Table of contents

1 Introduction

This document describes the Computer Emergency Response Team (CERT) engagement that Fox-IT performed for Eindhoven University of Technology (hereinafter: TU/e) during the period from 11 January 2025 until 11 April 2025. This chapter starts with describing the incident background in Section 1.1. Section 1.2 introduces the investigation questions that were posed. The chapter finishes with a reading guide in Section 1.3.

1.1 Background

On Saturday 11 January 2025 at 21:55, SURFsoc[1] was alerted of potential malicious activity within the infrastructure of TU/e. Analysis of multiple consecutive alerts uncovered that the default domain administrator account and domain controller were involved. One of the alerts indicated the use of the CrackMapExec WMIExec module by the default domain administrator `account_hp1` on domain controller `SYSTEM_DC1_PROD`, configuring the domain controller to allow Windows Remote Assistance[2].

Following this analysis, SURFsoc decided to escalate the security incident to TU/e at 22:48. Contact with TU/e was established at 22:51, during which TU/e explained that they were already aware of potential malicious activity. According to protocol, SURFcert[3] was also informed of the incident at 23:06.

Meanwhile, SURFsoc informed Fox-IT's Computer Emergency Response Team (hereinafter: FoxCERT) at 23:20 of the potential incident that TU/e was facing. This allowed FoxCERT to prepare for an emergency call from TU/e. FoxCERT received the call from TU/e at 23:50. During this call, the aforementioned SURFsoc alerts at TU/e were assessed and determined to be of high importance and urgency, requiring immediate assistance to contain and remediate the now deemed serious security incident. An intake call between FoxCERT and TU/e took place approximately 25 minutes later on 12 January 2025 at 00:15.

During the intake, FoxCERT and TU/e verified earlier suspicions and determined that unauthorized access was gained to high privileged accounts and critical components of TU/e's IT infrastructure. To prevent further manual activity by the adversary, FoxCERT advised TU/e to immediately block all inbound and outbound network traffic and to terminate all current connections. TU/e was already prepared for this scenario, which resulted in swift execution that same night on 12 January 2025 at 01:17.

Upon TU/e's request, FoxCERT provided on-site assistance in the first week of the security incident. FoxCERT first arrived on-site on 12 January 03:00; the early morning after the security incident was identified. Fox-IT joined a crisis response meeting 10 minutes later, at 03:10. In this meeting, the implemented containment measures and more details of the security incident were discussed. For Fox-IT, this crisis response meeting marked the start of an extensive CERT engagement between Fox-IT and TU/e.

[1] SURFsoc is a SIEM-based Security Operations Center service tendered by SURF, delivered by Fox-IT. Read more at https://www.surf.nl/diensten/surfsoc.

[2] The Windows-RemoteAssistance-Exe component allows a user to receive hands-on-keyboard assistance from another person on a different location. Read more at https://learn.microsoft.com/en-us/windows-hardware/customize/desktop/unattend/microsoft-windows-remoteassistance-exe.

[3] SURFcert is SURF's Computer Security Incident Response Team (CSIRT) and collaborates with FoxCERT and SURF-members with a FoxCERT retainer. Read more at https://www.surf.nl/diensten/surfcert and https://www.fox-it.com/nl-en/protection-detection-and-response/incident-response/.

1.2 Objectives

Fox-IT was tasked to conduct an investigation to provide answers on the following four investigation questions:

1. What happened?
2. How did it happen?
3. What is the scope of the compromise?
4. What data was accessed by the adversary?

While obtaining preliminary answers on the investigation questions, Fox-IT continuously used the gathered information to additionally provide TU/e with mitigation steps to remediate the compromise.

1.3 Reading guide

This document describes how the CERT engagement was organised, which investigation approach was taken, what findings were made and conclusions based on the findings. This is divided over the remaining chapters as follows:

Chapter 2 describes the approach and methodology used in the investigation.
Chapter 3 provides the detailed findings.
Chapter 4 contains the conclusions based on these findings.

Appendix A provides the Indicators of Compromise.

Dates and times mentioned in this report represent the time in Central European Time zone (CET/CEST), unless stated otherwise.

This document regularly refers to tactics, techniques, and procedures (TTPs) as described in the MITRE ATT&CK framework.[4] Such references are put in square brackets ([ ]), e.g.: Exploit Public-Facing Application [T1190].

Usernames, hostnames, and IP addresses have been altered to obfuscate the true names, but they are known to the relevant parties.

[4] The MITRE ATT&CK framework is a framework to which adversary activity can be mapped created by The MITRE corporation. Read more at https://attack.mitre.org/.

2 Emergency response approach

This chapter describes the multipronged approach that Fox-IT followed during the emergency response phase of the incident. Section 2.1 describes the investigation approach. Section 2.2 covers the containment measures provided during the CERT engagement, whereas Section 2.3 describes the provided mitigation measures.

2.1 Investigation approach

This section describes the approach of the investigation. Subsection 2.1.1 describes the four main investigation tracks that were setup. Hereafter, Subsection 2.1.2 provides the investigation collection methods that were used.

2.1.1 Investigation tracks

Fox-IT employed four investigation tracks to focus the investigation on answering the most important questions in a time-efficient manner. The remainder of this subsection describes the four tracks in more detail.

Track 1: Identifying patient zero & initial foothold

The first track focused on the identification of the initial point of entry in TU/e's infrastructure. It primarily followed a "follow-the-evidence principle"; starting from an initially identified malicious activity and tracing that back to its origin.

This track aimed to provide insights that feed into the mitigation of the incident. The rationale behind this, is that the adversary (or another adversary) could potentially start a new attack, if the initial point of entry is not identified and mitigated.

Track 2: Identifying the route to highest level of access

The second track focussed on identifying the highest level of access the adversary managed to obtain. This typically involves identifying adversary activity from several categories in the MITRE ATT&CK framework, such as Privilege Escalation [TA0004], Credential Access [TA0006] and Lateral Movement [TA0008].

