jamk

Phishing Simulations in Preventing Business Email Compromise Attacks

Makkonen Joni

Karjalainen Pauliina

Bachelor’s thesis

May 2025

Bachelor's Degree programme in Information and Communications Technology

jamk | Jyväskylän ammattikorkeakoulu
University of Applied Sciences

jamk

Description

Makkonen, Joni & Karjalainen, Pauliina

Phishing Simulations in Preventing Business Email Compromise Attacks

Jyväskylä: University of Applied Sciences, May 2025, 47 pages

Bachelor's Degree programme in Information and Communications Technology. Bachelor's thesis.

Permission for open access publication: Yes

Language of publication: English

Abstract

Phishing emails pose an increasingly serious threat to the cybersecurity of organizations, as more and more Business Email Compromise incidents are triggered by successful phishing attacks. The aim of this thesis was to investigate how phishing simulations implemented by organizations affect employees' ability to recognize and block genuine phishing emails. In addition, the aim was to determine whether there are any noticeable changes in the simulation results over a 3-year time frame. The aim was to compare the results between different organizations, understand the core factors in the success of phishing emails, and seek recommendations to improve the effectiveness of the simulations.

The thesis was based on equally anonymized data collected from phishing simulations of three different organizations. Each organization regularly implemented several campaigns per month for approximately three years. The simulation messages used varying themes, and the resulting failure rate of users' simulations was calculated as the ratio of clicks on the simulation link to the number of people who saw the simulation message.

The results obtained showed that in each organization the failure rate of simulations was high in the initial phase, but after it decreased significantly during the first year, the development levelled off after 6-12 months. A common phenomenon was also observed in the increase in the failure rate during the summer and Christmas holidays. The differences in development, either for better or worse, between organizations were minor, which suggests that the phenomena underlying the phishing are universal. These phenomena include, for example, seasonal variation and people's sensitivity to error under rush and pressure.

The conclusion of the thesis is that regular phishing simulations reduce human error and help people recognize the correct phishing messages, especially during the first year. However, after this initial phase, training employees based solely on phishing simulations does not maintain learning and reduce failure rates in the longer term. To ensure the best benefit of simulations, organizations should complement the learning experience with, for example, regular information about current threats and interactive exercises. In addition, organizations can combine various micro-trainings, for example, in situations where an employee has clicked on a link in a simulation message.

Keywords/tags (subjects)

Phishing simulations, Phishing, Business email compromise (BEC), Phishing awareness training, Cybersecurity awareness training, Email Security

Miscellaneous (Confidential information)

-

jamk

Kuvailulehti

Makkonen, Joni & Karjalainen, Pauliina

Kalastelusimulaatiot Business Email Compromise-hyökkäysten torjumisessa ja ehkäisemisessä

Jyväskylä: Jyväskylän ammattikorkeakoulu. Toukokuu 2025, 47 sivua.

Tieto- ja viestintätekniikan tutkinto-ohjelma. Opinnäytetyö AMK.

Julkaisun kieli: englanti

Julkailusulupa avoimessa verkossa: kyllä

Tiivistelmä

Kalasteluviestit muodostavat yhä isomman uhan organisaatioiden kyberturvallisuudelle, koska entistä useammat Business Email Compromise – poikkeama käynnistyy onnistuneella kalasteluhyökkäyksellä. Tämän opinnäytetyön tavoitteena oli selvittää, että miten organisaatioiden toteuttamat kalastelusimulaatiot vaikuttavat työntekijöiden kykyyn tunnistaa sekä torjua aidot kalasteluviestit. Tämän lisäksi tavoitteena oli selvittää, että onko noin 3 vuoden aikajanalla havaittavissa muutoksia simulaatioiden tuloksissa. Tavoitteena oli verrata tuloksia eri organisaatioiden välillä, ymmärtää ydintekijät kalasteluiden onnistumisessa ja etsiä suositusehdotuksia parantamaan simulaatioiden tehokkuutta.

Opinnäytetyö perustui kolmen eri organisaation yhdenvertaisesti anonymisoituihin kalastelusimulaatioista kerättyihin aineistoihin. Jokainen organisaatio toteutti säännöllisesti noin kolmen vuoden ajan useita kampanjoita kuukaudessa. Simulaatioviesteissä käytettiin vaihtelevia teemoja ja tuloksena saatu käyttäjien simulaatioiden epäonnistumisprosentti (failure rate) laskettiin simulaatiolinkin klikkausten suhteena simulaatioviestin nähneisiin henkilöihin. Tällä tavoin opinnäytetyö sai vertailukelpoista ja luotettavaa ainestoa, joka analysoitiin vertailevilla menetelmillä ottaen huomioon myös kausivaihtelut, kuten kesä- ja joululomien aikaiset muutokset.

Saadut tulokset osoittivat, että jokaisessa organisaatiossa simulaatioiden epäonnistumisprosentti oli alkuvaiheessa korkea, mutta ensimmäisen vuoden aikana sen laskiessa merkittävästi, kehitys tasaantui 6–12 kuukauden jälkeen. Yhteisenä ilmiönä havaittiin myös kesä – ja joululomien aikaan sijoittuvat nousut epäonnistumisprosentissa. Organisaatioiden välillä erot kehityksessä joko parempaan tai huonompaan olivat vähäisiä, joka viittaa siihen, että kalasteluiden taustalla vaikuttavat ilmiöt ovat universaaleja. Näitä ilmiöitä ovat esimerkiksi kausivaihtelu sekä ihmisten virheherkkyys kiireen ja paineen alla.

Opinnäytetyön johtopäätöksenä säännölliset kalastelusimulaatiot vähentävät ihmisvirheitä sekä auttavat ihmisiä tunnistamaan oikeat kalasteluviestit varsinkin ensimmäisen vuoden aikana. Kuitenkin tämän alkuvaiheen jälkeen pelkkä kalastelusimulaatioihin pohjautuva työntekijöiden kouluttaminen ei ylläpidä oppimista ja laske epäonnistumisprosenttia pidemmällä aikavälillä. Simulaatioiden parhaan hyödyn varmistamiseksi organisaatioiden tulisi täydentää oppimiskokemusta esimerkiksi säännöllisillä tiedottamisilla ajankohtaisista uhkista ja interaktiivisilla harjoituksilla. Näiden lisäksi organisaatiot voivat yhdistää erilaisia mikrokoulutuksia esimerkiksi tilanteisiin, joissa työntekijä on klikannut simulaatioviestin linkkiä.

Avainsanat (asiasanat)

Kalastelusimulaatiot, Tietojenkalastelu, BEC-hyökkäys, Tietoturvakoulutus, Sähköpostiturvallisuus

Miscellaneous (Confidential information)

-

Contents

Figures

1 Introduction

Business Email Compromise (BEC) refers to cyberattacks in which threat actors impersonate trusted individuals via email to deceive employees and gain unauthorized access to sensitive data, user accounts, or financial assets. These attacks often require minimal technical skill—sometimes just a single click is enough to cause a major breach. According to IBM's Cost of a Data Breach Report 2024, the average cost of a successful BEC incident was $4.88 million USD, highlighting its significant financial impact on organizations (IBM, 2024).

BEC attacks often involve the impersonation of executives, IT personnel, colleagues, or other trusted entities, typically leveraging urgency and authority to pressure victims into acting quickly. This threat has been amplified by the rise of remote work, where employees are more isolated, depend heavily on email communication, and frequently lack immediate ways to verify suspicious requests. The absence of in-person oversight, increased reliance on personal or unsecured devices, and a diluted security culture in distributed settings further contribute to organizational vulnerability. As a result, BEC represents a significant and growing threat to businesses operating in hybrid or fully remote environments.

Cybercriminals have increasingly adopted Adversary in the middle (AiTM) techniques to bypass Multi Factor authentication (MFA), which is now a standard security practice in many organizations (National Cyber Security Center of Finland, 2024). At the same time, the growing use of AI by threat actors has further intensified the threat of BEC, enabling more convincing phishing content, adaptive social engineering strategies, and large-scale automation. These advancements—particularly when combined with tactics like AiTM—make it increasingly difficult for traditional security controls to detect and prevent such attacks.

In response to these evolving threats, many organizations have adopted phishing simulations as a training tool to improve employee awareness and encourage appropriate responses to suspicious emails. For instance, Microsoft provides a phishing simulation platform that delivers realistic scenarios designed to expose users to various phishing tactics (Microsoft, 2024a). These simulations aim to strengthen cybersecurity awareness—a key defence against BEC—by addressing human error, which remains one of the most common vulnerabilities in email-based attacks.

While phishing simulations are widely used to raise employee awareness, their long-term impact on behavior change remains uncertain—and in some cases, may even diminish effectiveness over time (Lain, Kostiainen, & Capkun, 2022). Understanding how such training influences user behavior across extended periods is crucial, particularly as organizations continue to depend on simulations to mitigate human-related security risks. This thesis examines the long-term effects of phishing simulations across three organizations, aiming to identify behavioral patterns and trends in awareness development over a three-year period.

2 Research Ideology

The research approach adopted in this study is grounded in practical relevance and real-world application, with the goal of examining how phishing simulation failure rates evolve over time in different organizational contexts. The methodology combines both quantitative and qualitative elements to allow a deeper understanding of behavioral trends and organizational differences. Emphasis is placed on data-driven analysis and ethical data collection practices, ensuring that the research remains transparent, replicable, and applicable to real-world cybersecurity awareness initiatives.

2.1 Research Objectives and Questions

This research examines how phishing simulations contribute to improving organizational resilience against BEC attacks. As phishing remains one of the most common and effective entry points for threat actors (Verizon, 2024), many organizations have implemented simulation-based training to increase employee awareness and reduce the likelihood of human error. Despite their widespread use, the long-term effects of these simulations on employee behavior—particularly across different organizational environments—remain insufficiently understood.

The study is guided by the following primary research question:

> *How do phishing simulation failure rates evolve over time, and what organizational patterns can be observed?*

To explore this question, the research sets out three main objectives. First, it aims to track trends in phishing simulation failure rates over a three-year period to better understand how employee behavior changes over time. Second, it compares these trends across three organizations of varying sizes and sectors to examine how organizational context may influence outcomes. Third, the study seeks to identify recurring patterns or notable improvements in simulation performance that could inform the development of more effective cybersecurity awareness strategies and long-term training approaches.

2.2 Research Approach and Philosophy

This study adopts a mixed-methods approach, combining quantitative analysis of phishing simulation data with qualitative insights drawn from organizational context. The primary aim is to explore behavioral trends over time, rather than to test a hypothesis or develop predictive models. For this reason, the research design is exploratory and descriptive, focusing on uncovering patterns and differences in how phishing simulation results evolve across different organizational environments.

In line with this, the research follows a pragmatic approach, focusing on real-world relevance and practical results rather than following one strict research method. By using failure rate data from three organizations over a three-year period, the goal is to gain useful insights that are based on actual data and take the organizational context into account.

2.3 Data Collection and Analysis

This study draws on data collected from three organizations, each of which had implemented phishing simulations as part of their broader cybersecurity awareness programs. The dataset covers the first three years of simulation use within each organization. Simulated phishing emails were delivered randomly and frequently, with multiple campaigns taking place each month. The primary metric analyzed is the failure rate, defined as the percentage of employees who clicked on a link contained in a simulated phishing email.

All three organizations utilized the same phishing simulation platform, though the platform's name is withheld to ensure confidentiality. Data was provided in the form of line graphs that illustrated

changes in failure rates over time. These visuals were subsequently redrawn to anonymize the content and remove any identifying company information. For the purposes of analysis, the organizations are referred to as Company A, Company B, and Company C throughout the thesis.

The participating organizations represent different industries and vary in both internal cybersecurity practices and organizational size, ranging from approximately 5,000 to 10,000 employees. The data was collected exclusively for academic purposes, and no personal or sensitive employee information was accessed at any stage. Individual user actions were not analyzed; only aggregated results, as presented in the original visualizations, were used in the study.

