. G. Barto, Reinforcement learning: An introduction, MIT press, 2018.

[9] M. Angelini, N. Prigent, G. Santucci, Percival: proactive and reactive attack and response assessment for cyber incidents using visual analytics, in: 2015 IEEE Symposium on Visualization for Cyber Security (VizSec), IEEE, 2015. URL: http://dx.doi.org/10.1109/VIZSEC.2015.7312764. doi:10.1109/vizsec.2015.7312764.

[10] S. Jajodia, S. Noel, Topological Vulnerability Analysis, Springer US, 2009, p. 139–154. URL: http://dx.doi.org/10.1007/978-1-4419-0140-8_7. doi:10.1007/978-1-4419-0140-8_7.

[11] X. Ou, S. Govindavajhala, A. W. Appel, MulVAL: A logic-based network security analyzer, in: 14th USENIX Security Symposium (USENIX Security 05), USENIX Association, Baltimore, MD, 2005. URL: https://www.usenix.org/conference/14th-usenix-security-symposium/mulval-logic-based-network-security-analyzer.
[12] MITRE ATT&CK, Spearphishing attachment, https://attack.mitre.org/techniques/T1566/001/, 2023. URL: https://attack.mitre.org/techniques/T1566/001/, accessed on December 2024.
[13] MITRE ATT&CK, User execution: Malicious file, https://attack.mitre.org/techniques/T1204/002/, 2023. URL: https://attack.mitre.org/techniques/T1204/002/, accessed on December 2024.