This track aimed to determine how the adversary obtained the highest level of access (known as domain administrator access) in TU/e's Active Directory domains. Section 3.3 describes how the adversary likely gained the highest level of privileges and Section 3.6 dives deeper into the extent of data access.

Track 3: Identifying command & control (C2) and persistence

The third track focused on the identification of command and control (C2) and persistence mechanisms. These mechanisms allow an adversary to send instructions and maintain access to the infrastructure respectively. Identification of C2 and persistence mechanisms is a prerequisite for successful remediation.

This track aimed to determine how the adversary performed their actions and maintained access once they had the highest level of access. More information and findings can be found in Section 3.4.

Track 4: Identifying data access

The data access track focused on identifying what data the adversary had gained access to. This includes, but is not limited to:

- Data that was on screen and may have been used immediately by the adversary to progress to their goals.
- Network and/or domain discovery [TA0007] data that was collected for analysis.
- Data that was collected [TA0009] and exfiltrated [TA0010].

This track aimed to first determine if important and/or sensitive data was potentially accessed by the adversary. If so, the track aimed to identify signs of exfiltration of such data.

2.1.2 Collection of investigation material

Fox-IT collected investigation material from several sources. The main collection methods were:

- **Collecting light-weight investigation packages via Acquire**. Acquire is a data acquisition tool based on the Dissect[5] framework. The acquisition was primarily performed on the ESXi hypervisor's NFS data store. In some cases, Acquire was executed from the running operating system.
- **Copying of full (virtual) disk images**. This was done by creating a copy of raw (virtual) disks. This data collection method was used whenever the investigation required more detailed information than the light-weight investigation package could provide or when creating such a package was not feasible.
- **Copying FortiGate firewall logs**. The logs were collected from the FortiAnalyzer management system.

2.2 Recommendations for containment

During the first hours of the incident, Fox-IT recommended TU/e to implement several containment measures. These recommendations aimed to deny the adversary access to TU/e infrastructure and to prevent potential automated spread of malware[6]. Table 1 shows the containment measures that were recommended.

Table 1: Recommendations for containment provided by FoxCERT during the first hours of the CERT engagement.

Date/time	Recommendation	Rationale
12-1-2025 00:15	Disable in/outbound traffic from/to the network and terminate established connections from/to the network. Except EDR/SIEM telemetry	Deny the adversary access to the network and prevent further spread and/or impact
12-1-2025 00:15	Deny new VPN connections and terminate established connections	Deny the adversary access to the network and prevent further spread and/or impact
12-1-2025 00:15	Isolate systems which were accessed by the adversary	Contain current impacted systems and prevent potential automated propagation through the network
12-1-2025 00:15	Reset passwords of high privileged accounts and revoke issued Kerberos tickets and reset	Mitigate risk of (future) use of high privileged accounts by adversary
12-1-2025 00:15	Reset KRBTGT passwords of domain controllers twice	Mitigate risk of use of golden tickets generated by the adversary
12-1-2025 00:15	Secure and check the integrity of backups	Ensure that backups are available in the event where system restores are necessary

[5] Dissect is an open-source forensic framework developed by Fox-IT. Read more at https://dissect.tools/.

[6] Malware that spreads itself without manual input is called a 'worm'. Read more at https://learn.microsoft.com/en-us/defender-endpoint/malware/worms-malware.

2.3 Recommendations for remediation

Fox-IT provided TU/e with technical and tactical input for the remediation activities during the CERT engagement. TU/e, together with Fox-IT, decided to implement the containment measures listed in Table 2.

Table 2: Technical and tactical recommendations for remediation provided by FoxCERT during the CERT engagement.

Date	Recommendations	Rationale
13-1-2025	Rebuild and/or restore compromised systems to known-good state	Rebuilding systems from scratch or restoring a system to a known-good state allows for safe recovery of systems and is preferred over manual eradication of affected systems
13-1-2025	Review domain controller configuration by (Fox-IT) red-team	Identify and mitigate high security risks to prevent future incidents
13-1-2025	Onboard EDR in Security Operations Center	Increase proactive security monitoring on systems
14-1-2025	Scan systems on presence of malware, isolate and restore/rebuild compromised systems	Prevent malware artefacts from resurfacing and the adversary from re-entering the network or reach actions on objectives
15-1-2025	Review possibility to implement an (emergency) Intrusion Detection System	Further increase proactive security monitoring on a network level
17-1-2025	Isolate unmanaged and unmonitored systems	Decrease attack surface and therefore security risks
17-1-2025	Deny outbound connections to common remote desktop tooling	Decrease usage of command and control commonly used by adversaries

2.3.1 Recovery strategy for compromised systems

Whenever a system is (potentially) compromised, Fox-IT recommends restoring the system from a known safe state. The general instructions to do so were as follows:

1. Install the system from a known safe state or rebuild from scratch it if no safe state exists.
2. Apply all security updates.
3. Install anti-virus and EDR software.
4. Install the required additional software for the system to perform its tasks.

3 Investigation results

This chapter describes the results from the investigation tracks. Intermediate conclusions are summarised at the end of each section or subsection, if applicable. The conclusions based on all findings are provided in chapter 4.

This chapter starts with a schematic overview of the findings in Section 3.1. Section 3.2 provides the findings regarding the first moment of adversary activity. Hereafter, Section 3.3 elaborates on the findings that most likely explain how the adversary raised their privileges in TU/e's network. Section 3.4 describes the activity performed by the adversary after gaining the highest privileges in the network. Because of the adversary's large extent of access, Section 3.5 dives deeper into the scope of the compromise, whereas Section 3.6 elaborates on the potential data access by the adversary. Finally, Section 3.7 sketches a general profile of the threat actor based on the identified activity.

3.1 Overview of findings

Figure 1 depicts a general timeline with key findings of the investigation. Each finding is appointed a tactic according to the MITRE ATT&CK framework in red, such as "Initial Access" and "Lateral Movement". In blue, TU/e's containment measure of disconnecting the network from the internet is depicted.