The analysis focused on identifying trends in phishing simulation failure rates over time within each organization, as well as comparing these trends across all three companies. The objective was to detect general patterns, fluctuations, and potential improvements in performance, along with any differences that could be linked to organizational context. The data was manually transferred into spreadsheet software (e.g., Excel) and visualized using line graphs to support trend analysis.

2.4 Ethical Considerations and Use of AI in Writing

This research was conducted in full accordance with the ethical guidelines of JAMK University of Applied Sciences (2024), which emphasize research integrity, data protection, proper referencing, original authorship, and appropriate permissions for research involving organizations or individuals. Written permission was obtained from all participating organizations to use their phishing simulation data for academic purposes, including general references to company size. At the request of the organizations, their fields of operation are not disclosed.

All data used in the study was anonymized—either provided in anonymized form or anonymized by the researcher before analysis. Visual materials were redrawn to remove any company-specific details such as logos, internal terminology, or branding. Organizational identifiers introduced earlier in the thesis are used consistently to preserve confidentiality. No identifiable or sensitive personal data was accessed, and individual-level user actions were not analysed.

Throughout the research process, honesty, transparency, and academic integrity were maintained. The research design, data analysis, visualizations, and conclusions are entirely the work of the author. Where external tools were used, such as for language refinement or formatting, all output was critically reviewed and edited by the author to ensure originality.

AI-assisted tools were used during the writing process of this thesis in a limited and non-substantive capacity. Their use was restricted to tasks such as language refinement, structural clarity, and formatting consistency. These tools were not used to generate content, conduct analysis, or interpret results. All theoretical framework development, data analysis, visualizations, and conclusions were independently produced by the author. Where AI-generated suggestions were applied to improve clarity or readability, the final wording was reviewed and approved by the author.

3 Theoretical Framework

The theoretical framework of this thesis draws from current academic research on phishing simulations, cybersecurity awareness training, and employee behavior in the context of BEC. The literature review includes both supportive and critical perspectives, highlighting not only the effectiveness of simulation-based training but also its limitations, psychological impacts, and emerging trends such as AI-driven phishing strategies.

Alongside the literature, key concepts and terminology relevant to phishing attacks, user awareness, and organizational vulnerability are introduced to provide a conceptual foundation for the analysis. This combination of theoretical insights and clearly defined concepts supports the research focus on how phishing simulation failure rates evolve over time and vary across organizational environments.

3.1 Literature Review

The literature review explores how phishing simulations are applied within organizations to enhance cybersecurity awareness and reduce employee susceptibility to social engineering attacks. Eight peer-reviewed studies were selected for their relevance, methodological rigor, and the diversity of perspectives they offer on the long-term effectiveness of simulation-based training programs.

A targeted review approach was adopted rather than a systematic literature review, focusing on identifying high-relevance academic sources. Studies were primarily located through Google Scholar and IEEE Xplore using search terms such as “phishing simulations,” “awareness training,” “user behaviour,” and “business email compromise.” To ensure the material reflects the current state of the field, only research published within the past three years was considered. Preference was given to large-scale studies conducted in real-world organizational settings.

The selected literature provides both quantitative and qualitative insights into employee behaviour, training strategies, psychological factors, and evolving phishing techniques. To support thematic analysis, the findings are categorized into three groups: studies demonstrating the effectiveness of phishing simulations, studies presenting mixed or critical results, and emerging research focused on psychological dynamics and technological developments. A summary at the end of the section highlights key findings and clarifies the research gap this thesis aims to address.

3.1.1 Studies Demonstrating Simulation Effectiveness

Several studies provide strong evidence that phishing simulations, when implemented consistently and thoughtfully, can lead to measurable improvements in employee awareness and reductions in failure rates. These findings often come from large-scale organizational environments and offer empirical insight into how employees respond to different simulation formats, threat types, and training interventions.

In a higher education context, Ciupe and Orza (2023) conducted a large-scale phishing simulation campaign at the Technical University of Cluj-Napoca, targeting around 20,000 users, including students, faculty, and administrative staff. Using Microsoft’s attack simulation tools, the study implemented a tiered series of phishing scenarios—ranging from generic beginner-level emails to advanced, role-specific attacks. The results indicated that certain phishing methods, such as drive-by URLs and OAuth consent phishing, were particularly effective, while user reporting behavior remained relatively low. Despite these challenges, the campaign contributed to increased awareness and led to institutional recommendations aimed at strengthening secure communication practices. The study emphasized the value of realistic, context-aware simulations in reinforcing cybersecurity culture across diverse user groups.

A separate study by Sutter et al. (2022) approximately 31,000 participants and 140 phishing simulations conducted at Swiss Federal Institute of Technology in Zurich. The research compared multiple training formats, including embedded training, feedback rubrics, and a control group with no training. Results showed clear behavioral improvements over time, with even minimal training yielding better outcomes than no training. Only 0.64% of users were identified as repeat offenders, while mobile device usage emerged as a significant risk factor, contributing to over one-third of credential submissions. Notably, the study introduced a machine learning model capable of predicting email difficulty using semantic and emotional characteristics. These findings support the use of data-driven, role-specific training programs that enhance phishing resilience while respecting ethical boundaries.

Together, these studies highlight the importance of repetition, adaptive content design, and organizational context in effective phishing awareness training. They suggest that simulations can serve not only as a tool for vulnerability detection but also as a catalyst for behavioral change when supported by well-designed interventions and institutional follow-through.

3.1.2 Mixed or Cautionary Findings

Although phishing simulations are widely implemented in organizational settings, several studies have questioned their long-term effectiveness and ability to maintain user engagement. Rather than dismissing simulation-based training, these studies offer a more nuanced perspective—highlighting limitations, inconsistent outcomes, and the critical role of training design, content delivery, and user experience.

In a 15-month field study involving over 14,000 employees, Lain et al. (2022) examined the impact of phishing training, warning messages, and reporting systems within a large public sector organization. The results showed that 32% of participants clicked on at least one phishing email, and 25% took high-risk actions such as entering credentials. Notably, training interventions failed to reduce these numbers over time, and in some cases, susceptibility increased among those who received voluntary embedded training—suggesting a potential desensitization effect. Conversely, simple email warnings proved consistently effective, regardless of the amount of detail provided. The study also found that employee-based reporting, supported by a “report” button and feedback loop, enabled rapid and accurate threat identification with minimal operational burden.

Ho et al. (2025) expanded this critique through a randomized controlled trial at UC San Diego Health, involving more than 19,000 employees across ten simulation campaigns. Their study evaluated five training models, including annual cybersecurity training and various forms of embedded training. Annual training showed no measurable benefit, and embedded training led to only a modest 1.7% reduction in failure rates. Engagement levels were low: over half of users exited the training immediately, and fewer than 24% completed it. Alarmingly, users who underwent multiple static training sessions were more likely to fail future simulations, possibly due to training fatigue. Only interactive training formats, and only when fully completed, showed small but statistically significant improvements—underscoring the importance of tailoring both content and delivery.

Taking a comparative approach, Davis and Grant (2022) reviewed ten studies contrasting traditional phishing simulations with gamified phishing education tools. While simulations yielded mixed results, gamified methods—such as interactive quizzes and scenario-based games—consistently outperformed them in terms of engagement, knowledge retention, and behavior change. One program, What.Hack, improved post-test phishing detection by 36.7%. Gamified training was also found to reduce anxiety and build user confidence. The review concluded that although simulations are scalable and easy to deploy, gamified learning may offer more sustainable impact, particularly for cultivating long-term behavioral change.

Taken together, these findings suggest that phishing simulations, while valuable for assessing and mapping user behavior, may not be sufficient on their own to produce lasting change. The research emphasizes the need for well-designed, interactive, and role-sensitive training interventions—ideally supported by complementary strategies such as simplified warnings and crowdsourced detection mechanisms.

3.1.3 Psychological and Technological Perspectives

Beyond assessing the effectiveness of phishing simulations, recent research has begun to explore deeper dimensions of phishing susceptibility—shifting focus from surface-level behavioral metrics (such as click rates) to psychological responses, individual traits, and the influence of emerging

technologies like AI-generated phishing content. This growing body of work reflects a broader effort to understand not only what users do, but why they behave in certain ways, and how phishing threats may evolve in response to technological and organizational change.

Schöps et al. (2024) examined the emotional and perceptual impact of simulated phishing campaigns in a large manufacturing organization. Drawing on both quantitative surveys and qualitative interviews, the study found that employees who clicked on phishing emails experienced significantly higher stress levels and lower self-efficacy compared to those who reported the emails. Common emotional reactions included shame, guilt, and frustration—particularly among employees in sensitive roles. While most participants acknowledged simulations as effective awareness tools, those who failed were more likely to perceive the experience as punitive rather than educational. The study emphasized the importance of psychological safety, constructive feedback, and open communication in the design of training programs, arguing that behavioral outcomes alone fail to capture the full effect of awareness efforts.

Taking an individual-level perspective, Beu et al. (2023) investigated how factors such as job satisfaction, tenure, and cognitive bias influence susceptibility to phishing. In a corporate setting, the study found that employees with low satisfaction and loyalty scores were significantly more likely to fall for phishing attempts—more so than those with lower detection accuracy. Newer employees were also more vulnerable, likely due to limited familiarity with organizational communication patterns. Demographic factors such as age and gender were not strongly associated with phishing risk. These findings suggest that human resource data may be a valuable input into cybersecurity strategies, and that user vulnerability is shaped as much by organizational context and engagement as by technical knowledge or training exposure.

From a technological standpoint, Brachten (2025) conducted a field study involving phishing emails generated by large language models (LLMs), such as GPT-4o. The results showed that AI-personalized emails—particularly those customized at the department or role level—outperformed generic messages in terms of click-through and conversion rates. Interestingly, highly personalized emails aimed at individuals were less effective, possibly due to increased user suspicion. Even smaller or less advanced LLMs were able to generate persuasive content, lowering the technical barrier for attackers. The study highlights the growing sophistication of AI-assisted phishing,

calls for updated awareness strategies that go beyond visual detection, and raises ethical concerns about the misuse of generative AI in cybersecurity contexts.

Collectively, these studies reflect an important expansion in phishing awareness research—moving from a narrow focus on actions and outcomes to a more holistic understanding of the psychological, organizational, and technological factors that shape user behavior. They underscore the need for adaptive and empathetic training approaches that account for emotional impact, organizational culture, and the growing threat of AI-enhanced social engineering.

3.1.4 Summary and Research Gap

The reviewed literature demonstrates that phishing simulations can play an important role in raising cybersecurity awareness and reducing employee susceptibility to email-based attacks. Studies by Ciupe and Orza (2023) and Sutter et al. (2022) provide compelling evidence for the effectiveness of simulations, particularly when combined with repetition and contextual feedback. However, other research presents a more cautious perspective. Findings from Lain et al. (2022) and Ho et al. (2025) raise concerns about low user engagement, training fatigue, and diminishing returns, while Davis and Grant (2022) suggest that gamified approaches may offer a more sustainable impact in certain organizational contexts.

More recent research has shifted focus toward deeper psychological and technological influences. Studies by Schöps et al. (2024), Beu et al. (2023), and Brachten (2025) highlight the role of employee stress, organizational attitudes, and emerging threats such as AI-personalized phishing in shaping user behavior. These findings suggest that traditional training models may no longer be sufficient on their own and should evolve to address emotional, behavioral, and technological complexities.

Despite the breadth of existing research, few studies have explored how phishing simulation failure rates evolve over extended periods within real-world organizations. Even fewer have compared results across organizations of different sizes and industries. Rather than aiming to propose definitive solutions, this thesis seeks to contribute to the field by analyzing long-term phishing simulation outcomes across three anonymized organizations. The goal is to better understand how

employee awareness develops over time and how organizational context may influence behavioural responses to simulation-based training.

