



**UNIVERSITY
OF TURKU**

Assessing the Relevance of Deception Technologies in Modern Cybersecurity Environments

Cyber Security

Master's Degree Programme in Information and Communication Technology

Department of Computing, Faculty of Technology

Master of Science in Technology Thesis

Author:

Pekka Puska

Supervisors:

Professor Seppo Virtanen (Department of Computing)

Doctoral Researcher Saku Lindroos (Department of Computing)

June 2026

The originality of this thesis has been checked in accordance with the University of Turku quality assurance system using the Turnitin Originality Check service.

Master of Science in Technology Thesis
Department of Computing, Faculty of Technology
University of Turku

Subject: Cyber Security

Programme: Master's Degree Programme in Information and Communication Technology

Author: Pekka Puska

Title: Assessing the Relevance of Deception Technologies in Modern Cybersecurity

Environments

Number of pages: 76 pages, 7 appendix pages

Date: June 2026

This thesis examines the relevance, adoption, and operational positioning of deception technologies in modern cybersecurity environments. As threat actors become increasingly adaptive and capable of blending malicious activity with legitimate system behaviour, traditional defensive approaches based primarily on detection and response face growing challenges. The purpose of this study is to examine where deception technologies fit within contemporary cybersecurity architectures and how they complement existing defensive controls.

The theoretical section reviews deception technologies, threat actor behaviour, traditional security controls, and emerging adaptive approaches to cyber deception. Particular attention is given to the relationship between deception technologies and existing capabilities such as Security Information and Event Management (SIEM) and Endpoint Detection and Response (EDR) platforms.

The empirical component of the study was conducted through a questionnaire distributed to experienced cybersecurity practitioners. A total of 47 survey invitations were issued, resulting in nine completed responses. Although the response rate was limited, the respondent group consisted of highly experienced cybersecurity practitioners representing a variety of professional backgrounds and operational environments.

The findings indicate that deception technologies are generally regarded as relevant within modern cybersecurity environments. Respondents most frequently associated deception with early compromise detection, adversary behaviour analysis, lateral movement visibility, and the reduction of uncertainty during security investigations. At the same time, deception technologies were primarily viewed as complementary capabilities rather than replacements for traditional security controls. The results also highlighted the importance of integration, governance, operational complexity, and organizational maturity in determining the successful adoption of deception technologies.

The study concludes that deception technologies have progressed beyond purely experimental concepts and are increasingly recognized as a valuable component of layered cybersecurity architectures. However, their practical value depends less on technical feasibility and more on effective integration, operationalization, and governance within existing security environments. The findings further suggest that the primary challenge facing deception technologies is no longer demonstrating their potential value, but enabling organizations to consistently realize that value in practice.

Keywords: cybersecurity, deception technologies, cyber deception, threat detection, cybersecurity architecture, threat actors, SIEM, EDR

Table of contents

1	Introduction	1
1.1	Research Questions.....	2
1.2	Structure of the Thesis	3
1.3	Use of Artificial Intelligence Tools	3
2	Threat Landscape and Attacker Behaviour	5
2.1	Threat Actor Behaviours and Tactics	5
2.2	Common Techniques and Tools Used During Lateral Movement	8
2.3	Indicators of Compromise	11
2.4	Threat Actor Motivations	15
3	Deception technologies.....	20
3.1	Introduction to Deception Technologies	20
3.2	Placement of Deception Assets	23
3.3	Adaptive and AI-Driven Deception.....	26
4	Traditional Security Controls and Deception Technologies	31
4.1	Security Information and Event Management (SIEM)	31
4.2	Endpoint Detection and Response (EDR)	32
4.3	Comparison with Traditional Security Measures	33
5	Empirical Study: Practitioner Perspectives	36
5.1	Research Design and Methodology	36
5.2	Respondent Profile	38
5.3	Results	40
5.3.1	Perceived Relevance of Deception Technologies	40
5.3.2	Architectural Positioning: Feature, Standalone Tool, or Integrated Control	42
5.3.3	Perceived Operational Applications of Deception Technologies.....	43
5.3.4	Defensive Philosophy: Immediate Containment or Controlled Observation	45
5.3.5	Adoption Barriers and Practical Concerns	48
5.4	Deeper Interpretation of Findings	50
5.4.1	Detection as the Primary Value Driver	50
5.4.2	Theory–Practice Gap.....	52

5.4.3	Importance of Integration.....	53
5.4.4	Organizational Maturity and Perceived Relevance	55
5.4.5	Defensive Philosophy and Attacker Observation	57
5.4.6	Future Potential of AI-Driven Deception	59
5.5	Synthesis of Findings	60
5.5.1	Implications for the Research Questions	62
5.6	Survey Limitations	64
6	Discussion.....	66
6.1	Interpreting the Findings in Context.....	66
6.2	Gap Between Theory and Practice.....	68
6.3	Practical Implications	70
6.4	Limitations of the Study.....	71
6.5	Directions for Future Research	73
7	Conclusion	75
8	References.....	77
9	Appendix A.....	80

1 Introduction

Cybersecurity strategies have evolved considerably over the past two decades. Earlier approaches focused heavily on prevention and perimeter protection, with the assumption that unauthorized access could be prevented through sufficiently strong controls. However, both academic research and industry observations have increasingly recognized that preventing all intrusions is unrealistic in complex digital environments (Hutchins et al., 2011; NIST, 2024). As a result, defensive thinking has gradually shifted toward resilience, visibility, and the ability to detect and respond before an intrusion develops into a significant incident.

One challenge identified in modern cybersecurity environments is that malicious activity does not always appear obviously malicious. Attackers increasingly rely on legitimate credentials, administrative tools, remote management protocols, and trusted communication channels to conduct operations inside compromised networks (MITRE, 2024). Consequently, the problem is often not a lack of telemetry but rather the difficulty of distinguishing malicious activity from normal operational behaviour. This challenge becomes particularly notable during the early stages of an intrusion, where indicators may be subtle and easily overlooked.

Traditional defensive technologies continue to form the foundation of most security architectures. Security Information and Event Management (SIEM) platforms aggregate and correlate data from multiple sources, while Endpoint Detection and Response (EDR) solutions provide visibility into activity occurring on individual hosts. Together, these technologies support detection, investigation, and response activities across modern environments. However, several studies have noted that many detection approaches still depend on indirect indicators and behavioural correlations that may require analysts to piece together multiple events before malicious intent becomes visible (Shackleford, 2023; MITRE, 2024). This can create delays, particularly during the early stages of an intrusion.

Deception technologies approach this problem from a different perspective. Rather than relying solely on observation of existing activity, they introduce controlled assets into the environment that are designed to attract attacker interaction. These assets may include decoy systems, deceptive credentials, synthetic datasets, or simulated services that appear legitimate but have no operational purpose. Interactions with such assets can provide high-confidence indicators of malicious activity because legitimate users typically have little reason to access them (Bowen et al., 2009; Fraunholz et al., 2018).

Research interest in deception technologies has increased considerably in recent years. Previous studies have explored topics such as honeypots, honeytokens, deception environments, and adversarial uncertainty as mechanisms for improving detection and collecting threat intelligence (Almeshekeh & Spafford, 2014; Rowe & Rrushi, 2016; Fraunholz et al., 2018). Much of this work focuses on technical implementation, architecture, and theoretical effectiveness. Comparatively less attention has been given to how cybersecurity practitioners perceive deception technologies, how they are integrated into operational environments, and what factors influence their adoption in practice. This creates a gap between the technical potential described in literature and the realities of organizational deployment.

The purpose of this thesis is to examine the relevance of deception technologies in modern cybersecurity environments. The study combines a review of existing literature with an empirical survey targeting cybersecurity practitioners. Rather than focusing solely on technical implementation, the objective is to understand how deception technologies are positioned within contemporary security architectures, how practitioners perceive their value relative to traditional controls, and which factors appear to influence their adoption and operational use.

1.1 Research Questions

The primary objective of this thesis is to examine the relevance of deception technologies in modern cybersecurity environments and to understand how these technologies are currently positioned within organizational security architectures. While the technical principles of deception have been discussed extensively in both academic and practitioner literature, less attention has been given to how cybersecurity professionals perceive their practical value and role in operational environments.

To address this gap, the study focuses on three research questions.

The first research question examines the overall relevance of deception technologies in contemporary cybersecurity environments. Although deception-based approaches have received increasing attention in recent years, their practical importance relative to other defensive capabilities remains unclear.

The second research question explores how cybersecurity practitioners position deception technologies in relation to traditional security controls such as Security Information and Event Management (SIEM) systems, Endpoint Detection and Response (EDR) platforms, identity security solutions, and network segmentation. Understanding this relationship helps clarify whether deception is viewed as a complementary capability, a niche technology, or a more strategic component of defensive architecture.

The third research question focuses on adoption and operational use. Even when organizations recognize potential value in deception technologies, implementation may be affected by factors such as integration requirements, operational complexity, governance considerations, available expertise, and the ability to measure effectiveness. Identifying these factors helps explain the current maturity and practical deployment of deception technologies. Together, these research questions provide the framework for both the literature review and the empirical study presented later in this thesis.

1.2 Structure of the Thesis

The remainder of this thesis is organized as follows. Chapter 2 examines attacker behaviour, credential abuse, privilege escalation, lateral movement, and other techniques relevant to understanding the defensive role of deception technologies. Chapter 3 reviews deception technologies, their underlying concepts, implementation approaches, and operational considerations. Chapter 4 discusses deception technologies in relation to traditional defensive controls, including SIEM and EDR platforms, as well as automation and orchestration capabilities. Chapter 5 presents the empirical study, including the survey methodology, respondent demographics, results, and analysis. Chapter 6 discusses the findings in relation to existing literature and the research questions. Finally, Chapter 7 concludes the thesis, outlines key findings, and identifies areas for future research.

1.3 Use of Artificial Intelligence Tools

Artificial intelligence (AI) tools were used in a limited supporting role during the preparation of this thesis. Their use was restricted to language assistance, idea exploration, draft review, and preliminary information retrieval. All research design decisions, literature selection, analysis, interpretation of findings, and conclusions were conducted by the author.

Grammarly was used as a language support tool to assist with grammar, spelling, readability, and consistency of written English. Suggestions provided by the tool were reviewed and accepted, modified, or rejected at the author's discretion.

Google Search and Google's AI-assisted search features (AI Overviews) were used during the literature review and background research phases to identify potentially relevant academic publications, industry reports, technical documentation, and publicly available cybersecurity resources. Sources identified through these tools were independently reviewed by the author before inclusion in the thesis. All references included in this thesis were selected and evaluated by the author.

ChatGPT (OpenAI) was used as a support tool for brainstorming, reviewing draft text, improving structure and readability, identifying areas requiring clarification, and discussing the interpretation and presentation of research findings. The tool was also used to assist with language refinement and proofreading. Any suggestions provided by ChatGPT were critically reviewed by the author before incorporation into the thesis.

No AI tool was used to generate research data, survey responses, empirical findings, or the conclusions presented in this study. The author remained fully responsible for the accuracy of the content, the interpretation of the results, and all final decisions regarding the structure and content of the thesis.

2 Threat Landscape and Attacker Behaviour

2.1 Threat Actor Behaviours and Tactics

Understanding attacker behaviour is a crucial part of modern cybersecurity research because defensive controls are often designed around expected adversary actions. Although individual attacks vary according to objectives, resources, and technical capability, previous research has identified recurring patterns that appear across different threat actor groups and attack campaigns (Hutchins et al., 2011; Strom et al., 2018). Frameworks such as MITRE ATT&CK provide a structured method for categorizing these activities into tactics and techniques, allowing defenders to better understand how attackers progress through an intrusion lifecycle (Strom et al., 2018).

While attack paths are rarely identical, many successful intrusions follow broadly similar stages. These commonly include initial access, credential acquisition, privilege escalation, persistence, lateral movement, and actions performed to achieve the attacker's objective (Hutchins et al., 2011). The sequence may vary depending on the environment and attacker goals, but the overall progression remains recognizable across a wide range of incidents. Understanding these stages is particularly relevant in the context of deception technologies because opportunities for detection and intervention frequently emerge as attackers transition between phases of an intrusion.

One of the most commonly observed methods of obtaining initial access remains phishing. Phishing attacks attempt to persuade users to open malicious attachments, visit fraudulent websites, or disclose authentication credentials. Despite substantial investment in awareness training and email security technologies, phishing continues to be one of the most frequently reported intrusion vectors in modern breach investigations (Verizon, 2024). Its effectiveness stems largely from exploiting human behaviour rather than technical vulnerabilities. Within the MITRE ATT&CK framework, phishing-related activities are commonly associated with the Initial Access tactic and, in some cases, Credential Access when authentication information is targeted directly (Strom et al., 2018).

Once access has been established, attackers frequently attempt to acquire additional credentials or increase their level of privilege within the environment. One technique that has received considerable attention in both academic and practitioner literature is Kerberoasting.

To understand Kerberoasting, it is first necessary to understand the role of the Kerberos authentication protocol.

Kerberos is a ticket-based authentication protocol widely used in Microsoft Active Directory environments to authenticate users and services (Microsoft Corporation, 2024). Rather than transmitting passwords directly across the network, Kerberos relies on encrypted tickets to verify identity and authorize access to resources. When a user requests access to a service, the authentication process generates a service ticket that allows access without repeatedly exposing credentials. This approach improves both security and usability by reducing the need for repeated authentication requests while limiting the exposure of sensitive authentication data (Microsoft Corporation, 2024).

Kerberoasting exploits this legitimate authentication process. In a Kerberoasting attack, an attacker who has already obtained access to a domain requests Kerberos service tickets associated with service accounts. These tickets contain encrypted data derived from the service account's password. The attacker can then extract the ticket and attempt to crack the password offline using password-cracking tools. If successful, the attacker may gain access to privileged service accounts without generating the types of authentication failures typically associated with brute-force attacks (Metcalf, 2015; Harmj0y, 2017).

What makes Kerberoasting particularly challenging from a defensive perspective is that requesting service tickets is a normal and expected part of Kerberos operation. As a result, the activity may initially appear indistinguishable from legitimate user behaviour. This highlights a broader challenge in modern cybersecurity environments: attackers increasingly rely on legitimate administrative functions and trusted protocols rather than obviously malicious actions. Consequently, identifying malicious intent often requires additional context beyond the individual event itself (Metcalf, 2015; Harmj0y, 2017).

From the perspective of deception technologies, monitored decoy service accounts and deceptive credentials may provide opportunities to detect this type of activity. If an attacker attempts to enumerate or request tickets associated with decoy accounts that have no legitimate operational purpose, the interaction itself can become a high-confidence indicator of malicious intent. This demonstrates how deception technologies can complement traditional monitoring approaches by creating visibility into activities that might otherwise blend into normal system behaviour.

The significance of Kerberoasting lies in its ability to abuse legitimate functionality rather than exploit software vulnerabilities. The initial ticket requests are valid operations within the Kerberos protocol and may appear indistinguishable from normal activity. Consequently, detecting malicious use often requires behavioural analysis or additional contextual information rather than reliance on individual authentication events alone (Harmj0y, 2017). This characteristic makes Kerberoasting particularly relevant in the context of deception technologies. Because attackers frequently catalog service accounts and request service tickets during credential acquisition activities, defenders can deploy monitored decoy service accounts and deceptive SPNs to generate high-confidence indicators of malicious behaviour.

Following credential acquisition, attackers commonly seek to establish persistence within the environment. Persistence refers to techniques that allow continued access after the original entry point has been discovered or removed. Examples include creating additional user accounts, modifying authentication mechanisms, installing backdoors, or altering system configurations to permit future access (MITRE, 2024). These activities are often designed to blend into normal administrative operations, making them difficult to identify without detailed monitoring. Within the MITRE ATT&CK framework, these activities are categorized under the Persistence and Privilege Escalation tactics (Strom et al., 2018).

As intrusions progress, attacker behaviour typically evolves in response to both opportunities and defensive controls. Activities may transition from credential theft and reconnaissance toward lateral movement, command and control communication, data collection, and ultimately data exfiltration. Lateral movement refers to the process of moving between systems within a compromised environment in order to access additional resources or reach higher-value targets. Command and control mechanisms enable attackers to maintain communication with compromised systems and issue instructions remotely. Data exfiltration involves the unauthorized transfer of information outside the targeted environment (MITRE, 2024).

This progression aligns closely with the intrusion lifecycle described by Hutchins et al. (2011) in the Cyber Kill Chain framework. Rather than viewing individual techniques as isolated events, the framework emphasizes how attackers move through interconnected stages while adapting their approach based on defensive responses and environmental conditions. This aspect is particularly useful when evaluating defensive technologies because it highlights multiple opportunities to interrupt attacker activity before objectives are achieved.

From a defensive perspective, this creates a significant challenge. Organizations must decide where to focus detection and mitigation efforts across a complex attack lifecycle.

Concentrating too heavily on a single stage, such as initial access, may improve visibility in that area while leaving gaps elsewhere. Conversely, attempting to provide equal coverage across every phase of an intrusion can quickly become impractical in larger environments. As a result, many defensive strategies focus on identifying points where attacker progression can be disrupted most effectively. This principle underpins both traditional defensive controls and deception technologies, which seek to detect, delay, or redirect attacker activity before critical systems or sensitive information are reached.

