Acta Informatica Pragensia
2027, Volume 16, Issue 1, Forthcoming article
https://doi.org/10.18267/j.aip.334

PRAGUE UNIVERSITY OF ECONOMICS AND BUSINESS

Article Open Access

Initial Access Threat Assessment Methodology for Internet-Exposed Assets

Oleksii Shuba [1], Ondrej Kainz [2], Ivana Semanová [2], Miroslav Murin [2], Peter Feciľak [2], Ervín Šimko [3]

[1] Siemens Healthineers, Košice, Slovakia
[2] Department of Computers and Informatics, Faculty of Electrical Engineering and Informatics, Technical University of Košice, Košice, Slovakia
[3] Management and Information Science Department, Academy of the Police Force in Bratislava, Bratislava, Slovakia

Corresponding author: Ondrej Kainz (ondrej.kainz@tuke.sk)

Editorial Record

First submission received:
June 6, 2026

Revisions received:
July 31, 2026
August 31, 2026

Accepted for publication:
September 27, 2026

Academic Editor:
Zdenek Smutny
Prague University of Economics and Business, Czech Republic

This article was accepted for publication by the Academic Editor upon evaluation of the reviewers' comments.

How to cite this article:
Shuba, O., Kainz, O., Semanová, I., Murin, M., Feciľak, P., & Šimko, E. (2027). Initial Access Threat Assessment Methodology for Internet-Exposed Assets. *Acta Informatica Pragensia*, 16(1), Forthcoming article. https://doi.org/10.18267/j.aip.334

Copyright:
© 2027 by the author(s). Licensee Prague University of Economics and Business, Czech Republic. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution License (CC BY 4.0).

Abstract

Background: Internet-facing enterprise assets represent a critical part of the external attack surface because they can be discovered, fingerprinted, and targeted by attackers before initial access is established. Although MITRE ATT&CK provides a structured taxonomy of adversarial behavior, organizations still need a repeatable way to transform relevant techniques into prioritized and executable security assessments.

Objective: The objective was to develop and demonstrate a structured MITRE ATT&CK based threat assessment methodology that supports the identification, clustering, prioritization, assessment, and improvement planning of adversarial techniques. Initial Access threats targeting internet-exposed enterprise assets were used as the demonstration domain.

Methods: The methodology was designed as a four-phase assessment lifecycle consisting of (1) baseline analysis, clustering, and preliminary prioritization, (2) Cyber Threat Intelligence-supported assessment prioritization, (3) assessment design and execution, and (4) improvement plan creation. Its feasibility was demonstrated in a single enterprise environment using selected MITRE ATT&CK Enterprise v18.1 techniques from the Reconnaissance, Initial Access, and Credential Access tactics.

Results: The demonstration produced assessment-oriented threat clusters and identified Network/Service Exploitation, Brute Force & Exploitation, and Valid Accounts Abuse as high-priority clusters. A perimeter discoverability assessment showed low external discoverability, 95% average EASM inventory completeness (77.39% completeness for IP ranges), 5–6 times higher domain coverage by the open-source toolchain than by external TPRM platforms, and a 257.54% active-IP attribution ratio for one external source. The findings were transformed into an improvement plan with traceable remediation items.

Conclusion: The proposed methodology provides a practical bridge between MITRE ATT&CK, Cyber Threat Intelligence, security assessment execution, and improvement planning. The demonstration shows that the approach can produce specific and traceable outputs under real enterprise conditions and can be adapted to other ATT&CK tactics or assessment domains. Because it covers one organization and one executed sub-assessment, it establishes feasibility rather than general effectiveness.

Index Terms

MITRE ATT&CK; Initial access; Cyber threat intelligence; Perimeter security; Cybersecurity assessment; Threat assessment; Methodology.

1 INTRODUCTION

The continuing digitalization of enterprise environments has significantly increased the importance of securing systems and services exposed to the public internet.

Modern organizations depend on externally accessible web applications, public APIs, VPN gateways, remote access services, cloud interfaces, email infrastructure, and identity-related services to support daily business operations. While these assets are essential for connectivity, remote work, customer interaction, partner integration, and cloud-based service delivery, they also expand the organization's external attack surface and create additional entry points that can be discovered, fingerprinted, and targeted by attackers (Gelernter et al., 2024). This makes the protection of internet-exposed assets one of the central challenges of modern enterprise cybersecurity.

In the context of this study, the internet-facing perimeter is understood as the set of externally reachable assets that can be accessed, interacted with, or enumerated from the public internet. Such assets may include public-facing applications, VPN and RDP gateways, cloud-accessible APIs, device management interfaces, externally reachable authentication portals, and third-party connected services (Husák & Sadlek, 2025). The more systems and services are exposed to the internet, whether intentionally through business requirements or unintentionally through misconfiguration, legacy deployment, or incomplete asset management, the greater the opportunity for attackers to identify potential initial access paths. The exposure of these assets is particularly significant because it places the first stage of technical compromise at the boundary between the external threat environment and the internal enterprise network.

From a defender's perspective, the internet-facing perimeter is difficult to secure because it is both dynamic and distributed. New assets can appear through cloud deployments, development environments, third-party integrations, mergers, temporary services, or forgotten infrastructure. At the same time, attackers can use automated tools and public data sources to continuously discover domains, subdomains, IP ranges, exposed services, login portals, software versions, TLS certificates, and leaked technical metadata. As a result, the security of the internet-facing perimeter depends not only on vulnerability management, but also on the organization's ability to maintain a complete and up-to-date view of its external exposure.

This problem is intensified by the fact that exposed assets do not need to be directly vulnerable in order to support adversarial activity. Even passive or low-interaction reconnaissance can provide valuable information about technologies in use, authentication surfaces, naming conventions, organizational relationships, or outdated infrastructure.

Therefore, modern perimeter defense requires a structured understanding of which externally applicable adversarial techniques are relevant, how they interrelate, and how they should be assessed by defenders. This need provides the practical motivation for developing a methodology that transforms the MITRE ATT&CK knowledge base (MITRE Corporation, 2025) into an assessment-oriented workflow focused on external threats.

This study makes five contributions. First, it defines a **reproducible clustering process**: the cluster categories are derived from the ATT&CK tactics included in the assessment scope, the clusters are formed within those categories according to six stated criteria, and the assignment rule for techniques that relate to more than one cluster is given, so that another team can repeat the grouping or revise it on a stated basis (RQ2). Second, it defines the **prioritization logic** as a written artifact rather than an analyst judgement: score definitions for each criterion in both prioritization steps, fixed rules for combining the individual scores into a priority, and secondary criteria for the case in which several clusters receive the same priority (RQ4). Third, it shows the **integration of Cyber Threat Intelligence with assessment planning** by carrying specific threat reporting through to the factor scores of a cluster and then to the key questions of the assessment that was executed (RQ3). Fourth, it defines the **end-to-end traceability** and the record templates behind it, connecting an ATT&CK technique to a cluster, a cluster to an assessment, an assessment to a finding, and a finding to an owned and verified improvement item (RQ5). Fifth, it applies this to **Initial Access against internet-exposed assets** (RQ1), a scope for which the established frameworks are either broader, as with full-scope red-team engagements, or oriented toward evaluating security products rather than the organization's own posture. The complete lifecycle is then demonstrated in a large enterprise environment, together with the measurement definitions used to express perimeter discoverability and inventory completeness.

1.1 Objectives

The main objective is to develop and demonstrate a structured MITRE ATT&CK-based threat assessment methodology that enables organizations to identify, group, prioritize, and assess adversarial techniques relevant to their security posture. Although the methodology is demonstrated on Initial Access threats targeting internet-

exposed enterprise assets, its workflow is designed as a repeatable assessment process that can be adapted to other ATT&CK tactics or assessment domains according to organizational needs.

To achieve this objective, a relevant subset of MITRE ATT&CK techniques and sub-techniques is first defined according to the selected assessment scope. These techniques are then organized into logical threat clusters based on their operational similarities, attack objectives, and defensive relevance. Cyber Threat Intelligence is integrated into the methodology to support prioritization of the most relevant and probable threat clusters. The prioritized clusters are subsequently transformed into specific security assessments, executed in a real enterprise environment, and interpreted to identify strengths, weaknesses, and improvement opportunities.

The expected result is a usable methodology with clearly defined phases, inputs, outputs, and transition logic between analysis, practical assessment, and improvement planning. The enterprise demonstration shows that the proposed approach can be applied under real enterprise conditions and can produce specific, traceable, and practical outputs relevant to the assessed organization.

1.2 Research questions

The development of the methodology is guided by a set of research questions that address the identification, organization, prioritization, and assessment of adversarial techniques targeting internet-facing enterprise assets. These questions provide a structured basis for connecting MITRE ATT&CK technique selection, clustering, Cyber Threat Intelligence, assessment prioritization, and practical security evaluation.

The main research question is formulated as follows:

Main RQ: How can an organization systematically identify, prioritize, and assess adversarial techniques targeting its internet-facing perimeter, in alignment with the MITRE ATT&CK framework, to strengthen its security posture and defensive capabilities against Initial Access threats?

To address this main question, the following supporting research questions are defined:

1. **RQ1.** Which MITRE ATT&CK techniques are relevant to Initial Access threats targeting internet-facing enterprise assets?
2. **RQ2.** How can the identified techniques be organized into meaningful and coherent clusters?
3. **RQ3.** In what ways can Cyber Threat Intelligence support the identification of the most relevant and probable attack clusters?
4. **RQ4.** What criteria and methods can be applied to prioritize the identified clusters according to their relevance and risk?
5. **RQ5.** How can the organization's security posture be systematically assessed against the prioritized attack clusters?

2 BACKGROUND

2.1 Evolution of cyberattacks, automation, and the cybercrime-as-a-service model

Cyberattacks have evolved from relatively isolated and technically simple activities into organized, scalable, and service-driven operations. Earlier generations of attacks were often associated with standalone malware, opportunistic exploitation, or direct compromise of individual systems. With the growth of internet connectivity and enterprise digitalization, attackers increasingly shifted toward remote exploitation of networked systems, public-facing applications, exposed authentication services, and cloud-based infrastructure (Admass et al., 2024; Cotrina et al., 2024).