3.2 Business Email Compromise and Phishing Techniques

Phishing attacks play a central role in BEC incidents by targeting human behaviour and exploiting trust-based communication systems. Understanding the techniques used in these attacks, from traditional credential harvesting to advanced methods like AiTM proxies, is essential for developing effective defence strategies. The conceptual foundation of phishing within BEC includes both widely used and increasingly sophisticated tactics, all of which aim to manipulate recipients into disclosing sensitive information or granting unauthorized access. Recognizing these patterns provides important context for why organizations implement phishing simulations as a proactive cybersecurity measure.

3.2.1 Business Email Compromise

BEC is a sophisticated form of cybercrime in which attackers gain unauthorized access to an organization's internal systems by exploiting human behavior, typically through phishing. These campaigns aim to deceive employees into disclosing login credentials, approving unauthorized transactions, or exposing sensitive information, often without triggering traditional technical defenses. Unlike attacks that rely on malware or brute-force techniques, BEC leverages social engineering to exploit trust and workplace routines.

Phishing remains the most common entry point in BEC. Messages are crafted to closely resemble legitimate communications, often impersonating executives, IT staff, or colleagues. These emails frequently rely on psychological triggers such as urgency, authority, or fear of provoking action—such as approving a payment, opening a malicious file, or clicking on a fraudulent link.

In many cases, attackers do not need to break in—they simply log in using stolen or purchased credentials. Once access is gained, they can move laterally across the network, monitor conversations, and time their attacks for maximum effect. This shift in strategy, from system exploitation to behavioral manipulation, significantly increases the stealth and success rate of BEC.

According to the Federal Bureau of Investigation (2023), BEC-related fraud has been reported globally, with losses exceeding $50 billion between 2013 and 2022. Fraudulent transactions are commonly routed through financial institutions in Hong Kong, China, the UK, Mexico, and Singapore.

As organizations recognize that phishing is the primary vector for BEC, attention has increasingly turned toward building user awareness and resistance. Since these attacks often bypass technical controls, the development of behavioral defenses—such as phishing simulations—has become a critical component of modern cybersecurity strategy.

3.2.2 Common Phishing Techniques and Delivery Methods

Phishing attacks often rely on credential harvesting, where users are lured into entering login information on fake websites. These emails typically contain links directing recipients to webpages that mimic legitimate login portals, such as Microsoft 365 or internal systems. The intent is to deceive users into entering usernames and passwords, which attackers then collect for unauthorized access to corporate environments (Ironscales, n.d.). Messages often use urgency—such as claiming an account is expiring or compromised—to compel users to act quickly.

Another commonly used tactic is email spoofing, where the attacker forges the sender's email address to appear as a trusted source, such as a CEO or IT department. This method is particularly effective in BEC because it reduces scepticism and increases the likelihood of users responding or following malicious instructions. As Fortinet (2024) notes, spoofed emails significantly enhance the credibility of phishing messages, raising their success rates.

Malicious attachments represent another frequent phishing vector. Files disguised as invoices, reports, or job applications may contain links or macros that lead to credential harvesting or malware installation. Attackers often label these documents with urgent or routine-sounding filenames to encourage opening—such as "Invoice_Due_Today.pdf." Although email filters often block malware, cleverly disguised attachments can bypass detection and remain highly effective in BEC attacks (Phriendly Phishing, n.d.).

In addition, phishing campaigns often benefit from Open-Source Intelligence (OSINT). By gathering publicly available data—such as job titles, contact details, and organizational charts—from platforms like LinkedIn or company websites, attackers can tailor phishing messages to appear highly specific and credible. This personalization increases the perceived legitimacy of the message, making recipients more likely to engage (SANS Cyber Defense, 2024).

While these techniques form the foundation of most phishing campaigns, attackers have also developed more advanced strategies to bypass even the most vigilant defences and security tools.

3.2.3 Advanced Threat Actor Methods

Modern phishing campaigns increasingly use advanced techniques designed to bypass even robust security defences such as MFA. Among the most dangerous of these is the AiTM attack. In this setup, a proxy server is positioned between the victim and the legitimate login page. When the victim enters their credentials and completes MFA, the proxy intercepts not only the login details but also the session cookies that authenticate the user. These session cookies allow attackers to hijack active sessions and impersonate the user without needing to re-enter credentials or MFA tokens (Weinert, 2024).

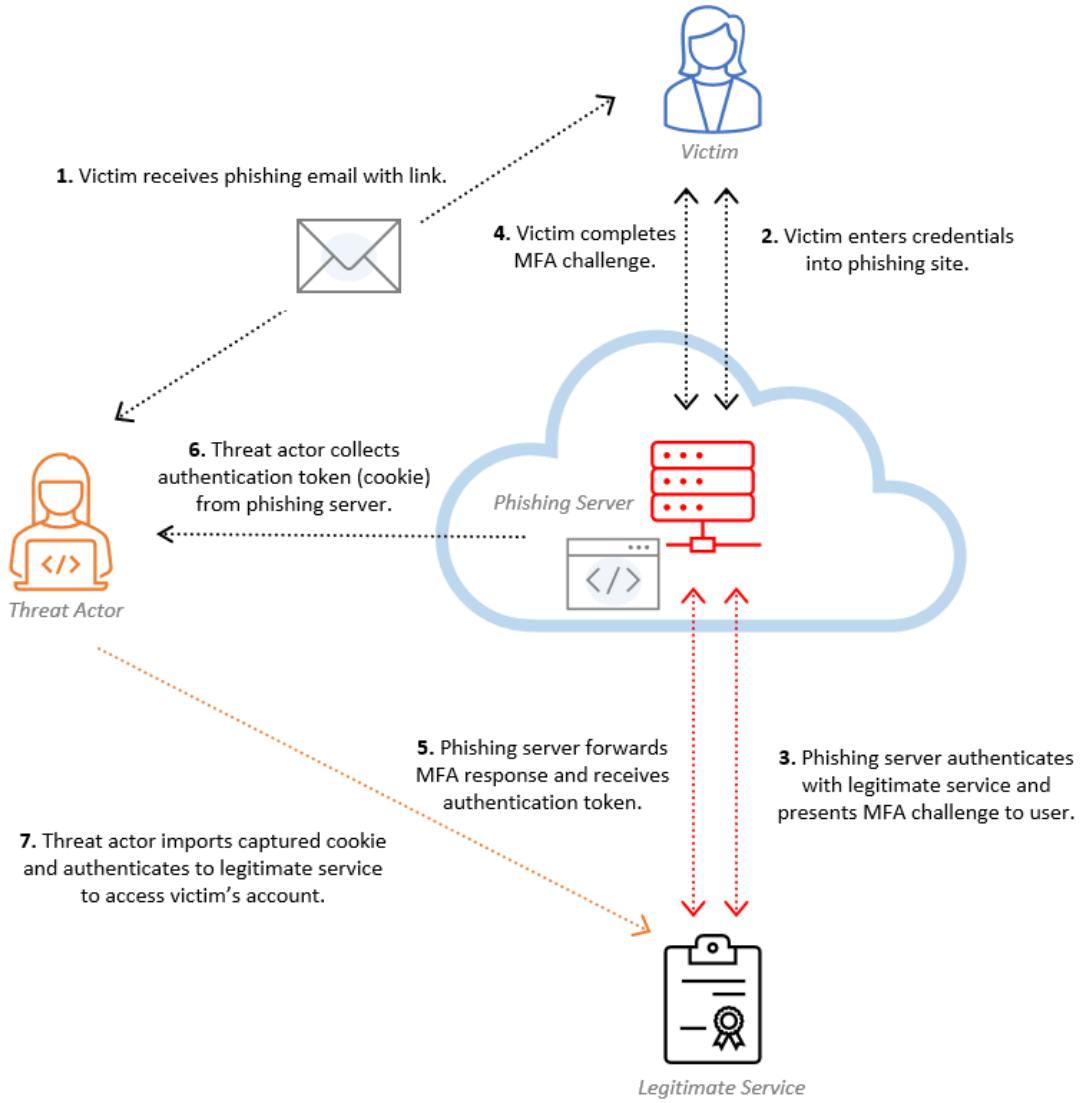

Figure 1. Sample attack diagram showing the MFA bypass process (adapted from AON, n.d).

While MFA is considered a best practice in modern cybersecurity, attackers have devised several ways to circumvent it. For instance, once inside a system, an attacker may register a new MFA device to maintain persistent access, even if the original user changes their password.

Another emerging tactic is known as MFA fatigue or push bombing. In this form of social engineering, the attacker floods the victim's device with repeated MFA push notifications, hoping the target will approve one out of annoyance or confusion. These attacks are especially effective in organizations that use mobile-based MFA apps, where users become desensitized to repeated

prompts. Hoxhunt (2023) reports a growing number of high-profile breaches involving MFA fatigue, highlighting the urgent need for more user-centric education and layered authentication approaches.

3.2.4 AI in Phishing Campaings

The integration of AI into phishing operations has significantly increased the sophistication, scalability, and contextual relevance of phishing campaigns. Modern AI-powered tools can generate messages that mimic corporate language, grammatical patterns, and internal communication styles. These messages often include company-specific terminology and are crafted to closely resemble authentic emails, increasing the likelihood of user interaction (Burgess, 2024).

Beyond content generation, AI is also used to automate reconnaissance. By scanning publicly available information such as LinkedIn profiles, press releases, and company websites, attackers can dynamically tailor phishing emails to specific roles or departments. This context-driven customization enhances believability—especially when the email appears to come from a known executive or internal contact.

The risks of AI-generated phishing are not merely theoretical. In a notable case from Finland, two men orchestrated a sophisticated fraud campaign using generative AI tools. According to Harju (2025) in Finnish newspaper called Helsingin Sanomat, the perpetrators sent thousands of text messages in Finnish, posing as trusted governmental organizations such as “Suomi.fi” or “Omavero.fi.” These messages falsely claimed that the recipients had unpaid bills and directed them to fraudulent websites designed to capture bank credentials.

AI tools such as ChatGPT were reportedly used to craft these messages, enhancing their credibility and linguistic fluency. Victims were further manipulated through follow-up phone calls in which the attackers posed as bank employees, using voice scripts that were later presented as evidence in court. The case underscores the growing threat of AI-assisted phishing and the urgent need for advanced detection and education measures to counter this evolving vector.

3.3 Phishing Simulations in Organizations

Phishing simulations are widely recognized as a practical method for addressing human vulnerabilities in cybersecurity. By replicating real-world social engineering scenarios in a controlled environment, these exercises enable organizations to evaluate employee readiness, strengthen awareness, and promote safer behavior in digital communication. As phishing continues to be a primary vector in BEC attacks, simulations play a central role in improving organizational resilience.

The use of phishing simulations extends beyond training; they also serve as diagnostic tools that help security teams identify risk-prone groups, monitor behavioral trends, and tailor interventions. Their effectiveness depends on the relevance of the scenarios, the quality of feedback provided, and how well they align with everyday communication patterns. While the benefits of simulation-based training are well documented, their implementation is not without challenges. Concerns around user stress, desensitization, and long-term behavioral change have raised important questions about sustainability and design quality.

A balanced examination of their purpose, strategic value, and limitations provides critical insight into how simulations contribute to organizational cybersecurity and what factors influence their success.

3.3.1 Purpose and Function

Phishing simulations serve a dual role within organizational cybersecurity: they educate employees while providing measurable insights into user behavior. By mimicking real-world phishing attacks in a risk-free environment, these exercises help individuals learn to identify malicious content and respond with caution. At the same time, they allow organizations to assess vulnerabilities across teams and departments, identifying patterns and areas that require further attention.

Effective simulations are tailored to reflect specific communication styles and the threat landscape relevant to each role. For example, finance staff may be tested with fake invoice scams, while IT teams might receive simulations involving account reset requests. Many platforms provide instant feedback or brief microlearning sessions following a failed simulation, turning mistakes into teachable moments and reinforcing lessons in context.