2.2 Common Techniques and Tools Used During Lateral Movement

Research on cyber intrusions has consistently shown that attackers rarely rely on a single technique throughout an attack. Instead, adversaries typically combine multiple methods as they progress through different phases of an intrusion. Frameworks such as MITRE ATT&CK categorize these activities into tactics and techniques, providing a structured view of how attackers move from initial compromise toward their ultimate objectives (Strom et al., 2018). Understanding these relationships is vital when evaluating deception technologies, as deception is most effective when deployed at points where attacker behaviour becomes predictable.

One technique that has received considerable attention in recent years is Living Off the Land (LotL). The term refers to the use of legitimate administrative tools and operating system functionality for malicious purposes. Rather than introducing custom malware, attackers exploit tools that are already present within the environment, including PowerShell, CertUtil, Windows Management Instrumentation (WMI), Bash, or Bitsadmin (MITRE Corporation, 2024). CertUtil is a Microsoft utility intended for certificate management, while Bitsadmin was originally designed to manage background file transfers. Because these tools are commonly used by administrators, malicious activity may appear similar to routine operational behaviour.

The growing use of LotL techniques illustrates a broader challenge for defenders. Detection mechanisms that focus primarily on identifying malicious binaries or known signatures may provide limited visibility when attackers rely on legitimate software. Previous studies have

noted that this overlap between legitimate and malicious activity can significantly complicate detection efforts (MITRE Corporation, 2024). Consequently, defenders increasingly rely on behavioural indicators and contextual information rather than individual tool executions alone.

Deception technologies offer a different approach to this problem. Rather than attempting to identify every instance of PowerShell or command-line activity, deception environments introduce monitored assets that should not normally attract attention. Examples include decoy scripts, simulated administrative tools, or controlled execution paths. Interactions with these assets may provide a stronger indication of malicious intent because legitimate users generally have no operational reason to access them. In this context, deception complements traditional monitoring by reducing uncertainty surrounding potentially suspicious behaviour (Almeshekah & Spafford, 2014).

Credential theft frequently follows reconnaissance and initial access activities. One of the most widely known tools associated with credential acquisition is Mimikatz, which is capable of extracting authentication credentials and password hashes from Windows memory (MITRE Corporation, 2024). Once credentials have been obtained, attackers may employ techniques such as Pass-the-Ticket or Golden Ticket attacks. Pass-the-Ticket involves reusing legitimate Kerberos authentication tickets to access systems without requiring the original password. Golden Ticket attacks are more severe and involve the creation of forged Kerberos Ticket Granting Tickets, potentially allowing attackers to impersonate users and maintain long-term access within an Active Directory environment.

These techniques demonstrate how attackers frequently exploit trusted authentication mechanisms rather than software vulnerabilities. As a result, activity associated with credential misuse may appear legitimate when viewed in isolation. This characteristic has made credential-based attacks an area of particular interest for deception technologies. Researchers have shown that decoy credentials, honeytokens, and monitored service accounts can create highly reliable indicators of malicious activity because legitimate users should never interact with them during normal operations (Fraunholz et al., 2018). When such credentials are used, defenders gain visibility into attacker behaviour without relying solely on anomaly detection.

Following credential acquisition, attackers often attempt to move laterally across the network. Lateral movement refers to the process of accessing additional systems after an initial compromise in order to expand access, locate valuable resources, or reach strategic objectives. Common protocols used during lateral movement include Remote Desktop Protocol (RDP) and Server Message Block (SMB). RDP allows a user to remotely control another computer through a graphical desktop interface, while SMB is primarily used for file sharing, printer access, and communication between systems within a network. Because both protocols are commonly used for legitimate administration and day-to-day business operations, distinguishing malicious use from normal activity can be difficult (MITRE Corporation, 2024).

Deception technologies seek to make lateral movement more visible by introducing monitored systems into likely attack paths. These systems may take the form of full honeypots, lightweight service simulations, or decoy hosts designed to resemble operational assets. From the attacker's perspective, they appear to be ordinary systems connected to the network. From the defender's perspective, any interaction becomes a potentially valuable indicator of malicious activity. The objective is not necessarily to prevent movement entirely, but to create points where attacker actions become observable and attributable (Almeshekah & Spafford, 2014).

More advanced deception deployments extend beyond individual hosts. Honeypots simulate single services or systems, whereas honeynets consist of multiple interconnected deceptive systems designed to resemble realistic network environments. These environments may contain simulated file shares, databases, web applications, or business systems populated with synthetic but believable data. Research has shown that the effectiveness of such environments depends heavily on realism and consistency with surrounding infrastructure (Fraunholz et al., 2018). If the environment appears artificial or disconnected from normal operations, experienced attackers may identify it as deceptive and avoid further interaction.

The relationship between attacker progression and deception deployment is particularly important. Intrusions often evolve through identifiable stages, beginning with initial access and progressing toward credential acquisition, lateral movement, and ultimately data access or exfiltration. Several researchers have argued that deception technologies are most effective when aligned with these stages rather than deployed as isolated security controls (Almeshekah & Spafford, 2014; Hutchins et al., 2011). For example, deceptive credentials may support

early-stage detection, decoy systems may expose lateral movement activity, and simulated data repositories may reveal exfiltration attempts. As illustrated in Figure 1, deception technologies can be positioned throughout the intrusion lifecycle, providing visibility into activities ranging from reconnaissance and credential access to lateral movement and data access.




Figure 1. Relationship Between Attacker Progression and Deception Visibility

Source: Adapted from Hutchins et al. (2011), MITRE ATT&CK (2024), and Almeshekah and Spafford (2014)

This perspective highlights a key characteristic of deception technologies. Their value does not necessarily derive from broad coverage across all attack techniques. Instead, they are most effective when positioned at locations where attacker behaviour becomes predictable and where interaction itself provides meaningful information. Viewed in this way, deception functions less as a standalone security layer and more as a mechanism for increasing visibility across critical stages of an intrusion lifecycle.

2.3 Indicators of Compromise

The identification of Indicators of Compromise (IoCs) is a fundamental component of modern cybersecurity operations. An IoC refers to observable evidence suggesting that a system, network, or account may have been involved in malicious activity. Examples include unusual authentication events, unexpected network connections, suspicious process execution, or

interactions with known malicious infrastructure. While individual indicators may appear insignificant when viewed in isolation, previous studies have demonstrated that their value often emerges when they are analyzed collectively as part of a broader attack sequence (Hutchins et al., 2011; MITRE Corporation, 2024).

One of the challenges associated with IoC-based detection is that many attacker activities resemble legitimate operational behaviour. Modern intrusions frequently rely on valid credentials, administrative tools, and trusted communication channels rather than exploiting obvious technical vulnerabilities. As a result, defenders are often required to distinguish malicious activity from legitimate system administration, making contextual analysis increasingly essential (MITRE Corporation, 2024).

Credential-based attacks provide a useful example of this challenge. Techniques such as Pass-the-Ticket and Golden Ticket attacks abuse the Kerberos authentication protocol to gain or maintain access within Active Directory environments. Although these attacks often leave traces within authentication logs, the resulting indicators may not be immediately recognizable as malicious. Inconsistencies in ticket usage, unusual authentication sequences, or ticket requests originating from unexpected systems may indicate malicious activity, but individual events can often appear legitimate when viewed separately (Metcalf, 2015; Harmj0y, 2017).

Similar observations apply to Service Principal Name (SPN) enumeration. As discussed earlier, SPNs are identifiers used by Kerberos to associate network services with specific service accounts. Attackers frequently enumerate SPNs when preparing for Kerberoasting attacks, making unusual increases in SPN-related queries a potentially valuable indicator of reconnaissance activity. However, SPN requests also occur during normal operations. Their significance therefore depends heavily on surrounding context and supporting evidence. This illustrates a broader challenge in intrusion detection: indicators rarely provide certainty on their own and often require correlation with additional observations before meaningful conclusions can be drawn.

At the endpoint level, credential theft activity may generate more direct indicators. One frequently monitored target is the Local Security Authority Subsystem Service (LSASS), a Windows process responsible for managing authentication credentials and security policies. Credential harvesting tools such as Mimikatz often attempt to access LSASS memory in order

to extract passwords, password hashes, or authentication tickets (MITRE Corporation, 2024). While unusual LSASS access patterns may indicate malicious activity, legitimate administrative tools can occasionally produce similar behaviour, creating uncertainty for defenders.

Research on deception technologies suggests that controlled deceptive assets can help reduce this uncertainty by introducing interactions that should not occur during normal operations (Almeshekah & Spafford, 2014; Fraunholz et al., 2018). Examples include monitored credential stores, decoy service accounts, or deceptive authentication artifacts designed specifically to attract attacker attention. When such assets are accessed, the resulting indicators often provide stronger evidence of malicious intent than traditional behavioural anomalies alone. Rather than relying exclusively on deviations from expected behaviour, defenders gain visibility through interactions that have no legitimate business purpose.

At the network level, identifying attacker activity can be even more challenging. Lateral movement frequently occurs through trusted relationships between systems, particularly in environments with limited network segmentation. Traffic generated during lateral movement may therefore appear similar to routine administration or file-sharing activity. Indicators often emerge not from individual connections, but from changes in communication patterns. Connections between systems that do not normally interact, repeated authentication attempts across multiple hosts, or large-scale discovery activity may all indicate an attacker exploring the environment (MITRE Corporation, 2024).

The Domain Name System (DNS), which translates hostnames into network addresses, can also provide valuable indicators. Queries targeting non-existent but plausible internal systems may reveal reconnaissance activity as attackers attempt to identify resources that can be accessed. Deception technologies can exploit this behaviour by introducing monitored hostnames and services into the environment. Interactions with these deceptive assets may reveal attacker activity that would otherwise blend into normal network traffic (Almeshekah & Spafford, 2014).

IoCs are also frequently observed during ransomware incidents. Previous investigations have shown that ransomware attacks often follow a progression involving initial access, privilege escalation, lateral movement, and eventually large-scale file operations prior to encryption (Verizon, 2024). Significant increases in file access activity, bulk file modifications, or mass

file renaming may therefore indicate that an attack has progressed toward its final stages. Deception technologies can provide additional visibility by redirecting attacker activity toward monitored datasets or decoy file shares, allowing defenders to observe behaviour without exposing operational data. Table 1 presents examples of indicators that may emerge during different stages of an intrusion and illustrates how deception technologies can be aligned with those activities to improve visibility and detection.

Table 1. Example Progression of Indicators of Compromise During an Intrusion

Intrusion Stage	Typical Indicators	Example Deception Asset
Initial Access	Phishing activity	Decoy credentials
Credential Access	SPN enumeration, Kerberos anomalies	Honeytokens
Lateral Movement	RDP and SMB connections	Decoy hosts
Discovery	Network scanning	Honeypots
Data Access / Exfiltration	File access anomalies	Decoy data repositories

As shown in Table 1, indicators of compromise often evolve alongside attacker progression through an environment. Early-stage indicators may include phishing activity and credential-related anomalies, while later stages increasingly involve lateral movement, network discovery, and attempts to access sensitive information. By placing deceptive assets along these likely attack paths, defenders may gain earlier visibility into malicious activity and obtain additional context regarding attacker behaviour.

Research has demonstrated that indicators rarely appear as isolated events. Instead, they typically emerge as part of a sequence of related activities that reflect attacker progression through an environment (Hutchins et al., 2011). An intrusion may begin with a phishing email, followed by authentication anomalies, credential acquisition, lateral movement, and ultimately attempts to access or extract sensitive information. While individual indicators may be explainable in isolation, their combined progression often reveals a much clearer picture of attacker intent.

This perspective is particularly relevant when evaluating deception technologies. Rather than functioning solely as additional alert sources, deception assets can provide context for interpreting existing indicators. Interactions with decoy accounts, deceptive credentials, monitored file shares, or simulated systems can strengthen confidence in otherwise

ambiguous observations and help defenders distinguish genuine threats from routine operational activity (Bowen et al., 2009; Fraunholz et al., 2018).

There are, however, practical, and ethical considerations associated with collecting and analyzing IoC data. Monitoring activities must be designed to avoid unnecessary collection of legitimate user information, and deceptive assets should remain clearly separated from operational resources. Effective implementations therefore seek to maximize visibility while minimizing the risk of exposing sensitive information or disrupting legitimate business processes.

Ultimately, effective IoC identification extends beyond passive observation. It involves understanding how indicators relate to one another, how they fit within the broader intrusion lifecycle, and how the environment itself can be structured to make malicious activity more visible. By combining traditional monitoring with carefully deployed deception assets, organizations can improve both the quality of detection and their understanding of attacker behaviour.

2.4 Threat Actor Motivations

Understanding attacker motivation is an important aspect of cybersecurity because motivations frequently influence both target selection and operational behaviour. While threat actors are often grouped into broad categories, research suggests that motivations are rarely static and may evolve throughout an intrusion campaign (Almeshekeh & Spafford, 2014; Fraunholz et al., 2018). An operation that begins as information gathering may later transition into financial extortion, data theft, or disruptive activity if new opportunities emerge. Consequently, classifying an attacker solely based on an initial objective may provide only a partial understanding of the threat.

Previous studies have shown that attacker behaviour is influenced by a combination of motivation, capability, available resources, and perceived risk (Pawlick et al., 2019). These factors affect not only the techniques used during an attack but also the level of persistence, operational security, and adaptability demonstrated by the attacker. Understanding these differences is valuable when designing defensive measures because different motivations often result in different patterns of behaviour.

Table 2 summarizes common threat actor categories, their primary motivations, typical objectives, and examples of how deception technologies may be aligned with those objectives. While individual actors and campaigns may vary considerably, the categories provide a useful framework for understanding how attacker motivations can influence defensive planning. Table 2. Relationship Between Threat Actor Motivations and Typical Objectives

Threat Actor Type	Primary Motivation	Typical Objectives	Example Deception Opportunity
Financially Motivated Actors	Financial gain	Ransomware, fraud, extortion, data theft	Decoy financial records, fake databases, monitored credentials
Ideologically Motivated Actors	Political or social influence	Defacement, disruption, misinformation	Public-facing decoy services, monitored web assets
State-Sponsored Actors	Intelligence collection, strategic advantage	Defacement, disruption, misinformation	High-interaction deception environments, decoy research repositories
Reputation-Driven Actors	Recognition, technical challenge	Defacement, disruption, misinformation	Technically attractive honeypots, challenge-oriented decoys
Insider Threats	Financial, personal, or organizational motives	Data theft, sabotage, unauthorized access	Technically attractive honeypots, challenge-oriented decoys
Opportunistic Actors	Easily obtainable gain	Exploitation of exposed systems and weak controls	Internet-facing honeypots and exposed decoy services

Source: Adapted from Almeshekah and Spafford (2014), Fraunholz et al. (2018), and Verizon (2024).

As illustrated in Table 2, threat actors pursue a variety of objectives ranging from financial gain and intelligence collection to ideological influence and personal recognition. These differences are relevant from a defensive perspective because they affect the types of assets attackers are likely to target and the techniques they may employ. Consequently, deception

strategies can often be tailored to align with specific attacker interests, increasing the likelihood of interaction with deceptive assets.

Financially motivated threat actors are among the most frequently encountered adversaries in modern cybersecurity incidents. Their activities are typically focused on generating economic benefit through ransomware operations, financial fraud, extortion, or theft of valuable information. The Verizon Data Breach Investigations Report consistently identifies financial gain as one of the most common drivers behind cyberattacks (Verizon, 2024). These actors often prioritize efficiency and scalability, relying on techniques such as phishing, credential theft, and exploitation of exposed services to gain access. Because their objective is usually measurable financial return, operations tend to focus on assets that can be monetized directly.

Ideologically motivated actors operate under different objectives. Rather than seeking financial gain, these groups often pursue political, social, or ideological outcomes. Activities may include website defacement, information manipulation, service disruption, or public disclosure of sensitive information. Success is frequently measured through visibility and influence rather than direct access to organizational assets. From a defensive perspective, these motivations may create opportunities to deploy deception assets that appear publicly major while remaining isolated from operational systems.

State-sponsored actors present a more complex challenge. The prior research indicates that nation-state operations frequently prioritize intelligence collection, strategic advantage, or long-term access rather than immediate disruption (Hutchins et al., 2011). Such campaigns may remain active for extended periods before any observable objective becomes apparent. The emphasis is often placed on maintaining persistence, avoiding detection, and collecting information over time. Consequently, deception technologies intended to detect these actors must be sufficiently realistic to withstand extended observation and validation attempts.

Another category consists of reputation-driven actors. These individuals or groups may seek recognition, technical prestige, or demonstration of capability rather than direct financial or political outcomes. Their behaviour often includes experimentation with different techniques, exploration of technically challenging environments, and attempts to compromise systems perceived as difficult targets. Although their objectives may appear less structured than those of financially motivated or state-sponsored actors, their activities can still create substantial security risks. Research suggests that technically attractive deception environments may be

particularly effective against this category because they encourage interaction while remaining isolated from operational systems (Fraunholz et al., 2018).