Current cyber operations are commonly multi-stage and combine reconnaissance, automated scanning, exploitation of known or zero-day vulnerabilities, credential abuse, misuse of trusted relationships, and malware-free techniques based on legitimate tools or valid accounts. This evolution has reduced the effectiveness of purely static perimeter defenses and increased the need for proactive, intelligence-driven assessment of exposed assets.

A major factor behind this development is the Cybercrime-as-a-Service model. Malware, exploit kits, stolen credentials, phishing infrastructure, ransomware operations, and compromised access can be rented, purchased, or

reused by actors with different levels of technical sophistication. Initial Access Brokers are especially relevant in this context because they specialize in obtaining access to organizations and selling it to other threat actors (Darktrace, 2024; Lerver, 2025). As a result, Initial Access has become both a technical phase of intrusion and a monetized service within the cybercriminal ecosystem.

For defenders, this means that internet-facing assets are continuously exposed to automated and economically motivated adversarial activity. Organizations are not defending only against individual attackers, but against an ecosystem that systematically searches for vulnerable services, weak authentication mechanisms, misconfigured systems, and exposed credentials.

2.2 Initial Access as a critical phase of the attack lifecycle

Initial Access represents the point at which external adversarial activity becomes the first successful access in a target environment. In MITRE ATT&CK, Initial Access is represented as tactic TA0001 and includes techniques such as Exploit Public-Facing Application, External Remote Services, Valid Accounts, and other methods that enable attackers to enter an environment (MITRE Corporation, 2025). In the context of internet-exposed assets, Initial Access is achieved when an attacker successfully abuses an exposed application, service, interface, or credential to gain entry from outside the organization.

This phase is strategically important because it determines whether an attempted intrusion remains external or develops into a broader security incident. Once access is established, the attacker may continue with execution, privilege escalation, persistence, lateral movement, credential access, data exfiltration, or ransomware deployment. Preventing or detecting Initial Access is therefore one of the most effective opportunities to disrupt the intrusion lifecycle before deeper compromise occurs.

Threat reporting from 2024–2025 confirms the continued importance of Initial Access vectors. The Verizon Data Breach Investigations Report 2025 states that exploitation of vulnerabilities accounted for approximately 20% of breaches as initial access vector, with a 34% year-over-year increase driven in part by zero-day exploitation of edge devices and VPNs (Verizon, 2025). CrowdStrike's Global Threat Report 2025 highlights that 52% of vulnerabilities observed in 2024 were related to Initial Access and that access broker advertisements increased by 50% compared with 2023 (CrowdStrike, 2025). IBM X-Force further reports that many high-impact vulnerabilities had public exploit code available or were actively exploited, with 60% of IBM's top ten vulnerabilities having public exploit availability less than two weeks after disclosure (IBM, 2025).

The relevance of this problem is even stronger for the internet-facing perimeter. Credential abuse of exposed services, exploitation of vulnerable public-facing systems, and the expansion of externally accessible infrastructure through cloud, SaaS, and remote work all increase the likelihood of successful Initial Access. This makes the assessment of internet-exposed assets a critical part of enterprise defensive strategy.

2.3 Gap between MITRE ATT&CK theory and enterprise assessment practice

MITRE ATT&CK provides a structured and widely adopted knowledge base for describing adversarial tactics, techniques, and procedures (Al-Sada et al., 2024; MITRE Corporation, 2025). It enables consistent terminology for threat analysis, detection engineering, adversary emulation, and mapping of defensive controls (Bonhard et al., 2025). However, the existence of such a taxonomy does not automatically solve the practical problem of how an enterprise should assess its own security posture against relevant techniques.

The central challenge is the gap between theoretical identification of ATT&CK techniques and their practical application in a specific organizational context. Enterprises may understand that techniques such as Exploit Public-Facing Application, External Remote Services, Valid Accounts or Brute Force are generally relevant. However, they still need a systematic way to determine which techniques apply to their internet-facing perimeter, how these techniques should be grouped into meaningful assessment areas, which of them should be prioritized based on current Cyber Threat Intelligence (CTI), and how they should be transformed into executable assessments.

This gap is especially visible because perimeter security activities are often distributed across several functions, including vulnerability management, attack surface management, CTI, penetration testing, red teaming, identity

security, and monitoring. Without a structured transition between these activities, organizations may struggle to connect adversarial behavior, technical exposure, assessment design, and improvement planning.

Therefore, this study proposes a methodology that applies MITRE ATT&CK in practice for the assessment of Initial Access exposure at the internet-facing perimeter. Instead of using ATT&CK only as a descriptive taxonomy, the proposed approach uses it as a basis for selecting externally relevant techniques, organizing them into logical clusters, prioritizing them through CTI, designing practical assessments, and deriving improvement actions from the results.

3 RELATED WORK

Existing approaches provide several ways to apply adversary behavior, threat intelligence, and security testing, but they differ from the methodology proposed in this study in terms of scope, purpose, and level of abstraction. VECTR is a purple-teaming and adversary-emulation platform used to document, track, reproduce, score, and improve adversary simulations (Security Risk Advisors, 2025). It supports MITRE ATT&CK alignment, defense success metrics, reporting, and tracking of red-team and blue-team testing activities. However, VECTR is primarily focused on supporting and measuring purple-team exercises (Security Risk Advisors, 2025). In contrast, the proposed methodology covers the full assessment lifecycle: selecting relevant ATT&CK techniques, grouping them into assessment-oriented threat clusters, prioritizing them with Cyber Threat Intelligence, transforming the prioritized clusters into specific security assessments, executing the selected assessments, interpreting their findings, and translating the results into improvement actions. Beyond purple-team platforms and proceduralised red-team frameworks, recent academic work surveys how MITRE ATT&CK has been applied across security domains (Al-Sada et al., 2024); this methodology builds on that body of work by translating ATT&CK into a repeatable, enterprise-oriented assessment lifecycle.

Threat-intelligence-led red-teaming frameworks, such as TIBER-EU and CBEST, provide a more formal approach to testing cyber resilience. TIBER-EU is a European framework for threat-intelligence-based ethical red teaming (European Central Bank, 2024), while CBEST provides a structured assessment model for financial-sector firms and financial market infrastructures (Bank of England, 2024). The TIBER-EU framework is implemented at national level by individual central banks; the Czech implementation, TIBER-CZ, is operated by the Czech National Bank for entities that are critical to the Czech financial system (Czech National Bank, 2025). These frameworks define roles, phases, deliverables, threat intelligence requirements, red-team execution, and remediation activities. They are highly mature, but they are broader, more formal, and typically oriented toward full-scope red-team engagements. The proposed methodology is narrower in operational use, but more modular in analytical design: it can be applied to a selected ATT&CK tactic or assessment domain, such as Initial Access, Credential Access, Execution, or other tactics, depending on the organization's priorities and assessment scope.

MITRE's Center for Threat-Informed Defense and ATT&CK Evaluations also represent important related work. Threat-Informed Defense emphasizes the systematic application of knowledge about attacker methods and techniques to improve defensive capabilities (MITRE Corporation, n.d.-b), while ATT&CK Evaluations use transparent, threat-informed purple-team methodology to evaluate cybersecurity solutions in the context of ATT&CK (MITRE Corporation, n.d.-a). These initiatives support the broader adoption of ATT&CK in security practice, but their primary orientation is different. CTID provides community-driven resources and guidance for threat-informed defense (MITRE Corporation, n.d.-b), whereas ATT&CK Evaluations focus mainly on assessing security products or solution capabilities (MITRE Corporation, n.d.-a). The proposed methodology instead provides a repeatable process for internal organizational assessment: it helps security teams move from ATT&CK techniques to clusters, from clusters to priorities, from priorities to executable assessments, and from assessment findings to improvement actions.

ATT&CK is also applied inside operational security products. Network detection and response tools such as GREYCORTEX Mendel align the activity they detect with ATT&CK tactics and techniques (GREYCORTEX, 2025), and the open-source security platform Wazuh maps its detection rules and alerts to ATT&CK identifiers and presents coverage by tactic and technique (Wazuh, n.d.). These products apply the framework on the detection side, after activity has already reached the environment monitored. The methodology proposed here applies it earlier, to

decide which techniques should be assessed, in what order, and by what means, which is a complementary use of the same knowledge base.

On the academic side, recent work has targeted the automated extraction and mapping of ATT&CK tactics and techniques from unstructured cyber threat intelligence using transformer-based multi-label classifiers (Li et al., 2024). Such methods address the upstream data-processing side of CTI; the methodology proposed in this study addresses the complementary downstream side, namely how prioritized threat clusters are turned into executable security assessments and traceable improvement actions.

4 MATERIALS AND METHODS

This study follows a design science research approach, in which the artifact is the proposed four-phase threat assessment methodology and the enterprise demonstration serves as its evaluation (Hevner et al., 2004; Peffers et al., 2007). Following the evaluation logic of design science research, the enterprise application reported below is a demonstration of feasibility. Its purpose is to show that the artifact can be executed end to end under real conditions and that each phase produces the input the next phase requires. It is not an effectiveness study, and it does not support conclusions about how the methodology would perform across organizations.

The proposed methodology is designed as a structured ATT&CK-based assessment lifecycle that transforms adversarial techniques into practical security assessments and improvement actions. The methodology is generic in design and can be adapted to different MITRE ATT&CK tactics or assessment domains. For demonstration purposes, it is demonstrated on Initial Access threats targeting internet-exposed enterprise assets.

The methodology consists of four sequential phases:

1. Baseline analysis, clustering, and preliminary prioritization
2. Cyber Threat Intelligence analysis and assessment prioritization
3. Assessment design and execution
4. Improvement plan creation

Each phase produces outputs that serve as inputs for the following phase. This structure supports traceability from the original ATT&CK techniques to threat clusters, from clusters to prioritized assessments, from assessments to findings, and from findings to improvement actions. The methodology and its iterative nature are represented in Figure 1.