The importance of phishing simulations is amplified by the growing sophistication of cyberattacks. As phishing tactics increasingly bypass technical safeguards like spam filters or even multi-factor authentication, human awareness becomes the last line of defense. Regular simulation-based training encourages a culture of vigilance, helping to embed cybersecurity awareness into everyday workflows.

Recent statistics reinforce the scale of the problem. According to Cartier (2025), 73% of reported cyber incidents in 2024 were linked to BEC. These figures underscore the critical role of phishing simulations in mitigating human error and preparing organizations for evolving threats. When simulations are run regularly and thoughtfully, they transform employees from passive targets into active participants in cyber defense.

3.3.2 Benefits and Organizational Value

Phishing simulations offer clear benefits for organizations seeking to enhance cybersecurity resilience, especially against BEC attacks. One of the most recognized advantages is the boost in employee awareness. By exposing users to realistic phishing attempts in a risk-free environment, simulations help employees recognize suspicious content, reducing click rates and promoting safer behavior. Research by Schöps et al. (2024) demonstrates that when thoughtfully executed, simulations can significantly improve users' self-efficacy and threat detection skills.

The long-term effectiveness of simulations lies in their ability to reinforce behavioral change. When combined with consistent repetition and personalized feedback, phishing simulations help normalize cybersecurity mindfulness. This study contributes to this understanding by analyzing long-term trends across multiple organizations to assess how repeated exposure influences employee behavior over time.

Beyond their educational value, simulations play a strategic diagnostic role. Organizations can use failure rate data to identify high-risk roles, departments, or time periods, allowing them to tailor awareness campaigns accordingly. This data-driven approach not only improves internal defenses but also supports compliance with widely recognized frameworks such as ISO/IEC 27001 and the NIST Cybersecurity Framework, which emphasize ongoing security training and risk management.

Many organizations have integrated phishing simulations into the operations of their Security Operation Center (SOC). SOC teams increasingly oversee the design and deployment of custom simulation campaigns tailored to specific departments or threat profiles. This integration enables more sophisticated awareness programs, aligning closely with evolving threat landscapes and organizational priorities. Simulations supported by SOC intelligence allow for precise vulnerability assessments and continuous improvement of overall security posture.

By embedding cybersecurity awareness directly into everyday communication workflows, phishing simulations bridge the gap between policy and practice. They empower employees to act as a line of defense and enable organizations to adapt dynamically to an ever-changing threat environment.

3.3.3 Challenges and Limitations

Despite their growing popularity, phishing simulations face several limitations that can affect their overall effectiveness. One prominent issue is simulation fatigue, which arises when employees become disengaged due to frequent or repetitive scenarios. Without variation, personalization, or strategic timing, simulations may feel routine or frustrating—reducing attentiveness and lowering the potential for meaningful learning.

Another concern is the false sense of security that simulations may create. Employees who repeatedly succeed in simulated scenarios might develop overconfidence in their phishing detection skills. This can be problematic, as real phishing emails often feature more convincing language, personalization, and psychological manipulation than many simulation platforms can currently replicate.

Emotional and psychological responses also play a role in how users engage with simulations. As shown by Schöps et al. (2024), employees who fall for simulated phishing emails may experience stress, embarrassment, or frustration—especially if they believe their performance is being monitored or judged. While this heightened emotional state can promote vigilance, it may also reduce the long-term effectiveness of training if feedback is not delivered in a constructive and supportive manner.

The effectiveness of simulation-based training programs has also been questioned in recent studies. Ho et al. (2025), in a large-scale analysis of a U.S. healthcare organization, found that traditional formats—such as static training modules or annual awareness sessions—did not result in significant long-term improvements in phishing resistance. Although slight reductions in failure rates were observed, the overall protective value was limited, particularly when employee engagement with the content was low.

Moreover, many simulation platforms offer limited feedback and analytics, which restricts organizations from gaining deep insights into behavioral patterns or progress over time. Without detailed and actionable data, it’s difficult for organizations to adjust their awareness strategies or assess the true impact of their training efforts.

Taken together, these challenges highlight the importance of designing phishing simulations that are engaging, context-aware, and supported by thoughtful feedback. Without such considerations, even well-intentioned simulation programs risk becoming ineffective or counterproductive.

3.3.4 Commercial phishing simulation platforms

There are multiple commercial phishing simulation platforms used by organizations, here is a few examples. With Microsoft Defender for Office 365's attack simulation training, organizations can design and execute realistic phishing and credential harvesting campaigns entirely in the cloud. Administrators choose from built-in templates (credential harvest, malware attachment, drive-by URL, OAuth consent, QR code, etc.) or import custom payloads, then target specific users or groups using Azure AD-driven distribution and geo-aware delivery windows. Campaigns automatically collect real-time metrics (click-through rates, "Report Phishing" submissions in Outlook, compromised sessions, training completions) and display them alongside other alerts in the Microsoft 365 Security Center. Because it uses existing Exchange Online mailboxes and Defender telemetry, there's no need for on-premises servers or additional agents, and results are unified with your broader Defender for Office 365 visibility (Microsoft, 2024c).

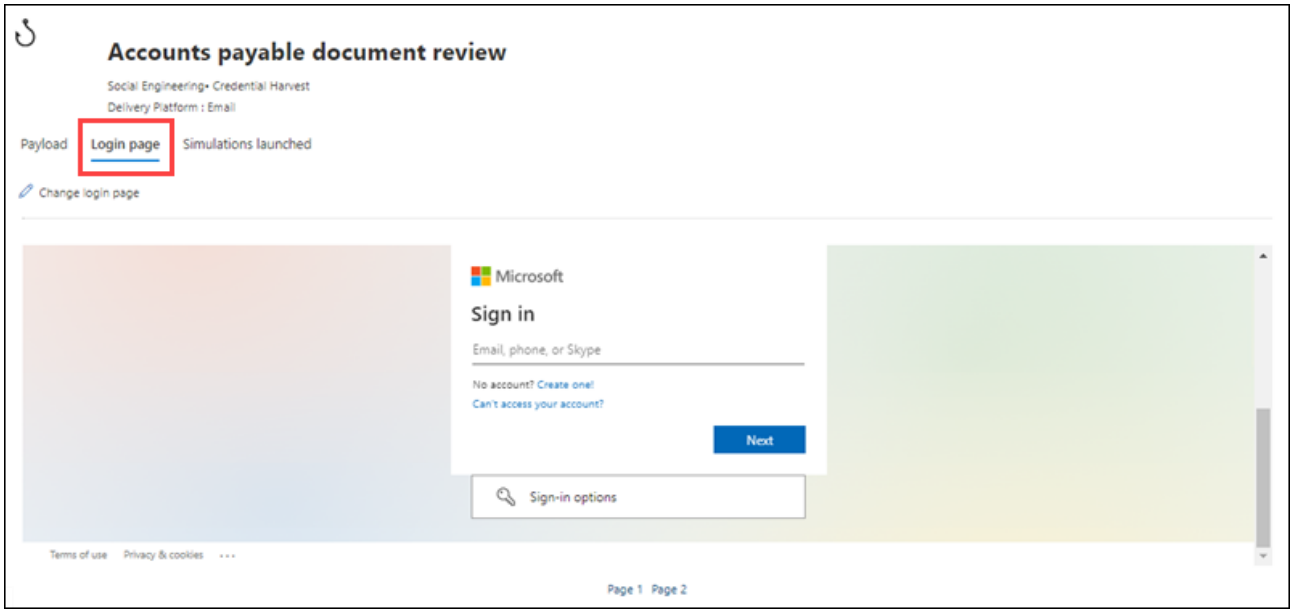

Figure 2. Generating phishing simulation login page (Microsoft, 2024c).

One another popular PDR is Cofense PhisMe. It delivers highly realistic, customizable phishing scenarios-complete with branded content, credential harvesting pages, malicious attachments, and parasitic URLs-directly to users' inboxes. Administrators can schedule campaigns, target specific cohorts via Azure AD or CSV imports, and choose from a library of global and tenant-specific payloads. PhishMe integrates with Outlook via a one-click "Report Phish" add-in and feeds reports into Cofense Triage for rapid human-verified analysis. Verified threats are automatically clustered and can be quarantined enterprise-wide via Cofense Vision or fed into SIEM and SOAR systems via rich telemetry and APIs. Key metrics-such as click-through rates, report volumes, time-to-report, training assignments, and completion rates-are tracked in the cloud-hosted Cofense portal, eliminating the need for on-premises infrastructure. This unified approach leverages intelligence from over 25 million global participants to continuously refine scenarios and accelerate incident response, reducing manual effort and improving resilience across the enterprise. (Cofence, n.d.).

Third one, KnowBe4 Security Awareness Training provides a comprehensive phishing simulation platform that allows administrators to design and deploy highly customizable phishing campaigns. Through an intuitive console, organizations can schedule tests, choose from hundreds of up-to-date templates-or create their own-that mimic real-world attack scenarios (e.g., credential harvesting, billing scams), and target specific user groups imported directly from Active Directory. When users interact with a simulated phish, the system provides immediate, just-in-time feedback

and training modules to reinforce learning. The platform also leverages AI-driven risk assessments to profile user susceptibility and recommend personalized training content. (Knowbe4, n.d.)

3.4 Summary of Theoretical Insights

Phishing simulations are increasingly used by organizations to address the growing threat of BEC. The theoretical background explored various aspects of these threats, highlighting how attackers leverage both basic and advanced phishing tactics—such as email spoofing, credential harvesting, and the use of AI-generated content—to exploit human behaviour and bypass traditional security controls.

The reviewed literature suggests that phishing simulations can support organizations in raising awareness, reinforcing behavioural change, and identifying vulnerabilities. When tailored to specific roles and followed by meaningful feedback, these exercises can strengthen employee readiness against evolving phishing methods. However, existing research also emphasizes certain limitations, including training fatigue, low engagement, and the risk of overestimating the impact of simulations without ongoing development and context-sensitive delivery.

The conceptual insights underline the complexity of BEC threats and the need for flexible, data-informed defence strategies. Particularly relevant to this thesis is the observation that while phishing simulations are widely adopted, few studies have tracked their long-term effectiveness across different organizational contexts. The theoretical perspectives outlined here provide a foundation for exploring how simulation outcomes develop over time and whether observable patterns can inform more sustainable awareness practices.

This understanding forms the basis for the empirical section of the thesis, where simulation results from three organizations are analysed to uncover behavioural trends and assess the value of simulations as a long-term awareness tool.

4 Implementation and Analysis

The analysis focuses on phishing simulation data collected from three anonymized Finnish organizations over a period of three years. By examining failure rates—defined as the percentage of users who clicked on phishing links—user behavior development is researched in responses to ongoing training.

The findings are interpreted through comparative analysis, highlighting similarities and differences between the organizations. While the data lacks detailed metadata about user roles or simulation content, it still reveals meaningful patterns related to training consistency, organizational differences, and behavioral fluctuations. These observations contribute to the thesis objective of understanding how phishing simulation outcomes evolve over time and under varying conditions.

4.1 Data Overview

In the context of phishing simulations, failure rate refers to the percentage of users who receive a simulated phishing message and proceed to click the embedded URL. If a user does not view the simulation message, the incident is not recorded as a failure but instead categorized under the missed rate. This thesis focuses exclusively on failure rates, as they offer a more direct and consistent measure for evaluating simulation effectiveness across different organizations.

Phishing simulation messages can be divided into two primary categories: personalized and non-personalized. In personalized scenarios, for instance, an IT employee might receive an email regarding urgent software updates or antivirus alerts. Because these topics are aligned with the recipient's daily responsibilities, the likelihood of engaging with the message—and therefore failing the simulation—is significantly higher.

In contrast, non-personalized messages typically take a more general form, such as simulated LinkedIn notifications or missed message alerts. While these lack job-specific relevance, they are still capable of eliciting a substantial failure rate across diverse employee groups.

While phishing simulation platforms offer both personalized and non-personalized message templates, this analysis was conducted without access to the specific content of individual messages.