Insider threats differ from external threat actors because they begin with some level of legitimate access. The underlying motivation may still be financial, personal, ideological, or retaliatory, but the attack path differs significantly from traditional external intrusions. Instead of obtaining access through exploitation or credential theft, insiders misuse privileges that have already been granted. This creates a unique detection challenge because insider activity often resembles legitimate behaviour. Deception technologies may provide additional visibility by introducing monitored resources that should not normally attract attention from authorized personnel.

Opportunistic actors represent another distinct category. Unlike highly structured threat groups, opportunistic attackers frequently target systems based on accessibility rather than strategic value. Their behaviour is often driven by exposed services, weak credentials, or publicly known vulnerabilities. Although these actors may appear less sophisticated, their actions can still result in considerable disruption. Research has shown that opportunistic attackers frequently follow predictable attack paths, making them suitable candidates for detection through honeypots and other deceptive assets designed to attract automated scanning and exploitation attempts (Almeshekeh & Spafford, 2014).

A vital observation across all categories is that motivation should not be viewed as a fixed attribute. Threat actors may alter their objectives in response to opportunities encountered during an intrusion. Information initially collected for intelligence purposes may later become valuable for extortion. Similarly, a financially motivated intrusion may evolve into broader disruption if attackers encounter resistance or seek additional leverage. This dynamic nature limits the usefulness of static classification models and reinforces the need for continuous behavioural assessment.

The relationship between motivation and behaviour is particularly relevant in the context of deception technologies. Different attacker motivations often influence which assets are likely to attract attention and how those assets are evaluated. Interactions with deceptive resources can therefore provide insights not only into what an attacker is doing, but also into the objectives that may be driving those actions. This additional context can improve defensive

decision-making by helping analysts distinguish between different forms of malicious activity and prioritize responses accordingly.

There are, however, practical limitations to this approach. Effective deception environments require a balance between realism and security. Assets that appear artificial may fail to attract meaningful interaction, while environments that are excessively realistic may introduce operational or compliance risks. Maintaining this balance remains one of the central challenges associated with the deployment of deception technologies (Fraunholz et al., 2018).

Ultimately, understanding attacker motivation is less about assigning permanent labels and more about improving expectations regarding likely behaviour. While motivations alone cannot predict future actions with certainty, they provide a useful framework for understanding how attackers select targets, adapt techniques, and pursue objectives.

Deception technologies build upon this understanding by introducing controlled opportunities for observation at points where attacker behaviour is most likely to reveal intent.

3 Deception technologies

3.1 Introduction to Deception Technologies

Deception technologies represent a proactive approach to cybersecurity that differs from traditional monitoring and detection methods. Rather than relying solely on observation of existing systems and user activity, deception technologies introduce controlled assets into an environment with the purpose of influencing attacker behaviour and generating observable interactions. Research in this area has increasingly focused on how deceptive assets can influence attacker decision-making, increase uncertainty, and create opportunities for earlier detection of malicious activity (Almeshekah & Spafford, 2014; Jajodia et al., 2016).

The fundamental principle regarding deception technologies is that interactions with deceptive assets should not occur frequently during legitimate operations. Decoy credentials, simulated systems, fictitious datasets, and deceptive network services are intentionally designed to appear authentic while remaining isolated from production resources. Because these assets do not serve a legitimate business function, interactions with them can often be interpreted with a higher degree of confidence than traditional anomaly-based alerts (Fraunholz et al., 2018).

Unlike conventional security controls that primarily seek to identify suspicious behaviour after it has occurred, deception technologies attempt to influence attacker decision-making during an intrusion. Previous research has suggested that attackers frequently operate under conditions of incomplete information and uncertainty. Deception seeks to exploit this uncertainty by presenting believable but misleading information that influences how attackers prioritize targets, allocate resources, and navigate through an environment (Almeshekah & Spafford, 2014; Pawlick et al., 2019).

Several forms of deception technologies have emerged within modern cybersecurity architectures. Honeypots are among the most widely recognized examples and consist of systems designed to provoke attacker interaction while recording associated activity. More advanced implementations may involve honeynets, which are collections of interconnected deceptive systems intended to simulate realistic network environments. Other forms of deception include honeytokens, which are digital artifacts such as credentials, documents, or

authentication tokens that generate alerts when accessed, and decoy credentials designed to expose unauthorized credential use (Fraunholz et al., 2018).

The value of deception technologies extends beyond detection alone. Because deceptive assets are positioned along likely attack paths, they may reveal reconnaissance activity, credential harvesting attempts, lateral movement, or data access behaviour before operational systems are affected. For example, attackers performing Kerberoasting often enumerate service accounts and request Kerberos service tickets during the credential acquisition phase of an intrusion. The deployment of monitored decoy service accounts can therefore create opportunities to identify such activity before legitimate accounts are compromised (Metcalf, 2015; Harmj0y, 2017).

Research has also shown that deception technologies may contribute to reducing false positives when compared with traditional monitoring approaches. Conventional detection mechanisms frequently rely on identifying deviations from expected behaviour, which can result in large volumes of alerts requiring validation. This can lead to Alert Fatigue, where analysts become overwhelmed by the volume of alerts and may overlook genuinely important security events, allowing some deviations to go unnoticed (Proofpoint, 2026). Interactions with deceptive assets, however, often provide stronger evidence of malicious intent because legitimate users typically have no operational reason to access them (Fraunholz et al., 2018).

The effectiveness of deception technologies depends heavily on realism. If deceptive assets appear artificial or inconsistent with the surrounding environment, sophisticated attackers may identify and avoid them. Consequently, successful implementations seek to replicate normal operational characteristics, including naming conventions, network relationships, access patterns, and data structures. At the same time, deceptive assets must remain sufficiently isolated to prevent unintended exposure of operational systems or sensitive information (Jajodia et al., 2016).

The application of deception technologies varies across different phases of an intrusion. Depending on the attacker's objectives and stage of progression, different forms of deceptive assets can provide visibility into specific attacker activities. Table 3 summarizes how common deception mechanisms can be aligned with different stages of the intrusion lifecycle and illustrates the types of visibility that may be gained through their deployment.

Table 3. Relationship Between Attacker Actions and Deception-Based Visibility

Intrusion Stage	Typical Attacker Activity	Example Deception Asset	Defensive Value
Reconnaissance	System discovery	Honeypot	Early detection
Credential Access	Credential harvesting	Decoy credentials	High-confidence alert
Lateral Movement	Host-to-host access	Decoy systems	Visibility into attacker progression
Data Access	File and database discovery	Decoy data repositories	Insight into attacker objectives
Exfiltration	Data collection and transfer	Decoy documents	Detection of theft attempts

Source: Adapted from Hutchins et al. (2011), Almeshekah and Spafford (2014), and Fraunholz et al. (2018).

As shown in Table 3, deception technologies can provide value throughout multiple stages of an intrusion rather than serving a single detection purpose. Early-stage deception assets, such as honeypots, may reveal reconnaissance activity, while decoy credentials and deceptive systems can provide visibility into credential abuse and lateral movement. At later stages of an intrusion, deceptive datasets and document repositories may help identify attacker objectives and potential data theft attempts. This demonstrates that deception technologies can contribute to visibility across the intrusion lifecycle and support defensive decision-making at multiple points during an attack.

Although deception technologies do not replace traditional security controls, they provide a complementary mechanism for increasing visibility across different stages of an intrusion lifecycle. Their primary contribution resides with creating opportunities for observation where attacker interaction itself becomes meaningful. Rather than relying exclusively on detecting anomalies after they occur, defenders gain the ability to introduce controlled uncertainty into the environment and observe how adversaries respond. This shift from passive observation toward active influence represents one of the defining characteristics of modern deception-based cybersecurity approaches, pushing the decision-making challenge on to the malicious actor.

3.2 Placement of Deception Assets

The placement of deception assets is one of the most essential design decisions in deception-based cybersecurity architectures. Previous research has indicated that the effectiveness of deception technologies depends not only on the quality of individual deceptive elements but also on where those elements are positioned within an environment (Almeshekah & Spafford, 2014; Fraunholz et al., 2018). If deceptive assets are placed in locations that attackers are unlikely to encounter, they provide limited defensive value. Conversely, assets positioned along expected attack paths are more likely to generate meaningful interactions and reveal attacker behaviour.

Research generally distinguishes between endpoint-oriented and network-oriented deception approaches. Although both seek to create opportunities for observation and detection, they operate at different levels of the environment and target different stages of attacker activity.

Endpoint-based deception focuses on systems where attackers are expected to operate after initial access has been established. These systems include user workstations, application servers, administrative hosts, and other endpoints that may contain credentials, configuration data, or references to additional resources. Rather than attempting to prevent unauthorized access directly, endpoint-based deception introduces controlled artifacts designed to attract attacker attention and generate observable interactions.

Examples of endpoint deception include decoy credentials, deceptive configuration files, synthetic documents, and references to non-existent resources. These artifacts are intended to appear legitimate while remaining disconnected from operational assets. The underlying assumption is that attackers frequently search compromised systems for information that can facilitate privilege escalation, lateral movement, or access to additional resources.

Consequently, interactions with deceptive artifacts may provide early indicators of malicious activity that would otherwise remain difficult to distinguish from legitimate system behaviour (Fraunholz et al., 2018).

The relevance of endpoint-based deception has increased alongside the growing use of LotL techniques. As discussed previously, LotL approaches rely on legitimate operating system tools and administrative utilities rather than custom malware. Because these tools are commonly used for routine administration, traditional detection mechanisms may struggle to distinguish malicious activity from normal operations. Deceptive artifacts positioned within

expected system locations can provide an alternative detection mechanism by creating interactions that should not occur during legitimate use (MITRE Corporation, 2024).

Credential-focused deception represents a particularly common implementation. Attackers frequently target authentication data during post-compromise activities, using techniques such as credential harvesting, Kerberoasting, or privilege escalation. The deployment of monitored credentials, decoy service accounts, or deceptive authentication artifacts can therefore provide visibility into activities that might otherwise blend into normal system behaviour. Research has suggested that interactions involving such assets often produce high-confidence indicators because legitimate users have no operational reason to access them (Bowen et al., 2009; Fraunholz et al., 2018).

Network-based deception extends these principles beyond individual systems and applies them across broader network environments. Rather than focusing on artifacts located on specific endpoints, network-oriented approaches introduce deceptive systems, services, and resources into areas where attackers are expected to conduct reconnaissance, lateral movement, or resource discovery.

Common examples include honeypots, honeynets, simulated databases, deceptive file shares, and controlled application services. These assets are designed to resemble legitimate components of the environment while remaining isolated from production infrastructure. Their purpose is not simply to attract attention but to provide visibility into how attackers explore and navigate a network (Jajodia et al., 2016).

Several studies have argued that network-based deception is particularly valuable during reconnaissance and lateral movement phases because these activities often require attackers to identify systems, enumerate services, and discover relationships between resources (Almeshekeh & Spafford, 2014). By introducing believable but controlled targets into these processes, defenders can gain insight into attacker objectives and progression through the environment.

Identity infrastructure frequently plays a central role in network-based deception strategies. Modern enterprise environments rely heavily on centralized authentication services and directory systems. Because attackers commonly use compromised credentials to move between systems, deceptive resources integrated into authentication workflows can expose unauthorized access attempts that may otherwise appear legitimate. This approach shifts

detection away from identifying anomalies alone and toward identifying interactions that should not occur under normal operating conditions.

The effectiveness of both endpoint-based and network-based deception depends heavily on context. Deceptive assets must reflect the characteristics of the environment in which they are deployed. Resources that appear plausible within a healthcare organization may differ significantly from those expected within a financial institution, manufacturing company, or government environment. Previous research has emphasized that successful deception environments align with attacker expectations regarding naming conventions, system roles, data structures, and operational workflows (Fraunholz et al., 2018).

Table 4 summarizes several core design principles commonly associated with effective deception deployments. These principles influence both the credibility of deceptive assets and their ability to generate meaningful visibility into attacker behaviour.

Table 4. Core Design Principles of Deception Technologies.

Design Principle	Purpose	Example Implementation
Placement	Increase likelihood of attacker interaction	Endpoint and network decoys
Realism	Increase credibility	Synthetic datasets, realistic services
Isolation	Prevent operational risk	Segmented environments
Adaptivity	Respond to attacker behaviour	Dynamic deception assets

Source: Adapted from Almeshekah and Spafford (2014), Fraunholz et al. (2018), and Jajodia et al. (2016).

As Table 4 illustrates, successful deception environments rely on more than the deployment of individual decoys. Placement influences the likelihood of attacker interaction, realism supports credibility, isolation reduces operational risk, and adaptivity enables deception assets to respond to changing attacker behaviour. Together, these principles help ensure that deception technologies remain both effective and safe to operate within production environments.

An important observation within the literature is that deception assets rarely operate effectively in isolation. The most successful implementations often combine multiple deceptive elements into coherent environments that support one another. A deceptive document may reference a hostname associated with a decoy system, which in turn may

provide access to additional synthetic resources. This interconnected approach increases plausibility and allows defenders to observe attacker behaviour across multiple stages of an intrusion rather than through a single interaction (Jajodia et al., 2016).

Ultimately, the placement of deception assets is intended to maximize opportunities for meaningful observation while minimizing operational risk. Endpoint-based deception provides visibility into activities occurring on compromised systems, whereas network-based deception extends visibility into attacker movement and resource discovery across the broader environment. Together, these approaches form a foundation upon which other deception design principles, including realism, isolation, and adaptivity, can be built.

3.3 Adaptive and AI-Driven Deception

Recent research has increasingly explored how deception technologies can move beyond static deployments toward more adaptive and responsive environments. Traditional deception implementations typically rely on preconfigured assets that remain largely unchanged after deployment. While such approaches can provide valuable detection capabilities, their effectiveness may decrease over time as attackers become more experienced in identifying and validating deceptive environments (Fraunholz et al., 2018). As a result, researchers have begun examining how automation, behavioural analytics, and artificial intelligence (AI) can be used to improve the realism and adaptability of deception systems (Pawlick et al., 2019; Walter et al., 2021).

In the context of deception technologies, AI-driven deception refers to the use of automated decision-making, behavioural analysis, or predictive models to modify deceptive assets in response to observed activity. Rather than relying exclusively on predefined configurations, adaptive systems seek to adjust aspects of the environment based on how attackers interact with it. The objective is not to predict attacker behaviour with complete accuracy, but to increase uncertainty and maintain the credibility of deceptive assets as an intrusion progresses.

Previous studies have suggested that attackers frequently adapt their techniques in response to defensive controls and environmental conditions (Hutchins et al., 2011). Consequently, deception environments that remain entirely static may become less effective if they fail to

reflect changes in attacker behaviour. Adaptive approaches attempt to address this limitation by introducing controlled variation into the environment. Adjustments may include the creation of additional deceptive resources, modification of existing assets, or changes in how information is presented to an attacker. These changes are intended to maintain plausibility while preserving the controlled nature of the environment.

One area where adaptive deception may provide value is during transitions between different attack phases. Intrusions rarely follow a completely fixed path, and attackers frequently modify their approach when initial objectives are unsuccessful. Activity that begins as credential acquisition may later transition into lateral movement, privilege escalation, or attempts to access sensitive information. Adaptive deception systems seek to identify these behavioural shifts and introduce deceptive elements that remain consistent with the attacker's apparent objectives. In practice, this may involve presenting additional systems, credentials, services, or datasets that appear relevant to the observed activity while remaining isolated from operational resources.

Healthcare and financial environments provide useful examples of how context influences adaptive deception. In healthcare organizations, suspicious activity associated with ransomware preparation may justify the deployment of synthetic patient records or deceptive clinical resources designed to attract further interaction while protecting operational systems. Similarly, financial environments may employ synthetic transaction data, decoy financial records, or deceptive repositories that appear valuable to attackers seeking sensitive information. Although implementation details differ between sectors, the underlying principle remains consistent: deceptive assets are aligned with attacker expectations while remaining separated from production resources.

The credibility of deception environments represents another area where adaptive approaches may provide advantages. Research has shown that sophisticated attackers frequently attempt to validate systems before committing substantial effort to an intrusion path (Fraunholz et al., 2018). Such validation may involve examining system behaviour, evaluating response times, reviewing historical data, or assessing relationships between systems. If deceptive assets appear unrealistic or inconsistent with the surrounding environment, attackers may identify and avoid them as well as become more alert to the existence of deceptions in the environment. Adaptive deception attempts to mitigate this risk by introducing controlled

variation over time, allowing deceptive assets to evolve in a manner that resembles legitimate operational systems.

This adaptation does not necessarily require large-scale changes as even minute adjustments may significantly influence perceived realism. Examples include updating synthetic log entries, modifying references between systems, generating realistic background activity, or adjusting metadata associated with deceptive resources. Individually, these modifications may appear minor, but collectively they can contribute to a more believable environment.

Adaptive deception may also influence attacker decision-making rather than merely observing attacker activity. Previous research has argued that deception can alter how attackers prioritize targets and allocate resources by introducing uncertainty into the decision-making process (Almeshekah & Spafford, 2014; Pawlick et al., 2019). Adaptive systems extend this concept by dynamically presenting resources that appear increasingly attractive based on observed behaviour. The objective is not to force attacker actions, but to increase the likelihood that attackers voluntarily engage with controlled assets rather than operational systems.