Four terms are used throughout the description of the phases and are defined here, because they name records that the methodology produces and are not concepts taken from ATT&CK. A *cluster* is a group of techniques that are assessed together because they share the properties listed in Table 2. A *cluster category* is the set of clusters that belong to the same ATT&CK tactic within the assessment scope, named from the defender's point of view. A *preliminary priority* is the relevance verdict assigned to a cluster in Phase 1 on the basis of open technical analysis, and it determines which clusters proceed to CTI enrichment. An *assessment priority* is the priority assigned to a cluster in Phase 2 on the basis of CTI-supported analysis, and it determines which clusters proceed to execution. The two priorities are separate attributes of the same cluster and are not required to agree: the second is produced from different evidence and may raise or lower the first.

4.1 Assessment scope

The demonstration scope focuses on adversarial techniques that can be applied against internet-exposed enterprise assets. Internet-exposed assets refer to systems and services that are reachable from the public internet and contribute to the external attack surface of an organization (Husák & Sadlek, 2025). Such assets may include web servers, web applications, public APIs, VPN gateways, remote access services, email infrastructure, cloud interfaces, authentication portals, and externally reachable management or administrative services.

The demonstration scope includes selected techniques from the Reconnaissance, Initial Access, and Credential Access tactics of MITRE ATT&CK (MITRE Corporation, 2025). Reconnaissance techniques are included because they support or directly enable exploitation of exposed assets. Selected Credential Access techniques are included when they can be performed externally against exposed authentication interfaces and can directly support the

establishment of initial access. Techniques related to phishing, social engineering, internal lateral movement, persistence, command and control, or other post-exploitation activities are excluded from the demonstration scope.

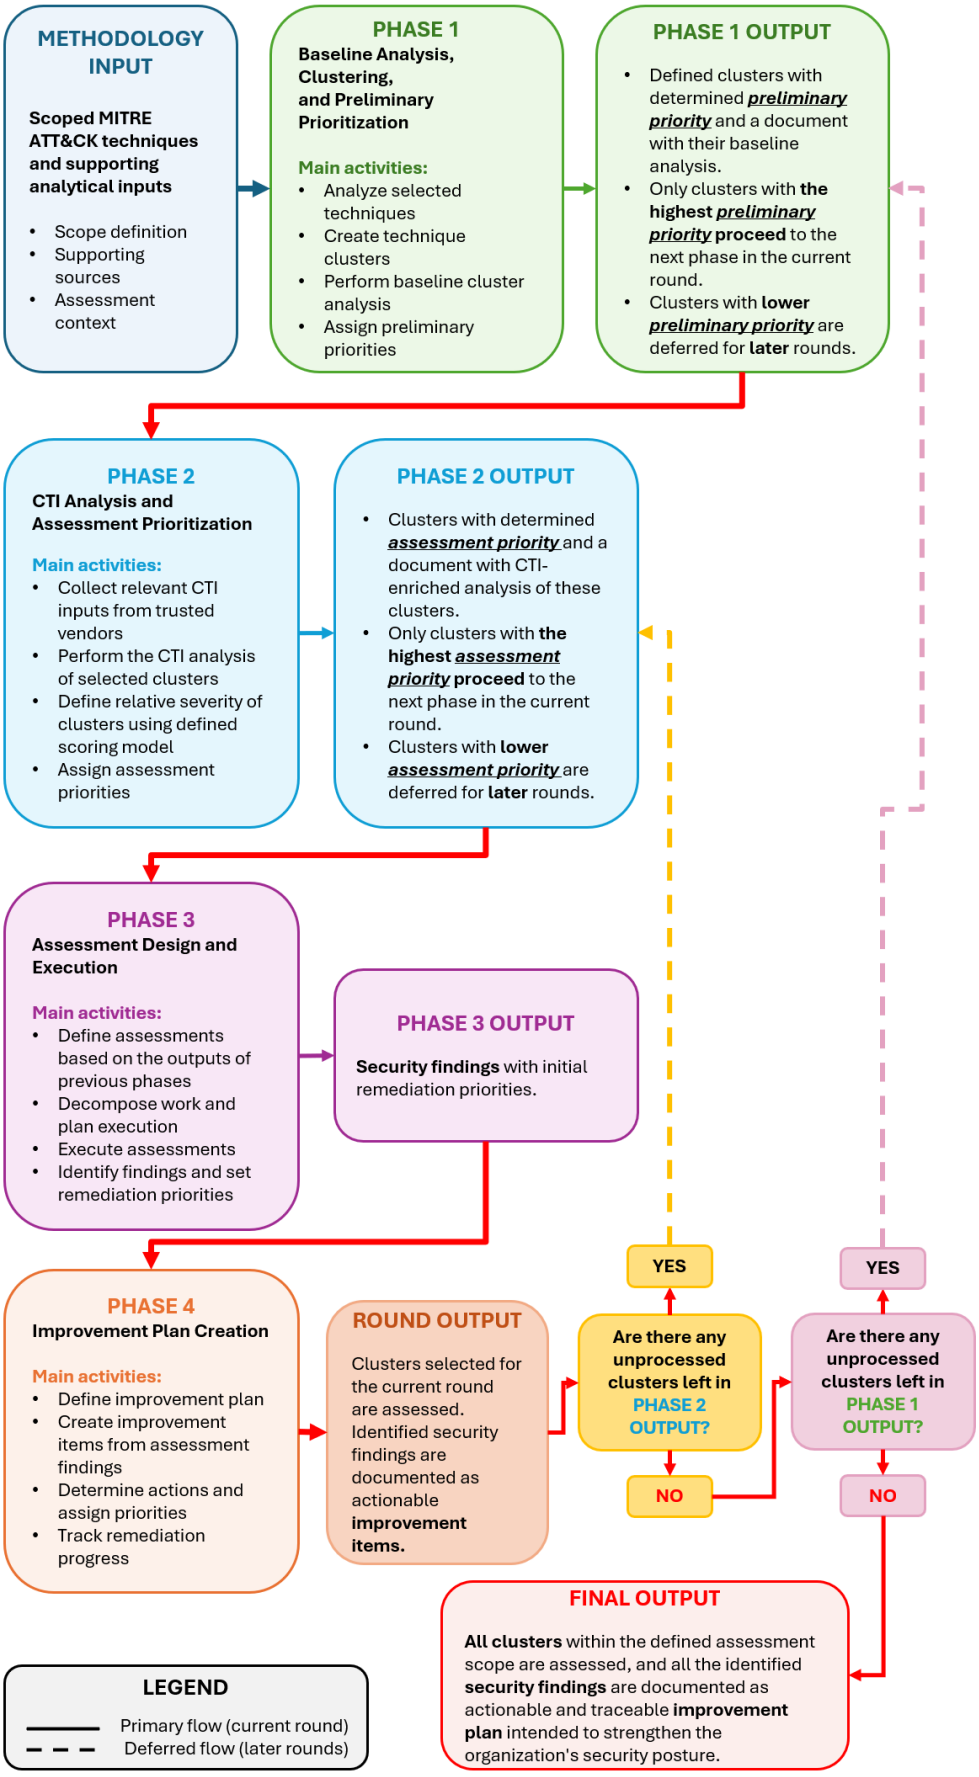


Figure 1. *Diagram of the proposed Threat Assessment Methodology.*

The selection of ATT&CK techniques and sub-techniques for the demonstration scope was performed based on four inclusion criteria:

- **External applicability:** The technique can be executed remotely against internet-exposed systems, services, or authentication interfaces from outside the organization.
- **Initial Access relevance:** The technique directly establishes, supports, or enables the first unauthorized access in the target environment.
- **Assessment suitability**: The technique is observable, interpretable, and assessable from the defender's perspective in the context of enterprise perimeter security.
- **Scope consistency**: The technique does not primarily belong to excluded categories such as social engineering, user-driven vectors, persistence, lateral movement, command and control, or other post-exploitation activities.

The technique selection was based on MITRE ATT&CK Enterprise v18.1, which was the version in force when the scope was defined and the demonstration was carried out in November 2025 (MITRE Corporation, 2025). Version 19 was released in April 2026, after the demonstration had been completed, and was therefore not used. Because the methodology operates on techniques and their identifiers rather than on version-specific content, it can be re-executed against a later version without a change to the process.

Based on these criteria, the demonstration scope includes the following techniques from the Reconnaissance, Initial Access, and Credential Access tactics represented in Table 1.

Table 1. *Scope of MITRE ATT&CK TTPs. Source: Adapted from MITRE Corporation (2025).*

Tactic (ID)	Technique (ID)	Sub-technique (ID)
Reconnaissance (TA0043)	Active Scanning (T1595)	Scanning IP Blocks (T1595.001)
		Vulnerability Scanning (T1595.002)
		Wordlist Scanning (T1595.003)
	Gather Victim Org Information (T1591)	Business Relationships (T1591.002)
		Identify Roles (T1591.004)
	Search Open Websites/Domains (T1593)	Search Engines (T1593.002)
		Code Repositories (T1593.003)
	Search Victim-Owned Websites (T1594)	--
	Gather Victim Host Information (T1592)	Software (T1592.002)
		Client Configurations (T1592.004)
	Gather Victim Identity Information (T1589)	Credentials (T1589.001)
		Email Addresses (T1589.002)
		Employee Names (T1589.003)
	Gather Victim Network Information (T1590)	Domain Properties (T1590.001)
		DNS (T1590.002)
		Network Trust Dependencies (T1590.003)
		Network Topology (T1590.004)
		IP Addresses (T1590.005)
		Network Security Appliances (T1590.006)
	Search Open Technical Databases (T1596)	DNS/Passive DNS (T1596.001)
		WHOIS (T1596.002)
		Digital Certificates (T1596.003)
		CDNs (T1596.004)
		Scan Databases (T1596.005)
Initial Access (TA0001)	Content Injection (T1659)	--
	Exploit Public-Facing Application (T1190)	--

Tactic (ID)	Technique (ID)	Sub-technique (ID)
	External Remote Services (T1133)	--
	Supply Chain Compromise (T1195)	Compromise Software Dependencies and Development Tools (T1195.001)
		Compromise Software Supply Chain (T1195.002)
	Trusted Relationship (T1199)	--
	Valid Accounts (T1078)	Default Accounts (T1078.001)
		Domain Accounts (T1078.002)
		Local Accounts (T1078.003)
		Cloud Accounts (T1078.004)
Credential Access (TA0006)	Exploitation for Credential Access (T1212)	--
	Forge Web Credentials (T1606)	Web Cookies (T1606.001)
		SAML Tokens (T1606.002)
	Brute Force (T1110)	Password Guessing (T1110.001)
		Password Cracking (T1110.002)
		Password Spraying (T1110.003)
		Credential Stuffing (T1110.004)

4.2 Phase 1: Baseline analysis, clustering, and preliminary prioritization

The first phase establishes the analytical baseline for the assessment. A key activity in this phase is clustering. Instead of assessing each technique individually, related techniques are grouped into threat clusters. A cluster represents a logical group of techniques that share functional, operational, or contextual similarity.