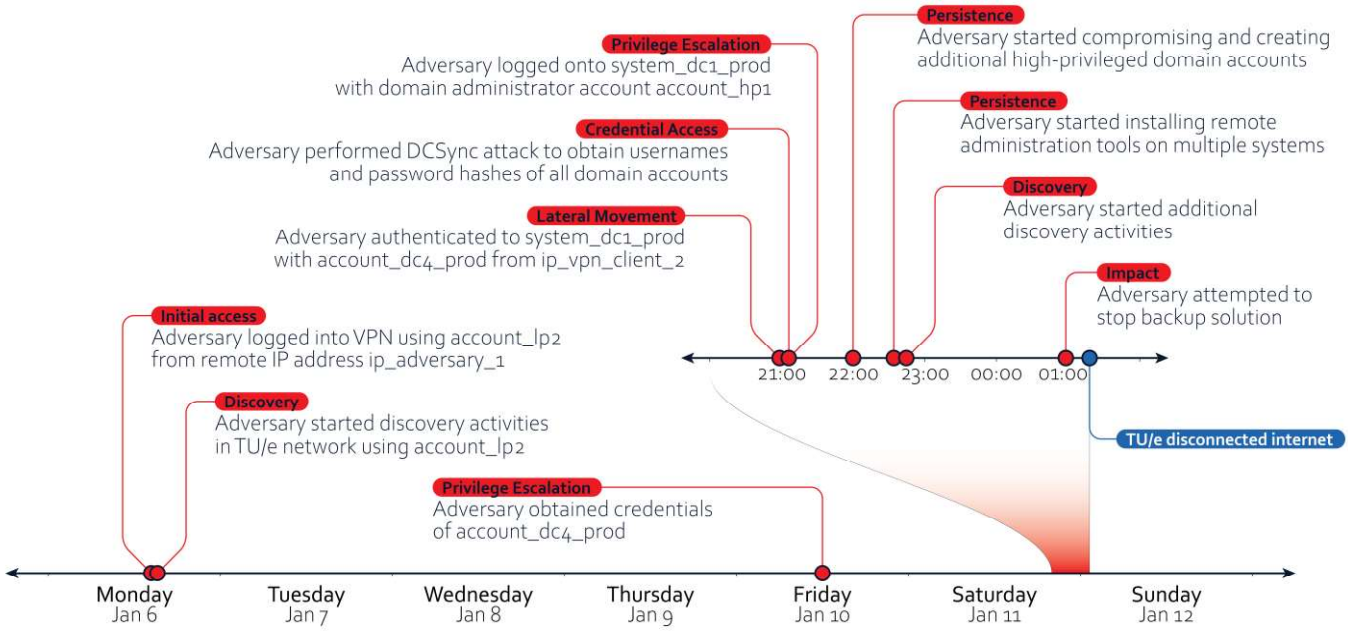

Figure 1: Overview of the incident timeline of the key findings of the investigations.

3.2 Initial foothold and early discovery actions

The investigation performed by Fox-IT uncovered that the adversary gained access to the network of TU/e via its remote work or Virtual Private Network (VPN) solution. Details of this initial access are described in Subsection 3.2.1. Hereafter, Subsection 3.2.2 explains how the adversary likely gained this access.

3.2.1 First adversary activity using legitimate user accounts on TU/e's VPN

On 6 January 2025 at 13:57 logs from TU/e's VPN solution show the account `account_lp1` failed to authenticate from the remote IP address `ip_adversary_1`. Almost ten minutes later, at 14:08, this same IP address was used to successfully log into the account of `account_lp2`. At 14:13, another account, that of `account_lp3`, was also successfully logged into from this same IP address. The usage of a single IP address to log into multiple accounts is an indicator for suspicious activity.

An hour later, at 15:14, the account `account_lp2` again logged in from the same IP address. During this VPN session, the account started connecting to multiple systems in TU/e's network. According to the available logs, these connections were atypical for the account `account_lp2`. Furthermore, the rapid succession of these connections indicates that these authentications were performed in an automated fashion.

The atypical and automated authentications from the account `account_lp2` to multiple systems within TU/e's network that followed, raised the suspicious nature of these VPN sessions. Furthermore, the IP address used to login to the VPN belonged to a hosting provider, which is not common for login actions to a VPN for legitimate users. Because of these suspicious characteristics, Fox-IT links the IP address `ip_adversary_1` and all related activity to the adversary.

Based on the related activity, Fox-IT identified two additional remote IP addresses that can be linked to the adversary. These IP addresses are `ip_adversary_2` and `ip_adversary_3`. The adversary used these IP addresses to connect to TU/e's VPN. Table 11 in Appendix A shows a list of all VPN sessions that Fox-IT related to the adversary.

Fox-IT identified suspicious successful VPN sessions to the accounts prior to the timestamps mentioned. However, these logins could not be directly related to the adversary. From 6 January 2025 onward, the activities of the adversary could unambiguously be linked to the malicious activity five days later. For this reason, Fox-IT considers 6 January 2025 at 14:08 as the start of this incident; the first moment the adversary successfully logged in. However, it should be kept in mind that the same adversary, or other actors, might have had access prior to that moment.

3.2.2 Adversary likely gained access to end user accounts via leaked credentials

The means through which an adversary gains access to legitimate VPN accounts are commonly either by using valid leaked credentials, or opportunistically trying combinations of usernames and passwords. For both means to be viable, the VPN should ideally not enforce multi-factor authentication (MFA). TU/e confirmed that their VPN solution did not enforce MFA.

In case of opportunistically trying combinations of usernames and passwords, one would typically see a high number of failed login attempts in the VPN logs. However, these were not present, making it more likely that the adversary had foreknowledge of the credentials.