Cloud environments have created additional opportunities for adaptive deception because resources can be provisioned, modified, and removed with relatively low operational overhead. Virtual systems, application interfaces, and storage resources can be generated dynamically without affecting production services. At the same time, this flexibility introduces new challenges. Changes that occur too frequently or without sufficient consistency may reduce credibility and reveal the presence of deception. Maintaining realism therefore remains a critical requirement regardless of the underlying technology platform.

Information sharing may further enhance adaptive deception capabilities. Observations collected from one environment can potentially alert to adjustments in another, reducing the time required to respond to emerging attacker techniques. Although this concept remains an active area of research, several authors have suggested that collaborative approaches may improve the overall effectiveness of deception-based defense by accelerating learning across multiple environments (Walter et al., 2021).

Despite these potential benefits, adaptive deception introduces additional challenges. The increased use of automation and AI requires careful governance to ensure that deceptive assets remain separated from operational systems and do not inadvertently expose sensitive information. Synthetic data must remain clearly distinguishable from production data from a

governance perspective, even when the environment is designed to appear realistic to an attacker. Maintaining this balance becomes increasingly essential as deception systems become more autonomous and dynamic.

Another consideration is that excessive adaptation may itself become observable. If environments change too rapidly or in ways that do not reflect normal operational behaviour, attackers may become suspicious. Consequently, adaptive deception is not intended to create perfectly responsive environments. Instead, its purpose is to introduce sufficient variability to reduce predictability while preserving realism and maintaining defender control over the environment.

Current research suggests that adaptive and AI-driven deception represents an important area of future development within cybersecurity. While traditional deception technologies remain valuable for detection and observation, adaptive approaches may provide additional opportunities to influence attacker behaviour, maintain credibility over longer periods, and improve the overall effectiveness of deception-based defense strategies (Pawlick et al., 2019; Walter et al., 2021).

Table 5 presents a simplified conceptual model of adaptive deception. The model illustrates how observed attacker activity may be analyzed and used to dynamically adjust deception assets in response to changing attacker behaviour.

Table 5. Conceptual Model of Adaptive Deception

Observed Activity	System Analysis	Adaptive Response
Credential harvesting	Behaviour assessment	Deploy decoy credentials
Network reconnaissance	Pattern identification	Present decoy systems
Lateral movement	Context evaluation	Generate additional deceptive assets
Data discovery	Objective estimation	Present synthetic datasets

Source: Adapted from Almeshekah and Spafford (2014), Fraunholz et al. (2018), and NIST (2024).

Table 5 presents adaptive deception as an iterative process in which observed attacker activity informs subsequent defensive actions. In this model, indicators such as credential harvesting, reconnaissance, lateral movement, or data discovery activities may trigger adjustments to the deception environment. The purpose of these adjustments is not to predict attacker behaviour

with certainty, but to ensure that deceptive assets remain credible and relevant throughout the intrusion lifecycle.

This approach differs from traditional static deception by enabling defensive measures to evolve in response to changing attacker behaviour. Consequently, adaptive deception may improve the longevity and effectiveness of deception environments while creating additional opportunities to observe attacker intent and decision-making.

4 Traditional Security Controls and Deception Technologies

4.1 Security Information and Event Management (SIEM)

Security Information and Event Management (SIEM) platforms have become a central component of modern security operations. Their primary purpose is to aggregate and analyze security-relevant data from multiple sources, including servers, endpoints, network devices, authentication services, and security applications. By consolidating information from across an environment, SIEM platforms aim to provide a broader understanding of security events than would be possible through isolated log analysis alone (Shackleford, 2023; NIST, 2024).

The value of a SIEM is largely determined by the quality and completeness of the telemetry it receives. Research has consistently highlighted that SIEM platforms do not generate visibility independently; rather, they depend on data collected by other security controls and infrastructure components (Shackleford, 2023). Consequently, gaps in logging, monitoring, or telemetry collection may significantly reduce the effectiveness of subsequent analysis and correlation.

One frequently discussed challenge is contextual interpretation. SIEM platforms are designed to identify relationships between events that may individually appear benign. However, many modern attacks are intentionally designed to resemble legitimate activity. This is particularly evident in attacks that rely on valid credentials, administrative utilities, or LoTL techniques. In such cases, individual events may not appear sufficiently suspicious to trigger meaningful alerts, even though they form part of a broader intrusion sequence (MITRE Corporation, 2024).

Another limitation relates to scale. Enterprise environments generate large volumes of operational and security data, much of which may have limited investigative value. As alert volumes increase, analysts must prioritize which events warrant investigation. Previous studies have identified alert fatigue and prioritization challenges as persistent issues within security operations centers, particularly when detection relies heavily on indirect indicators requiring extensive correlation and validation (Shackleford, 2023).

The relationship between SIEM and deception technologies is therefore largely centered on signal quality. Traditional SIEM analysis often attempts to infer malicious intent from combinations of events and behavioural anomalies. Deception technologies approach the

problem differently by introducing assets that should not attract legitimate interaction. Examples include decoy credentials, deceptive file shares, synthetic datasets, and monitored systems. Interactions with such assets may provide indicators that are more explicit and easier to interpret than conventional anomalies because they remove part of the ambiguity associated with normal operational activity (Almeshekah & Spafford, 2014; Fraunholz et al., 2018).

When integrated into SIEM workflows, deception-generated events can complement traditional telemetry rather than replace it. The SIEM continues to provide aggregation, correlation, and historical visibility, while deception introduces high-confidence signals that may strengthen investigative conclusions. This relationship illustrates an important distinction: deception technologies do not eliminate the need for centralized monitoring platforms, but they may improve the quality and interpretability of the information those platforms receive.

4.2 Endpoint Detection and Response (EDR)

Endpoint Detection and Response (EDR) solutions provide visibility directly at the endpoint level, where attacker activity often first becomes observable. Unlike SIEM platforms, which rely primarily on aggregated telemetry, EDR systems operate on individual devices and monitor processes, file activity, authentication events, memory access, and network communications in near real time (NIST, 2024).

This proximity allows EDR solutions to identify activities that may be difficult to observe through centralized logging alone. Credential theft provides a useful example. Techniques targeting the LSASS, which stores authentication-related information in Windows environments, may generate only limited visibility through conventional log sources. Endpoint-level monitoring, however, can provide detailed information regarding process behaviour, memory access, and execution context (MITRE Corporation, 2024).

The value of EDR extends beyond detection. Modern platforms frequently support response capabilities including process termination, file quarantine, network isolation, and automated containment actions. These capabilities enable organizations to reduce attacker dwell time by responding directly at the point where suspicious activity is observed.

Despite these advantages, EDR solutions face challenges similar to those encountered by SIEM platforms. Many detections still require contextual interpretation, particularly when attackers rely on legitimate administrative tools or valid credentials. Activities associated with reconnaissance, persistence, or lateral movement may appear legitimate when viewed in isolation, requiring analysts to evaluate broader behavioural patterns before reaching conclusions.

Deception technologies can enhance endpoint visibility by introducing controlled assets designed to reveal unauthorized activity. Examples include decoy credentials, deceptive files, monitored registry entries, and synthetic data artifacts. Because legitimate users should rarely interact with these assets, they may generate more reliable indicators of malicious behaviour than behavioural anomalies alone (Bowen et al., 2009; Fraunholz et al., 2018).

The combination of EDR and deception technologies is particularly relevant in environments where false positives carry significant operational consequences. Interactions with deceptive assets provide additional context that can improve confidence in detection decisions and reduce uncertainty during investigations. Consequently, deception may serve not as a replacement for endpoint monitoring, but as a mechanism for strengthening the interpretation of endpoint-level activity.

4.3 Comparison with Traditional Security Measures

Traditional security controls remain essential components of modern cybersecurity architectures. Firewalls, intrusion detection systems, antivirus platforms, SIEM solutions, and EDR technologies provide the foundation upon which most defensive strategies are built. Their effectiveness has been demonstrated across a wide range of threat scenarios, particularly when addressing known threats, policy violations, and observable malicious behaviour (NIST, 2024).

However, several researchers have argued that modern attackers increasingly exploit legitimate functionality rather than relying exclusively on traditional malware or clearly malicious actions (MITRE Corporation, 2024). Credential misuse, administrative tool abuse, and LoTL techniques frequently generate activity that appears consistent with normal

operations. As a result, distinguishing malicious behaviour from legitimate behaviour has become an increasingly essential challenge.

Deception technologies address this challenge through a fundamentally different detection philosophy. Traditional controls primarily seek to identify suspicious behaviour after it occurs, relying on signatures, behavioural models, rules, or anomaly detection. Deception technologies instead introduce controlled assets specifically designed to attract interaction. The resulting indicators are often easier to interpret because they are based on interactions that should not occur during legitimate operations (Almeshekah & Spafford, 2014; Fraunholz et al., 2018).

Table 6 compares several characteristics commonly associated with SIEM platforms, EDR solutions, and deception technologies. The comparison highlights how these capabilities differ in their sources of visibility, detection methods, and contribution to security operations.

Table 6. Comparison of traditional security controls and deception technologies.

Capability	SIEM	EDR	Deception Technologies
Primary Data Source	Aggregated logs	Endpoint telemetry	Controlled attacker interactions
Detection Method	Correlation and analytics	Behaviour and process monitoring	Interaction with deceptive assets
Visibility into Attacker Intent	Indirect	Partial	Often more direct
False Positive Risk	Moderate to High	Moderate	Typically lower
Response Capability	Limited	Extensive	Usually integrated with other controls
Primary Strength	Centralized visibility	Endpoint visibility	High-confidence indicators

Source: Adapted from Almeshekah and Spafford (2014), Fraunholz et al. (2018), and NIST (2024).

The comparison presented in Table 6 demonstrates that deception technologies differ from traditional security controls in the way they generate and interpret signals. While SIEM platforms aggregate and correlate large volumes of data and EDR solutions focus on endpoint behaviour, deception technologies derive visibility from interactions with assets that should

not normally be accessed. As a result, the resulting alerts may provide stronger indications of malicious intent and are often associated with lower false-positive rates.

Rather than competing with traditional security controls, deception technologies are generally positioned as complementary capabilities. SIEM platforms continue to provide centralized visibility and correlation, while EDR solutions deliver detailed endpoint telemetry and response capabilities. Deception technologies contribute by reducing ambiguity and providing additional context for interpreting suspicious activity. Their primary value lies not in replacing existing controls, but in strengthening the quality of information available to those controls.

Viewed in this way, deception technologies represent an additional layer within a broader defense strategy. Traditional controls remain responsible for prevention, monitoring, and response, while deception technologies contribute visibility, context, and insight into attacker behaviour. The combination of these approaches reflects a broader shift in cybersecurity toward layered and complementary defensive architectures rather than reliance on any single technology.

5 Empirical Study: Practitioner Perspectives

5.1 Research Design and Methodology

The empirical component of this thesis was conducted using a questionnaire targeted at cybersecurity practitioners. The purpose of the questionnaire was not to produce statistically representative results for the cybersecurity industry as a whole, but rather to gather informed perspectives from individuals with practical experience in security-related roles and environments. The focus was on understanding how deception technologies are perceived in practice, where practitioners identify their value, and which factors appear to influence their adoption and operational use.

A questionnaire-based approach was selected because deception technologies remain unevenly adopted across organizations. Existing literature provides extensive discussion of technical implementation, architectural design, and theoretical defensive advantages, but these sources do not necessarily explain how deception technologies are viewed by professionals responsible for day-to-day security operations, architecture, governance, or incident response. The questionnaire therefore served as a mechanism for connecting the conceptual and technical discussion presented in earlier chapters with contemporary practitioner perspectives.

The questionnaire, which can be found as Appendix A, consisted of both structured and open-ended questions. Structured questions were used to identify measurable patterns across respondents, including perceived relevance, architectural positioning, preferred use cases, adoption barriers, and future expectations. Open-ended questions were included to allow respondents to elaborate on their views and provide context that could not be captured through predefined answer options alone. This approach was considered particularly substantial because cybersecurity environments vary considerably in terms of organizational maturity, sector, risk profile, and operational responsibilities.

Respondent selection followed a purposive expert sampling approach. Participants were selected from the researcher's professional network or were recommended by other cybersecurity professionals during the research process. The objective was not to achieve demographic representativeness, but to ensure that responses were obtained from individuals with substantial cybersecurity experience and exposure to different operational perspectives. Effort was made to include respondents from varied backgrounds, including security

operations, architecture, governance, consulting, risk management, and leadership roles. This diversity was considered vital because deception technologies intersect multiple areas of cybersecurity rather than belonging exclusively to a single technical discipline.

The questionnaire was distributed to 47 cybersecurity practitioners, resulting in nine completed responses. This corresponds to a response rate of approximately 19.1%. While the number of responses remains relatively limited, the respondent group consisted exclusively of experienced cybersecurity practitioners. The purpose of the study was therefore to capture expert viewpoints rather than general industry awareness. Consequently, the findings should be interpreted as indicative of practitioner perspectives rather than statistically representative of the broader cybersecurity population.

The design of the questionnaire sought to balance flexibility with analytical value. Questions were intentionally formulated to allow respondents to express nuanced opinions while still producing responses that could be compared and analyzed systematically. This was particularly important because deception technologies are often discussed differently depending on organizational context, technical maturity, and operational objectives. The combination of structured and open-ended responses enabled both quantitative analysis and qualitative interpretation.

Anonymity was another essential consideration during data collection. Several potential participants expressed reservations about providing identifiable information such as names, organizational affiliation, or detailed professional background. This reluctance is not uncommon within the cybersecurity field. Professionals working with security operations, incident response, critical infrastructure, or regulated environments frequently operate under confidentiality requirements and may be restricted from publicly discussing security practices or defensive capabilities. Even when surveys do not request sensitive operational information, a culture of discretion often influences participation decisions. To encourage candid responses and reduce participation barriers, respondents were not required to provide identifying information beyond what was necessary for analytical purposes.

The analysis was conducted using an interpretive approach that combined quantitative and qualitative observations. Quantitative responses were used to identify recurring patterns and trends across respondents, while open-text responses provided additional context regarding the reasoning behind those patterns. Rather than focusing solely on individual responses, the

objective was to examine what the collective survey responses indicate about the current maturity, relevance, and positioning of deception technologies within contemporary cybersecurity practice.

This methodological approach aligns with the broader objective of the thesis. The study seeks not only to examine how deception technologies function technically, but also to understand how cybersecurity practitioners evaluate them, where they fit within existing defensive architectures, and what factors influence their adoption in real-world environments.

5.2 Respondent Profile

A total of nine completed responses were received from cybersecurity practitioners. Although the number of respondents was relatively small, the experience level of the participant group was notably high. All respondents reported at least eight years of professional cybersecurity experience, indicating that the findings reflect views formed through substantial exposure to security operations, architecture, governance, incident response, risk management, or related cybersecurity disciplines. The distribution of professional experience is presented in Figure 2.




Figure 2. Distribution of Respondent Cybersecurity Experience

Figure 8 illustrates the professional experience of the respondent group. The results imply a highly experienced sample, with eight respondents (88.9%) reporting more than 13 years of cybersecurity experience and four respondents (44.4%) reporting more than 20 years of experience. No respondents reported less than eight years of professional cybersecurity experience.

The experience profile of the respondent group is particularly relevant because deception technologies are often evaluated differently depending on operational responsibilities and organizational perspective. Security operations personnel may prioritize alert quality, detection confidence, and incident response efficiency. Security architects may focus on integration requirements, technical placement, and compatibility with existing security controls. Security leaders and governance-focused practitioners may place greater emphasis on risk management, organizational maturity, and strategic alignment.

Consequently, the findings presented in this chapter should be interpreted as informed practitioner perspectives rather than general awareness or opinion-based responses. The high level of professional experience represented within the respondent group provides valuable insight into how deception technologies are currently viewed by individuals who regularly work with security operations, architecture, governance, and defensive decision-making.

The respondent group therefore represents more than a single technical perspective. The collected responses reflect viewpoints from practitioners who have experience with different aspects of cybersecurity programs and defensive architectures. This diversity is valuable because deception technologies frequently span multiple areas of cybersecurity practice, including detection engineering, threat hunting, security operations, architecture, governance, and intelligence collection.

The experience distribution also provides valuable context for interpreting the survey findings. Responses should not be understood as general opinions regarding cybersecurity technologies, but rather as informed assessments from individuals who regularly work with defensive security capabilities. Consequently, observations regarding adoption barriers, operational value, integration challenges, and future development trends may carry greater significance than would be expected from a similarly sized sample drawn from a broader audience.

Another notable observation is that several respondents elected not to provide personally identifiable information beyond the information necessary for analysis. This reflects a common characteristic of cybersecurity research involving industry practitioners as mentioned before. Individuals working in security-sensitive environments frequently avoid unnecessary attribution due to confidentiality requirements, professional discretion, or organizational policies governing public discussion of defensive capabilities. The willingness of respondents to participate while remaining anonymous may therefore be interpreted as a reflection of professional norms within the cybersecurity field rather than a lack of engagement with the subject matter.

Taken together, the respondent profile suggests that the findings presented in this chapter should be viewed as expert practitioner perspectives rather than general market sentiment. While the sample size limits statistical generalization, the depth of professional experience represented within the respondent group provides a valuable foundation for exploring how deception technologies are currently perceived, adopted, and positioned within modern cybersecurity environments.