Therefore, distinctions between message types are not analysed directly in this study. To preserve confidentiality, company names, branding, and identifiable data were removed or redrawn. All datasets were anonymized using the same process, enabling consistent and objective comparison across cases. The primary goal is to evaluate the effectiveness of phishing simulations, rather than draw attention to individual organizations.

4.2 Method of Analysis

The phishing simulation data analysed in this thesis was collected from three anonymized organizations, each of which had implemented phishing awareness training using simulated emails. The dataset specifically covers the first three years after each organization introduced phishing simulations into their security awareness programs. This timeframe was selected to allow a consistent comparison of early-phase performance trends and training effectiveness.

Each organization used the same simulation platform, which produced visual performance data in the form of monthly line graphs. These graphs depicted the failure rate—defined as the percentage of users who clicked on a phishing link after receiving and viewing the simulation message. This metric serves as the primary indicator of user susceptibility in this study. Other metrics, such as message reporting or missed message rates, were not consistently available and were excluded from the analysis.

To preserve confidentiality, the original visuals were anonymized and redrawn by the researcher. All company-specific details, such as logos, internal terminology, or branding, were removed. The organizations are referred to in this thesis as Company A, Company B, and Company C.

The data was manually transferred into spreadsheet software (e.g., Microsoft Excel) and visualized as line graphs to detect overall trends in simulation performance. The analysis focused on identifying behavioral patterns—such as consistent decline in failure rates, performance spikes, or periods of stagnation—within each organization and across all three cases.

It is important to note that the researcher did not have access to the content or design of individual phishing simulation messages. Therefore, this analysis does not distinguish between personalized and generic simulations, nor does it assess specific message themes. The focus remains on broader behavioral trends rather than message-level performance.

By applying a uniform analytical approach and focusing on comparable early training periods, this method enables meaningful cross-case interpretation while respecting data limitations and ethical boundaries.

4.3 Company A

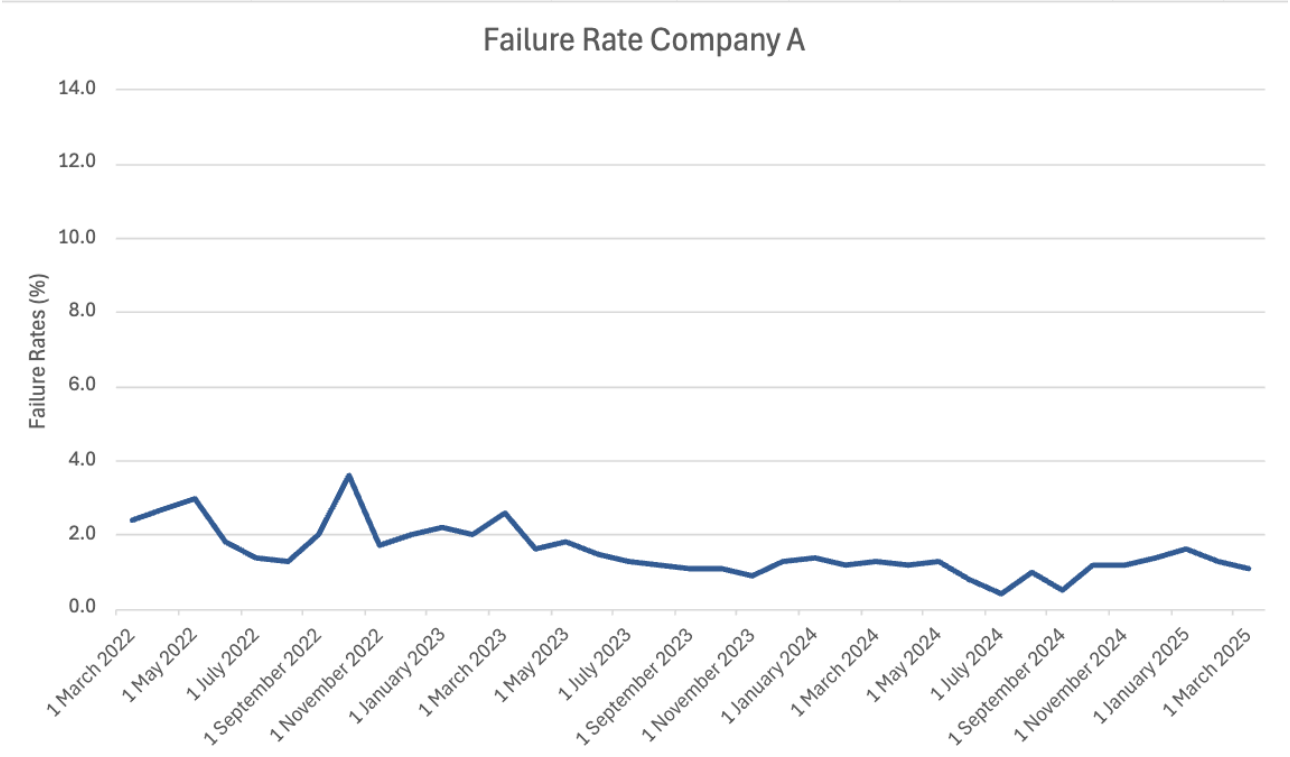

Figure 3. Company A failure rate.

The failure rate data for Company A spans from March 2022 to March 2025 and reflects employee performance in simulated phishing scenarios over a three-year period. The failure rate, defined as the percentage of users who clicked on a link within a phishing simulation email, is used as a key metric to assess training effectiveness and employee awareness.

The data reveals a gradual downward trend in failure rates, suggesting that phishing simulations may have had a positive impact on user behavior, particularly in the initial year. The highest failure peaks appear in May 2022 and December 2022, with the latter likely influenced by seasonal factors such as end-of-year holidays. This is consistent with earlier literature highlighting how psychological conditions—such as stress, distraction, or reduced staffing during holidays—can increase user vulnerability to phishing (Schöps et al., 2024).

From early 2023 onward, the failure rate stabilizes between approximately 1.2% and 2%, with minor fluctuations. Despite the continuous implementation of simulations, the data does not show a significant reduction in failure rates beyond the first year. This plateau suggests a possible saturation point in simulation effectiveness or reduced user engagement over time—issues raised in prior research (Ho et al., 2025).

While the general decline indicates some success, the relatively slow progress raises questions about the long-term impact of phishing simulations. For example, the failure rate in February 2025 is comparable to July 2023, despite over 18 months of continued training. This implies that additional measures, such as content variation, personalized simulations, or gamified training modules, may be required to maintain engagement and reinforce learning outcomes.

In summary, the simulation program at Company A appears to have led to initial improvements, but the diminishing effect over time highlights the need for evolving training strategies. These results align with broader findings that emphasize the importance of adaptive, user-centred approaches in sustaining long-term behavioral change.

4.4 Company B

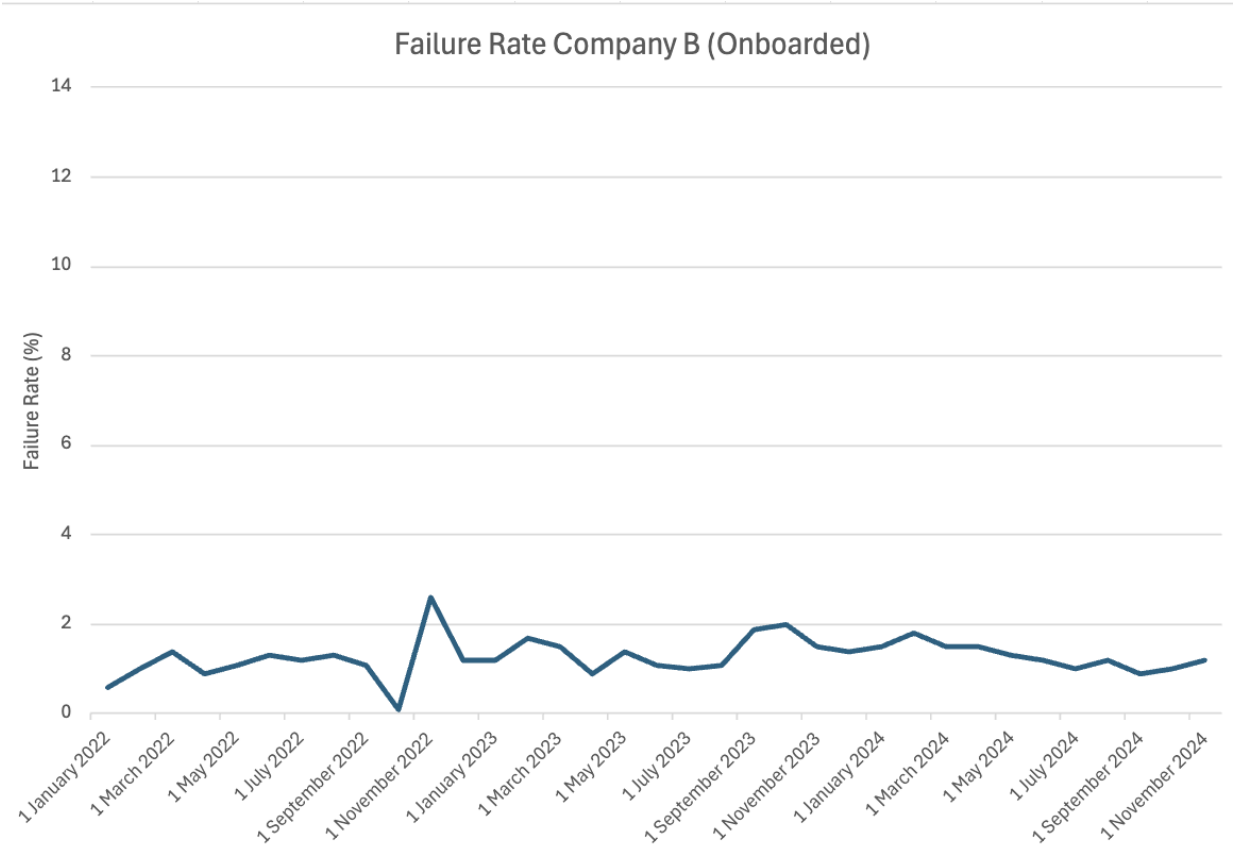

Figure 4. Company B (Onboarded) failure rate.

The phishing failure rate for onboarded users at Company B covers the period from January 2022 to November 2024 and shows a relatively consistent pattern with minor fluctuations. The term

“onboarded” refers to employees who were already part of the company and were not new hires during the simulation period.

The data begins with a moderate failure rate around 0.8%, which increases slightly and stabilizes between 1% and 2% over time. A notable dip occurs in November 2022, where the failure rate briefly drops to zero. While this outlier could be due to limited participation or technical anomalies

in data collection, it is immediately followed by a spike in January 2023, suggesting possible timing effects around the end-of-year holidays or a change in the difficulty of simulation content.

Unlike Company A, Company B does not show a clear downward trend. Instead, the data remains relatively flat after the first year, with minor short-term increases and decreases. This plateau suggests that phishing simulations did not produce strong cumulative improvements in employee behavior among the onboarded group. This pattern aligns with earlier findings by Ho et al. (2025), who observed that simulations alone may not sustain long-term behavioral change without continuous adaptation and reinforcement.

Interestingly, the consistency in results also reflects a stable training environment where failure rates do not worsen over time. This stability might indicate that baseline awareness is being maintained, even if no significant improvement is achieved. It also raises questions about the type and variation of training content provided—whether it remained static or evolved to reflect newer threats.

In summary, the onboarded users at Company B exhibited relatively steady phishing failure rates across the three-year period. While this suggests some level of resilience, the absence of a downward trajectory may indicate limitations in training effectiveness. This group’s data reinforces the broader theme that simulations must be continuously engaging, context-specific, and well-integrated into organizational culture to produce measurable long-term improvement.