Clustering is carried out on two levels. The upper level consists of *cluster categories*, which are taken directly from the ATT&CK tactics included in the assessment scope. In the demonstration scope, Reconnaissance (TA0043) gives the category Reconnaissance and Discovery, Initial Access (TA0001) gives Exploitation and Execution, and Credential Access (TA0006) gives Credential Access and Abuse. The categories therefore inherit the structure of the framework and do not introduce a competing classification, which keeps the result aligned with ATT&CK and makes the top level of the grouping reproducible by construction.

The lower level consists of the *clusters* themselves. Clusters are formed within each category by grouping the selected techniques and sub-techniques according to the criteria listed in Table 2. Cluster names are descriptive labels chosen to express the shared attack objective and target surface of the grouped techniques; they are defender-facing labels and are not ATT&CK terms.

The criteria are applied in combination rather than in isolation, because the meaning of each criterion depends on the assessment scope. For that reason, Table 2 states, for each criterion, both its role in grouping and its interpretation in the context of an Initial Access assessment. Where a technique relates to more than one cluster, it is assigned to the cluster in which its dominant role in the given assessment context is most clearly expressed. This rule prevents the same technique from appearing in several clusters and keeps the later phases analytically clear.

Clustering is an analytical step performed by people, not an automated procedure, so the methodology prescribes how it is conducted and recorded. An initial grouping is proposed by an analyst within the assessment team, using the documented outputs of the organization's threat intelligence and security assessment functions as input, and is then reviewed by the assessment team as a whole. In this study the initial grouping was proposed by the analyst who carried out the assessment; during the review that followed, techniques that did not match the definition of the cluster to which they had been assigned were removed. The procedure is therefore neither a single analyst working alone nor an automated rule-based mapping: the criteria in Table 2 constrain the grouping, the analyst applies them to the selected techniques, and the team review corrects the result. The resulting clusters, the criteria applied to each of them, and the per-cluster analysis are recorded in the baseline cluster analysis, which is the documented output of this phase and the basis on which the grouping can be reproduced, audited, or revised in a later iteration.

Table 2. *Criteria for cluster creation.*

Criterion	Use in grouping	Meaning in the Initial Access context
Attack objective	Techniques with the same immediate objective are grouped together even if their technical implementation differs.	The immediate attacker goal, such as external discovery, credential abuse, exploitation of exposed services, or abuse of a trust relationship.
Attack vector	Techniques targeting the same exposure channel are grouped together to reflect similar defender-facing risk.	The path through which the technique reaches the target, for example web applications, remote access services, exposed authentication interfaces, cloud control planes, or third-party access paths.
Operational similarity	Techniques with comparable execution logic are grouped together to reflect similar attacker behavior.	Similarity in how the technique is carried out, including reconnaissance behavior, credential abuse, exploitation workflow, or service interaction pattern.
Target surface and asset type	Techniques affecting comparable classes of exposed assets are grouped together.	The class of internet-exposed asset the technique is typically used against, such as VPN gateways, web applications, APIs, identity services, email infrastructure, or cloud services.
Type of weakness abused	Techniques abusing the same general class of weakness are grouped together to support later assessment design.	The dominant weakness the technique relies on, such as a software vulnerability, a misconfiguration, weak authentication, credential reuse, or exposed information.
Role in the Initial Access process	Techniques with a similar role in the progression toward a foothold are grouped together.	Whether the technique is mainly enabling, preparatory, or directly compromising within the Initial Access context.

The created clusters of ATT&CK techniques and sub-techniques are analyzed according to their tactical purpose, operational characteristics, exploited attack vectors, expected target surface, availability of public tooling or guidance, and relevance to the selected assessment scope.

For each cluster, the analysis captures the associated ATT&CK techniques, attack phase, exploited attack vectors, operational logic, representative real-world cases, detection opportunities, prevention or mitigation opportunities, availability of public guidance or tooling, and commonness of the target surface. This creates a structured baseline that supports both preliminary prioritization and later assessment design.

Preliminary prioritization is performed to determine which clusters proceed to deeper CTI analysis first. Its result is not a final assessment priority but a relevance verdict that expresses how far a cluster appears worth analyzing within the current assessment context. The verdict rests on five criteria, each scored High, Medium, or Low according to the definitions in Table 3.

Table 3. *Criteria and score definitions for preliminary prioritization.*

Criterion	Meaning in the Initial Access context	Score definitions
Observed real-world abuse	Whether the techniques in the cluster are publicly documented in incidents or campaigns.	**High**: repeatedly observed in real-world attacks. **Medium**: some documented use but limited or less consistent. **Low**: little or no publicly documented evidence of use.
Availability of public exploitation guidance or tooling	Whether public exploit code, proof-of-concept material, offensive tools, or technical walkthroughs exist.	**High**: public tooling or practical guidance is readily available. **Medium**: partial or limited material exists. **Low**: little or no public exploitation guidance is available.
Commonness of the target surface	How common the targeted asset type, service, or exposure pattern is in an enterprise environment.	**High**: a widely deployed or commonly exposed surface. **Medium**: a moderately common or sector-specific surface. **Low**: a niche or rarely exposed surface.
Direct relevance to the assessed domain	How far the cluster contributes to the attack objective of the assessed domain, here Initial Access.	**High**: contributes directly to the assessed objective. **Medium**: supports it but requires additional conditions or complementary techniques. **Low**: relates only indirectly or serves a mainly preparatory or supporting role.

Criterion	Meaning in the Initial Access context	Score definitions
Potential impact	The likely security impact of successful exploitation, in terms of unauthorized access, exposure of sensitive systems or data, and enablement of further attacker activity.	**High**: may directly enable significant compromise of exposed assets, sensitive access paths, or critical business services. **Medium**: gives a meaningful attacker advantage, but with more limited reach or under additional conditions. **Low**: limited immediate impact, or mainly preparatory value.

The five criteria are read together rather than independently, and two of them act as qualifying conditions. Direct relevance to the assessed domain and potential impact decide whether a cluster can reach a high verdict at all, because a cluster that neither contributes to the assessed objective nor produces impact is not a useful first assessment target regardless of how widely its techniques are used. The verdict is assigned as follows:

- **Low** when both qualifying criteria are scored Low, or when at least three of the five criteria are scored Low.
- **High** when both qualifying criteria are at least Medium and at least one of them is High, and in addition either four or more of the five criteria are High, or three are High and the remaining two are Medium.
- **Medium** in all other combinations.

The rule is stated in this order because the Low condition overrides the others: a cluster whose techniques are widely used, well documented, and supported by public tooling still receives a Low preliminary priority if it neither contributes directly to the assessed objective nor produces impact on its own. This is what separates the reconnaissance clusters from the rest in the demonstration reported below. The result of this phase is a set of defined threat clusters with assigned preliminary priorities.

4.3 Phase 2: CTI analysis and assessment prioritization

The second phase enriches the selected clusters with Cyber Threat Intelligence, which supports evidence-based interpretation of adversary behavior and improves prioritization of relevant threat scenarios (Ainslie et al., 2023; Shin & Lowry, 2020). The purpose is to determine which clusters are most relevant, probable, and impactful in the current range of threats. Compared with Phase 1, which relies mainly on general technical analysis, Phase 2 provides a more contextual view based on observed adversary behavior, threat reporting, campaigns, exploitation trends, and known threat actor adoption.

CTI sources are selected according to relevance, credibility, recency, and level of technical detail. Preference is given to reputable threat intelligence vendors and established intelligence-oriented sources, especially materials that describe observed campaigns, threat actor usage, operational details, targeted sectors, prevalence indicators, or technique-specific statistics. Depending on the assessment context, the CTI input may be general, sector-specific, or organization-specific, for example when telemetry or contextual reports from an enterprise threat intelligence platform are available. This selection logic ensures that prioritization is based on realistic adversarial behavior rather than only on theoretical technique descriptions.

The CTI analysis template contains several fields: **cluster name**, **associated ATT&CK techniques**, **attack phase, prevalence and frequency**, **complexity and requirements**, **risk and impact evaluation**, **assessment prioritization scoring**, **detection and mitigation recommendations**, and **technical exploitation details**. The scoring fields are used to determine the final assessment priority of each cluster.

Assessment prioritization is performed based on three qualitative scoring factors:

1. **Prevalence and Frequency**: Evaluates how commonly the cluster appears in recent threat reporting, campaigns, incidents, or threat actor activity.
2. **Complexity and Requirements:** Evaluates attacker effort, prerequisites, required access conditions, tooling, infrastructure, and technical knowledge.
3. **Risk and Impact Evaluation:** Evaluates the potential consequences of successful abuse, including compromise potential and enablement of further attacker activity.

Overall Assessment Priority: Combines the proposed individual factor scores into a final qualitative priority for assessment planning.

Each factor receives a score of High, Medium, or Low according to the definitions in Table 4. Scores are not assigned as a free judgement: each score is taken from the fields of the CTI analysis template that are designated as scoring fields, and each is recorded together with a written justification and the sources it rests on. The justification is part of the phase output, so a score can be reviewed, challenged, or reproduced by a reader who did not perform the analysis.

Evidence of all four kinds listed for a High score is not required, but a High score is supported by the majority of them, and the specific sources behind each score are recorded in the cluster analysis so that the basis of the score stays visible after the fact.