5.3 Results

5.3.1 Perceived Relevance of Deception Technologies

The responses generally indicated a positive perception of deception technologies, although they were not typically viewed as indispensable security controls. Respondents consistently considered deception technologies to be relevant, but their perceived value appeared to depend on the operational environment and organizational context. Compared with established controls such as SIEM, EDR, identity security, and network segmentation, deception technologies were more commonly viewed as complementary capabilities rather than foundational components of a security architecture. Figure 3 summarizes the distribution of respondent relevance scores.

Figure 3 illustrates that all respondents assigned a positive relevance score to deception technologies. Four respondents (44.4%) assigned the highest score of seven, while three respondents (33.3%) assigned a score of six and two respondents (22.2%) assigned a score of five. No respondent assigned a score below five on the seven-point scale.




Figure 3. Perceived relevance of deception technologies in modern cybersecurity environments. All respondents rated deception technologies as relevant, with 77.8% assigning a score of six or seven on a seven-point scale.

Several themes emerged from a closer examination of the responses. There was very little outright negativity toward deception technologies overall. Instead, reservations were primarily associated with practical considerations rather than disagreement with the underlying concept. Several respondents identified operational complexity, integration challenges, and governance considerations as potential barriers to adoption. While organizational maturity was not directly assessed by the questionnaire, these responses suggest that successful implementation may depend on broader organizational capabilities beyond the technology itself.

That is also interesting from a broader perspective. A large part of the literature presents deception as a way of shifting uncertainty toward the attacker and forcing interaction inside controlled environments (Almeshekeh & Spafford, 2014; Rowe & Rrushi, 2016). The questionnaire responses did not really challenge that thinking. If anything, they suggested that the concept itself is fairly well understood. The bigger issue appears to be how comfortably it fits into existing operational security work.

In practice, relevance alone does not automatically translate into adoption. Security teams already operate with large numbers of alerts, platforms, integrations, and overlapping controls. Under those conditions, even technologies that are seen as useful may remain

secondary if their value is difficult to demonstrate clearly or if implementation introduces additional operational complexity.

5.3.2 Architectural Positioning: Feature, Standalone Tool, or Integrated Control

Respondents were asked how deception technologies should be positioned within a cybersecurity architecture. As presented in Table 7, respondents most commonly viewed deception technologies as an integrated but distinct component of the broader security architecture. This position was selected by four respondents (44.4%), while three respondents (33.3%) considered the role of deception to be context dependent. Two respondents (22.2%) viewed deception primarily as a capability embedded within a broader security platform. No respondents identified deception as a standalone solution, and none considered it to have no meaningful relevance within modern cybersecurity environments.

Table 7. Architectural Positioning of Deception Technologies

Architectural Position	Number of Respondents	Percentage
Integrated but distinct capability	4	44.4%
Context dependent	3	33.3%
Platform capability	2	22.2%
Standalone solution	0	0.0%
No meaningful relevance	0	0.0%

The results suggest broad agreement that deception technologies should not operate independently from the wider security ecosystem. Instead, respondents consistently positioned deception alongside existing defensive capabilities such as Security Information and SIEM, EDR, identity security platforms, and orchestration tools.

This observation helps explain how cybersecurity practitioners evaluate the practical role of deception technologies within existing security architectures. Rather than replacing established security controls, deception is primarily viewed as a complementary capability that strengthens existing detection and response processes. This observation is consistent with the broader trend toward layered defensive architectures, where multiple technologies contribute different forms of visibility and defensive value.

At the same time, the results reveal an important distinction. While respondents generally favored integration, they did not appear to support the idea that deception should exist merely as a minor feature within another platform. The most common response was that deception should remain an integrated but distinct capability. This suggests that respondents recognize value in maintaining a dedicated deception function while still embedding it within existing operational workflows.

Previous research has similarly argued that the value of deception extends beyond individual alerts or isolated detection events. Deception technologies can introduce uncertainty, delay attacker progression, influence attacker decision-making, and support intelligence collection activities (Almeshekeh & Spafford, 2014; Rowe & Rrushi, 2016). However, the survey results indicate that cybersecurity practitioners may currently prioritize the operational integration of deception over these broader strategic objectives.

Taken together, the survey responses indicate that deception technologies are increasingly viewed as part of the security ecosystem rather than as independent security products. Their value appears strongest when integrated into existing monitoring, detection, and response processes while still retaining the unique characteristics that differentiate deception from traditional security controls.

5.3.3 Perceived Operational Applications of Deception Technologies

Respondents were asked to identify the operational applications where they believed deception technologies provide the greatest value. The results are summarized in Figure 4.

Figure 4 illustrates the operational use cases most frequently associated with deception technologies. Early compromise detection emerged as the most commonly selected application, followed by adversary behaviour analysis and lateral movement visibility. In contrast, strategic applications such as threat intelligence collection and detection validation were selected less frequently.




Figure 4. Perceived operational applications of deception technologies.

More specifically, seven respondents (77.8%) identified early compromise detection as a primary use case, while six respondents (66.7%) highlighted adversary behaviour analysis. Lateral movement visibility was selected by five respondents (55.6%), and four respondents (44.4%) associated deception technologies with defensive misdirection. Threat intelligence collection and detection validation received fewer selections, with two respondents (22.2%) and one respondent (11.1%) respectively.

These findings suggest that cybersecurity practitioners primarily associate deception technologies with practical detection and visibility benefits rather than with broader intelligence-gathering or adversary influence objectives. This observation aligns with several other findings in the survey, which indicate that respondents tend to view deception technologies through an operational rather than a strategic lens.

The results demonstrate that respondents primarily associate deception technologies with improving detection and visibility rather than influencing attacker behaviour directly. Early compromise detection, credential misuse identification, and lateral movement visibility all represent operational security challenges that remain difficult to address using conventional monitoring approaches alone.

The responses point toward that deception technologies are primarily positioned as complementary capabilities rather than standalone security solutions. Respondents appear to view deception primarily as a mechanism for making attacker activity more visible. The emphasis is therefore placed on detection enhancement rather than strategic manipulation of attacker behaviour.

The practical focus observed in the survey aligns with many of the challenges discussed earlier in this thesis. Modern attackers frequently rely on legitimate administrative tools, valid credentials, and trusted communication channels. Such activity may appear normal when viewed through traditional monitoring systems. Deception technologies address this challenge by creating controlled assets that should not attract legitimate interaction. As a result, interactions with deceptive resources may provide stronger indicators of malicious activity than behavioural anomalies alone.

The findings also reveal a notable difference between theoretical and operational perspectives. Academic literature frequently presents deception as a broader strategic capability capable of introducing uncertainty, delaying attackers, influencing decision-making, and supporting intelligence collection (Almeshekah & Spafford, 2014; Rowe & Rrushi, 2016). While respondents acknowledged some of these benefits, they were considerably less prominent than detection-focused use cases.

This difference does not diminish the value of deception technologies. Instead, it may reflect how cybersecurity capabilities are typically adopted in practice. Organizations often prioritize technologies that address immediate operational challenges before exploring more advanced strategic applications. Consequently, deception appears to enter security programs primarily as a detection capability before evolving into a broader intelligence and adversary engagement tool.

The results therefore suggest that deception technologies are currently valued most for their ability to improve visibility and reduce uncertainty within existing detection workflows. More strategic applications remain recognized, but they do not yet appear to be the primary driver of adoption among cybersecurity practitioners.

5.3.4 Defensive Philosophy: Immediate Containment or Controlled Observation

Respondents were asked whether deception technologies should primarily support immediate containment of attackers or controlled observation for intelligence collection purposes. The responses revealed a division in defensive philosophy. Three respondents (33.3%) favoured immediate detection and removal, three respondents (33.3%) supported a context-dependent

approach, and three respondents (33.3%) considered controlled observation to be appropriate under certain conditions.

Table 8 summarizes respondents' preferred response approaches following detection through deception technologies.

Table 8. Preferred response approach following detection through deception technologies. Responses were evenly distributed between immediate containment, controlled observation, and context-dependent approaches, indicating that practitioners do not share a single view regarding how deception-generated detections should be handled operationally.

Preferred Response Approach	Number of Respondents	Percentage
Immediate detection and removal	3	33.3%
Context-dependent approach	3	33.3%
Controlled observation	3	33.3%

The distribution presented in Table 8 suggests that there is no clear consensus regarding how deception-generated detections should be handled operationally. Instead, respondents appear to balance the potential intelligence value of observation against the risks associated with allowing attacker activity to continue, even within controlled environments.

This observation provides insight into how cybersecurity practitioners define the purpose of deception technologies within modern security operations. If deception is used solely to trigger faster containment, it remains closely aligned with traditional detection and response processes. However, if deception is used to observe attacker behaviour within controlled environments, it can support intelligence collection, behavioural analysis, and a deeper understanding of adversary objectives.

The responses therefore suggest that deception technologies may serve multiple purposes depending on organizational priorities, risk tolerance, and operational maturity. Rather than representing a single defensive philosophy, deception appears capable of supporting both rapid containment and intelligence-driven approaches to incident response.

The findings help clarify how cybersecurity practitioners define the purpose of deception technologies within modern security operations. If deception is used solely to trigger faster containment, it remains closely aligned with traditional detection and response processes. However, if deception is used to observe attacker behaviour within controlled environments,

it can support intelligence collection, behavioural analysis, and a deeper understanding of adversary objectives.

The respondents who favored immediate containment appear to reflect a more risk-averse operational perspective. Rapid removal reduces uncertainty and minimizes the possibility of unintended consequences. This approach may be particularly appropriate in environments where service availability, safety, or operational continuity are critical concerns. Healthcare organizations, industrial environments, public services, and critical infrastructure operators may have limited tolerance for prolonged observation activities due to the potential impact of operational disruption.

Conversely, respondents supporting controlled observation appear to recognize the intelligence value that deception technologies can provide. Advanced intrusions frequently unfold over extended periods and often involve multiple stages including reconnaissance, credential acquisition, persistence, lateral movement, and data access. Observing attacker behaviour within controlled environments may therefore provide valuable information regarding objectives, tactics, techniques, and procedures. Previous incident response research has shown that sophisticated intrusion campaigns frequently involve adaptation and progression over time rather than isolated malicious events (Mandiant, 2023; Hutchins et al., 2011).

Perhaps the most notable result is that one-third of respondents adopted a context-dependent position. This suggests that cybersecurity practitioners do not necessarily view containment and observation as mutually exclusive approaches. Instead, the appropriate response may depend on factors such as organizational risk tolerance, asset criticality, confidence in containment mechanisms, and the maturity of the security program.

This interpretation aligns closely with the broader literature on cyber deception. Research has suggested that the value of deception extends beyond detection and may include intelligence collection, adversary analysis, and the creation of uncertainty for attackers (Almeshekeh & Spafford, 2014; Rowe & Rrushi, 2016). However, realizing these benefits requires environments that are sufficiently isolated and governed to support controlled observation without introducing unacceptable risk.

The findings therefore suggest that governance and segmentation play a central role in the practical application of deception technologies. Whether observation is considered acceptable

depends largely on the organization's ability to maintain separation between deceptive assets and operational systems. In this context, the question is not simply whether attackers should be observed or removed, but whether the environment provides sufficient control to make observation both safe and valuable.

5.3.5 Adoption Barriers and Practical Concerns

Respondents were also asked to identify the factors they perceived as the most notable barriers to the adoption of deception technologies. The results are summarized in Table 9.

Table 9. Perceived adoption barriers and concerns related to deception technologies

Adoption Barrier	Number of Respondents	Percentage
Operational complexity	5	55.6%
Legal ambiguity	3	33.3%
Exposure of defensive strategy	3	33.3%
Escalation or misattribution risk	3	33.3%
Resistance or lack of trust	2	22.2%
Ethical concerns	2	22.2%

Table 9 illustrates the primary concerns associated with the adoption of deception technologies. Operational complexity emerged as the most frequently identified barrier, followed by legal ambiguity, exposure of defensive strategy, and concerns related to escalation or misattribution.

More specifically, five respondents (55.6%) identified operational complexity as a major concern. Legal ambiguity, exposure of defensive strategy, and the risk of escalation or misattribution were each selected by three respondents (33.3%). Resistance or lack of organizational trust and ethical concerns were selected by two respondents (22.2%) respectively.

These results suggest that the principal barriers to adoption are operational rather than conceptual. While respondents generally recognized the value of deception technologies, concerns regarding implementation complexity, governance requirements, and potential

organizational impacts appear to influence how readily such capabilities are adopted in practice.

The results suggest that the primary obstacles to adoption are not related to a lack of perceived value. Instead, respondents appear to be concerned with the practical realities of deploying, managing, and governing deception technologies within operational environments.

Operational complexity emerged as the most substantial concern. Effective deception involves more than simply deploying decoy systems or deceptive data. Deceptive assets must be sufficiently realistic to attract attacker attention, while remaining isolated from production systems to avoid introducing unnecessary risk. In addition, organizations must continuously monitor deceptive assets, maintain their credibility, and ensure that alerts generated by those assets are incorporated into existing security processes. These requirements increase both implementation effort and ongoing operational overhead.

The findings also highlight the importance of integration. Modern security operations typically rely on multiple technologies working together. SIEM platforms aggregate and analyze security events from across the environment, while EDR solutions monitor activity on individual devices. SOAR platforms are often used to coordinate and automate response activities. If deception-generated alerts are not integrated into these existing workflows, their practical value may be extensively reduced. Several respondents therefore appeared to view successful integration as a prerequisite for meaningful adoption rather than as a secondary consideration.

Another recurring theme was uncertainty regarding how the value of deception technologies should be measured. Traditional security controls are often evaluated through relatively straightforward metrics such as malware detections, blocked attacks, detection coverage, or response times. Deception technologies can be more difficult to evaluate because some of their benefits are indirect. For example, a deception deployment may improve visibility into attacker behaviour, delay attacker progression, or increase analyst confidence during investigations. While these outcomes may provide substantial defensive value, they are often more difficult to quantify using conventional performance measures. This may help explain why organizations recognize the potential value of deception while still hesitating to invest in wider deployment.

Legal and ethical concerns were identified less frequently than operational concerns, but they remain relevant. Deception technologies intentionally create resources that appear genuine while serving defensive purposes. Although this approach is generally accepted within cybersecurity practice, it can raise questions regarding monitoring, privacy, synthetic data usage, and organizational accountability. Such considerations may become particularly important in highly regulated sectors where compliance obligations influence technology adoption decisions.

The responses therefore suggest that successful adoption depends not only on technical capability but also on governance, operational maturity, and organizational readiness. The challenge does not appear to be convincing cybersecurity practitioners that deception can provide value. Rather, the challenge lies in demonstrating how that value can be implemented, managed, measured, and governed within existing security programs.

These findings are broadly consistent with previous research, which has identified integration, realism, maintenance requirements, and operational management as recurring challenges associated with deception deployments (Fraunholz et al., 2018; Jajodia et al., 2016). The survey results therefore reinforce the view that adoption barriers are primarily practical rather than conceptual. Deception technologies are generally recognized as useful, but organizations may struggle to operationalize them effectively without sufficient resources, processes, and supporting infrastructure.

5.4 Deeper Interpretation of Findings

5.4.1 Detection as the Primary Value Driver

One of the most consistent observations emerging from the survey is that respondents primarily associate deception technologies with improving detection rather than fundamentally changing defensive strategy. This interpretation is supported by the results presented in Section 5.3.3, where early compromise detection, adversary behaviour analysis, and lateral movement visibility were identified as the most central operational applications of deception technologies.

This finding is noteworthy because much of the academic literature describes deception as a broader strategic capability. Previous research has argued that deception can influence

attacker decision-making, introduce uncertainty, delay progression, and support intelligence collection activities (Almeshekah & Spafford, 2014; Rowe & Rrushi, 2016). While these capabilities remain relevant, they were not the aspects most frequently emphasized by respondents. Instead, the survey indicates that cybersecurity practitioners primarily value deception for its ability to improve visibility into attacker activity and strengthen existing detection processes.

This should not be interpreted as a negative finding. New cybersecurity capabilities are often adopted because they address specific operational challenges before their broader strategic implications are fully realized. The presented survey results suggest that deception technologies currently derive much of their perceived value from addressing practical problems that security teams encounter in daily operations. Examples include identifying credential misuse, improving visibility into lateral movement, reducing uncertainty during investigations, and increasing confidence in detection outcomes.

The findings also highlight a distinction between strategic potential and operational perception. A capability may provide strategic advantages while still being evaluated according to its day-to-day operational benefits. In practice, security teams are often responsible for managing large volumes of alerts, investigating suspicious activity, and prioritizing limited resources. Within this context, a technology that improves alert quality or provides more reliable indicators of malicious activity may be perceived primarily as a detection enhancement, even if its theoretical applications extend considerably further.

This interpretation aligns closely with the discussion presented in Chapter 4 regarding SIEM and EDR technologies. Traditional security controls generate large volumes of telemetry from endpoints, network devices, authentication systems, and applications. Although this information provides valuable visibility, determining which events indicate genuine malicious activity often requires extensive analysis and correlation. Deception technologies differ because they introduce controlled assets that should not normally attract legitimate interaction. As a result, interactions with deceptive resources frequently provide more direct indicators of potentially malicious activity than anomaly-based detections alone (Fraunholz et al., 2018).