Furthermore, Fox-IT's Threat Intelligence Team found credentials in a publicly available credential leak document for one of the two accounts to which the adversary successfully logged on. For the other account, the Threat Intelligence Team found traces that at least two known data breaches contain information about the account. These findings for the two accounts raise the likeliness that information about the accounts was available to the adversary prior to the login actions. Therefore, Fox-IT considers it most likely that the adversary gained access to TU/e's VPN by using leaked credentials of the two successfully compromised accounts.

Fox-IT considers 6 January 2025 at 14:08 the start of this incident. At this moment, the adversary successfully logged into TU/e's VPN with the account account_lp2 from an uncommon IP address. Five minutes later, at 14:13, this IP address was also used to log into the account account_lp3. Based on the suspicious nature and the reuse of the IP address together with the suspicious follow-up activity, Fox-IT linked the IP address to the adversary. Fox-IT considers it most likely that the adversary gained access to the two accounts via leaked credentials.

3.3 Privilege escalation and implications of resulting full domain control

The investigation identified that the adversary obtained the highest privileges within a Microsoft Windows network, known as domain administrator enterprise administrator privileges, in the entire TU/e Active Directory domain forest. This includes the domains `DOMAIN_2` (the campus domain) as well as the `DOMAIN_1` (the root domain).

As part of the investigation, Fox-IT was not able to find irrefutable evidence that proves exactly when and how the adversary obtained domain administrator credentials. However, Fox-IT identified multiple indirect traces that give indications on both the when and how questions. This section elaborates on these indirect traces and the hypothesis to how the privilege escalation to enterprise administrator privileges took place.

The section starts with Subsection 3.3.1, indicating how it was known from the start of the engagement that the adversary obtained enterprise administrator privileges. Subsection 3.3.2 dives into the indirect traces to explain the most likely moment and method of obtaining these privileges.

3.3.1 Elaboration on SOC alerts that resulted in detection of the adversary with highest privileges

As stated in Section 1.1, the incident was detected based on malicious activities on a domain controller. In total 63 alerts were linked to the adversary on the evening and night of respectively January 11 and January 12. A selection of these alerts is shown in Table 3. Alerts, later linked to the adversary, indicated that a domain administrator account was used to perform reconnaissance and privilege escalation. This raised immediate suspicions about a potential domain compromise. These suspicions were quickly confirmed after reviewing the available alerts and adversary activity.

Table 3: Overview of the most relevant alerts received by SURFsoc.

Timestamp	Hostname	Username	Alert name	Severity
2025-01-11 23:11:58	system_srv1.campus.domain_2.nl	account_hp1	NCC-MITRE-T1098-001 Member added to Sensitive Group	High
2025-01-11 22:43:03	system_srv2.campus.domain_2.nl	account_hp1	ESCU_a51bfe1a-94f0-48cc-b4e4-16a110145893 Attacker Tools On Endpoint	Critical
2025-01-11 21:54:56	system_dc1_prod.campus.domain_2.nl	ACCOUNT_DC1_PROD	NCC-COMMANDLINE-WINDOWS-001 Domain Administrator Discovery	High
2025-01-11 21:20:51	system_dc1_prod.campus.domain_2.nl	account_hp1	NCC-COMMANDLINE-WINDOWS-001 CrackMapExec WMIExec	Critical

The account account_hp1, on which the adversary activity was performed, was the built-in default domain administrator account which was designated as the break-glass account. This account also had enterprise administrator privileges. The password of this account was exclusively stored on multiple physical locations as a security measure. Furthermore, TU/e elaborated that the account should not be in use and was blended in with other accounts by giving it a non-descriptive name. These attributes were a great help in quickly determining the account to be compromised by an adversary.

3.3.2 Privilege escalation to a domain administrator account

As stated in Section 3.2.1, the adversary connected to TU/e's network through a VPN solution of TU/e. On 11 January 2025 at 19:59 a successful authentication was registered on SYSTEM_DC4_PROD using the system account ACCOUNT_DC4_PROD from IP address ip_vpn_client_1. This IP address was assigned to a VPN session associated to the adversary. Twenty seconds later, a DCSync [T1003.006] attempt originated from the same VPN IP address. The DCSync attack was identified by Microsoft Defender on SYSTEM_DC4_PROD and was classified by Microsoft Defender as unsuccessful.

At 20:59, an hour after the previous attempt, another successful authentication was registered on one of TU/e's domain controllers. This time on SYSTEM_DC1_PROD using the system account ACCOUNT_DC4_PROD. This time from address IP ip_vpn_client_2 , again assigned to a VPN session of the adversary. Three seconds later, another DCSync attack was performed from adversary's VPN IP address ip_vpn_client_2. This time the attack was successful, as indicated by another Microsoft Defender alert.

Fox-IT considered two hypotheses of attack paths that likely attributed to the successful DCSync. After careful consideration, one of the hypotheses was dismissed as it became clear that it was technically impossible to have attributed to the DCSync.

Fox-IT continued its investigation based on the remaining hypothesis that focussed on a coercion attack and subsequent cracking of NTLMv1 challenge/response hashes. The following subsection describes the moments that lead up to the successful DCSync attack and coercion attack.

3.3.3 Coercion attack and NTLMv1 authentications

Fox-IT investigated the authentication methods used between domain controllers. As shown in Table 4, it became apparent that leading up to 11 January, multiple domain controller computer accounts were authenticated using the NTLMv1 method. This includes several authentications (marked in red) from domain controller computer accounts to multiple domain controllers, all originating from IP addresses assigned to adversary VPN sessions.

Table 4: Selection of unique authentications (Windows Event ID 4624) using the NTLMv1 authentication method.