The direction of the Complexity and Requirements factor is defender-oriented and is stated explicitly because it is the opposite of the attacker-side reading. A High score means that the cluster is easy for an attacker to use, needs few prerequisites, and is supported by widely available tooling or public guidance, all of which raise the risk for the defender. A Low score means that the cluster demands significant effort, specialized knowledge, or uncommon conditions.

Table 4. *Score definitions for the three assessment prioritization factors.*

Factor	High	Medium	Low
Prevalence and Frequency	The cluster appears repeatedly in recent vendor and incident reporting, is subject to observed exploitation in the wild, is used by more than one distinct threat actor group and is relevant to the sector of the assessed organization.	The cluster is observed in practice, but less consistently or less frequently, or the evidence comes from a narrow set of sources.	The cluster appears rarely or only sporadically, or there is limited public evidence of current use.
Complexity and Requirements	The cluster is relatively easy to apply, requires limited prerequisites, or is supported by widely available tooling or guidance.	The cluster requires moderate attacker effort, some prerequisites, or more specific conditions.	The cluster requires significant effort, specialized knowledge, uncommon conditions, or many dependencies.
Risk and Impact Evaluation	Successful exploitation may lead to significant compromise, exposure of sensitive assets, or further attacker activity with serious security consequences.	Successful exploitation may give a meaningful attacker advantage, but with more limited reach or effect.	Successful exploitation is likely to have limited security impact or mainly supporting value.

The three factor scores are combined into an overall assessment priority using the fixed rule in Table 5. The rule is deterministic, so two analysts who assign the same factor scores obtain the same overall priority, and the step from evidence to priority can be checked independently of the analyst who performed it.

Table 5. *Combination of factor scores into the overall assessment priority.*

Combination of individual scores	Overall assessment priority
3 × High	High
2 × High + 1 × Medium	High
2 × High + 1 × Low	Medium
1 × High + 2 × Medium	Medium
1 × High + 1 × Medium + 1 × Low	Medium
1 × High + 2 × Low	Medium
3 × Medium	Medium
2 × Medium + 1 × Low	Medium
1 × Medium + 2 × Low	Low
3 × Low	Low

Because the model has three levels, several clusters can receive the same overall priority. The methodology therefore defines secondary criteria that determine the order in which equally rated clusters enter Phase 3. They are applied in sequence and stop as soon as the order is decided:

1. The number of factors scored High, which favours the cluster that is strong on every factor rather than on two of them.
2. The cluster's Phase 1 score for commonness of the target surface, as defined in Table 3.
3. The cluster's Phase 1 score for availability of public exploitation guidance or tooling, as defined in Table 3.

The second and third criteria reuse two of the preliminary prioritization criteria rather than introducing new ones. Both scores were assigned in Phase 1 and are recorded in the baseline cluster analysis, so breaking a tie requires no new evidence and no new judgement; it uses attributes that the process has already produced. If clusters remain equal after these criteria, the order is agreed by the assessment team and the reason is recorded in the prioritization output. Clusters that are not selected first are not discarded: they stay in the working set and are taken up in the following assessment round. The secondary criteria change the order of work, not the priority itself, so the overall rating of a cluster remains the value produced by Table 5.

4.4 Phase 3: Assessment design and execution

The third phase transforms prioritized clusters into specific security assessments. The objective is to convert CTI-enriched cluster analysis into practical assessment work that can be executed by security teams. Each assessment definition contains the assessment objective, input sources, key questions, expected evidence, execution approach, and interpretation logic.

The assessment design is formulated from the defender's perspective. The key question is not only whether an adversarial technique exists, but whether the organization can discover, prevent, detect, or reduce the relevant attack path in practice. For example, the Network/Service Exploitation cluster can be decomposed into assessments focused on web application vulnerabilities, web application misconfigurations, logic or design flaws, remote-access service weaknesses, exposure to common security scanning tools, and abuse of integrated command-shell functionality.

The assessment definition template supports consistent execution across different clusters. It ensures that each assessment has a clear purpose, defined input sources, practical key questions, and an expected output. The key questions should be specific enough to guide execution, but broad enough to support different technologies and environments. It includes the following fields:

- **Assessment name:** Name of the assessment derived from the prioritized cluster.
- **Input:** Short description of the problem or defensive capability to be assessed.
- **Input sources:** Documents, CTI outputs, inventories, tools, logs, or other data sources used for execution.
- **Key questions:** Defender-oriented questions that guide the assessment.
- **Execution approach:** Documented practical steps, checks, comparisons, or testing activities performed during the assessment with attached results.
- **Findings:** Strengths, weaknesses, and additional observations identified during execution.
- **Interpretation**: Explanation of the security meaning and relevance of the findings.

Assessment execution may use multiple sources of evidence, including asset inventories, attack surface management platforms, vulnerability management data, CTI reports, scanner outputs, logging and monitoring data, configuration reviews, and controlled testing results. The exact execution method depends on the selected cluster, available data, operational constraints, and organizational risk tolerance.

4.5 Phase 4: Improvement plan creation

The fourth phase converts assessment findings into an improvement plan. The purpose is to ensure that the methodology produces usable outcomes rather than only analytical results. Findings are categorized as strengths, weaknesses, or additional observations. Strengths represent effective controls or positive security conditions. Weaknesses represent gaps that may increase exposure or reduce defensive capability. Additional observations capture relevant findings that do not clearly fall into either category but may support future assessment or monitoring.

Weaknesses are converted into improvement items. Each improvement item should describe the affected area, the identified issue, the proposed remediation or investigation activity, the expected security benefit, and the priority for implementation. Prioritization can be supported by available resources, dependency on other teams, urgency of the identified risk and an impact-effort matrix represented by Figure 2 (Athuraliya, 2026). The impact-effort matrix is a widely used practitioner heuristic for ranking remediation actions.

The improvement plan creates traceability between the assessment lifecycle and security remediation. Each improvement item can be linked back to an assessment finding, each finding to an executed assessment, each assessment to a prioritized cluster, and each cluster to the original ATT&CK techniques. This traceability supports communication with stakeholders and helps justify why specific remediation activities are prioritized.

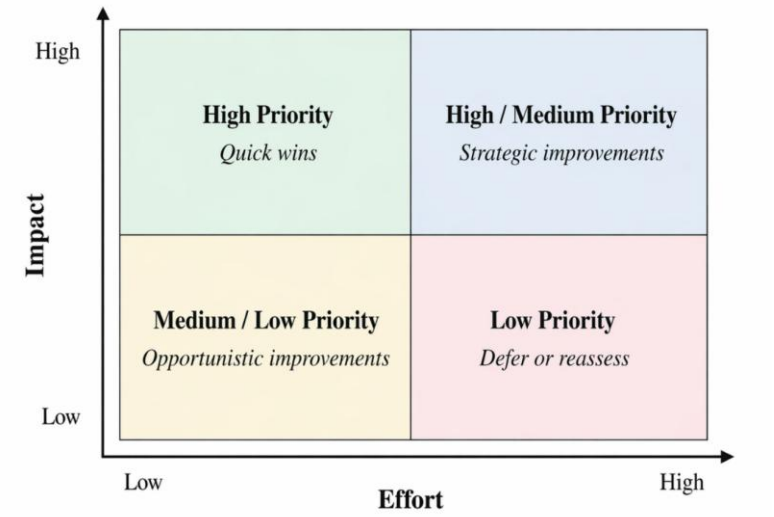


***Figure 2**. Impact-effort matrix for prioritization of improvement items. Source: Adapted from (Athuraliya, 2026).*

The methodology also defines what happens to an improvement item after it has been created, because an improvement plan supports remediation only if its items are owned and tracked. Each item is assigned an owner, which is the team or role that owns the affected asset, tool, or control. Depending on the finding, the owner may sit inside the assessing team, in another security function such as the security operations or detection engineering team, or with the service owner of the affected system. Handover to that owner is a transfer of responsibility and not the end of the process: the item is registered as a work item in the organization's tracking system, linked to the parent improvement plan, and given a status attribute that records its progress. The assessing team follows the item up with the owner until it is closed.

Closure requires verification rather than a statement that the work has been done. Where the original assessment can be repeated, the item is verified by re-executing the relevant part of the assessment and comparing the result with the recorded baseline. Where re-execution is not technically feasible, verification rests on confirmation by the owning team together with the supporting evidence they provide, and the fact that re-testing was not possible is recorded. Items that remain open, or that are only partly resolved, are carried into the next assessment round. This makes the methodology iterative rather than a single exercise: the improvement plan of one round becomes an input to the baseline of the next.

5 RESULTS

5.1 Phase 1: Baseline analysis, clustering, and preliminary prioritization

The first phase produced a clustering of the selected MITRE ATT&CK techniques and sub-techniques relevant to the demonstration scope. Instead of treating each technique as an isolated assessment object, the techniques were grouped into broader threat clusters according to their operational similarity, attack objective, and relevance from the defender's perspective. The resulting clusters are presented in Table 6. The grouping followed the two-level

derivation defined in Methods. The three cluster categories correspond to the three ATT&CK tactics in the demonstration scope, and the ten clusters were formed inside those categories using the criteria in Table 2.

Two examples show how the criteria decided the result. Exploit Public-Facing Application (T1190) and External Remote Services (T1133) were placed in the same cluster, Network/Service Exploitation, because they share the attack objective of obtaining access through an exposed service, they act on the same class of target surface, and they rely on the same classes of weakness, namely unpatched software and weak service configuration, even though their technical execution differs. In the opposite direction, the reconnaissance techniques were separated into four clusters instead of one. They share a preparatory role in the Initial Access process, which alone would place them together, but they differ in attack vector and target surface: collection of host, identity, and network information about the target; collection of information about the organization, its roles, and its business relationships; collection from open technical databases, search engines, and victim-owned websites; and active scanning together with searches of public code repositories. Each of these is assessed by different means and against different evidence, so the criteria of attack vector and target surface were given precedence over the shared role.

The dominant-role rule was applied where a technique could be argued into more than one cluster. Searching public code repositories (T1593.003), for example, is a search of an open website by attack vector, but its dominant role in this assessment context is the discovery of technical exposure, so it was assigned to Technical Footprinting and Vulnerability Discovery rather than to Open-Source Intelligence.