From this perspective, the primary value of deception may not be the generation of additional security data, but the improvement of signal quality. Rather than increasing visibility in a

purely quantitative sense, deception technologies appear to help security teams distinguish meaningful activity from the larger volume of operational noise that characterizes modern cybersecurity environments. This observation is reflected throughout the survey results and represents one of the strongest themes emerging from the empirical findings.

Taken together, the results suggest that cybersecurity practitioners currently view deception technologies less as a standalone strategic capability and more as a practical mechanism for improving the effectiveness of existing detection and response processes. Whether this perception changes as deception technologies mature remains an open question, but the findings indicate that operational value currently represents the primary driver of perceived relevance.

5.4.2 Theory–Practice Gap

One of the more interesting observations emerging from the survey is the apparent difference between how deception technologies are described in academic literature and how they are currently perceived in operational environments. While respondents consistently recognized the value of deception technologies, the use cases they emphasized were generally more focused on detection, visibility, and operational support than on strategic adversary manipulation.

As discussed in earlier chapters, the literature frequently presents deception as a capability that can influence attacker decision-making, introduce uncertainty, delay progression, support intelligence collection, and shape adversary behaviour (Almeshekah & Spafford, 2014; Rowe & Rrushi, 2016). These concepts represent some of the most distinctive characteristics of deception-based defense and are often cited as advantages that differentiate deception from traditional security controls.

The survey results suggest a somewhat narrower operational focus. Respondents most frequently associated deception technologies with early compromise detection, adversary behaviour analysis, credential misuse detection, and visibility into lateral movement. Although these applications remain valuable, they are generally aligned with improving existing detection and response processes rather than fundamentally altering the relationship between attacker and defender.

This observation is important because it highlights that the relevance of deception technologies is not determined solely by their technical capabilities. Adoption is also influenced by organizational priorities, available resources, operational maturity, governance requirements, and existing security processes. A capability may offer substantial theoretical advantages while still being utilized in a more limited manner if organizations lack the resources, confidence, or operational framework necessary to support broader implementation.

The findings therefore suggest that many organizations currently approach deception as an enhancement to existing security controls rather than as a standalone adversary engagement capability. This interpretation is consistent with the survey results presented earlier, where respondents repeatedly emphasized integration with SIEM, EDR, and broader security operations workflows. The value of deception appears to be strongest when it contributes to established defensive processes rather than when it operates independently of them.

Importantly, this should not be interpreted as evidence that the academic literature is unrealistic. A more likely explanation is that operational adoption remains at an earlier stage of maturity than the theoretical models discussed in the literature. Organizations often adopt new security capabilities incrementally, beginning with narrowly defined operational objectives before expanding into more advanced use cases. Deception technologies may currently be following a similar pattern.

If this interpretation is correct, the survey findings may indicate that operational adoption is still evolving. The broader strategic applications described in the literature remain relevant, but they do not yet appear to be the primary driver of practical deployment. At present, deception technologies seem to occupy a position between specialized detection capability and broader strategic security function. How this position develops over time will likely depend on improvements in integration, automation, governance, and organizational confidence in deception-based approaches.

5.4.3 Importance of Integration

Another important interpretation emerging from the survey is that the value of deception technologies appears to depend heavily on integration with existing security processes and

technologies. This observation is supported by several findings presented earlier in the chapter. Respondents frequently positioned deception as an integrated capability rather than a standalone solution, while operational complexity and integration challenges were identified as some of the most significant barriers to adoption.

This finding is considerable because it shifts the discussion away from the technology itself and toward the environment in which the technology operates. The question is not simply whether an organization should deploy deception technologies, but whether it can effectively incorporate them into existing security operations. In practice, the value of a deception deployment is determined not only by the quality of the deceptive assets but also by how effectively the resulting information can be used by defenders.

The survey results suggest that cybersecurity practitioners evaluate deception within the context of broader defensive workflows rather than as an isolated capability. Modern security operations typically involve multiple layers of monitoring, analysis, investigation, and response. Alerts generated by deception technologies must therefore be incorporated into these existing processes if they are to produce meaningful operational value.

For example, a deceptive credential that triggers an alert may provide useful information on its own. However, the value of that alert increases considerably when it can be correlated with additional evidence such as endpoint activity, authentication events, network connections, threat intelligence data, or known attacker techniques. In this context, deception functions less as an independent detection mechanism and more as a source of high-confidence indicators that strengthen analytical decision-making.

This interpretation is also consistent with the findings presented in Chapter 4. SIEM platforms and EDR solutions are designed to collect and analyze large volumes of security telemetry. One of the challenges associated with these technologies is determining which events represent genuine threats and which represent normal operational activity. Deception technologies contribute value by introducing events that are often easier to interpret because they originate from interactions with assets that should not normally be accessed. The resulting alerts may therefore provide stronger contextual signals than many traditional anomaly-based detections.

The findings also suggest that deception technologies should not be evaluated solely as a product category. Instead, they should be viewed as one component within a broader

defensive operating model. Organizations with mature security operations, established incident response procedures, and effective integration capabilities may be better positioned to realize the benefits of deception. Conversely, organizations lacking these supporting processes may experience additional complexity without achieving the same level of operational value.

This observation may also help explain why respondents generally viewed deception as relevant while simultaneously identifying practical barriers to adoption. The challenge does not appear to be a lack of confidence in the concept itself. Rather, the challenge lies in transforming deception-generated information into actionable intelligence that supports security operations. This requires integration, governance, operational processes, and sufficient organizational maturity.

Taken together, the survey results indicate that successful deception deployments depend as much on the surrounding security architecture as on the deception technology itself. The effectiveness of deception therefore appears to be closely linked to an organization's ability to integrate, operationalize, and act upon the information generated by deceptive assets. In this sense, deception may be best understood not as a standalone security capability, but as a force multiplier that enhances the effectiveness of existing detection and response processes.

5.4.4 Organizational Maturity and Perceived Relevance

The survey findings suggest that the perceived value of deception technologies is influenced by the maturity of the organization deploying them. While respondents generally considered deception technologies relevant, they did not consistently position them as foundational security capabilities. Instead, many respondents viewed deception as a specialized capability that becomes increasingly valuable as organizational security programs mature.

This interpretation is supported by the findings presented earlier in the chapter. Although respondents assigned high relevance scores to deception technologies, the majority did not describe them as a core requirement for all organizations. Instead, deception was more frequently positioned as a complementary capability that supports existing security controls and enhances visibility within established defensive processes.

One possible explanation is that deception technologies often rely on capabilities that are already in place. Effective deployment typically requires monitoring infrastructure, incident response processes, governance mechanisms, and integration with broader security operations. Organizations lacking these foundational capabilities may struggle to derive meaningful value from deception, regardless of the quality of the underlying technology.

For less mature organizations, relatively simple forms of deception may provide practical benefits with limited operational overhead. Examples include honeytokens, deceptive credentials, decoy user accounts, or monitored file shares. These techniques can generate useful security signals while requiring comparatively little infrastructure or administrative effort. In such environments, deception may serve primarily as a detection enhancement rather than a broader strategic capability.

More mature organizations may be able to adopt deception in more sophisticated ways. Examples include dedicated deception environments, controlled observation of attacker activity, adaptive deception mechanisms, and intelligence-driven use of deceptive assets. These approaches generally require stronger governance, greater operational maturity, and the ability to integrate deception into existing detection, response, and threat intelligence workflows.

This interpretation also helps explain why several respondents described deception as context dependent. The responses do not necessarily suggest uncertainty regarding the value of deception technologies themselves. Rather, they suggest recognition that successful deployment depends heavily on the conditions surrounding the technology. Organizational maturity, staffing levels, security architecture, governance structures, and risk tolerance all influence how effectively deception can be implemented and maintained.

The findings therefore indicate that adoption barriers are not primarily driven by a lack of awareness. All respondents considered deception technologies relevant to some degree, and none dismissed the concept outright. Instead, the obstacles appear to be associated with practical implementation factors such as integration complexity, ownership, governance, operational resources, and uncertainty regarding how value should be measured. These factors become increasingly paramount as organizations move from experimental or limited deployments toward broader operational use.

This observation aligns with previous research suggesting that successful deception deployments require more than technical implementation alone. Factors such as integration, operational management, realism, and organizational support have repeatedly been identified as critical determinants of long-term success (Fraunholz et al., 2018; Jajodia et al., 2016). The survey findings therefore suggest that the relevance of deception technologies cannot be evaluated independently of the environment in which they are deployed.

Taken together, the results indicate that deception technologies may be best understood as maturity-dependent capabilities. Their value does not appear to be uniform across all organizations. Instead, the benefits that can be realized from deception are influenced by the broader maturity of the security program and the organization's ability to operationalize, govern, and integrate deception into existing defensive processes.

5.4.5 Defensive Philosophy and Attacker Observation

The findings related to defensive philosophy reveal a valuable aspect of deception technologies that extends beyond technical implementation. Unlike many traditional security controls, deception introduces questions regarding how organizations should respond once attacker activity has been identified. The survey results demonstrated no clear consensus on this issue, with respondents divided between immediate containment, controlled observation, and context-dependent approaches.

The responses reveal a contrast between traditional incident response principles, which generally prioritize immediate containment, and the intelligence-gathering opportunities that deception technologies may provide through controlled observation of attacker behaviour. The objective is generally to reduce risk as quickly as possible by limiting attacker access and preventing further progression within the environment. Deception technologies introduce an additional possibility. Rather than immediately removing an attacker, organizations may choose to observe activity within a controlled environment in order to collect intelligence regarding objectives, techniques, and behaviour.

The survey results suggest that cybersecurity practitioners do not view this decision uniformly. This is not unexpected. The appropriate response may vary considerably depending on organizational context, asset criticality, regulatory requirements, and confidence

in containment mechanisms. In environments where safety, operational continuity, or regulatory compliance are primary concerns, immediate containment may remain the preferred approach. Conversely, organizations with mature security programs and well-isolated deception environments may be more willing to conduct limited observation activities when doing so provides valuable intelligence.

This observation aligns with previous research on cyber deception, which has highlighted the intelligence-gathering potential of controlled engagement with adversaries (Almeshekah & Spafford, 2014; Rowe & Rrushi, 2016). At the same time, the literature consistently emphasizes the importance of maintaining strict separation between deceptive environments and operational systems. Without effective segmentation and governance, the risks associated with observation may outweigh the potential benefits.

The survey findings therefore suggest that deception technologies introduce considerations that are not typically present in traditional detection and response workflows. Once a deception-generated alert occurs, the organization may need to decide whether the greatest value lies in immediate containment or in obtaining additional information regarding attacker behaviour. This decision is influenced not only by technical factors but also by organizational policy, risk tolerance, governance requirements, and operational maturity.

An essential implication is that deception should not be viewed solely as a security engineering capability. The decision to observe or immediately remove an attacker involves broader organizational considerations regarding acceptable risk, intelligence value, and defensive objectives. Consequently, deception technologies may require closer alignment between technical teams, incident response personnel, security leadership, and governance functions than many traditional security controls.

Taken together, the survey results indicate that deception technologies do more than generate additional detection opportunities. They also introduce new decision points into the incident response process. Whether organizations choose containment, observation, or a combination of both depends largely on the context in which deception is deployed and the objectives the organization seeks to achieve through its defensive strategy.

5.4.6 Future Potential of AI-Driven Deception

Although the survey did not specifically measure the adoption of AI-driven deception technologies, the findings provide useful insight into how such capabilities may be perceived in the future. Several of the themes identified throughout the questionnaire—including the desire for higher-quality detection signals, improved visibility into attacker behaviour, and stronger integration with existing security processes—are closely aligned with the objectives commonly associated with adaptive and AI-driven deception.

As discussed in Chapter 3, AI-driven deception refers to the use of automation, behavioural analysis, and adaptive mechanisms to modify deceptive assets based on observed activity. The objective is not simply to deploy static decoys, but to create environments that can respond to attacker behaviour and maintain credibility over time. In theory, this may improve realism, reduce predictability, and provide defenders with more relevant information regarding attacker actions.

The survey findings indirectly support the relevance of these objectives. Respondents consistently emphasized the importance of detection quality, visibility, and operational usefulness. These are areas where adaptive deception may offer advantages compared to traditional static deployments. By adjusting deceptive assets in response to observed behaviour, adaptive approaches may be able to generate more relevant interactions and provide additional context for security investigations.

At the same time, the survey findings also highlight why widespread adoption may remain limited in the near term. Many of the barriers identified by respondents become even more noteworthy when automation and adaptive decision-making are introduced. Operational complexity, integration requirements, governance concerns, and uncertainty regarding measurable value were already identified as challenges affecting deception adoption. AI-driven deception does not eliminate these challenges and may, in some cases, increase them.

For example, adaptive deception requires reliable telemetry, effective segmentation, strong governance controls, and confidence in automated decision-making processes. Organizations must be able to trust that changes made by the system remain aligned with defensive objectives and do not introduce unintended risk. As the level of automation increases,

questions regarding oversight, accountability, and operational control become increasingly important.

This interpretation suggests that AI-driven deception should not be viewed as an immediate mainstream capability. The survey results indicate that many organizations are still addressing more fundamental questions regarding where deception technologies fit within security architectures, how they should be integrated into operational workflows, and how their value should be measured. These challenges are likely to influence the adoption of adaptive deception just as they currently influence the adoption of more traditional deception approaches.

Nevertheless, the data suggest that AI-driven deception aligns with the broader direction of cybersecurity development. Security environments continue to increase in scale and complexity, while attackers become more adaptive and difficult to distinguish from legitimate activity. In such environments, the ability to dynamically adjust deceptive assets and provide more contextually relevant signals may become increasingly valuable.

Taken together, the survey results suggest that AI-driven deception is best understood as an emerging capability rather than an established operational practice. Its underlying objectives appear consistent with the needs identified by cybersecurity practitioners, but its widespread adoption will likely depend on improvements in integration, governance, automation, and organizational maturity. Consequently, AI-driven deception may represent the future direction of deception technologies, even if that future has not yet been fully realized in current operational environments.

5.5 Synthesis of Findings

Taken together, the survey results point toward that deception technologies occupy an intermediate position within modern cybersecurity environments. They are no longer viewed as purely experimental or theoretical concepts, yet they have not reached a level of adoption where they are consistently regarded as foundational security controls. All respondents considered deception technologies relevant to some degree, but most positioned them as complementary or specialized capabilities rather than as core elements of every security program.

One of the clearest themes emerging from the survey is the value of deception as a mechanism for improving detection quality. Respondents most frequently associated deception technologies with early compromise detection, adversary behaviour analysis, credential misuse detection, and visibility into lateral movement. These findings indicate that deception is primarily valued for its ability to reduce uncertainty and provide stronger indicators of potentially malicious activity within existing security operations.

At the same time, the responses point toward a difference between the broader strategic potential described in the literature and the way deception technologies are commonly perceived in practice. Academic research frequently presents deception as a means of influencing attacker behaviour, creating uncertainty, supporting intelligence collection, and introducing asymmetry between attacker and defender (Almeshekah & Spafford, 2014; Rowe & Rrushi, 2016). The survey results indicate that cybersecurity practitioners currently emphasize more operational applications focused on detection and visibility. This does not necessarily contradict the literature, but it suggests that practical adoption may be concentrated on narrower use cases than those commonly discussed in theoretical models.

A second recurring theme is that the effectiveness of deception technologies appears closely linked to organizational context. Respondents consistently highlighted factors such as integration, governance, operational complexity, and organizational maturity. These observations suggest that deception technologies derive much of their value from the environment in which they are deployed. Without effective integration into monitoring, detection, and response processes, deception risks becoming an isolated source of alerts. When supported by mature operational processes and governance structures, however, deception may become a meaningful component of a layered defensive architecture.

The findings also indicate that relevance and necessity should not be treated as equivalent concepts. Although respondents generally viewed deception as valuable, most did not describe it as an essential capability for all organizations. Instead, the perceived value of deception appeared to increase alongside organizational maturity and the ability to operationalize the technology effectively. This observation helps explain why deception was frequently described as context dependent despite receiving consistently positive relevance assessments.

Finally, the results suggest that the future role of deception technologies is likely to be increasingly integrated rather than independent. Respondents repeatedly emphasized the importance of incorporating deception into existing security workflows, monitoring platforms, and response processes. Consequently, the long-term evolution of deception may be less about establishing a separate category of security tools and more about embedding deception-based capabilities into broader security operations, identity systems, endpoint monitoring platforms, and threat intelligence processes.

Overall, the survey findings indicate that deception technologies are moving toward a more established role within cybersecurity practice. However, their adoption and effectiveness remain closely influenced by organizational maturity, integration capabilities, governance structures, and the ability to translate deception-generated information into actionable security outcomes.

5.5.1 Implications for the Research Questions

The purpose of the empirical study was to examine how cybersecurity practitioners perceive the relevance, role, and adoption of deception technologies within modern cybersecurity environments. The findings provide insight into each of the three research questions defined in Chapter 1.

RQ1: How Relevant Are Deception Technologies in Modern Cybersecurity Environments?