Timestamp	Account name	Hostname	Source IP
2025-01-10 14:51:54	ACCOUNT_DC4_PROD	SYSTEM_DC2_PROD	ip_vpn_client_3
2025-01-10 14:35:54	ACCOUNT_DC2_PROD	SYSTEM_DC4_PROD	ip_vpn_client_3
2025-01-10 14:34:15	ACCOUNT_DC1_PROD	SYSTEM_DC2_PROD	ip_vpn_client_3
2025-01-10 14:32:46	ACCOUNT_DC3_PROD	SYSTEM_DC2_PROD	ip_vpn_client_3
2025-01-06 18:59:16	ACCOUNT_DC3_PROD	SYSTEM_DC1_PROD	ip_system_dc3_prod
2024-12-30 10:27:21	ACCOUNT_DC1_PROD	SYSTEM_DC3_PROD	ip_system_dc1_prod_1
2024-12-22 11:16:00	ACCOUNT_DC2_PROD	SYSTEM_DC3_PROD	ip_system_dc2_prod

Authentications with computer accounts originating from any host other than its associated host are highly suspicious and is indicative of a potential compromise. It is likely that the adversary performed a coercion attack followed by cracking NTLMv1 challenge/response hashes. With a coercion attack, the adversary attempts to trick a host to authenticate to their host instead of the intended target host. This is done by poisoning host discovery protocols. Subsequently, the adversary can crack the challenge/response hash that it recorded during the coercion attack.

The NTLMv1 authentication method was accepted on most domain controllers because the '`lmcompatibilitylevel`' setting was set to allow NTLMv1 authentications. An overview of the allowed authentication methods per domain controller can be found in Appendix A.3, Table 12.

Irrefutable evidence is absent, as traces of raw communication between the adversary and TU/e's systems (also known as packet captures), were not present to conclusively determine whether the adversary executed a coercion attack.

> *Through circumstantial evidence, Fox-IT considers it likely that between 6 January 2025 and 11 January 2025, the adversary coerced multiple domain controllers into downgrading and authenticating to the adversary via the NTLMv1 authentication protocol and cracked the challenge/response hashes.*

3.3.4 Validating if TU/e domain infrastructure could be attacked via DCSync

To validate Microsoft Defender's detection of the successful DCSync, Fox-IT investigated if TU/e's infrastructure provided the conditions to allow a DCSync attack. A DCSync attack leverages benign protocols and services that allow synchronization between domain controllers. A successful DCSync attack allows the adversary to retrieve all password hashes stored on a specific domain controller. An adversary can then re-use these hashes in a pass-the-hash attack to authenticate to computers and services without the need for the actual password.

To successfully perform a DCSync, the adversary:

- should be able to communicate from the VPN subnet(s) to one of the domain controllers
- should have obtained credentials of an account with domain replication rights
- should be able to authenticate to one of the domain controllers

Based on successful authentications with a domain controller computer account[7] from the VPN subnet, it showed that TU/e's infrastructure met these conditions. Table 5 shows details of these authentications and summarises that all three conditions were met. Mere seconds before the successful DCSync attack, the adversary authenticated with account `ACCOUNT_DC4_PROD` on `SYSTEM_DC1_PROD` from adversary's VPN IP `ip_vpn_client_2`, as shown in Table 5. Note that the authentication method (NTLM) and source IP address on the rows marked red, stand out from legitimate authentication behaviour.

[7] Domain controller computer accounts have domain replication rights by default. These accounts can easily be identified as the username carries the host name of the domain controller and the $ suffix.

Table 5: Legitimate and malicious successful authentications from ACCOUNT_DC4_PROD.

Timestamp	Hostname	User	Source IP	Authentication method
11 January 2025 at 20:58:09	SYSTEM_DC4_PROD	ACCOUNT_DC4_PROD	ip_system_dc4_prod	Kerberos
11 January 2025 at 20:58:41	SYSTEM_ROOTDC2_PROD	ACCOUNT_DC4_PROD	ip_system_dc4_prod	Kerberos
11 January 2025 at 20:59:12	SYSTEM_DC4_PROD	ACCOUNT_DC4_PROD	ip_vpn_client_2	NTLM V2
11 January 2025 at 20:59:17	SYSTEM_DC2_PROD	ACCOUNT_DC4_PROD	ip_system_dc4_prod	Kerberos
11 January 2025 at 20:59:22	SYSTEM_DC4_PROD	ACCOUNT_DC4_PROD	ip_vpn_client_2	NTLM V2
11 January 2025 at 20:59:28	SYSTEM_DC3_PROD	ACCOUNT_DC4_PROD	ip_system_dc4_prod	Kerberos
11 January 2025 at 20:59:28	SYSTEM_DC1_PROD	ACCOUNT_DC4_PROD	ip_vpn_client_2	NTLM V2
11 January 2025 at 20:59:45	SYSTEM_DC1_PROD	ACCOUNT_DC4_PROD	ip_vpn_client_2	NTLM V2

These traces confirm that the adversary somehow managed to obtain or crack the password of account ACCOUNT_DC4_PROD, allowing the adversary to imitate a DCSync on SYSTEM_DC1_PROD.

Traces of successful domain replication events on SYSTEM_DC1_PROD, that could have validated Microsoft Defender's DCSync alert, were non-existent. This is due to the audit policy configuration on SYSTEM_DC1_PROD for 'Directory Service Replication' and 'Details Directory Service Replication' events, which only logged failure events, as shown in Table 6. The audit policy configuration of all domain controllers can be read in Appendix A.4, Table 13.

Table 6: Domain controller domain replication audit policies.

Hostname	Event log category	Event log name	Value
SYSTEM_DC1_PROD	DS Access	Detailed Directory Service Replication	Failure
SYSTEM_DC1_PROD	DS Access	Directory Service Replication	Failure

Fox-IT considers it likely that on 11 January 2025 at 20:59 the adversary successfully executed a DCSync attack to host SYSTEM_DC1_PROD by obtaining computer account credentials from host SYSTEM_DC4_PROD.

With the DCSync attack being successful, the adversary now obtained all NTLM hashes of all accounts managed on domain controller SYSTEM_DC1_PROD. This includes the NTLM hashes of all user accounts, including accounts with domain administrator or similar high privileges. This allowed the adversary to authenticate to any domain joined computer, its underlying services, and data with ease, via a pass-the-hash attack, without knowing the actual password of an account.