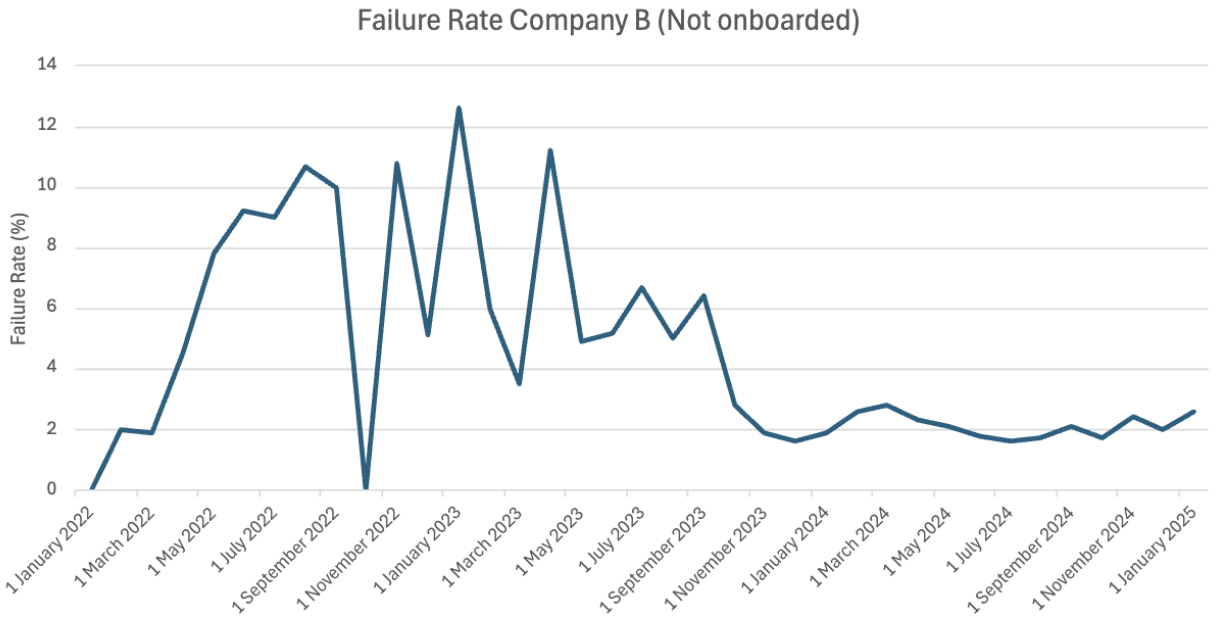

Figure 5. Company B (Not onboarded) failure rate.

The phishing failure rates for Company B's non-onboarded users—employees who did not participate in a structured onboarding process for phishing awareness—show significant volatility and elevated levels of susceptibility, particularly during the first two years of the observation period.

From an initial failure rate of 0% in early 2022, the rate escalated rapidly to over 11% by September of the same year, followed by repeated spikes—some exceeding 12%—throughout 2023. These fluctuations suggest pronounced inconsistency in user behavior, which may be attributed to a lack of systematic training or varying exposure to phishing content. Additionally, the identity of this user group remains unclear. It is unknown whether the individuals belong to a specific department, consist of new hires, or represent other internal classifications. Likewise, there is no visibility into the types or difficulty of the phishing simulations used in this segment.

One plausible explanation for the pronounced variability in the data is the possibility of pauses or irregularities in the simulation schedule. Temporary suspension of training may lead to lapses in

user readiness, allowing failure rates to rise sharply once simulations resume. This pattern is consistent with broader findings that underscore the importance of continuous reinforcement in awareness training to sustain behavioral improvements.

After peaking in mid-2023, the failure rate began to decline and stabilize between 2% and 3% during the latter part of the timeline. This may suggest that other informal learning mechanisms, peer influence, or broader organizational awareness were gradually internalized, leading to improved caution despite the absence of onboarding.

Comparing these results with the onboarded user group reveals a consistent trend: users who underwent structured awareness training performed more reliably and maintained lower failure rates over time. These results align with research by Beu et al. (2023), which highlights the role of training exposure, organizational integration, and employee engagement in shaping cybersecurity behavior.

This analysis reinforces the strategic value of onboarding-based phishing awareness programs. In the absence of structured early interventions, user behavior appears more erratic and risk-prone, and the organization may be left more vulnerable to social engineering threats.

4.5 Company C

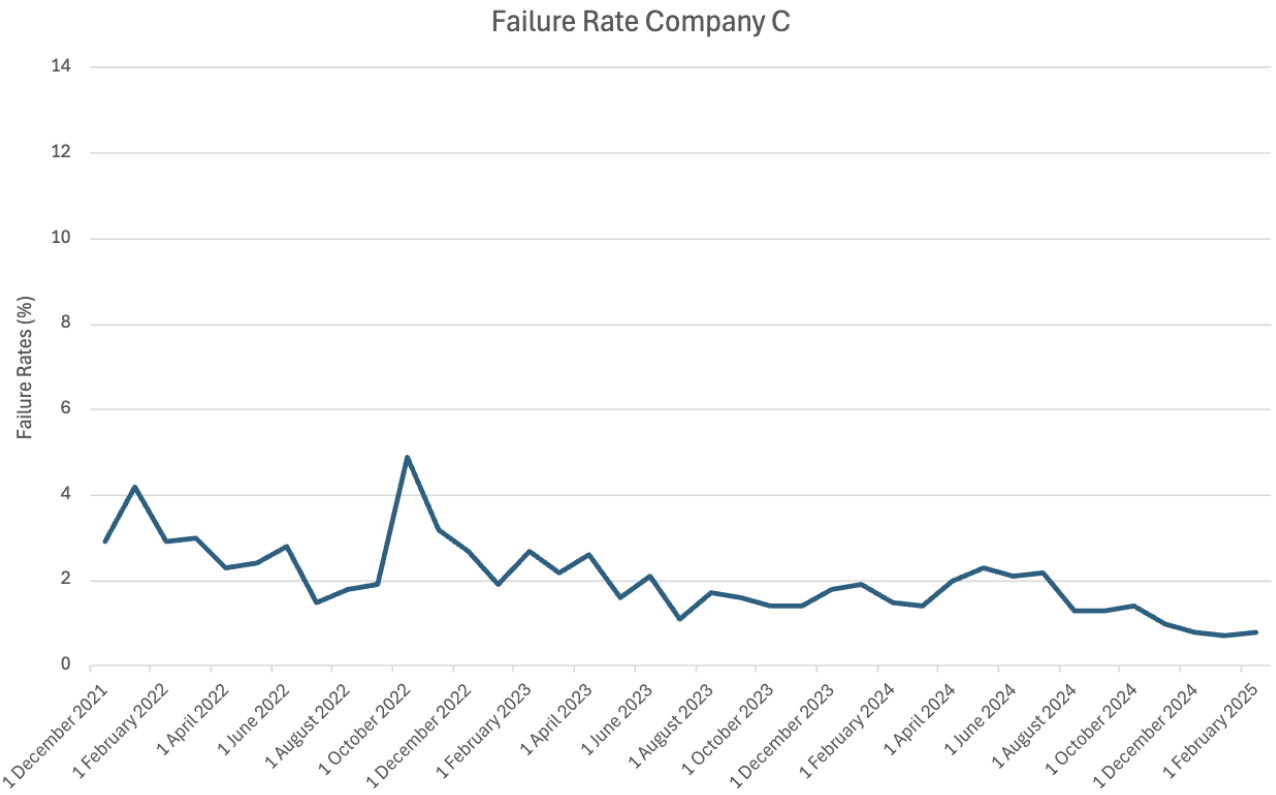

Figure 6. Company C failure rate.

The phishing simulation failure rates for Company C display a relatively stable and moderate trend over a three-year period, with notable improvement in the later stages of the timeline. The data begins in December 2021 with a failure rate slightly above 3%, followed by early fluctuations peaking near 4% in early 2022 and again in early 2023. Despite these initial spikes, the overall trajectory suggests a gradual decline in failure rates over time.

Throughout the middle of the dataset—from mid-2022 through mid-2023—the failure rates hover between 2% and 3%. This consistency suggests a degree of user familiarity with phishing simulations, possibly indicating either the presence of ongoing awareness training or a maturing security culture. A temporary spike near January 2023 breaks this pattern, though it quickly normalizes in subsequent months. As in other companies, the cause of such fluctuations may relate to unobserved factors such as seasonal timing, varying simulation difficulty, or campaign pauses.

From early 2024 onward, Company C demonstrates a clear downward trend. By early 2025, failure rates drop below 1%, marking a significant improvement from earlier levels. This reduction may reflect the cumulative effect of training, internal communication, or iterative exposure to phishing campaigns. However, without knowing the precise nature of the simulations—such as their frequency, targeting strategy, or difficulty—it is not possible to conclusively determine the source of the improvement.

Compared to Companies A and B, Company C appears to have maintained a relatively stable and moderately successful phishing awareness effort. While early progress was gradual, the data indicates tangible improvement over the long term. This pattern aligns with findings from Ho et al. (2025) and Sutter et al. (2022), which emphasize that measurable improvement in employee behavior often requires consistent reinforcement over extended periods.

In conclusion, Company C's phishing simulation data suggests that sustained, long-term exposure to training—even in the absence of dramatic early gains—can contribute to meaningful risk reduction. The steady decline in failure rates underscores the value of patient, continuous awareness efforts in building organizational resilience against phishing threats.

4.6 Comparative Summary Across Companies

A cross-comparison of the three participating organizations reveals several distinct patterns and commonalities in phishing simulation outcomes. While all companies implemented multi-year simulation campaigns, their results differ significantly in terms of performance stability, improvement trends, and behavioral response patterns. These differences offer valuable insights into how organizational context, onboarding practices, and training consistency may influence the long-term effectiveness of simulation-based phishing awareness programs.

Company A displayed a clear but slow improvement over time. The failure rate steadily decreased during the three-year period, with two prominent spikes—likely linked to seasonal factors such as holidays—followed by a consistent downward trajectory. Despite this gradual decline, the rate of improvement slowed considerably in the latter half of the timeline, raising questions about train-

ing saturation and the need for more dynamic content. Company A's data suggests that while sustained simulation exposure can yield incremental improvements, long-term behavioral change may plateau without adaptive strategies.

Company B provided a unique opportunity to compare two distinct employee groups: those who were onboarded with security training and those who were not. The onboarded group maintained low and stable failure rates, rarely exceeding 2%. In contrast, the non-onboarded group experienced significantly higher and more volatile failure rates, with several spikes reaching above 10%. These findings support existing research that emphasizes the importance of initial onboarding and role-specific security orientation (Beu et al., 2023). The data also suggests that inconsistent training schedules or pauses in simulation delivery may contribute to sharp behavioral regressions.

Company C exhibited the most consistent and gradual reduction in failure rates, beginning with moderate levels and trending steadily downward to below 1% by early 2025. This improvement, though slow, was stable and less volatile than in other organizations. Company C's performance suggests that a long-term, consistent simulation strategy—even in the absence of dramatic shifts—can be effective in reducing employee susceptibility to phishing. It also illustrates the importance of patience and repetition in simulation-based learning.