The survey findings indicate that deception technologies are generally considered relevant by cybersecurity practitioners. All respondents assigned a positive relevance score, and none viewed deception as having no meaningful role within cybersecurity environments. The perceived value was particularly strong in areas where organizations struggle to identify lateral movement, detect credential misuse, or distinguish malicious activity from normal operational behaviour.

At the same time, the findings suggest that relevance should not be interpreted as universal necessity. While respondents consistently recognized value in deception technologies, most did not position them as essential security capabilities for every organization. Instead,

relevance appeared to depend on factors such as organizational maturity, security architecture, operational requirements, and risk profile.

RQ2: How Do Cybersecurity Practitioners Position Deception Technologies Relative to Traditional Defensive Controls?

The results suggest that cybersecurity practitioners do not view deception technologies as replacements for traditional security controls. Instead, deception was consistently positioned as a complementary capability that enhances existing detection and response processes.

Respondents most frequently associated deception with improved visibility, higher-quality detection signals, and greater confidence when investigating suspicious activity. This position aligns closely with the broader findings presented throughout the study. Traditional controls such as SIEM, EDR, network segmentation, and identity security remain fundamental components of defensive architectures, while deception appears to provide additional context and reduce uncertainty surrounding attacker behaviour.

The findings therefore suggest that practitioners view deception as part of a layered defense strategy rather than as a standalone security solution.

RQ3: What Factors Influence the Adoption and Operational Use of Deception Technologies?

The survey findings indicate that the primary barriers to adoption are practical rather than conceptual. Respondents generally accepted that deception technologies can provide value, but they also identified several challenges that may limit broader implementation.

The most frequently cited concerns related to operational complexity, integration requirements, governance considerations, and uncertainty regarding how value should be measured. These responses suggest that organizations are often less concerned about whether deception works and more concerned about how deception can be deployed, managed, integrated, and justified within existing security programs.

The results also indicate that organizational maturity plays an important role. Organizations with established monitoring capabilities, mature incident response processes, and effective governance structures may be better positioned to realize the benefits of deception technologies than organizations that are still developing foundational security capabilities.

Taken together, the results suggest that the adoption of deception technologies depends not only on technical capability, but also on the broader organizational environment in which the technology is deployed.

5.6 Survey Limitations

Several limitations should be considered when interpreting the findings of this study.

The first limitation relates to the size of the respondent group. Although the survey was distributed to 47 cybersecurity practitioners, only nine completed responses were received, resulting in a response rate of approximately 19.1%. While the respondents represented a highly experienced group of cybersecurity practitioners, the sample size remains relatively small. Consequently, the findings should not be interpreted as statistically representative of the broader cybersecurity industry. Instead, they should be understood as indicative of how experienced practitioners currently perceive deception technologies and their role within cybersecurity environments.

A second limitation concerns the nature of the data itself. The findings are based on professional experience, perception, and judgment rather than direct measurement of technical effectiveness. Respondents evaluated deception technologies through the lens of their own operational experience, organizational context, and exposure to cybersecurity practices. Such perspectives provide valuable insight into real-world adoption and practical relevance, but they do not directly measure detection performance, attacker behaviour, or defensive effectiveness under controlled conditions.

A third limitation relates to the scope of the study. The questionnaire was intentionally designed to examine perceptions of relevance, adoption, and operational use rather than to compare specific vendors, products, deployment models, or technical architectures. As a result, the findings provide insight into how deception technologies are viewed and positioned in practice, but they do not determine which implementations are most effective or appropriate in particular environments.

The findings may also be influenced by the professional backgrounds of the respondents. Although efforts were made to include practitioners with different experiences and responsibilities, cybersecurity professionals often evaluate technologies through the

requirements and challenges of their own operational environments. Organizations with different levels of maturity, sector-specific requirements, or regulatory obligations may therefore arrive at different conclusions regarding the value and applicability of deception technologies.

Future research could address these limitations by expanding the respondent pool, incorporating additional organizational perspectives, or combining practitioner surveys with case studies, controlled experiments, or longitudinal observations of deception deployments. Such approaches could provide deeper insight into how deception technologies perform in practice and how their effectiveness evolves over time.

Despite these limitations, the empirical study provides useful insight into how experienced cybersecurity practitioners currently perceive deception technologies. While the findings cannot be generalized across the broader cybersecurity industry, they offer valuable context for understanding how deception technologies are viewed, adopted, and positioned within contemporary security environments.

6 Discussion

6.1 Interpreting the Findings in Context

The findings of this study suggest that deception technologies have progressed beyond the purely theoretical stage and are increasingly recognized as a legitimate component of modern cybersecurity practice. At the same time, the results indicate that their practical role remains inconsistent across organizations. Although respondents generally viewed deception technologies as relevant, there was less agreement regarding how they should be positioned within defensive architectures and how extensively they should be used in operational environments.

One of the most notable observations is that respondents primarily discussed deception through practical security challenges rather than through broader strategic concepts. The strongest associations were with early compromise detection, lateral movement visibility, credential misuse detection, and improved confidence in identifying malicious activity. These findings suggest that deception technologies are currently valued most for their ability to improve visibility and reduce uncertainty within existing security operations.

This interpretation differs somewhat from the emphasis found in portions of the academic literature. Research on cyber deception frequently highlights objectives such as influencing attacker decision-making, introducing uncertainty, shaping adversary behaviour, and creating asymmetry between attacker and defender (Almeshekeh & Spafford, 2014; Pawlick et al., 2019; Rowe & Rrushi, 2016). While respondents did not reject these ideas, they rarely emerged as primary motivations for deployment. Instead, the empirical results suggest that organizations are more likely to adopt deception technologies to address immediate operational challenges than to pursue broader strategic objectives.

This difference should not necessarily be interpreted as a contradiction between theory and practice. A more plausible explanation is that cybersecurity capabilities often mature incrementally. Technologies are commonly introduced to solve specific operational problems before their broader strategic applications are fully explored. In this respect, deception technologies appear to follow a pattern similar to many other defensive capabilities. Their adoption is initially driven by practical security needs, while more advanced applications emerge as organizations gain experience, confidence, and operational maturity.

The findings also reinforce the importance of organizational context. Throughout the survey, respondents repeatedly highlighted integration, governance, operational complexity, and security maturity as factors influencing the success of deception deployments. This observation is consistent with previous research suggesting that deception effectiveness depends not only on technical implementation but also on how well deceptive assets are incorporated into broader security processes (Fraunholz et al., 2018; Jajodia et al., 2016). The value of deception therefore appears to be closely connected to the environment in which it operates rather than to the technology alone.

Another important observation concerns the relationship between deception and traditional security controls. The findings suggest that cybersecurity practitioners generally do not view deception as a replacement for technologies such as SIEM, EDR, network segmentation, or identity security. Instead, deception is primarily perceived as a complementary capability that enhances the effectiveness of existing controls by generating higher-confidence indicators of potentially malicious activity. This supports the argument that the future role of deception is likely to be increasingly integrated within broader security architectures rather than developing as an independent security discipline.

The findings also provide some insight into the future development of deception technologies. Respondents consistently identified challenges related to ambiguity, attacker adaptation, and the difficulty of distinguishing malicious behaviour from legitimate activity. These challenges align closely with the objectives of adaptive and AI-driven deception approaches discussed earlier in the thesis. Although such capabilities remain relatively immature, the survey results suggest that the underlying problems they seek to address are already recognized by cybersecurity practitioners. Consequently, adaptive deception may represent a natural progression in the evolution of deception technologies, particularly within organizations that possess the operational maturity required to support more advanced implementations.

Taken together, the results suggest that deception technologies currently occupy a transitional position within cybersecurity practice. They are no longer viewed as experimental concepts, yet they have not reached the level of standardization associated with more established security controls. Their future relevance is therefore likely to depend less on demonstrating technical feasibility and more on addressing practical concerns related to integration, governance, operational management, and measurable value. In this respect, the empirical findings support the broader conclusion that the challenge facing deception technologies is no

longer whether they can provide value, but how that value can be consistently realized within modern cybersecurity environments.

6.2 Gap Between Theory and Practice

One of the most notable observations emerging from this study is the apparent difference between the theoretical potential of deception technologies and their practical implementation within operational environments. The literature frequently presents deception as a strategic capability capable of influencing attacker behaviour, creating uncertainty, supporting intelligence collection, and altering the balance of information between attacker and defender (Almeshekeh & Spafford, 2014; Pawlick et al., 2019; Rowe & Rrushi, 2016). The survey findings, however, suggest that organizations currently apply deception in a more focused and operational manner.

Respondents most commonly associated deception technologies with improving detection quality, identifying credential misuse, increasing visibility into lateral movement, and reducing uncertainty during investigations. These use cases are important and provide tangible operational value, but they represent a narrower application of deception than many of the broader strategic models described in the academic literature.

Several factors may help explain this difference. One explanation is that cybersecurity technologies are rarely adopted in their most advanced form from the outset. Organizations often implement new capabilities incrementally, beginning with clearly defined operational objectives before expanding into more complex use cases. From this perspective, deception technologies may currently be following a natural maturity path in which detection-oriented deployments serve as an entry point to broader strategic applications.

The survey findings also highlight the importance of integration. Deception-generated alerts are most valuable when they can be incorporated into existing monitoring, investigation, and response workflows. If interactions with deceptive assets remain isolated from SIEM, EDR, or incident response processes, even highly reliable indicators may fail to generate meaningful defensive outcomes. This observation reinforces the argument that deception derives much of its value from its relationship with the surrounding security architecture rather than from the deceptive assets themselves.

Another contributing factor appears to be organizational risk tolerance. A recurring theme throughout the survey was the tension between intelligence collection and immediate containment. While observing attacker activity within controlled environments may provide valuable insight into objectives, tactics, and behaviour, it also introduces additional operational considerations. Organizations responsible for critical services, industrial operations, healthcare systems, or other high-consequence environments may have limited tolerance for prolonged observation activities, regardless of their potential intelligence value. Consequently, the practical use of deception may remain constrained even when its theoretical benefits are well understood.

The findings also suggest that organizational maturity plays a substantial role in determining how deception technologies are used. Mature security programs are generally better positioned to support the governance, monitoring, integration, and operational processes required for advanced deception deployments. Organizations with less developed security capabilities may instead focus on simpler deception techniques that provide immediate detection benefits while minimizing operational overhead. This may help explain why respondents frequently described deception as context dependent despite assigning consistently positive relevance ratings.

Importantly, the observed difference between theory and practice should not be interpreted as a contradiction. The survey findings do not suggest that the strategic concepts discussed in the literature lack value. Rather, they indicate that practical adoption currently emphasizes those aspects of deception that provide the most immediate and measurable operational benefits. The broader strategic applications remain relevant, but they appear to be less common in everyday defensive practice.

Taken together, the data indicates that the gap between theory and practice is best understood as a difference in maturity rather than a disagreement regarding value. The academic literature often describes what deception technologies may ultimately enable, whereas practitioners are primarily focused on how those capabilities can be applied within existing operational environments. As organizations continue to develop their security programs and improve integration between defensive technologies, the distinction between theoretical potential and operational use may gradually become less pronounced.

6.3 Practical Implications

The findings of this thesis suggest several practical implications for organizations considering the adoption of deception technologies. Most notably, deception appears to deliver the greatest value when implemented as a complement to existing security controls rather than as an independent defensive capability. Throughout the survey, respondents consistently positioned deception alongside technologies such as SIEM, EDR, identity security controls, and incident response processes rather than as a replacement for them. This indicates that deception is most effective when integrated into a broader security architecture and operational workflow.

A key practical benefit identified in both the literature and the survey findings is the ability of deception technologies to improve signal quality. Traditional security environments often generate large volumes of telemetry, alerts, and behavioural anomalies that require interpretation before meaningful conclusions can be drawn. Deception technologies differ because they create assets and interactions that should not normally be accessed during legitimate operations. As a result, interactions with deceptive resources may provide stronger indications of potentially malicious activity and reduce the uncertainty associated with many conventional detection methods.

However, the findings also suggest that this value is highly dependent on implementation quality. Simply deploying deceptive assets does not guarantee operational benefit. Effective deception requires realistic design, clear separation from production systems, ongoing maintenance, and integration with monitoring and response processes. If deceptive assets are unrealistic, poorly maintained, or disconnected from operational workflows, their value may be significantly reduced. In some cases, ineffective implementations may even contribute additional operational complexity without providing corresponding security benefits.

The findings further indicate that organizations should evaluate deception technologies within the context of their overall security maturity. Respondents repeatedly identified integration, governance, operational complexity, and measurement challenges as important factors influencing adoption. Consequently, organizations with established security operations, mature monitoring capabilities, and well-defined incident response procedures may be better positioned to realize the benefits of deception than organizations that are still developing foundational security capabilities.

Another practical implication concerns how organizations assess the value of security technologies. Traditional evaluation methods often focus on metrics such as detection rates, blocked attacks, response times, or alert volumes. While these measures remain important, the survey responses indicate that they may not fully capture the value provided by deception technologies. Some benefits are indirect and may include increased visibility into attacker behaviour, improved confidence during investigations, earlier identification of compromise, or greater understanding of attacker objectives. These outcomes may be operationally meaningful even when they are difficult to quantify using conventional performance measures.

The findings also have implications for incident response and security governance. Deception technologies introduce situations in which organizations may need to decide whether to prioritize immediate containment or controlled observation of attacker activity. This decision extends beyond technical implementation and requires consideration of organizational risk tolerance, regulatory obligations, asset criticality, and intelligence objectives. As a result, successful use of deception may require closer coordination between security operations, security leadership, governance functions, and incident response teams than is typically required for many traditional security controls.

Finally, the findings suggest that organizations should approach deception as part of a long-term capability development strategy rather than as a standalone technology acquisition. The effectiveness of deception appears to depend less on the individual product selected and more on the organization's ability to integrate, operationalize, govern, and act upon the information generated by deceptive assets. Consequently, the practical success of deception technologies is likely to be determined as much by organizational readiness and operational maturity as by the technology itself.

6.4 Limitations of the Study

The findings presented in this thesis should be interpreted in light of several limitations. While these limitations do not invalidate the results, they influence the extent to which the findings can be generalized beyond the contexts represented within the study.

The most significant limitation relates to the size and composition of the respondent group. The empirical study was based on nine completed responses obtained from a purposively selected group of experienced cybersecurity practitioners. Although the respondents possessed substantial professional experience, the sample remains relatively small and cannot be considered statistically representative of the broader cybersecurity community. Organizations operating in different sectors, regulatory environments, or maturity levels may therefore evaluate deception technologies differently.

A second limitation concerns the nature of the evidence collected. The study focuses primarily on practitioner perceptions, experiences, and professional judgments rather than direct measurement of technical effectiveness. This distinction is important because perceived value and measurable performance are not necessarily identical. A technology may be viewed positively despite limited operational impact, or conversely, a technically effective capability may remain underutilized if organizations struggle to integrate it into existing processes. The findings therefore provide insight into how deception technologies are currently understood and positioned within practice rather than definitive evidence of their technical effectiveness.

The scope of the study also influences the interpretation of the results. The research was intentionally focused on relevance, adoption, and operational positioning rather than on detailed comparisons between specific vendors, products, or deployment models. Consequently, the findings provide a broad view of how deception technologies are perceived within cybersecurity environments, but they do not identify which implementations, architectures, or deployment strategies are most effective under particular conditions.

An additional consideration is that deception technologies themselves continue to evolve. The field has developed considerably over the past decade, particularly through increased integration with cloud environments, security orchestration platforms, threat intelligence capabilities, and adaptive approaches. As a result, practitioner perceptions captured during this study represent a snapshot of a technology area that remains under active development. Future advances may influence both adoption patterns and perceptions of value.

These limitations should be considered when interpreting the findings presented in Chapters 5 and 6. Nevertheless, the study provides useful insight into how deception technologies are currently understood by experienced cybersecurity practitioners and how those perceptions compare with the broader academic literature. The findings therefore contribute to

understanding not only the technical relevance of deception technologies, but also the practical and organizational factors that influence their adoption in real-world cybersecurity environments.

Future research could expand upon these findings through larger-scale surveys, sector-specific studies, longitudinal case studies, or controlled evaluations of deception deployments. Such approaches would help clarify how practitioner perceptions relate to measurable operational outcomes and how deception technologies evolve as organizations continue to mature their defensive capabilities.

6.5 Directions for Future Research

The findings of this study identify several areas where additional research would be valuable. Although deception technologies have received increasing attention in both academic and professional literature, a number of questions remain regarding their practical implementation, effectiveness, and long-term role within cybersecurity architectures.

One area that would benefit from further investigation is the integration of deception technologies into existing security operations. The survey findings consistently highlighted integration as one of the most important factors influencing adoption and operational value. Future research could therefore examine how deception technologies can be incorporated more effectively into SIEM, EDR, SOAR, identity security, and threat intelligence platforms. Developing more standardized integration models may help reduce operational complexity and improve the consistency of deception deployments across organizations.

A second area concerns empirical measurement of effectiveness. While the theoretical benefits of deception technologies are well documented within the literature, comparatively little publicly available research examines measurable outcomes in operational environments. Future studies could investigate metrics such as detection speed, attacker dwell time, incident response efficiency, analyst workload, and the quality of intelligence generated through deception-based approaches. Such research would provide a stronger empirical foundation for evaluating the practical value of deception technologies.