About eight minutes later, at 21:07, a successful authentication from the adversary's VPN IP was registered on SYSTEM_DC1_PROD with the default domain administrator account account_hp1.

After the adversary gained the highest privileges within the DOMAIN_2 and DOMAIN_1 domain, full control was achieved over both domains. At any time, the adversary could theoretically have deployed ransomware across all domain joined systems. This notion played an important role in the decision-making process.

Fox-IT considers the Active Directory domains configured on SYSTEM_DC1_PROD to be compromised since 11 January 2025 at 21:07. This is the moment where the adversary utilised earlier obtained high privileged credentials of the default domain administrator account to authenticate to SYSTEM_DC1_PROD.

3.4 Post-privilege escalation activities by the adversary

This section describes in detail the activities after the adversary gained full control of the TU/e domain. Subsection 3.4.1 focuses on the discovery activities from the adversary. Subsection 3.4.2 describes persistence activities by the adversary with remote administration tools, whereas Subsection 3.4.3 describes persistence in terms of accounts created by the adversary. Lastly, Subsection 3.4.4 describes the adversary activity involving TU/e's backup solution.

3.4.1 Discovery activities by the adversary

The program Advanced IP Port Scanner[8] was executed by the adversary on 11 January 2025 at 22:43 on the system `system_srv2` and the program SoftPerfect Network Scanner[9] was executed on 11 January 2025 at 22:53 on system `tfe290`. These types of programs collect information about reachable systems in the network and retrieve information about the network devices.

On 11 January 2025 at 23:56, the adversary executed a command related to the program ShareFinder[10] with the break-glass account on the system `SYSTEM_SRV4`. ShareFinder is a program that allows a user to discover accessible shared drives on the network. The code block below shows the actual command retrieved from the PowerShell logs on the system `SYSTEM_SRV4`.

```
Invoke-ShareFinder -CheckShareAccess -Verbose | Out-File -Encoding ascii
C:\programdata\found_shares.txt
```

The executed command searched for shares in the network and wrote the results to the file `C:\programdata\found_shares.txt`. Because the information inside this file could reveal information about what information the adversary was able to retrieve, Fox-IT made efforts to retrieve this file. However, the file was no longer present.

Adversaries use programs like Advanced IP Port Scanner, SoftPerfect Network Scanner, and ShareFinder to explore the environment and find as many connected devices as possible. Therefore, Fox-IT considers it highly likely that the adversary used these tools to perform discovery activities in TU/e's network.

3.4.2 Persistence by the adversary with remote administration tools

Fox-IT found traces of two different remote administration tools that were installed and used by the adversary: AnyDesk and TeamViewer. Adversaries make use of these tools to maintain access to systems in the network. If a system where such a tool is installed can connect to the internet, these tools then enable an adversary to remotely log in to the system without requiring access through a VPN.

Fox-IT identified traces of the adversary using AnyDesk on four systems and TeamViewer on a total of three systems (two additional). For the systems `system_srv3` and `system_srv1`, on which AnyDesk was installed, Fox-IT was able to identify successful AnyDesk connections from the firewall logs. This implicates that for these two servers it is highly likely that the adversary used AnyDesk to control them from a remote location. Table 7 shows a summary for the systems on which AnyDesk and/or TeamViewer were installed.

[8] Advanced IP Scanner is a free network scanner tool owned by Famatech Corporation. Read more at https://www.advanced-ip-scanner.com/.

[9] SoftPerfect Network Scanner is a tool to scan IPv4 and IPv6 in a network owned by SoftPerfect Pty Ltd. Read more at https://www.softperfect.com/products/networkscanner/.

[10] ShareFinder is a free tool to discover file shares on a network. Read more at https://github.com/darkoperator/Veil-PowerView/blob/master/PowerView/functions/Invoke-ShareFinder.ps1.

Table 7: Systems for which remote administration tools were installed by the adversary.

Host	Traces of remote desktop tooling	Creation time
system_dc1_prod	Anydesk	12 January 2025 at 00:44
system_rootdc2_prod	Anydesk	12 January 2025 at 00:23
system_srv3	Anydesk/TeamViewer	11 January 2025 at 23:29/11 January 2025 at 23:58
system_srv1	Anydesk	11 January 2025 at 23:27
system_ws1	TeamViewer	11 January 2025 at 23:32
system_ws2	TeamViewer	11 January 2025 at 22:36

3.4.3 Persistence by the adversary with additional and new domain accounts

To not only maintain access but also maintain high-privileged access, adversaries tend to create additional accounts with high-privileges. This way, an adversary can revert to one of these additional accounts when others are disabled or reset. Fox-IT identified the creation of two new high-privileged accounts by the adversary, namely `account_hp4` and `account_hp5`. Table 8 provides an overview of the high-privileged accounts that were either compromised or created by the adversary.

Table 8: Overview of high-privileged accounts that were either compromised or created by the adversary.

Accounts	Description	Time of creation or first time compromised
DOMAIN_2\account_hp1	Compromised by Adversary	11 January 2025 at 21:07
DOMAIN_2\account_hp2	Compromised by Adversary	11 January 2025 at 22:00
DOMAIN_2\account_hp3	Compromised by Adversary	11 January 2025 at 22:01
DOMAIN_2\account_hp4	Created by Adversary	11 January 2025 at 22:46
DOMAIN_2\account_hp5	Created by Adversary	11 January 2025 at 23:11

3.4.4 Suspicious activities involving TU/e's backup solution

Fox-IT identified that the adversary interacted with Veeam on system `system_srv5` on 12 January 2025 at 00:52. The code block below shows a representation of a Defender log on system `system_srv5`. The log shows access was attempted via the Veeam application with the break-glass account.

```
"Command execution: ""Veeam.Backup.Satellite.exe"" ""DOMAIN_2_account_hp1_Console_system_srv5_06b09421-5d4c-4c8c-9287-aedb4c27f53a"""
```

Five minutes later, at 00:57, a PowerShell command was logged on `system_srv5` that showed that the adversary tried to stop Veeam. The code block below shows this exact command.

```
$SqlServerName = (Get-ItemProperty -Path $VeaamRegPath -ErrorAction Stop).SqlServerName`
```

Multiple similar commands were logged where only the variable name `SqlServerName` was replaced with `SqlInstanceName` and `SqlDatabaseName`.