A. ICT infrastructure and Modeled Attack

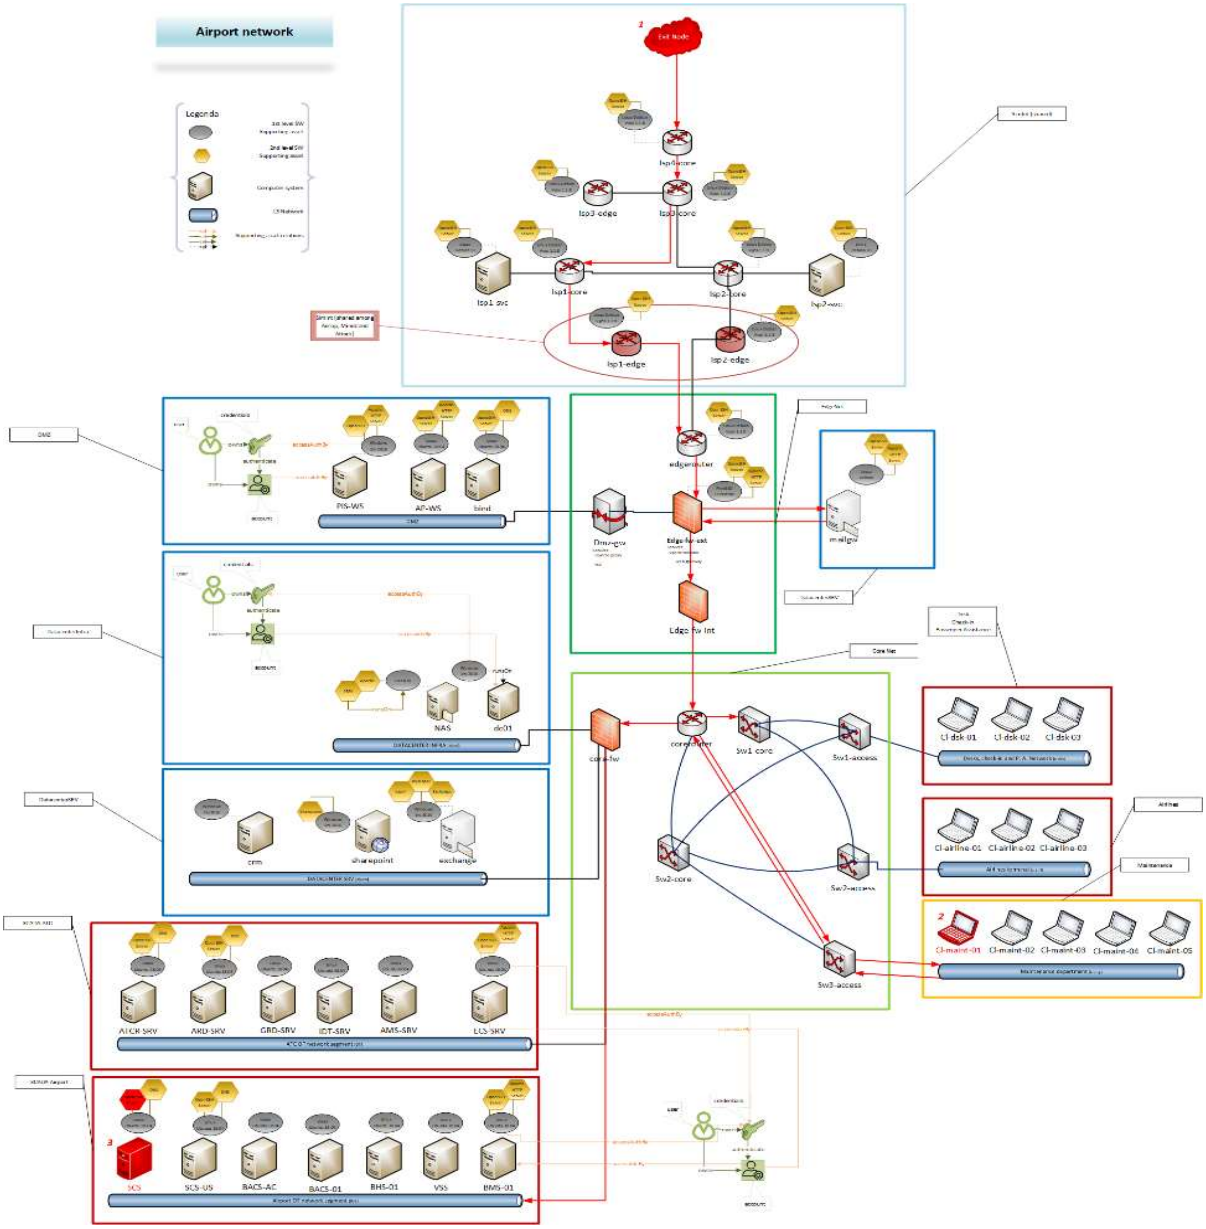

Figure 1: ICT Infrastructure under attack for Cyber-DRIVE output validation.

Table 2: Translation table of a reference kill chain into Cyber-DRIVE propagation rules.

PROPAGATION RULE	DESC (TITLE)	TECH ASSET TYPE	PROPERTIES	THREAT ADVANTAGE STATE (Required / Obtained)			
				PRIVILEGE	PERSISTENCE	EXT COMM /C2C	PAY LOAD
Type: MITRE ATT&CK Rule Desc: T1566.001	Phishing: Spearphishing Attachment	From: Software	riaSpecificData.trafficProfileTDPList[].type = EmailMessaging riaSpecificData.trafficProfileTDPList[].direction = Outbound	~ any ~	~ any ~	~ any ~	~ any ~
		To: ClientSoftware	riaSpecificData.trafficProfileTDPList[].type = EmailMessaging riaSpecificData.trafficProfileTDPList[].direction = Inbound	~ any ~	~ any ~	~ any ~	Obt: 1
Type: Sys-Legit User Action Rule Desc: SLUA.001	File Download	From: ClientSoftware	riaSpecificData.trafficProfileTDPList[].type = EmailMessaging riaSpecificData.trafficProfileTDPList[].direction = Inbound	~ any ~	~ any ~	~ any ~	Req: 1
		To: ComputerClientSystem	riaSpecificData.hasStorage = True	~ any ~	~ any ~	~ any ~	Obt: 1
Type: MITRE ATT&CK- Rule Desc: T1204.002	User Execution: Malicious File	From: ComputerClientSystem	riaSpecificData.hasStorage = True	~ any ~	~ any ~	~ any ~	Req: 1
		To: ~ self ~		~ any ~	~ any ~	~ any ~	Obt: 2
Type: MITRE ATT&CK Rule Desc: T1071	C2C: Application Layer Protocol	From: ComputerClientSystem, ComputerSystem	computerSystem.osPlatform = Windows	~ any ~	~ any ~	~ any ~	Req: 1
		To: ~ self ~		~ any ~	~ any ~	Obt: 1	~ any ~
Type: MITRE ATT&CK Rule Desc: T1547.010	Boot or Logon Autostart Execution: Port Monitors	From: ComputerClientSystem, ComputerSystem	computerSystem.osPlatform = Windows	~ any ~	~ any ~	~ any ~	Req: 1
		To: ~ self ~		~ any ~	Obt:1	~ any ~	~ any ~
Type: CVE EXPLOIT Rule Desc: CVE-2024-6387	Exploit SSH CVE xxx	From: ComputerClientSystem, ComputerSystem		~ any ~	~ any ~	~ any ~	~ any ~
		Via: Software	trafficProfileTDPList[].type = SSH trafficProfileTDPList[].direction = Inbound activeCVE[].id =CVE-2024-6387	~ any ~	~ any ~	~ any ~	~ any ~
		To: ComputerClientSystem, ComputerSystem	~ any ~	Obt:2	~ any ~	~ any ~	~ any ~
Type: MITRE ATT&CK Rule Desc: T1489	Service Stop	From: ComputerClientSystem, ComputerSystem	~ any ~	Req: 2	~ any ~	~ any ~	~ any ~
		To: Software	~ any ~	~ any ~	~ any ~	~ any ~	~ any ~
Type: MITRE ATT&CK Rule Desc: T1485	Data Destruction	From: ComputerClientSystem, ComputerSystem	~ any ~	Req: 2	~ any ~	~ any ~	~ any ~
		To: Software	~ any ~	~ any ~	~ any ~	~ any ~	~ any ~