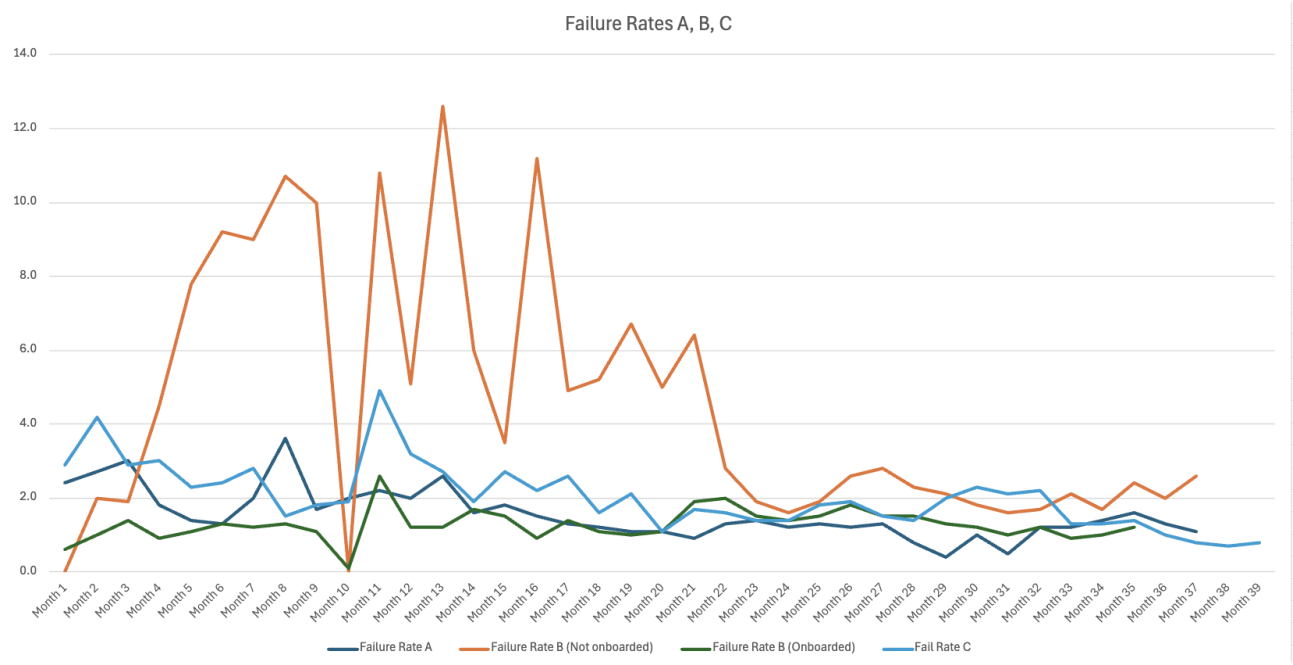

Figure 7. Failure rates from all companies.

5 Discussion

The results of the analysis highlight how phishing simulation performance evolved across three organizations over a three-year period. While all companies showed some reduction in failure rates, the pace and stability of improvement varied, often influenced by onboarding practices, training continuity, and organizational context. These findings offer insight into the real-world effectiveness of simulation-based awareness programs and raise further questions about training design, user engagement, and the limitations of current practices.

5.1 Interpretation of Results

The analysis of simulation data from three organizations offers important insights into how phishing simulation performance evolves over time and under varying conditions. Rather than focusing on specific monthly variations, this section interprets the broader behavioral and organizational patterns that emerged, linking them to the theoretical framework and prior research.

Training Engagement and Behavioral Change

All organizations showed a general decline in failure rates over the course of repeated simulations, supporting the argument that consistent exposure helps reinforce cybersecurity awareness. This aligns with findings by Sutter et al. (2022), who demonstrated that even minimal training—if sustained—can lead to improved employee caution. However, improvements tended to slow or plateau, indicating that repetition alone may not be enough to sustain engagement. Literature from Ho et al. (2025) further suggests that when simulation content becomes repetitive or lacks interactivity, users may disengage or develop overconfidence, limiting long-term behavioral change.

Onboarding as a Behavioral Foundation

Company B's data, which included both onboarded and non-onboarded user groups, underscores the significance of early exposure to phishing awareness training. The non-onboarded group exhibited more fluctuation and generally higher failure rates, suggesting that initial orientation into phishing awareness practices may have a compounding effect over time. This observation is supported by Beu et al. (2023), who identified job tenure and organizational connectedness as important predictors of phishing susceptibility.

Contextual Influences and Campaign Gaps

All companies experienced periodic spikes or regressions in failure rates, often aligning with holiday seasons or training pauses. These contextual variables—such as seasonal workload, employee availability, or cognitive fatigue—may significantly influence simulation outcomes. Schöps et al. (2024) emphasize the psychological dimension of simulations, including how stress, timing, and perceived scrutiny impact user reactions. These findings highlight the need for simulation schedules and content to be sensitive to broader organizational rhythms.

Design and Delivery Implications

The comparative analysis also points to key considerations for simulation design. While all three organizations made progress, the magnitude and consistency of improvement varied—indicating that effectiveness may be enhanced when simulations are tailored to user roles, deployed consistently, and delivered in psychologically safe ways. Theoretical perspectives (e.g., Davis & Grant, 2022) suggest that gamification, personalization, and adaptive feedback loops could provide a more engaging experience and help prevent behavioral stagnation.

5.2 Explanations for Variability

The variability observed in phishing simulation failure rates across the participating organizations can be attributed to a combination of structural, behavioral, and contextual factors. Rather than reflecting random fluctuation, these changes often reveal patterns rooted in organizational practices, user behavior, and training strategies.

One plausible explanation is the impact of training consistency. As the data showed, periods of regular simulation deployment were often followed by gradual improvement in employee performance. This supports previous findings by Sutter et al. (2022), who emphasized that even minimal but repeated training had measurable benefits over time. Conversely, interruptions or long gaps in training — as seen in some company datasets — appear to coincide with elevated failure rates or stagnant progress. This aligns with Ho et al. (2025), who noted that sustained user engagement is critical for simulation-based training to be effective.

A second factor involves employee onboarding and training coverage. The striking differences between onboarded and non-onboarded groups in one organization suggest that early exposure to phishing simulations may build resilience from the outset. This observation is supported by Beu et al. (2023), who highlighted the importance of organizational tenure and early exposure to training in reducing susceptibility to phishing.

Seasonal and situational factors also likely contributed to the spikes observed at specific points — particularly around holiday periods such as December or summer breaks. These findings echo the behavioral insights discussed by Schöps et al. (2024), who reported increased stress and reduced attentiveness as contributing factors to higher failure rates during periods of organizational disruption.

Finally, differences in simulation design — including message content, personalization, and difficulty — likely played a role, although the lack of metadata in this study limits any firm conclusions. Still, earlier work by Sutter et al. (2022) and Davis & Grant (2022) illustrates that simulation structure significantly influences user response, suggesting that variability may partly stem from how closely simulated emails resemble real-world threats.

In sum, variability in failure rates is not solely the result of user inconsistency or organization maturity for external phishing threats, but reflects a complex interplay between training regularity, organizational context, user psychology, and simulation strategy. Understanding this complexity is essential for designing more adaptive and sustainable awareness programs.

5.3 Recommendations and Implications for Practice

The findings of this study highlight several actionable insights for organizations seeking to enhance their cybersecurity resilience through phishing simulations.

First, the importance of training consistency cannot be overstated. Organizations that maintained a steady cadence of simulations—without long pauses—demonstrated more stable reductions in failure rates. This suggests that phishing simulations should be viewed not as one-time interventions, but as an ongoing process embedded in the organization’s broader cybersecurity strategy.

Regular exposure reinforces behavioral patterns and helps mitigate the risk of skill atrophy over time.

Second, onboarding programs should include phishing awareness training from the very beginning. The sharp contrast observed between onboarded and non-onboarded employees click rates underscores how initial exposure can significantly shape long-term behavior. Integrating simulations into onboarding workflows ensures that new employees adopt security-conscious habits early, rather than being left vulnerable during their most impressionable period.

Third, simulations must be adaptive and engaging. While general phishing messages can provide broad coverage, personalized simulations that reflect job roles and real organizational context are more likely to improve attentiveness and learning outcomes. However, care must be taken to avoid simulation fatigue—training content should be varied, and feedback mechanisms should be constructive rather than punitive.

Finally, data from simulations should inform broader awareness strategies. Rather than treating simulations as isolated tests, results should be used to identify risk patterns, tailor communications, and align with other security measures such as incident response planning. Involving the SOC in simulation design, as discussed in earlier chapters, can further enhance the realism and operational relevance of training.

In conclusion, effective simulation-based training requires a long-term, data-informed, and user-centered approach. Organizations that commit to continuous improvement, adapt training to evolving threats, and engage employees constructively are more likely to reduce human risk and build a culture of cybersecurity mindfulness.

5.4 Study Limitations

While this study offers valuable insights into the evolution of phishing simulation failure rates across three organizations, several limitations must be acknowledged. These factors may have influenced the depth, accuracy, and generalizability of the findings.

First, the lack of metadata significantly restricted the scope of analysis. The available data consisted solely of failure rate trends over time, without information on the specific content, frequency, or type of phishing simulations used. As a result, it was not possible to assess how message complexity, personalization, or scenario variation impacted user performance—factors that are known to influence susceptibility in previous research.

Second, the absence of demographic or contextual information (such as employee roles, tenure, or departmental breakdowns) limited the ability to analyze user-specific or role-specific vulnerabilities. While literature has shown that individual and organizational characteristics can shape simulation outcomes, this thesis could only focus on aggregate trends at the organizational level.

Third, some variations in simulation strategy may have occurred between organizations, but without standardized training frameworks or detailed documentation, these differences could not be systematically accounted for. The lack of information on simulation pauses, onboarding program content, or feedback mechanisms further complicated efforts to interpret changes in failure rates.

Finally, the sample size—limited to three anonymized companies—provides only a partial view of the broader organizational landscape. While these cases offer meaningful contrast in terms of training consistency and onboarding structure, additional organizations across different sectors would strengthen the reliability and generalizability of the conclusions.

Despite these constraints, the structured approach to data analysis and the triangulation with literature strengthen the trustworthiness of the results. However, readers should interpret the findings as exploratory rather than definitive, and future research should aim to address these limitations through richer datasets and broader organizational participation.

5.5 Trustworthiness of the Study

Ensuring the credibility and reliability of this research was a priority throughout the study. Several measures were taken to strengthen its trustworthiness, particularly in how the data was handled, interpreted, and presented.

The primary data consisted of real-world phishing simulation results gathered over a three-year period from three separate organizations. This provided a strong empirical basis for observing employee behavior in practical cybersecurity scenarios. Unlike survey-based or lab-based research, the dataset reflected actual interactions with simulated phishing messages, adding authenticity to the findings.

A key strength of the study lies in the consistency of how failure rates were defined and measured across all participating organizations. Failure was only recorded when a user both saw the simulation message and clicked the embedded phishing link. This standardized definition ensured that results were not based on subjective memory or post-event reporting but rather on clear, measurable actions. As a result, the comparisons between companies are grounded in a common metric, improving the validity of cross-case analysis.

In addition, the longitudinal structure of the data allowed the study to explore how user behavior evolved over time—not just as a snapshot but across multiple years of exposure. This temporal depth enabled the detection of seasonal variations, the effects of organizational onboarding practices, and the impact of potential pauses or changes in campaign strategies.

To ensure confidentiality and ethical integrity, all company names and identifying characteristics were removed. Data was anonymized at the source or by the researcher, and visuals were redrawn to remove internal labels or logos. The analysis avoided speculation in areas where metadata—such as employee roles or the nature of simulation content—was not available.

While the sample size was limited to three organizations, the combination of empirical data, consistent definitions, and transparent analysis methods contributes to the reliability of the results. Limitations, such as the lack of contextual metadata, were acknowledged throughout the thesis, and no overgeneralization was made beyond the scope of the available data.

In conclusion, the study was conducted with a careful and transparent approach to data interpretation, supported by consistent measurement practices and a focus on real-world user behavior. These efforts help establish a trustworthy foundation for the insights presented and provide a basis for future research.

6 Conclusion

The concluding section consolidates the findings of this research and reflects on their broader significance in the context of phishing simulations and organizational cybersecurity awareness. The analysis brings together observed trends, highlights key behavioral patterns, and examines their implications for practice. Attention is also given to the limitations of the study, areas for future research, and the broader value of simulation-based training as a long-term strategy for strengthening human resilience against Business Email Compromise threats.

6.1 Summary of Main Findings

This study examined how phishing simulation failure rates evolve over time across three Finnish organizations and explored patterns that may explain differences in user behavior. The results suggest that simulation-based training has the potential to reduce susceptibility to phishing, especially when implemented consistently and as part of a broader awareness strategy.