Fox-IT found traces that the adversary performed additional discovery activities after gaining full control over the environment. Traces show that the adversary installed remote administration tools on six systems to expand their persistence methods. Furthermore, Fox-IT found traces that the adversary interacted with TU/e's backup solution.

3.5 Scope of compromise

As described in Section 3.3, the adversary gained full control over TU/e's domain. In essence, this meant that the adversary was able to navigate and access all systems and underlying information within the domain. However, this does not necessarily mean that the adversary accessed all systems of TU/e. This section provides a more detailed scope of the compromise.

Based on the adversary activity identified, Fox-IT divided all the systems within scope into three categories. The first category is called "hands-on-keyboard" and contains the systems on which the adversary logged in and performed (manual) actions. The second category contains the systems on which the adversary logged in, but did not create any traces of follow-up activity. This category is called "accessed only". The last category is called "no activity" and contains the remainder of the systems on which no traces of adversary activity were found.

In summary, Fox-IT found that 91 systems in total contained traces of adversary activity. On fourteen of these systems, Fox-IT found traces of hands-on-keyboard activity by the adversary. On the remaining 77 of these 91 systems, Fox-IT only identified traces of some form of authentication performed by the adversary. Table 9 contains the overview of the number of systems within each category for the scope of compromise.

Table 9: Scope of compromise divided into three categories "hands-on-keyboard", "accessed only" and "no activity".

Category	Number of systems
Hands-on-keyboard	14
Accessed only	77
No activity	259

For readability, Fox-IT does not provide a detailed list of system in this document. The related document named "`Timeline_Armstrong.xlsx`" contains a detailed overview of all the systems and the identified traces per system.

Fox-IT identified that the adversary interacted with at least 91 of the 350 systems. On fourteen of these systems, Fox-IT found traces of hands-on-keyboard activity by the adversary. On the remaining 77 systems that were interacted with, Fox-IT only identified traces of some form of authentication performed by the adversary.

3.6 Data access and analysis of potential data exfiltration

This section dives deeper into the extent of data access that the adversary had and investigates traces that could reveal any sign of data collection and exfiltration. Subsection 3.6.1 explains the range of access the adversary had to TU/e's data. Hereafter, Subsection 3.6.2 describes the traces that Fox-IT searched for to find potential signs of data exfiltration and the results thereof.

3.6.1 The adversary's range of access to TU/e's data

As described in Section 3.3, the adversary obtained the highest level of privileges within TU/e's DOMAIN_2 and DOMAIN_1 domains. This level of privileges can be leveraged to get access to all computers and servers within the compromised domains. Therefore, the adversary could access all unencrypted data that was stored on these systems.

In some cases, specific data on systems may be stored encrypted. Examples are encrypted databases created by a password manager or password protected documents. This data can only be accessed with knowledge about the decryption key and/or password. This encrypted data is therefore not directly accessible by an adversary with the highest level of access. However, the level of access does allow an adversary to use multiple techniques to intercept the decryption key and/or password, such as using keyloggers. These techniques are not always successful and therefore this data is, in general, less likely to be accessed by an adversary.

Because of the enterprise administrator privileges, the adversary could in theory access at least all unencrypted data on TU/e's systems. The remainder of this section elaborates on the search for signs of data exfiltration performed by Fox-IT and the results thereof. However, one should keep in mind that the absence of traces does not necessarily mean that no data has been exfiltrated.

3.6.2 Search for traces of potential data exfiltration based on multiple forensic data sources

Because of the extent of access to TU/e's data, Fox-IT made additional efforts to find any traces of data exfiltration by the adversary in multiple sources. Adversaries can exfiltrate data to use it as leverage to extort their victims. Because adversaries do not usually know what specific data is most useful for this extortion means, they tend to exfiltrate a broad scope of data.

Fox-IT looked for traces of the following two tactics used by the adversary to determine if data exfiltration took place:

- traces of data collection
- traces of data exfiltration

The remainder of this subsection explains these tactics together with their related potential traces and concludes if any of these traces were found within the investigation data.

Traces of data collection

The data collection tactic [TA0009] is generally comprised of the data archiving and staging techniques. An adversary may either manually or automatically [T1119] search for data of interest. Data is then often archived/compressed [T1560] and staged [T1074] to facilitate swift data exfiltration.

Investigative efforts were directed to discover traces of commonly used file archiving and compression filetypes such as zip, rar and gz. Efforts were also directed to discover the use of software that facilitates automated data collection. Fox-IT found no traces of (compressed) archives or software that indicate that the adversary collected and staged data for exfiltration.

Traces of data exfiltration

The data exfiltration tactic [TA0010] constitutes techniques to transfer data to an external location under the control of the adversary. In general, the adversary may use their Command and Control (C2) channel [T1041], specialised exfiltration programs [T1048] or web services [T1567] to exfiltrate the data.

First, Fox-IT determined the amount of data that was transferred to the three IP addresses known to be used by the adversary and discussed in Section 3.2. These three IP addresses can be regarded as the C2 channels of the adversary. Based on the firewall logs and TU/e's network data provided by SURF, Fox-IT determined that in total approximately 2.1 gigabytes of data were transferred to the IP addresses in the timeframe of January 5 to 12 January 2025.

Because the log sources only contained metadata about the network traffic, Fox-IT was not able to determine the exact content of the traffic. However, Fox-IT considers it likely that a substantial part of the data is comprised of general connection data for VPN usage and sensitive information about TU/e's Active Directory such as system names, usernames and password hashes. The latter is assumed, because of the privilege escalation and discovery traces identified, for which it's likely that the adversary retrieved the (intermediate) information for analysis.