Table 6. *Technique clustering for Initial Access assessment.*

Cluster Category	Cluster	Technique	Technique MITRE ID
Reconnaissance and Discovery: Supporting activities used to directly enable initial access clusters/techniques.	**Identity & Network Mapping**	Gather Victim Host/Identity Information	T1592 / T1589
		Gather Victim Network Information	T1590
	Organizational Insight	Business Relationships	T1591.002
		Identify Roles	T1591.004
	Open-Source Intelligence (OSINT)	Search Open Technical Databases	T1596
		Search Engines	T1593.002
		Search Victim-Owned Websites	T1594
	Technical Footprinting & Vulnerability Discovery	Code Repositories	T1593.003
		Scanning IP Blocks	T1595.001
		Vulnerability Scanning	T1595.002
		Wordlist Scanning	T1595.003
Exploitation and execution: Direct exploitation of exposed systems or trusted mechanisms and execution of appropriate activities to directly gain initial access.	**Network/Service Exploitation**	Exploit Public-Facing Application	T1190
		External Remote Services	T1133
	Injection-based Exploitation	Content Injection	T1659
	Supply Chain & Trust Abuse	Trusted Relationship	T1199
		Compromise Software Supply Chain	T1195.002
		Compromise Software Dependencies and Development Tools	T1195.001
Credential Access and Abuse: Activities related to harvesting, brute-forcing, or misusing credentials to gain initial access.	**Brute Force & Exploitation**	Exploitation for Credential Access	T1212
		Brute Force	T1110
	Token & Session Hijacking	Web Cookies	T1606.001
		SAML Tokens	T1606.002
	Valid Accounts Abuse	Valid Accounts	T1078

Note: Technique names and IDs follow MITRE ATT&CK Enterprise v18.1 (MITRE Corporation, 2025).

The preliminary prioritization results are shown in Table 7. The reconnaissance-oriented clusters were assigned **Low** priority because they mainly support direct attacks and usually do not directly produce compromise by themselves. However, they remain important as enabling activities because they help attackers identify exposed assets, vulnerable services, authentication interfaces, and organizational context.

The **highest preliminary priority** was assigned to **Network/Service Exploitation**, **Brute Force & Exploitation**, and **Valid Accounts Abuse**. These clusters were prioritized because they represent direct or highly probable paths to unauthorized access through exposed applications, remote services, or authentication mechanisms. **Medium** priority was assigned to **Injection-based Exploitation**, **Supply Chain & Trust Abuse**, and **Token & Session Hijacking**, as these clusters may produce significant impact but are more dependent on specific technical or contextual conditions.

Table 7. *Preliminary priority of technique clusters.*

Cluster Category	Cluster	Preliminary priority
Reconnaissance and Discovery	Identity & Network Mapping	LOW
	Organizational Insight	LOW
	Open-Source Intelligence (OSINT)	LOW
	Technical Footprinting & Vulnerability Discovery	LOW
Exploitation and execution	**Network/Service Exploitation**	**HIGH**
	Injection-based Exploitation	MEDIUM
	Supply Chain & Trust Abuse	MEDIUM
Credential Access and Abuse	**Brute Force & Exploitation**	**HIGH**
	Token & Session Hijacking	MEDIUM
	Valid Accounts Abuse	**HIGH**

Overall, Phase 1 reduced the original technique-level scope to a smaller set of assessment-oriented clusters and identified the clusters that required deeper CTI-supported prioritization.

5.2 Phase 2: CTI analysis and assessment prioritization

The second phase focused on the clusters that received the highest preliminary priority in Phase 1. These clusters were enriched with Cyber Threat Intelligence to evaluate their real-world prevalence, attacker feasibility, and potential impact. The assessment prioritization results are summarized in Table 8.

The way the CTI evidence was used can be followed on the Network/Service Exploitation cluster. The reporting collected for this cluster showed that exploitation of public-facing applications and of external remote services accounted for a large share of the initial access activity observed during 2024, and that exploitation of public-facing applications was the leading initial access vector in the European and North American incident sets (IBM, 2025). The same body of reporting recorded fast weaponization of vulnerabilities in edge devices and network appliances, a rise in access-broker activity, and the use of these techniques by both financially motivated and state-linked groups (CrowdStrike, 2025; Verizon, 2025). Public exploit code and ready-made scanning templates were available for a large part of the vulnerabilities involved.

These observations were mapped to the three factors as follows. The consistent appearance of the cluster across campaigns, regions, and actor types gave a **High** score for Prevalence and Frequency. The availability of public exploit code, scanning templates, and access-as-a-service offerings, which lowers the effort needed to use the cluster, gave a **High** score for Complexity and Requirements in the defender-oriented direction defined in Table 4. The ability to reach core infrastructure directly from the internet, and the credential theft, persistence, and payload deployment that follow, gave a **High** score for Risk and Impact Evaluation.

The same evidence then shaped the assessment design in Phase 3. Because the reporting showed that exploitation activity is driven by systematic scanning of internet-exposed assets, and because an asset that the organization does not know about can be neither patched nor monitored, the leading key question for this cluster was formulated as a question about discoverability: how much of the perimeter can an external observer identify using commonly available reconnaissance, and how complete is the internal inventory when compared with that external view. That question defined the sub-assessment that was executed and reported below. It links a CTI observation to a specific piece of assessment work.

All three selected clusters received an overall **High** assessment priority. **Network/Service Exploitation** received High scores across all three factors. The CTI analysis also showed that these techniques are actively used by both financially motivated and state-linked actors and are supported by public exploit availability, rapid weaponization of vulnerabilities, and broad exposure of enterprise edge infrastructure.

Table 8. *Assessment prioritization of selected clusters.*

Cluster	Prevalence & Frequency	Complexity & Requirements	Risk & Impact Evaluation	Overall assessment priority
Network/Service Exploitation	High	High	High	**HIGH**
Brute Force & Exploitation	High	Medium	High	**HIGH**
Valid Accounts Abuse	High	Medium	High	**HIGH**

Because all three clusters reached the same overall priority, the secondary criteria defined in Phase 2 were applied. The first criterion was sufficient. Network/Service Exploitation was scored High on all three factors, while Brute Force & Exploitation and Valid Accounts Abuse were scored Medium on Complexity and Requirements, so Network/Service Exploitation entered Phase 3 first and the other two clusters remained in the working set for the following round. The outcome is consistent with the second criterion as well: the commonness of the target surface recorded for this cluster in Phase 1 covers public-facing applications and remotely accessible services, which in the assessed environment is the broadest exposed surface of the three, so the two criteria point to the same cluster rather than pulling in different directions.

A methodological observation from this phase is that the three-level qualitative scoring model provides a simple and explainable prioritization mechanism, but it has limited granularity when several clusters receive similar final priorities. For example, all three analyzed clusters were classified as High priority, although their operational characteristics differ. The secondary criteria resolve this case and make the resulting order explainable, but they work outside the three-level scale rather than inside it. A five-level or weighted scoring model would express the same differences within the scale itself and would remove the need for a separate ordering step; it remains an option for later iterations of the methodology.

5.3 Phase 3: Assessment design and execution

Based on the CTI-supported prioritization, the **Network/Service Exploitation** cluster was selected for detailed assessment design. This cluster was selected because it received High priority across all scoring factors and represents a direct path to compromise through public-facing applications, exposed remote services, and internet-accessible infrastructure.

The cluster was decomposed into a set of specific assessments listed below. These assessments translate the broader threat cluster into smaller evaluation areas that can be executed by security teams. Defined assessments:

- Exploitation of web application vulnerabilities.
- Exploitation of web application misconfigurations.
- Exploitation of web application logic/design flaws.
- Remote-access services misconfigurations and vulnerabilities exploitation.
- Security posture of Siemens Healthineers perimeter against well-known security scan/testing tools.
- Exploitation of integrated command shells to enable RCE.

As a representative executed sub-assessment, **Discoverability of Siemens Healthineers perimeter** was performed as part of the broader assessment **Exploitation of web application vulnerabilities**. The objective was to determine how easily the externally visible perimeter can be discovered using commercial internal and external visibility sources and selected open-source reconnaissance tools.

For confidentiality reasons, the names of some commercial and internal tools used during the assessment were anonymized in the publication. The anonymization does not affect the interpretation of the results, because the tools are discussed only according to their functional role in the assessment.

The assessment used the **corporate External Attack Surface Management (EASM) solution** as the internal attack surface management inventory source. **Two Third Party Risk Management (TPRM) solutions, TPRM-1** and **TPRM-2**, were used as external commercial reconnaissance sources, while the open-source toolchain included tools such as **subfinder**, **assetfinder**, **amass**, **dnsx**, and **recon-ng** modules. The **brute_suffix** and **brute_hosts recon-ng** modules were not executed because they could take up to 4 weeks to complete on the full scope of assets and could significantly load the assessment server.

Measurement definitions. The sub-assessment compared an internal and an external view of the same perimeter. The corporate EASM solution provided the internal view. The two TPRM platforms and the open-source toolchain provided the external view. The TPRM platforms are a usable proxy for an external observer in this setting because the organization does not share inventory data with them, so their results are produced only from externally observable information.

Let I be the set of assets recorded in the internal EASM inventory and let E_t be the set of assets reported by external source t. An asset is a *deviation* if it is reported by at least one external source and is absent from the internal inventory:

$$D = (\bigcup_t E_t) \setminus I \quad (1)$$

The *potential full scope* P is the union of the internal inventory and all deviations, and serves as the reference set against which both views are measured:

$$P = I \cup D \quad (2)$$

Inventory completeness of the internal source and external visibility of source t are then

$$C = \frac{|I|}{|P|}, \qquad V_t = \frac{|E_t|}{|P|} \quad (3)$$

The relative domain coverage of the open-source toolchain against a commercial source is the ratio $V_{\text{toolchain}}/V_t$, and the share of deviations attributable to the toolchain is $|D \cap E_{\text{toolchain}}|/|D|$, computed for domains.