Longitudinal case studies would also represent a valuable contribution. Many existing studies focus on technical implementation or conceptual frameworks rather than sustained operational

use. Examining deception deployments over extended periods would provide insight into issues such as maintenance requirements, organizational acceptance, operational effectiveness, and the evolution of attacker behaviour in response to deception-based controls. Such studies may also help clarify whether the benefits described in theoretical research can be consistently realized in practice.

The findings further suggest that organizational context plays a substantial role in shaping how deception technologies are deployed and perceived. Future research could therefore examine sector-specific applications in greater detail. Industries such as healthcare, finance, critical infrastructure, manufacturing, and public administration operate under different regulatory requirements, risk profiles, and operational constraints. Understanding how these factors influence deception adoption could provide more targeted guidance for practitioners and decision-makers.

Another promising area concerns adaptive and AI-driven deception. Although the survey findings suggest that these capabilities have not yet achieved widespread operational adoption, they align closely with several challenges identified by respondents, including ambiguity, attacker adaptation, and the need for higher-quality detection signals. Future research could explore how adaptive deception systems perform in operational environments, how such systems can be governed safely, and whether automation can improve the effectiveness and scalability of deception deployments without introducing unacceptable levels of risk.

Finally, future studies could examine the relationship between organizational maturity and deception adoption. The findings of this thesis suggest that deception technologies may provide different forms of value depending on the maturity of the security program in which they are deployed. Further research could investigate whether certain maturity levels, governance models, or operational capabilities are associated with more successful deception implementations.

Taken together, these research directions suggest that the next stage of deception research should move beyond questions of technical feasibility and focus more heavily on operational effectiveness, organizational adoption, governance, and measurable outcomes. Such work would help bridge the remaining gap between theoretical potential and practical application identified throughout this thesis.

7 Conclusion

This thesis examined the relevance of deception technologies in modern cybersecurity environments, with particular attention to how they are perceived and positioned in operational practice. The objective was not only to explore the technical foundations of deception technologies, but also to understand how cybersecurity practitioners evaluate their role within contemporary defensive architectures. To achieve this, the study combined a review of academic and technical literature with an empirical survey of experienced cybersecurity practitioners.

The findings indicate that deception technologies have progressed beyond the status of purely experimental or theoretical concepts. Both the literature and the survey responses suggest that deception can provide meaningful value within modern cybersecurity environments, particularly in areas such as early compromise detection, credential misuse identification, lateral movement visibility, and the reduction of uncertainty during security investigations. These findings support the view that deception technologies have established a legitimate role within the broader cybersecurity landscape.

At the same time, the study suggests that deception technologies are generally positioned as complementary rather than foundational security capabilities. Respondents consistently viewed deception as most effective when integrated with existing controls such as SIEM, EDR, identity security mechanisms, and incident response processes. Rather than replacing traditional security controls, deception appears to strengthen their effectiveness by providing higher-confidence indicators of potentially malicious activity.

A key observation emerging from the research is the difference between the broader strategic potential described in the literature and the way deception technologies are commonly applied in practice. Academic research frequently emphasizes the ability of deception to influence attacker behaviour, create uncertainty, support intelligence collection, and alter the dynamics between attacker and defender. In contrast, respondents primarily associated deception with operational objectives related to detection, visibility, and investigation. This suggests that the practical adoption of deception technologies currently emphasizes immediate defensive value more strongly than broader strategic applications.

The findings also indicate that adoption is influenced less by awareness of deception technologies and more by practical implementation considerations. Integration requirements,

operational complexity, governance concerns, measurement challenges, and organizational maturity all emerged as important factors affecting deployment decisions. As a result, the success of deception technologies appears to depend not only on technical capability, but also on an organization's ability to operationalize, govern, and integrate deception within existing security processes.

Taken together, the results suggest that deception technologies currently occupy a transitional position within cybersecurity practice. They are no longer viewed as niche or experimental concepts, yet they have not achieved the level of standardization associated with more established security controls. Their practical value appears strongest when used to improve detection quality, reduce ambiguity, and support existing defensive operations.

The study further suggests that the future development of deception technologies is likely to be closely linked to improvements in integration, automation, and operational maturity. As security environments become increasingly complex and attackers continue to blend malicious activity with legitimate system behaviour, the ability to generate higher-quality signals and improve defender visibility may become increasingly valuable. In this context, deception technologies appear well positioned to play a growing role within layered defensive architectures.

Ultimately, the findings indicate that the primary question surrounding deception technologies is no longer whether they can provide value. Rather, the more significant challenge lies in how organizations can effectively integrate, manage, and operationalize that value within real-world cybersecurity environments. The answer to that challenge will likely determine the future role of deception technologies within modern cybersecurity practice.

8 References

- Almeshekah, M. H. (2015). *Using deception to enhance security: A taxonomy, model, and novel uses* (Doctoral dissertation). Purdue University, West Lafayette, IN, United States. Retrieved from https://docs.lib.purdue.edu/open_access_dissertations/1334/ (Accessed June 11, 2026).
- Almeshekah, M. H., & Spafford, E. H. (2014). Planning and integrating deception into computer security defenses. In *Proceedings of the 2014 New Security Paradigms Workshop (NSPW '14)* (pp. 127–138). Victoria, British Columbia, Canada. New York, NY, United States: Association for Computing Machinery (ACM). <https://doi.org/10.1145/2683467.2683482> (Accessed June 1, 2026).
- Bowen, B. M., Hershkop, S., Keromytis, A. D., & Stolfo, S. J. (2009). Baiting inside attackers using decoy documents. In S. Jajodia & J. Lopez (Eds.), *Security and Privacy in Communication Networks: 5th International Conference, SecureComm 2009, Athens, Greece, September 14–18, 2009, Revised Selected Papers* (pp. 51–70). Berlin, Germany: Springer. https://doi.org/10.1007/978-3-642-05284-2_4 (Accessed June 1, 2026).
- Fraunholz, D., Duque Anton, S., Lipps, C., Reti, D., Krohmer, D., Pohl, F., Tammen, M., & Schotten, H. D. (2018). Demystifying deception technology: A survey. *arXiv preprint arXiv:1804.06196*. Ithaca, NY, United States: Cornell University. Retrieved from <https://arxiv.org/abs/1804.06196> (Accessed June 11, 2026).
- Harmj0y. (2017, January 17). *Roasting AS-REPs*. Harmj0y. Retrieved from <https://blog.harmj0y.net/activedirectory/roasting-as-reps/> (Accessed June 11, 2026).
- Harmj0y. (2019, February 20). *Kerberoasting revisited*. SpecterOps. Alexandria, VA, United States. Retrieved from <https://posts.specterops.io/kerberoasting-revisited-d434351bd4d1> (Accessed June 11, 2026).
- Hutchins, E. M., Cloppert, M. J., & Amin, R. M. (2011). Intelligence-driven computer network defense informed by analysis of adversary campaigns and intrusion kill chains. In *Proceedings of the 6th International Conference on Information Warfare and Security (ICIW 2011)* (pp. 113–125). Washington, DC, United States. Reading, United Kingdom: Academic Conferences International Limited.
- Jajodia, S., Subrahmanian, V. S., Swarup, V., Wang, C., & Wang, X. S. (Eds.). (2016). *Cyber deception: Building the scientific foundation*. Cham, Switzerland: Springer

- International Publishing. <https://doi.org/10.1007/978-3-319-32699-3> (Accessed June 11, 2026).
- Mandiant. (2023). *M-Trends 2023 Special Report*. Reston, VA, United States: Mandiant, Inc. Retrieved from <https://www.mandiant.com/resources/reports/m-trends-2023-special-report> (Accessed June 11, 2026).
- Metcalf, S. (2015). *Red vs. blue: Modern Active Directory attacks, detection, and protection*. Black Hat USA 2015. Las Vegas, NV, United States: Black Hat Conferences. Retrieved from <https://blackhat.com/docs/us-15/materials/us-15-Metcalf-Red-Vs-Blue-Modern-Active-Directory-Attacks-Detection-And-Protection.pdf> (Accessed June 11, 2026).
- Metcalf, S. (2015). *Red vs. blue: Modern Active Directory attacks, detection, and protection* [White paper]. Black Hat USA 2015. Las Vegas, NV, United States: Black Hat Conferences. Retrieved from <https://blackhat.com/docs/us-15/materials/us-15-Metcalf-Red-Vs-Blue-Modern-Active-Directory-Attacks-Detection-And-Protection-wp.pdf> (Accessed June 11, 2026).
- Microsoft Corporation. (2025). *Kerberos authentication overview in Windows Server*. Redmond, WA, United States: Microsoft Corporation. Retrieved from <https://learn.microsoft.com/en-us/windows-server/security/kerberos/kerberos-authentication-overview> (Accessed June 11, 2026).
- MITRE Corporation. (2024). *MITRE ATT&CK framework*. Bedford, MA, United States: MITRE Corporation. Retrieved from <https://attack.mitre.org/> (Accessed June 11, 2026).
- National Institute of Standards and Technology. (2024). *The NIST Cybersecurity Framework (CSF) 2.0* (NIST Cybersecurity White Paper 29). Gaithersburg, MD, United States: U.S. Department of Commerce. <https://doi.org/10.6028/NIST.CSWP.29> (Accessed June 1, 2026).
- Pawlick, J., Colbert, E., & Zhu, Q. (2019). A game-theoretic taxonomy and survey of defensive deception for cybersecurity and privacy. *ACM Computing Surveys*, 52(4), Article 82, 1–28. New York, NY, United States: Association for Computing Machinery (ACM). <https://doi.org/10.1145/3337772> (Accessed June 1, 2026).
- Proofpoint. (n.d.). *What is alert fatigue in cybersecurity?* Sunnyvale, CA, United States: Proofpoint, Inc. Retrieved from <https://www.proofpoint.com/us/threat-reference/alert-fatigue> (Accessed June 11, 2026).

- Rowe, N. C., & Rrushi, J. (2016). *Introduction to cyberdeception*. Cham, Switzerland: Springer International Publishing. <https://doi.org/10.1007/978-3-319-41187-3> (Accessed June 11, 2026).
- Strom, B. E., Applebaum, A., Miller, D. P., Nickels, K. C., Pennington, A. G., & Thomas, C. B. (2018). *MITRE ATT&CK: Design and philosophy* (Technical Report MP180360). Bedford, MA, United States: MITRE Corporation. Retrieved from <https://www.mitre.org/news-insights/publication/mitre-attck-design-and-philosophy> (Accessed June 11, 2026).
- Verizon. (2024). *2024 Data Breach Investigations Report*. New York, NY, United States: Verizon Business. Retrieved from <https://www.verizon.com/business/resources/reports/2024-dbir-data-breach-investigations-report.pdf> (Accessed June 11, 2026).
- Walter, E., Ferguson-Walter, K., & Ridley, A. (2021). Incorporating deception into CyberBattleSim for autonomous defense. *arXiv preprint arXiv:2108.13980*. Ithaca, NY, United States: Cornell University. Retrieved from <https://arxiv.org/abs/2108.13980> (Accessed June 11, 2026).

9 Appendix A

PRACTITIONER SURVEY QUESTIONNAIRE

University of Turku – Expert Perspectives Survey

This questionnaire was conducted as part of a Master's thesis examining the strategic, operational, and future relevance of deception technologies in modern cybersecurity environments.

Deception technologies are cybersecurity tools that place fake assets (such as decoy servers, credentials, or data) within a network to detect attackers. Because legitimate users should never interact with these assets, any interaction may indicate malicious activity and provide early detection opportunities with low false-positive rates.

SECTION 1: BACKGROUND INFORMATION (OPTIONAL)

Q1. Name (Optional)

Q2. Current Role / Job Title (Optional)

Q3. Country of Work (Optional)

Q3a. Did you leave any of the above fields unanswered?

- ☐ Yes
- ☐ No

Q4. Primary reason for not providing background details

- ☐ Personal privacy considerations
- ☐ Organizational policy restrictions
- ☐ Legal or regulatory concerns
- ☐ Preference for anonymity in academic research
- ☐ Not relevant to my responses
- ☐ Other

Q4a. Optional clarification

SECTION 2: PROFESSIONAL CONTEXT

Q5. Years of professional experience in cybersecurity

- Less than 3 years
- 3–7 years
- 8–12 years
- 13–20 years
- More than 20 years

Q6. Primary operational environment

- Commercial / Enterprise
- Government / Public Sector
- Military / Defense
- Critical Infrastructure
- Consulting / Advisory
- Academic / Research
- Mixed / Multiple

Q7. Primary security focus areas (Select all that apply)

- Security Operations (SOC / Blue Team)
- Incident Response / DFIR
- Threat Intelligence / CTI
- Security Architecture / Engineering
- Risk Management / GRC
- Red Team / Adversary Simulation

- Security Leadership / Strategy

SECTION 3: FAMILIARITY AND MATURITY

Q8. Your level of familiarity with deception technologies

Scale:

1 = Conceptual awareness only

5 = Deep practical and architectural expertise

Q9. Organizational maturity in using deception technologies

- Deception is operationalized and integrated into SOC processes
- Deception is deployed primarily as a detection or tripwire capability
- Deception has been piloted or tested
- Deception has been evaluated but not adopted
- Deception is not currently considered

Q9a. Optional context or clarification

SECTION 4: STRATEGIC RELEVANCE

Q10. Overall relevance of deception technologies in modern cybersecurity

Scale:

1 = Not relevant

7 = Strategically critical

Q11. Deception technologies primarily provide value at which level?

- Tactical (alerts, tripwires, detection)
- Operational (incident response, threat hunting)
- Strategic (adversary understanding, long-term defense posture)
- Equally across all levels

- Strongly context-dependent

SECTION 5: ARCHITECTURAL POSITIONING OF DECEPTION TECHNOLOGIES

Q12. Indicate your level of agreement with the following statements.

Scale:

Strongly Disagree – Disagree – Neutral – Agree – Strongly Agree

- a) Deception technologies meaningfully reduce attacker dwell time.
- b) Deception enables insights not achievable through traditional controls.
- c) Deception shifts defenders from a reactive to a proactive posture.
- d) Deception effectiveness depends heavily on attacker sophistication.
- e) Deception value is underestimated by organizational leadership.

Q13. How do you primarily conceptualize the role of deception technologies within a cybersecurity architecture?

- Deception technologies have no meaningful relevance in modern cybersecurity environments
- Deception technologies are best delivered as a feature or capability within broader security platforms (e.g., XDR, SIEM, EDR)
- Deception technologies should exist as a stand-alone specialized solution
- Deception technologies provide the most value as an integrated but distinct security control interoperating with other systems
- The relevance and form of deception technologies are highly context-dependent

Q13a. Optional explanation or qualification

SECTION 6: USE CASES AND STRATEGIC INTENT

Q14. Which objectives best justify the use of deception technologies? (Select up to three)

- Early compromise detection

- Lateral movement containment
- Adversary behaviour analysis
- Long-term threat intelligence collection
- Attribution and investigation support
- Defensive misdirection / attacker manipulation
- Validation of detection and response capabilities

Q15. Which defensive philosophy best aligns with your view on deception technologies?

- Immediate detection and removal of intruders
- Controlled engagement to understand attacker intent
- Selective engagement based on risk and context
- Deception should avoid prolonged attacker presence

Q16. During which attack phases does deception provide disproportionate value? (Select all that apply)

- Reconnaissance
- Initial access
- Credential access
- Lateral movement
- Command and control
- Post-compromise persistence

SECTION 7: RISK, ETHICS AND GOVERNANCE

Q17. How do you assess the overall risk-benefit trade-off of deception technologies?

- Benefits significantly outweigh risks
- Benefits outweigh risks in controlled environments

- Benefits and risks are highly context-dependent
- Risks often outweigh benefits
- Risks are not well understood

Q18. Primary concerns related to deception adoption (Select all that apply)

- Legal or regulatory ambiguity
- Ethical considerations of intentional deception
- Risk of escalation or misattribution
- Operational complexity
- Exposure of defensive strategy
- Organizational resistance or lack of trust

Q19. To what extent should deception activities be explicitly governed?

- Formal policies and legal review required
- Internal guidelines sufficient
- Case-by-case judgment by security leadership
- Minimal governance required
- Unsure

SECTION 8: FUTURE OUTLOOK

Q20. Expected evolution of deception technologies over the next 3–5 years

- Decline due to alternative detection approaches
- Stable niche capability
- Increasing adoption as a complementary control
- Becoming a standard SOC component
- Becoming strategically essential for advanced defense

Q21. Capabilities most critical for next-generation deception technologies (Select all that apply)

- Adaptive, AI-driven deception behaviour
- Deep integration with XDR / SOC platforms
- Cloud-native and hybrid infrastructure coverage
- Automated response and orchestration
- Stronger analytical and intelligence outputs
- Regulatory and compliance alignment

SECTION 9: SYNTHESIS AND REFLECTION

Q22. How would you ultimately position deception technologies in a mature security program?

- Core security capability
- Strategic enhancer of detection and response
- Specialized capability for mature organizations
- Situational or experimental control
- Not a meaningful contributor

Q23. Optional final reflection

What do you consider the most misunderstood or overlooked aspect of deception technologies?