Secondly, Fox-IT looked for traces of exfiltration programs known to be used by adversaries. Of such programs, Fox-IT solely found traces of remote administration tools, as mentioned in Subsection 3.4.2. These tools contain the ability to transfer files. Based on the firewall logs, Fox-IT did not find any traces that suggested that the adversary used these tools to perform large-scale data exfiltration.

Lastly, Fox-IT looked at signs of web services for data exfiltration used by the adversary. Fox-IT found no traces that indicated that the adversary used such web services to perform large-scale data exfiltration.

In general, Fox-IT would like to stress that the absence of traces does not imply that strictly no large-scale data exfiltration has taken place. However, the absence of traces makes it less likely that the adversary engaged in activities that involved exfiltrating a substantial amount data from TU/e's network.

Fox-IT found no traces of large-scale data exfiltration in the investigation data within the period of the incident, between 6 and 12 January 2025. Fox-IT did find traces that make it likely that the adversary exfiltrated sensitive information from TU/e's Active Directory, such as usernames and password hashes.

3.7 Threat Actor

This section aims to sketch a general threat actor profile of the adversary. To do so, this section uses the identified TTPs to categorize the adversary's profile as much as possible in Subsection 3.7.1. Subsection 3.7.2 provides a comment on the significance of Cyrillic characters found in commands executed by the adversary.

3.7.1 General threat actor profile based on identified TTPs

Based on the identified TTPs of the adversary described in the preceding sections of this chapter, Fox-IT considers it likely that the adversary fits the profile of a ransomware threat actor. The initial access via TU/e's VPN with an existing account, the usage of well-known off-the-shelf tooling for lateral movement and persistence, and the attempt to stop TU/e's backup solution, are in line with precursors to a full domain compromise followed by ransomware encryption.

Moreover, the techniques applied by the adversary after gaining full control over TU/e's domain resulted in multiple security alerts received by SURFsoc. This showed that the adversary attached little value to solely using techniques that ensured the activities remained unnoticed. Advanced threat actors, however, invest a lot of effort in staying under the radar. Therefore, it's unlikely that the adversary fits the profile of an advanced threat actor.

Although the TTPs allowed Fox-IT to determine this general threat actor profile, they did not allow for pinpointing the exact (ransomware) threat actor.

3.7.2 Comment on significance of Cyrillic characters in commands performed by the adversary

Within the investigated commands performed by the adversary, Fox-IT found traces of Cyrillic characters. The code block below shows a command executed by the account `account_hp1` on 12 January 2025 at 00:58.

```
net  group "Domain Admins" /domain   - узнать ДА
```

The Cyrillic characters in the command translate to "Find out DA". This makes it highly likely that these characters are meant as a comment to explain that this command is used to find domain administrator accounts. The presence of these characters is no conclusive evidence with regards to the geographical origin of the adversary.

Fox-IT was not able to determine the exact threat actor. However, Fox-IT considers it likely that the adversary fits the profile of a ransomware actor. The used TTPs and their off-the-shelf and non-stealthy nature contribute to this likeliness.

4 Conclusions

Based on the findings from the conducted investigation, Fox-IT formulates the following answers on the investigation questions as posed in Section 1.2.

1 What happened?

On 6 January 2025 at 14:08, the adversary successfully logged into TU/e's VPN with the account `account_lp2`. Five minutes later, at 14:13, the adversary logged in with the account `account_lp3` from the same IP address. After gaining this initial access, the adversary performed discovery activities in TU/e's network. On 11 January at 21:07, the adversary managed to escalate their privileges by obtaining access to TU/e's break-glass account. Having the highest privileges, the adversary continued expanding their foothold in TU/e's network. On 12 January 2025 at 00:52, the adversary attempted to stop TU/e's backup solution. Due to TU/e disconnecting their network from the internet on 12 January 2025 at 01:17, the adversary was no longer able to connect to TU/e's network. This therefore immediately stopped the attack.

2 How did it happen?

Fox-IT considers it likely that the adversary obtained leaked credentials of the accounts `account_lp2` and `account_lp3`. With these credentials, the adversary was able to connect to TU/e's network through the VPN solution for which no multi factor authentication was required. From this VPN, the adversary was able to communicate with domain controllers and other services. Though irrefutable evidence is absent, Fox-IT considers it likely that the adversary coerced domain controller `SYSTEM_DC4_PROD` into downgrading to the NTLMv1 authentication protocol and authenticating to the adversary. Fox-IT considers it likely that this allowed the adversary to obtain and crack a NTLMv1 challenge response from computer account `ACCOUNT_DC4_PROD` of `SYSTEM_DC4_PROD`.

With the adversary having obtained the password of a computer account of one of the domain controllers, the adversary successfully executed a DCSync attack to domain controller `SYSTEM_DC1_PROD`. This attack exposed hashed passwords of all accounts present in the Active Directory of this domain controller. Fox-IT considers it highly likely that the adversary then obtained the highest possible privileges within the domain by using the hash of the default domain administrator account `account_hp1` to authenticate to other services.

3 What is the scope of the compromise?

Fox-IT Identified that the adversary obtained enterprise administrator privileges via the break-glass account in the domain of TU/e. Theoretically the entire domain, including all assets in that domain, should be considered compromised. However, the adversary did not target all assets in TU/e's domain. Fox-IT identified traces of the adversary on a total of 91 systems. On fourteen of these systems, Fox-IT found traces of hands-on-keyboard activities by the adversary. On the remaining 77 of these systems, Fox-IT only identified traces of some form of authentication performed by the adversary without any follow-up activity.

4 What data was accessed by the adversary?

Because of the enterprise administrator privileges, the adversary could in theory access at least all unencrypted data on TU/e's systems. However, Fox-IT did not find traces of large-scale data exfiltration within the period of the incident, 6 and 12 January 2025 in the investigation data. Fox-IT did find traces that make it likely that the adversary exfiltrated sensitive information from TU/e's Active Directory, such as usernames and password hashes.