All three companies experienced a general decline in failure rates during the observation period, indicating that repeated exposure to phishing simulations can lead to improved employee awareness. However, the rate of improvement varied. Company C, which maintained consistent simulation frequency, showed the most stable downward trend, whereas interruptions or inconsistent training cycles—such as those seen in Company A and the non-onboarded group in Company B—correlated with spikes or irregular patterns.

Onboarding emerged as a critical factor. Company B's employees who had received phishing training from the start consistently outperformed those who had not, suggesting that early exposure plays a key role in shaping long-term behavior. Additionally, seasonal patterns and campaign timing appeared to influence failure rates, with notable increases during holiday seasons.

While progress was evident, the data also revealed a tendency toward performance plateaus over time. This highlights the need for continuous improvement, varied training content, and ongoing monitoring to sustain awareness and engagement.

Together, these findings emphasize that phishing simulations can be effective tools for reducing human risk—but only when they are thoughtfully designed, consistently delivered, and contextually informed.

6.2 Answer to the Research Question

This thesis is set out to answer to question:

How do phishing simulation failure rates evolve over time, and what organizational patterns can be observed?

The results show that phishing simulation failure rates generally decline over time when training is implemented consistently. All three companies exhibited reduced failure rates across the three-year period, although the pace and stability of improvement varied. Notably, organizations that maintained regular simulation activity—without long pauses—achieved more consistent reductions, suggesting that continuity is a key factor in long-term behavioral change.

Clear organizational patterns also emerged. For example, Company B demonstrated that onboarding employees with phishing awareness training from the outset leads to significantly better outcomes than introducing it later. Furthermore, seasonal spikes—such as those around holidays—highlight the impact of timing and contextual factors on user performance.

In sum, failure rates tend to decrease over time, but the extent of improvement depends heavily on how training is structured, introduced, and maintained within each organization.

6.3 Suggestions for Future Research

While this thesis provides insight into the long-term development of phishing simulation failure rates across three organizations, several areas remain open for further exploration.

One key direction for future research is the role of simulation content and complexity. This study lacked access to metadata on the specific phishing scenarios used, limiting analysis of how message type, personalization level, or thematic relevance may influence user behavior. Future studies

could benefit from a more granular dataset that includes simulation characteristics and user role mapping, enabling deeper behavioral analysis.

Second, additional research is needed on the psychological and emotional dimensions of simulation training. Prior studies (e.g., Schöps et al., 2024) have highlighted the impact of stress, shame, and self-efficacy on training outcomes. Including employee feedback or qualitative assessments could offer valuable perspectives on how users experience simulated threats and how that affects their learning and trust in security programs.

Third, future work could explore organizational culture and leadership influence on simulation outcomes. For instance, the visibility of executive support, internal communication tone, and the presence of reinforcement mechanisms (e.g., recognition or follow-up training) may all contribute to sustained improvements.

Additionally, research should address the impact of AI-generated phishing content, which is rapidly changing the threat landscape. As attackers adopt large language models (LLMs) to craft more convincing messages, it is critical to understand how well current training prepares users to recognize these new threats.

Another valuable direction involves studying the durability of behavioral improvements—specifically, how long reduced failure rates persist after training has ended. This would support evidence-based planning for the timing and frequency of simulations.

Furthermore, expanding the dataset to include a greater number of organizations from different sectors and countries could help validate whether the patterns observed in this study are generalizable. Comparative studies across diverse settings would enhance understanding of how organizational size, security maturity, and sector-specific risks shape simulation performance.

Finally, incorporating broader behavioral indicators—such as time-to-report or false-positive rates—would allow future research to go beyond failure rates and better assess overall user awareness and threat response capacity.

6.4 Final Reflections

This thesis set out to examine how phishing simulation failure rates evolve over time and what organizational patterns can be observed across three real-world cases. While the scope was focused and the data anonymized, the findings offer meaningful insight into how consistent training, onboarding practices, and contextual factors shape employee behavior in response to phishing threats.

One of the central takeaways is the value of longitudinal data collected through consistent simulation practices. Rather than relying on self-reported behavior or cross-sectional surveys, this study leveraged platform-based click-rate metrics, allowing for objective comparison and temporal tracking. Although the dataset did not include message-level metadata, the standardization of failure rate definitions across all companies ensured that the comparisons rested on shared and reliable indicators.

At the same time, the findings reflect the inherent complexity of human behavior in cybersecurity contexts. Simulations can reduce susceptibility, but their long-term impact varies based on how they are delivered, who receives them, and whether the organizational environment supports learning. The occasional rise in failure rates—even after extended training periods—demonstrates that behavioral change is not linear. It is influenced by timing, training pauses, external stressors, and perhaps even the evolving tactics of attackers.

This research contributes to a growing body of work emphasizing the importance of user-centered approaches in cybersecurity. While technical defenses remain essential, the human layer continues to be both a vulnerability and a strategic asset. Future efforts should integrate behavioral science, training design, and feedback mechanisms to foster not only awareness but also confidence and trust among users.

In reflecting on this thesis process, even with limited data, structured analysis can yield practical and academically valuable insights. The study underscores the importance of measuring what matters, questioning assumptions about training effectiveness, and viewing cybersecurity as a continuous learning journey rather than a one-time intervention.

References

Aon. (n.d.). *Bypassing MFA: A forensic look at Evilginx2 phishing kit*. Aon Cyber Labs. Retrieved December 12, 2024, from https://www.aon.com/cyber-solutions/aon_cyber_labs/bypassing-mfa-a-forensic-look-at-evilginx2-phishing-kit/

Beu, N., Jayatilaka, A., Zahedi, M., Babar, M. A., Hartley, L., Lewinsmith, W., & Baetu, I. (2023). *Falling for phishing attempts: An investigation of individual differences that are associated with behavior in a naturalistic phishing simulation*. Computers & Security, 131, 103313. https://doi.org/10.1016/j.cose.2023.103313

Brachten, C. E. W. (2025). *Impact of AI personalization on email clicks and conversions: Insights from a real-world AI-personalized phishing simulation*. USENIX. https://studenttheses.uu.nl/bitstream/handle/20.500.12932/48815/Thesis_redacted.pdf

Burgess, M. (2024, August 8). M*icrosoft's Copilot raises new phishing and data extraction concerns*. Wired. Retrieved December 24, 2024, from https://www.wired.com/story/microsoft-copilot-phishing-data-extraction/

Cartier, M. (2025, January 10). *Business email compromise statistics: 2025 edition*. Hoxhunt. Retrieved from https://hoxhunt.com/blog/business-email-compromise-statistics

Ciupe, A., & Orza, B. (2023). *Reinforcing cybersecurity awareness through simulated phishing attacks: Findings from an HEI case study*. In 2023 IEEE International Conference on Multimedia Systems and Applications. IEEE. https://doi.org/10.1109/ICERECT56837.2022.10060595

Cofense. (n.d.). *Solution brief: Cofense phishing detection and response platform.* Cofense. Retrieved April 28, 2025, from https://cofense.com/product-overview/

Davis, N., & Grant, E. S. (2022). Si*mulated phishing training exercises versus gamified phishing education games.* In 2022 Fourth International Conference on Emerging Research in Electronics, Computer Science and Technology (ICERECT). IEEE. https://doi.org/10.1109/ICERECT56837.2022.10060595

Federal Bureau of Investigation. (2023, June 9). *Business email compromise: The $50 billion scam.* Internet Crime Complaint Center. Retrieved November 7, 2024, from https://www.ic3.gov/PSA/2023/PSA230609

Fortinet. (2024). *Email spoofing.* Fortinet CyberGlossary. Retrieved from https://www.fortinet.com/resources/cyberglossary/email-spoofing

Harju, J. (2025, March 11). *”Teillä on avoin kulu, hoitakaa se viipymättä” – tällaisia viestejä löytyi kymmenittäin luksushotelliin majoittuneiden miesten puhelimista.* Helsingin Sanomat. Retrieved April 4, 2025, from https://www.hs.fi/suomi/art-2000011076281.html

Ho, G., Mirian, A., Luo, E., Tong, K., Lee, E., Liu, L., Longhurst, C. A., Dameff, C., Savage, S., & Voelker, G. M. (2025). *Understanding the efficacy of phishing training in practice.* In 2025 IEEE Symposium on Security and Privacy (SP) (pp. 76–76). Retrieved from https://people.cs.uchicago.edu/~grantho/papers/oakland2025_phishing-training.pdf

Hoxhunt. (2023, April 18). *MFA fatigue: What it is and how to prevent it.* Hoxhunt. Retrieved November 7, 2024, from https://hoxhunt.com/blog/mfa-fatigue

IBM. (2024). *Cost of a Data Breach Report 2024.* IBM Security. Retrieved from https://www.ibm.com/downloads/documents/us-en/107a02e94948f4ec

Ironscales. (n.d.). *Credential harvesting.* Ironscales. Retrieved April 23, 2024, from https://ironscales.com/glossary/credential-harvesting

JAMK University of Applied Sciences. (2024). *Research ethics guidelines* (L. Liimatainen, J. Hautamäki, E. Kirjalainen, M. Kokko, K. Korhonen, S. Laitinen-Väänänen, K. Norvapalo, A. Törn-Laapio, & S. Hyvätti, Eds.). https://www.jamk.fi/en/media/41521

KnowBe4. (n.d.). *Adaptable, AI-driven security awareness training.* KnowBe4. Retrieved April 28, 2025, from https://www.knowbe4.com/products/security-awareness-training

Lain, D., Kostiainen, K., & Capkun, S. (2022). *Phishing in organizations: Findings from a large-scale and long-term study.* In 2022 IEEE Symposium on Security and Privacy (SP). https://doi.org/10.3929/ethz-b-000588856

Microsoft. (2024a). *Simulate a phishing attack with attack simulation training.* Microsoft Defender for Office 365 documentation. Retrieved November 7, 2024, from https://learn.microsoft.com/en-us/defender-office-365/attack-simulation-training-simulations

Microsoft. (2024c, October 22). *Simulate a phishing attack with attack simulation training. Microsoft Defender for Office 365 documentation.* Retrieved April 28, 2025, from https://learn.microsoft.com/en-us/defender-office-365/attack-simulation-training-simulations

National Cyber Security Center of Finland. (2024). *Increasing number of M365 data breaches utilise AiTM phishing*. Retrieved November 7, 2024, from https://www.kyberturvallisuuskeskus.fi/en/ncsc-news/instructions-and-guides/increasing-number-m365-data-breaches-utilise-aitm-phishing

Phriendly Phishing. (n.d.). Dan*gerous email attachments: What to look out for and how to stay safe.* Phriendly Phishing. Retrieved April 23, 2025, from https://www.phriendlyphishing.com/blog/dangerous-email-attachment

SANS Cyber Defense. (2024, October 24). E*dge out your next BEC adversary: Investigation and prevention strategies using OSINT [Video].* YouTube. https://www.youtube.com/watch?v=J01dyi5YFFw

Schöps, M., Gutfleisch, M., Wolter, E., & Sasse, M. A. (2024). *Simulated stress: A case study of the effects of a simulated phishing campaign on employees' perception, stress and self-efficacy.* In Proceedings of the 33rd USENIX Security Symposium. USENIX. https://www.usenix.org/conference/usenixsecurity24/presentation/schöps

Sutter, T., Bozkir, A. S., Gehring, B., & Berlich, P. (2022). *Avoiding the hook: Influential factors of phishing awareness training on click-rates and a data-driven approach to predict email difficulty perception.* IEEE Access, 10, 100540–100563. https://doi.org/10.1109/ACCESS.2022.3207272

Verizon. (2024). *Data Breach Investigations Report (DBIR).* Verizon Business. Retrieved from https://www.verizon.com/business/resources/reports/2024-dbir-data-breach-investigations-report.pdf

Weinert, A. (2024*). Defeating adversary-in-the-middle phishing attacks*. Microsoft Tech Community. https://techcommunity.microsoft.com/blog/microsoft-entra-blog/defeating-adversary-in-the-middle-phishing-attacks/1751777