The comparison was carried out separately for domains, IP addresses, and IP ranges, because the sources differ in how completely they cover each asset type, and separately for two inventory scopes: the full inventory, which includes active, inactive, and semi-related assets, and the active-only inventory. For the active-only scope the deviation set of the full scope was reused, so that assets which are visible internally under a status other than active are not counted as deviations and do not distort the result. Domain inputs from the commercial external sources were first reduced to second-level and top-level domain pairs and deduplicated, and the open-source toolchain was run on that normalized list, so that all sources were compared on the same input basis.

For the active-only IP comparison an additional ratio was used, because that comparison is made against the active internal inventory rather than against the potential full scope:

$$A_t = |E_t|/|I_{active}| \quad (4)$$

where I_{active} is the active-only internal inventory. A value above 100% means that the external source attributes more IP addresses to the organization than the active internal reference set contains. This is referred to as over-attribution: it signals that the source may be assigning cloud, semi-related, or unrelated addresses to the organization, and it requires attribution review before the source is used for perimeter visibility. The ratio is not interpreted as external visibility or coverage, which cannot exceed 100% when computed against the potential full scope. The main quantitative results are summarized in Table 9.

Table 9. *Quantitative results of the perimeter discoverability assessment.*

Metric	Result	Interpretation
Average corporate EASM tool inventory completeness	95%	High internal inventory completeness, but not complete.
Internal ASM tool IP range completeness	77.39%	Weaker completeness for IP ranges compared with the overall average.

Metric	Result	Interpretation
External discoverability by individual tools	Less than 20% for each tool, except TPRM-1 IP ranges at 51.02%	Overall low discoverability of the perimeter from external sources.
Toolchain domain coverage compared with TPRM-1 and TPRM-2	5–6 times higher	Open-source reconnaissance toolchain provided significantly stronger domain discovery.
TPRM-2 active-IP attribution ratio	257.54%	Substantial mismatch between external attribution and the active internal reference set; signals possible over-attribution and requires attribution review.
Share of identified domain deviations found by the open-source tools toolchain	More than 95%	Most domain deviations were discovered through the toolchain-based approach.

The results show that the overall discoverability of the Siemens Healthineers internet-facing perimeter can be considered **LOW**. This is a positive finding because external tools and common reconnaissance sources did not provide broad visibility into the perimeter. At the same time, the assessment identified two important weaknesses. First, the corporate EASM tool inventory was highly complete but still did not cover all potentially relevant exposed assets (~5% were not covered). Second, TPRM-2 produced an active-IP attribution ratio of 257.54%, a substantial mismatch with the active internal reference set that signals possible over-attribution. These findings are summarized in Table 10.

Table 10. *Security findings and initial remediation priorities.*

Finding type	Finding	Description	Severity	Initial remediation priority
Strength	Low discoverability of the perimeter	The perimeter showed low discoverability, with less than 20% visibility for each tool except TPRM-1 IP ranges at 51.02%.	--	--
Weakness	Incomplete corporate EASM inventory	Average inventory completeness was 95%, except IP ranges at 77.39%, meaning that some assets may remain outside monitoring and scanning coverage.	Medium	Medium
Weakness	Weak TPRM-2 IP attribution quality	TPRM-2 produced an active-IP attribution ratio of 257.54%, which signals over-attribution of cloud, semi-related, or unrelated IPs.	Low	Low

5.4 Phase 4: Improvement plan creation

The final phase transformed the weaknesses identified during the executed assessment into a documented improvement plan. The purpose of this phase was to demonstrate that assessment findings can be converted into traceable follow-up records suitable for remediation planning and further handling by responsible teams. Since the methodology was demonstrated on one representative sub-assessment, the improvement plan was intentionally limited in scope and contained only the issues confirmed during that assessment. The created improvement plan is represented in Table 11. It serves as a parent record for the improvement items derived from the identified weaknesses.

Table 11. *Created improvement plan.*

Field	Value
Improvement Plan title	Initial Access – perimeter threats
Improvement Plan owner	Person / Team
Improvement Plan description	Tracking and managing remediation activities related to identified weaknesses in the organization's security posture against Initial Access threats targeting the internet-facing perimeter.

Within this improvement plan, each confirmed weakness was represented by a separate improvement item. The assigned priority of each item followed the initial remediation priority established during Phase 3, while attributes

such as owner, status, and planned actions were specified according to practical remediation needs. The created improvement items are summarized in Table 12.

Table 12. *Created improvement items.*

Field	Improvement Item 1	Improvement Item 2
Item title	EASM deviations	TPRM-2 IP attribution quality
Item owner	Person / Team	Person / Team
Item priority	Medium	Low
Status	Handed over	Handed over
Input	Multiple deviations were identified: assets attributed to Siemens Healthineers but not present in the corporate EASM tool inventory. These were detected by external scoring platforms and the custom reconnaissance toolchain.	TPRM-2 reported a significantly higher number of IPs than the internally maintained active corporate EASM tool inventory, which signals possible over-attribution.
Problem	The current EASM inventory does not fully reflect the externally observable perimeter, creating possible blind spots in monitoring and scanning coverage.	TPRM-2 IP-based visibility contains substantial noise and cannot be used as a fully reliable source without additional validation and filtering.
Proposed actions	Review collected deviations, validate them as true or false positives, enrich the inventory of the corporate EASM tool with confirmed true positives, and reduce the recurrence of similar deviations where possible.	Review TPRM-2 IP attribution logic, define validation or filtering rules, and exclude clearly unrelated IPs from practical perimeter visibility evaluation.

In this way Phase 4 demonstrates that the findings produced by previous phases can be turned into practical and traceable records suitable for planning and performing remediation activities. The resulting improvement plan and respective improvement items therefore represent the final output of the first round of the proposed methodology within the scope of this practical demonstration.

Both items were at the status *handed over* at the time of writing: the deviation review was transferred to the owner of the EASM tool and the attribution issue to the owner of the external platform integration, in both cases as a linked work item in the corporate tracking system. The deviation review is verifiable by re-execution, because repeating the discoverability comparison recomputes completeness against the same reference set and shows directly whether confirmed deviations have been absorbed into the inventory. The attribution issue is not verifiable in the same way, because the filtering logic sits with the external provider, so its closure depends on confirmation and evidence from the owning team. Both items remain open inputs to the next assessment round.

6 DISCUSSION

The results of the demonstration should be read on two levels, and the two levels should be kept separate. At the level of the methodology, the demonstration shows that every phase produced the output that the next phase required, that the scoring and ordering rules were sufficient to select a cluster from a set of equally rated ones, and that the traceability chain held from the selected ATT&CK techniques through clusters and assessment design to findings and tracked improvement items. That is what the enterprise application was designed to test. At the level of the environment, the reported values describe the state of one organization's tooling at one point in time. They show that the method produces measurable and specific results, which is itself a property of the method, but the values themselves are not evidence about how well the method performs in general and should not be read as benchmark figures for other organizations. An organization with a different asset inventory, different external providers, and a different exposure profile would produce different numbers from the same process.

The results demonstrate that the proposed methodology converts MITRE ATT&CK techniques into a systematic assessment workflow with traceable outputs. The process connected technique selection, clustering, CTI-supported prioritization, assessment execution, and improvement planning. This traceability is important because it allows security teams to explain why a specific assessment was selected, what findings it produced, and how those findings were converted into improvement actions.

The clustering performed in Phase 1 reduced the complexity of working with individual ATT&CK techniques. Instead of assessing techniques separately, related techniques were grouped into broader threat clusters that better reflect how security teams usually evaluate attack paths in practice. The CTI-supported prioritization in Phase 2 then helped identify which clusters represented the most relevant and impactful assessment targets. The results also showed the limit of a three-level scale: all selected clusters received an overall High priority, so the scale alone could not order them. The secondary criteria defined in Phase 2 resolved the case, and the first of them was sufficient, but the fact that an additional step was needed is itself a finding about the scoring model. A five-level or weighted model would carry that differentiation inside the scale.

The executed discoverability assessment produced both positive and negative findings. The overall low discoverability of the Siemens Healthineers perimeter represents a defensive strength, as common external sources did not provide broad visibility into the perimeter. At the same time, the corporate EASM inventory was not fully complete, with the EASM inventory near but below full completeness, and IP range coverage much lower (see Results). This shows that even mature attack surface management may still contain visibility gaps. The open-source reconnaissance toolchain also proved useful, providing 5–6 times higher domain coverage than the TPRM platforms and identifying more than 95% of domain deviations. In contrast, the 257.54% active-IP attribution ratio reported by TPRM-2 showed a substantial attribution mismatch and possible over-attribution, confirming that external visibility platforms should be validated before being used as authoritative perimeter sources.

The workflow is useful for security teams, CTI analysts, vulnerability management, attack surface management, and stakeholders responsible for remediation planning.

From an organizational perspective, the methodology supports security governance by giving managers and security teams a traceable line from attacker techniques to prioritized assessments and remediation. Because each improvement item can be traced back to a specific ATT&CK technique and a business-relevant exposure, the approach helps security, vulnerability management, attack surface management, and risk-owning stakeholders agree on priorities and justify resource allocation. It is therefore relevant as both a technical procedure and a decision-support process for managing enterprise exposure.

Set against the approaches reviewed earlier, the difference is one of coverage rather than of subject. Purple-team platforms such as VECTR record, score, and reproduce simulated attacks that have already been chosen. Threat-intelligence-led frameworks such as TIBER-EU, CBEST, and their national implementations prescribe the roles, phases, deliverables, and reporting of full-scope red-team engagements. ATT&CK Evaluations examine how security products respond to a selected set of techniques, and detection products map their own alerts to the framework after the activity has occurred. None of these defines how an organization decides which techniques to assess in the first place, how those techniques should be grouped so that the assessment work becomes manageable, how the grouping is prioritized against current threat reporting, or how the resulting findings are carried through to owned and verified remediation. The methodology proposed here addresses this part of the work, at a scale that a single security team can run without commissioning a full red-team engagement.

Several limitations must be acknowledged. The demonstration was performed in a single enterprise environment and on a limited assessment scope. One cluster of the three prioritized ones was taken into Phase 3, and within one of the six assessments defined for that cluster a single sub-assessment was executed, so the demonstration covers one path through the lifecycle rather than the whole of it. In line with design science research, this enterprise demonstration serves as a proof-of-concept evaluation of the proposed artifact rather than a multi-organization empirical generalization (Hevner et al., 2004; Peffers et al., 2007). The clustering and scoring steps are documented procedures carried out by analysts, so two teams working from the same evidence may still produce different groupings or scores; the criteria and the written justifications make such differences visible and reviewable, but they do not remove them. Multi-organization and multi-tactic evaluations remain a direction for future work. Raw asset data and detailed internal results could not be disclosed due to confidentiality requirements, so only aggregated and anonymized findings are reported. Automation of Phase 2, using ATT&CK-aware classifiers over unstructured CTI (Li et al., 2024), is a natural extension that could reduce manual effort in CTI-driven prioritization.

Despite these limitations, the methodology is transferable because its workflow is not dependent on a specific organization, tool, or ATT&CK tactic. The same process can be adapted to other assessment scopes, different CTI

sources, and other ATT&CK tactics. Future work should validate the methodology across additional organizations, sectors, and tactics.

7 CONCLUSION

The proposed methodology addresses the practical gap between MITRE ATT&CK as a structured taxonomy of adversarial behavior and its application in enterprise security assessment. It provides a repeatable process for selecting relevant ATT&CK techniques, grouping them into meaningful threat clusters, prioritizing them with Cyber Threat Intelligence, transforming them into executable assessments, and converting assessment findings into improvement actions.

The methodology was demonstrated on Initial Access threats targeting internet-exposed enterprise assets. The demonstration showed that the four-phase workflow can be applied under real enterprise conditions and can produce traceable outputs across all phases. The selected techniques were grouped into assessment-oriented clusters, high-priority clusters were identified through CTI-supported prioritization, one representative assessment was executed, and the identified weaknesses were transformed into an improvement plan and corresponding improvement items.

The demonstration produced specific results. The selected sub-assessment showed that the overall discoverability of the Siemens Healthineers internet-facing perimeter can be considered low. At the same time, it identified two specific weaknesses requiring follow-up action: incomplete domain coverage in the inventory of the corporate EASM solution, with approximately 5% of domains missing, and weak IP attribution quality in one of the tools used for external perimeter scanning.

Each research question is addressed within the proposed methodology: RQ1 by the selection of MITRE ATT&CK techniques relevant to internet-exposed assets (Phase 1). RQ2 by the clustering of related techniques into assessment-oriented groups (Phase 1). RQ3 and RQ4 by the CTI-supported prioritization with explicit scoring factors (Phase 2), and RQ5 by the assessment design, execution, and finding-to-improvement traceability (Phases 3 and 4).

This study resulted in the design and practical demonstration of a structured MITRE ATT&CK-based threat assessment methodology for internet-exposed enterprise assets. The performed demonstration showed that the proposed methodology can be applied under real enterprise conditions and can transform the analysis of relevant ATT&CK techniques into prioritized cybersecurity assessments whose execution produces specific security findings, subsequently formalized in practical follow-up plans that support further remediation and strengthening of the enterprise security posture. Establishing how the methodology performs across organizations, sectors, and ATT&CK tactics requires further studies, which is the direction indicated for future work.

ADDITIONAL INFORMATION AND DECLARATIONS

Funding: This work was supported by the Cultural and Educational Grant Agency (KEGA) of the Ministry of Education, Science, Research and Sport of the Slovak Republic under the project No. 004ŽU-4/2024.

Conflict of Interests: O.S. is employed by Siemens Healthineers, whose environment was used for the enterprise demonstration of the proposed methodology. The remaining authors declare no conflict of interest.

Author Contributions: O.S.: Conceptualization, Methodology, Software, Validation, Investigation, Writing – original draft, Visualization. O.K.: Conceptualization, Methodology, Supervision, Writing – review & editing. I.S.: Supervision, Project administration, Resources, Writing – review & editing. M.M.: Methodology, Investigation, Writing – review & editing. P.F.: Methodology, Supervision, Writing – review & editing. E.Š.: Project administration, Funding acquisition, Writing – review & editing.

Statement on the Use of Artificial Intelligence Tools: Artificial intelligence tools (Claude by Anthropic, model versions Sonnet 4.5 / Opus 4) were used by the authors only for language editing, sentence-level rephrasing, and improving the structure and clarity of the manuscript text. All conceptual work, methodology design, validation, analysis, and interpretation of results are the authors' own. The authors take full responsibility for the content of the manuscript, including any AI-assisted text.

Data Availability: The data that support the findings of this study are not publicly available due to confidentiality requirements of the enterprise environment used for the demonstration, only aggregated and anonymized results are reported in this article.

REFERENCES

Admass, W. S., Munaye, Y. Y., & Diro, A. A. (2024). Cyber security: State of the art, challenges and future directions. *Cyber Security and Applications*, 2, 100031. https://doi.org/10.1016/j.csa.2023.100031

Ainslie, S., Thompson, D., Maynard, S., & Ahmad, A. (2023). Cyber-threat intelligence for security decision-making: A review and research agenda for practice. *Computers & Security*, 132, 103352. https://doi.org/10.1016/j.cose.2023.103352

Al-Sada, B., Sadighian, A., & Oligeri, G. (2024). MITRE ATT&CK: State of the art and way forward. *ACM Computing Surveys*, 57(1), Article 12. https://doi.org/10.1145/3687300

Athuraliya, A. (2026). Impact effort matrix: What it is and how to use it effectively. *Creately*. https://creately.com/guides/impact-effort-matrix/

Bank of England. (2024). CBEST threat intelligence-led assessments: Implementation guide for CBEST participants. https://www.bankofengland.co.uk/financial-stability/operational-resilience-of-the-financial-sector/cbest-threat-intelligence-led-assessments-implementation-guide

Bonhard, S. K., Villalta, P. G., Rosés, O., & Pegueroles, J. (2025). A review of tactics, techniques, and procedures (TTPs) of MITRE framework for business email compromise (BEC) attacks. *IEEE Access*, 13, 50761–50776. https://doi.org/10.1109/ACCESS.2025.3552523

Cotrina, L., León, P., Reyes, C., Arbulú Ballesteros, M., Guzmán, M., Arbulú Castillo, J., Acosta, R., & Morales, A. (2024). Cyber crimes: A systematic review of evolution, trends, and research approaches. *Journal of Educational and Social Research*, 14, 96. https://doi.org/10.36941/jesr-2024-0124

CrowdStrike. (2025). CrowdStrike 2025 global threat report. https://go.crowdstrike.com/rs/281-OBQ-266/images/CrowdStrikeGlobalThreatReport2025.pdf

Czech National Bank. (2025). TIBER-CZ implementation document (Version 1.0). *Czech National Bank*. https://www.cnb.cz/export/sites/cnb/cs/dohled-financni-trh/dora-digitalni-provozni-odolnost-financniho-trhu/tiber-implemetation-guide-cz.pdf

Darktrace. (2024). Annual threat report 2024. https://cdn.prod.website-files.com/626ff4d25aca2edf4325ff97/67b5c5f85bac8f34e1b44d35_Darktrace%20Annual%20Threat%20Report%202024.pdf

European Central Bank. (2024). What is TIBER-EU? https://www.ecb.europa.eu/paym/cyber-resilience/tiber-eu/html/index.en.html

Gelernter, N., Schulmann, H., & Waidner, M. (2024). External attack-surface of modern organizations. In *Proceedings of the 19th ACM Asia Conference on Computer and Communications Security (ASIA CCS '24)*, (pp. 589–604). ACM. https://doi.org/10.1145/3634737.3656295

GREYCORTEX. (2025). GREYCORTEX Mendel: Network detection and response. https://www.greycortex.com/mendel

Hevner, A. R., March, S. T., Park, J., & Ram, S. (2004). Design science in information systems research. *MIS Quarterly*, 28(1), 75–105. https://doi.org/10.2307/25148625

Husák, M., & Sadlek, L. (2025). Attack surface management: State of the art and operational challenges. In *2025 IEEE 11th International Conference on Network Softwarization (NetSoft)*, (pp. 1–6). IEEE. https://doi.org/10.1109/NetSoft64993.2025.11080588

IBM. (2025). IBM X-Force 2025 threat intelligence index. https://www.ibm.com/reports/threat-intelligence

Lerver, D. (2025). State of the underground 2025: Key trends shaping cyber risk today. Bitsight. https://www.bitsight.com/blog/state-of-the-underground-2025

Li, L., Huang, C., & Chen, J. (2024). Automated discovery and mapping ATT&CK tactics and techniques for unstructured cyber threat intelligence. *Computers & Security*, 140, 103815. https://doi.org/10.1016/j.cose.2024.103815

MITRE Corporation. (n.d.-a). ATT&CK evaluations. https://evals.mitre.org/

MITRE Corporation. (n.d.-b). Center for Threat-Informed Defense. https://ctid.mitre.org/

MITRE Corporation. (2025). MITRE ATT&CK framework (Version 18.1). https://attack.mitre.org/versions/v18/

Peffers, K., Tuunanen, T., Rothenberger, M. A., & Chatterjee, S. (2007). A design science research methodology for information systems research. *Journal of Management Information Systems*, 24(3), 45–77. https://doi.org/10.2753/MIS0742-1222240302

Security Risk Advisors. (2025). VECTR documentation. https://docs.vectr.io/

Shin, B., & Lowry, P. B. (2020). A review and theoretical explanation of the "Cyberthreat-Intelligence (CTI) capability" that needs to be fostered in information security practitioners and how this can be accomplished. *Computers & Security*, 92, 101761. https://doi.org/10.1016/j.cose.2020.101761

Verizon. (2025). 2025 data breach investigations report. https://www.verizon.com/business/resources/reports/2025-dbir-data-breach-investigations-report.pdf

Wazuh. (n.d.). MITRE ATT&CK framework. https://documentation.wazuh.com/current/user-manual/ruleset/mitre.html

Acta Informatica Pragensia is published by the Prague University of Economics and Business, Czech Republic | eISSN: 1805-4951